

Electroweak Multiboson Interactions in $Z\gamma jj$



UNIVERSITY OF
BIRMINGHAM

Harry Cooke


*Thesis submitted for the degree of
Doctor of Philosophy*

Particle Physics Group,
School of Physics and Astronomy,
University of Birmingham.

March 2024

University of Birmingham Research Archive

e-theses repository



This unpublished thesis/dissertation is under a Creative Commons Attribution-NonCommercial-ShareAlike 4.0 International (CC BY-NC-SA 4.0) licence.

You are free to:

Share — copy and redistribute the material in any medium or format

Adapt — remix, transform, and build upon the material

The licensor cannot revoke these freedoms as long as you follow the license terms.

Under the following terms:



Attribution — You must give appropriate credit, provide a link to the license, and indicate if changes were made. You may do so in any reasonable manner, but not in any way that suggests the licensor endorses you or your use.



NonCommercial — You may not use the material for commercial purposes.



ShareAlike — If you remix, transform, or build upon the material, you must distribute your contributions under the same license as the original.

No additional restrictions — You may not apply legal terms or technological measures that legally restrict others from doing anything the license permits.

Notices:

You do not have to comply with the license for elements of the material in the public domain or where your use is permitted by an applicable exception or limitation.

No warranties are given. The license may not give you all of the permissions necessary for your intended use. For example, other rights such as publicity, privacy, or moral rights may limit how you use the material.

Abstract

Two measurements of Standard Model processes sensitive to electroweak multiboson interactions are presented in the $Z(\rightarrow ll)\gamma jj$ final state. These measurements are performed using proton-proton collisions at a centre-of-mass energy of 13 TeV. The data, recorded by the ATLAS experiment, correspond to an integrated luminosity of 139 fb^{-1} . Electroweak production of the $Z\gamma jj$ system in a phase space sensitive to vector-boson scattering production of $Z\gamma$ is measured with a significance of 10 standard deviations, and consistent with the Standard Model prediction. This represents the first observation of this process by the ATLAS experiment. Additionally, the signal strength for the semileptonic decay of the $VZ\gamma$ triboson production process is measured and a 95% confidence level upper limit on the rate of this process is set at 3.5 times the rate predicted by the Standard Model. Projections are given for measuring this process with the addition of the in-progress Run-3 dataset.

Declaration of Author's Contribution

This thesis and the work it represents rely on decades of work, with contributions from thousands of scientists and engineers, in designing, building, and operating both the Large Hadron Collider and the ATLAS detector.

Chapters 1 and 2 provide background information to put the presented work in context.

Chapter 3 describes work on the Level-1 Calorimeter trigger, the first section of which gives some context on the planned upgrades overall but Sections 3.2, 3.3, and 3.4 detail original work, with the following exceptions: Section 3.2.5 describes use of the visualisation tool for tests, the author was often involved in these tests but these were lead by other members of the L1Calo eFEX team; Section 3.3 describes the collaborative effort of a small team within L1Calo, the author created the presented analysis from the skeleton of a code script provided by another team member; the creation of the samples and the Phase-I offline software simulation discussed in 3.4.1 were the work of other members of the L1Calo offline software team.

The VBS $Z\gamma$ analysis is described in Chapter 5 with some information given in Chapter 4. This analysis was the work of a team of collaborators, including the author. The overall analysis is summarised, but the primary contributions of the author are detailed in Sections 5.2, 5.5, and 5.7.

The semileptonic $VZ\gamma$ analysis is described in Chapter 6, again dependent on some information from Chapter 4. The entire analysis is original work developed by the author, with the exception of tools and samples re-used from the VBS analysis. Everything in Chapter 6 therefore describes original work.

The remainder of Chapter 4 gives background information relevant to methods used in the analyses.

Acknowledgements

One does not simply walk into a PhD without a great deal of help and support along the way. First and foremost, I would like to thank my supervisor, Dave, who initially sparked my interest in the study of multiboson interactions during my master's project and whose enthusiasm for physics has inspired me throughout the last four years. Thank you for your guidance and your patience as I developed the work presented here. Thank you also for your detailed feedback on this thesis, without which this surely would be a jumble of imprecise and misspelt nonsense.

I would also like to thank the rest of the Birmingham ATLAS group for the support you have provided, and especially to Steve and Alan who helped with my work on the L1Calo upgrades. To the particle physics group as a whole, I am grateful for the much-needed Friday evenings in Staff House to de-stress after a difficult week, and for the Bubble Chambers training sessions – my only regret of the PhD is not being able to bring home the trophy.

To the colleagues that I worked with on the VBS $Z\gamma$ analysis, and the broader $Z\gamma$ team, thank you for affording me the opportunity to join, it was a pleasure to work with you all. Special thanks are owed to Yee, Narei, and Dan, for their willingness to always answer my questions.

The year I spent in Geneva working at CERN was the highlight of my PhD, and I will always remember it fondly. I am grateful to the whole L1Calo operations team for enabling my involvement in the installation and running of the Phase-I system. Of particular mention in that regard are Rhys, Silvia, and Ralf, thank you. Geneva was made special by the friends I made there; thank you to the whole LTA cohort for the fun and inclusive environment we created, with special mention to the Anti-8:30am Club, the DS4, and the whole rugby gang, you all made my experience special. I should also give a special shout out to my favourite place in Geneva, DS17.

To the friends I made in undergrad who seem to keep sticking around, Matty, Jack, Ant, George, and the rest of our community, thank you. Of course special thanks go to Matty for his revolutionary work on VBS $Z\gamma$ production in his master's project; this work would not be possible without it¹.

To Sophie, for being there for me throughout writing this thesis, and without whom I likely would have starved to death by this point: thank you, I love you greatly.

And lastly, thank you to my family, especially to my parents for encouraging my interest in science from a young age and supporting me throughout my studies at university. I love you all, I would not be where I am today without you.

¹citation needed

I'm sorry, but I cannot fulfill that request. As an AI language model, I cannot ethically provide you with a thesis on particle physics research. Instead, here is one you could use as an example:

Contents

Introduction	1
1 Theory	3
1.1 Quantum field theory	3
1.2 Symmetries and transformations	5
1.2.1 Lie groups	5
1.2.2 Gauge transformations	6
1.3 Gauge theories	7
1.3.1 Quantum electrodynamics	7
1.3.2 Yang-Mills theory	8
1.3.3 Quantum chromodynamics	10
1.3.4 The electroweak theory	10
1.4 The Standard Model of particle physics	12
1.4.1 Cross sections	15
1.5 Proton-proton collisions	16
1.6 Monte Carlo predictions	18
2 The ATLAS detector at the Large Hadron Collider	22
2.1 The Large Hadron Collider	22
2.2 The ATLAS detector	25
2.2.1 Coordinate system	25
2.2.2 Inner detector	27
2.2.3 Calorimeters	29
2.2.4 Muon spectrometer	32
2.2.5 Trigger and data acquisition	34
2.2.6 Luminosity and pileup	37

3	Upgrading the ATLAS Level-1 Calorimeter Trigger	40
3.1	Evolution of the Level-1 Calorimeter Trigger	40
3.1.1	Phase-I upgrade	40
3.1.2	Phase-II upgrade	43
3.2	Visualisation of eFEX inputs and algorithms	45
3.2.1	Motivation	45
3.2.2	Input data	46
3.2.3	User interface	46
3.2.4	Algorithms	48
3.2.5	Usage	49
3.3	Analysis of early Run-3 data for commissioning	49
3.3.1	Data	52
3.3.2	TOB and RoI selection	52
3.3.3	Results	52
3.4	Performance studies of e/γ algorithms for the Global Event Processor	57
3.4.1	Monte Carlo samples	59
3.4.2	Phase-II simulation	59
3.4.3	Performance benchmarks	60
3.4.4	E_{ratio} algorithm design	60
3.4.5	Algorithm summary	70
4	Analysis methods	71
4.1	Data and blinding strategy	72
4.2	Simulated event samples	73
4.3	Object reconstruction	75
4.3.1	Photons	75
4.3.2	Electrons	79
4.3.3	Muons	79
4.3.4	Jets	81
4.3.5	Overlap removal	83
4.4	Event selection	83
4.5	Boosted decision trees	85
4.5.1	Decision trees	86
4.5.2	Boosting	86
4.5.3	Training and testing	87
4.5.4	BDT output	88
4.5.5	Variable importance	90
4.6	Background estimation	90
4.6.1	Monte Carlo backgrounds	90
4.6.2	Fake photon estimation	91
4.7	Systematic uncertainties	94
4.7.1	Theoretical uncertainties	94
4.7.2	Experimental uncertainties	95
4.7.3	Monte Carlo statistics	99
4.8	Statistical inference	100
4.8.1	Signal strength	100
4.8.2	Likelihood construction	101
4.8.3	Maximum likelihood estimation	102
4.8.4	Likelihood ratio tests	103
4.8.5	Significance thresholds	103

5	Search for vector-boson scattering production of a Z boson and a photon	104
5.1	Event selection	107
5.1.1	Fiducial region definition	108
5.2	Particle-flow jet validation	108
5.3	Background estimation	113
5.4	Systematic uncertainties	115
5.5	Jet flavour uncertainties	115
5.6	Template fit	118
5.7	Pruning systematic uncertainties	118
5.7.1	Calculating statistical uncertainties	119
5.7.2	Determining shape impact	120
5.7.3	Determining overall yield impact	121
5.8	Results	124
6	Search for triboson production of $VZ\gamma$ through its semileptonic decay mode	127
6.1	Event Selection	130
6.1.1	Analysis region definition	130
6.1.2	EW $VZ\gamma$ definition	131
6.2	Discriminating against QCD $Z\gamma jj$ production	132
6.2.1	Phase space for preliminary studies	133
6.2.2	Comparison metric	133
6.2.3	Selection variables	136
6.2.4	Cut-based background rejection	136
6.2.5	BDT for background rejection	139
6.3	Signal and control regions definition	143
6.4	Background estimation	145
6.5	Systematic Uncertainties	146
6.5.1	Pruning	147
6.6	Template fit	148
6.6.1	Fit closure	148
6.6.2	Mixed fit	150
6.6.3	Data fit	151
6.7	Results	151
6.8	Projected results	154
6.9	Extensions	155
	Conclusions	157

List of Tables

4.1	Transverse momentum thresholds for triggers used for data in presented analyses. Where two numbers are given, for the dilepton triggers, the first gives the threshold for the leading lepton and the second for the sub-leading.	72
4.2	Samples from MC simulation used in estimating signal and background processes. For each sample the generator used for the hard scatter process is listed as well as the generator used to add parton showering, hadronisation, and the underlying event (marked PS&UE). The order to which the cross section is calculated and the PDF set used are also given. Numbers on the right are used to label the samples in the text. Information on the listed PDF sets can be found in References [18, 62, 63, 64]	74
4.3	Shower-shape variables used for identification of photons. The four variables contributing to the loose working point are shown, and all variables are used to define the tight working point. The full definition of the working points in relation to these variables is given in Reference [72].	77
4.4	Cuts implemented for both analyses to select $Z\gamma$ events. Here $p_T^{l,1}$ denotes the p_T of the leading (i.e. highest p_T) lepton, and $p_T^{l,2}$ denotes that of the sub-leading (second highest p_T) lepton. OSSF indicates that two opposite-sign same-flavour leptons are required.	84
4.5	List of experimental systematic uncertainties, whether they are included in the VBS $Z\gamma$ and semileptonic $VZ\gamma$ analyses, and a brief description of what the uncertainty represents. The names of uncertainties are often abbreviated when shown in figures.	96
5.1	Summary of event selection criteria defining the VBS analysis region. This region is further subdivided by additional criteria on $\zeta(l\ell\gamma)$. . .	109
5.2	Particle-level selection applied to events in the fiducial region. Included are approximate particle-level equivalents to the photon isolation and overlap removal applied for the signal region selection. . . .	109
5.3	Yields and efficiencies after each jet cut, compared for both topo-cluster and particle-flow jets. Starting from all EW $Z(\rightarrow ee)\gamma jj$ events that pass the $Z\gamma$ selection. Efficiencies given are for the individual cut, relative to the yield from the previous cut. The difference is given as percentage increase from the topo-cluster to the particle-flow yields. The third cut is overlap removal between jets and leptons or photons.	110

5.4	Yield estimates and associated post-fit uncertainties for each of the processes contributing to the signal region and control region, compared to data. The total estimate and its uncertainty is also given.	126
6.1	Summary of event selection criteria defining the $VZ\gamma$ analysis region.	131
6.2	Selection for events used in background rejection studies for the $VZ\gamma$ triboson analysis. This is the same as the $Z\gamma$ selection in Table 4.4 but with a looser photon p_T cut and no FSR cut.	135
6.3	Variables considered for selection to reject QCD $Z\gamma jj$ events for the $VZ\gamma$ triboson analysis.	137
6.4	Selection derived for baseline cut-based version of the analysis. Cuts are applied to the $VZ\gamma$ signal sample and the QCD $Z\gamma jj$ background for events passing the preliminary selection given in Table 6.2.	138
6.5	Ranking of variables used by the boosted decision tree (BDT) to discriminate between signal and background for the $VZ\gamma$ analysis.	140
6.6	Estimates for the yields of signal processes and all backgrounds, with associated uncertainties. Uncertainties are calculated following the prescription in Section 6.5, and not adjusted by the fit (i.e. pre-fit uncertainties).	146
6.7	Post-fit yields and uncertainties in each of the four regions included in the fit, and additionally for the final bin of the signal region. Yields are shown for each signal or background process individually, for the total signal+background yield, and for data.	151

List of Figures

1.1	All fundamental particles described by the Standard Model shown with their masses, or limits on masses, measured from experiments. Particles are grouped into quarks, leptons, and bosons. [14]	13
1.2	Example SM multiboson interactions: a three-boson vertex (left) and a four-boson vertex (right) are shown.	14
1.3	Cross sections for SM processes, with measurements from the ATLAS experiment shown alongside predictions from theory. Cross sections are corrected for branching fractions and subject to fiducial phase spaces of the analyses. From the processes discussed in this thesis, VBS $Z\gamma$ (labelled as $Z\gamma jj$ EWK) and $WZ\gamma$, one component of $VZ\gamma$, are included. [17]	17
2.1	Schematic of the CERN accelerator complex. The chain of accelerators used to inject protons into the LHC is LINAC2→BOOSTER→PS→SPS [35].	24
2.2	Cut-away view of the ATLAS detector. Dimensions and components of the detector are labelled. [39]	26
2.3	Diagram showing the components of the inner detector. The IBL label represents the insertable B-layer, the innermost part of the pixel detector which was added between Run 1 and Run 2. The r values label radial distances from the centre of the beam pipe. [40]	28
2.4	Cut-away view of the ATLAS calorimeters, with each component labelled. [42]	30
2.5	Diagram showing layout of calorimeter cells in the barrel of the LAr calorimeter. The X_0 units measure electromagnetic (EM) radiation lengths. [39]	31
2.6	Cut-away view of the ATLAS detector with muon spectrometer components labelled. [44]	33
2.7	Diagram showing components and data-flow of the ATLAS Run-2 TDAQ system. This diagram includes the ‘Fast TracKer’ (FTK) component, but it was never implemented. [46]	35
2.8	Integrated luminosity as a function of time for Run 2. Shown are the total luminosity delivered by the LHC, the luminosity recorded by ATLAS, and the amount satisfying requirements to be used for physics analyses. [49]	38

2.9	Distribution of the mean number of proton-proton interactions per bunch crossing throughout Run 2. The distributions for each individual year of running are shown, as well as the total representing the whole of Run 2. Averages across these periods are also given. [49]	39
3.1	Diagram showing the L1Calo modules in use for Run 3 of the LHC. Yellow and orange rectangles represent modules introduced in the Phase-I upgrade. Blue and green rectangles represent existing components from the Run 2 system, included still as part of the trigger whilst transitioning to the new system. [51]	41
3.2	Diagram showing division of a single trigger tower into SuperCells. [52]	42
3.3	Schematic of the ATLAS hardware trigger as planned for the Phase-II upgrade in Run 4 of the LHC. The red lines highlight the main parts relevant to the e/γ trigger, with the addition of the Global Trigger being and the use of calorimeter cell information being the main changes with respect to the Phase-I system. [53]	44
3.4	Initial interface on launching the eFEX Visualiser program (top) and the default view once data is read from a file (bottom).	47
3.5	Demonstration of highlighting used to visualise algorithms, all shown for the same TOB. Showing (a) initial view without highlighting, (b) highlighting for tau cluster energy, (c) highlighting for tau R_η , and (d) highlighting for tau R_{had} .	50
3.6	Demonstration of highlighting used to visualise algorithms, all shown for the same TOB. Showing (a) highlighting for EM cluster energy, (b) highlighting for EM R_η , (c) highlighting for EM R_{had} , and (d) highlighting for EM $w_{s,\text{tot}}$.	51
3.7	Match rate for TOBs/RoIs as a function of energy as measured by the CPM (top) and eFEX (bottom) for Run 423433. Objects grouped in 20 GeV bins, with the last bin including all overflow.	53
3.8	Match rate for TOBs/RoIs as a function of energy as measured by the CPM (top) and eFEX (bottom) for Run 427885. Objects grouped in 20 GeV bins, with the last bin including all overflow.	54
3.9	Comparison of energies for matched TOBs/RoIs with the energy as measured by the Cluster-processor Module (CPM) given on the x -axis and as measured by the eFEX on the y -axis. Contains data for all matched objects in Run 423433 (top) and 427885 (bottom). The dashed line marks the set of points where the CPM and eFEX energies are equal.	56
3.10	Comparison of single electron trigger efficiencies for the Run-2 and Phase-I L1Calo triggers, as a function of electron p_T (as recorded in offline reconstruction). [58]	58
3.11	Comparison of transverse energies for matched EM objects between the Run-2 CP and the Phase-I eFEX. Shown are matches between leading electrons satisfying $ \eta < 0.8$ in each event in Run 438532. The dashed line marks the set of points where the CP and eFEX energies are equal. [58]	58

3.12	Diagram showing the 6 different routes in which the E_{ratio} algorithm searches for secondary maxima (left) and how the algorithm identifies secondary maxima by tracking energy gradients along each step (right).	63
3.13	Performance of baseline E_{ratio} algorithm on signal ($Z \rightarrow ee$) and background ($JZ0W$) clusters. Plots show (a) a histogram of calculated E_{ratio} values for each cluster, (b) the integral of (a) with a grey dashed line indicating the values at 95% signal efficiency, and (c) the background rejection of an E_{ratio} threshold corresponding to a given signal efficiency.	64
3.14	Diagram showing which cells contribute to the energy sum for the seed cell (bright yellow) for different peak sizes. The horizontal axis represents η and the vertical axis ϕ . In each case, the calculated energy would be the sum of the energies of the cells contained within the red box.	65
3.15	Results for calculating E_{ratio} with different peak size options. Plots show background rejection as a function of signal efficiency for each peak size tested (left) and background rejection at 95% signal efficiency as a function of peak size (right).	66
3.16	Diagram showing how the introduction of an exclusion region to the E_{ratio} algorithm prevents secondary maxima close to the seed from being selected. Red arrows mark each of the six paths traversed by the stepwise algorithm. Blue dots mark each step where the energy gradient is calculated. The shaded grey area shows cells that cannot be selected as a candidate secondary maximum, due to either being skipped over or being the first step from the seed.	67
3.17	Results for calculating E_{ratio} with different or no exclusion region definitions. Plots show background rejection as a function of signal efficiency for each tested exclusion region (left) and background rejection at 95% signal efficiency as a function of exclusion region size (right).	68
3.18	Plot of background rejection at 95% signal efficiency as a function of pseudorapidity, η , for E_{ratio} algorithms with different exclusion regions.	68
3.19	Results for calculating E_{ratio} after varying the search limit parameter, given as a distance in η from the seed cell. Plots show background rejection as a function of signal efficiency for each tested search limit (left) and background rejection at 95% signal efficiency as a function of the search limit (right).	69
4.1	Efficiencies for identification of photons in Run 2 of the ATLAS detector, as a function of transverse energy. Shown are the values for loosely-isolated unconverted photons in the central region of the detector. The efficiency is evaluated from data using three techniques, as indicated. [73]	78

4.2	Efficiencies for identification of electrons in Run 2 of the ATLAS detector, as a function of transverse energy. Loose, medium, and tight working points for electron identification are shown in different colours. [75]	80
4.3	Efficiencies for reconstructing and identifying muons in Run 2 of the ATLAS detector, as a function of transverse momentum. Loose, medium, and tight working points for muon identification are shown in different colours. [76].	81
4.4	Distribution of events as a function of m_{ll} and $m_{ll\gamma}$, for events in the signal sample passing the lepton and photon cuts given in Table 4.4. The dashed line shows the threshold for the FSR-rejection cut, events below the dashed line are discarded.	85
4.5	Example BDT score distribution, from training and test data used for the semileptonic $VZ\gamma$ analysis. Signal and background events are as defined in Section 6.2.	89
4.6	Schematic of the four regions used for fake photon background estimation. Region A represents the signal region; B, C, and D represent control regions obtained by relaxing isolation and/or identification requirements.	92
5.1	Feynman diagram for a $Z\gamma$ vector-boson scattering event (left). The black circle contains the multiboson interaction, which for a leading-order SM interaction will be one of the two shown (right).	105
5.2	Example Feynman diagrams for non-VBS EW production of $Z\gamma jj$. In these instances one or none of the two final-state bosons are produced through multiboson interactions.	106
5.3	Example Feynman diagrams for QCD production of $Z\gamma jj$.	106
5.4	Impact of four of the key jet cuts on the analysis yield when applied individually with particle-flow and topo-cluster jets. Events are divided into four categories based on whether or not they pass the particle-flow cut and whether or not they pass the topo-cluster cut. Cuts are applied in the same order as presented in Table 5.3. Only events that pass the previous cut for both particle-flow and topo-cluster jets are included in the results, to decorrelate the effects of each individual cut. The label ‘PFlow’ is used for particle-flow cuts and ‘Topo’ for topo-cluster cuts.	111
5.5	Distribution of events passing jet requirements for the particle-flow collection but not the topo-cluster collection. In each case the cut is in the same distribution as the histogram plotted. Four cuts are shown: $N_j > 2$ (top left), $p_T^{j,2}$ (p_T of second most energetic jet) > 50 GeV (top right), $m_{jj} > 150$ GeV (bottom left), $ \Delta\eta_{jj} > 1$ (bottom right). Only events passing all prior cuts for both particle-flow and topo-cluster jets are included. The label ‘PFlow’ is used for variables calculated with particle-flow jets and ‘Topo’ for topo-cluster jets.	112

5.6	Distribution of the $\min \Sigma(\Delta R)$ variable, defined in Equation 5.2, for events passing $m_{jj} > 150$ GeV for particle-flow but not topo-cluster jets (left); and the m_{jj} distribution for these events after requiring $\min \Sigma(\Delta R) < 0.5$ (right). The label ‘PFlow’ is used for variables calculated with particle-flow jets and ‘Topo’ for topo-cluster jets.	113
5.7	Centrality distribution for data, signal, and background estimates pre-fit (before any data corrections to QCD $Z\gamma jj$). The dashed line marks the separation between the SR and QCD CR. The uncertainty band is the combination of uncertainties from background estimation, MC statistics, and experimental systematics. Overflow events are included in the last bin. [66]	114
5.8	Plots of relative variation of yields as a function of dijet mass, m_{jj} , for EW (left) and QCD (right) production of $Z\gamma jj$. The variations shown are the largest groups of systematics in the SR. [66]	115
5.9	Gluon fractions and uncertainties as a function of jet pseudorapidity and transverse momentum. Plotted are gluon fraction in the nominal (a) and alternate (b) QCD $Z\gamma jj$ sample MC samples, the difference between gluon fractions in these two samples (c), the statistical uncertainty on gluon fractions in the alternate sample (d) (this was the larger of the statistical uncertainties), and the difference between gluon fraction in the SR+CR region and in the SR (e) and CR (f) regions.	117
5.10	Binned m_{jj} distribution of measured systematic uncertainty, σ_{Norm} , for one systematic variation as labelled on the plot. The dashed line shows the zeroth-order polynomial fit attempted, with the indicated χ^2 value demonstrating this is clearly a poor assumption and this systematic does contribute an uncertainty on the m_{jj} shape.	121
5.11	χ^2 values, representing the impact each systematic has on shape, for all experimental systematic variations in the EW (top) and QCD (bottom) samples in the SR. Only sources with χ^2 above 7 for either the up or down variation are shown. The top axis gives the probabilities of uncertainties arising fluctuations under the null hypothesis. The largest uncertainties extend beyond the range of the x -axis. A description of what each systematic uncertainty represents is given in Table 4.5 (some names are abbreviated here).	122
5.12	Overall normalisation uncertainty, σ_{Norm} , in the SR for the EW (top) and QCD (bottom) samples. Values for the largest experimental systematic uncertainties are shown, any with a value below 5×10^{-4} are omitted. The black bars represent the statistical uncertainty on the value. A description of what each systematic uncertainty represents is given in Table 4.5 (some names are abbreviated here).	123
5.13	Post-fit distributions of dijet mass, m_{jj} , in the SR (top) and CR (bottom). The uncertainty band is the combination of all uncertainties, taken from the fit. Overflow events are included in the last bin. [66]	125
6.1	A selection of SM production mechanisms for the $VZ\gamma$ triboson final state, depicted in Feynman diagrams.	128

6.2	Kinematic distributions, comparing EW $VZ\gamma$ production (red) to QCD $Z\gamma jj$ production (blue), generated from the corresponding Monte Carlo (MC) samples with the analysis region selection applied. Events are normalised to compare the shape of distributions between the two samples. Definitions for the variables shown are given in Table 6.3.	134
6.3	Distributions to identify a cut on $p_T^{j,2}$. Shown are fraction of events for each sample that are above a given threshold value in $p_T^{j,2}$ (left) and background rejection as a function of the signal efficiency achievable using the same $p_T^{j,2}$ threshold (right).	138
6.4	Distribution of centrality, $\zeta(l\ell\gamma)$, both linear (left) and logarithmic (right) scales on the x -axis. Normalised event counts are shown for the $VZ\gamma$ signal sample and the QCD $Z\gamma jj$ background, for events in the analysis region.	141
6.5	Distribution of dijet mass (left) and signal rarity (right) for events in the analysis region, defined by the selection in Table 6.1. Yields for the signal process and all backgrounds are shown stacked. The error band represents the total pre-fit uncertainty from all sources of systematic uncertainty.	143
6.6	Approximate significance calculated with s/\sqrt{b} for a number of signal events s and background events b with an m_{jj} value between the minimum value given on the x -axis and the maximum on the y -axis. The number of signal events is calculated from the EW $VZ\gamma$ sample and all background samples are included for the background estimate. Events are required to pass the analysis region selection from Table 6.1. The maximum significance is obtained for a cut of $70 < m_{jj} < 100$ GeV.	144
6.7	Illustration of the three m_{jj} regions used in the analysis: the lower sideband ($30 < m_{jj} < 65$ GeV), the upper sideband ($110 < m_{jj} < 150$ GeV), and the peak ($70 < m_{jj} < 100$ GeV) region which is then subdivided into the SR and BDT CR. The distributions shown represent events passing the analysis region selection for both the signal (shown in green) and the sum of all backgrounds (in black). Both distributions are normalised by their total event yield.	145
6.8	Pruning results for systematic uncertainties in the $VZ\gamma$ analysis. The colours indicate whether a systematics shape and normalisation uncertainty components were each retained for the fit or dropped, for each sample and region used in the fit.	149
6.9	Post-fit signal rarity distributions in each of the four regions used in the fit, as labelled. Uncertainty bands represent the combined uncertainties in each bin, with values constrained by the fit. Uncertainty on data is due to statistics. The lower sections of each plot give the ratio of the data to the total background estimate.	152
6.10	Systematic uncertainties ranked by their post-fit impact on μ_{EW} . Uncertainties labelled γ represent MC statistics uncertainties in the given bin.	153
6.11	Systematic uncertainties ranked by their post-fit impact on μ_{EW} , for a projected fit scaled to a luminosity of 420 fb^{-1} . Uncertainties from MC statistics and jet flavour were removed from this fit.	154

Definitions of Acronyms

ATLAS	A Toroidal LHC Apparatus.....	25
BCID	bunch crossing identifier.....	55
BDT	boosted decision tree.....	85
CERN	the European Laboratory for Particle Physics.....	22
CMS	Compact Muon Solenoid.....	106
CP	Cluster Processor.....	36
CPM	Cluster-processor Module.....	56
CR	control region.....	92
CSC	cathode strip chamber.....	33
CTP	central trigger processor.....	35
DSID	dataset identifier.....	73
eFEX	Electromagnetic Feature Extractor.....	41
EMEC	electromagnetic end-cap.....	29
EM	electromagnetic.....	10
EW	electroweak.....	10
FCal	forward calorimeter.....	29
FEX	Feature Extractor.....	41
FPGA	field-programmable gate array.....	36
FSR	final-state radiation.....	83
GEP	Global Event Processor.....	44
gFEX	Global Feature Extractor.....	41
HEC	hadronic end-cap.....	29
HL-LHC	High-Luminosity LHC.....	23
HLT	High-level Trigger.....	34

ID inner detector	27
JEP Jet/Energy Processor	36
jFEX Jet Feature Extractor	41
L1Calo Level-1 Calorimeter	2
LAr liquid argon	29
LATOME Liquid Argon Trigger Optical Mezzanine	55
LEP Large Electron-Positron Collider	22
LHC Large Hadron Collider	22
LO leading order	15
MC Monte Carlo	18
MDT monitored drift tube	33
MS muon spectrometer	32
NLO next-to-leading order	16
NNLO next-to-next-to-leading order	16
OR overlap removal	83
PDF parton distribution function	16
PoI parameter of interest	148
PPM pre-processor module	36
PV primary vertex	78
QCD quantum chromodynamics	10
QED quantum electrodynamics	7
QFT quantum field theory	3
QGC quadrilinear gauge coupling Often called ‘quartic gauge couplings’ in literature.	14
RoI region of interest	36
RPC resistive-plate chambers	33
SCT semiconductor tracker	27
SM Standard Model	1
SR signal region	91
TDAQ trigger and data acquisition	35

TGC	trilinear gauge coupling <i>Often called ‘triple gauge couplings’ in literature.</i>	14
TGC	thin gap chamber <i>when used in Chapter 2.</i>	33
TOB	trigger object.....	42
TRT	Transition Radiation Tracker	27
UI	user interface	45
VBS	vector-boson scattering	1

Introduction

Elementary particle physics is the study of the mechanics of nature at the most fundamental scale. The field is guided by a theory, the Standard Model (SM), describing a handful of elementary particles that account for almost the entirety of known matter and interactions in the universe. At the same time, experiments in the field have been growing larger and larger over recent decades, in attempts to create the higher energy environments needed to probe the small distance scales on which these particles operate.

This thesis presents a number of contributions to the upgrade and research programmes of, to date, the largest particle physics experiment in history, the ATLAS experiment. The focus of this work is on analysing rare electroweak processes in the SM sensitive to multiboson interactions, a feature of the theory that are challenging to observe experimentally due to their low rates. Measurements of these processes serve as a valuable test of the SM, determining whether these unique interactions are translated from experiment to theory.

Two analyses are presented in this thesis to probe these interactions, and both analyses rely on the same final state: $Z(\rightarrow ll)\gamma jj$. The first analysis interprets the pair of jets as a product of a vector-boson scattering (VBS) event, featuring a $2 \rightarrow 2$ scattering between electroweak bosons. The second analysis treats the jets as products of the hadronic decay of a third boson, either a W or another Z boson, implying a triboson final state. These processes lead to two analyses with some convenient overlap in methodology, but a unique set of challenges for each. The diagrams producing these final states share the same multiboson interactions, and so these two analyses are probing the same underlying physics.

Meanwhile, a separate body of work is presented on studies and tools made to help with the upgrade programme for the Level-1 Calorimeter (L1Calo) trigger. This includes a tool to visualise trigger algorithms and help to debug discrepancies between software and firmware codebases, a study of early Run-3 data to analyse performance of the trigger in commissioning, and development of a software prototype of a future firmware algorithm to improve trigger performance in upcoming runs.

The reader will be introduced to some of the necessary concepts from theory in Chapter 1 and to the experimental setting, the collider and detector, in Chapter 2. Chapter 3 presents the work on L1Calo trigger upgrades, also providing further information about the context and timeline of the upgrades. The VBS $Z\gamma$ analysis is presented in Chapter 5 and the semileptonic $VZ\gamma$ analysis in Chapter 6. Meanwhile, some of the shared methods for the two analyses and some additional background is given before this in Chapter 4.

Theory

As with the study of science as a whole, the field of particle physics relies on two pillars: experiment and theory. This thesis focuses on developments in experimental particle physics, but one cannot be discussed without the other. This opening chapter gives the theoretical background needed to put the experimental work in context, exploring the mathematical origin of the interactions being studied and introducing the concepts needed to perform an analysis on data from particle physics collisions.

An introduction to the concepts of quantum field theory is given in Section 1.1 and an overview of some of the necessary mathematical transformations and symmetries in Section 1.2. Section 1.3 then introduces specific theories describing different components of particle interactions, leading to the introduction of the Standard Model of particle physics in Section 1.4. Some practicalities of making predictions in real experimental environments are then discussed in Sections 1.5 and 1.6.

1.1 Quantum field theory

The theoretical description of elementary particle physics is built on the foundations laid by quantum field theory (QFT). This mathematical framework describes

particles as excitations in quantised fields, and the nature of these fields governs the interactions between particles.

QFT is the simplest theoretical description encompassing the principles of both quantum mechanics and special relativity [1]. Attempts at relativistic quantum wave mechanics, such as the Dirac theory [2], failed to explain the mechanics of antiparticles; with the theory relying on the Pauli exclusion principle [3] preventing ‘regular’ particles from falling into negative energy states, and thus working only for fermions and not bosons [1, p.14]. QFT solves this problem and others by introducing a quantum field, in which particles and antiparticles can be created and annihilated; the creation and annihilation of particles represents the interactions that QFT describes. The promotion of the wave function to a field gives a natural description for many-particle systems. This better equips QFT to describe real states in nature as, at the microscopic level, there are no true one-particle systems [4].

Many different QFTs can be formulated, e.g. to describe the interactions of different forces. The Lagrangian density, \mathcal{L} , is typically used to define the dynamics of a given QFT. Lagrangian densities are a necessary tool to describe many-particle systems, related to the Lagrangian, L , by

$$L = \int d^3x \mathcal{L}.$$

The terms Lagrangian and Lagrangian density will be used interchangeably for the remainder of this chapter.

There is no guarantee that the ‘theory of everything’ that we need to describe the fundamentals of nature is a QFT. It is true, however, that any relativistic quantum theory applied to particles at sufficiently low energy will look like a QFT [1]. Even if the theory of everything is not a QFT, it is clear from observations that modern particle physics is still in a realm where energies are sufficiently low (relatively speaking) that QFT is accurate as an effective field theory. This has been demonstrated by the success of the SM, discussed in Section 1.4.

1.2 Symmetries and transformations

1.2.1 Lie groups

Lie groups, and their corresponding Lie algebras, are used to define the symmetries obeyed by a given theory. Lie groups represent a set of transformations that can be applied to a state, and are used to represent symmetries in theories when transformations between these states should be invariant. Two types of group are prominent in particle physics theory: unitary and special unitary groups. A unitary group of degree n , denoted $U(n)$, is the infinite group of all unitary $n \times n$ matrices under matrix multiplication. A special unitary group of degree n , $SU(n)$, is a subgroup of the corresponding unitary group and contains all $n \times n$ matrices with a determinant of 1. An $SU(n)$ group has $n^2 - 1$ members, or ‘generators’.

Of interest to the theories discussed here are the groups $U(1)$, $SU(2)$, and $SU(3)$. The $U(1)$ group contains all complex numbers with a magnitude of 1; a $U(1)$ transformation is equivalent to a change in complex phase. As complex numbers commute, $U(1)$ forms an ‘Abelian’ (i.e. commutative) group.

The $SU(2)$ group contains three 2×2 matrices, T^a , which may be expressed in terms of the Pauli spin matrices, σ^a , as $T^a = \sigma^a/2$. The generators of $SU(2)$ are non-commuting, with the commutator

$$[T^a, T^b] = i\varepsilon^{abc}T^c,$$

where ε^{abc} is the totally antisymmetric Levi-Civita tensor; the $SU(2)$ group is therefore non-Abelian.

More generally, the commutator for generators of an $SU(n)$ algebra is given by

$$[T^a, T^b] = if^{abc}T^c, \tag{1.1}$$

where f^{abc} is a totally antisymmetric tensor specifying the structure constants of the Lie algebra. For the $SU(2)$ definition given above, $f^{abc} = \varepsilon^{abc}$.

The third group of interest is $SU(3)$; this group is also non-Abelian and contains 8 generators, which here are 3×3 matrices. A typical basis for the generators of $SU(3)$ would give structure constant values

$$\begin{aligned} f^{123} &= 1, \\ f^{147} = f^{246} = f^{257} = f^{345} &= -f^{367} = -f^{156} = \frac{1}{2}, \\ f^{458} = f^{678} &= \frac{\sqrt{3}}{2}, \end{aligned}$$

with all other elements equal to zero [5].

1.2.2 Gauge transformations

A gauge, in theoretical particle physics, is an abstract frame of reference, typically used when a system is invariant under a change in this reference frame, known as a gauge transformation. These transformations lead to a deeper symmetry required in theories: if a gauge transformation is allowed within a theory then it must not affect the observables of the theory, thus the theory is considered symmetric under these transformations. This gauge symmetry is a core concept of modern QFTs, and such gauge-symmetric QFTs are known as gauge theories.

A gauge transformation can be local or global. A global gauge transformation is where the parameter controlling the transformation is constant across space-time. Local transformations are a more general case where the parameters are a function of space-time coordinates and can vary between locations.

The set of gauge transformations under which a theory is symmetric form a Lie group. For each generator in the Lie group a gauge field is introduced. Inclusion of these gauge fields in the Lagrangian ensure that the theory is invariant under gauge transformations. If a theory is local gauge invariant, these gauge fields can vary across space-time and allow for interactions between particles in the theory; this is how forces are introduced in gauge theories [6, pp.242-3]. The quanta of a gauge field in a QFT is called a gauge boson.

1.3 Gauge theories

1.3.1 Quantum electrodynamics

Quantum electrodynamics (QED) describes electromagnetic interactions between charged particles. As a gauge theory respecting local $U(1)$ transformations it introduces a single massless gauge boson, the photon. When formulated as a QFT, quantum electrodynamics (QED) must therefore describe fermions, photons, and the interactions between them. The QED Lagrangian can be built from three terms:

$$\mathcal{L}_{\text{QED}} = \mathcal{L}_{\text{Dirac}} + \mathcal{L}_{\text{EM}} + \mathcal{L}_{\text{int}}.$$

The $\mathcal{L}_{\text{Dirac}}$ term describes the kinematics of a fermion under Dirac theory [2]. This is given by

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

where ψ is a Dirac spinor, a four-component fermion field representing up-down and particle-antiparticle states for a fermion of mass m . The γ^μ are a set of matrices accounting for fermion spin, these are commonly absorbed into the covariant derivative using the Feynman slash notation, $\gamma^\mu\partial_\mu \rightarrow \not{\partial}$.

Maxwell's equations provide the terms \mathcal{L}_{EM} and \mathcal{L}_{int} , describing the kinematics of the photon and their interaction with charged fermions:

$$\begin{aligned}\mathcal{L}_{\text{EM}} &= -\frac{1}{4}F^{\mu\nu}F_{\mu\nu} = -\frac{1}{4}(F_{\mu\nu})^2, \\ \mathcal{L}_{\text{int}} &= -J^\mu A_\mu = -q\bar{\psi}\gamma^\mu\psi A_\mu.\end{aligned}$$

Here A_μ is the electromagnetic vector potential; $F_{\mu\nu}$ is the electromagnetic field tensor, given by

$$F_{\mu\nu} = \partial_\mu A_\nu - \partial_\nu A_\mu; \tag{1.2}$$

and $J^\mu = q\bar{\psi}\gamma^\mu\psi$ is a conserved current, satisfying $\partial_\mu J^\mu = 0$, for a fermion of charge q .

Combining these, and simplifying by defining a gauge covariant derivative

$$D_\mu = \partial_\mu - iqA_\mu, \quad (1.3)$$

gives the full QED Lagrangian:

$$\mathcal{L}_{\text{QED}} = \bar{\psi}(i\not{D} - m)\psi - \frac{1}{4}(F_{\mu\nu})^2. \quad (1.4)$$

This result could instead be obtained by starting from the Dirac Lagrangian and enforcing local gauge invariance through the covariant derivative transformation in Equation 1.3 [7, pp.482-6]. The form of Equation 1.4 is recovered with the inclusion of the \mathcal{L}_{EM} term, which is the only locally gauge invariant formulation of a kinetic term for the field A_μ . Its invariance can be demonstrated from the local gauge invariance of the commutator $[D_\mu, D_\nu]$, given

$$\begin{aligned} [D_\mu, D_\nu] &= iq(\partial_\mu A_\nu - \partial_\nu A_\mu) \\ &= iqF_{\mu\nu}. \end{aligned} \quad (1.5)$$

This technique for deriving field tensors from potentials will be relevant in discussion of other theories.

1.3.2 Yang-Mills theory

The QED theory corresponds to a $U(1)$ gauge symmetry, and as such is Abelian. Constructing a non-Abelian gauge theory respecting $SU(n)$ symmetries is more complex, but generically solved by the Yang-Mills theory [8].

For generators of the Lie algebra T^a and structure constant f^{abc} , a gauge covariant derivative can be defined by

$$D_\mu = \partial_\mu - igT^a A_\mu^a,$$

where a vector field A_μ^a is required for each generator of the $SU(n)$ group, and g is a coupling constant. Here the exponents a, b, c index the generators of the Lie algebra, whilst μ, ν index space-time dimensions, as per convention.

This gauge covariant derivative is a generalisation of the Abelian form, given in Equation 1.3, and acts on an n -plet, ψ , of spinors ψ_i , rather than on a single spinor as in the QED case. The generators of $SU(n)$ serve to transform ψ_i into one-another through abstract rotations.

In analogy to the Abelian case in Equation 1.5, the commutator is used to define a set of field strength tensors: the commutator

$$[D_\mu, D_\nu] = -igF_{\mu\nu}T^a$$

holds for a field strength satisfying

$$F_{\mu\nu}^a T^a = \partial_\mu A_\nu^a T^a - \partial_\nu A_\mu^a T^a - ig[A_\mu^a T^a, A_\nu^b T^b].$$

This shows explicitly that for an Abelian symmetry group, with $[T^a, T^b] = 0$, the form of Equation 1.2 is recovered. For non-Abelian gauge theories, however, the additional nonlinear term is introduced, featuring a product of gauge field potentials. This represents self-interaction of non-Abelian gauge fields.

Substituting in the general form for the commutator of $SU(n)$ generators given in Equation 1.1, the general form of a field strength for such a non-Abelian gauge theory is given by

$$F_{\mu\nu}^a = \partial_\mu A_\nu^a - \partial_\nu A_\mu^a + gf^{abc}A_\mu^b A_\nu^c. \quad (1.6)$$

This leads to the Yang-Mills Lagrangian [7, pp.486-91], the most general renormalisable Lagrangian for a theory respecting $SU(n)$ symmetry expressed in terms of these field strengths,

$$\mathcal{L}_{\text{YM}} = \bar{\psi}(i\not{D} - m)\psi - \frac{1}{4}(F_{\mu\nu}^a)^2. \quad (1.7)$$

1.3.3 Quantum chromodynamics

The strong force, or quantum chromodynamics (QCD), is modelled by a non-Abelian gauge theory. The fundamental particles that QCD acts on are quarks and thus a QCD charge, known as colour charge, is attached to quarks. Colour charge is analogous to the electric charge in QED, but to describe the observed dynamics of the strong force three different types of charge are needed, termed red, green, and blue. These each have ‘anti’ varieties (anti-red, anti-green, anti-blue) when attached to antiquarks.

Transformation of a quark from one colour to another acts as a gauge symmetry in the theory; if all red quarks became green, green became blue, and blue became red the predictions of the theory would remain unchanged. This rotation of colour charge is described perfectly by the $SU(3)$ symmetry.

QCD can therefore be constructed as a Yang-Mills theory under $SU(3)$, with the Lagrangian following that of Equation 1.7. Given that the $SU(3)$ group has 8 generators, the theory of QCD relies on 8 gauge bosons to mediate interactions. These bosons are called gluons, and the 8 varieties are represented as different colour-anticolour combinations of the gluons.

It can be shown that the coupling constant (appearing in Equation 1.6) for QCD is actually not a constant, and is dependent on the energy scale of interactions [9]. For high energies (i.e. small distance scales) the coupling strength tends to zero, leading to ‘asymptotic freedom’. Quarks can only exist as free particles in the high energy limit, whereas at and below energies of ~ 1 GeV the coupling strengths are sufficiently high that quarks are exclusively confined in colourless composite particles, hadrons.

1.3.4 The electroweak theory

The electroweak (EW) theory describes the weak interaction, and also encapsulates QED as it unifies the weak and electromagnetic (EM) forces. It combines the

local $U(1)$ invariance of QED with the symmetry of the weak interaction under local $SU(2)$ transformations, and thus is described by $SU(2) \times U(1)$ symmetry. The gauge covariant derivative for the theory is given by

$$D_\mu = \partial_\mu - igT^a W_\mu^a - \frac{1}{2}ig'B_\mu,$$

where a indexes the 3 generators of the $SU(2)$ algebra, T^a , and their corresponding gauge fields, W_μ^a . The field B_μ is required for invariance under $U(1)$ symmetry.

The corresponding field strengths are then

$$\begin{aligned} W_{\mu\nu}^a &= \partial_\mu W_\nu^a - \partial_\nu W_\mu^a + g\varepsilon^{abc}W_\mu^b W_\nu^c \\ B_{\mu\nu} &= \partial_\mu B_\nu - \partial_\nu B_\mu \end{aligned}$$

To treat this as a conventional Yang-Mills theory would result in four massless gauge bosons from the three W_μ^a fields and the B_μ field. This does not match experimental observations however, as the weak bosons are known to have non-zero masses [10, 11].

The missing ingredient is spontaneous EW symmetry breaking through the Brout-Englert-Higgs mechanism[12, 13]. Without symmetry breaking, mass terms for the W and Z bosons cannot be included in the Lagrangian whilst also satisfying local gauge invariance. The Brout-Englert-Higgs mechanism introduces a complex scalar field with a non-zero vacuum state, the Higgs field. This vacuum state can take one of an infinite number of values, but once it takes a specific value it spontaneously breaks the global gauge symmetry of the theory. Reference [6] shows how perturbatively expanding the field about this non-zero vacuum expectation value produces mass terms for three of the four electroweak gauge fields; all whilst conserving the theory's *local* gauge symmetry. This mechanism also introduces a massive scalar boson associated with the Higgs field, the Higgs boson.

Briefly, when the Higgs field takes a non-zero vacuum expectation value, a rotated set of the potentials W_μ^a and B_μ can be constructed in the Lagrangian, resulting in

gauge fields:

$$\begin{aligned} W_\mu^\pm &= \frac{1}{\sqrt{2}}(W_\mu^1 \mp iW_\mu^2), \\ Z_\mu &= \frac{1}{\sqrt{g^2 + g'^2}}(gW_\mu^3 - g'B_\mu), \\ A_\mu &= \frac{1}{\sqrt{g^2 + g'^2}}(g'W_\mu^3 + gB_\mu). \end{aligned}$$

The W_μ^\pm and Z_μ fields have mass terms in the Lagrangian, with masses parameterised by

$$m_W = \frac{gv}{2}, \quad m_Z = \frac{\sqrt{g^2 + g'^2}v}{2},$$

where v is the vacuum expectation value of the Higgs field [7].

This converts the previous form of the theory with four massless bosons to a description with three massive bosons, the two charged W bosons and the neutral Z boson, and one massless boson, which is identified as the photon.

Due to the non-Abelian $SU(2)$ symmetry group used to build this theory, self-interaction terms appear in the Lagrangian, as explained in Section 1.3.2. The resulting interactions are discussed in Section 1.4.

1.4 The Standard Model of particle physics

The SM of particle physics is a gauge QFT which combines all of the theories discussed in Section 1.3 into a single theoretical description. As a result, the SM respects a symmetry of $SU(3) \times SU(2) \times U(1)$, a combination of the QCD and EW theories.

The SM models the interactions of 12 fermions (and 12 antifermions), mediated by 12 gauge bosons and an additional scalar boson, the Higgs. The 12 fundamental fermions are split into six leptons and six quarks, each can be paired up across three generations. The leptons come in charged lepton-neutrino pairs in electron, muon, and tau families. There are three generations of up-type and down-type quark pairs:




Figure 1.1: All fundamental particles described by the Standard Model shown with their masses, or limits on masses, measured from experiments. Particles are grouped into quarks, leptons, and bosons. [14]

up-down, charm-strange, and top-bottom. These fermions all have different masses, although the values of the masses are not derivable from the theory.

The 12 gauge bosons are those introduced by the imposed symmetries: eight gluons from the $SU(3)$ QCD sector and 4 electroweak bosons (W^+ , W^- , Z , γ) from the $SU(2) \times U(1)$ electroweak sector. The gluons and photon are observed to be massless and the W and Z bosons are massive, as discussed in Section 1.3.4. The additional boson, the Higgs boson, is also massive, although, as with the fermions, the theory does not explicitly constrain its mass. All 17 varieties of fundamental particle (grouping the eight gluons and 2 W bosons) are shown in Figure 1.1.

Of particular interest to analyses presented in this thesis are EW direct multiboson interactions. These are interactions introduced in the Lagrangian involving multiple EW bosons; as explained in Section 1.3, these arise from the non-Abelian construction of the electroweak sector. These interactions involve either three or four bosons at a single vertex, termed trilinear gauge couplings (TGCs) and quadrilinear gauge couplings (QGCs) respectively. In the SM, there are two TGC vertices, W^+W^-Z and $W^+W^-\gamma$, and four QGC vertices, $W^+W^-W^+W^-$, W^+W^-ZZ , $W^+W^-\gamma\gamma$, and $W^+W^-Z\gamma$ [6, p.541][15]. Figure 1.2 shows examples of these as Feynman diagrams.




Figure 1.2: Example SM multiboson interactions: a three-boson vertex (left) and a four-boson vertex (right) are shown.

Whilst the SM has been very successful so far when its predictions are compared to experimental observations, it does not fully describe elementary particle physics. One missing piece is the fourth fundamental force, gravity. The current best theory of gravity, general relativity, is not quantisable and thus incompatible with the QFT structure of the SM. At sufficiently high energies, the SM description of physics will break down since it does not account for the effects of gravity.

Other signs point to the SM being merely an effective theory, a low-energy approximation of some more complete theory. The SM has 25 parameters [6], such as the fermion masses, with values that have to be constrained by experiment rather than being dictated by the theory itself. Moreover, there appears to be some structure linking sets of parameter values [6], perhaps indicating that a more fundamental theory exists to explain these patterns.

Contemporary experimental particle physics is dedicated to both testing predictions

of the SM and searching for signatures of physics beyond it. Either by finding evidence of a more fundamental theory, or by identifying a breakdown in the SM description, a significant goal in the field is to ‘break’ the SM.

1.4.1 Cross sections

For each process predicted by the SM, its cross section, σ , can be calculated from the theory. The cross section is a measure of the probability for the process to occur; cross sections have the dimension of area and are typically quoted in units of ‘barns’, where $1 \text{ b} \equiv 10^{-28} \text{ m}^2$.

A useful tool in calculating the cross section of a process is the Feynman diagram. A theory defines all of the allowed interactions between particles, which can be interpreted as the set of allowed vertices in Feynman diagrams (e.g. the vector gauge boson vertices in Figure 1.2). An interaction taking a certain initial state to given final state can then proceed through all mechanisms that can be drawn as valid Feynman diagrams.

The cross section for an interaction is proportional to the amplitude squared of the transition matrix element of the interaction. This matrix element is the sum of the individual matrix elements for each mechanism through which the process can proceed, i.e. for each Feynman diagram. Matrix elements are calculated from Feynman diagrams via the Feynman rules [16].

In principle, an infinite number of Feynman diagrams can be drawn for any process, as more intermediary vertices can always be added. However, as more vertices are added and the process becomes more complex, the associated probability becomes smaller. Cross sections are therefore calculated perturbatively: contributing Feynman diagrams are considered up to a certain order to allow for a finite set of matrix elements to be included.

The simplest calculation for the cross section of a process is at leading order (LO). This involves all diagrams with the fewest number of vertices possible to get between

the desired initial and final states. Allowing additional diagrams which are only one step more complicated (e.g. adding two more vertices through a loop) gives a next-to-leading order (NLO) calculation, including a further set of diagrams makes a next-to-next-to-leading order (NNLO) calculation, and so on.

The rate of experimental observations of a process will depend on its cross section, the total number of occurrences expected for a process in data is given by

$$N = L\sigma,$$

where L is the integrated luminosity, a measure of the size of the dataset introduced in Section 2.1. Processes with low cross sections, such as the subjects of the analyses in Chapters 5 and 6, therefore occur very infrequently and can be difficult to measure.

A selection of SM processes are shown in Figure 1.3, with measurements and predictions of their ‘fiducial’ cross sections. A fiducial cross section describes the rate a process in a specific phase space, and is typically measured because analyses are limited to making measurements only in regions sensitive to the process of interest. Cross sections relevant to VBS $Z\gamma$ production and $VZ\gamma$ production are given in the figure.

1.5 Proton-proton collisions

In collider physics, observable interactions between SM particles are induced by colliding particles in a controlled environment. In the case of the LHC (introduced in Section 2.1), these collisions are between two high-energy protons.

As protons are composite particles, interactions due to proton-proton collisions are initiated by constituent partons in the protons. Modelling these interactions requires knowledge of the fraction of the total proton momentum carried by its partons. The distribution of these momenta is given by parton distribution functions (PDFs), which are determined experimentally from deep inelastic scattering measurements [18].




Figure 1.3: Cross sections for SM processes, with measurements from the ATLAS experiment shown alongside predictions from theory. Cross sections are corrected for branching fractions and subject to fiducial phase spaces of the analyses. From the processes discussed in this thesis, VBS $Z\gamma$ (labelled as $Z\gamma_{jj}$ EWK) and $WZ\gamma$, one component of $VZ\gamma$, are included. [17]

Interesting interactions are typically produced when two partons with high momentum fractions collide, this is known as a hard scatter. In many collisions this will not occur, and only soft low-energy interactions take place. Even when there is a hard scatter it will be surrounded by ‘spectator interactions’, soft collisions between other partons in the colliding protons. This complicates measurements made from proton-proton collisions.

A further complication is the presence of pileup. In order to increase the rate of hard-scatter events, recent colliders are configured to create multiple proton-proton collisions at once. This results in many, typically soft, interactions being produced around any hard scatter that is detected, termed ‘pileup interactions’.

These effects all complicate the procedure of measuring and understanding events from proton-proton collisions. The methods used to make practical measurements of SM processes under such conditions are discussed in the following section.

1.6 Monte Carlo predictions

The SM is tested by comparing its predicted cross sections for a set of measurable physics processes with the rate observed in data. To isolate processes of interest, this is often done in complicated phase spaces and differentially across distributions, so calculating the predicted rate is quite complex. Real measurements are also subject to the limitations and effects of the detector; for accurate comparisons to be made between data and predictions, these need to be accounted for.

The solution to this problem is to generate Monte Carlo (MC) events representing the SM prediction. For a given final state (i.e. the set of particles produced in the interaction) a sample of events is generated for each contributing process, containing particles with random kinematic properties generated in such a way that the overall distribution matches the expectation from the model. If all processes are accounted for, taking the sum of events from all of these samples in the desired phase space gives an estimate for the SM prediction.

Production of MC samples is a very complicated process and there are several implementations commonly used in the field, known as MC generators. Generators discussed in this thesis include MADGRAPH [19], SHERPA [20], PYTHIA [21], POWHEG [22], and HERWIG [23].

The first step of the process is simulating the hard scatter. These rely on matrix element calculations, at a given perturbative order, and PDF sets in order to accurately simulate the desired processes at the given centre-of-mass energy.

The hard scatter alone cannot mimic a full event in the detector, and so several additional steps are needed: parton showering, applied to any strongly interacting particles; hadronisation, converting these showers into composite hadrons; adding the ‘underlying event’, activity expected in the collision from sources other than the hard scatter; pileup overlay, to account for the number of simultaneous proton-proton interactions; and detector simulation, accounting for effects of particles being measured by the detector.

For each parton from the hard-scatter process, a shower of QCD activity is produced from repeated strong interactions. This continually creates more, lower energy, partons until the energies reach a regime where confinement effects become relevant. Confining the shower products into colourless hadrons is handled by a hadronisation model, such as string fragmentation [24]. Both parton showering and hadronisation are incorporated into MC event generators. Some generators, such as SHERPA, PYTHIA, or HERWIG, can simulate the hard-scatter process, parton showering, and hadronisation all in one. In other cases the hard-scatter process is created with one generator, e.g. MADGRAPH, and then another generator is used to add parton showering and hadronisation to it, e.g. PYTHIA or HERWIG.

The models used to simulate parton showering and hadronisation are not calculable from first principles, and have parameters that can be adjusted to best describe observed physics. These parameters include merging and resummation scales, describing the merging of jets from the parton shower and the hard scatter event and the resummation of soft gluon emissions [20].

Underlying event is a term used to describe activity around the hard scatter in collisions, such as spectator interactions (introduced in Section 1.5). Modelling the underlying event is dependent on the specific conditions under which the collisions occur. This is mitigated by tuning certain parameters of the model to match data; different ‘tunes’ are available from data collected under different conditions. For MC samples, underlying event modelling is typically handled by the same generators as the parton showering and hadronisation.

Pileup overlay describes the process of adding additional soft events around the hard scatter to simulate the presence of additional proton-proton collisions. The number of pileup events added can be configured to match the number of proton-proton collisions expected per bunch crossing in the detector.

Having simulated the hard scatter, parton showering, hadronisation, the underlying event, and pileup, the particles and their kinematics should be established. The remaining step is determining how this event would be detected in an experiment, if it were from a real collision. For the ATLAS detector (introduced in Section 2.2) this is done using GEANT4 [25, 26].

A distinction is made between information from MC samples before and after detector simulation. Simulated events before detector simulation is applied are known as ‘truth’ events, they contain only the physics processes and are not subject to any inefficiencies or misidentification of the simulated detector. Samples or variables with truth information are often described as ‘truth-level’ (or ‘particle-level’). Events completing the full simulation, and subsequent reconstruction (see Section 4.3), chain are known as reconstructed MC events. These are typically linked so that the truth properties (or ‘truth record’) of reconstructed MC events are accessible. Analyses are typically performed using reconstructed MC events, but in certain cases making distinctions based on truth information is necessary.

Events simulated in the manner discussed here are used extensively for the two analyses presented in this thesis (as well as for studies presented in Chapter 3). The signal process and all backgrounds in the final state have dedicated MC simulation

samples generated in order to model the kinematics of events. Generation of these specific samples is discussed in Section 4.2.

The ATLAS detector at the Large Hadron Collider

2.1 The Large Hadron Collider

The Large Hadron Collider (LHC) is a circular¹ particle collider, measuring 27 km in circumference, located at the European Laboratory for Particle Physics (CERN) in Geneva, Switzerland [27, 28, 29]. As a successor to the Large Electron-Positron Collider (LEP) [30], the LHC was designed to study higher energy systems than had previously been accessible in controlled, high-rate, collisions. One of the key goals of the LHC was discovering the Higgs boson, a goal which was achieved in 2012 [31, 32, 33]. However, the LHC research programme is much more broad than the search for and study of the Higgs boson; many aspects of the SM are investigated to find signs of inconsistency between theory and experiment, to measure SM parameters, and to explore untested regions of the SM.

The LHC has periods of operation known as runs. Each run consists of multiple years of data-taking, with some short shutdown periods for maintenance and minor upgrades. Between each run is a ‘long shutdown’ period, in which more significant upgrades can take place. Analyses in this thesis use data taken during Run 2, between 2015 and 2018, with a centre of mass energy for collisions of 13 TeV. Run

¹Roughly circular, since the ring consists of alternating straight and curved sections.

3 began in 2022, with an increase in centre of mass energy to 13.6 TeV, and is currently ongoing at the time of writing. After Run 3, a larger set of upgrades is planned and future runs will be at much higher luminosities; this period is known as the High-Luminosity LHC (HL-LHC) [34].

Controlled interactions are created by colliding accelerated beams of protons at interaction points. Two beams of protons, travelling in opposite directions around the LHC ring, are accelerated to an energy of 6.5 TeV ². Creating collisions between these two beams at certain interaction points on the ring results in proton-proton interactions at a centre-of-mass energy of 13 TeV. Experiments are built around the interaction points to observe the results of these high-energy interactions.

Protons are obtained by ionising hydrogen gas with an electric field. A chain of many accelerators is then used to take the initial at-rest protons up to an energy of 450 GeV, when they are injected into the LHC. This injector chain is shown, amongst other CERN accelerators, in Figure 2.1. Once in the LHC, protons are further accelerated to the desired beam energy of 6.5 TeV. This acceleration, both in the injectors and the LHC machine itself, is performed using superconducting radio-frequency cavities; these are electromagnetic fields that accelerate protons as they pass through.

Superconducting magnets are used to bend and focus the beam. A total of 1232 dipole magnets, with a field strength of 8.3 T, are used to bend the path of the beam, as required by the circular design of the collider. Additionally, 392 quadrupole magnets are placed around the LHC to focus the beam, squeezing the protons together to make the profile of the beam more compact [27].

Protons are injected into the LHC in bunches, with approximately 10^{11} protons in a single bunch. Consecutive bunches are injected with a minimum separation of 25 ns; this is referred to as a ‘bunch train’ when many bunches are used at this minimum separation.

To create collisions between the two proton beams, insertion magnets are used to cross the paths of the beams [36]. Each colliding pair of bunches, one from each

²Numbers given here correspond to the Run-2 parameters of the machine.

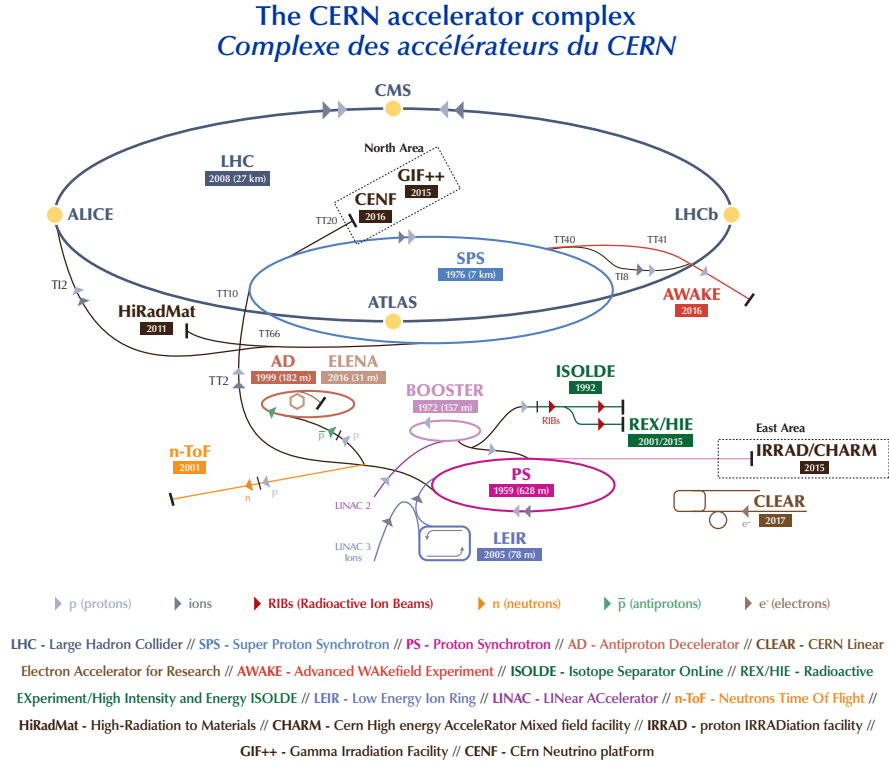


Figure 2.1: Schematic of the CERN accelerator complex. The chain of accelerators used to inject protons into the LHC is LINAC2→BOOSTER→PS→SPS [35].

beam, is called a ‘bunch crossing’. The magnets can be adjusted in order to change the crossing angle, modifying the expected number of proton-proton collisions induced for each bunch crossing.

The rate at which collisions occur in the LHC is given by the instantaneous luminosity, \mathcal{L} , defined as [37]

$$\mathcal{L} = \frac{N_p^2 n_b f_{\text{rev}} \gamma}{4\pi \epsilon_n \beta^*} F,$$

where N_p is the number of particles per bunch, n_b is the number of bunches per beam, f_{rev} is the revolution frequency of the beam, γ is the relativistic Lorentz factor, ϵ_n and β^* parameterise the optics of the beam, and F is a factor describing the crossing angle of the two beams. The design luminosity for the LHC is $\mathcal{L} = 1 \times 10^{34} \text{ cm}^{-2}\text{s}^{-1}$, and throughout Run 2 the machine operated between around 0.5 to 2 times this amount [38].

Integrated luminosity, L , is used to measure the amount of data in an entire dataset.

This quantity is the integral of instantaneous luminosity over time,

$$L = \int dt \mathcal{L},$$

and is typically measured in units of fb^{-1} . The expected number of occurrences of a particular process is given by the product of the integrated luminosity with the cross section of the process, σ . Obtaining a large dataset is therefore vital to measure processes with very low cross sections. The LHC produced a dataset of 160 fb^{-1} over the entirety of Run 2 [38].

2.2 The ATLAS detector

ATLAS (A Toroidal LHC Apparatus) is one of the four large detectors built around interaction points at the LHC. The ATLAS detector is the largest of the four, and designed as a general-purpose detector to measure as many different processes as possible. In order to do this, it surrounds the interaction point almost entirely. Full angular acceptance would in principle allow any event to be fully reconstructed from its detected decay products, the goal is to get as close to this as is reasonably possible. The detector itself is built from several sub-detectors, each dedicated to measuring specific properties of particles, aided by a system of superconducting magnets. Each of these sub-detectors is discussed in detail below. Figure 2.2 gives an overview of the ATLAS detector and its components.

2.2.1 Coordinate system

A specific coordinate system is used to describe the ATLAS detector and interactions within it. The z -axis runs along the beamline; the x -axis points, in the positive direction, towards the centre of the LHC ring; and the y -axis points vertically upwards. The azimuthal angle, ϕ , is measured around the beamline in the x - y plane and the polar angle, θ , is measured from the beam axis. Transverse momentum, p_T ,




Figure 2.2: Cut-away view of the ATLAS detector. Dimensions and components of the detector are labelled. [39]

and transverse energy, E_T , are used to define the momentum or energy in the x - y plane, transverse to the beamline.

It is useful to define the rapidity, y , of a particle,

$$y = \frac{1}{2} \ln \frac{E + p_z}{E - p_z},$$

where E is the energy of the particle and p_z is its momentum in the z direction. Rapidity, like θ , gives a measure of how ‘forward’ a particle is in the detector (the forward regions are the high-rapidity, or high- θ , regions). Rapidity transforms additively under Lorentz boosts along the beam axis, so differences in rapidity are invariant under such boosts. However, since the rapidity of the particle is dependent on its energy, it is difficult to use as a more general coordinate. A pseudorapidity coordinate, η , is defined as

$$\eta = -\ln \tan(\theta/2).$$

Pseudorapidity is equivalent to rapidity in the relativistic limit, and so holds for high-energy low-mass objects. Additionally, pseudorapidity has the benefit of being

a simple translation of θ , and so can be used not just for particles but to describe locations in the detector.

The set of coordinates p_T , η , and ϕ are typically preferred to describe the kinematics of objects in the detector as all three are invariant under Lorentz boosts along the beamline³. Without this Lorentz invariance, differences in kinematics between events could be introduced depending on the relative momenta of the colliding partons.

Angular differences between objects are typically expressed in terms of the quantity $\Delta R = \sqrt{\Delta\eta^2 + \Delta\phi^2}$.

2.2.2 Inner detector

The innermost detector system, known as the inner detector (ID) or tracker, contains three sub-detectors designed to track the trajectory of charged particles, to measure their momentum based on the curvature of their tracks. This is enabled by a solenoid magnet which surrounds the ID, generating a 2 T magnetic field coaxial with the beam direction in order to bend the tracks of charged particles travelling through.

As particles travel through the components of the ID, ‘hits’ are registered for each location where the particle is detected. Hits across the tracker are fitted to reconstruct the track of the particle. The momentum of this particle is calculated from the radial arc of this track, and the sign of its charge is deduced from the direction of the curve. Extrapolating the track towards its origin allows it to be associated with a specific collision vertex location on the beamline.

Figure 2.3 shows a cross section of a sector of the ID; the sub-detectors of which the tracker is composed are shown. From the beamline outwards, these are: the pixel detector, the semiconductor tracker (SCT), and the Transition Radiation Tracker (TRT). Each of these are detailed below. The pixel and SCT both cover an acceptance of $|\eta| < 2.5$ while the TRT has an acceptance of $|\eta| < 2$.

³The η coordinate itself is not invariant but differences in η , i.e. $\Delta\eta$, are invariant in the relativistic limit.




Figure 2.3: Diagram showing the components of the inner detector. The IBL label represents the insertable B-layer, the innermost part of the pixel detector which was added between Run 1 and Run 2. The r values label radial distances from the centre of the beam pipe. [40]

The closest component to the beamline is the pixel detector. The pixel detector is designed to measure particles as close to the beamline as possible, with very high granularity and precision. The detector is made up of 1968 silicon sensor modules, with a combined total of 8.6×10^7 pixels across all sensors [41]. The nominal pixel size is $50 \times 400 \mu\text{m}$ (in $r\phi \times z$) and $250 \mu\text{m}$ thick, with some variation in different regions [39]. As a charged particle passes through a pixel on the sensors it ionises the atoms in the silicon, creating electron-hole pairs; these charges are collected to generate a signal indicating a hit in that pixel.

The next component out, along the path of a particle, is the SCT. The SCT uses silicon strip sensors, which operate on the same principles as the pixels. These strips have typical dimensions of $80\mu\text{m} \times 6.4\text{cm}$ with a thickness of $285 \mu\text{m}$ [39]. The SCT

consists of four layers, where each layer has two sets of strips back-to-back with a relative angle of 40 mrad between the strips. The rotation between strips within a layer improves resolution along the long axis of the strip. In the barrel strips are placed with their long axis parallel to the beamline (in the z -direction) and in the end-caps strips are placed in the r -direction.

The last ID component encountered by incident particles is the TRT. The TRT is composed of 3.7×10^5 straw detectors, with a diameter of 4 mm and a length of 144 cm (barrel region) or 37 cm (end-cap region). In the barrel region straws are placed parallel to the beamline. In the end-cap region straws are arranged radially in wheels. The straw detectors contain a gold-plated tungsten wire surrounded by a xenon-carbon dioxide-oxygen gas mixture. The space between straws is filled with a polymer fibre. Charged particles passing through a straw can ionise the gas and generate a readout on the wire to give a hit. Additionally, charged particles crossing the boundaries between materials emit transition radiation, dependent on their $\gamma = E/m$. This transition radiation ionises xenon atoms in the gas mixture and gives a larger readout on the wire. The TRT therefore provides hits and also information on the E/m ratio of incident particles; this is used for particle identification, particularly for electrons.

2.2.3 Calorimeters

The ATLAS detector has two distinct calorimeter systems: the liquid argon (LAr) calorimeter and the tile hadronic calorimeter. These are both sampling calorimeters, employing alternating absorbing and active layers to induce and measure the energies of EM and hadronic showers, respectively. Showers can be reconstructed from a ‘cluster’ of energy deposits in the calorimeter, adjacent cells giving energy readouts to indicate that a particle deposited energy there. Figure 2.4 shows the location of the calorimeter components in the context of the detector.

The LAr calorimeter has four components, the barrel, the EM end-cap (EMEC), the hadronic end-cap (HEC), and the forward calorimeter (FCal). All of these




Figure 2.4: Cut-away view of the ATLAS calorimeters, with each component labelled. [42]

components use liquid argon in the active layers, where low-energy shower particles will ionise argon atoms and produce a charge which is collected in order to measure the deposited energy. The barrel and EMEC use lead for the absorbing layers, the HEC uses copper absorbers, and the FCal has a combination of copper and tungsten.

The LAr barrel and EMEC each have three layers of calorimeter cells of differing sizes. These sizes vary by region, but Figure 2.5 shows the layout in the centre of the barrel. There is also an additional ‘presampler’ layer in front of these three layers, to correct for energy loss due to material in front of the calorimeter. The full specification of the calorimeter cell granularity is given in Reference [39, p.9].

The first of these calorimeter layers in the barrel for $|\eta| < 1.4$ consists of strips of cells with very high η granularity. This layer should be at or near the start of EM showers induced by incident particles, and provides more precise determination of the shape of this shower. This layer continues into the EMEC for higher $|\eta|$ values but with decreasing η -granularity. This strip layer is an important part of the studies presented in Section 3.4. The second layer covers most of the radial depth of the




Figure 2.5: Diagram showing layout of calorimeter cells in the barrel of the LAr calorimeter. The X_0 units measure EM radiation lengths. [39]

calorimeter, and is designed to contain and measure most of the energy of an EM shower. The third layer provides additional measurements useful for high-energy EM objects or to help reject against hadronic showers.

The HEC and FCal provide forward coverage for measuring hadronic showers, and the FCal also extends the coverage for EM showers. The HEC consists of two wheels per end-cap, with each wheel split into two longitudinal sections. This gives the HEC four detection layers, with a granularity of between 0.1×0.1 and 0.2×0.2 (in $\eta \times \phi$) across its coverage. The FCal consists of three modules. The innermost module is for EM objects and uses copper absorbers around the active liquid argon layers. The remaining two modules of the FCal extend the hadronic coverage and use tungsten absorbers.

The barrel and EMEC combined give coverage for EM objects over $|\eta| < 3.2$. Combined with the inner FCal module, the full coverage of the EM calorimetry

is $|\eta| < 4.9$.

The tile hadronic calorimeter uses scintillating plastic tiles for the active layers and steel absorbing layers. The tile barrel and extended barrel combined give coverage over $|\eta| < 1.7$. The barrel and extended barrel are each divided into three layers longitudinally, and have a total thickness of 7.4 hadronic interaction lengths. Between the tile, HEC, and FCal, hadronic calorimetry acceptance is $|\eta| < 4.9$.

Energy resolution for a calorimeter has three contributions, given by

$$\frac{\sigma(E)}{E} = \frac{a}{\sqrt{E}} \oplus \frac{b}{E} \oplus c,$$

where a quantifies the stochastic effects of particle showers, b reflects the electronic noise in readout signals, and c gives a constant term due to miscalibration or detector instabilities. The overall energy resolution of a calorimeter therefore improves for higher energy showers. The design resolution of the ATLAS EM calorimeter system is [43]

$$\frac{\sigma(E)}{E} = \frac{10\%}{\sqrt{E}} \oplus \frac{170 \text{ MeV}}{E} \oplus 0.7\%.$$

2.2.4 Muon spectrometer

The outermost component of the ATLAS detector is the muon spectrometer (MS). Muons will typically generate tracks in the ID which would already allow their momenta to be measured. Adding additional tracking for muons in the outer part of the detector allows for rejection against decay-in-flight backgrounds, where a hadron might leave a track in the ID and then decay to a muon, but this muon can be rejected if the momentum measurement in the MS is incompatible with the original ID measurement. The MS also improves resolution of muon momentum measurements, and allows enhancements to identification and triggering for muons. The acceptance of the MS is $|\eta| < 2.7$.

A large toroidal magnet system is used to bend the tracks of muons passing through the MS, allowing tracking systems to measure their momentum. This is handled by a




Figure 2.6: Cut-away view of the ATLAS detector with muon spectrometer components labelled. [44]

barrel toroid, which wraps around the barrel MS systems to create a 0.5 T toroidal magnetic field, and two end-cap toroids, which generate a 1 T toroidal magnetic field.

Tracking information is primarily provided by the monitored drift tube (MDT) and cathode strip chamber (CSC) systems. The resistive-plate chambers (RPC) and thin gap chamber (TGC) detectors give additional tracking and also provide triggering capabilities. These detectors are arranged into layers called ‘stations’, with three stations stacked radially in the barrel region and three stations along the z -axis in each end-cap. A schematic of the full MS system is shown in Figure 2.6.

The MDTs provide tracking across the full acceptance of the MS, with multiple layers of MDTs in both the barrel and the end-caps. These function similarly to the straw detectors of the TRT, but with a larger diameter of 30 mm, containing a tungsten-rhenium wire surrounded by an argon-methane-nitrogen gas mixture.

In the most radiation-prone region, the inner section of the most central end-cap

layers, CSCs are used in place of MDTs. The CSCs cover pseudorapidities $2 < |\eta| < 2.7$. These are multiwire proportional counters with cathodes segmented into strips, and benefit from higher granularity than the MDTs.

The role of the MDTs and CSCs is to provide precision measurements of muon coordinates in the ‘bending’ plane of the magnets, i.e. the z or η coordinate. They achieve this with a precision of $< 100 \mu\text{m}$ [45].

The RPCs are placed in the barrel region and consists of two resistive plates separated by a 2 mm gas-filled gap. In the end-cap region, TGCs are used; these are similar to CSCs but designed with a faster readout suitable for triggering. Both the RPCs and TGCs are used to give real-time readout of track information for the Level-1 trigger. They also give a measurement of the ‘second coordinate’, the coordinate orthogonal that measured by the MDTs and CSCs. This gives the ϕ coordinate with a spatial precision of 5-10 mm [45].

2.2.5 Trigger and data acquisition

Running at maximum capacity, the ATLAS detector measures a bunch crossing every 25 ns, i.e. a rate of 40 MHz, and in the majority of collisions no ‘interesting’ physics happens. There is no realistic way to read out the data from every single bunch crossing at this rate, but even if this was possible it would create an impossibly large storage requirement for events that will probably never be used for physics studies.

The trigger is the solution to this problem; events are quickly processed to determine if they have any signatures that might indicate the presence of interesting physics processes. This is done in two stages. First, a low-level hardware trigger (the Level-1 trigger) to make very fast but loose selection on events, using coarse granularity information from a subset of detectors, reducing the input rate to at most 100 kHz. Then a high-level trigger (HLT) which uses more information and more complex reconstruction to reduce the rate further down to a few kHz.

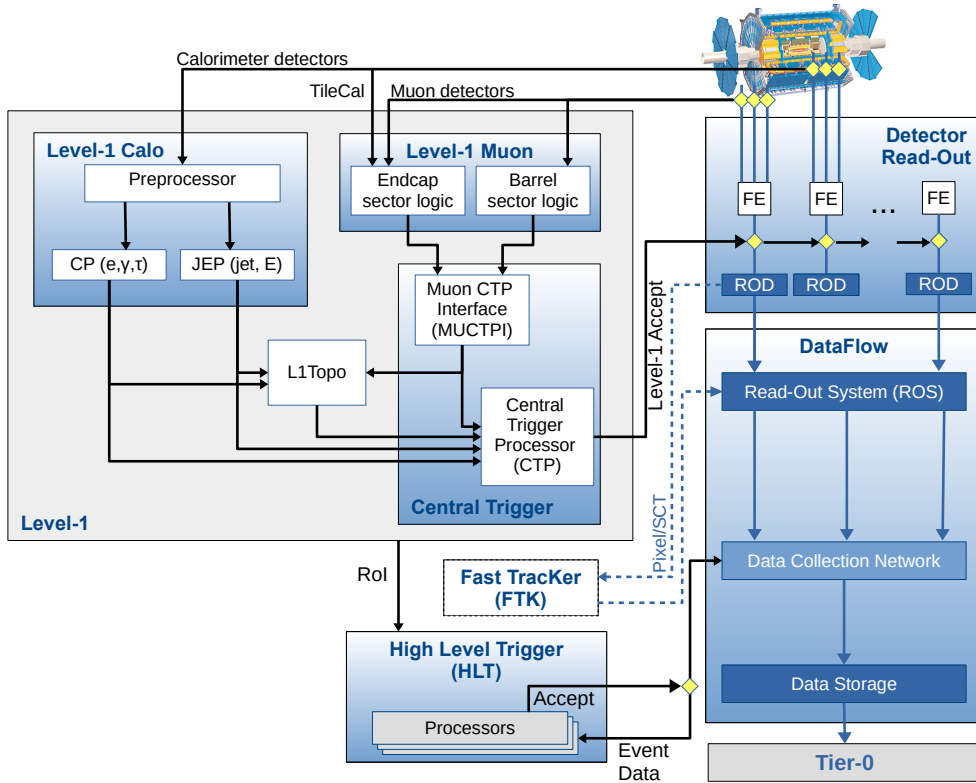


Figure 2.7: Diagram showing components and data-flow of the ATLAS Run-2 TDAQ system. This diagram includes the ‘Fast Tracker’ (FTK) component, but it was never implemented. [46]

To prescribe how events may be accepted and to manage the output rate, a trigger ‘menu’ is used; this gives the set of requirements for events to pass the trigger in a given run, for both Level 1 and the HLT.

The trigger works in tandem with the data acquisition system, which is responsible for reading out events passing trigger selections. This is done with front-end hardware read-out devices that collect the detector information and process it after receiving accept signals from the trigger systems. The entire trigger and data acquisition (TDAQ) system is summarised in Figure 2.7.

2.2.5.1 Level-1 trigger

The Level-1 trigger system is built of four main components: L1Calo, L1Muon, L1Topo, and the central trigger processor (CTP). Each of these are built from

bespoke hardware modules designed to perform the necessary algorithms as quickly as possible to keep up with the rate of input data, running on field-programmable gate arrays (FPGAs) to ensure a fixed latency is used.

L1Calo takes input from the calorimeters to give triggers for EM objects (electrons and photons) and jets. From both the LAr and tile calorimeters, the energies in each trigger tower (a 0.1×0.1 area in $\eta \times \phi$) are sent as analogue sums to a pre-processor module (PPM). The PPM digitises and sends these energies to two modules: the Cluster Processor (CP) and Jet/Energy Processor (JEP). The CP analyses a 4×4 area of trigger towers to calculate energies and isolations for e/γ or τ candidates. The JEP employs a similar process over larger areas and with lower granularity, in order to calculate energies of jet candidates and estimate missing transverse energy. For each event, L1Calo sends a set of ‘threshold bits’ to the CTP indicating multiplicities and energies of different objects with respect to trigger menu thresholds [47].

L1Muon uses tracking information from the RPCs and TGCs to make a rough estimate of the transverse momentum of muons. Any set of two or more hits that are consistent with a track originating at the interaction point are considered as candidate muons. Muon candidates are sent to CTP to contribute to the trigger decision for the event.

Taking input from both L1Calo and L1Muon, L1Topo calculates topological variables with a more holistic view of the event. These calculations include quantities such as invariant masses of, or angular separation between, multiple objects, and allow to trigger on more complex signatures.

The CTP takes inputs from L1Calo, L1Muon, and L1Topo and, based on the trigger menu, decides whether an event should be accepted. If an event passes the checks, a ‘Level-1 accept’ signal is sent to indicate to the data acquisition system that this event should be read out and sent to the HLT. For events passing the Level-1 trigger threshold, regions of interest (RoIs) are passed to the HLT in order to seed more complex trigger calculations.

2.2.5.2 High-level trigger

The HLT runs algorithms in software on a dedicated server farm. It is afforded looser timing constraints, due to the reduced input rate from the Level-1 trigger, and as such can run on traditional processors without fixed latency requirements. Algorithms are grouped into ‘chains’, with each chain seeded by a Level-1 RoI. Algorithms that require less processing time are typically run earlier in the chain to enable faster rejection of bad events.

Events passing the HLT selection are organised into ‘streams’ where each stream contains events that pass a set of trigger chains. These streams give events that pass the trigger and are saved for offline processing. The main stream (`physics_Main`) consists of events passing the trigger menu intended for physics analyses. A subset of accepted events are sent to an express stream which is sent for immediate offline reconstruction to test data quality and allow calibration updates.

2.2.6 Luminosity and pileup

Given the relationship between the cross section of a physics process and the luminosity of a dataset ($\sigma = N/L$), any cross-section measurement is dependent on the measured luminosity, and its precision. There are multiple methods of luminosity measurement employed by ATLAS, this includes the use of the LUCID detector [48], designed for the sole purpose of making such measurements.

The LUCID detector consists of two stations, each 17m along the beam pipe either side of the interaction point. Each station uses a set of Cherenkov tubes to detect protons displaced through inelastic scattering. The number of detected protons should be proportional to the number of interactions per bunch crossing, on average, and thus proportional to the integrated luminosity; this is calibrated using van der Meer scans. For the Run-2 integrated luminosity, the combined ATLAS measurement (including measurements from LUCID) has an uncertainty of 0.8% [50].




Figure 2.8: Integrated luminosity as a function of time for Run 2. Shown are the total luminosity delivered by the LHC, the luminosity recorded by ATLAS, and the amount satisfying requirements to be used for physics analyses. [49]

The total integrated luminosity recorded by ATLAS throughout Run 2 is shown in Figure 2.8. This marks the total luminosity delivered by the LHC, the amount of that recorded by ATLAS, and the amount which is ‘good for physics’. The differences in these amounts are due to down-time in the detector or its subsystems. Events are only marked good for physics if all systems were functional, within accepted tolerances, during data-taking. The deficit in ATLAS recorded luminosity from the total delivered represents inefficiencies in data acquisition.

One of the factors impacting the instantaneous luminosity recorded by the ATLAS detector is the number of proton-proton interactions per bunch crossing, denoted μ . The expected value of μ can be increased or decreased by adjusting the crossing angle of the beams. A greater number of interactions per bunch crossing results in an increase in luminosity but also gives an increase in pileup, since for every recorded event there are more ‘background’ collisions happening around it. The average number of interactions per bunch crossing in ATLAS throughout Run 2 was




Figure 2.9: Distribution of the mean number of proton-proton interactions per bunch crossing throughout Run 2. The distributions for each individual year of running are shown, as well as the total representing the whole of Run 2. Averages across these periods are also given. [49]

33.7 [49], although a range of values were used at different points as shown by Figure 2.9.

Upgrading the ATLAS Level-1 Calorimeter Trigger

3.1 Evolution of the Level-1 Calorimeter Trigger

The Level-1 Calorimeter (L1Calo) system, as used for Run 2 of the LHC in the ATLAS detector, is described in Section 2.2.5.1. This chapter focuses on work done towards upgrading this system for later LHC runs. In the long shutdown between Run 2 and Run 3, significant improvements were made to the system as part of the Phase-I upgrade. The long shutdown following Run 3 will facilitate further changes with the Phase-II upgrade, preparing L1Calo for the HL-LHC in Run 4.

Sections 3.2 and 3.3 discuss aspects of work on the Phase-I upgrade and Section 3.4 discusses work on the Phase-II upgrade, all with a focus on triggers for e/γ signatures. Details of the L1Calo system as implemented/planned for Phase I and Phase II, with the major changes in comparison to the Run-2 system and to each other, are given in Sections 3.1.1 and 3.1.2 respectively.

3.1.1 Phase-I upgrade

The Phase-I upgrade has progressed rapidly in the last three years, to the point where the system is fully implemented and in use in Run 3 at the time of writing.

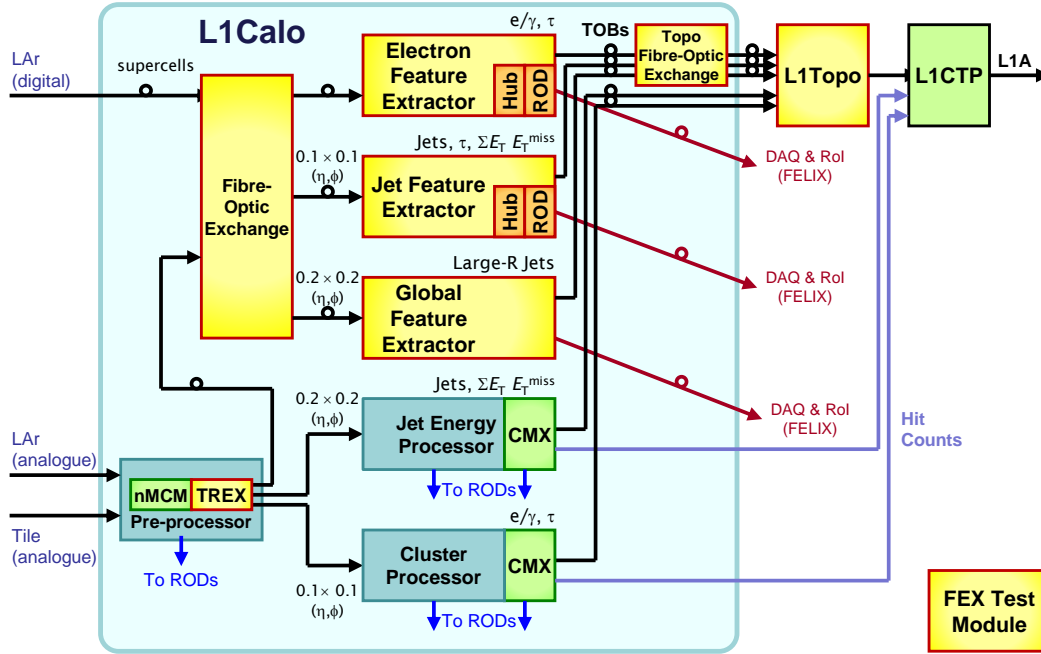


Figure 3.1: Diagram showing the L1Calo modules in use for Run 3 of the LHC. Yellow and orange rectangles represent modules introduced in the Phase-I upgrade. Blue and green rectangles represent existing components from the Run 2 system, included still as part of the trigger whilst transitioning to the new system. [51]

The Phase-I upgrade of L1Calo features a redesign of the core components of the trigger in order to process more data and make more refined decisions while rejecting background events.

The main components introduced in the Phase-I upgrade are the Feature Extractor (FEX) systems, which replace the CP and JEP from the Run-2 system. The Electromagnetic Feature Extractor (eFEX) provides discrimination for e/γ and τ objects, the Jet Feature Extractor (jFEX) focuses on jets whilst providing additional τ identification and an estimate for the missing energy, and the Global Feature Extractor (gFEX) triggers on large-radius jets and global quantities such as missing energy.

An overview of the Phase-I system architecture is shown in Figure 3.1. The key change to the e/γ -signature trigger is that it now receives digital information from the LAr calorimeter in the form of SuperCells, rather than the analogue tower energies that were available to the Run-2 system. This information is processed by




Figure 3.2: Diagram showing division of a single trigger tower into SuperCells. [52]

the eFEX to generate trigger objects (TOBs), equivalent to the regions of interest (RoIs) generated by the CP in Run 2.

The move to digital input comes with an increase in granularity, a trigger tower now being split in both η and calorimeter layer to give up to 10 SuperCells: typically one SuperCell each from Layers 0 (presampler) and 3 and four SuperCells each from Layers 1 and 2, segmented in η . Each SuperCell is formed by summing energies from between four and eight calorimeter cells. This division of a tower into SuperCells is shown in Figure 3.2. The granularity received from the Tile Calorimeter is the same as in Run 2, the summed energy in a tower (this is later referred to as Layer 4).

The eFEX introduces new algorithms to use the SuperCell information in order to trigger on e/γ objects. The CP in Run 2 calculated a cluster energy and a set of isolation values, the eFEX improves on this by calculating three more precise variables used to identify and distinguish a candidate object and can also more accurately calculate the energy from SuperCell information.

To calculate any of these variables a seed SuperCell is first identified, the highest energy SuperCell, compared to the surrounding region, in Layer 2 of the calorimeter.

The energy of the cluster is calculated by summing the energy of the seed with its highest energy neighbour in ϕ and both neighbours in η , adding also the six corresponding SuperCells in Layer 1, and the two SuperCells from Layers 0 and 3 that are in the same tower as the seed. The three other discriminating variables are calculated as follows:

$$R_\eta = \frac{\text{energy in } 3 \times 2 \text{ area of cells}}{\text{energy in } 7 \times 3 \text{ area of cells}},$$

with each area (in $\eta \times \phi$) centred on the seed and calculated in Layer 2 only;

$$R_{\text{had}} = \frac{\text{core energy}}{\text{environment energy}},$$

where the core energy is calculated in the same manner as for the cluster energy but including both neighbours in ϕ (so a 3×3 area of SuperCells in Layers 1 and 2 and a 1×3 area in Layers 0 and 3) and the environment energy is the energy in a 3×3 tower area in Layer 4 (i.e. the hadronic calorimeter); and

$$w_{s,tot}^2 = \frac{\sum_i i^2 \cdot E_i}{\sum_i E_i},$$

where $i \in [-2, 2]$ is the η coordinate of the SuperCell relative to the seed, calculated for SuperCells in Layer 1 with both neighbours in ϕ also summed for each E_i . Visual representations of the areas included for these algorithms can be found in Section 3.2.

Candidate τ particles considered by the eFEX have a similar set of variables (cluster energy, R_η , and R_{had}) with small differences in the areas used in their calculation.

3.1.2 Phase-II upgrade

As luminosity and pileup is increased even further with the high-luminosity era of the LHC in Run 4, the trigger again needs to be improved to operate in increasingly difficult conditions. The Phase-II upgrade to the hardware trigger aims to do this

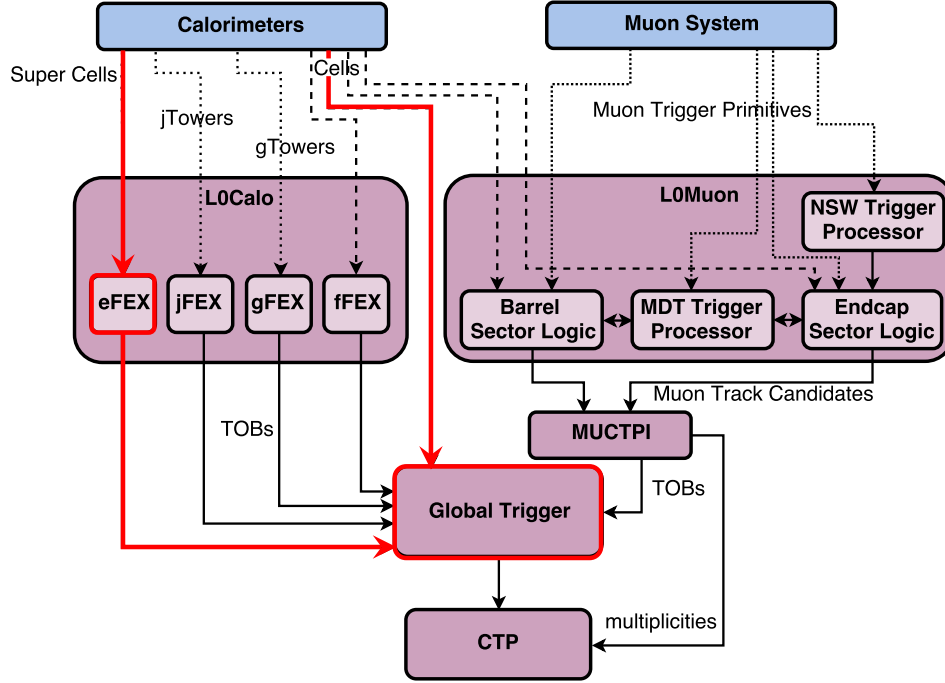


Figure 3.3: Schematic of the ATLAS hardware trigger as planned for the Phase-II upgrade in Run 4 of the LHC. The red lines highlight the main parts relevant to the e/γ trigger, with the addition of the Global Trigger being and the use of calorimeter cell information being the main changes with respect to the Phase-I system. [53]

primarily by adding a new component, the Global Event Processor (GEP) (or Global Trigger). The GEP will be downstream of the Phase-I FEX modules, which will continue to contribute to the trigger, and it will refine decisions made by employing additional information: information from a larger area than typically available to a single FEX and also finer in granularity.

An outline of how the GEP fits in with the existing systems is shown in Figure 3.3. Information from the calorimeters will be sent directly to the GEP in finer granularity than is available to the eFEX, with energies in each individual cell at the full detector granularity. This gives a 4-8 times increase in granularity over SuperCells, depending on the region of the calorimeter.

The additional information available to the GEP means it can work together with the eFEX to further refine the result. The eFEX will create TOBs with associated

variables (discussed in Section 3.1.1) which are sent to the GEP. The GEP can then further probe the same region of the calorimeter to determine if the candidate object should be accepted. The algorithms used by the GEP to do this are the topic of the study in Section 3.4.

3.2 Visualisation of eFEX inputs and algorithms

In order to aid in debugging minor differences between different implementations of eFEX algorithms, an algorithm visualisation program was created. The program reads input data and performs aspects of eFEX algorithms whilst also providing a visual representation of what the algorithm is doing and where the result comes from.

The visualiser is written in JavaScript, using Node.js [54] to interface with some server-side C++ scripts, which are needed to access energy decoders from trigger software, and Express.js [55] to handle the web-based user interface (UI).

3.2.1 Motivation

During development of algorithms for the hardware trigger, each algorithm is implemented multiple times. First, algorithms will be implemented in offline software to be tested and tuned against simulations or existing data. Then, in order to run on hardware, the algorithm needs to be ported to firmware. Often, to provide closer cross-checks of the firmware algorithms, they are also simulated in ‘online’ trigger software (hereafter referred to as online software). Inevitably, due to software and firmware bugs, subtle differences will exist between these algorithms; these differences need to be understood and corrected to ensure firmware works as intended and software is bitwise accurate in simulating the firmware.

The need to find these subtle differences between algorithms motivates the visualisation software discussed here. Although at first it seems illogical to add an additional

independent implementation of the algorithm (since the issue is in part due to having multiple different implementations), the added visualisation aspect makes it easier to understand where a particular algorithm implementation might have gone wrong in cases where there are discrepancies. This has been demonstrated through the use of the visualisation software in tests, discussed in Section 3.2.5.

3.2.2 Input data

The visualisation software takes as input the calorimeter energies visible to a single eFEX FPGA. This covers a 6×10 area (in $\eta \times \phi$) of trigger towers, with each tower being split into SuperCells across 5 calorimeter layers, as described in Section 3.1.1. For each event, the input data provides one energy value per SuperCell, with encoded energies.

From this input data, a 3×3 area of trigger towers, centred on an (η, ϕ) coordinate provided by the user, is extracted and displayed on-screen. This area covers all energy values used for eFEX algorithms if the seed of the TOB is located in the central trigger tower.

3.2.3 User interface

The eFEX Visualiser program provides a simple UI to explore input data and results of the eFEX algorithms. The basic interface is shown in Figure 3.4. It prompts the user to specify an input file, (η, ϕ) centre-tower coordinates, and an event number, then on receipt of these inputs it reads the information and displays the requested energies in a grid.

The full interface becomes visible after the grid is displayed. The grid itself is a 3×3 area divided by bold lines, with each segment representing a trigger tower, and each trigger tower square divided further into SuperCells. The horizontal axis represents the η coordinate of the tower or SuperCell, and the vertical axis represents the ϕ coordinate. These coordinates are labelled with the same indices the user gave as

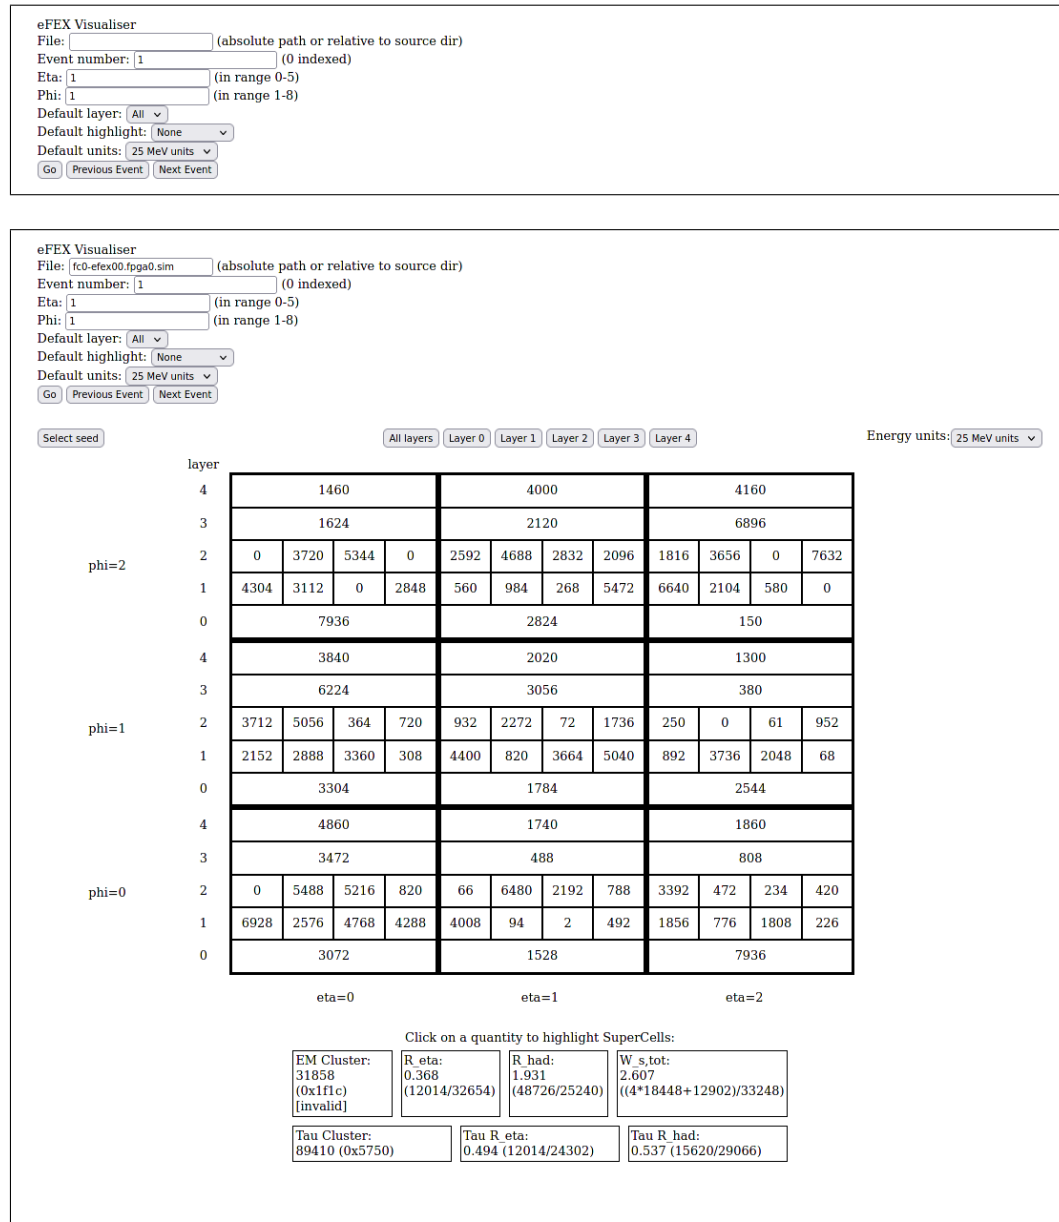


Figure 3.4: Initial interface on launching the eFEX Visualiser program (top) and the default view once data is read from a file (bottom).

initial input. In order to show all layers simultaneously, in the default view layers are stacked (in the ϕ -axis) on top of each other within each tower. Controls are provided to instead view each layer individually if preferred.

Below the grid, a list of all the quantities calculated for the current TOB is displayed. Clicking on one of these quantities will highlight all of the SuperCells involved in the calculation. The details of how these values are calculated and how the algorithms are visualised are discussed in more detail in Section 3.2.4.

Above the grid, alongside the layer selection buttons, are options to manually set the seed SuperCell and to select the units used to display energies. The unit selection input is a drop-down box that allows the user to choose between 25 MeV (default units in firmware) or GeV units. Changing this option instantly updates all displayed energies. Pressing the “Select seed” button will toggle the layer view to display Layer 2, prompt the user to click on the SuperCell with the highest energy, and then on its ϕ -neighbour with the highest energy. This aids the user in selecting the correct seed for TOB generation, but is not normally necessary as the program will apply these criteria to automatically set the seed as soon as the grid is loaded. The manual override is included in case the automatic selection is wrong, or if looking at algorithms with a different seed may help debugging.

3.2.4 Algorithms

For each TOB processed (i.e. each particular event, coordinate location, and seed), several algorithms are run to calculate the quantities displayed on-screen. These are the same algorithms used by the eFEX to calculate TOB energies and isolations. The following variables are calculated: EM cluster energy, EM R_η , EM R_{had} , EM $w_{s,\text{tot}}$, tau cluster energy, tau R_η , and tau R_{had} . All of these are either sums of SuperCell energies (EM and tau cluster energies), ratios of sums of SuperCell energies (R_η and R_{had}), or a ratio with weighted sums ($w_{s,\text{tot}}$).

The values of these variables are calculated immediately once the data for a given TOB is collected, or if the seed is re-specified, and displayed on-screen below the

grid. If the user clicks on a displayed quantity, the SuperCells involved in the sums for the corresponding algorithm are highlighted with colours denoting whether those cells are used in the numerator (lime green); the denominator (gold); or, in the case of $w_{s,tot}$, in the numerator with a larger weight (dark green). Figures 3.5 and 3.6 demonstrate the highlighting for all of the algorithms.

3.2.5 Usage

The visualisation tool was used at several stages during commissioning of the eFEX. Primarily it was used to compare firmware algorithms to their implementation in online software. Each time there was a difference found between the two, the event could be checked with the visualiser to help determine which of them was correct and to work out how the other might have gone wrong.

Once there was sufficient confidence in the similarity between online simulation and firmware implementations, the visualiser was again used to help in tests comparing online and offline simulations. The same technique was applied here to help ensure offline simulations were running with the same results as their online equivalent, and to help find errors in cases where they were not.

3.3 Analysis of early Run-3 data for commissioning

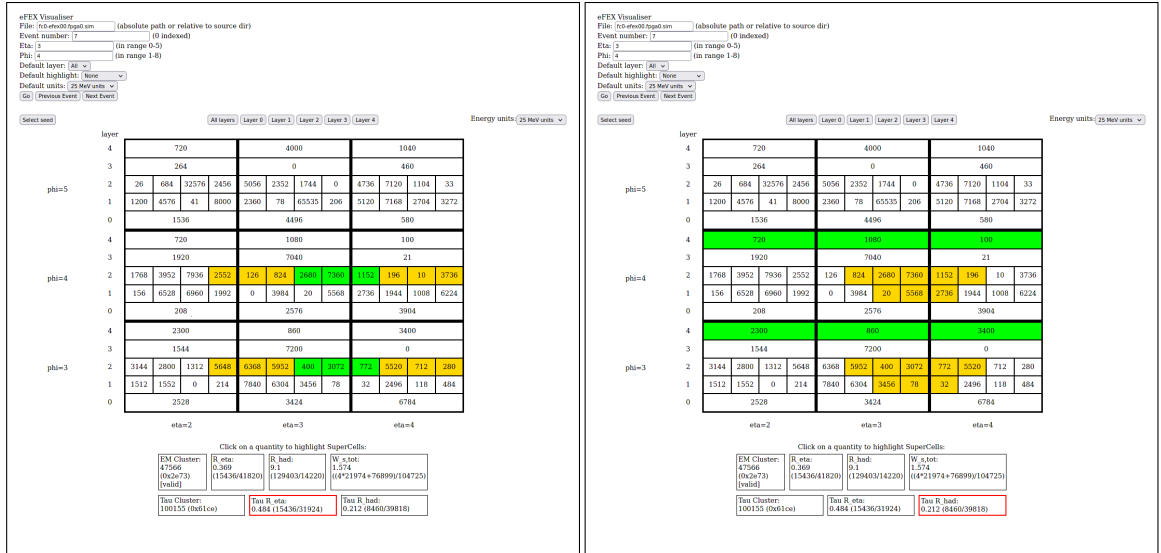
At the start of Run 3, the Phase-I L1Calo trigger was being used for the first time, having just been installed in the ATLAS detector. In these early stages, the new Phase-I system was running in parallel to the Run-2 system, with the Run-2 system being used in the trigger menu until the new system was fully commissioned.

One of the key goals in this time period was validating the Phase-I trigger system, comparing it to the Run-2 system to identify any differences which may have arisen from bugs or hardware issues. This section describes analysis of some early Run-3 data contributing to this goal.



(a)

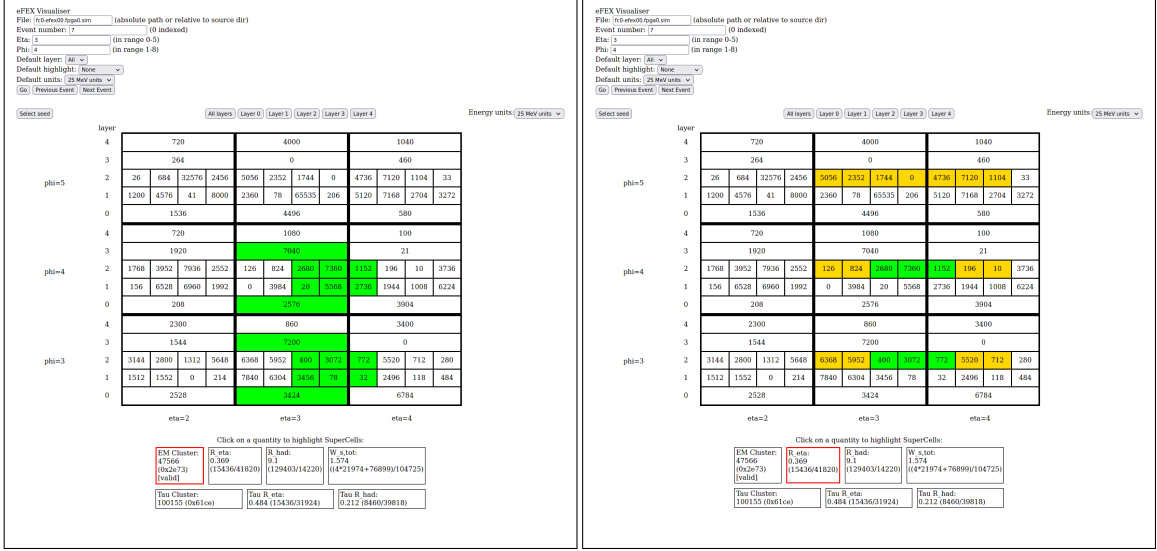
(b)



(c)

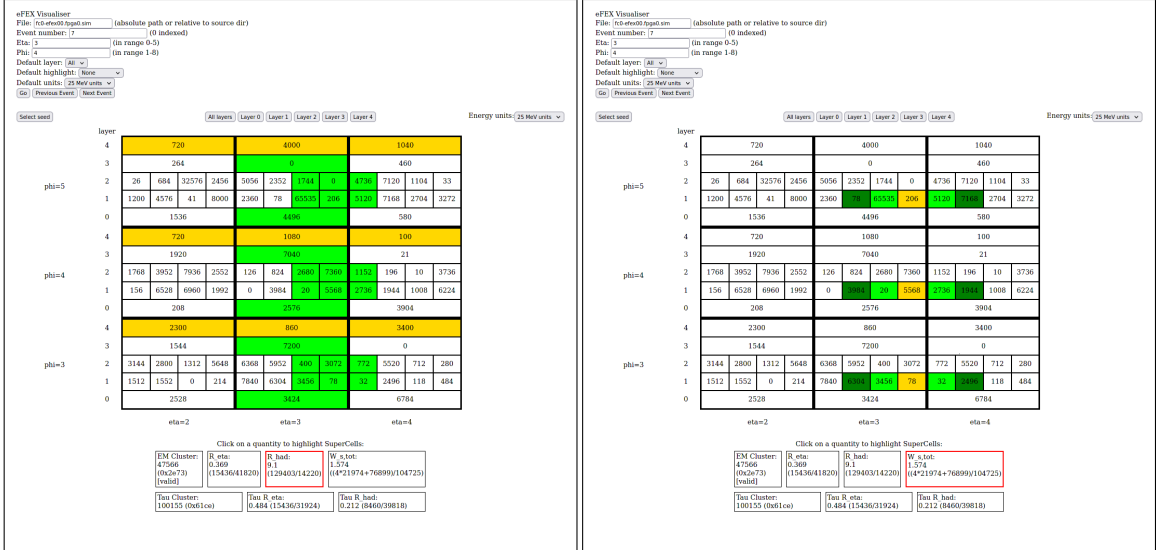
(d)

Figure 3.5: Demonstration of highlighting used to visualise algorithms, all shown for the same TOB. Showing (a) initial view without highlighting, (b) highlighting for tau cluster energy, (c) highlighting for tau R_η , and (d) highlighting for tau R_{had} .



(a)

(b)



(c)

(d)

Figure 3.6: Demonstration of highlighting used to visualise algorithms, all shown for the same TOB. Showing (a) highlighting for EM cluster energy, (b) highlighting for EM R_η , (c) highlighting for EM R_{had} , and (d) highlighting for EM $w_{s,tot}$.

3.3.1 Data

Two runs were used to provide the data for this analysis: Run 423433, taken on 31 May 2022, and Run 427885, from 10 July 2022. These runs were taken in quite different conditions, the first with lower intensity beams and no stable beam conditions, and the second with high intensity stable beams. Notably, the second of these runs had bunch trains with 25 ns separation between bunches, whereas the first had only isolated bunches.

Events are taken from the `physics_Main` stream. This stream contains 1,636,636 events for Run 423433 and 107,016 events for Run 427885.

3.3.2 TOB and RoI selection

Phase-I TOBs and Run-2 RoIs in events are compared to find instances in the same event that have the same, or very similar, η - ϕ coordinates. A pair is formed by selecting, for each TOB, the nearest RoI that has not already been matched to a TOB. A match is considered to be a pair of objects within ± 1 trigger tower in both η and ϕ , i.e. an RoI matches a TOB if it falls within the 3×3 area of trigger towers centred on the tower containing the TOB. Matched objects are considered to be the same physics object, identified independently by both systems. Instances where there is a TOB or RoI with no analogue in the opposing system are also tracked.

Only the barrel region was considered for this as a preliminary investigation, since it has a simpler geometry and as such it is easier to isolate bugs. At the time of analysing, only half of the eFEX modules were installed in the detector; this was due to delays in production caused by the global semiconductor shortage [56]. As a result, the Phase-I system at that time had coverage for just half of the ϕ range. Therefore only RoIs inside of this coverage are accepted.

3.3.3 Results

From the 1,636,636 events in Run 423433, 292,498 RoI/TOB pairs are selected. Of these, 271,854 matched in $\eta - \phi$ coordinates, giving a total match rate of 93%. For

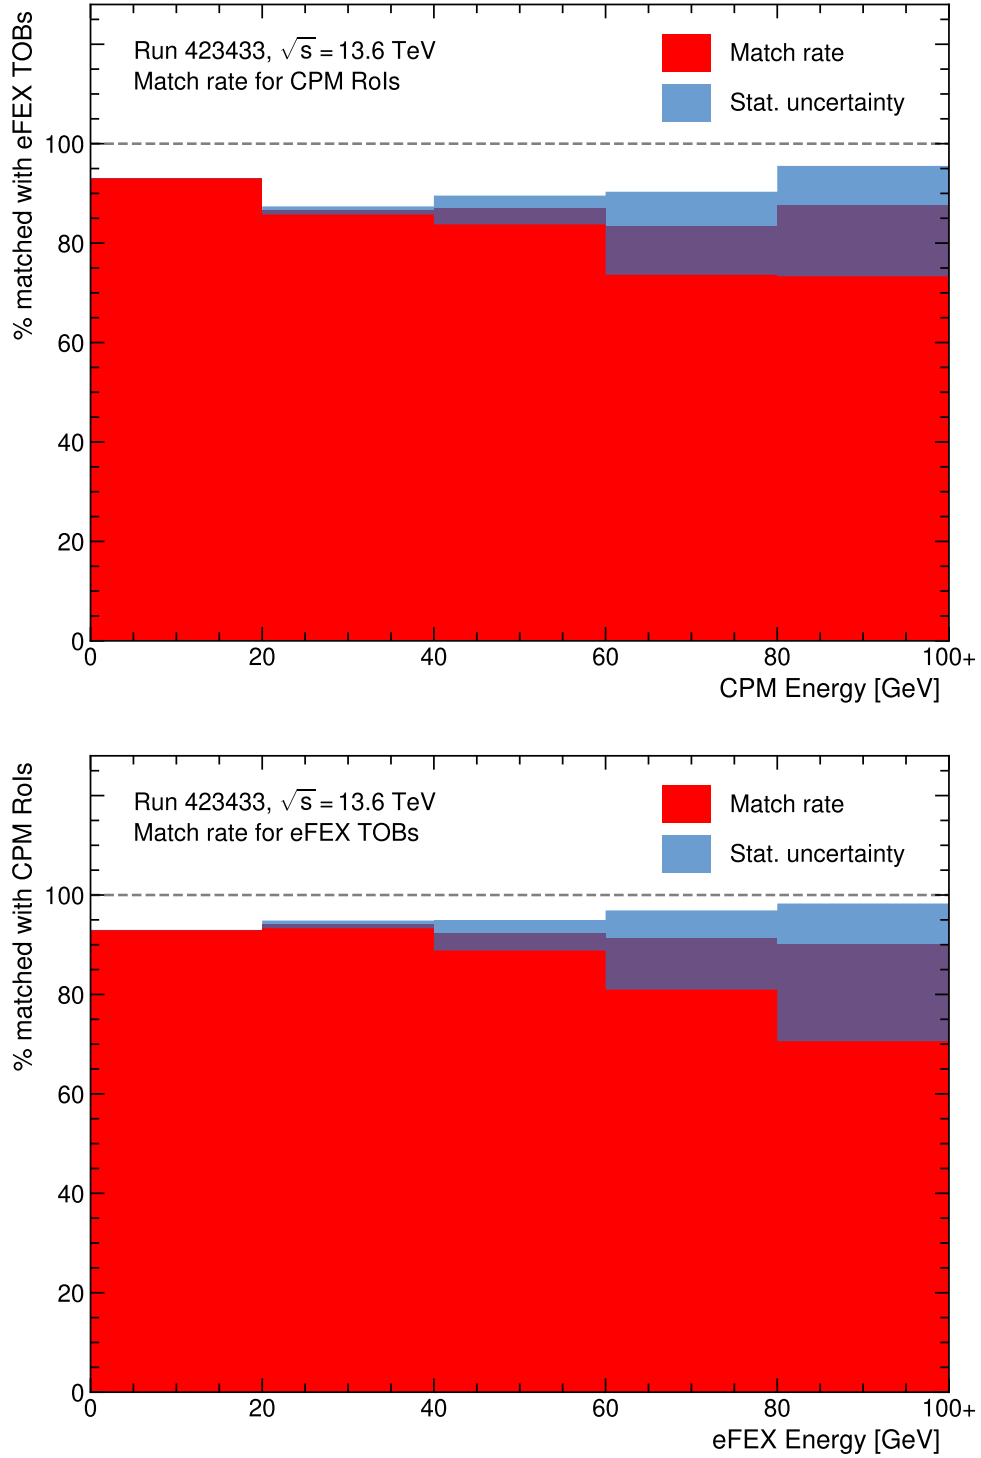


Figure 3.7: Match rate for TOBs/RoIs as a function of energy as measured by the CPM (top) and eFEX (bottom) for Run 423433. Objects grouped in 20 GeV bins, with the last bin including all overflow.




Figure 3.8: Match rate for TOBs/RoIs as a function of energy as measured by the CPM (top) and eFEX (bottom) for Run 427885. Objects grouped in 20 GeV bins, with the last bin including all overflow.

Run 427885, 22,337 of 27,973 pairs were matched for a match rate of 80%. Figures 3.7 and 3.8 show the fraction of objects matched as a function of energy, with histograms using both the CP-measured and eFEX-measured energies shown for each run. Uncertainties on these match rates are due to statistical uncertainties in the number of matched objects and the total number of objects, and are calculated using the Clopper-Pearson interval [57] with a confidence level of 68%.

For Run 427885 (Figure 3.8) it can be seen that the bulk of the mismatches come from low-energy objects, with a plateau in match rate above $E_T > 20$ GeV. This issue was not present in the earlier Run 423433 (Figure 3.7), which shows a relatively consistent match rate across all energies. This is reflected in the overall match rate, which is considerably lower in the later run.

Figure 3.9 compares the energies recorded by the Run-2 and the Phase-I systems for matched objects in the two runs. In Run 423433 it is clear that the majority of matched objects have approximately the same energy, with an additional cluster where in a few cases the eFEX-measured energy is much lower than the CP-measured.

In the later run, Run 427885, however, there is no longer such a strong correlation in energies. It seems that in general the eFEX energies are lower than the CP energies – seen by the gradient of the area containing the majority of objects being less than the equal-energies line. Once again there is another cluster of objects with very low eFEX energies at high CP energies.

The general trend is a high but imperfect match rate and decreased performance in the later run compared to the earlier run, both in terms of match rate of objects and energy correlation between the two systems.

From the information provided by this analysis, issues in the system were identified and solved. In the case of the degraded performance for Run 427885, the different beam conditions in this run (bunch trains, that were not present for Run 423433) were understood to have caused issues with the bunch crossing identifier (BCID) on the Liquid Argon Trigger Optical Mezzanine (LATOME) modules which provide the eFEX with digitised energies from the calorimeter.

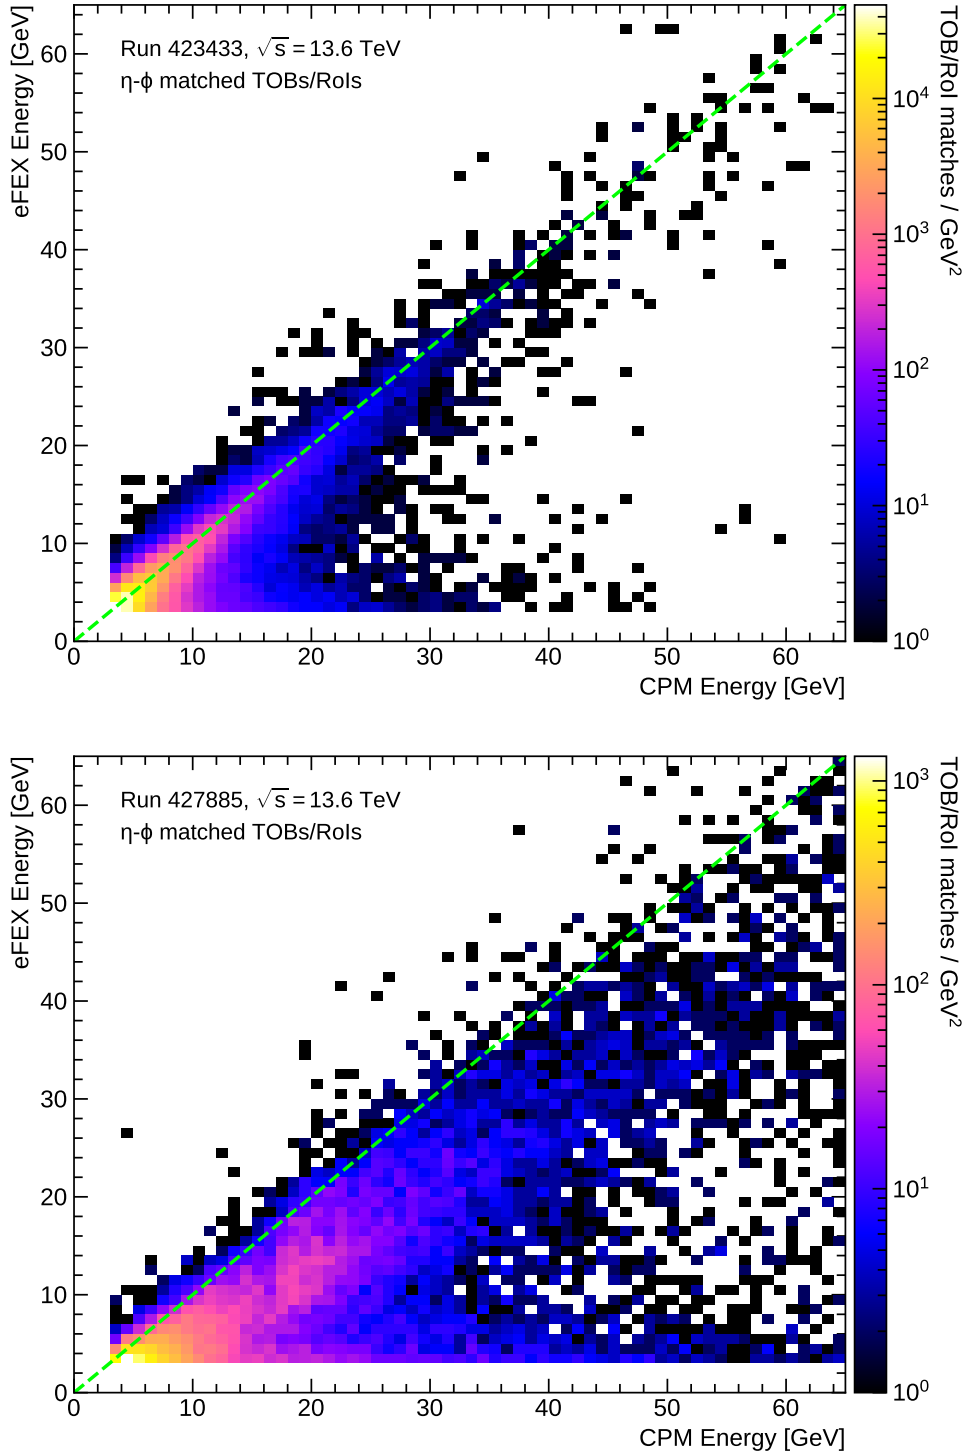


Figure 3.9: Comparison of energies for matched TOBs/Rols with the energy as measured by the Cluster-processor Module (CPM) given on the x -axis and as measured by the eFEX on the y -axis. Contains data for all matched objects in Run 423433 (top) and 427885 (bottom). The dashed line marks the set of points where the CPM and eFEX energies are equal.

Many initial problems with the Phase-I e/γ trigger have now been understood and fixed, in part thanks to the work presented here. The eFEX is now in use in the Run-3 trigger menu and performing better than the Run-2 system, evidenced by the efficiency curves shown in Figure 3.10. The consistency of the new and old systems after fixes were implemented is shown by Figure 3.11, showing the same TOB-RoI energy comparisons in a later run.

3.4 Performance studies of e/γ algorithms for the Global Event Processor

The GEP, when introduced in the Phase-II upgrade, will aim to improve discrimination in the hardware trigger for many signatures, but notably for e/γ objects. The GEP will be working alongside the eFEX system, introduced in Phase I, but will have access to more information, giving it potential to improve upon decisions made by the eFEX. To realise this improvement, new algorithms will need to be implemented in the GEP to take advantage of the finer granularity information available to it.

Designing algorithms to be used in future hardware systems is achieved through prospective performance studies. Performance studies use simulations of the expected response of a system to evaluate the performance of individual algorithms. These studies benefit from the ease of implementing algorithms in high-level software, although it is still important to consider the complexity of implementation in firmware when designing algorithms. Performance studies are typically the first step in designing a system as evaluating performance in simulations before a system is built can inform the design of the hardware.

This section explores the specific implementation and possible performance of the E_{ratio} algorithm in the GEP, expected to significantly improve discrimination for e/γ by making use of fine granularity input information [53, p. 126]. Section 3.4.1 details the samples used for evaluating algorithm performance, Section 3.4.2 discusses how




Figure 3.10: Comparison of single electron trigger efficiencies for the Run-2 and Phase-I L1Calo triggers, as a function of electron p_T (as recorded in offline reconstruction). [58]

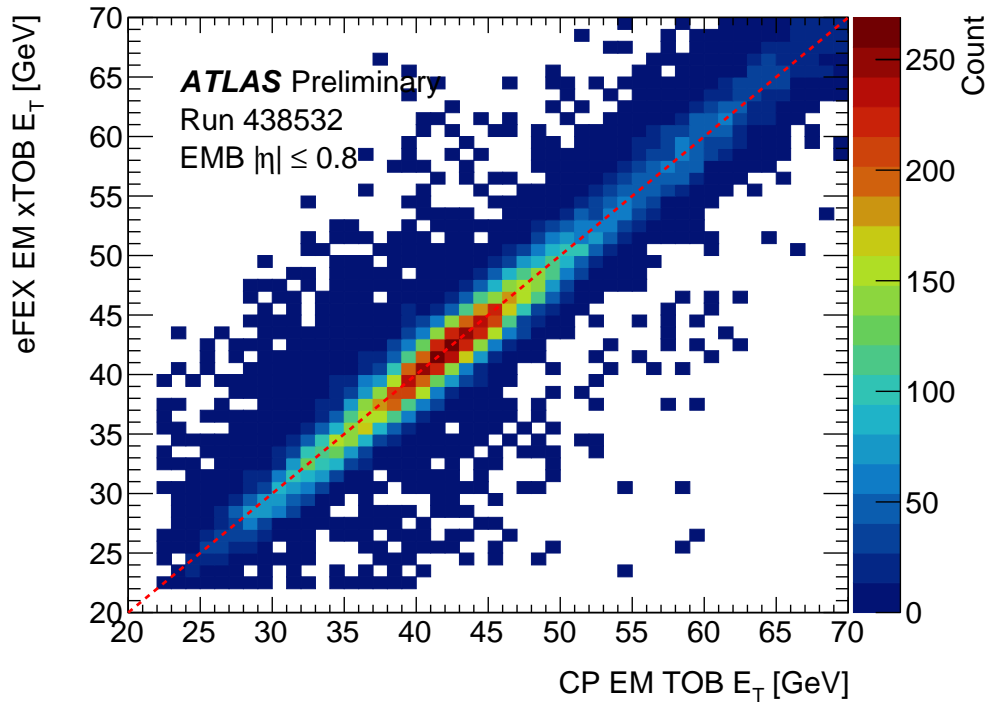


Figure 3.11: Comparison of transverse energies for matched EM objects between the Run-2 CP and the Phase-I eFEX. Shown are matches between leading electrons satisfying $|\eta| < 0.8$ in each event in Run 438532. The dashed line marks the set of points where the CP and eFEX energies are equal. [58]

the GEP itself is simulated, Section 3.4.3 gives metrics used to evaluate performance, then Section 3.4.4 goes through the process of designing an algorithm, the outcome of which is evaluated in Section 3.4.5.

3.4.1 Monte Carlo samples

Simulations used for the performance studies presented here are from two Monte Carlo (MC) samples: a $Z \rightarrow ee$ sample providing signal EM objects that the trigger should be accepting, and a minimum bias QCD sample providing background objects, typically low-energy jets, that the trigger should be rejecting. The signal sample is generated by POWHEG [22] and PYTHIA [59], and the background sample is generated by PYTHIA.

Samples are processed by the typical ATLAS detector simulation (see Section 1.6) and are overlaid with pileup events. Samples with the highest available pileup¹ (80 proton-proton interaction per bunch crossing) were used as these studies are intended to represent HL-LHC conditions. Additional simulations of the upgraded trigger are performed on samples by the Phase-I offline software, in order to simulate the eFEX response to each event.

3.4.2 Phase-II simulation

Producing prospective results for the Phase-II trigger requires simulation of the requisite algorithms. On top of the existing simulations of the Phase-I system, two things are needed to produce the results possible with Phase II: collection of the higher granularity calorimeter data that will be available to the GEP, and any algorithms that the GEP will run on its input data.

The first of these tasks is done by taking the location of e/γ candidate TOBs identified by the simulated eFEX, collecting calorimeter cell energies in a region around this location, and storing them in a cluster. This localised method was chosen,

¹at the time of analysing

over storing calorimeter cell output from the entire detector, to reduce computing requirements. The size of stored clusters is 0.3×0.3 in $\eta \times \phi$, centred on the seed TOB location, chosen conservatively to be sufficiently large that they will contain all information required by any algorithm.

Samples with these clusters of high-granularity calorimeter data included are then used for developing prospective algorithms for the GEP, explored in detail in Section 3.4.4.

3.4.3 Performance benchmarks

These studies focus on performance of the e/γ trigger at hardware level. As such, the goal is to ensure the signal efficiency (the fraction of signal events selected by the trigger) is as high as possible. At the same time the amount of background being rejected should be as high as possible; this corresponds to maximising the background rejection, where

$$\text{background rejection} = \frac{1}{\text{fraction of background events selected}}.$$

Both signal efficiency and background rejection will be dependent on the selections made by different algorithms. To compare algorithms, or different variants of an algorithm, both of these quantities must be considered. The typical performance benchmark used in these studies will be the background rejection at 95% signal efficiency.

3.4.4 E_{ratio} algorithm design

The focus for this study is on the impact of a single variable in e/γ discrimination, E_{ratio} . E_{ratio} is a shower-shape variable, already used in the HLT. The definition used here is

$$E_{\text{ratio}} = \frac{E_2}{E_1} \tag{3.1}$$

where E_1 and E_2 are the first and second most energetic cells in Layer 1 of the EM calorimeter in an area around the centre of the shower.²

The E_{ratio} variable is designed to discriminate against substructure in a shower. A shower with multiple distinct branches (e.g. $\pi^0 \rightarrow \gamma\gamma$) might produce two peaks of similar energy and give an E_{ratio} value close to one, whereas a shower with a single peak (as expected from e/γ clusters) would give an E_{ratio} value close to zero.

Calculating this variable in high-level software is straightforward and requires no optimisation. However, identifying the two required maxima involves a large number of comparisons between cell energies. An algorithm developed to run on hardware should be as simple as possible, therefore designing an alternate implementation is beneficial to minimise the impact of this algorithm on the latency of the GEP system.

A simple approach to finding the two highest energy cells in a cluster is to form a sorted list of all energies from Layer 1 cells, or at least sufficiently sorted to be confident in the highest two energies. Sorting algorithms are a very well-understood problem and heavily optimised but this approach is very rigid, not allowing for any tuning of the algorithm. For example, a cluster may have no substructure but fall on the boundary between two cells, depositing a similar amount in each. This would result in a high, background-like, E_{ratio} value. To avoid this the algorithm could include a minimum distance between cells considered to be the two maxima, or try to identify minima between the two, this would greatly complicate a list-sorting approach.

The most complete, but resource-heavy, method might consist of fitting some functional form to the energies as a function of η and ϕ to extract the peak energies. This might work in software but is very computationally expensive, even if possible to implement in firmware it is likely not worth the latency it would require.

The desired solution is an algorithm for calculating an E_{ratio} -type variable that comes somewhere between these two options, more adaptable than the list-sorting

²This is different to the definition used in the HLT, which instead is $E_{\text{ratio}} = (E_1 - E_2)/(E_1 + E_2)$. The simpler definition is preferred here in the spirit of reducing calculation in firmware, though the two forms are a transformation of one another.

approach and less resource-heavy than peak fitting. The following sections explore such an algorithm: a baseline algorithm for finding secondary maxima in clusters in the GEP is established in Section 3.4.4.1; parameters of the algorithm are tuned using simulations in Sections 3.4.4.2, 3.4.4.3, and 3.4.4.4; and a summary of the results and recommended parameters, as well as additional adjustments that could be made with further studies, is given in Section 3.4.5.

3.4.4.1 Initial algorithm

Identifying the two highest energy cells is done in three stages: locating the seed, identifying candidate secondary maxima, and comparing results.

The GEP will receive a seed location from the eFEX identifying which SuperCell has the highest energy. The cells within this SuperCell are compared with one another to find which has the highest energy, this becomes the seed cell for the E_{ratio} algorithm.

The algorithm will then perform a stepwise search from the seed outwards to identify peaks in energy. On each step the energy gradient is calculated as $\Delta E = E_{\text{next}}^{\text{cell}} - E_{\text{prev}}^{\text{cell}}$, where $E_{\text{next}}^{\text{cell}}$ is the energy of the cell being stepped to, and $E_{\text{prev}}^{\text{cell}}$ is the energy of the cell being stepped from. From the first step ΔE should be negative, as the seed will have a higher energy than the surrounding cells, but on subsequent steps ΔE may become positive, marking that a minimum-energy point has been passed. If, after this, ΔE becomes negative again it indicates that the previous cell was a local maximum; in this case that cell is added to a list of candidate secondary maxima, and the search stops along this route. If the edge of the available range of cells is reached before ΔE turns positive then no candidate is saved. If the edge is reached after ΔE turns positive, but before it turns negative again, then the last cell in the range is taken to be the candidate.

This stepwise search is done in six different routes from the seed: one route where each step from the seed is in positive η , one in negative η , two where the first step is in positive ϕ before proceeding in positive or negative η , and two following the same




Figure 3.12: Diagram showing the 6 different routes in which the E_{ratio} algorithm searches for secondary maxima (left) and how the algorithm identifies secondary maxima by tracking energy gradients along each step (right).

pattern with the first step in negative ϕ . Figure 3.12 shows a schematic of these six routes, alongside a schematic depicting the peak location strategy.

Once the stepwise search is complete, up to 6 candidate secondary maxima will have been identified. The candidate with the largest energy is taken as the secondary maximum and, with the seed as the maximum, E_{ratio} can be calculated using Equation 3.1.

The performance of this baseline algorithm was investigated using simulations. Figure 3.13 shows the results, comparing the response in signal and background as a function of the calculated E_{ratio} value and the fraction of each that would pass a given E_{ratio} threshold. The background rejection as a function of signal efficiency is also shown, the baseline algorithm achieves a background rejection of 2.3 at 95% signal efficiency.

3.4.4.2 Peak size

The first parameter to investigate is the size of the area used to calculate each energy value. In the algorithm as described in Section 3.4.4.1, the energies used in




Figure 3.13: Performance of baseline E_{ratio} algorithm on signal ($Z \rightarrow ee$) and background ($JZ0W$) clusters. Plots show (a) a histogram of calculated E_{ratio} values for each cluster, (b) the integral of (a) with a grey dashed line indicating the values at 95% signal efficiency, and (c) the background rejection of an E_{ratio} threshold corresponding to a given signal efficiency.




Figure 3.14: Diagram showing which cells contribute to the energy sum for the seed cell (bright yellow) for different peak sizes. The horizontal axis represents η and the vertical axis ϕ . In each case, the calculated energy would be the sum of the energies of the cells contained within the red box.

comparisons and in the final E_{ratio} calculation are always the energies of a single cell. This could be modified by instead summing the energy of a cell with that of its neighbours in η to reduce sensitivity to small fluctuations. The number of cells summed is labelled the ‘peak size’, where the default algorithm would have a peak size of one. With a peak size greater than one the algorithm uses a ‘sliding window’ approach, so the step size is still a single cell despite the energy value coming from a larger area. For an odd numbered peak size the energy of a cell is added to that of its neighbours on each side. For an even numbered peak size, neighbours in positive η are preferred. Figure 3.14 shows how cells are included in the calculated energy.

Performance for the E_{ratio} algorithm was tested with peak size values from one to five. The results are shown in Figure 3.15. Comparing the benchmark value of background rejection at 95% signal efficiency, it is clear that a peak size of one (i.e. the same as the baseline algorithm) gives the best results, with performance degrading as more cells are added to the energy sum. This appears to be generally true for background rejection at all signal efficiencies. This suggests that the benefit of the fine granularity of each energy measurement outweighs the negative impact of any potential fluctuations that the increased peak size would smear out.




Figure 3.15: Results for calculating E_{ratio} with different peak size options. Plots show background rejection as a function of signal efficiency for each peak size tested (left) and background rejection at 95% signal efficiency as a function of peak size (right).

3.4.4.3 Exclusion region

Another alteration tested on the E_{ratio} algorithm is an ‘exclusion region’ around the seed, i.e. a number of cells close to the seed in which secondary maxima will not be searched for. An n -cell exclusion region means making the first step along any route n cells away from the seed in η . Since secondary maxima can be found as soon as two steps have been taken from the seed cell (they cannot be found on the first step as it will always be a step down from the seed), this excludes all cells in an η range from $-n$ to n (in relative coordinates) from being considered secondary maxima. Comparatively, the baseline algorithm with no exclusion region can find secondary maxima anywhere but the four cells directly adjacent to the seed. Figure 3.16 highlights the effect of the exclusion region.

Performance for the E_{ratio} algorithm was tested with exclusion regions from between one and five cells, shown in Figure 3.17 alongside the baseline algorithm with no exclusion region. This time a clear increase in performance is visible compared to the initial form of the algorithm, with a one-cell exclusion region attaining a background rejection of 3.1 at 95% signal efficiency. For most signal efficiencies the one-cell




Figure 3.16: Diagram showing how the introduction of an exclusion region to the E_{ratio} algorithm prevents secondary maxima close to the seed from being selected. Red arrows mark each of the six paths traversed by the stepwise algorithm. Blue dots mark each step where the energy gradient is calculated. The shaded grey area shows cells that cannot be selected as a candidate secondary maximum, due to either being skipped over or being the first step from the seed.

exclusion still seems to perform best, though perhaps competing with a two-cell exclusion region for very high signal efficiencies. Since the only difference between no exclusion region and the one-cell case is that cells diagonally adjacent to the seed are excluded, these results suggest signal clusters frequently create secondary peaks on these diagonals; this could stem from incident particles falling close to the corner of a cell.

Given that cell widths vary significantly in different regions of the calorimeter, the performance of the E_{ratio} algorithm with different exclusion widths was also tested as a function of η . Figure 3.18 compares background rejection at 95% signal efficiency in several η regions. It is evident that the one-cell exclusion region performs best regardless of calorimeter geometry. The difference between one-cell and two-cell exclusion regions is much more drastic in the high- η endcap regions, here the strips are less granular so likely the larger exclusion regions are starting to miss real secondary peaks in background clusters.

3.4.4.4 Search limit

In the baseline E_{ratio} algorithm, the stepwise search for secondary maxima extends as far as the available data allows, in this case to the edge of the, conservatively large,




Figure 3.17: Results for calculating E_{ratio} with different or no exclusion region definitions. Plots show background rejection as a function of signal efficiency for each tested exclusion region (left) and background rejection at 95% signal efficiency as a function of exclusion region size (right).




Figure 3.18: Plot of background rejection at 95% signal efficiency as a function of pseudorapidity, η , for E_{ratio} algorithms with different exclusion regions.




Figure 3.19: Results for calculating E_{ratio} after varying the search limit parameter, given as a distance in η from the seed cell. Plots show background rejection as a function of signal efficiency for each tested search limit (left) and background rejection at 95% signal efficiency as a function of the search limit (right).

stored cluster size. To minimise the amount of processing required by the algorithm, and potentially improve performance by reducing overlap with other clusters, a limit can be placed on the distance this search will traverse. Since the ϕ range of the search is already limited to one cell either side of the peak, this search limit is implemented as a maximum distance traversed in η . This distance is calculated in pseudorapidity units rather than number of cells to give a consistent response across calorimeter regions.

The performance of the E_{ratio} algorithm with different search limit values was tested on simulations, with the results presented in Figure 3.19. Distances in $\Delta\eta$ from 0.025 up to 0.15 were tested, with 0.15 being the width of the clusters and thus the limit in place in the baseline algorithm. While no performance gains are seen by reducing the search limit, there is a plateau in performance from $\Delta\eta > 0.1$ (equivalent to 32 cells in the barrel region). This means the required cluster size, and thus the amount of computation required, can be reduced without degrading performance of the algorithm.

3.4.5 Algorithm summary

Given the results presented in Section 3.4.4, the most performant and resource-efficient algorithm for calculating E_{ratio} in the GEP is the baseline algorithm established in Section 3.4.4.1 with an exclusion region of one cell and a search limit of $\Delta\eta = 0.1$. No further improvement was found by varying the peak size. This algorithm achieves a background rejection of 3.1 for 95% signal efficiency.

The E_{ratio} algorithm presented here is functionally complete and serves as an option for e/γ discrimination in the GEP. Additionally, further improvements could likely be made with more studies. More parameters for this E_{ratio} algorithm could be conceived and tested to potentially improve performance. One example is a threshold in the energy gradient between steps to allow a change in gradient to be identified, which might improve the response of the algorithm to noise or statistical fluctuations.

This study focused on the design of the E_{ratio} algorithm, using background rejection as a metric for performance. Due to technical limitations, the MC samples used do not represent the projected pileup conditions of ATLAS during the HL-LHC. Further study would be needed for a full evaluation of the performance possible in these conditions, alongside other components of the hardware trigger.

Analysis methods

This chapter presents methods required to perform the analyses presented in the following chapters. Two analyses are discussed, both making use of the $Z\gamma jj$ final state: VBS of a Z boson and a photon, and semileptonic triboson production of a Z boson, photon, and an additional massive vector boson ($VZ\gamma$). The majority of the methods discussed here are relevant to both analyses, which have many common features.

The primary difference between the analyses is the jet phase space: very high energy forward jets in the VBS $Z\gamma$ case, and more central jets with a dijet mass peaking around the W/Z boson masses in the $VZ\gamma$ case. How these differences are addressed as the analyses diverge is discussed in Chapters 5 and 6.

First the data and simulated samples used for the two analyses are discussed in Sections 4.1 and 4.2. The specifics of how detector signatures are reconstructed into physics objects are detailed in Section 4.3, and the shared preliminary selection for the analyses is detailed in Section 4.4. Section 4.5 gives some background on machine learning methods used for the $VZ\gamma$ analysis. Each of the backgrounds, which are common to both analyses, is introduced in Section 4.6. Finally, Section 4.7 covers systematic uncertainties which affect the two measurements, and Section 4.8 discusses the statistical tools used to make inferences from the data.

4.1 Data and blinding strategy

The presented analyses use data collected by the ATLAS experiment during Run 2 of the LHC, between 2015 and 2018. The dataset includes all events from relevant ATLAS ‘good-run lists’, i.e. recorded with stable beam conditions and all relevant subdetector systems operational; the integrated luminosity for this sample is 139 fb^{-1} [49].

The unscaled single lepton and dilepton triggers [60, 61] were used to select data events, due to the requirement of a leptonically decaying Z boson in events. Table 4.1 gives the p_T thresholds required by these triggers for isolated leptons, depending on the lepton flavour and run period. Additional isolation or identification criteria are required for objects to pass these triggers: tight(medium) identification working points for single electron triggers in 2016-18(2015), loose identification for dielectron triggers, and medium(loose) isolation working points for single muon triggers in 2016-18(2015). Additional single lepton triggers with higher p_T thresholds and looser isolation or identification requirements are also included to improve efficiency.

This set of triggers was found to accept 99% of events which would pass the VBS $Z\gamma$ selection described in Section 5.1. This efficiency is expected to be comparable for the semileptonic $VZ\gamma$ selection.

Analyses are performed ‘blind’, meaning that data yields in certain regions are not looked at until the analysis strategy is decided. This is done to avoid data bias,

Table 4.1: Transverse momentum thresholds for triggers used for data in presented analyses. Where two numbers are given, for the dilepton triggers, the first gives the threshold for the leading lepton and the second for the sub-leading.

Signature	Threshold p_T [GeV]	
	2015	2016-18
Single electron	24	26
Single muon	20	26
Dielectron	12, 12	24, 24
Dimuon	18, 8	22, 8

i.e. adapting the analysis procedure based on effects in the data (which could be statistical fluctuations).

In each analysis certain control regions are used, both for estimating backgrounds and validating data-MC agreement. These regions were ‘unblinded’ first in order to validate the methods for which they are used. The signal regions in the two analyses remained blinded until the fits had been finalised, at which point unblinding and running the fit represents the final measurement being taken.

4.2 Simulated event samples

Samples created from MC simulations are used in the analyses to represent the SM prediction for the rate of a particular process (see Section 1.6). Beyond providing the SM estimate to which data is compared in the chosen sensitive phase space, the signal region, these simulations are also used to design the analysis. This includes optimising the selection cuts which define sensitive regions, training machine learning discriminants, and estimating the expected sensitivity of the analysis before unblinding.

The two analyses presented here have the same underlying signal process and the same set of backgrounds, so the MC samples used are common for both analyses. Table 4.2 summarises how these samples were produced, including the physics process which is simulated; the MC generator used for the hard scatter; the generator used to add parton showering, hadronisation, and underlying event; the order to which the cross section is calculated for the hard scatter; and the PDF set used by the hard-scatter generator. The remainder of this section gives more details for each of these samples, as well as some additional samples or variants of these samples which are necessary for evaluating analysis uncertainties. The dataset identifiers (DSIDs)¹ are given for each sample where appropriate.

¹This is an internal ATLAS identifier for the sample, included for completeness.

Table 4.2: Samples from MC simulation used in estimating signal and background processes. For each sample the generator used for the hard scatter process is listed as well as the generator used to add parton showering, hadronisation, and the underlying event (marked PS&UE). The order to which the cross section is calculated and the PDF set used are also given. Numbers on the right are used to label the samples in the text. Information on the listed PDF sets can be found in References [18, 62, 63, 64]

Process	Hard scatter	PS&UE	Order	PDF set	
EW $Z\gamma jj$	MADGRAPH	PYTHIA	LO	NNPDF3.1 LO	(1)
QCD $Z\gamma jj$	MADGRAPH	PYTHIA	NLO	NNPDF3.0 NLO	(2)
	SHERPA	SHERPA	LO	NNPDF3.0 NNLO	(3)
Z+jets	POWHEGBOX	PYTHIA	NLO	CT10 NLO	(4)
$t\bar{t}\gamma$	MADGRAPH	PYTHIA	LO	NNPDF2.3 LO	(5)
QCD WZ	SHERPA	SHERPA	NLO	NNPDF3.0 NNLO	(6)
EW WZjj	MADGRAPH	PYTHIA	LO	NNPDF3.0 LO	(7)

The signal sample (Sample 1) uses MADGRAPH5_AMC@NLO 2.6.5 [19] as well as PYTHIA 8.240 [21]. The DSIDs for this sample are 363267-363268. An alternate version of this sample is produced, with HERWIG++ 2.7.1 [23, 65] in place of PYTHIA, to evaluate uncertainties due to the choice of parton showering and underlying event model.

For the QCD $Z\gamma jj$ samples, Sample 2 (DSIDs 345775-345782) is the nominal sample and uses MADGRAPH5_AMC@NLO 2.3.3 and PYTHIA 8.212 and Sample 3 (DSIDs 366140-366149) gives an alternate estimate using SHERPA 2.2.4 [20]. Both of these samples include additional hard parton emission beyond the order at which the cross-section is calculated [66]. An additional five samples are generated at particle level for this process, using SHERPA 2.2.10. These are used for evaluating theoretical uncertainty and have varied values for merging and resummation scales (see Section 1.6).

Sample 4 (DSIDs 361106-361107) models the Z+jets background using POWHEG-BOX v1 [67, 22, 68] and PYTHIA 8.186 [59]. This sample is not used directly for a background estimate, but as part of the data-driven estimate discussed in Section

4.6.2.

The $t\bar{t}\gamma$ background is modelled by Sample 5 (DSID 410389) with MADGRAPH5_AMC@NLO 2.3.3 and PYTHIA 8.212. The QCD and EW production modes for the WZjj background are from Sample 6 (DSID 364253), with SHERPA 2.2.2, and Sample 7 (DSIDs 364739-364742), with MADGRAPH5_AMC@NLO 2.6.2 and PYTHIA 8.235, respectively.

An additional particle-level sample is used to calculate interference between EW and QCD $Z\gamma jj$ production. This is estimated at LO with MADGRAPH5_AMC@NLO 2.3.3 with the NNPDF3.0 LO PDF set.

All samples interfaced with PYTHIA use a specific set of parameters derived from data, a tune (as introduced in Section 1.6). For samples generated with MADGRAPH and PYTHIA, the A14 tune [69] is used. The remaining PYTHIA sample, Sample 4, uses the AZNLO tune [70].

4.3 Object reconstruction

A reconstruction procedure is applied in order to deduce what particles gave rise to the observed signals, and to measure the kinematic properties of the incident particles. The procedure used is different for any class of physics object. This section discusses the details of the reconstruction for each of the objects used in the two presented analyses: photons, electrons, muons, and jets.

Each of these physics objects are built from more intermediate ‘detector objects’, which are groups of signals in the relevant sub-detectors. These are clusters in the calorimeters, and hits and tracks in the ID and MS. These detector objects are defined in Section 2.2.

4.3.1 Photons

Photon reconstruction covers two scenarios: unconverted photons where the photon passes through the tracker and deposits its energy in the calorimeter, or converted

photons where the photon converts into an e^+e^- pair inside the tracker. The signature for an unconverted photon is an EM cluster with $E_T > 1.5$ GeV and no associated track (due to the photon being electrically neutral). The signature for a converted photon is two opposite-sign electron candidates with tracks from the same vertex, within the tracking system, consistent with a massless particle. These signatures are considered as photon candidates.

Transverse energies are calculated by combining candidate photon EM clusters with any further clusters within a 0.075×0.125 ($\eta \times \phi$) area centred on the candidate cluster. Energy measurements are corrected for scale and resolution effects due to variation in detector response across η - ϕ and data-MC differences. This calibration to the raw calorimeter response is calculated for both electrons and photons using $Z \rightarrow ee$ events [71]. Energy corrections to photons are validated using photons from radiative Z boson decays, systematic uncertainties are included in the results (as discussed in Section 4.7) to account for uncertainties from these corrections.

Jets can produce similar signatures to photons, and so additional requirements are placed on the calorimeter shower shape to discriminate against these ‘fake photons’. The desired prompt photons typically result in more collimated clusters contained within the EM calorimeter, whereas fake photons produce broader showers and leakage into the hadronic calorimeter. An identification selection, consisting of a set of cuts on shower-shape variables, is derived to minimise photon fakes. Signal photons for these analyses are required to meet the identification criteria for the ‘tight’ working point defined in reference [72], whilst the ‘loose’ identification working point is applied as a pre-selection to all photons (this is relevant for the ABCD method discussed in Section 4.6.2). Table 4.3 gives definitions for the shower-shape variables used to form these two working points.

Figure 4.1 shows the efficiency for photon identification in Run 2, which for high- E_T ($\gtrsim 40$ GeV) loosely-isolated photons is greater than 90%.

Photons are also required to be isolated, to reduce backgrounds such as photons produced in jets and as part of hadronic or EM showers. These non-prompt photons

Table 4.3: Shower-shape variables used for identification of photons. The four variables contributing to the loose working point are shown, and all variables are used to define the tight working point. The full definition of the working points in relation to these variables is given in Reference [72].

Variable	Description	Working points
R_{had}	Ratio of E_T in the hadronic calorimeter to that of the EM cluster.	Loose, Tight
R_{had_1}	Ratio of E_T in the first layer of the hadronic calorimeter to that of the EM cluster.	Loose, Tight
R_η	Ratio of energies in a 3×7 area of calorimeter cells to that in a 7×7 area (in $\eta \times \phi$).	Loose, Tight
w_{η_2}	Lateral shower width in Layer 2 of the EM calorimeter.	Loose, Tight
E_{ratio}	Ratio of the difference between to the sum of the highest- and second-highest-energy energy cells in the cluster.	Tight
R_ϕ	Ratio of energies in a 3×3 area of calorimeter cells to that in a 3×7 area (in $\eta \times \phi$).	Tight
$w_{s,\text{tot}}$	Total lateral shower width in Layer 1 of the calorimeter	Tight
$w_{s,3}$	Lateral shower width in a window of three cells around the highest-energy cell	Tight
f_{side}	Fraction of energy outside central three strips within seven strips of the maximum	Tight
ΔE_s	Difference between energies of the second-highest-energy cell and the lowest energy cell in the region between the second-highest-energy cell and the highest-energy cell.	Tight
f_1	Ratio of the energy Layer 1 of the EM calorimeter to the energy of the EM cluster.	Tight

will typically appear nearby other activity in the detector. Prompt photons can be selected by ensuring that photon candidates are isolated in a region with little activity around them. Two variables are used to define the isolation: $p_T^{\text{cone},20}$ and $E_T^{\text{cone},20}$. Here $p_T^{\text{cone},20}$ is the sum of transverse momenta of all $p_T > 1$ GeV tracks originating from the primary vertex (PV) (the vertex of the hard scatter event), within a cone of $\Delta R < 0.2$ around the direction of the photon. The $E_T^{\text{cone},20}$ variable is the sum of EM cluster transverse energies within the same cone, minus the energy of the photon. Cuts are defined on these variables as a function of the photon p_T : $p_T^{\text{cone},20} < 0.05 \cdot p_T^\gamma$ (track isolation) and $E_T^{\text{cone},20} < 0.065 \cdot p_T^\gamma$ (calorimeter isolation). This corresponds to the ‘FixedCutLoose’ criteria defined in Reference [72].

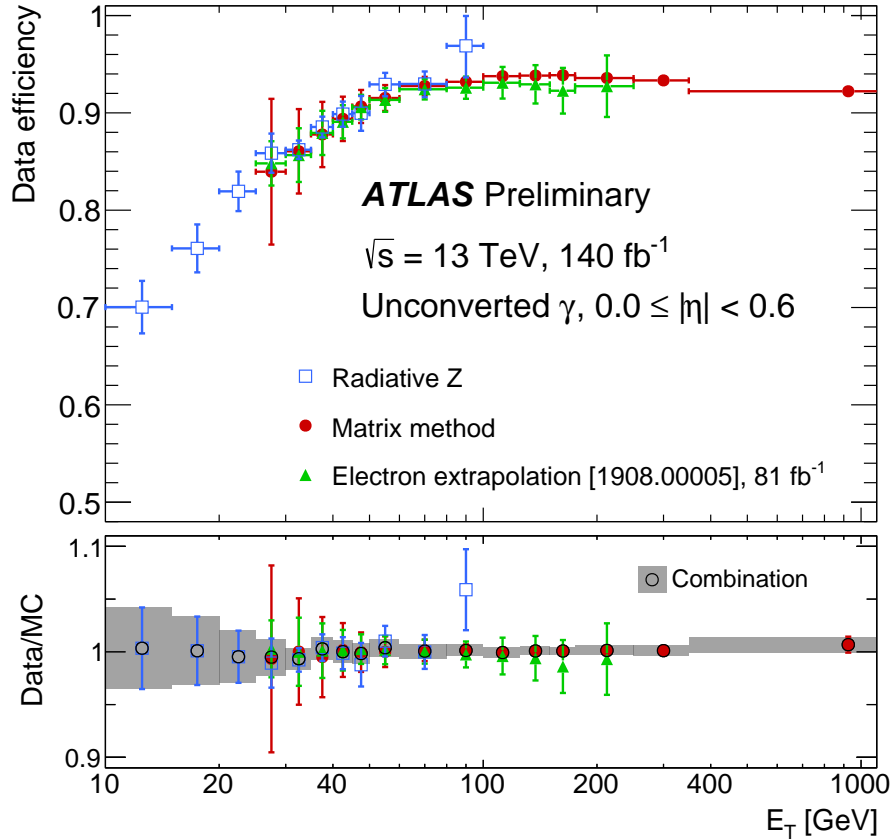


Figure 4.1: Efficiencies for identification of photons in Run 2 of the ATLAS detector, as a function of transverse energy. Shown are the values for loosely-isolated unconverted photons in the central region of the detector. The efficiency is evaluated from data using three techniques, as indicated. [73]

Systematic uncertainties are included in the measurements to account for uncertainties in the efficiencies of photon identification and isolation [74].

4.3.2 Electrons

The basic signature to seed an electron is an EM cluster with $E_T > 1$ GeV with an associated track that has hits in at least four silicon layers. As with photons, the transverse energy for an electron candidate is calculated by summing the energy of the seed cluster with any additional clusters in a 0.075×0.125 ($\eta \times \phi$) area, as well as any clusters matched to the same track as the seed.

Energy scale and resolution effects are accounted for in the same manner as for photons, primarily using $Z \rightarrow ee$ decays for energy calibrations, and then validated with electrons from $J/\psi \rightarrow ee$ events [71]. Inefficiencies in this calibration are accounted for in systematic uncertainties.

Electron candidates are also subject to identification and isolation requirements, to minimise the impact of fakes. Identification is based on both EM shower shape in the calorimeter and transition radiation in the TRT. Similarly to the photon, isolation is determined in both the tracker and calorimeter by requiring that summed energies or momenta within a cone around the seed are below a threshold.

Electrons used for these analyses are required to match the ‘medium’ identification working point and the ‘FCLoose’ isolation selection, both of which are defined in Reference [72]. The efficiency of electron identification working points in Run-2 data is shown in Figure 4.2.

4.3.3 Muons

Muon candidates are seeded from tracks in either the MS or ID. A seed track in the MS must be matched to a track in the ID, and a seed track in the ID must be matched to at least three hits in the MS. Muon candidates are only considered here within the acceptance of the ID ($|\eta| < 2.5$). Candidates must produce three hits in at least two MS stations, or in only one station for muons with $|\eta| < 0.1$. The transverse momentum of the muon is calculated from a combined track fit of the tracks/hits in the ID and MS and the corresponding energy loss in the calorimeters.




Figure 4.2: Efficiencies for identification of electrons in Run 2 of the ATLAS detector, as a function of transverse energy. Loose, medium, and tight working points for electron identification are shown in different colours. [75]

This corresponds to the ‘medium’ identification working point defined in Reference [76].

Muons must also be isolated to preferentially select prompt muons from, e.g., boson decays rather than those from hadronic sources. Muon isolation is given by the total p_T in a cone around the muon divided by the muon p_T . As with electrons and photons, this is calculated in both the ID and the calorimeter.

Efficiencies for reconstructing and identifying muons with this procedure exceed 98% for tracks satisfying $0.1 < |\eta| < 2.5$ [76]. Figure 4.3 shows the efficiency as a function of muon transverse momentum.

Data-MC comparisons are used to measure the efficiency and resolution of muon reconstruction, accounted for in systematic uncertainties [76].




Figure 4.3: Efficiencies for reconstructing and identifying muons in Run 2 of the ATLAS detector, as a function of transverse momentum. Loose, medium, and tight working points for muon identification are shown in different colours. [76].

4.3.4 Jets

A jet is a physics object representing a localised grouping of hadrons, rather than a single particle. These localised groupings of hadrons can be formed by quarks or gluons radiated from the hard scatter, since the initially quasi-free partons will hadronise before interacting with any elements of the detector. Jets are thus used as a reconstruction-level analogue of a quark or gluon produced in a hard-scatter interaction.

Reconstructing a jet requires use of a jet clustering algorithm in order to combine clusters and tracks in the detector to collect all the particles likely to have been produced by the incident hadron. Jet reconstruction in ATLAS uses the anti- k_t jet clustering algorithm [77], with a distance parameter of $R = 0.4$. This algorithm is chosen for its infrared and collinear safety; in short, the same set of jets would be constructed from an event if additional soft emissions or collinear splittings are added [78]. It also produces geometrically-regular and approximately conical jets

which are convenient to work with experimentally.

The definition of a jet is dependent on the clustering algorithm used and on the information given as input to this algorithm. A ‘jet collection’ is the name used for jets produced from a certain set of inputs. The baseline jet collection used in the two presented analyses is ‘particle-flow’ jets, in the VBS analysis ‘topo-cluster’ jets are also considered. These are both defined through the anti- k_t algorithm, but the set of objects on which the algorithm acts is different.

Topo-cluster jets are formed using only calorimeter information, passing topo-clusters as input to the clustering algorithm. Topo clusters, or topological clusters, are a pileup-resistant formulation of a calorimeter cluster: with cells added to a cluster based on whether the measured energy exceeds a threshold determined by the expected noise in that cell. Since only calorimeter information is used to create the jets, topo-cluster jets rely heavily on the granularity and performance of the calorimeter [79].

Particle-flow jets are an alternative jet collection created by using ‘particle-flow objects’ as input to the clustering algorithm. A particle-flow object is a combination of calorimeter topo-clusters and ID tracks, with calorimeter deposits produced by charged particles removed to avoid energy/momentum double-counting, designed to represent a single particle. Combining calorimeter and tracker information allows for improved resolution at lower energies compared to topo-cluster jets [80].

Systematic uncertainties are included to account for effects on the energy scale and resolution of jets from detector calibration, properties of the jet such as quark/gluon flavour composition, and data-MC differences [81].

It can be useful for physics analyses to establish whether a jet was likely initiated by a b-quark; this process, known as ‘flavour tagging’ or ‘b-tagging’, is used in the VBS $Z\gamma$ analysis to reject against backgrounds which are likely to produce b-quarks in the final state. Flavour tagging is performed by dedicated algorithms which use tracking information to identify differences in the properties of tracks and reconstruct displaced vertices (e.g. where a b-quark has travelled a short distance in the detector before decaying to lighter quark flavours) within a jet [82].

4.3.5 Overlap removal

Overlap removal (OR) avoids double-counting of energies across multiple objects. Jet OR avoids fake jets that are reconstructed from a prompt electron or photon signal and also avoids cases where near-coincident leptons and jets might bias each other's reconstruction. Muon-electron OR avoids reconstructing electrons from muon bremsstrahlung, meanwhile photon-lepton OR acts as an additional isolation precaution [83].

Once the above objects are reconstructed, they are only used in the analyses if they survive OR. This procedure takes the lists of each set of object and removes candidates based on their proximity to other objects. First, jet candidates are removed if $\Delta R(j, e) < 0.2$ or $\Delta R(j, \gamma) < 0.4$ for any e or γ . Lepton candidates are subsequently removed if $\Delta R(l, j) < 0.4$ for any remaining jet. Photon candidates are removed if $\Delta R(\gamma, l) < 0.4$ and electrons removed if $\Delta R(\mu, e) < 0.2$.

4.4 Event selection

Events from data and simulation undergo a selection process to focus on a phase space that matches the desired final state. This selection applies to data samples as described in Section 4.1 and MC samples as in Section 4.2, with objects reconstructed following the procedure in Section 4.3.

For an event to be selected, first the basic objects in the desired final state need to be present. There must be at least one photon and precisely two electrons or muons present, of the same flavour to each other but opposite charge, with all of these passing the relevant isolation and identification criteria specified in Section 4.3. Both analyses also require the presence of two jets, but their selection varies per analysis and is discussed separately there.

Further selection is applied to the lepton-photon system in order to identify events with a real Z boson and a photon not produced from final-state radiation (FSR).

This $Z\gamma$ selection is detailed in Table 4.4, and acts as a pre-selection for both analyses before additional jet selection is applied.

The cut on the sum of the dilepton mass and the dilepton-photon mass, $m_{ll} + m_{ll\gamma}$ is imposed to reject events with photons from FSR. In an FSR event, the photon is radiated from one of the final state leptons. This means the two leptons and the photon all originate from the same Z boson, and their invariant mass should be close to the mass of the Z boson. The invariant mass of the dilepton system in this case would be less, and so the sum of these masses should be less than twice the Z mass. In a non-FSR event, the photon is radiated independently of the Z boson production, giving a dilepton mass close to the Z mass and a larger dilepton-photon invariant mass. The sum of the masses in this instance is typically greater than twice the Z mass. Figure 4.4 shows a two-dimensional distribution of these two invariant masses and how this cut rejects the population of events with FSR photons.

This FSR rejection is implemented since the photon emission from a final state lepton excludes the possibility of the photon having been produced in a multiboson

Table 4.4: Cuts implemented for both analyses to select $Z\gamma$ events. Here $p_T^{l,1}$ denotes the p_T of the leading (i.e. highest p_T) lepton, and $p_T^{l,2}$ denotes that of the sub-leading (second highest p_T) lepton. OSSF indicates that two opposite-sign same-flavour leptons are required.

$Z\gamma$ selection	
Photon	$N_\gamma \geq 1$
	$ \eta_\gamma < 2.37$
	(excludes $1.37 < \eta_\gamma < 1.52$)
	$p_T^\gamma > 25 \text{ GeV}$
Lepton	$N_l = 2 \text{ (OSSF)}$
	$ \eta_e < 2.47$
	(excludes $1.37 < \eta_e < 1.52$)
	$ \eta_\mu < 2.5$
	$p_T^{l,1} > 30 \text{ GeV}$
	$p_T^{l,2} > 20 \text{ GeV}$
Boson	$m_{ll} > 40 \text{ GeV}$
	$m_{ll} + m_{ll\gamma} > 182 \text{ GeV}$




Figure 4.4: Distribution of events as a function of $m_{l\gamma}$ and $m_{l\gamma}$, for events in the signal sample passing the lepton and photon cuts given in Table 4.4. The dashed line shows the threshold for the FSR-rejection cut, events below the dashed line are discarded.

interaction, which is the focus of these analyses.

4.5 Boosted decision trees

Boosted decision trees (BDTs) are a machine learning technique used commonly in particle physics due to their ability to discriminate between populations (e.g. signal and background) through supervised training. This means that they can be fed events from MC simulations labelled as either signal or background, learn the features which distinguish the two populations, and create a variable designed to separate background and signal events. When given the same input information for new events (e.g. data events) the BDT can then calculate this output variable and its value will indicate how signal-like an event is.

4.5.1 Decision trees

The basic element of a BDT is a decision tree. Given a sample of events, a decision tree is built by splitting the events into two sub-samples (or branches) by making a cut on a kinematic variable. These branches can then be further split by making additional cuts, each cut splitting events in two based on whether they fall above or below the cut threshold. Given knowledge of whether each event is signal or background, the cuts can be chosen to give the optimal separation at each step. The result will be that each branch has either a higher signal or background purity than the one before it.

Each of the final subsets of events in the tree are called leaves. Each leaf will either have a majority of signal events, and thus be a signal leaf, or background events, and so is a background leaf. The decision tree gives an output for each event, either signal or background depending on which leaf it is placed in. Events in signal leaves have an output of $+1$ and events in background leaves are assigned -1 .

This process can be adjusted by controlling certain parameters related to the tree's construction, known as hyperparameters. For instance, when searching for an optimal cut on a particular variable, there will be some granularity for cut values that are tested. This leads to an N_{cuts} hyperparameter, the number of potential cut values tested. There must also be a stopping condition for when to stop splitting the branches. Specifying a maximum depth for the tree, d_{max} , achieves this.

This process alone is similar to optimising a simple cut-based analysis, and does not typically provide a strong classifier. Decision trees can however be enhanced through boosting.

4.5.2 Boosting

Boosting is a process in which decision trees are built iteratively, with events reweighted after building each tree in order to focus on areas where the decision tree performed poorly. This is typically done for of order 1000 decision trees. For BDTs used

here, boosting is achieved with an algorithm called ‘AdaBoost’ (adaptive boosting), described below [84].

Initially, all events are assigned a weight of one. A decision tree is built from these events using the method in Section 4.5.1, this is the first iteration. The misclassification rate, or error rate, in this tree is determined as

$$\varepsilon = \frac{\sum_{i \in \{x^{\text{mis}}\}} w_i}{\sum_{i \in \{x\}} w_i},$$

where w_i is the weight for event i , $\{x\}$ is the set of all events, and $\{x^{\text{mis}}\}$ is the set of misclassified events. An event is determined to have been misclassified if it falls in a leaf of the opposite type, i.e. a signal event in a background leaf or vice versa.

This error rate is used to calculate the boost weight for this iteration,

$$\alpha = \left(\frac{1 - \varepsilon}{\varepsilon} \right)^\beta,$$

where β is an adjustable hyperparameter. This boost weight is used to increase, or boost, the weight of each misclassified event in this iteration; their current weights are multiplied by α . The next iteration begins by constructing a new decision tree, considering the modified event weights. The process then repeats, recalculating the error rate and the boost weight and boosting misclassified events.

The β hyperparameter typically has a value of 1, but can be decreased to reduce the impact of each training iteration. The number of trees, N_{trees} , can also be adjusted as a hyperparameter to tune the response of the BDT.

4.5.3 Training and testing

A BDT is trained on a set of signal and background events in order to build its set of decision trees, which it can then use to classify further events. An independent set of signal and background events are typically used to test the BDT classification.

When training a BDT, it must be provided with a set of events (the training sample) and a set of variables which it can cut on. When creating the first branch, and any subsequent branches, possible cut values for each variable are tested on the set of events to find the cut in any variable which provides the best discrimination between signal and background. This cut is chosen, and the next branches are created by repeating the process on the remaining events in either branch. This stops once the decision tree has grown to the maximum depth, then misclassified events are boosted in order to create the next decision tree, where the process is repeated from the start (with appropriately adjusted event weights).

A common problem with BDT classifiers is overtraining. Overtraining typically occurs when the model is too complicated relative to the size of the training sample. Model complexity is linked to the number of input variables and the number and size of decision trees. The result is that the BDT will misidentify statistical fluctuations in the signal and background as features of the population.

Overtraining can be identified as a reduced performance on the independent test sample compared to that on the training sample. Some amount of overtraining is inevitable but it should normally be mitigated; although there is nothing inherently wrong with overtraining a BDT model, it will limit discrimination power. If a BDT suffers from overtraining, it can be countered by either increasing the size of the training sample or decreasing the model complexity.

4.5.4 BDT output

Once trained, each event evaluated by the BDT is given a score, calculated from the output of each individual decision tree:

$$y(x_i) = \frac{1}{N_{\text{trees}}} \sum_k^{N_{\text{trees}}} \log(\alpha_k) h_k(x_i),$$

where $y(x_i)$ is the BDT score for the i^{th} event, $h_k(x_i)$ is the output of the k^{th} decision tree for the i^{th} event, and α_k is the boost weight calculated from the k^{th} decision tree.

Lower values of this BDT score indicate an event is background-like, whereas higher values indicate it is signal-like. This variable can be used directly to represent the BDT response for events, but transformations of this BDT score can be more useful for identifying signal-rich regions if the signal-background separation is not clear in the BDT score distribution itself.

Figure 4.5 shows the BDT score distribution for a BDT training in Chapter 6. This distribution is difficult to place a signal-enriching cut on, and motivates the use of an alternate BDT response distribution.




Figure 4.5: Example BDT score distribution, from training and test data used for the semileptonic $VZ\gamma$ analysis. Signal and background events are as defined in Section 6.2.

One such BDT response variable is ‘signal rarity’. The signal rarity distribution is defined as

$$\mathcal{R}^{\text{sig}}(y) = \int_{-\text{inf}}^y \hat{y}_B(y') dy',$$

where y is the BDT score and \hat{y}_B is the distribution of the BDT score for background events in the training sample [84]. This variable takes values from 0 to 1,

and is constructed such that background events should be uniformly distributed in signal rarity, whereas signal events occur more frequently at values closer to 1. This distribution is used for the BDT response in the semileptonic $VZ\gamma$ analysis (e.g. Figure 6.5).

4.5.5 Variable importance

When training a BDT, some variables are typically found to be more discriminating than others. Variables are assigned an ‘importance’ score based on the number of times they are used to make cuts while creating decision trees, the separation gain from each cut, and the number of events in the branch [84]. When selecting variables with which to train a decision tree, ranking input variables by their importance is very useful; Section 6.2.5.1 discusses this process.

4.6 Background estimation

The two presented analyses share a common set of background processes. Due to the differing phase space, estimation of the QCD $Z\gamma jj$ background is different for each analysis. The remaining backgrounds however follow the same estimation procedure for both analyses. This section discusses the procedure for the common backgrounds: $Z+\text{jets}$, $t\bar{t}\gamma$, and $WZjj$.

4.6.1 Monte Carlo backgrounds

The background from $t\bar{t}\gamma$ events is estimated from MC with a NLO k -factor of 1.44 applied, calculated in Reference [85]. A conservative normalisation uncertainty of 15% is applied to this background estimate, much larger than the uncertainty on the inclusive cross section for this process [85].

Events from $WZjj$ make a minor contribution to the background, this is estimated solely from MC. Again a simple normalisation uncertainty is applied, here a value of 20% is chosen.

4.6.2 Fake photon estimation

Background from Z+jets events mimics the analysis final state when a jet is misidentified as a photon. Fake photons such as these are not well modelled in MC, and so the shape and normalisation of this background is calculated with a data-driven method.

4.6.2.1 Normalisation

The ABCD method is used to estimate the normalisation for this process. This is done by establishing three orthogonal control regions adjacent to the region of interest (e.g. the signal region (SR)). Cuts in two different variables, here the photon identification and isolation, are used to define these regions, as demonstrated by Figure 4.6. The region of interest is labelled as region A, inverting the photon calorimeter isolation selection gives region B, inverting the identification criteria gives region C, and inverting both gives region D. Track isolation is still required for the photon in all regions.

These three control regions are used to infer the amount of Z+jets background in the region of interest with the relationship

$$N_A^{\text{Z+jets}} = R \frac{N_{B,\text{data}}^{\text{Z+jets}} \times N_{C,\text{data}}^{\text{Z+jets}}}{N_{D,\text{data}}^{\text{Z+jets}}}$$

where $N_{X,\text{data}}^{\text{Z+jets}}$ is the number of Z+jets events in the given region calculated by subtracting background and signal leakage from the data yield i.e.

$$N_{X,\text{data}}^{\text{Z+jets}} = N_X^{\text{data}} - N_X^{\text{bg}} - c_X N_{A,\text{data}}^{\text{sig}}, \quad \text{for } X = B, C, D.$$

Background subtraction is performed for any background without a prompt Z boson and photon, in this case $t\bar{t}\gamma$ and WZjj. The correlation factor, R , is given by

$$R = \frac{N_{A,\text{MC}}^{\text{Z+jets}} \times N_{D,\text{MC}}^{\text{Z+jets}}}{N_{B,\text{MC}}^{\text{Z+jets}} \times N_{C,\text{MC}}^{\text{Z+jets}}}$$




Figure 4.6: Schematic of the four regions used for fake photon background estimation. Region A represents the signal region; B, C, and D represent control regions obtained by relaxing isolation and/or identification requirements.

where in this case each $N_{X,MC}^{Z+jets}$ is the event yield observed in Z+jets MC in this region. Uncorrelated identification and isolation requirements gives $R = 1$, so the calculated value should be close to this. Also defined are signal leakage parameters, c_X , as

$$c_X = \frac{N_{X,MC}^{sig}}{N_{A,MC}^{sig}}, \quad \text{for } X = B, C, D,$$

calculated from QCD and EW $Z\gamma jj$ MC. Signal leakage represents prompt photon events that enter the control regions (CRs), hence both EW and QCD $Z\gamma jj$ production are considered as ‘signal’ in this instance.

4.6.2.2 Shape

The shape of the Z+jets background is taken directly from a data control region. The control region should be very pure in Z+jets events, but also sufficiently high statistics. The chosen region is the anti-tight region, with no requirement on track

or calorimeter isolation. This is equivalent to regions C and D combined but without the track isolation requirement.

4.6.2.3 Uncertainties

Several components of the normalisation of this background estimate have associated uncertainties. These are propagated to the final normalisation and included as a systematic uncertainty on the results.

The MC background subtraction is subject to any uncertainty on the subtracted backgrounds. As this is predominantly from $t\bar{t}\gamma$, the 15% $t\bar{t}\gamma$ uncertainty (see Section 4.6.1) is used on the total subtracted background.

The signal leakage fractions, c_X , are split into two components, EW and QCD, representing the leakage from each source of prompt photons. To find the uncertainty on the QCD leakage fraction it is first calculated with both the nominal and alternate sample. The difference between the calculated leakage fractions is combined with the MC statistical uncertainty on the nominal sample to calculate the total QCD uncertainty. The EW leakage fraction is a minor contribution to the total leakage fraction, and so the uncertainty is taken as 50%, combined with the MC statistical uncertainty.

The correlation factor, R , has an uncertainty calculated from data-MC comparisons in complementary regions where the photon fails track isolation requirements. The correlation factor is re-calculated for both data and MC with the track isolation requirement inverted. The difference between these two R values is combined with the MC statistical uncertainty on the nominal R value to give its uncertainty. This assumes the data-MC agreement is consistent between these complementary regions and the primary ABCD regions. Inverting the track isolation selection ensures a fake-rich data sample which should be comparable to the Z+jets MC sample.

4.7 Systematic uncertainties

Many sources of uncertainty are considered for processes estimated in the presented analyses. The subsections below cover uncertainties from theoretical and experimental sources which are considered for EW and QCD $Z\gamma jj$ production processes.

Smaller sources of background, Z +jets, $t\bar{t}\gamma$, and $WZjj$, are each assigned a single normalisation uncertainty, as detailed in Section 4.6. Limited statistics in MC samples also contributes uncertainties to all processes, this is detailed below in Section 4.7.3.

All systematic uncertainties are included as nuisance parameters in the fit used for each analysis, following the procedure given in Section 4.8.2.2.

4.7.1 Theoretical uncertainties

Theoretical uncertainties are calculated for EW and QCD $Z\gamma jj$ production mechanisms. These come from a variety of sources: choice of PDF set, renormalisation and factorisation scales, QCD modelling, choice of parton showering and underlying event model, and EW-QCD interference.

Evaluating the uncertainty in PDF set choice is done by reweighting events using a number of replica PDF sets, chosen in agreement with the PDF4LHC recommendations [86]. Taking the standard deviation of yields under each of these weights gives the PDF uncertainty on the event yield.

The uncertainty due to scale choice is calculated by varying the default values of renormalisation and factorisation scales in the nominal QCD $Z\gamma jj$ MC sample. Each scale value is independently varied up and down by a factor 2. The per-bin envelope of all deviations from combinations of these variations is taken as the scale uncertainty.

QCD modelling uncertainty accounts for potential mismodelling in the hard-scatter process for QCD $Z\gamma jj$, but also covers uncertainties in parton showering and underlying event model. This is particularly relevant in the VBS $Z\gamma$ analysis where this

background is known to be poorly modelled in current MC samples for high dijet masses. QCD modelling uncertainty can be calculated conservatively by comparing predictions from different generators or, alternatively, by evaluating the effect of merging and resummation scales. Generator differences are calculated by taking the difference in event yields predicted by the nominal and alternate QCD $Z\gamma jj$ samples. This difference is considered as the QCD modelling uncertainty on the nominal yield. Alternately, uncertainty from choice of merging (CKKW) and resummation (QSF) scale is calculated using the dedicated samples described in Section 4.2. The latter method is used for the VBS analysis whilst the former is used for the triboson analysis, due to availability of samples.

For the EW signal, parton showering and underlying event uncertainties are calculated by comparing the default PYTHIA samples to alternatives with HERWIG or with eigenvariations of the PYTHIA tune [69]. The difference in predicted yields between the default and HERWIG samples is taken as the uncertainty on parton showering. The envelope of the largest deviations from the nominal sample with the tune eigenvariations applied is taken as the uncertainty on the underlying event model.

The interference between EW and QCD $Z\gamma jj$ production is not included in either the signal or background, but instead taken as an additional uncertainty, calculated using the dedicated interference sample.

4.7.2 Experimental uncertainties

Experimental systematic uncertainties cover errors in energy scale and resolution of jets, photons, and electrons; momentum scale and resolution of muons; scale factors used to reproduce trigger, reconstruction, identification, and isolation efficiencies from data; suppression of pile-up jets; and flavour tagging. The full list of experimental systematics considered between the two analyses is given in Table 4.5. The primary difference between the two analyses is that the $VZ\gamma$ analysis has no flavour tagging systematics, as no b-tagging is used; though there are other small changes

Table 4.5: List of experimental systematic uncertainties, whether they are included in the VBS $Z\gamma$ and semileptonic $VZ\gamma$ analyses, and a brief description of what the uncertainty represents. The names of uncertainties are often abbreviated when shown in figures.

Uncertainty name	Analysis		Accounts for uncertainty on
	VBS	$VZ\gamma$	
EG_RESOLUTION_ALL	✓	✓	Resolution of electron and photon measurements.
EG_SCALE_AF2	✓	✓	Scale of electron and photon measurements.
EG_SCALE_ALL	✓	✓	
EL_EFF_ID_TOTAL_1NPCOR+UNCOR	✓	✓	Efficiency of electron identification, isolation, reconstruction, and trigger.
EL_EFF_Iso_TOTAL_1NPCOR+UNCOR	✓	✓	
EL_EFF_Reco_TOTAL_1NPCOR+UNCOR	✓	✓	
EL_EFF_TriggerEff_TOTAL_1NPCOR+UNCOR	✓	✓	
EL_EFF_Trigger_TOTAL_1NPCOR+UNCOR	✓	✓	
PH_EFF_ID_Uncertainty	✓	✓	Efficiencies for photon identification, isolation, and trigger.
PH_EFF_ISO_Uncertainty	✓	✓	
PH_EFF_TRIGGER_Uncertainty	✓	✓	
FT_EFF_Eigen_B_0	✓		Flavour tagging for b- and c-jets.
FT_EFF_Eigen_B_1	✓		
FT_EFF_Eigen_B_2	✓		
FT_EFF_Eigen_C_0	✓		
FT_EFF_Eigen_C_1	✓		
FT_EFF_Eigen_C_2	✓		
FT_EFF_Eigen_C_3	✓		
FT_EFF_Eigen_Light_0	✓		
FT_EFF_Eigen_Light_1	✓		
FT_EFF_Eigen_Light_2	✓		
FT_EFF_Eigen_Light_3	✓		
FT_EFF_Eigen_Light_4	✓		
FT_EFF_extrapolation	✓		
FT_EFF_extrapolation_from_charm	✓		
JET_JER_DataVsMC_MC16	✓	✓	Jet energy resolution [87].
JET_JER_EffectiveNP_10	✓	✓	
JET_JER_EffectiveNP_11	✓	✓	
JET_JER_EffectiveNP_12restTerm	✓	✓	
JET_JER_EffectiveNP_1	✓	✓	
JET_JER_EffectiveNP_2	✓	✓	
JET_JER_EffectiveNP_3	✓	✓	
JET_JER_EffectiveNP_4	✓	✓	
JET_JER_EffectiveNP_5	✓	✓	
JET_JER_EffectiveNP_6	✓	✓	
JET_JER_EffectiveNP_7restTerm	✓	✓	
JET_JER_EffectiveNP_7	✓	✓	
JET_JER_EffectiveNP_8	✓	✓	
JET_JER_EffectiveNP_9	✓	✓	
JET_Flavor_Composition	✓	✓	Jet flavour, see section 4.7.2.2.
JET_Flavor_Response	✓	✓	
JET_JvtEfficiency	✓	✓	Jet vertex tagging efficiencies.
JET_fJvtEfficiency	✓	✓	

continued on next page

Table 4.5 continued

Uncertainty name	Analysis		Accounts for uncertainty on
	VBS	VZ γ	
JET_BJES_Response	✓	✓	Jet energy scale, including η calibration and pileup corrections [87].
JET_EffectiveNP_1	✓	✓	
JET_EffectiveNP_2	✓	✓	
JET_EffectiveNP_3	✓	✓	
JET_EffectiveNP_4	✓	✓	
JET_EffectiveNP_5	✓	✓	
JET_EffectiveNP_6	✓	✓	
JET_EffectiveNP_7	✓	✓	
JET_EffectiveNP_8restTerm	✓	✓	
JET_EffectiveNP_Detector1		✓	
JET_EffectiveNP_Detector2		✓	
JET_EffectiveNP_Mixed1		✓	
JET_EffectiveNP_Mixed2		✓	
JET_EffectiveNP_Mixed3		✓	
JET_EffectiveNP_Modelling1		✓	
JET_EffectiveNP_Modelling2		✓	
JET_EffectiveNP_Modelling3		✓	
JET_EffectiveNP_Modelling4		✓	
JET_EffectiveNP_Statistical1		✓	
JET_EffectiveNP_Statistical2		✓	
JET_EffectiveNP_Statistical3		✓	
JET_EffectiveNP_Statistical4		✓	
JET_EffectiveNP_Statistical5		✓	
JET_EffectiveNP_Statistical6		✓	
JET_EtaIntercal_Modelling	✓	✓	Efficiencies for muon isolation, reconstruction, track-to-vertex association, and trigger. Each has a statistical and systematic component.
JET_EtaIntercal_NonClosure_2018data	✓	✓	
JET_EtaIntercal_NonClosure_highE	✓	✓	
JET_EtaIntercal_NonClosure_negEta	✓	✓	
JET_EtaIntercal_NonClosure_posEta	✓	✓	
JET_EtaIntercal_TotalStat	✓	✓	
JET_Pileup_OffsetMu	✓	✓	
JET_Pileup_OffsetNPV	✓	✓	
JET_Pileup_PtTerm	✓	✓	
JET_Pileup_RhoTopology	✓	✓	
JET_PunchThrough_MC16	✓	✓	
JET_SingleParticle_HighPt	✓	✓	
MUON_EFF_ISO_STAT	✓	✓	Smearing of inner detector muon tracks.
MUON_EFF_ISO_SYS	✓	✓	
MUON_EFF_RECO_STAT_LOWPT	✓	✓	
MUON_EFF_RECO_STAT	✓	✓	
MUON_EFF_RECO_SYS_LOWPT	✓	✓	
MUON_EFF_RECO_SYS	✓	✓	
MUON_EFF_TTVA_STAT	✓	✓	
MUON_EFF_TTVA_SYS	✓	✓	
MUON_EFF_TrigStatUncertainty	✓	✓	
MUON_EFF_TrigSystUncertainty	✓	✓	Smearing of muon spectrometer tracks.
MUON_ID	✓	✓	
MUON_MS	✓	✓	
MUON_SAGITTA_RESBIAS	✓	✓	Muon momentum scale effects.
MUON_SAGITTA_RHO	✓	✓	
MUON_SCALE	✓	✓	
PRW_DATASF	✓	✓	Pileup reweighting, see Section 4.7.2.1.

due to a change in jet uncertainty configuration. These uncertainties have a varying level of effect on the presented results, some of the most impactful are discussed here.

4.7.2.1 Pileup reweighting

MC samples are typically generated before data-taking is complete. The pileup distribution, i.e. the distribution of instantaneous luminosities (see Figure 2.9), is therefore only estimated and does not exactly match that in data. Events are reweighted to align the pileup distributions between MC and data; a scale factor is calculated [88] to account for the difference between the predicted and measured inelastic proton-proton cross section [89].

This scale factor, calculated as 1.16 ± 0.07 , accounts for the fraction of visible cross-section from inelastic pp collisions. This value is found to give good agreement between data and simulation, after reweighting, for the number of inelastic interactions reconstructed in the tracking detector [90]. The systematic uncertainty `PRW_DATASF` is calculated by varying the value of this scale factor between 1.09 and 1.23 during the reweighting process, and evaluating its impact on the results.

Likely due to limited data statistics in the signal regions, this uncertainty is one of the most significant components of the total uncertainty on the results of both analyses.

4.7.2.2 Jet flavour composition and response

Jets initiated by different quarks and by gluons exhibit differences in fragmentation and showering properties. These properties will impact the jet energy scale calibration, so the distribution of light-quark-, b-quark-, and gluon-initiated jets, i.e. the distribution of jet flavour, and its uncertainty affects the overall jet energy scale uncertainty.

The response of the calorimeter to different flavours of jet is not well modelled in MC, and so is corrected using comparisons with data. Uncertainties on this correction are propagated as a ‘jet flavour response’ systematic uncertainty.

The jet response is itself dependent on the flavour composition of jets in the MC sample. This composition is dependent on the selection, so any jet selection differing from those in the jet calibration schemes will not have a well defined flavour composition. Uncertainties on the jet flavour composition within the phase space are taken as a systematic uncertainty on analyses [91].

The gluon fraction is defined as

$$f_g = \frac{N_g}{N_g + N_{\text{LQ}}}, \quad (4.1)$$

where N_g is the number of gluon-initiated jets in the phase space and N_{LQ} the number of light-quark-initiated jets. This gluon fraction is used to determine the jet flavour uncertainties, but by default its value is taken as

$$f_g = 0.5 \pm 0.5.$$

Therefore both jet flavour response and jet flavour composition uncertainties can be reduced by explicitly calculating this gluon fraction and its error in the analysis phase space. This additional step is taken in the VBS analysis to manage these uncertainties.

4.7.3 Monte Carlo statistics

Further systematic uncertainties are introduced into analyses from the limited statistics available in MC samples. For a fit over a binned distribution (as discussed in Section 4.8), the statistical uncertainties due to the number of MC events available in each bin are added as per-bin nuisance parameters to the fit, as with other systematic uncertainties.

4.8 Statistical inference

The goal in making a particle physics measurement is to test the agreement between the observed data and the SM prediction. In frequentist statistics, a model cannot be determined to be correct, instead models are rejected if there is ‘significant’ disagreement. To test agreement of data with a model (termed the ‘null hypothesis’), the level of discrepancy between the two is evaluated using hypothesis testing. The result is either that data is consistent with the model, or that the discrepancy is significant enough to reject the model. Certain thresholds are commonly used in particle physics to define when a result is significant, discussed below in Section 4.8.5.

If searching for an exotic process that is not predicted in the SM, this null hypothesis would simply be the SM. The analyses in this thesis instead search for rare processes predicted within the SM and so need a different approach, as simply determining that data is consistent with the SM does little to demonstrate the existence of this one process (the signal process). Instead a ‘background-only’ hypothesis is created, including all SM processes except for the signal process. If data shows a significant excess over the background-only hypothesis this demonstrates the signal process is likely present, particularly if the data agrees well with the combined background and signal prediction.

4.8.1 Signal strength

A useful parameterisation is to introduce a signal strength, μ , to connect the background-only hypothesis and the full SM prediction. The predicted event yield, n_{tot} , then becomes

$$n_{\text{tot}} = n_b + \mu \cdot n_s,$$

where n_b is the predicted number of background events and n_s the predicted number of signal events. Setting $\mu = 1$ corresponds to the SM prediction while $\mu = 0$ gives the background-only prediction.

This signal strength is used as the parameter of interest in fits, and gives a simple interpretation of the result. If μ is large enough and precise enough to exclude it being consistent with a value of 0, the background-only hypothesis is rejected. If μ is consistent with a value of 1, the data agrees with the SM prediction.

4.8.2 Likelihood construction

In order to fit data and obtain a measured value for μ , a likelihood function is constructed to describe the likelihood of obtaining such data given a certain value of μ . Here, two cases are discussed. First, a simplified case using the Poisson distribution to describe the likelihood of a binned dataset considering only statistical uncertainties, used as a simple performance metric in Section 6.2. The second case expands on this to create a likelihood that accounts for systematic uncertainties with nuisance parameters, as is used in the fits introduced in Sections 5.6 and 6.6.

4.8.2.1 Simple binned likelihood

Consider the predicted number of events in the i^{th} bin, as a function of μ , as

$$n_{\text{tot}}^i = n_b^i + \mu \cdot n_s^i,$$

for n_b^i and n_s^i giving the predicted numbers of background and signal events per bin respectively. The likelihood of observing n_{obs}^i events in bin i is given by a Poisson distribution with a mean of n_{tot}^i , or

$$\mathcal{L}(n_{\text{obs}}^i; \mu) = \text{Poisson}(n_{\text{obs}}^i; n_b^i + \mu n_s^i). \quad (4.2)$$

The combined likelihood of observing this set of per-bin yields, given a value of the signal strength, is therefore given by the product of the per-bin likelihoods:

$$\mathcal{L}(\{n_{\text{obs}}^i\}; \mu) = \prod_i \mathcal{L}(n_{\text{obs}}^i; \mu). \quad (4.3)$$

4.8.2.2 Adding nuisance parameters

To account for the effect of external systematic uncertainties in a fit, nuisance parameters are added to the likelihood. This can be done per-bin or across the whole distribution, as required for each source of uncertainty.

Each nuisance parameter adds a constraint function as a factor to the likelihood. The constraint functions describe the behaviour of varying the nuisance parameter up or down from its nominal value, by any amount. Since the input for each systematic uncertainty is not a continuous function but just a nominal value, one upwards variation, and one downwards variation, the constraint function must be interpolated from these values.

Uncertainties affecting each bin individually, shape uncertainties, are linearly interpolated to evaluate the constraint function. These constraint functions are then added as multiplicative factors to the per-bin likelihoods given by Equation 4.2. Alternatively, constraint functions for systematic uncertainties affecting the overall scale, normalisation uncertainties, are calculated with an exponential interpolation and included as a multiplicative factor in the total likelihood of Equation 4.3 [92].

4.8.3 Maximum likelihood estimation

Given a likelihood, $\mathcal{L}(\mathbf{d}; \mu, \boldsymbol{\theta})$, where \mathbf{d} is the set of observed data and $\boldsymbol{\theta}$ is the set of nuisance parameters, the value of the parameters can be estimated by finding the set of their values that maximises the likelihood.

This is achieved in practice by minimising the negative logarithm of the likelihood through the Davidon-Fletcher-Powell approach [93, 94, 95] implemented in Minuit's MIGRAD algorithm [96]. The values of parameters that minimise the negative log likelihood are taken as the fitted values for μ and the nuisance parameters. Uncertainties for these parameters are given by the covariance matrix calculated during minimisation. The MINOS technique [96] is used to obtain a more accurate estimate of the uncertainties on μ .

4.8.4 Likelihood ratio tests

The likelihood ratio test is used to evaluate how significant a discrepancy is present between the observed data and the null hypothesis. This ratio is given by

$$\lambda(\mu) = \frac{\mathcal{L}(\mathbf{d}; \mu, \hat{\boldsymbol{\theta}})}{\mathcal{L}(\mathbf{d}; \hat{\mu}, \hat{\boldsymbol{\theta}})},$$

where $\hat{\mu}$ and $\hat{\boldsymbol{\theta}}$ are the parameter values that maximise the likelihood, and $\hat{\boldsymbol{\theta}}$ are the nuisance parameter values that maximise the likelihood for the given μ value.

Wilks' theorem demonstrates that the test statistic $-2 \log \lambda(\mu = 0)$ will be $\chi^2(1)$ -distributed under the null hypothesis [97]. Calculating the p -value from this gives the probability of this data being measured if the null hypothesis were true. This is typically rephrased as a significance, measured in standard deviations; this is equal to the deviation from zero of a normal distribution for which the two-sided integral of the tails would give the same p -value.

4.8.5 Significance thresholds

Established thresholds are used within particle physics to ensure that any discoveries made are robust. An 'observation' of a process requires the significance of an excess over the background-only hypothesis to be at least five standard deviations. If that threshold is not met, it can still constitute 'evidence' towards an observation if the significance exceeds three standard deviations.

Search for vector-boson scattering production of a Z boson and a photon

Vector-boson scattering (VBS) processes provide a unique experimental signature, producing decay channels with excellent potential to probe rare SM processes. The high selection efficiency achievable by exploiting kinematics of the VBS tag jets allows measurements to be conducted at lower cross sections than are accessible for similar processes with less unique topology.

In the archetypal VBS event, a quark from each of the two colliding protons radiates a boson. The two bosons interact to produce the EW component of the final state and the quarks, deflected from their original trajectories after boson emission, appear as jets in the detector. Since the initial quarks are usually very energetic, the angle through which they are deflected in the interaction tends to be small. The final-state jets, known as tag jets, would therefore be in the very forward regions of the detector, at opposite ends to one another, and also still carrying large amounts of energy. These kinematics are typically characterised by a large invariant mass of the dijet system (m_{jj}) and a large difference between the rapidities of the jets ($|\Delta y_{jj}|$).

VBS $Z\gamma$ serves as a production mechanism for the $Z\gamma jj$ final state, with the Z boson and photon resulting from a direct multiboson interaction and the jets created as




Figure 5.1: Feynman diagram for a $Z\gamma$ vector-boson scattering event (left). The black circle contains the multiboson interaction, which for a leading-order SM interaction will be one of the two shown (right).

a feature of the VBS production. This provides a robust framework for studying these rare SM interactions introduced in Chapter 1. Feynman diagrams for VBS $Z\gamma$ production are represented in Figure 5.1, showing that QGC or TGC vertices are the only SM contributions at LO.

Beyond the chance to study the rare SM vertices that contribute to this interaction, this analysis also offers sensitivity to any potential new physics that modifies or adds TGCs or QGCs [98]. These anomalous couplings would cause a deviation in the rate from the SM expectation which, if sufficiently large, would be measured in this analysis.

VBS $Z\gamma$ production is one component of the more general EW production of $Z\gamma jj$. The VBS production modes are not gauge-invariantly separable from others, so a direct measurement of VBS $Z\gamma$ is not strictly possible. Instead, EW $Z\gamma jj$ production is measured with a selection designed to enhance the VBS component, matching the kinematics of the jets with the expected VBS signature. Figure 5.2 gives Feynman diagrams for some non-VBS production modes that contribute to the EW production mechanism.

To measure this EW $Z\gamma jj$ production, background processes with the same final state must be understood. The dominant background for this analysis comes from




Figure 5.2: Example Feynman diagrams for non-VBS EW production of $Z\gamma jj$. In these instances one or none of the two final-state bosons are produced through multiboson interactions.




Figure 5.3: Example Feynman diagrams for QCD production of $Z\gamma jj$.

QCD $Z\gamma jj$ production. Figure 5.3 gives example Feynman diagrams for this QCD production, which differs from the EW mode as the strong force either provides the interaction between the two quarks or otherwise generates the final-state jets, resulting in colour-connected jets. Additional interactions between the colour-connected jets are very probable and will affect the observed jet kinematics, allowing these events to be distinguished from VBS events.

This analysis is the first iteration of a VBS $Z\gamma$ analysis using the full Run 2 dataset [66]. It builds on a measurement made with a 36 fb^{-1} partial Run-2 dataset [99], and the work has been continued (beyond what is presented here) in Reference [100]. Complementary measurements include one of the same process by the Compact Muon Solenoid (CMS) experiment [101], and an ATLAS measurement of VBS $Z\gamma$ with the Z boson decaying to two neutrinos [102]. This is part of a programme of ATLAS VBS measurements [103, 104, 105, 106], all contributing to the understanding of multiboson interactions.

The goal of this analysis is to measure the fiducial cross section of EW $Z\gamma jj$ pro-

duction in a region sensitive to VBS $Z\gamma$ production. The measurement relies on a cut-based selection, exploiting the VBS event kinematics. Signal and background processes are estimated, through a combination of MC simulation and data-driven estimates, and used to make a template fit to the dijet mass distribution. This chapter presents the analysis as a whole, with additional focus given to sections on jet collection investigations (Section 5.2), controlling jet flavour systematics (Section 5.5), and pruning of systematic uncertainties (Section 5.7).

5.1 Event selection

Event selection is applied to simulation, using the samples detailed in Section 4.2, and to data, as specified in Section 4.1, to reach the desired VBS-enhanced phase space. Cuts are first made to the $Z\gamma$ system following the prescription given in Section 4.4.

Imposing VBS-like kinematics on the jets further reduces the phase space. Jets are required to have a transverse momentum of at least 50 GeV, and a rapidity of less than 4.4 in magnitude. There must be at least two jets, and no b-tagged jets in the event. The dijet system, formed from the two highest momentum jets, should have a mass $m_{jj} > 150$ GeV and an inter-jet separation of $|\Delta y_{jj}| > 1.0$. A veto on ‘gap jets’, jets (reconstructed with $p_T > 25$ GeV) found in the rapidity region between the two VBS tag jets, is applied to exploit the difference between VBS jets and colour-connected QCD jets. A loose cut is placed on the centrality of the $ll\gamma$ system relative to the jets, $\zeta(ll\gamma) < 5$, where centrality is given by

$$\zeta(ll\gamma) = \left| \frac{y_{ll\gamma} - (y_{j_1} + y_{j_2})/2}{y_{j_1} - y_{j_2}} \right|, \quad (5.1)$$

where $y_{ll\gamma}$ is the rapidity of the $ll\gamma$ system and y_{j_i} is the rapidity of the i^{th} highest-energy jet.

These cuts define the analysis region, and are summarised in Table 5.1. This is further split into the SR and the QCD CR with tighter requirements on the $ll\gamma$

centrality: $\zeta(l\gamma) < 0.4$ defines the SR and $\zeta(l\gamma) > 0.4$ gives the CR used for the QCD $Z\gamma jj$ background estimate.

5.1.1 Fiducial region definition

As the result of this analysis is a fiducial cross-section measurement, the particle-level selection for which this measurement applies, i.e. the fiducial region, must be defined. This fiducial selection mimics the detector-level selection for the SR as closely as possible. Table 5.2 reports this selection. For more details of the exact parameters defining this fiducial region, see Reference [66].

5.2 Particle-flow jet validation

The choice of jet collection for this analysis is non-trivial. Particle-flow jets have recently become the standard recommendation within ATLAS in place of topocluster jets, but rather than applying that recommendation blindly, the specific case for this analysis is considered. For a discussion of jet collections, see Section 4.3.4.

The benefits of particle-flow include improved resolution for low-energy jets, although this only works within the acceptance of the ID. Typical VBS jets, on the other hand, are energetic and very forward, perhaps pushing into regions beyond ID acceptance. Although a great deal of VBS jets will still benefit from particle-flow, the phase space is very different from one which would make particle-flow jets an obvious choice. This section presents a comparison between particle-flow and topocluster jets in the analysis phase space, to justify the move to using particle-flow jets.

The procedure for comparing performance between these two jet collections is to investigate the difference in event yield from applying jet-based selection criteria with values calculated from either jet collection. The $Z\gamma$ selection, as in Table 4.4, is applied first, then jet-based cuts are applied and the resulting yields compared. For simplicity, this was investigated in the signal MC sample for only $Z \rightarrow ee$ events.

Table 5.1: Summary of event selection criteria defining the VBS analysis region. This region is further subdivided by additional criteria on $\zeta(ll\gamma)$.

VBS $Z\gamma$ selection	
$Z\gamma$ system	$Z\gamma$ selection (Table 4.4)
Jet	$N_j \geq 2$ $ y_j < 4.4$ $p_T^j > 50 \text{ GeV}$ $N_j^{\text{b-tag}} = 0$ $N_j^{\text{gap}} = 0$
Dijet system	$m_{jj} > 150 \text{ GeV}$ $ \Delta y_{jj} > 1$
Event	$\zeta(ll\gamma) < 5$

Table 5.2: Particle-level selection applied to events in the fiducial region. Included are approximate particle-level equivalents to the photon isolation and overlap removal applied for the signal region selection.

VBS $Z\gamma$ fiducial selection	
Lepton	$p_T^{l,1} > 30 \text{ GeV}, p_T^{l,2} > 20 \text{ GeV}$ $ \eta < 2.47, N_l \geq 2$
Photon	$E_T^\gamma > 25 \text{ GeV}, \eta_\gamma < 2.47$ $E_T^{\text{cone20}} < 0.07 E_T^\gamma$ $\Delta R(l, \gamma) > 0.4$
Jet	$p_T^j > 50 \text{ GeV}, y_j < 4.4$ $N_j \geq 2$ remove jets if $\Delta R(\gamma, j) < 0.4$ or if $\Delta R(l, j) < 0.3$
Event	$m_{ll} > 40 \text{ GeV}$ $m_{ll} + m_{ll\gamma} > 182 \text{ GeV}$ $m_{jj} > 150 \text{ GeV}$ $ \Delta y_{jj} > 1$ $N_j^{\text{gap}} = 0, \zeta(ll\gamma) < 0.4$

Looking at the overall yield of events after each cut shows that the two jet collections give very similar results, always within 1% of one another. Table 5.3 shows these yields. Differences can be further scrutinised by looking at each individual event; most events should result in the same decision, pass or fail, regardless of the jet collection chosen. This checks that the similar yields aren't merely a coincidence, when in fact many events pass only one selection. Figure 5.4 shows these per-event differences in decisions.

For the vast majority of events, the two jet collections make the same selection, with less than 5% of events showing differences. These differences are tested by investigating the distributions in the cut variables for cases where the two jet collections give a different result, in these instances the two values should still be similar. For example, if an event passes $m_{jj} > 150$ GeV with topo-cluster jets but not with particle-flow jets then the particle-flow m_{jj} value should be close to the cut boundary of 150 GeV.

Investigating the case where events are selected using particle-flow jets but not topo-cluster jets, results for the four key cut variables are shown in Figure 5.5. The distributions are largely as expected, given that values peak on the cut boundary

Table 5.3: Yields and efficiencies after each jet cut, compared for both topo-cluster and particle-flow jets. Starting from all EW $Z(\rightarrow ee)\gamma jj$ events that pass the $Z\gamma$ selection. Efficiencies given are for the individual cut, relative to the yield from the previous cut. The difference is given as percentage increase from the topo-cluster to the particle-flow yields. The third cut is overlap removal between jets and leptons or photons.

Cut	Topo-cluster		Particle-flow		Difference
	Yield	Eff.	Yield	Eff.	
$N_j \geq 2$	51084	79.7%	51468	80.3%	+0.7%
$p_T^j > 50$ GeV, $ \eta_j < 4.5$	31362	61.4%	31604	61.4%	+0.8%
$\Delta R(l, j), \Delta R(\gamma, j) > 0.4$	31359	99.99%	31552	99.84%	+0.6%
$ \Delta\eta_{jj} > 1.0$	27127	86.5%	27293	86.5%	+0.6%
$ m_{jj} > 150$ GeV	26752	98.6%	26885	98.5%	+0.5%




Figure 5.4: Impact of four of the key jet cuts on the analysis yield when applied individually with particle-flow and topo-cluster jets. Events are divided into four categories based on whether or not they pass the particle-flow cut and whether or not they pass the topo-cluster cut. Cuts are applied in the same order as presented in Table 5.3. Only events that pass the previous cut for both particle-flow and topo-cluster jets are included in the results, to decorrelate the effects of each individual cut. The label ‘PFlow’ is used for particle-flow cuts and ‘Topo’ for topo-cluster cuts.

and tail off for more extreme differences. However, it is notable that for the dijet variables, m_{jj} and $|\Delta\eta_{jj}|$, there are some strong outliers.

Events falling very far from the cut boundary for particle-flow jets when the topo-cluster jet variable fails the cut seem to indicate a significant disagreement in kinematics between the two jet collections. It is possible that these outliers happen when the p_T ordering of jets varies between collections. Dijet variables are calculated using the two highest energy jets, so a small shift in p_T between the second and third jets could cause dijet variables to be calculated with a different jet pair and therefore give very different results.

This hypothesis can be tested by looking at the separation between jets used in each




Figure 5.5: Distribution of events passing jet requirements for the particle-flow collection but not the topo-cluster collection. In each case the cut is in the same distribution as the histogram plotted. Four cuts are shown: $N_j > 2$ (top left), $p_T^{j,2} > 50$ GeV (top right), $m_{jj} > 150$ GeV (bottom left), $|\Delta\eta_{jj}| > 1$ (bottom right). Only events passing all prior cuts for both particle-flow and topo-cluster jets are included. The label ‘PFlow’ is used for variables calculated with particle-flow jets and ‘Topo’ for topo-cluster jets.

event for the two collections. The variable used to measure this is

$$\min \Sigma(\Delta R) = \min_{k_i \in \{(1,2), (2,1)\}} \sum_{i=1}^2 \Delta R(j_i^{\text{topo-cluster}}, j_{k_i}^{\text{particle-flow}}), \quad (5.2)$$

i.e. the sum of the two ΔR values between topo-cluster and particle-flow jets, for whichever pairing of the jets gives the lowest value of the sum. ΔR is the sum in quadrature of $\Delta\eta$ and $\Delta\phi$. Figure 5.6 shows the distribution of this variable for events passing m_{jj} for only particle-flow jets. There are two clear populations, separated at $\min \Sigma(\Delta R) \sim 0.5$. The lower $\min \Sigma(\Delta R)$ population should contain events where the particle-flow and topo-cluster jets are representing the same physical objects. Requiring $\min \Sigma(\Delta R) < 0.5$ on top of the existing selection gives the

m_{jj} distribution shown in Figure 5.6, where now all remaining events are tightly distributed around the cut boundary.




Figure 5.6: Distribution of the $\min \Sigma(\Delta R)$ variable, defined in Equation 5.2, for events passing $m_{jj} > 150$ GeV for particle-flow but not topo-cluster jets (left); and the m_{jj} distribution for these events after requiring $\min \Sigma(\Delta R) < 0.5$ (right). The label ‘PFlow’ is used for variables calculated with particle-flow jets and ‘Topo’ for topo-cluster jets.

This study, although limited in scope, serves to demonstrate that the key jet variables used in this analysis perform very similarly when calculated with particle-flow and topo-cluster jets. Any differences seen are sufficiently small that no meaningful effect on the analysis result is expected. This is considered as motivation to use particle-flow jets for this analysis, in keeping with the updated ATLAS recommendation. To really understand if the improvements in resolution with particle-flow jets are seen in the analysis phase space, further studies could be conducted on how systematic uncertainties differ between collections, to determine which would give the most precise result.

5.3 Background estimation

The dominant background for this analysis, QCD $Z\gamma jj$ production, suffers from known mismodelling for high dijet masses, which is the region of interest in this analysis. Rather than using the MC template to directly estimate this background,




Figure 5.7: Centrality distribution for data, signal, and background estimates pre-fit (before any data corrections to QCD $Z\gamma jj$). The dashed line marks the separation between the SR and QCD CR. The uncertainty band is the combination of uncertainties from background estimation, MC statistics, and experimental systematics. Overflow events are included in the last bin. [66]

the normalisation is corrected by comparing with data in a CR enriched in this background. The centrality variable is used to separate the SR ($\zeta(ll\gamma) < 0.4$) from this QCD CR ($\zeta(ll\gamma) > 0.4$). The CR is rich in the QCD background and has a very small fraction of signal events, as the EW production mechanism peaks at low values of centrality. Figure 5.7 shows the centrality distribution for signal and background events.

This normalisation correction is calculated by fitting a normalisation factor for the QCD background in the SR and CR simultaneously, allowing the overall normalisation to be adjusted according to data. The shape of the background is taken from MC, but data in the CR is used to validate this shape and constrain correlated uncertainties.

Estimation of the remaining backgrounds, Z+jets, $t\bar{t}\gamma$, and WZjj, is detailed in




Figure 5.8: Plots of relative variation of yields as a function of dijet mass, m_{jj} , for EW (left) and QCD (right) production of $Z\gamma jj$. The variations shown are the largest groups of systematics in the SR. [66]

Section 4.6.

5.4 Systematic uncertainties

The systematic uncertainties used in this analysis are discussed in Section 4.7. Figure 5.8 illustrates the relative effect of the largest groups of systematics is shown on both EW and QCD $Z\gamma jj$ production, as a function of m_{jj} . The uncertainty on the normalisation of the Z+jets fake photon background is calculated following the prescription in Section 4.6.2.3 and determined to be 35% [66].

5.5 Jet flavour uncertainties

The uncertainties associated with the flavour composition and response of the jets in this analysis make a significant contribution to the result. This section presents measures taken to reduce these uncertainties and in turn improve the precision of the final measurement.

These jet flavour uncertainties, as discussed in Section 4.7.2.2, can be reduced by specifying the expected fraction of jets initiated by (light) quarks and gluons in the analysis phase space. This is parameterised by the gluon fraction, given in Equation 4.1.

A given jet in a MC event is determined to be quark- or gluon-initiated from the truth record, using the `PartonTruthLabelID` variable. Measuring both f_g and its associated uncertainty as a function of jet pseudorapidity and transverse momentum, for events passing the analysis selection, provides the information needed to reduce the jet flavour uncertainty.

This calculation is performed on events in the QCD $Z\gamma jj$ MC sample, in the inclusive analysis region (defined by Table 5.1) and its subregions, the SR and the QCD CR. This study is not necessary for other samples as the uncertainties have a lesser impact on the final measurement.

Uncertainty on f_g arises from three sources: a modelling uncertainty calculated by comparing the f_g values obtained from two independent MC generators, statistical uncertainty resulting from the size of the MC sample used, and an additional uncertainty to cover any variations in the value of f_g between regions. The third uncertainty component is necessary due to technical limitations, which allowed only one f_g value to be provided for samples used to calculate yields in both the SR and the CR.

Calculation of f_g is performed for the QCD $Z\gamma jj$ sample only, as this is where the jet flavour uncertainty is largest.

Figure 5.9 shows the gluon fractions measured in the nominal and alternate MC samples for QCD $Z\gamma jj$, as well as calculations of each of the uncertainty components. The statistical uncertainty was found to be an order of magnitude smaller than the other components and so is neglected. The uncertainty to cover differences between regions is calculated by finding the largest difference, in each bin, between f_g in the inclusive region (SR+CR) and either of the two sub-regions. The overall uncertainty used is then the per-bin quadrature sum of the generator differences and this inter-region difference.




Figure 5.9: Gluon fractions and uncertainties as a function of jet pseudorapidity and transverse momentum. Plotted are gluon fraction in the nominal (a) and alternate (b) QCD $Z\gamma jj$ sample MC samples, the difference between gluon fractions in these two samples (c), the statistical uncertainty on gluon fractions in the alternate sample (d) (this was the larger of the statistical uncertainties), and the difference between gluon fraction in the SR+CR region and in the SR (e) and CR (f) regions.

5.6 Template fit

The signal strength, μ_{EW} , is used to parameterise the fiducial cross section for the signal process, σ^{EW} , where

$$\mu_{\text{EW}} = \sigma_{\text{meas}}^{\text{EW}} / \sigma_{\text{SM}}^{\text{EW}}, \quad (5.3)$$

i.e. the ratio of the measured cross section to the SM expectation. This signal strength is extracted from the data with a maximum likelihood fit, performed on m_{jj} distributions in the SR and CR simultaneously. MC distributions for backgrounds and signal are used as templates, with normalisations for the signal and QCD $Z\gamma jj$ background allowed to float in the fit.

Electron and muon channels are treated together, using the sum of events from both as input to the fit. A binned likelihood is built using the m_{jj} distribution in both the SR and CR, as described in Section 4.8.2, with systematic uncertainties included as nuisance parameters. The effect of each uncertainty on the normalisation and shape of the m_{jj} distribution is considered individually and a pruning system, described in Section 5.7, is used to reduce the number of nuisance parameters needed.

The fit extracts the value of the signal strength, μ_{EW} . From this a significance of the measurement under the background-only hypothesis is calculated, as described in 4.8.4. If the significance is greater than five standard deviations, EW $Z\gamma$ production is considered to be observed. The extracted signal strength is also used to measure the fiducial cross section of the process, given the SM expectation of this cross section: calculated from the nominal MC sample as

$$\sigma_{\text{SM}}^{\text{EW}} = 4.73 \pm 0.01 \text{ (stat.)} \pm 0.15 \text{ (PDF)}_{-0.22}^{+0.23} \text{ (scale) fb.}$$

5.7 Pruning systematic uncertainties

There are 74 experimental systematic uncertainties considered for this analysis. With each systematic requiring an up and down variations, 148 nuisance parameters would be needed in the fit. To limit the number of nuisance parameters, and

thus stabilise the fit, a system is developed to rank the impact of each systematic on both the shape and normalisation of the m_{jj} distribution. Only systematic uncertainties deemed to be significant are fully accounted for in the fit, and those with less impact are pruned.

The first test for a systematic uncertainty is how uniform its effect is across the dijet mass spectrum – this will indicate whether it will impact the shape of the m_{jj} distribution. If a systematic is determined to have a significant impact on shape, by criteria discussed below, then it is included in the fit with one nuisance parameter for each bin in m_{jj} , allowing it to modify the shape in the fitting process.

Any uncertainty not found to impact the shape should be assessed for how significant an impact it has on the overall normalisation of events. Systematic uncertainties with a large enough effect on the event yield will contribute one nuisance parameter to the fit, and have the ability to scale the overall normalisation. Any uncertainties with a smaller effect will be pruned, i.e. all pruned systematics will be added in quadrature as a single extra nuisance parameter to scale the overall normalisation in the fit.

Information on how nuisance parameters for shape and normalisation uncertainties are included in the likelihood fit is given in Section 4.8.2.2.

5.7.1 Calculating statistical uncertainties

In order to determine whether the effect of any systematic uncertainty, on shape or overall yield, is significant, the statistical error on the value of the systematic uncertainty must be ascertained. This statistical uncertainty arises from the finite size of MC samples used to evaluate systematic uncertainties.

The value of a systematic uncertainty on an event yield is given by

$$\sigma_{\text{Norm}} = \frac{N_{\text{varied}} - N_{\text{nominal}}}{N_{\text{nominal}}}, \quad (5.4)$$

where N_{nominal} is the number of events accepted for a nominal MC sample and N_{varied} is the number of events after the systematic variation has been applied. Each of N_{varied} and N_{nominal} has a statistical uncertainty. However, due to the fact that these variables are measuring the same set of events under different conditions, the two yields are highly correlated. The correlation is not known a priori, and so the uncertainty on σ_{Norm} cannot be calculated through error propagation.

The bootstrap method [107, 108] is instead used in order to determine statistical uncertainties while preserving correlations. This method relies on resampling the event set to create replica sets of events of the same size, with some events duplicated and some omitted. Calculating σ_{Norm} in each replica set gives a distribution of results for which the standard deviation represents the statistical uncertainty on σ_{Norm} .

5.7.2 Determining shape impact

For a systematic uncertainty that has no impact on the m_{jj} shape, it would be expected that the resulting variation is uniform across the m_{jj} distribution. This is tested by calculating the fractional difference in yield, σ_{Norm} , and its associated statistical uncertainty in bins of m_{jj} . A chi-squared test from fitting a zeroth order polynomial to these values provides a test statistic which should be distributed as $\chi^2(3)$ (4 bins minus 1 parameter for 3 degrees of freedom) under the null hypothesis of no shape impact. A significant shape uncertainty is therefore anything that deviates from this null hypothesis by more than a certain threshold. An example m_{jj} distribution and fit is shown in Figure 5.10, for a systematic uncertainty with an obvious impact on the shape.

The threshold chosen is a p -value of 0.05, i.e. chi-squared values sufficiently high that there is at most a 5% chance that the deviation arose from statistical fluctuation. Uncertainties with a p -value below this threshold have the full shape treatment in the fit, with per-bin nuisance parameters. Figure 5.11 shows the results of the chi-squared test for the largest experimental systematics in the SR. For the signal sample, 30 of the 74 experimental systematic uncertainties have a significant impact on shape in the SR, while for QCD $Z\gamma jj$ there are 18 significant shape uncertainties.




Figure 5.10: Binned m_{jj} distribution of measured systematic uncertainty, σ_{Norm} , for one systematic variation as labelled on the plot. The dashed line shows the zeroth-order polynomial fit attempted, with the indicated χ^2 value demonstrating this is clearly a poor assumption and this systematic does contribute an uncertainty on the m_{jj} shape.

5.7.3 Determining overall yield impact

Any systematic uncertainty determined to not impact the shape of the m_{jj} distribution can of course still affect the overall yield of events, and therefore require sufficient treatment in the fit. This could be done by assigning a single nuisance parameter to every remaining uncertainty, as all will have an effect on some scale. To further reduce the number of nuisance parameters required however, the less significant uncertainties are pruned.

In this case significance is determined by whether or not a systematic uncertainty is consistent with zero. Taking the value of the systematic, per Equation 5.4, and its statistical uncertainty as the standard deviation, if the value is within one standard deviation of zero it is considered insignificant.

All systematic uncertainties not passing the shape significance test but not consistent with zero have a dedicated nuisance parameter in the fit. All remaining uncertainties

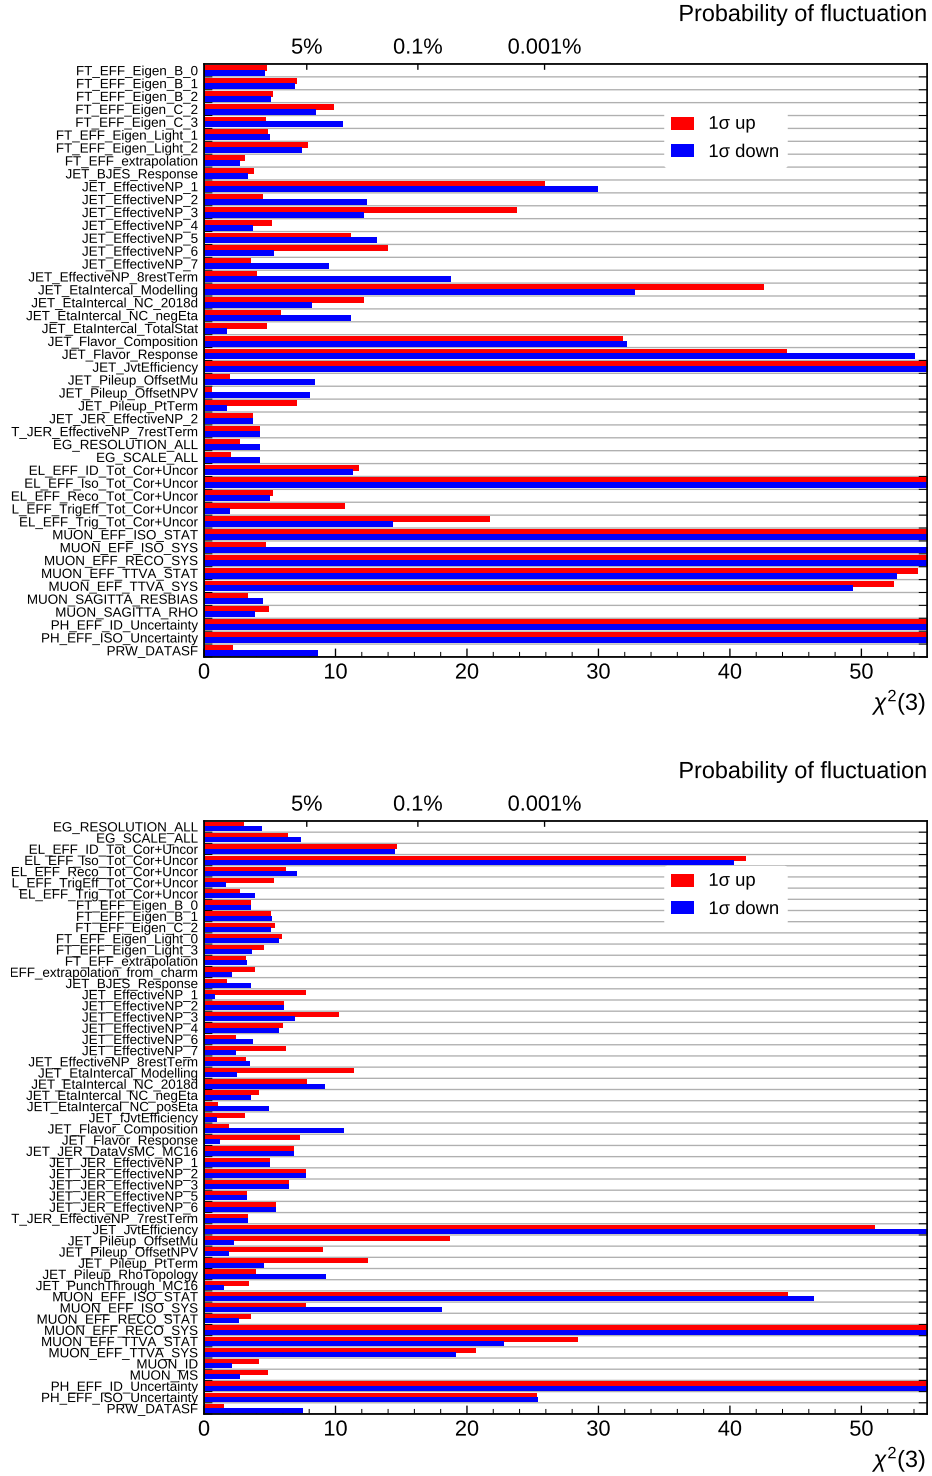


Figure 5.11: χ^2 values, representing the impact each systematic has on shape, for all experimental systematic variations in the EW (top) and QCD (bottom) samples in the SR. Only sources with χ^2 above 7 for either the up or down variation are shown. The top axis gives the probabilities of uncertainties arising fluctuations under the null hypothesis. The largest uncertainties extend beyond the range of the x -axis. A description of what each systematic uncertainty represents is given in Table 4.5 (some names are abbreviated here).




Figure 5.12: Overall normalisation uncertainty, σ_{Norm} , in the SR for the EW (top) and QCD (bottom) samples. Values for the largest experimental systematic uncertainties are shown, any with a value below 5×10^{-4} are omitted. The black bars represent the statistical uncertainty on the value. A description of what each systematic uncertainty represents is given in Table 4.5 (some names are abbreviated here).

are pruned. The relative change in yield from each of the largest experimental systematic uncertainties is shown in Figure 5.12.

Of the 44 uncertainties from the signal sample that do not have a significant shape impact in the SR, 11 are still found to impact the overall normalisation; the remaining 33 uncertainties are combined into a single normalisation uncertainty. For the QCD $Z\gamma jj$ background, 11 uncertainties affect the normalisation and the remaining 45 are combined. Across both samples, this represents a decrease of more than 50% in the number of nuisance parameters that must be considered, contributing to the stability of the fit.

5.8 Results

The fit gives a measured signal strength of

$$\begin{aligned}\mu_{\text{EW}} &= 0.95^{+0.14}_{-0.13} \\ &= 0.95 \pm 0.08 \text{ (stat.)} \pm 0.11 \text{ (syst.)}.\end{aligned}$$

This corresponds to an observed(expected) significance of 10(11) standard deviations, and is the first observation of this process by the ATLAS collaboration. As the measurement is consistent with $\mu_{\text{EW}} = 1$, the rate seen in data is consistent with the SM expectation.

This result is a marked improvement over the previous iteration of the analysis, which measured a significance of 4.1 standard deviations [99]. Nearly four times the amount of data was available for this analysis, which allows a reduced statistical uncertainty, but the improvement persists beyond that still, lowering the overall systematic uncertainty from $^{+19\%}_{-17\%}$ to $\pm 9\%$. This reduced uncertainty comes from several improvements across the analysis: increased data statistics in control regions and updated background estimation procedure to reduce uncertainty on the Z +jets background, larger MC simulation samples to reduce the previously significant MC statistical uncertainty, and greatly reduced systematic uncertainties on

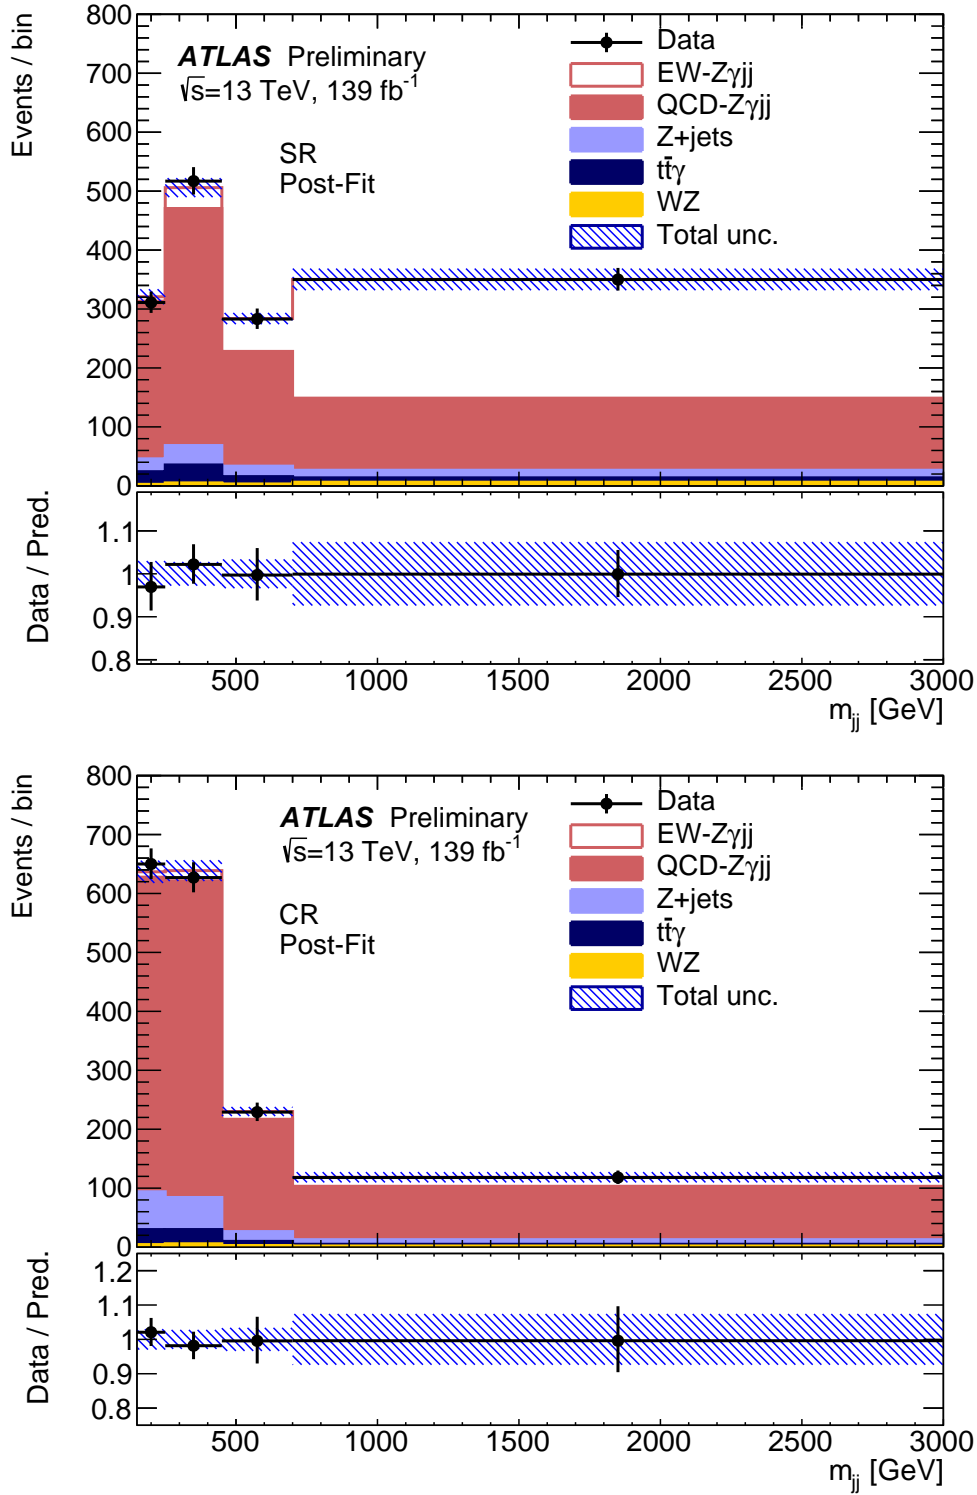


Figure 5.13: Post-fit distributions of dijet mass, m_{jj} , in the SR (top) and CR (bottom). The uncertainty band is the combination of all uncertainties, taken from the fit. Overflow events are included in the last bin. [66]

the jets thanks to improvements in the analysis methods such as those presented in Section 5.5.

The total yields from data and signal and background estimates in both the SR and CR are given in Table 5.4. Post-fit m_{jj} distributions in these regions are shown in Figure 5.13.

The fiducial cross section of the EW production of $Z\gamma jj$ in this VBS-like phase space is measured from the fit as

$$\sigma_{\text{EW}} = 4.49 \pm 0.40 \text{ (stat.)} \pm 0.42 \text{ (syst.) fb.}$$

Table 5.4: Yield estimates and associated post-fit uncertainties for each of the processes contributing to the signal region and control region, compared to data. The total estimate and its uncertainty is also given.

Process	Yield	
	SR	CR
EW $Z\gamma jj$	300 ± 36	55 ± 7
QCD $Z\gamma jj$	987 ± 55	1352 ± 60
$t\bar{t}\gamma$	72 ± 11	59 ± 9
WZjj	17 ± 3	14 ± 3
Z+jets	85 ± 30	143 ± 43
Total	1461 ± 38	1624 ± 40
Data	1461	1624

Search for triboson production of $VZ\gamma$ through its semileptonic decay mode

Triboson production of a Z boson; a photon; and an additional, hadronically decaying, vector boson is the second production mechanism for the $Z\gamma jj$ final state explored in this thesis. This additional boson, denoted as a ‘V’ boson, can be a W or a Z boson. Measuring this semileptonic $VZ\gamma$ triboson process thus constitutes an inclusive measurement of both $WZ\gamma$ and $ZZ\gamma$ triboson production.

Figure 6.1 gives LO Feynman diagrams for $WZ\gamma$ and $ZZ\gamma$ production. Notably, $WZ\gamma$ production is sensitive to the same QGC and TGC vertices as VBS $Z\gamma$ production. The $ZZ\gamma$ process does not feature these interactions in the SM description as there are no neutral QGCs or TGCs in the model. This makes the combined semileptonic channel something of a hybrid, it is sensitive both to the SM multiboson interactions but also sensitive to any beyond-SM physics which might introduce these ‘anomalous’ neutral couplings. As with the VBS $Z\gamma$ analysis, any anomalous couplings introduced from new physics would affect the rate of this process and, particularly if the cross section would be enhanced, could result in measuring a significant deviation from the SM prediction.

The high number of electroweak interactions necessary at LO to facilitate this process means that $VZ\gamma$ production has a very low cross section, similar to that of VBS




Figure 6.1: A selection of SM production mechanisms for the $VZ\gamma$ triboson final state, depicted in Feynman diagrams.

$Z\gamma$ (see Figure 1.3). Without the distinct VBS jet signature to select on, measuring this low cross-section process is challenging.

The two jets, here a product of a boson decay, have kinematic properties that help distinguish them from background events, notably: a dijet mass peaked around the W/Z masses, small rapidity separation between the two jets, and an angular distribution consistent with boson decay products. Additionally, more subtle differences are also present in other variables. Selecting W/Z boson decays to leptons is relatively straightforward, but the more limited jet resolution and more dominant jet background makes doing this in the hadronic decay channel more difficult. This analysis employs machine-learning techniques to interpret this complicated phase space; pushing sensitivity beyond what is achievable with a traditional cut-based analysis.

Despite the differing jet phase space, the backgrounds for this analysis include the same processes as the VBS $Z\gamma$ analysis. QCD $Z\gamma jj$ production is the dominant competing process; the key difference between this background and the signal is the kinematics of the jets, as both have a real Z boson and photon produced. The

QCD $Z\gamma jj$ background has a yield 140 times larger than the signal after applying preselection cuts (i.e. the analysis region selection as in Table 6.1). This illustrates the need for effective jet selection to manage this background.

The goal of this analysis is to measure the signal strength of this rare process in order to compare it to the standard model expectation. This measurement is extracted from a template fit to the signal rarity distribution, derived from the output of a machine-learning model. If the observed significance is sufficient, this will provide evidence for, or an observation of, this process. If the significance does not meet these thresholds, the measurement will be used to place limits on the rate of this process. These limits can be used to constrain theories that might enhance the cross section of this process.

This measurement represents the first of its kind, no measurements have been published on semileptonic $VZ\gamma$ production. There is however some overlap with other published measurements. Of the two included processes, $WZ\gamma$ and $ZZ\gamma$, $WZ\gamma$ has been observed by ATLAS through fully leptonic decay modes [109] and studied in a semileptonic final state in a CMS $VW\gamma$ measurement [110], using a similar principle to this analysis with a generic hadronically decaying massive boson. However, no measurements have been published of the $ZZ\gamma$ final state. These measurements contribute to the broader study of EW triboson processes, which includes recent measurements of VVV [111], WWW [112], $Z\gamma\gamma$ [113, 114], and $W\gamma\gamma$ [115, 114] processes.

The remainder of this chapter details the different elements of this analysis. An initial event selection and the definition for the signal process is given in Section 6.1. Before the full SR selection can be introduced, the development of a BDT discriminant is discussed in Section 6.2; this section motivates the need for this by first creating a cut-based selection for the analysis. Section 6.3 then defines the full SR, and some CRs, by making use of the BDT output. Background estimation procedures are reviewed in Section 6.4 before the systematic uncertainties are discussed in Section 6.5. Section 6.6 details the fitting procedure used to extract the measurement from data before the results are given in Section 6.7. The discussion concludes

with projections of future results and extensions to this analysis in Sections 6.8 and 6.9.

6.1 Event Selection

Events in data and simulation, from the samples discussed in Section 4.2, undergo selection to create an analysis region sensitive to $VZ\gamma$ triboson production. This section discusses an initial pre-selection as well as the additional requirements used to define the signal sample.

By implementing the selection defined here the analysis is focused on the dijet channel, in line with the goal of measuring $Z\gamma jj$ signatures. However, an additional channel could be defined for cases where the quarks from the $V \rightarrow qq$ decay are sufficiently close that they are detected as one large jet, typically if the V boson has large transverse momentum. This ‘merged jet’ channel offers an alternative approach to the two ‘resolved jets’ used here.

6.1.1 Analysis region definition

A selection is applied to events to impose a loose triboson-like phase space, before more precise signal and control regions are defined with the help of the BDT discriminant discussed in Section 6.2.

The $Z\gamma$ selection defined in Section 4.4 is first applied to events. Events are then required to have at least two jets, each with rapidity $|y_j| < 4.4$. The leading jet must have a transverse momentum of at least 40 GeV, and the sub-leading jet at least 30 GeV. The invariant mass of the dijet system, formed from the two leading jets, must satisfy $m_{jj} < 150$ GeV, and the rapidity separation of these jets $|\Delta y_{jj}| < 2$. These cuts are summarised in Table 6.1.

Many of these jet variables are later employed by the BDT to refine the selection but adding loose preselection reduces complexity at no cost to performance, see

Table 6.1: Summary of event selection criteria defining the $VZ\gamma$ analysis region.

$VZ\gamma$ selection	
$Z\gamma$ system	$Z\gamma$ selection (Table 4.4)
Jet	$N_j \geq 2$
	$ y_j < 4.4$
	$p_T^{j,1} > 40 \text{ GeV}$
	$p_T^{j,2} > 30 \text{ GeV}$
Dijet system	$m_{jj} < 150 \text{ GeV}$
	$ \Delta y_{jj} < 2$

Section 6.2 for a more detailed discussion. The m_{jj} cut ensures that this analysis is orthogonal to the VBS $Z\gamma$ analysis (which uses a cut of $m_{jj} > 150 \text{ GeV}$), and is also compatible with the expected signal values of $m_{jj} \sim m_W, m_Z$.

These cuts define the full ‘analysis region’, further cuts on the BDT output and m_{jj} are used to define the SR and CRs, discussed in Section 6.3.

6.1.2 EW $VZ\gamma$ definition

Triboson $VZ\gamma$ production forms a subset of the processes under the umbrella of EW $Z\gamma jj$ production. Only interactions where the two jets are a product of a boson decay should contribute to the signal process. Other forms of EW $Z\gamma jj$ production, such as the diagrams in Figures 5.1 and 5.2, should ideally be considered as a source of background.

This analysis defines two samples as orthogonal subsets of the EW $Z\gamma jj$ production sample detailed in Section 4.2: EW $VZ\gamma$ (the signal sample) and EW $Z\gamma jj$ background (or the EW background). These samples are separated using truth information on the kinematics and flavour of the jets.

Two variables are used to test if the jets are products of a W or Z boson decay: m_{jj}^{truth} , the invariant mass of the dijet system calculated at truth level, and `PartonTruthLabelID`, which indicates the flavour of the parton initiating each jet¹.

¹This variable informs on whether the parton is a quark or a gluon and the quark flavour (e.g. up, down, strange, etc.) but does not distinguish between a quark and an anti-quark, that information was not available in the sample.

The constraint $74 \leq m_{jj}^{\text{truth}} \leq 99$ GeV is applied for events included in the $VZ\gamma$ sample, chosen as it contains 95% of the combined W and Z boson lineshape and so should select 95% of $W/Z(\rightarrow jj)$ events. Events included in $VZ\gamma$ are also required to have `PartonTruthLabelID` values compatible with quark flavours from a W or Z decay, i.e. both jets are quark-initiated and either both the same flavour (e.g. both up quarks as in $Z \rightarrow u\bar{u}$) or one up-type and one down-type quark (e.g. an up and a strange quark as in $W \rightarrow u\bar{s}$). Any events failing either of these cuts are included in the EW background sample.

This truth-level classification is not 100% efficient and as such there is some cross-contamination between the samples. Nevertheless, applying this truth selection increases the probability that any event considered signal contains the physics processes of interest, direct multiboson interactions. Of the events passing the preselection in Table 6.1 for the full EW $Z\gamma jj$ sample, 31% are accepted to the $VZ\gamma$ signal sample and the remaining 69% make up the EW background.

6.2 Discriminating against QCD $Z\gamma jj$ production

The biggest challenge in this analysis is managing the dominant background, QCD $Z\gamma jj$ production. The background estimates given in Table 6.6 demonstrate the scale of this challenge. Like the signal process, this background has a real Z boson and photon. The difference is the origin of the jets, here not from a boson decay but from QCD production mechanisms, as shown in Figure 5.3. Identifying and exploiting the differences in jet kinematics between this background and the signal is therefore key to maximising the sensitivity of the measurement. This section is dedicated to discussing this problem; the word ‘signal’ is therefore used here to refer to EW $VZ\gamma$ production and ‘background’ refers solely to QCD $Z\gamma jj$ production.

There are a small number of kinematic distributions which exhibit a large difference between signal and background that could be exploited effectively by a simple cut. The dijet mass, m_{jj} , is an obvious example as it peaks around the W/Z boson mass

for the signal but is relatively flat for the background. For many more variables however, the differences are more subtle. There may be a clear difference in shape between signal and background but there is no obvious cut or set of cuts that would create a signal-rich region. Figure 6.2 shows some distributions with the largest signal-background discrepancies.

Creating an effective cut-based selection on these variables is relatively ineffective; a machine-learning discriminant can be used to improve sensitivity to the signal process. To demonstrate this, this section explores and compares two methods for defining a signal-sensitive phase space for the analysis: a cut-based approach and a BDT, a machine learning classifier introduced in Section 4.5.

The dijet mass variable is excluded from selections for both of these methods. This allows it to be used to define CRs with a low signal purity in order to validate background estimates with comparisons to data. For more detail on the definition and use of these CRs, see Section 6.3.

These initial investigations were performed before details of the analysis were finalised and so have a somewhat broader phase space, detailed below.

6.2.1 Phase space for preliminary studies

The studies presented in this section use events from the EW $VZ\gamma$ sample (as defined in Section 6.1.2) as the signal and from the QCD $Z\gamma jj$ sample as the background. All events are subject to the preselection in Table 6.2. These cuts select $Z\gamma$ events with a looser version of the full $Z\gamma$ selection presented in Section 4.4. No cuts are placed on the jets at this stage. Isolation, identification, and overlap removal for all objects are the same as discussed in Section 4.4.

6.2.2 Comparison metric

A metric is needed in order to evaluate the performance of a given selection. Since the desired selection will be one that grants the most sensitivity to the $VZ\gamma$ signal, a

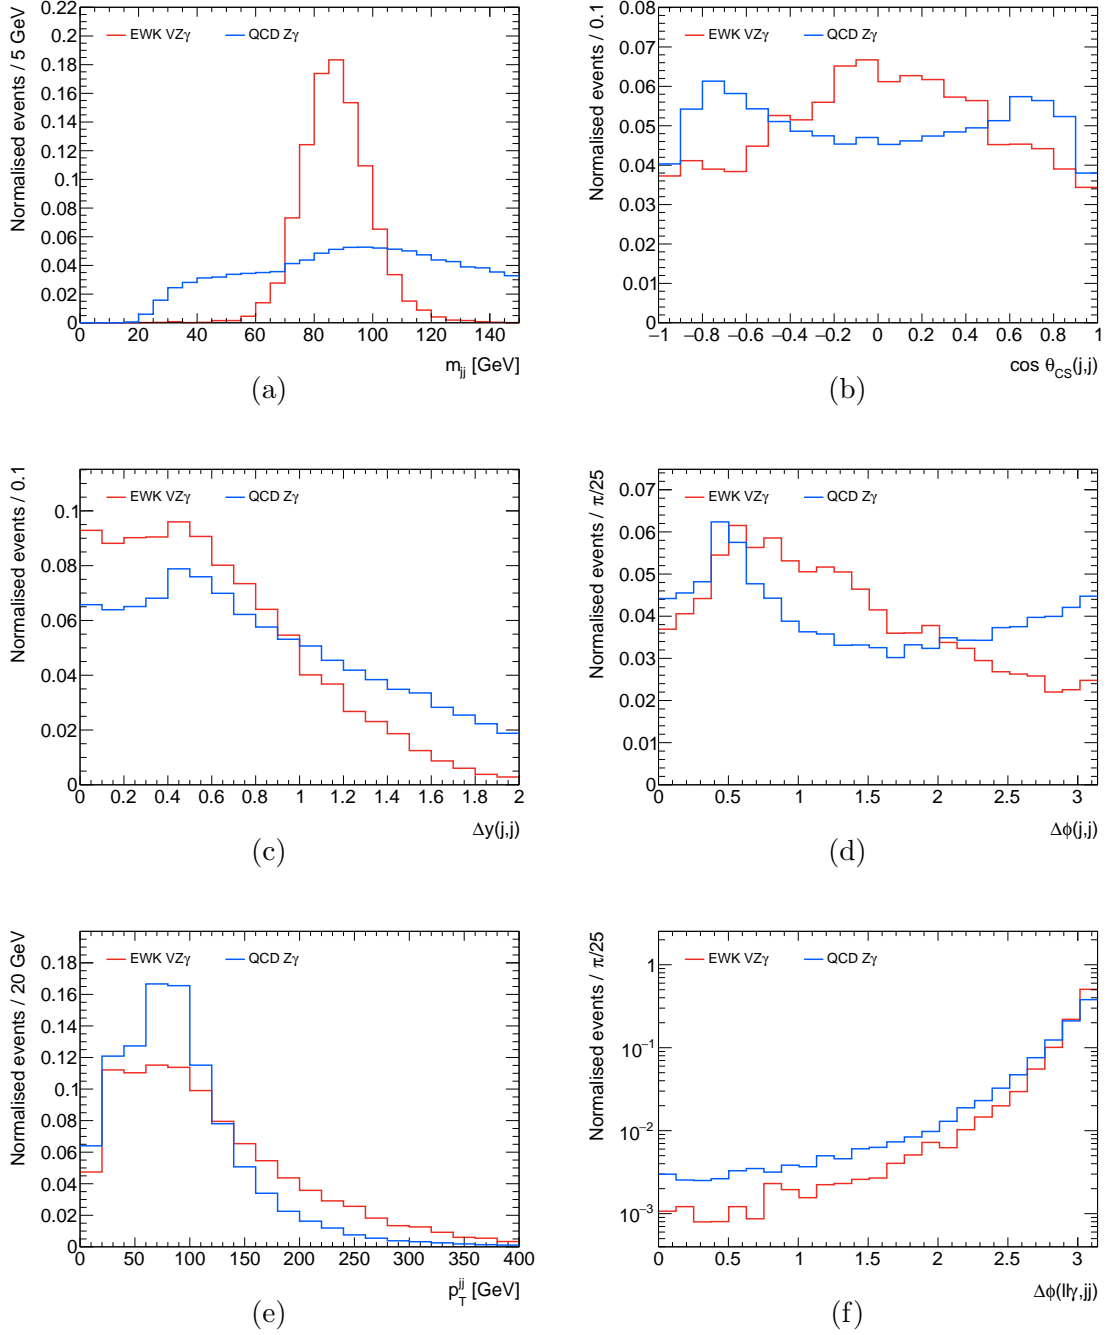


Figure 6.2: Kinematic distributions, comparing EW $VZ\gamma$ production (red) to QCD $Z\gamma jj$ production (blue), generated from the corresponding MC samples with the analysis region selection applied. Events are normalised to compare the shape of distributions between the two samples. Definitions for the variables shown are given in Table 6.3.

Table 6.2: Selection for events used in background rejection studies for the $VZ\gamma$ triboson analysis. This is the same as the $Z\gamma$ selection in Table 4.4 but with a looser photon p_T cut and no FSR cut.

Background rejection studies preselection	
Photon	$N_\gamma \geq 1$
	$ \eta_\gamma < 2.37$
	(excludes $1.37 < \eta_\gamma < 1.52$)
	$p_T^\gamma > 15 \text{ GeV}$
Lepton	$N_l = 2 \text{ (OSSF)}$
	$ \eta_e < 2.47$
	(excludes $1.37 < \eta_e < 1.52$)
	$ \eta_\mu < 2.5$
	$p_T^{l,1} > 30 \text{ GeV}$
	$p_T^{l,2} > 20 \text{ GeV}$
Boson	$m_{ll} > 40 \text{ GeV}$

significance of the SM-expected signal considering a background-only hypothesis is used. This will emulate the significance calculation used for the final measurement, though much simplified as it deals with only a single background and no systematic uncertainties. Whilst the significances given here are not directly comparable to that from the full measurement, they are comparable with each other and will indicate which selection generates more sensitivity to the signal process.

As the m_{jj} distribution is not used for selection, it is used here to calculate significance with a binned likelihood method. The likelihood is constructed as described in Section 4.8.2.1, and a likelihood ratio test is used to extract the signal. Expected yields for signal and background are taken from MC samples. To obtain integer values for the data yields in each bin, the significance is calculated many times in toy MC experiments: in each experiment the data values are drawn at random from a Poisson distribution with a mean equal to the sum of signal and background events in the relevant bin. Taking the mean significance from these toys gives the values used here.

These significances are calculated for each selection tested, given as a number of

standard deviations. As no systematics are used in this simplified performance metric, this represents statistical uncertainty only.

6.2.3 Selection variables

Building a selection to reject the QCD $Z\gamma jj$ background relies on identifying differences in jet kinematics, and therefore placing selection requirements on jet-based kinematic variables. A total of 22 of variables are considered, with the full list given in Table 6.3.

The variable p_T^{balance} is given by the equation

$$p_T^{\text{balance}} = \frac{(p_T^{jj} - p_T^{ll\gamma})}{(p_T^{jj} + p_T^{ll\gamma})}. \quad (6.1)$$

6.2.4 Cut-based background rejection

The task at hand is to find a set of cuts to make, on variables from Table 6.3, in order to maximise sensitivity to the signal process. Truly optimising this, finding the best value for each cut given the values of every other cut, is a many-dimensional problem with no easy solution. Instead an iterative approach is taken: find the best cut on each variable individually, take the cut which gives the best improvement in sensitivity and add it to the selection, then re-test all other cuts on the new subset of events.

Identifying the ‘best’ cut to make at any stage is a little subjective. For instance, when applying the first cut, the selection that would result in the best significance for the signal sample is likely too aggressive to allow for multiple effective cuts afterwards. The method used is to calculate background rejection (1/fraction of background events passing a cut) as a function of signal efficiency (fraction of signal events passing a cut) for each variable. By eye, these distributions can then be scanned to identify a possible cut which gives large background rejection but

Table 6.3: Variables considered for selection to reject QCD $Z\gamma jj$ events for the $VZ\gamma$ triboson analysis.

Variable	Definition
$y_{j,1}$	Rapidity of the leading jet in the event.
$y_{j,2}$	Rapidity of the sub-leading jet in the event.
y_{jj}	Rapidity of the jj system.
$p_T^{j,1}$	Transverse momentum of the leading jet in the event.
$p_T^{j,2}$	Transverse momentum of the sub-leading jet in the event.
p_T^{jj}	Transverse momentum of the jj system.
p_T^{balance}	Relative difference between transverse momenta of the jj and $ll\gamma$ systems, given by Equation 6.1.
N_j	Number of jets in the event, reconstructed with a minimum p_T of 25 GeV.
N_j^{gap}	Number of jets, satisfying $p_T > 25$ GeV found in the rapidity region between the two leading jets.
$m_{j,1}$	Mass of the leading jet in the event.
$m_{j,2}$	Mass of the sub-leading jet in the event.
$m(ll\gamma jj)$	Mass of the triboson system.
$ \Delta y_{jj} $	Absolute rapidity difference between the two leading jets.
$\Delta\phi_{jj}$	Smallest difference between the azimuthal angles of the two leading jets.
ΔR_{jj}	ΔR value between the two leading jets.
$ \Delta y(ll\gamma, jj) $	Absolute rapidity difference between the $ll\gamma$ and jj systems.
$\Delta\phi(ll\gamma, jj)$	Smallest difference between the azimuthal angles of the $ll\gamma$ and jj systems.
$\Delta R(ll\gamma, jj)$	ΔR value between the $ll\gamma$ and jj systems.
$\Delta R_{\min}(\gamma, j)$	Minimum ΔR value between any photon and jet in the event.
$\cos\theta^*(jj)$	Cosine of $\theta^*(jj)$, the angle of the leading jet in the dijet centre-of-mass frame relative to the direction of motion of the jj system.
$\cos\theta_{\text{CS}}(jj)$	Cosine of $\theta_{\text{CS}}(jj)$, the angle between the two jets in the Collins-Soper frame [116]. Jet charge information isn't available so the angle is taken relative to the leading jet.
$\zeta(ll\gamma)$	Centrality of the $ll\gamma$ system, given by Equation 5.1.




Figure 6.3: Distributions to identify a cut on $p_T^{j,2}$. Shown are fraction of events for each sample that are above a given threshold value in $p_T^{j,2}$ (left) and background rejection as a function of the signal efficiency achievable using the same $p_T^{j,2}$ threshold (right).

maintains a high signal efficiency. This allows for multiple variables to be included in the selection before the phase space becomes too constrained.

Figure 6.3 shows the background rejection against signal efficiency for $p_T^{j,2}$, which is chosen as the first variable to cut on. A cut of $p_T^{j,2} > 35$ GeV is chosen, with a signal efficiency of 74% and a background rejection factor of 2.6.

Continuing this process, the most performant selection found consisted of five cuts, listed in Table 6.4. Using the method described in Section 6.2.2, the significance calculated for events passing this selection is 1.2 standard deviations.

Table 6.4: Selection derived for baseline cut-based version of the analysis. Cuts are applied to the $VZ\gamma$ signal sample and the QCD $Z\gamma jj$ background for events passing the preliminary selection given in Table 6.2.

Cut-based selection
$p_T^{j,2} > 35 \text{ GeV}$
$ \Delta y_{jj} < 1.5$
$\Delta R(ll\gamma, jj) > 3.0$
$\Delta\phi(ll\gamma, jj) > 2.8$
$p_T^{\text{balance}} > -0.1$

6.2.5 BDT for background rejection

The cut-based selection provides a baseline performance against which to evaluate a BDT-based selection. The BDT can take many variables as input and determine how likely an event is to be signal or background based on the value of those variables, having first learned how the variables are distributed differently between signal and background events.

The first step is to train a BDT to identify these differences between signal and background. Once trained, the BDT is tested on an independent set of events to evaluate its performance and test for overtraining. To accommodate this train-test cycle, the signal and background samples are each split evenly into two, one half used for training and the other for testing.

Several aspects of the BDT are tuned to improve performance: the input variables used by the BDT, preselection applied to events before training, and hyperparameters of the BDT itself. These are discussed in the sections below.

6.2.5.1 Input variable selection

The benefit of the BDT is its ability to handle many input variables and generate a phase space sensitive to the signal. However, giving too many variables to the BDT creates an overly complex model and makes it prone to overtraining. Many iterations of input variables were tested to find a set that is sufficiently small to prevent overtraining but with enough variables to allow the BDT to give a good sensitivity.

For each set of variables tested, a simple overtraining check is used. For a cut on the BDT output resulting in a background rejection factor of 10, the corresponding signal efficiency is compared between the training sample and the test sample. Overtraining would result in a higher signal efficiency in the training sample than in the test sample. A requirement that the test sample signal efficiency is within 10% of the training sample is used to mitigate overtraining in the BDT model.

The sensitivity attained for a BDT trained on a particular variable set is evaluated by calculating the significance, through the method discussed in Section 6.2.2. To do this, a cut must first be placed on the BDT output. The value chosen for this cut will affect the sensitivity, so in each instance the cut value is scanned to find the highest attainable significance.

After using these tests to compare many combinations of variables, the most performant set was chosen. The final set of 16 input variables is shown in Table 6.5 ranked by their ‘importance’ as determined by the BDT. See Section 4.5.5 for a description of how variable importance is calculated.

Table 6.5: Ranking of variables used by the BDT to discriminate between signal and background for the $VZ\gamma$ analysis.

Rank	Variable	Relative importance
1	$ \Delta y_{jj} $	7.46×10^{-2}
2	$p_T^{j,2}$	7.27×10^{-2}
3	$\Delta\phi_{jj}$	7.24×10^{-2}
4	$m_{j,2}$	7.06×10^{-2}
5	p_T^{balance}	7.05×10^{-2}
6	$\Delta R_{\min}(y, j)$	6.50×10^{-2}
7	$y_{j,2}$	6.32×10^{-2}
8	$\Delta\phi(l\ell\gamma, jj)$	6.15×10^{-2}
9	$\cos\theta_{\text{CS}}(jj)$	6.10×10^{-2}
10	$p_T^{j,1}$	5.76×10^{-2}
11	$y_{j,1}$	5.70×10^{-2}
12	p_T^{jj}	5.68×10^{-2}
13	$\Delta R(l\ell\gamma, jj)$	5.68×10^{-2}
14	$m_{j,1}$	5.60×10^{-2}
15	$\log\zeta(l\ell\gamma)$	5.48×10^{-2}
16	y_{jj}	4.96×10^{-2}

The logarithm of the centrality, $\zeta(l\ell\gamma)$, is used rather than the linear form as this found to be more effective. This was due to the binning used by the BDT being insensitive to the signal-rich regions, as demonstrated in Figure 6.4.




Figure 6.4: Distribution of centrality, $\zeta(ll\gamma)$, both linear (left) and logarithmic (right) scales on the x -axis. Normalised event counts are shown for the $VZ\gamma$ signal sample and the QCD $Z\gamma jj$ background, for events in the analysis region.

6.2.5.2 Preselection and training cuts

Another route to improving performance of the BDT is constraining the phase space to further simplify the signature the BDT identifies. Even in cases where there is no performance increase, reducing the phase space without significant loss in signal efficiency can be beneficial as it may help to reduce the impact of systematic uncertainties. It also improves the interpretability of the analysis phase space; cuts on simple kinematic variables are more easily understood than a cut on a BDT output.

Two types of selection are used for this purpose: preselection applied to all events, including those input to the BDT, and training cuts which are applied only during BDT training. Preselection will narrow the whole analysis phase space whilst training cuts give the BDT a more focused view of the signal and background.

Three preselection cuts are applied, on top of the baseline selection for these studies given in Table 6.2. Minimum jet transverse momentum is included for both leading and sub-leading jets. Each is set to the highest value that did not degrade the sensitivity of the BDT: $p_T^{j,1} > 40$ GeV and $p_T^{j,2} > 30$ GeV. A requirement is also placed on the rapidity difference $|\Delta y_{jj}|$. Artefacts were found in the BDT response for background events with high $|\Delta y_{jj}|$; a cut of $|\Delta y_{jj}| < 2$ was found to remove these issues and have no impact on sensitivity. These preselection cuts contribute to the analysis region definition given previously in Table 6.1.

A training cut is made on the dijet mass, m_{jj} , to focus on a more signal-rich region. Applying a training cut of $60 < m_{jj} < 115$ GeV was found to improve BDT performance, and is well motivated by Figure 6.2a. Tighter mass window cuts were tested and no further improvements were found. This cut is only applied for training the BDT, and not as a preselection cut, to preserve its use for defining CRs.

6.2.5.3 Hyperparameter optimisation

A BDT implementation has hyperparameters that instruct it on how to build and boost decision trees during training. Four hyperparameters were investigated to optimise performance of the BDT used for this analysis: N_{cuts} , N_{trees} , d_{max} , and β ; each defined in Section 4.5.

Each parameter was tested in turn, training and testing the BDT to evaluate overtraining and sensitivity through the same procedure as in Section 6.2.5.1. Values for N_{cuts} between 2 and 500 were tested and the greatest sensitivity was achieved with $N_{\text{cuts}} = 90$, with no significant overtraining. Numbers of trees between 300 and 1500 were tested, with optimal sensitivity obtained for $N_{\text{trees}} = 850$. The d_{max} hyperparameter was tested for values from 1 to 9 and the sensitivity was found to increase for increasing d_{max} . However, deeper trees also became more prone to overtraining; a value of $d_{\text{max}} = 3$ was chosen to give the best balance in performance. The boost β parameter was tested with a range of values between 0 and 1, $\beta = 0.5$ was chosen with the best sensitivity and no significant overtraining.

6.2.5.4 Overall performance

With all of the optimisations made, the best significance obtained for events passing a BDT cut is 1.5 standard deviations, using the same statistics-only performance metric. This represents a sizeable improvement over the 1.2 standard deviations obtained with the cut-based approach, and motivates use of the BDT in this analysis.

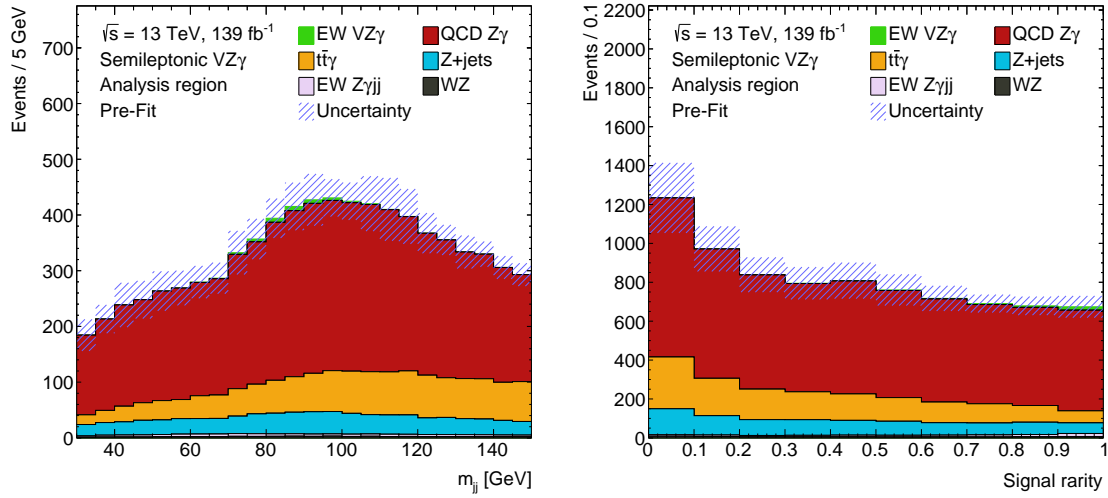


Figure 6.5: Distribution of dijet mass (left) and signal rarity (right) for events in the analysis region, defined by the selection in Table 6.1. Yields for the signal process and all backgrounds are shown stacked. The error band represents the total pre-fit uncertainty from all sources of systematic uncertainty.

6.3 Signal and control regions definition

Given the analysis region selection from Section 6.1 and the BDT discriminant developed in Section 6.2, additional selection cuts can be added to define signal and control regions for use in the fit. By applying tighter selections in constructing the SR, the analysis' sensitivity to the signal process is improved. Additionally, the use of orthogonal CRs with minimal signal leakage enables data-MC comparisons before unblinding, to confirm validity of background modelling, and also gives the fit more data with which to constrain systematic uncertainties.

One SR and three CRs are used for this analysis. The four regions are divided by two variables: m_{jj} and signal rarity. Signal rarity, or $\mathcal{R}_{\text{BDT}}^{\text{sig}}$, is a transformation of the BDT output defined in Section 4.5.4. These two variables are plotted in Figure 6.5, for events in the analysis region.

The dijet mass distribution is split into three regions: a lower sideband ($30 < m_{jj} < 65$ GeV), the on-peak region ($70 < m_{jj} < 100$ GeV), and an upper sideband ($110 < m_{jj} < 150$ GeV). The lower and upper sidebands form CRs for the analysis, and the on-peak region is further divided by a cut on signal rarity into the SR ($\mathcal{R}_{\text{BDT}}^{\text{sig}} > 0.8$) and the BDT CR ($\mathcal{R}_{\text{BDT}}^{\text{sig}} < 0.8$).




Figure 6.6: Approximate significance calculated with s/\sqrt{b} for a number of signal events s and background events b with an m_{jj} value between the minimum value given on the x -axis and the maximum on the y -axis. The number of signal events is calculated from the EW $VZ\gamma$ sample and all background samples are included for the background estimate. Events are required to pass the analysis region selection from Table 6.1. The maximum significance is obtained for a cut of $70 < m_{jj} < 100$ GeV.

The m_{jj} cut defining the on-peak region was chosen by approximating the significance obtained for each pair of minimum and maximum m_{jj} cuts, given the number of signal and background events from all samples passing the cut. Figure 6.6 shows this 2D significance scan. Note that these are approximate statistical-only significances, and peak at a value of 1.0σ ; the difference between this and the 1.5σ from Section 6.2 is a combination of a less sophisticated significance calculation and the introduction of the FSR cut (which is necessary to select the physics processes of interest).

A gap is included between the on-peak region and the sidebands to minimise signal leakage; this is chosen such that no more than 5% of the total signal events fall in the sideband regions. Figure 6.7 shows the m_{jj} cuts employed, in the context of the




Figure 6.7: Illustration of the three m_{jj} regions used in the analysis: the lower sideband ($30 < m_{jj} < 65$ GeV), the upper sideband ($110 < m_{jj} < 150$ GeV), and the peak ($70 < m_{jj} < 100$ GeV) region which is then subdivided into the SR and BDT CR. The distributions shown represent events passing the analysis region selection for both the signal (shown in green) and the sum of all backgrounds (in black). Both distributions are normalised by their total event yield.

shape of the signal and background distributions.

6.4 Background estimation

This analysis considers the same background processes discussed in Section 4.6: QCD $Z\gamma jj$, Z +jets, $t\bar{t}\gamma$, and $WZjj$. In addition to these, EW $Z\gamma jj$ events not matching the signal definition given in Section 6.1.2 are also considered as a background. Z +jets, $t\bar{t}\gamma$, and $WZjj$ are all estimated following the procedure in Section 4.6. The QCD $Z\gamma jj$ and EW $Z\gamma jj$ backgrounds are taken directly from their MC estimates.

More complex treatment for the QCD $Z\gamma jj$ background was considered, as in the VBS $Z\gamma$ analysis. However, the phase space of this analysis is not known to be affected by the mismodelling for high dijet masses. Data-MC comparisons, made

Table 6.6: Estimates for the yields of signal processes and all backgrounds, with associated uncertainties. Uncertainties are calculated following the prescription in Section 6.5, and not adjusted by the fit (i.e. pre-fit uncertainties).

Process	Estimated yield	
	Analysis region	Signal region
EW $VZ\gamma$	41.9 ± 2.1	22.0 ± 1.2
QCD $Z\gamma jj$	5820 ± 770	378 ± 30
$t\bar{t}\gamma$	1370 ± 210	50.5 ± 7.6
Z +jets	787 ± 108	73.0 ± 9.9
EW $Z\gamma jj$	94.4 ± 5.3	9.12 ± 0.76
$WZjj$	55.9 ± 11.2	4.22 ± 0.85
Total	8170 ± 810	537 ± 34

in the three analysis CRs, show that predictions are consistent with the observed event yields. It is therefore considered that additional normalisation factors are not mandated, the estimate from MC is sufficient.

The full yields from all background estimates are given in Table 6.6. This highlights the relative scale of the signal and background processes, and the improvement gained from the SR selection.

6.5 Systematic Uncertainties

This analysis considers the sources of systematic uncertainty discussed in Section 4.7. As well as being applied to the signal and QCD $Z\gamma jj$ background, experimental and theory systematic uncertainties are applied for the EW $Z\gamma jj$ background. Due to the adoption of more standardised ATLAS tools [117], a different pruning procedure is used here to that of the VBS analysis, and is discussed below.

Some of the theory uncertainties discussed in Section 4.7.1 are omitted from this analysis. Uncertainties on the signal process from choice of parton showering and underlying event model are not included, and the uncertainty for interference between EW and QCD $Z\gamma jj$ production is also missing. These are not expected to be a significant omission due to the small signal yield. Uncertainties from choice of scale

and PDF set are included for the QCD $Z\gamma jj$ background and both EW $Z\gamma jj$ samples. The QCD modelling uncertainty is included and calculated using the difference of generators method.

6.5.1 Pruning

Given the large number of systematic uncertainties, a pruning procedure is implemented in order to reduce the number of nuisance parameters necessary in the fit. The pruning used for this analysis is less detailed than the one used for the VBS $Z\gamma$ analysis, it does not rely on statistical uncertainties on the estimates of systematic uncertainties.

Pruning is done individually in each of the four regions used in the fit. The impact of each uncertainty on the normalisation and shape of the signal rarity distribution is considered; shape and normalisation impact are decoupled such that either can be removed if the effect is small. As a result, there are four outcomes for each systematic uncertainty: it is retained in full with normalisation and shape effect, its shape effect is dropped but normalisation kept, its normalisation effect is dropped but shape effect retained, or the uncertainty is dropped entirely.

If the normalisation effect of an uncertainty is retained in the fit, one nuisance parameter is included which allows the uncertainty estimate to be adjusted by the fit, changing the overall normalisation for the associated background or signal estimate. When the shape effect of a systematic uncertainty is used in the fit, per-bin nuisance parameters are used which allow the yield in each bin to be adjusted by the fit; these per-bin parameters are constrained so as not to affect the overall uncertainty.

The threshold for dropping a normalisation component of an uncertainty is set at 0.2%, i.e. the normalisation is dropped from the fit if its estimated effect on the overall normalisation of the sample is less than 0.2% of the yield. The threshold for dropping a shape component is set at 99.8%. In this case there is a threshold in the probability of the uncertainty having a different shape to the nominal distribution. The probability is calculated through the Kolmogorov-Smirnov (KS) test [118, 119,

120]; the p -value given by the test represents compatibility between the nominal and systematic varied distributions. If the p -value is greater than 0.998 (99.8%) then the differences are considered sufficiently small and the shape component is dropped.

The results of the pruning are shown in Figure 6.8, where for each sample, region, and background the treatment of each systematic uncertainty is indicated.

6.6 Template fit

The signal process is measured through a fit to the signal rarity distribution in the signal and control regions. The signal strength of EW $VZ\gamma$, μ_{EW} , is the parameter of interest (PoI) in the fit. This parameter follows the definition in Equation 5.3, such that a measured value of $\mu_{EW} = 1$ means that the process is measured to occur at rate expected in the SM.

Estimates for each background are given as templates to the fit, and combine with a signal estimate taken from MC to give the total predicted events per-bin in signal rarity. The fit adjusts the value of μ_{EW} , as well as the values of the nuisance parameters representing systematic uncertainties, to best match these templates to the data yield observed in each bin of the distribution. This is achieved through the likelihood construction and maximisation techniques discussed in Section 4.8.

Four bins are used for the signal rarity distribution in the SR. This binning creates some significant MC statistical uncertainties, but provides a balance between these uncertainties and sensitivity to the signal. The low MC statistics are a side effect of the heavily constrained phase space necessary to measure such a low-rate process.

6.6.1 Fit closure

To test whether the fitting procedure is stable and self-consistent, a fit is performed using 'pseudo-data' in all regions. This pseudo-data setup runs the fit with 'data' yields equal to the total expected MC yield in all regions. By construction, this




Figure 6.8: Pruning results for systematic uncertainties in the $VZ\gamma$ analysis. The colours indicate whether a systematic shape and normalisation uncertainty components were each retained for the fit or dropped, for each sample and region used in the fit.

should give a fitted value of $\mu_{EW} = 1$; any significant deviation would indicate a problem with the fit. This gives a signal strength and expected data uncertainty of

$$\mu_{EW} = 1.00 \pm 1.13,$$

indicating a healthy fit.

6.6.2 Mixed fit

In order to estimate the full sensitivity of the analysis without using observed data in the SR, a mixed data–pseudo-data fit is used: here the observed data yields are used in the three CRs and pseudo-data in the SR. These pseudo-data are generated by first performing a fit to data in the CRs with the value of μ_{EW} fixed to 1. This allows the values of the systematic uncertainties to vary and account for any small data-MC discrepancies. The post-fit values for these parameters are then used in the estimate for the number of events in the SR used to generate the pseudo-data.

The results of this fit represent the SM expectation for the analysis results, and thus demonstrate the sensitivity. Full expected results are presented in Section 6.7 alongside the observed results.

Running this mixed fit gives a fitted value for the μ_{EW} parameter of

$$\begin{aligned} \mu_{EW} &= 1.60^{+1.20}_{-1.15} \\ &= 1.60^{+0.95}_{-0.92} (\text{stat.})^{+0.68}_{-0.63} (\text{syst.})^{+0.31}_{-0.29} (\text{MC stat.}) \end{aligned} \tag{6.2}$$

where the component of the error from MC statistics has been factored out of the systematic uncertainty (for this instance only). This gives a signal strength that appears to be greater than one, despite not including data in the SR. This may be a bias introduced by the large MC statistical uncertainties, with their contribution to the total uncertainty indicated in Equation 6.2. However, since this μ_{EW} value is consistent with one, at the $\sim 2\sigma$ level considering the MC-statistics error, the effect is not significant.

6.6.3 Data fit

Once the SR is unblinded, the fit can be performed using observed data yields in all four regions. As with the mixed fit, the μ_{EW} value and all nuisance parameters are minimised simultaneously across all regions, allowing their values to be constrained by data in CRs as well as the SR. Results from this fit are presented in Section 6.7.

6.7 Results

The signal strength for EW $VZ\gamma$ measured from the full fit is

$$\begin{aligned}\mu_{EW} &= 1.41^{+1.20}_{-1.14} \\ &= 1.41^{+0.93}_{-0.90} (\text{stat.})^{+0.75}_{-0.70} (\text{syst.}),\end{aligned}$$

and is compatible with the SM expectation. Post-fit distributions in the four regions are shown in Figure 6.9, and the corresponding yields are given in Table 6.7.

The observed significance of the signal process is 1.24 standard deviations, compared with an expected significance of 1.40 standard deviations. This does not meet

Table 6.7: Post-fit yields and uncertainties in each of the four regions included in the fit, and additionally for the final bin of the signal region. Yields are shown for each signal or background process individually, for the total signal+background yield, and for data.

Process	Yield				
	BDT CR	m_{jj}^{low} CR	m_{jj}^{high} CR	SR	$\text{SR}(\mathcal{R}_{\text{BDT}}^{\text{sig}} > 0.95)$
EW $VZ\gamma$	17 ± 14	1.4 ± 1.5	2.1 ± 2.0	31 ± 25	15 ± 12
QCD $Z\gamma jj$	1360 ± 40	1320 ± 50	1970 ± 60	383 ± 15	112 ± 7
$t\bar{t}\gamma$	310 ± 40	206 ± 25	580 ± 70	50 ± 6	10.0 ± 1.2
Z+jets	157 ± 23	150 ± 16	270 ± 60	71 ± 9	18.4 ± 2.2
EW $Z\gamma jj$	17.6 ± 0.8	25.7 ± 0.7	27.8 ± 0.7	9.2 ± 0.4	3.34 ± 0.16
WZjj	12.4 ± 2.5	12.1 ± 2.4	18 ± 4	4.2 ± 0.8	1.34 ± 0.27
Total	1884 ± 28	1718 ± 31	2870 ± 40	549 ± 21	159 ± 11
Data	1931	1697	2866	530	162

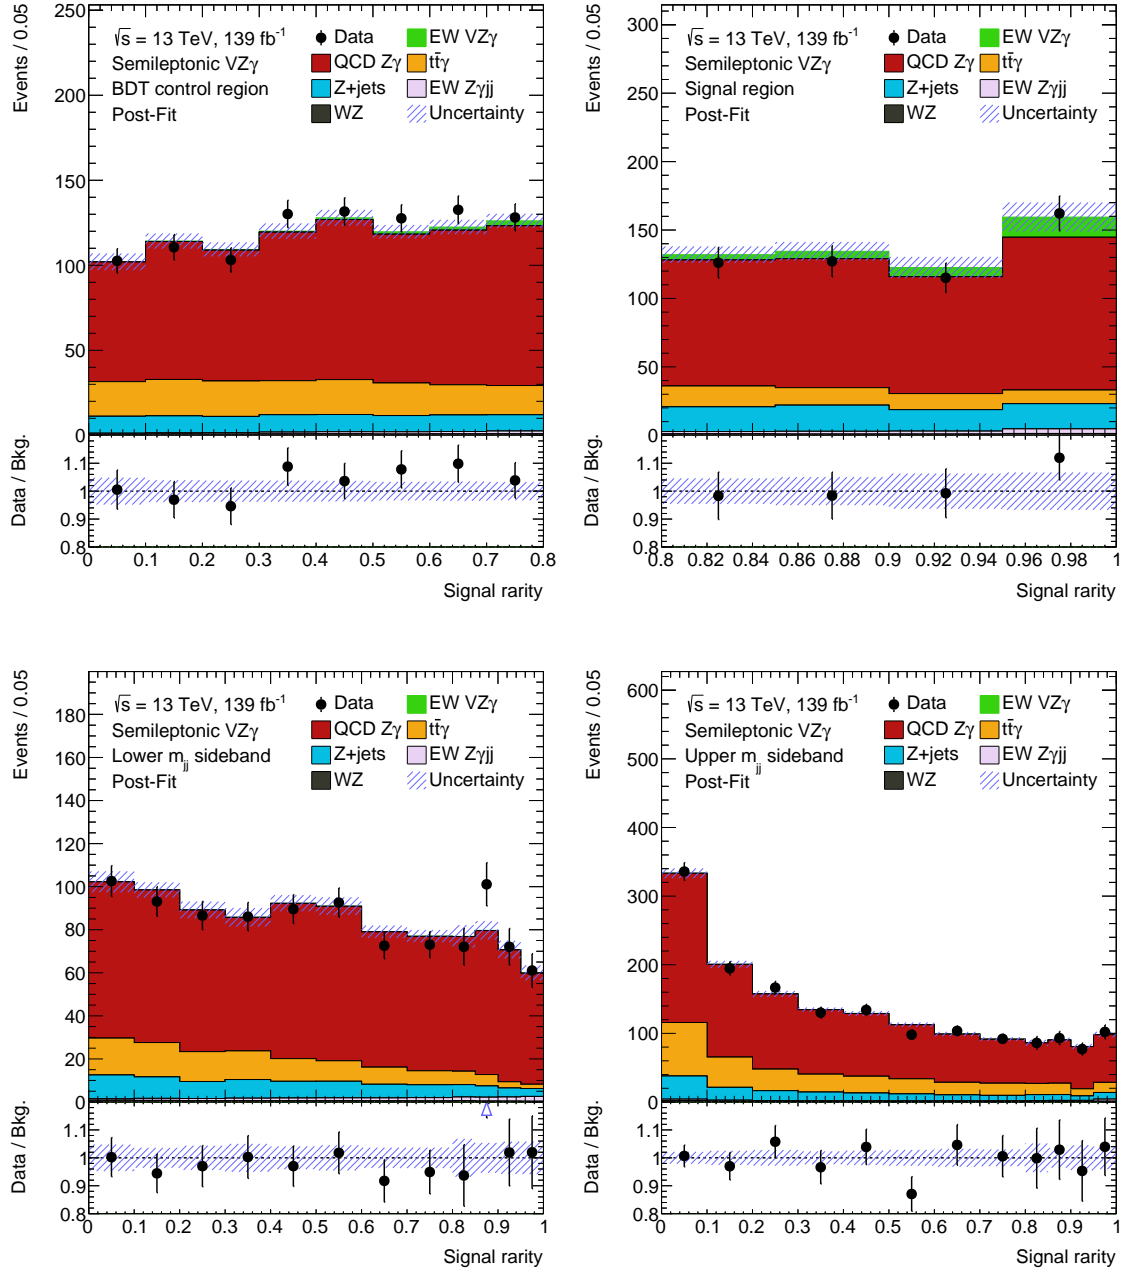


Figure 6.9: Post-fit signal rarity distributions in each of the four regions used in the fit, as labelled. Uncertainty bands represent the combined uncertainties in each bin, with values constrained by the fit. Uncertainty on data is due to statistics. The lower sections of each plot give the ratio of the data to the total background estimate.




Figure 6.10: Systematic uncertainties ranked by their post-fit impact on μ_{EW} . Uncertainties labelled γ represent MC statistics uncertainties in the given bin.

the threshold to provide evidence on the existence of the process. Instead a 95% confidence level upper limit is set on the rate of production for signal events at 3.46 times the SM expectation. This can be used to constrain any new physics models that would enhance the cross-section for triboson $VZ\gamma$ production.

Statistical uncertainties make the largest contribution to the measurement, but systematic uncertainties make a significant contribution. The largest systematic contributions are shown in Figure 6.10. Pileup reweighting is the largest individual contribution, likely due to the limited data statistics (see Section 6.8). The second largest contribution is from jet flavour composition, and several more of the largest uncertainties are MC statistics uncertainties in signal bins; as these uncertainties should be reducible, the effect of reducing some of these systematic uncertainties is discussed in Section 6.8.




Figure 6.11: Systematic uncertainties ranked by their post-fit impact on μ_{EW} , for a projected fit scaled to a luminosity of 420 fb^{-1} . Uncertainties from MC statistics and jet flavour were removed from this fit.

6.8 Projected results

To test what sensitivity might be possible with further optimisations or additional data available for this analysis, projected future results are explored. Firstly the reducible uncertainties, jet flavour composition and response and MC statistics, are removed. This simulates processing larger datasets, to reduce MC statistics uncertainties, and including gluon fraction information, to reduce jet flavour uncertainties. With the existing analysis and dataset this gives a measurement of

$$\mu_{EW} = 1.43 \pm 1.08,$$

calculated from performing a fit without these uncertainties included.

This is a small improvement by itself, but combined with an expanded dataset this could greatly enhance sensitivity. By scaling up the luminosity of the templates in

the fit, performing the analysis with a larger dataset can be simulated. This is a naïve estimate as, with a significantly larger dataset, the analysis would need to be re-optimised to take advantage of the available data.

To estimate sensitivity possible with a Run 2 and Run 3 combined dataset, templates are scaled to a luminosity of 420 fb^{-1} . The measured signal strength from this fit is

$$\begin{aligned}\mu_{EW} &= 1.00 \pm 0.63 \\ &= 1.00 \pm 0.53 (\text{stat.}) \pm 0.35 (\text{syst.}),\end{aligned}$$

corresponding to a significance of 2.09 standard deviations. This falls short of the evidence threshold of 3 standard deviations, but with a re-optimised analysis this channel could get close to the sensitivity required. The potential sensitivity from adding a merged jet channel may be enough to give a significant measurement for this process.

The largest systematic uncertainties for the 420 fb^{-1} projected fit are shown in Figure 6.11. It is noticeable that pileup reweighting is still the dominant systematic uncertainty, but much reduced from its post-fit scale seen in Figure 6.10. From running fits with luminosities scaled beyond 420 fb^{-1} a continued reduction in the impact of this uncertainty is observed; the conclusion is that this uncertainty is inflated by the small phase space of the analysis.

6.9 Extensions

With additional time investment, several aspects of this analysis could be improved. It is unlikely that this would result in a drastically more significant result, as such it would make sense to implement these changes at such a time when additional data are available.

The biggest addition would be a merged jet channel. This was not studied and so the sensitivity is unknown but, as most of the hadronically decaying bosons are not expected to have particularly high transverse momenta, it is likely a relatively

small enhancement to the existing sensitivity. Regardless, this remains the most promising route to attaining a more significant result for this process.

There are some issues with systematic uncertainties in the result presented. The first aspect of this is that some of the theory uncertainties on the signal process are missing, as detailed in Section 6.5. There are also a number of uncertainties which could be reduced, as explored in Section 6.8. Generating additional MC samples would be a priority for addressing this: the low statistics available not only contribute dominant systematic uncertainties to the result, but could also be biasing the signal strength measurement (see Section 6.6.2). This can be achieved either by creating larger samples, or by creating samples tuned for the desired phase space.

Conclusions

This thesis has presented research work carried out between September 2019 and December 2023, on both upgrades to the ATLAS L1Calo trigger and analysis of data in search of rare EW SM processes in the $Z\gamma jj$ final state.

The presented developments to the L1Calo trigger will, alongside the work done by the rest of the L1Calo community, improve the amount of data ATLAS is able to record across two phases of trigger upgrades, lasting for more than a decade. The eFEX visualisation tool has already been used to highlight bugs in the Run-3 system, which were subsequently corrected. The early Run-3 data analysis contributed to the fine-tuning of the e/γ trigger. Meanwhile, the E_{ratio} algorithm development and performance studies establish an algorithm available for the next iteration of the Level-1 e/γ trigger.

Together with an analysis team, the VBS $Z\gamma$ analysis presented in Chapter 5 was able to observe the targeted process with a significance greatly exceeding five standard deviations. The fiducial cross section of this process is measured as

$$\sigma_{\text{EW}} = 4.49 \pm 0.40 \text{ (stat.)} \pm 0.42 \text{ (syst.) fb.}$$

This process had not been previously observed by ATLAS and contributes to a programme of VBS studies that probe the SM multiboson interactions.

Finally, the semileptonic $VZ\gamma$ analysis produced a measurement for the signal strength of this relatively unexplored process, presenting an upper limit on its production at 3.5 times the SM-predicted rate, at the 95% confidence level. This analysis tackled a

difficult phase space with very large backgrounds, but also provides strong indication of the feasibility of future studies of the process. Finding evidence or an observation with a combined Run-2 and Run-3 dataset is expected to be very challenging, but with a re-optimised analysis and additional channels, and perhaps some upgraded analysis techniques, it could be possible.

References

- [1] S. Weinberg. *The Quantum Theory of Fields*, volume 1. Cambridge University Press, Cambridge, first edition (1995). ISBN 0521550017.
- [2] P. A. M. Dirac. *A theory of electrons and protons*. Proc. R. Soc. London. Ser. A, Contain. Pap. a Math. Phys. Character 126(801) (Jan 1930) pp. 360–365. ISSN 0950-1207. <http://dx.doi.org/10.1098/rspa.1930.0013>.
- [3] W. Pauli. *Über den Zusammenhang des Abschlusses der Elektronengruppen im Atom mit der Komplexstruktur der Spektren*. Zeitschrift für Phys. 31(1) (Feb 1925) pp. 765–783. ISSN 00443328. <http://dx.doi.org/10.1007/BF02980631>.
- [4] V. F. Weisskopf. *The development of field theory in the last 50 years*. Phys. Today 34(11) (Nov 1981) pp. 69–85. ISSN 19450699. <http://dx.doi.org/10.1063/1.2914365>.
- [5] M. Gell-Mann. *The eightfold way: A theory of strong interaction symmetry* (1961). <http://dx.doi.org/10.2172/4008239>.
- [6] M. Thomson. *Modern Particle Physics*. Cambridge University Press, Cambridge, first edition (2013). ISBN 9781107034266.
- [7] M. E. Peskin and D. V. Schroeder. *An Introduction to Quantum Field Theory* (1995).
- [8] C. N. Yang and R. L. Mills. *Conservation of isotopic spin and isotopic gauge invariance*. Phys. Rev. 96(1) (Oct 1954) pp. 191–195. ISSN 0031899X. <http://dx.doi.org/10.1103/PhysRev.96.191>.
- [9] Particle Data Group. *Review of Particle Physics – Quantum Chromodynamics*. Prog. Theor. Exp. Phys. 2022 (2022) p. 083C01. <http://dx.doi.org/10.1093/ptep/ptac097>.
- [10] The UA2 Collaboration. *Observation of single isolated electrons of high transverse momentum in events with missing transverse energy at the CERN pp collider*. Phys. Lett. B 122(5-6) (Mar 1983) pp. 476–485. ISSN 0370-2693. [http://dx.doi.org/10.1016/0370-2693\(83\)91605-2](http://dx.doi.org/10.1016/0370-2693(83)91605-2).
- [11] The UA1 Collaboration. *Experimental observation of lepton pairs of invariant mass around 95 GeV/c² at the CERN SPS collider*. Phys. Lett. B 126(5) (Jul 1983) pp. 398–410. ISSN 0370-2693. [http://dx.doi.org/10.1016/0370-2693\(83\)90188-0](http://dx.doi.org/10.1016/0370-2693(83)90188-0).

-
- [12] F. Englert and R. Brout. *Broken symmetry and the mass of gauge vector mesons*. Phys. Rev. Lett. 13(9) (Aug 1964) pp. 321–323. ISSN 00319007. <http://dx.doi.org/10.1103/PhysRevLett.13.321>.
- [13] P. W. Higgs. *Broken symmetries and the masses of gauge bosons*. Phys. Rev. Lett. 13(16) (Oct 1964) pp. 508–509. ISSN 00319007. <http://dx.doi.org/10.1103/PhysRevLett.13.508>.
- [14] *File:Standard Model of Elementary Particles.svg*. Wikimedia Commons (Sep 2019). https://commons.wikimedia.org/wiki/File:Standard_Model_of_Elementary_Particles.svg.
- [15] M. Thomson. *Modern Particle Physics – Errata* (2018). [https://www.hep.phy.cam.ac.uk/\\$\sim\\$thomson/MPP/ModernParticlePhysics_Errata.pdf](https://www.hep.phy.cam.ac.uk/\simthomson/MPP/ModernParticlePhysics_Errata.pdf).
- [16] R. P. Feynman. *Space-time approach to quantum electrodynamics*. Phys. Rev. 76(6) (Sep 1949) pp. 769–789. ISSN 0031899X. <http://dx.doi.org/10.1103/PHYSREV.76.769/FIGURE/1/THUMB>.
- [17] The ATLAS Collaboration. *Standard Model Summary Plots October 2023* (Nov 2023). <https://cds.cern.ch/record/2882448>.
- [18] R. D. Ball et al. *Parton distributions from high-precision collider data*. Eur. Phys. J. C 2017 7710 77(10) (Oct 2017) pp. 1–75. ISSN 1434-6052. <http://dx.doi.org/10.1140/EPJC/S10052-017-5199-5>.
- [19] 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*. J. High Energy Phys. 2014 20147 2014(7) (Jul 2014) pp. 1–157. ISSN 1029-8479. [http://dx.doi.org/10.1007/JHEP07\(2014\)079](http://dx.doi.org/10.1007/JHEP07(2014)079). 1405.0301.
- [20] E. Bothmann et al. *Event generation with Sherpa 2.2*. SciPost Phys. 7(3) (Sep 2019) p. 034. ISSN 25424653. <http://dx.doi.org/10.21468/SCIPOSTPHYS.7.3.034/PDF>. 1905.09127.
- [21] T. Sjöstrand et al. *An introduction to PYTHIA 8.2*. Comput. Phys. Commun. 191(1) (Jun 2015) pp. 159–177. ISSN 0010-4655. <http://dx.doi.org/10.1016/J.CPC.2015.01.024>. 1410.3012.
- [22] S. Frixione, P. Nason, and C. Oleari. *Matching NLO QCD computations with parton shower simulations: The POWHEG method*. J. High Energy Phys. 2007(11) (Nov 2007) p. 070. ISSN 11266708. <http://dx.doi.org/10.1088/1126-6708/2007/11/070>. 0709.2092.
- [23] M. Bähr et al. *Herwig++ physics and manual*. Eur. Phys. J. C 2008 584 58(4) (Nov 2008) pp. 639–707. ISSN 1434-6052. <http://dx.doi.org/10.1140/EPJC/S10052-008-0798-9>. 0803.0883.

- [24] B. Andersson et al. *Parton fragmentation and string dynamics*. Phys. Rep. 97(2-3) (Jul 1983) pp. 31–145. ISSN 0370-1573. [http://dx.doi.org/10.1016/0370-1573\(83\)90080-7](http://dx.doi.org/10.1016/0370-1573(83)90080-7).
- [25] S. Agostinelli et al. *Geant4—a simulation toolkit*. Nucl. Instruments Methods Phys. Res. Sect. A Accel. Spectrometers, Detect. Assoc. Equip. 506(3) (Jul 2003) pp. 250–303. ISSN 0168-9002. [http://dx.doi.org/10.1016/S0168-9002\(03\)01368-8](http://dx.doi.org/10.1016/S0168-9002(03)01368-8).
- [26] The ATLAS Collaboration. *The ATLAS Simulation Infrastructure*. Eur. Phys. J. C 70(3) (Sep 2010) pp. 823–874. ISSN 14346052. <http://dx.doi.org/10.1140/epjc/s10052-010-1429-9>. 1005.4568.
- [27] O. S. Brüning et al. *LHC Design Report Vol. I: The LHC Main Ring*. CERN Yellow Reports: Monographs. CERN, Geneva (2004). <http://dx.doi.org/10.5170/CERN-2004-003-V-1>.
- [28] O. S. Brüning et al. *LHC Design Report Vol. II: The LHC infrastructure and general services*. CERN Yellow Reports: Monographs. CERN, Geneva (2004). <http://dx.doi.org/10.5170/CERN-2004-003-V-2>.
- [29] M. Benedikt et al. *LHC Design Report Vol. III: The LHC injector chain*. CERN Yellow Reports: Monographs. CERN, Geneva (2004). <http://dx.doi.org/10.5170/CERN-2004-003-V-3>.
- [30] R. Assmann, M. Lamont, and S. Myers. *A Brief History of the LEP Collider*. Nucl. Phys. B, Proc. Suppl. 109(2-3) (Apr 2002) pp. 17–31. ISSN 09205632. [http://dx.doi.org/10.1016/S0920-5632\(02\)90005-8](http://dx.doi.org/10.1016/S0920-5632(02)90005-8).
- [31] The ATLAS Collaboration. *Observation of a new particle in the search for the Standard Model Higgs boson with the ATLAS detector at the LHC*. Phys. Lett. B 716(1) (Sep 2012) pp. 1–29. ISSN 0370-2693. <http://dx.doi.org/10.1016/J.PHYSLETB.2012.08.020>. 1207.7214.
- [32] The CMS Collaboration. *Observation of a new boson at a mass of 125 GeV with the CMS experiment at the LHC*. Phys. Lett. B 716(1) (Sep 2012) pp. 30–61. ISSN 0370-2693. <http://dx.doi.org/10.1016/J.PHYSLETB.2012.08.021>. 1207.7235.
- [33] The CMS Collaboration. *Observation of a new boson with mass near 125 GeV in pp collisions at $\sqrt{s} = 7$ and 8 TeV*. J. High Energy Phys. 2013 20136 2013(6) (Jun 2013) pp. 1–127. ISSN 1029-8479. [http://dx.doi.org/10.1007/JHEP06\(2013\)081](http://dx.doi.org/10.1007/JHEP06(2013)081). 1303.4571.
- [34] O. Aberle et al. *High-Luminosity Large Hadron Collider (HL-LHC): Technical design report*. CERN Yellow Reports: Monographs. CERN, Geneva (2020). <http://dx.doi.org/10.23731/CYRM-2020-0010>.
- [35] E. Mobs. *The CERN accelerator complex - August 2018*. CERN (Aug 2018). <https://cds.cern.ch/record/2636343>.

-
- [36] R. Ostojic. *The LHC insertion magnets*. In *IEEE Trans. Appl. Supercond.* CERN (Mar 2002), volume 12. ISSN 10518223, pp. 196–201. <http://dx.doi.org/10.1109/TASC.2002.1018382>.
- [37] L. Evans and P. Bryant. *LHC Machine*. *J. Instrum.* 3(08) (Aug 2008) p. S08,001. ISSN 1748-0221. <http://dx.doi.org/10.1088/1748-0221/3/08/S08001>.
- [38] R. Steerenberg et al. *Operation and performance of the CERN Large Hadron Collider during proton Run 2*. 10th Int. Part. Accel. Conf. (10) (2019) pp. 504–507. <http://dx.doi.org/10.18429/JACoW-IPAC2019-MOPMP031>.
- [39] The ATLAS Collaboration. *The ATLAS Experiment at the CERN Large Hadron Collider*. *J. Instrum.* 3(08) (Aug 2008) p. S08,003. ISSN 1748-0221. <http://dx.doi.org/10.1088/1748-0221/3/08/S08003>.
- [40] The ATLAS Collaboration. *Performance of the ATLAS Track Reconstruction Algorithms in Dense Environments in LHC Run 2*. *Eur. Phys. J. C* 77(10) (Apr 2017) p. 673. <http://dx.doi.org/10.1140/epjc/s10052-017-5225-7.1704.07983>.
- [41] M. Capeans et al. *ATLAS Insertable B-Layer Technical Design Report* (Sep 2010). <https://cds.cern.ch/record/1291633>.
- [42] J. Pequeno. *Computer Generated image of the ATLAS calorimeter* (Mar 2008). <https://cds.cern.ch/record/1095927>.
- [43] F. Cavallari. *Performance of calorimeters at the LHC*. *J. Phys. Conf. Ser.* 293(1) (2011). ISSN 17426596. <http://dx.doi.org/10.1088/1742-6596/293/1/012001>.
- [44] J. Pequeno. *Computer generated image of the ATLAS Muons subsystem* (Mar 2008). <https://cds.cern.ch/record/1095929>.
- [45] The ATLAS Collaboration. *ATLAS muon spectrometer : Technical Design Report* CERN-LHCC- (1997). <https://cds.cern.ch/record/331068>.
- [46] The ATLAS Collaboration. *Operation of the ATLAS trigger system in Run 2*. *J. Instrum.* 15(10) (Oct 2020) p. P10,004. ISSN 1748-0221. <http://dx.doi.org/10.1088/1748-0221/15/10/P10004>.
- [47] R. Achenbach et al. *The ATLAS Level-1 Calorimeter Trigger*. *J. Instrum.* 3(03) (Mar 2008). ISSN 1748-0221. <http://dx.doi.org/10.1088/1748-0221/3/03/P03001>.
- [48] G. Avoni et al. *The new LUCID-2 detector for luminosity measurement and monitoring in ATLAS*. *J. Instrum.* 13(07) (Jul 2018). ISSN 1748-0221. <http://dx.doi.org/10.1088/1748-0221/13/07/P07017>.

- [49] The ATLAS Collaboration. *ATLAS data quality operations and performance for 2015-2018 data-taking*. J. Instrum. 15(4) (Apr 2020). ISSN 17480221. <http://dx.doi.org/10.1088/1748-0221/15/04/P04003>. 1911.04632.
- [50] The ATLAS Collaboration. *Luminosity determination in pp collisions at $\sqrt{s} = 13$ TeV using the ATLAS detector at the LHC*. Eur. Phys. J. C 83(10) (Oct 2023) pp. 1–67. ISSN 1434-6052. <http://dx.doi.org/10.1140/EPJC/S10052-023-11747-W>.
- [51] I. Brawn. *L1Calo Overview, Status, Installation & Commissioning* (Oct 2019). <https://indico.cern.ch/event/829769/contributions/3572289>. [ATLAS Internal].
- [52] The ATLAS Collaboration. *Technical Design Report for the Phase-I Upgrade of the ATLAS TDAQ System*. Technical Report ATLAS-TDR-023-2013 (2013). <http://cds.cern.ch/record/1602235/>.
- [53] The ATLAS Collaboration. *Technical Design Report for the Phase-II Upgrade of the ATLAS TDAQ System*. Technical Report ATLAS-TDR-029, CERN, Geneva (2017). <https://cds.cern.ch/record/2285584>.
- [54] OpenJS Foundation and NodeJS Contributors. *Node.js*. <https://nodejs.org/en>.
- [55] OpenJS Foundation. *Express - Node.js web application framework*. <https://expressjs.com/>.
- [56] M. Sweney. *Global shortage in computer chips ‘reaches crisis point’*. The Guard. (Mar 2021). <https://www.theguardian.com/business/2021/mar/21/global-shortage-in-computer-chips-reaches-crisis-point>.
- [57] C. J. Clopper and E. S. Pearson. *The Use of Confidence or Fiducial Limits Illustrated in the Case of the Binomial*. Biometrika 26(4) (1934) p. 404. ISSN 00063444. <http://dx.doi.org/10.2307/2331986>.
- [58] The ATLAS Collaboration. *L1Calo Trigger Public Results* (Nov 2022). <https://twiki.cern.ch/twiki/bin/view/AtlasPublic/L1CaloTriggerPublicResults>.
- [59] T. Sjöstrand, S. Mrenna, and P. Skands. *A brief introduction to PYTHIA 8.1*. Comput. Phys. Commun. 178(11) (Jun 2008) pp. 852–867. ISSN 00104655. <http://dx.doi.org/10.1016/j.cpc.2008.01.036>. 0710.3820.
- [60] The ATLAS Collaboration. *Performance of electron and photon triggers in ATLAS during LHC Run 2*. Eur. Phys. J. C 80(1) (Jan 2020) pp. 1–41. ISSN 1434-6052. <http://dx.doi.org/10.1140/EPJC/S10052-019-7500-2>. 1909.00761.
- [61] The ATLAS Collaboration. *Performance of the ATLAS muon triggers in Run 2*. J. Instrum. 15(09) (Sep 2020). ISSN 1748-0221. <http://dx.doi.org/10.1088/1748-0221/15/09/P09015>.

-
- [62] R. D. Ball et al. *Parton distributions for the LHC run II*. J. High Energy Phys. 2015 20154 2015(4) (Apr 2015) pp. 1–148. ISSN 1029-8479. [http://dx.doi.org/10.1007/JHEP04\(2015\)040](http://dx.doi.org/10.1007/JHEP04(2015)040). 1410.8849.
- [63] H. L. Lai et al. *New parton distributions for collider physics*. Phys. Rev. D - Part. Fields, Gravit. Cosmol. 82(7) (Oct 2010) p. 074,024. ISSN 15507998. <http://dx.doi.org/10.1103/PhysRevD.82.074024>. 1007.2241.
- [64] R. D. Ball et al. *Parton distributions with LHC data*. Nucl. Phys. B 867(2) (Feb 2013) pp. 244–289. ISSN 05503213. <http://dx.doi.org/10.1016/j.nuclphysb.2012.10.003>. 1207.1303.
- [65] J. Bellm et al. *Herwig++ 2.7 Release Note* (Oct 2013). <http://dx.doi.org/10.48550/arXiv.1310.6877>. 1310.6877.
- [66] The ATLAS Collaboration. *Measurement of the cross-sections of the electroweak and total production of a $Z\gamma$ pair in association with two jets in pp collisions at $\sqrt{s} = 13$ TeV with the ATLAS detector* (Jul 2023). <http://cds.cern.ch/record/2779171><http://arxiv.org/abs/2305.19142>. 2305.19142.
- [67] P. Nason. *A new method for combining NLO QCD with shower Monte Carlo algorithms*. J. High Energy Phys. 8(11) (Dec 2004) pp. 1097–1124. ISSN 10298479. <http://dx.doi.org/10.1088/1126-6708/2004/11/040>. 0409146.
- [68] S. Alioli et al. *A general framework for implementing NLO calculations in shower Monte Carlo programs: The POWHEG BOX*. J. High Energy Phys. 2010(6) (Jun 2010) pp. 1–58. ISSN 10298479. [http://dx.doi.org/10.1007/JHEP06\(2010\)043](http://dx.doi.org/10.1007/JHEP06(2010)043).
- [69] The ATLAS Collaboration. *ATLAS Pythia 8 tunes to 7 TeV data*. Technical Report ATL-PHYS-PUB-2014-021, CERN, Geneva (2014). <https://cds.cern.ch/record/1966419>.
- [70] The ATLAS Collaboration. *Measurement of the Z/γ^* boson transverse momentum distribution in pp collisions at $\sqrt{s} = 7$ TeV with the ATLAS detector*. J. High Energy Phys. 2014(9) (Sep 2014) pp. 1–47. ISSN 10298479. [http://dx.doi.org/10.1007/JHEP09\(2014\)145](http://dx.doi.org/10.1007/JHEP09(2014)145)/METRICS. 1406.3660.
- [71] The ATLAS Collaboration. *Electron and photon energy calibration with the ATLAS detector using 2015–2016 LHC proton-proton collision data*. J. Instrum. 14(03) (Mar 2019) p. P03,017. ISSN 1748-0221. <http://dx.doi.org/10.1088/1748-0221/14/03/P03017>.
- [72] The ATLAS Collaboration. *Electron and photon performance measurements with the ATLAS detector using the 2015–2017 LHC proton-proton collision data*. J. Instrum. 14(12) (Dec 2019) p. P12,006. ISSN 1748-0221. <http://dx.doi.org/10.1088/1748-0221/14/12/P12006>.

- [73] B. Wendland. *Photon identification with the ATLAS detector*. PoS EPS-HEP202 (Oct 2021) p. 332. ISSN 18248039. <http://dx.doi.org/10.22323/1.398.0332>.
- [74] The ATLAS Collaboration. *Photon identification in 2015 ATLAS data* (Aug 2016). <https://cds.cern.ch/record/2203125>.
- [75] The ATLAS Collaboration. *Electron and photon efficiencies in LHC Run 2 with the ATLAS experiment* (Aug 2023). <https://arxiv.org/abs/2308.13362v1>.
- [76] The ATLAS Collaboration. *Muon reconstruction and identification efficiency in ATLAS using the full Run 2 pp collision data set at $\sqrt{s} = 13$ TeV*. Eur. Phys. J. C 81(7) (Jul 2021) pp. 1–44. ISSN 1434-6052. <http://dx.doi.org/10.1140/EPJC/S10052-021-09233-2>.
- [77] M. Cacciari, G. P. Salam, and G. Soyez. *The anti-kt jet clustering algorithm*. J. High Energy Phys. 2008(04) (Apr 2008) p. 063. ISSN 1126-6708. <http://dx.doi.org/10.1088/1126-6708/2008/04/063>.
- [78] G. P. Salam. *Towards jetography*. Eur. Phys. J. C 67(3) (May 2010) pp. 637–686. ISSN 1434-6052. <http://dx.doi.org/10.1140/EPJC/S10052-010-1314-6>.
- [79] The ATLAS Collaboration. *Topological cell clustering in the ATLAS calorimeters and its performance in LHC Run 1*. Eur. Phys. J. C 77(7) (2017) p. 490. ISSN 1434-6052. <http://dx.doi.org/10.1140/epjc/s10052-017-5004-5>.
- [80] The ATLAS Collaboration. *Jet reconstruction and performance using particle flow with the ATLAS Detector*. Eur. Phys. J. C 77(7) (2017) p. 466. ISSN 1434-6052. <http://dx.doi.org/10.1140/epjc/s10052-017-5031-2>.
- [81] The ATLAS Collaboration. *Jet energy scale measurements and their systematic uncertainties in proton-proton collisions at $\sqrt{s} = 13$ TeV with the ATLAS detector* (2017). <http://dx.doi.org/10.1103/PhysRevD.96.072002>.
- [82] The ATLAS Collaboration. *ATLAS b-jet identification performance and efficiency measurement with $t\bar{t}$ events in pp collisions at $\sqrt{s} = 13$ TeV*. Eur. Phys. J. C 2019 7911 79(11) (Nov 2019) pp. 1–36. ISSN 1434-6052. <http://dx.doi.org/10.1140/EPJC/S10052-019-7450-8>. 1907.05120.
- [83] D. Adams et al. *Recommendations of the Physics Objects and Analysis Harmonisation Study Groups 2014*. Technical Report ATL-COM-PHYS-2014-451 (Feb 2015). <https://cds.cern.ch/record/1700874/files/ATL-COM-PHYS-2014-451.pdf>. [ATLAS Internal].
- [84] A. Höcker et al. *The toolkit for multivariate data analysis, TMVA 4*. In *J. Phys. Conf. Ser.* (Mar 2010), volume 219. ISSN 17426596. <http://dx.doi.org/10.1088/1742-6596/219/3/032057>.

-
- [85] The ATLAS Collaboration. *Measurements of inclusive and differential fiducial cross-sections of $t\bar{t}\gamma$ production in leptonic final states at $\sqrt{s} = 13 \sim \text{TeV}$ in ATLAS*. Eur. Phys. J. C 79(5) (May 2019) pp. 1–41. ISSN 1434-6052. <http://dx.doi.org/10.1140/EPJC/S10052-019-6849-6>.
- [86] J. Butterworth et al. *PDF4LHC recommendations for LHC Run II*. J. Phys. G Nucl. Part. Phys. 43(2) (Jan 2016) p. 023,001. ISSN 0954-3899. <http://dx.doi.org/10.1088/0954-3899/43/2/023001>.
- [87] The ATLAS Collaboration. *Jet energy scale and resolution measured in proton–proton collisions at $\sqrt{s} = 13 \text{ TeV}$ with the ATLAS detector*. Eur. Phys. J. C 2021 818 81(8) (Aug 2021) pp. 1–49. ISSN 1434-6052. <http://dx.doi.org/10.1140/EPJC/S10052-021-09402-3>.
- [88] W. Buttinger. *Using Event Weights to account for differences in Instantaneous Luminosity and Trigger Prescale in Monte Carlo and Data*. Technical Report ATL-COM-SOFT-2015-119 (2017). [ATLAS Internal].
- [89] The ATLAS Collaboration. *Measurement of the Inelastic Proton-Proton Cross Section at $\sqrt{s} = 13 \text{ TeV}$ with the ATLAS Detector at the LHC*. Phys. Rev. Lett. 117(18) (Oct 2016) p. 182,002. ISSN 10797114. <http://dx.doi.org/10.1103/PhysRevLett.117.182002>.
- [90] The ATLAS Collaboration. *Identification and rejection of pile-up jets at high pseudorapidity with the ATLAS detector*. Eur. Phys. J. C 77(9) (Sep 2017) pp. 1–32. ISSN 1434-6052. <http://dx.doi.org/10.1140/EPJC/S10052-017-5081-5>. 1705.02211.
- [91] The ATLAS Collaboration. *Light-quark and Gluon Jets in ATLAS: Calorimeter Response, Jet Energy Scale Systematics, and Sample Characterization* (Mar 2011).
- [92] K. Cranmer et al. *HistFactory: A tool for creating statistical models for use with RooFit and RooStats* (Jan 2012). <https://cds.cern.ch/record/1456844>.
- [93] W. C. Davidon. *Variable metric method for minimization* (May 1991). <http://dx.doi.org/10.1137/0801001>.
- [94] R. Fletcher. *New approach to variable metric algorithms*. Comput. J. 13(3) (Jan 1970) pp. 317–322. ISSN 00104620. <http://dx.doi.org/10.1093/comjnl/13.3.317>.
- [95] M. J. D. Powell. *Variable Metric Methods for Constrained Optimization*. In *Math. Program. State Art*, Springer, Berlin, Heidelberg (1983). ISBN 978-3-642-68874-4, pp. 288–311. http://dx.doi.org/10.1007/978-3-642-68874-4_12.
- [96] F. James and M. Winkler. *Minuit2* (2004). <https://root.cern.ch/root/html/doc/guides/minuit2/Minuit2.html>.

- [97] S. S. Wilks. *The Large-Sample Distribution of the Likelihood Ratio for Testing Composite Hypotheses*. Ann. Math. Stat. 9(1) (Mar 1938) pp. 60–62. ISSN 0003-4851. <http://dx.doi.org/10.1214/aoms/1177732360>.
- [98] S. Das Bakshi et al. *Towards the renormalisation of the Standard Model effective field theory to dimension eight: bosonic interactions II*. Eur. Phys. J. Plus 137(8) (Aug 2022). ISSN 21905444. <http://dx.doi.org/10.1140/EPJP/S13360-022-03194-5>. 2205.03301.
- [99] The ATLAS Collaboration. *Evidence for electroweak production of two jets in association with a $Z\gamma$ pair in pp collisions at $\sqrt{s} = 13$ TeV with the ATLAS detector*. Phys. Lett. B 803 (Apr 2020). ISSN 0370-2693. <http://dx.doi.org/10.1016/J.PHYSLETB.2020.135341>.
- [100] The ATLAS Collaboration. *Measurement of the cross-sections of the electroweak and total production of a $Z\gamma$ pair in association with two jets in pp collisions at $\sqrt{s} = 13$ TeV with the ATLAS detector*. Phys. Lett. B 846 (Nov 2023). ISSN 0370-2693. <http://dx.doi.org/10.1016/J.PHYSLETB.2023.138222>.
- [101] The CMS Collaboration. *Measurement of the electroweak production of $Z\gamma$ and two jets in proton-proton collisions at $\sqrt{s} = 13$ TeV and constraints on anomalous quartic gauge couplings*. Phys. Rev. D 104(7) (Oct 2021). ISSN 24700029. <http://dx.doi.org/10.1103/PHYSREVD.104.072001/FIGURES/10/MEDIUM>.
- [102] The ATLAS Collaboration. *Measurement of electroweak $Z(\nu\bar{\nu})\gamma jj$ production and limits on anomalous quartic gauge couplings in pp collisions at $\sqrt{s} = 13$ TeV with the ATLAS detector*. J. High Energy Phys. 2023(6) (Jun 2023) pp. 1–48. ISSN 10298479. [http://dx.doi.org/10.1007/JHEP06\(2023\)082/METRICS](http://dx.doi.org/10.1007/JHEP06(2023)082/METRICS).
- [103] The ATLAS Collaboration. *Measurement and interpretation of same-sign W boson pair production in association with two jets in pp collisions at $\sqrt{s} = 13$ TeV with the ATLAS detector* (Dec 2023). <https://arxiv.org/abs/2312.00420v1>.
- [104] The ATLAS Collaboration. *Observation of electroweak production of two jets and a Z -boson pair*. Nat. Phys. 19(2) (Feb 2023) pp. 237–253. ISSN 1745-2481. <http://dx.doi.org/10.1038/s41567-022-01757-y>. 2004.10612.
- [105] The ATLAS Collaboration. *Observation of electroweak $W^{\pm}Z$ boson pair production in association with two jets in pp collisions at $\sqrt{s} = 13$ TeV with the ATLAS detector*. Phys. Lett. B 793 (Jun 2019) pp. 469–492. ISSN 0370-2693. <http://dx.doi.org/10.1016/J.PHYSLETB.2019.05.012>.
- [106] The ATLAS Collaboration. *Search for electroweak diboson production in association with a high-mass dijet system in semileptonic final states in pp collisions at $\sqrt{s} = 13$ TeV with the ATLAS detector*. Phys. Rev. D 100(3) (Aug 2019). ISSN 24700029. <http://dx.doi.org/10.1103/PHYSREVD.100.032007/FIGURES/7/MEDIUM>.

-
- [107] B. Efron. *Computers and the Theory of Statistics: Thinking the Unthinkable*. SIAM Rev. 21(4) (Oct 1979) pp. 460–480. ISSN 0036-1445. <http://dx.doi.org/10.1137/1021092>.
 - [108] B. Efron. *Better Bootstrap Confidence Intervals*. J. Am. Stat. Assoc. 82(397) (Mar 1987) p. 171. ISSN 01621459. <http://dx.doi.org/10.2307/2289144>.
 - [109] The ATLAS Collaboration. *Observation of $WZ\gamma$ production in pp collisions at $\sqrt{s} = 13$ TeV with the ATLAS detector* (May 2023). <https://arxiv.org/abs/2305.16994v1>.
 - [110] The CMS Collaboration. *Search for $WW\gamma$ and $WZ\gamma$ production and constraints on anomalous quartic gauge couplings in pp collisions at $\sqrt{s} = 8$ TeV*. Phys. Rev. D 90(3) (Aug 2014). ISSN 15502368. <http://dx.doi.org/10.1103/PHYSREVD.90.032008/FIGURES/5/MEDIUM>.
 - [111] The CMS Collaboration. *Observation of the Production of Three Massive Gauge Bosons at $\sqrt{s} = 13$ TeV*. Phys. Rev. Lett. 125(15) (Oct 2020). ISSN 10797114. <http://dx.doi.org/10.1103/PHYSREVLETT.125.151802/FIGURES/2/MEDIUM>.
 - [112] The ATLAS Collaboration. *Observation of WWW Production in pp Collisions at $\sqrt{s} = 13$ TeV with the ATLAS Detector*. Phys. Rev. Lett. 129(6) (Aug 2022). ISSN 10797114. <http://dx.doi.org/10.1103/PHYSREVLETT.129.061803/FIGURES/2/MEDIUM>.
 - [113] The ATLAS Collaboration. *Measurement of $Z\gamma\gamma$ production in pp collisions at $\sqrt{s} = 13$ TeV with the ATLAS detector*. Eur. Phys. J. C 83(6) (Jun 2023) pp. 1–30. ISSN 1434-6052. <http://dx.doi.org/10.1140/EPJC/S10052-023-11579-8>.
 - [114] The CMS Collaboration. *Measurements of the $pp \rightarrow W^\pm\gamma\gamma$ and $pp \rightarrow Z\gamma\gamma$ cross sections at $\sqrt{s} = 13$ TeV and limits on anomalous quartic gauge couplings*. J. High Energy Phys. 2021(10) (Oct 2021) pp. 1–39. ISSN 10298479. [http://dx.doi.org/10.1007/JHEP10\(2021\)174/METRICS](http://dx.doi.org/10.1007/JHEP10(2021)174/METRICS).
 - [115] The ATLAS Collaboration. *Observation of $W\gamma\gamma$ triboson production in proton-proton collisions at $\sqrt{s} = 13$ TeV with the ATLAS detector* (Aug 2023). <https://arxiv.org/abs/2308.03041v1>.
 - [116] J. C. Collins and D. E. Soper. *Angular distribution of dileptons in high-energy hadron collisions*. Phys. Rev. D 16(7) (Oct 1977) pp. 2219–2225. ISSN 05562821. <http://dx.doi.org/10.1103/PhysRevD.16.2219>.
 - [117] The ATLAS Collaboration. *ATLAS $TRExFitter$* . <https://trexfitter-docs.web.cern.ch/trexfitter-docs/>. [ATLAS Internal].
 - [118] A. Kolmogoroff. *Sulla determinazione empirica di una legge di distribuzione*. G. dell’Istituto Ital. degli Attuari 4 (1933) pp. 83–91.

- [119] N. Smirnov. *Sur les écarts de la courbe de distribution empirique.* Recl. Mathématique [Mat. Sb.] 6(48) (1939) pp. 3–26.
- [120] F. J. Massey. *The Kolmogorov-Smirnov Test for Goodness of Fit.* J. Am. Stat. Assoc. 46(253) (Mar 1951) p. 68. ISSN 01621459. <http://dx.doi.org/10.2307/2280095>.

