

Direct measurement of the Higgs boson mass, natural width, and cross section times branching ratio to four leptons using a per-event lineshape in the Higgs to ZZ to four lepton decay channel with the ATLAS detector

by

Graham Cree

A thesis submitted to the
Faculty of Graduate and Postdoctoral Affairs
in partial fulfillment of the requirements for the degree of

Doctor of Philosophy, Particle Physics

Department of Physics
Carleton University
Ottawa-Carleton Institute for Physics
Ottawa, Ontario
April, 2017

© Graham Cree, 2017



Abstract

The discovery of the Higgs boson by the ATLAS and CMS Collaborations in 2012 remains the crowning achievement of the Large Hadron Collider (LHC) physics programme. Five years since its discovery, Run 2 at the LHC is underway and producing more data than ever before, allowing measurements of the Higgs boson beyond the reach of Run 1. Precise measurement of the Higgs boson’s properties help guide particle physicists in understanding the Standard Model, and what lays beyond.

This thesis presents a measurement of the Higgs boson mass, natural width, and cross section times branching ratio in the $H \rightarrow ZZ^{(*)} \rightarrow 4\ell$ decay channel using the full 2015+2016 combined dataset from Run 2 at the LHC, totalling 36.1 fb^{-1} of p - p collisions at centre-of-mass energy $\sqrt{s} = 13 \text{ TeV}$. The analysis is performed using a technique developed by the author, called the *per-event response* method. The technique is designed to produce a more precise, accurate, and model-independent measurement of the Higgs boson properties grounded directly on the performance of the ATLAS detector.

Using this technique, the Higgs boson mass was measured to be $m_H = 124.61^{+0.44}_{-0.44}(\text{stat.})^{+0.06}_{-0.06}(\text{syst.}) \text{ GeV}$. The total Higgs boson production cross section times branching ratio to four electrons or muons in p - p collisions at $\sqrt{s} = 13 \text{ TeV}$ was measured to be $\sigma_H \times \text{BR}(H \rightarrow 4\ell) = 8.55^{+1.23}_{-1.15}(\text{stat.})^{+0.63}_{-0.52}(\text{syst.}) \text{ fb}$. An upper limit was set on the natural width of the Higgs as $\Gamma_H < 3.65(\text{stat.}) \oplus 0.19(\text{syst.}) \text{ GeV}$ at 95% confidence. These measurements are consistent with previous measurements and with the Standard Model. This result is the first ever measurement of Higgs boson mass using ATLAS at $\sqrt{s} = 13 \text{ TeV}$, exceeding the precision of any $H \rightarrow 4\ell$ m_H measurement made to this point.

Also presented are the author’s contributions to the precision assembly of the small-strip Thin Gap Chambers to be installed in ATLAS in 2019, and to the ATLAS reconstruction software, both essential to continued detector operation and the many physics analyses yet to come from ATLAS.

Acknowledgments

As the other students graduated and the office evolved, it rarely occurred to me that my turn would come as well. For having gotten this far, I owe several people my sincere gratitude:

My supervisor, Thomas Koffas, for his guidance, advice, and loyal advocacy on my behalf.

To Christos Anastopoulos, William Leight, Ioannis Nomidis, for an abundance of useful discussions, and for tolerating my pickiness over shared analysis code.

To the other members of the Carleton student office, David Di Valentino, Jim Lacey, Tom McCarthy, Rick Ueno, Stephen Weber, Kate Whalen, for categorically rejecting the miserable grad student trope, and helping me do the same. The last five years have not been without challenges, but for the most part I was just having fun.

To my parents, for a lifetime (mine) of love and support (theirs), without which I would certainly not be here. To my grandmother, who sometimes hears news from the LHC before I do, and to the rest of my family, whose support has been invaluable.

Finally, to Emily, who has been an unwavering ally and my ultimate confidante through this whole affair. I would be neither where I am, nor who I am, without you.

I would also like to acknowledge the support of the Natural Sciences and Engineering Research Council of Canada.

Table of Contents

Abstract	ii
Acknowledgments	iii
Table of Contents	iv
List of Tables	viii
List of Figures	x
List of Abbreviations	xvi
1 Introduction	1
2 The Standard Model	4
2.1 Particle Content	4
2.1.1 Composite particles	6
2.2 Interactions	8
2.2.1 Strong interaction	9
2.2.2 Electroweak interaction	9
2.2.3 Z boson decay	10
2.3 The Higgs boson	12
2.3.1 The Higgs mechanism	12
2.3.2 Higgs production and decay	16
2.3.3 The $H \rightarrow ZZ^{(*)} \rightarrow 4\ell$ channel	23
2.4 Beyond the Standard Model	24
3 ATLAS and the LHC	26
3.1 The LHC	26
3.1.1 From canister to collision	27

3.1.2	Collisions	28
3.2	The ATLAS Experiment	31
3.2.1	The ATLAS coordinate system	32
3.2.2	Inner Detector	35
3.2.3	Calorimeters	38
3.2.4	Muon Spectrometer	39
3.3	New Small Wheel	43
3.3.1	sTGC	43
4	Reconstruction	54
4.1	Data and Simulation	54
4.2	Kinematic formulae	56
4.3	Trigger	60
4.4	Electrons and photons	61
4.4.1	Gaussian Sum Filter	61
4.5	Muons	70
4.5.1	Silicon Associated Forward Muons	72
4.6	Jets	79
4.7	Missing transverse energy	79
4.8	Reconstruction summary	80
5	$H \rightarrow ZZ^{(*)} \rightarrow 4\ell$ Event Selection	81
5.1	Pre-Selection	82
5.2	Event-level selection	82
5.3	Lepton selection	84
5.3.1	Electron selection	84
5.3.2	Muon selection	84
5.4	Jet and photon selection	86
5.5	Overlap removal	86
5.6	4ℓ selection	86
5.7	Additional Criteria	88
5.7.1	ΔR and cross-pairing	88
5.7.2	Isolation	88
5.7.3	d_0 significance and vertexing	89
5.8	FSR recovery	89
5.9	Z Mass Constraint	90
5.10	Invariant mass resolution	92

5.11	Analysis categories	98
6	Background Estimation and Modelling	99
6.1	Standard Model ZZ^*	100
6.2	$VVV, t\bar{t} + Z$	104
6.3	$Z + \text{Jets}, t\bar{t}, WZ$	107
6.3.1	$Z + ee$	107
6.3.2	$Z + \mu\mu$	114
6.3.3	$Z + \text{Jets}, t\bar{t}, WZ$ PDF	114
6.4	Backgrounds summary	116
7	Per-Event Response Signal Model	117
7.1	Signal Normalization	117
7.2	Template Method	119
7.2.1	Signal template smoothing	120
7.2.2	Signal template morphing	120
7.3	Per-Event Response method	127
7.3.1	Per-Event Response PDF	128
7.3.2	Describing Single Lepton Energy Response	130
7.3.3	PDF simplification	136
7.3.4	Gaussian sum parameter transformation	140
7.3.5	Response after Z Mass Constraint	142
7.3.6	Distribution of resolutions	146
7.4	Choosing a signal model	151
8	Systematic Uncertainties	153
8.1	Parameterization of systematic uncertainties	154
8.2	Sources of systematic uncertainty	155
8.2.1	Theory systematics	156
8.2.2	Experimental systematics	157
8.2.3	Analysis Systematics	158
8.3	Systematic affecting event yields	158
8.4	Systematics affecting scale and resolution	165
9	Statistics and Property Measurement	172
9.1	The method of maximum likelihood	172
9.2	Tests using an effective dataset	177

9.2.1	Tests of expected resolution	178
9.2.2	Tests of model closure	183
9.3	Effects of systematics	185
9.4	Tests using pseudo-experiments	189
9.4.1	Fit results in pseudo-experiments	191
9.4.2	Toy-by-toy comparison	198
9.5	Validation using the $Z \rightarrow 4\ell$ resonance	201
9.5.1	Per-event response and the $Z \rightarrow 4\ell$ lineshape	202
9.5.2	Property measurement of the $Z \rightarrow 4\ell$ resonance	204
9.6	Results with data	209
10	Conclusion	220
Appendix A	Backgrounds supplementary plots	222
Appendix B	Extra Discussion of the Signal Model	226
B.1	Stability of Template interpolation	226
B.2	On the definition of x_i	228
B.3	Response model fits	228
B.4	Response model parameters	237
B.5	Model Independence of the PER signal model	242
Appendix C	Statistics and Measurement Supplement	244
C.1	Data distributions of $m_{4\ell}^{\text{fsr}}$	244
C.2	Post-fit PDF	247
C.3	More results from pseudo-experiments	250
C.4	$Z \rightarrow 4\ell$ analysis using $m_{4\ell}^{\text{fsr}}$	261
C.5	Response distribution of data events	266
List of	References	275

List of Tables

1	Standard Model Fermions	5
2	Standard Model Bosons	5
3	SM Higgs production fractions	20
4	Branching Ratio $\text{BR}(H \rightarrow ZZ^*)$	21
5	Particle Boosts and Lifetimes	59
6	Trigger Rates	60
7	Muon Reconstruction Resolution by Type	77
8	Trigger List	83
9	Invariant mass resolution	93
10	$Z + ee$ Background Estimate	113
11	Backgrounds summary table	116
12	Expected signal yield	118
13	Unconstrained PER transformation parameters	141
14	Constrained PER transformation parameters	145
15	Dominant systematics for signal yield	164
16	Dominant systematics for signal scale	166
17	Dominant systematics for signal resolution	169
18	Effective dataset results	182
19	MC statistical fluctuation of m_H	183
20	Residual bias of m_H	184
21	Systematics summary on $m_H, \sigma_{H \rightarrow 4\ell}, \Gamma_H$	189
22	Summary table of toys fitting m_H	192
23	Summary table of toys fitting $\sigma_{H \rightarrow 4\ell}$	196
24	Summary table of toys fitting Γ_H	197
25	Improvement from PER	201
26	$Z \rightarrow 4\ell$ yields for $m_{4\ell}^{\text{cons}}$	207
27	$Z \rightarrow 4\ell$ constrained fit results	209
28	Event yields in data	210

29	Fit results in data	218
30	Lepton response model example fits	229
31	Lepton response model parameters table	237
32	$Z \rightarrow 4\ell$ yields for $m_{4\ell}^{\text{fsr}}$	261
33	$Z \rightarrow 4\ell$ unconstrained fit results	265
34	Per-event PDFs figures summary	266

List of Figures

1	Proton Substructure	7
2	Strong interaction diagrams	9
3	Electroweak interaction diagrams	10
4	Z boson lineshape	11
5	Higgs Potential	13
6	Higgs interaction diagrams	17
7	ggF Feynman diagram	18
8	VBF Feynman diagram	18
9	VH Feynman diagram	19
10	ttH Feynman diagram	19
11	Higgs production cross section versus m_H	20
12	Higgs branching ratio versus m_H	22
13	Higgs lineshape	23
14	$H \rightarrow ZZ^{(*)} \rightarrow 4\ell$ Feynman diagram	23
15	LHC Accelerator Chain	29
16	LHC Collision Event	31
17	ATLAS Detector	32
18	ATLAS Coordinate System	33
19	Pseudorapidity Conversion	34
20	ATLAS Inner Detector	36
21	ATLAS Inner Detector in r - $ z $	36
22	ATLAS Calorimeter System	38
23	ATLAS Muon Spectrometer	41
24	ATLAS Muon Spectrometer in r - $ z $	42
25	NSW Triggering	44
26	sTGC and NSW assembly	45
27	sTGC Alignment Mechanism	47
28	sTGC Alignment Microscope	48

29	sTGC Alignment Algorithm	48
30	Alignment Measurement Calibration	49
31	Alignment Measurement Orientation	50
32	Alignment Microscope Orientation	50
33	Alignment Measurement Example	51
34	sTGC Alignment Measurement Locations	52
35	Gaussian sum approximation of Bethe-Heitler	63
36	Gaussian Sum Filter propagation	64
37	(q/p) pull for different N_{red}	66
38	Tuning of GSF parameter N_{red}	67
39	Auto-vectorization	68
40	GSF CPU Improvements	69
41	Pixel and SCT hits by muon type	73
42	Muon Types	74
43	Muon Kinematics by Type	75
44	Muon Resolution by Type	76
45	Muon Resolution vs Hits	78
46	Particle Reconstruction in ATLAS	80
47	Impact Parameters d_0, z_0	85
48	Di-lepton invariant mass cuts	88
49	m_{12} distributions	92
50	$m_{4\ell}$ resolution	94
51	$m_{4\ell}$ resolution with FSR	95
52	$m_{4\ell}$ resolution with FSR + ZMC	96
53	Inclusive $m_{4\ell}$ resolution	97
54	$q\bar{q} \rightarrow ZZ^{(*)}$ Feynman diagrams	100
55	$gg \rightarrow ZZ^{(*)}$ Feynman diagram	100
56	$q\bar{q} \rightarrow ZZ^{(*)}$ PDF of $m_{4\ell}^{\text{cons}}$	102
57	$gg \rightarrow ZZ^{(*)}$ PDF of $m_{4\ell}^{\text{cons}}$	103
58	Resonant $q\bar{q} \rightarrow ZZ^{(*)}$ Feynman diagram	104
59	VVV Feynman diagrams	105
60	$t\bar{t} + Z$ Feynman diagrams	105
61	VVV, $t\bar{t} + Z$ PDF of $m_{4\ell}^{\text{cons}}$	106
62	n_{InnerPix} templates from the $Z + X$ CR	109
63	Data/MC comparison of $Z + X$ CR	110
64	X electron efficiency vs. p_T	112

65	Fit on data of the $3\ell + X$ CR	113
66	$Z + \mu\mu$ CRs	115
67	$Z + \text{Jets}, t\bar{t}, WZ$ PDF of $m_{4\ell}^{\text{cons}}$	116
68	Signal Yield Parameterization	119
69	Template smoothing 4μ	121
70	Template smoothing $4e$	122
71	Template smoothing $2\mu 2e$	123
72	Template smoothing $2e 2\mu$	124
73	Template morphing	126
74	Unconstrained muon response fits (1/2)	133
75	Unconstrained electron response fits (1/2)	134
76	χ^2/n distribution of unconstrained lepton response model	135
77	Lepton response model σ_1	136
78	Comparison with CP Resolutions	137
79	Gaussian sum collapsing	139
80	Unconstrained transformation parameters ggF 125 GeV	143
81	Unconstrained transformation parameters ggF 124 GeV	144
82	Unconstrained transformation parameters ggF 126 GeV	144
83	Unconstrained transformation parameters VBF 125 GeV	144
84	Constrained transformation parameters ggF 125 GeV	147
85	Constrained transformation parameters ggF 124 GeV	148
86	Constrained transformation parameters ggF 126 GeV	148
87	Constrained transformation parameters VBF 125 GeV	148
88	Response PDF examples	149
89	Distribution of σ_{eff}	150
90	FlexibleInterpVar used for systematics	155
91	Theory systematics	160
92	Normalization systematics for $H \rightarrow 4\ell$ events	161
93	Normalization systematics for $q\bar{q} \rightarrow ZZ^{(*)}$ events	162
94	Normalization systematics for $Z + \text{Jets}, t\bar{t}, WZ$ events	163
95	Signal systematic morphing	163
96	Scale systematics for $H \rightarrow 4\ell$ events	167
97	Resolution systematics for $H \rightarrow 4\ell$ events	168
98	Shape systematics for $q\bar{q} \rightarrow ZZ^{(*)}$	171
99	Scans of m_H in effective dataset	179
100	Scans of $\sigma_{H \rightarrow 4\ell}$ in effective dataset	180

101	Scans of Γ_H in effective dataset	181
102	Systematic pulls on m_H	186
103	Systematic pulls on $\sigma_{H \rightarrow 4\ell}$	187
104	Systematic pulls on Γ_H	188
105	Toy fit results m_H	193
106	Toy fit results $\sigma_{H \rightarrow 4\ell}$	194
107	Toy fit results Γ_H	195
108	Toys 2D correlation	199
109	Toys expected difference	200
110	Toys one-to-one accuracy	200
111	$Z \rightarrow 4\ell$ efficiency function	203
112	$Z \rightarrow 4\ell$ PER model on simulated data	205
113	$Z \rightarrow 4\ell$ likelihood scans on effective dataset using $m_{4\ell}^{\text{cons}}$	206
114	$Z \rightarrow 4\ell$ likelihood scans on data using $m_{4\ell}^{\text{cons}}$	207
115	$Z \rightarrow 4\ell$ constrained PER model post-fit	208
116	Data distribution of $m_{4\ell}^{\text{cons}}$	212
117	$m_{4\ell}^{\text{cons}}$ vs σ_{eff} in data	213
118	Scans of m_H in data	215
119	Scans of $\sigma_{H \rightarrow 4\ell}$ in data	216
120	Scans of Γ_H in data	217
121	$q\bar{q} \rightarrow ZZ^{(*)}$ PDF of $m_{4\ell}^{\text{fsr}}$	223
122	$gg \rightarrow ZZ^{(*)}$ PDF of $m_{4\ell}^{\text{cons}}$	224
123	$VVV, t\bar{t} + Z$ PDF of $m_{4\ell}^{\text{fsr}}$	225
124	Template interpolation check	227
125	Unconstrained muon response fits (2/2)	230
126	Unconstrained electron response fits (2/2)	231
127	Constrained muon response fits (1/2)	232
128	Constrained muon response fits (2/2)	233
129	Constrained electron response fits (1/2)	234
130	Constrained electron response fits (2/2)	235
131	χ^2/n distribution of constrained lepton response model	236
132	Lepton response model unconstrained σ_1	237
133	Lepton response model unconstrained σ_2	238
134	Lepton response model unconstrained σ_3	238
135	Lepton response model unconstrained μ_1	238
136	Lepton response model unconstrained μ_2	239

137	Lepton response model unconstrained μ_3	239
138	Lepton response model unconstrained w_1	239
139	Lepton response model unconstrained w_2	239
140	Lepton response model constrained σ_1	240
141	Lepton response model constrained σ_2	240
142	Lepton response model constrained σ_3	240
143	Lepton response model constrained μ_1	240
144	Lepton response model constrained μ_2	241
145	Lepton response model constrained μ_3	241
146	Lepton response model constrained w_1	241
147	Lepton response model constrained w_2	241
148	Kinematic Model Independence of PER	243
149	Data distribution of $m_{4\ell}^{\text{fsr}}$	245
150	$m_{4\ell}^{\text{fsr}}$ vs σ_{eff} in data	246
151	Post-fit PDF of $m_{4\ell}^{\text{cons}}$	248
152	Post-fit PDF of $m_{4\ell}^{\text{fsr}}$	249
153	Toy fit results m_H for constrained PER model	250
154	Toy fit results m_H for Template model	251
155	Toy fit results m_H for unconstrained PER model	251
156	Toy fit results $\sigma_{H \rightarrow 4\ell}$ for constrained PER model	252
157	Toy fit results $\sigma_{H \rightarrow 4\ell}$ for Template model	252
158	Toy fit results $\sigma_{H \rightarrow 4\ell}$ for unconstrained PER model	253
159	Toy fit results Γ_H for constrained PER model	253
160	Toy fit results Γ_H for unconstrained PER model	254
161	Toys 2D Correlation of $\hat{m}_H$ by category	255
162	Toys 2D Correlation of $\hat{\sigma}_{H \rightarrow 4\ell}$ by category	256
163	Toys Expected Difference of $\hat{m}_H$ by category	257
164	Toys Expected Difference of $\hat{\sigma}_{H \rightarrow 4\ell}$ by category	258
165	Toys One-to-one accuracy fitting m_H by category	259
166	Toys One-to-one accuracy fitting $\sigma_{H \rightarrow 4\ell}$ by category	260
167	$Z \rightarrow 4\ell$ likelihood scans on effective dataset using $m_{4\ell}^{\text{fsr}}$	262
168	$Z \rightarrow 4\ell$ likelihood scans on data using $m_{4\ell}^{\text{fsr}}$	263
169	$Z \rightarrow 4\ell$ unconstrained PER model post-fit	264
170	Per-event signal PDF of 4μ events (1/2)	267
171	Per-event signal PDF of 4μ events (2/2)	268
172	Per-event signal PDF of $4e$ events (1/2)	269

173	Per-event signal PDF of $4e$ events (2/2)	270
174	Per-event signal PDF of $2\mu 2e$ events (1/2)	271
175	Per-event signal PDF of $2\mu 2e$ events (2/2)	272
176	Per-event signal PDF of $2e 2\mu$ events (1/2)	273
177	Per-event signal PDF of $2e 2\mu$ events (2/2)	274

List of Abbreviations

Abbreviation	Meaning	First introduced
ALICE	A Large Ion Collider Experiment	3.1
ATLAS	A Toroidal LHC ApparatuS	3.2
BH	Bethe-Heitler	4.4.1
BR	branching ratio	2.3.2.2
BW	Relativistic Breit-Wigner	2.2.3
CB	Crystal Ball	7.3
CT	Calo-tagged (muon)	4.5
CB	Combined (muon)	4.5
CDF	cumulative distribution function	4.4.1
CL	confidence level	9.1
CP	Combined Performance	7.3.2
CR	Control Region	6.3.1
CERN	European Organization for Nuclear Research	3.1
CMS	Compact Muon Solenoid	3.1
c.o.m.	centre of momentum	3.2.1.1
CSC	cathode strip chamber	3.2.4
DM	Dark Matter	2.4
DQ	Data Quality	5.1
ECAL	Electromagnetic Calorimeter	3.2.3.1

EW	Electroweak	2.2.2
EWSB	Electroweak Symmetry Breaking	2.3.1
FCal	Forward Calorimeter	3.2.3.2
FIV	FlexibleInterpVar	8
FSR	Final State Radiation	5.8
ggF	gluon gluon fusion	2.3.2.1
GRL	Good Runs List	5.1
HCAL	Hadronic Calorimeter	3.2.3.2
HEC	Hadronic End-Cap Calorimeter	3.2.3.2
HF	Heavy Flavour	6.3.2
HLT	High-Level Trigger	4.3
HS	Hard Scatter	3.1.2
IBL	Insertable B-Layer	3.2.2.1
ID	Inner Detector	3.2.2
L1	Level 1 Trigger	4.3
LAr	Liquid Argon	3.2.3.1
LHC	Large Hadron Collider	3.1
LHCb	LHC-beauty	3.1
LO	leading order	2.2
MC	Monte-Carlo	4.1
MDT	monitored drift tube	3.2.4
ME	Extrapolated (muon)	4.5
MM	MicroMegas	3.3
NLL	negative log likelihood	9.1
NLO	next-to-leading order	2.2
NNLO	next-to-next-to-leading order	2.2

N3LO	(next-to) ³ leading order	2.2
MPI	multiple-parton-interaction	3.1.2
NP	Nuisance Parameter	8.2
NSW	New Small Wheel	3.3
<i>p-p</i>	proton-proton	3.1
<i>p-$\bar{p}$</i>	proton-antiproton	3.1.1
<i>Pb-Pb</i>	lead-lead	3.1
PDF	probability density function	6
PER	Per-Event Response	7.3
POI	parameter of interest	9.1
PS	Proton Synchrotron	3.1.1
PSB	Proton Synchrotron Booster	3.1.1
PU	Pileup	3.1.2
QCD	Quantum Chromodynamics	2.2.1
RF	radio-frequency	3.1.1
RPC	resistive plate chamber	3.2.4
SiAF	silicon-associated forward (muon)	4.5
SCT	Semiconductor Tracker	3.2.2.3
SFOS	same flavour opposite sign	5.6
SM	Standard Model	2
SPS	Super Proton Synchrotron	3.1.1
ST	Segment-tagged (muon)	4.5
sTGC	small-strip thin gap chamber	3.3
TGC	thin gap chamber	3.2.4
TRT	Transition Radiation Tracker	3.2.2.4
ttH	top fusion	2.3.2.1

UE	Underlying Event	3.1.2
UL	upper limit	9.1
VBF	vector boson fusion	2.3.2.1
VH	associated production	2.3.2.1
ZMC	Z Mass Constraint	5.9

Chapter 1

Introduction

The Higgs Boson, predicted in 1964 [1], proved elusive to searches for nearly half a century. In the wake of its eventual discovery in 2012 [2][3], it was often heralded as the last missing piece of the Standard Model. Since the discovery, focus has shifted to the measurement of its properties. Some properties are free parameters of nature for which our theory gives no prediction. Others can be tested against theory to help determine whether the observed particle is in fact the Higgs boson predicted by the Standard Model, or some manifestation of an underlying theory.

The main objective of this thesis is to present measurements of the mass, natural width, and cross section times branching ratio to four electrons or muons of the Higgs boson using a technique developed by the author. The mass, m_H , is a free parameter of nature, while the other two measurements, $\sigma_H \times \text{BR}(H \rightarrow 4\ell)$ and Γ_H , are predicted by the Standard Model, once the mass is known. The measurements are made using the Higgs' decay to two Z bosons which subsequently decay to pairs of electrons or muons, called the $H \rightarrow ZZ^{(*)} \rightarrow 4\ell$ channel. Other hardware and software contributions to the ATLAS Experiment from the author are presented as well.

Chapter 2 presents the Standard Model of particle physics - our modern description of the nature of particles and their interactions. A large focus lies on the Higgs boson and its role in the Standard Model.

Chapter 3 presents the impressive hardware used to produce and detect Higgs bosons - the Large Hadron Collider and the ATLAS Detector, respectively. Also presented is an upgrade project for the future of the ATLAS Detector, and an approach used to assure precision alignment of its components developed by the author (in Section 3.3.1.1).

Chapter 4 moves on to the software level, showing how various particles are reconstructed in the ATLAS Detector. Emphasis is placed on the reconstruction of the

electrons and muons used this analysis, in particular the author’s own contributions are highlighted inline in Section 4.4.1 and Section 4.5.1.

Chapter 5 begins the presentation of the $H \rightarrow ZZ^{(*)} \rightarrow 4\ell$ analysis with the event selection, showing how candidate events are selected out of the enormous numbers of collisions in the detector. The described selection was implemented by the author, but developed and executed by other members of the HZZ working group within ATLAS.

Chapter 6 presents the various techniques used to estimate the different sources of background surviving the event selection. Particular emphasis is placed on the $Z + ee$ background estimation (Section 6.3.1) developed in part by the author.

Chapter 7 introduces the method used to measure the Higgs boson mass, natural width, and cross section times branching ratio to four electrons or muons, the so-called *per-event response* method. The chapter focuses on the signal model: how uncertainties at particle-level reconstruction are propagated upwards to the four-lepton candidate level, and combined with the theoretically motivated lineshape to produce a final expected signal distribution for each event observed. Also presented is an alternate method used as a cross-check, called the *template* method. The full contents of this chapter are the work of the author.

Chapter 8 presents the various sources of systematic uncertainty considered in the analysis. The implementation and impact of each uncertainty on the measurements presented in this thesis is the sole work of the author, although the sources of uncertainty combine the work of many other ATLAS collaborators.

Chapter 9 completes the presentation of the per-event response measurement method, incorporating expected yields and descriptions of all sources of signal and background. The chapter includes a description of the statistical methods used, and several validation steps of the measurement procedure. The chapter closes by presenting the final measurements of the Higgs boson mass, natural width, and cross section times branching ratio to four electrons or muons. All results in this chapter are the work of the author, though comparisons are made with results by others, and will be noted as such.

Chapter 10 concludes this thesis with a summary of the results and outlook.

Several appendices supplement the main body of work. Appendix A supplements Chapter 6 by describing the distribution of background events used for a cross-check. Appendix B supplements Chapter 7 with supporting plots and descriptions of the per-event response and template methods. Appendix C supports Chapter 9 with additional plots and checks.

Various publications have emerged in relation to the work of this thesis. Elements of the author's work on the muon reconstruction software entered into [4], though the primary project described in Section 4.5.1 did not (it was completed after this publication). The $H \rightarrow 4\ell$ event selection and background estimation techniques were described as they are used in this thesis in [5]; both evolve over time as new techniques are developed, and so previous iterations with some differences can be found in [6], and [7]. The per-event response technique (in an earlier version) was also described in [7] where it was used to set a direct limit on the Higgs width, and as a cross-check on the Higgs mass measurement. The mass measurement was combined with CMS in [8]. The author also presented searches for a high-mass Higgs boson with ATLAS in [9], which includes the $H \rightarrow ZZ^{(*)} \rightarrow 4\ell$ channel, along with other di-boson decay channels.

Note that the LHC and ATLAS do not run continuously - they operate for multi-year periods of data taking, split by multi-year periods of shut down for upgrades and maintenance. The two such periods to date are called Run 1 (2011-2013), and Run 2 (2015-present). Though the two different runs have differences right from the hardware level to the final analysis, all descriptions and data in this thesis correspond to Run 2, unless specifically noted. The measurements presented in this thesis were performed on the to-date ATLAS Run 2 dataset, 36.1 fb^{-1} of p - p collisions at $\sqrt{s} = 13 \text{ TeV}$.

Chapter 2

The Standard Model

Particle physics is the study of fundamental particles and their interactions. Through this lens we gain an understanding about the nature of the universe, from its very origins to the interactions we make use of every day.

Our description of what is fundamental has evolved over time. Historically indivisible particles have been revealed to be composite - made up of something more fundamental and governed by a new interaction heretofore unseen. Our modern description of the fundamental particles and their interactions is called the Standard Model of Particle Physics (SM).

This chapter will describe the content of the SM. Section 2.1 discusses the particle content, while Section 2.2 describes how these particles interact. The main focus of this thesis is a measurement of the Higgs boson, and so Section 2.3 is dedicated to discuss its role in the SM, how they will be produced and observed at the LHC, and some of its properties. Section 2.4 explores some of the weak points of the SM, and why measuring the properties of the Higgs boson can be useful in resolving them.

The complete mathematical formulation of the SM is beyond the scope of this thesis, and is used only where it is essential for understanding the content.

2.1 Particle Content

The fundamental particles of SM are divided into *fermions*, and *bosons*. Fermions may be thought of as matter particles, and bosons as force-carriers. The fermion content of the SM is summarized in Table 1, and the boson content in Table 2.

Fermions are spin-1/2 particles. They are divided in *leptons* and *quarks* based on their interactions. The leptons come in pairs of a negatively charged lepton (ℓ^-) and

Table 1: Fermion content of the SM, based on [10]. The top mass is updated using a recent measurement [11]. Note that the corresponding anti-particles for each of the listed particles has an equal spin and mass, but opposite charge. The uncertainties on the quark masses are much larger than those of the leptons.

	symbol	common name	spin	charge	mass [MeV]
<i>quarks</i>					
up-type quarks	u	up	1/2	+2/3	3
	c	charm	1/2	+2/3	1200
	t	top	1/2	+2/3	172, 440
down-type quarks	d	down	1/2	-1/3	7
	s	strange	1/2	-1/3	120
	b	bottom	1/2	-1/3	4, 300
<i>leptons</i>					
charged leptons	e^-	electron	1/2	-1	0.510999
	μ^-	muon	1/2	-1	105.659
	τ^-	tau	1/2	-1	1776.99
neutrinos	ν_e	electron neutrino	1/2	0	$< 10^{-6}$
	ν_μ	muon neutrino	1/2	0	$< 10^{-6}$
	ν_τ	tau neutrino	1/2	0	$< 10^{-6}$

Table 2: Boson content of the SM, based on [10]. The Higgs boson properties are the latest ATLAS+CMS combination [8].

symbol	common name	spin	mass [GeV]	interacts with
H	Higgs boson	0	125.09	$H, W^\pm, Z, \ell, q$
g	gluon	1	0	g, q
$W^\pm$	W boson	1	80.42	$H, W^\pm, Z, \gamma, \ell, \nu_\ell, q$
Z	Z boson	1	91.19	$H, W^\pm, \ell, \nu_\ell, q$
γ	photon	1	0	$W^\pm, \ell, q$

its corresponding neutrino (ν_ℓ). There are three such generations, ordered by increasing mass m_ℓ (while neutrinos definitively have mass [12][13], we do not know their mass hierarchy, and their masses are not included in the SM). The three generations of leptons are called the electron, muon, and tau ($\ell = e, \mu, \tau$). Quarks also come in three

generations of “up-type” (charge $+2/3$) and “down-type” (charge $-1/3$) quarks. The three generations of quarks are the up and down ($q = u, d$), the charm and strange (c, s), and the top and bottom (t, b). For each of these 12 particles, there is a corresponding anti-particle. The only one with a special name is the anti-electron, called the positron. The others are simply denoted as $\ell^+, \bar{\nu}_\ell, \bar{q}$ according to their type. In some contexts the particle-anti-particle distinction is either implied or unimportant. For example in the decay $Z \rightarrow \mu^+ \mu^-$, including the charge designation is redundant so it is often omitted.

In the language of the SM, fermions are separated into left(L) and right-handed(R) chirality. The left-handed fermions are described mathematically as doublets, and the right-handed fermions are described as singlets. There are no right-handed neutrinos. For the first generation fermions, this looks like

$$\begin{aligned} E_L &= \begin{pmatrix} \nu_e \\ e^- \end{pmatrix}_L, e_R, \\ Q_L &= \begin{pmatrix} u \\ d \end{pmatrix}_L, u_R, d_R, \end{aligned} \tag{1}$$

and the second and third generation can be written by the obvious substitution of $e \rightarrow \nu, \tau$, $u \rightarrow c, t$, and $d \rightarrow s, b$. This structure will return later in Section 2.3.1.

The bosons are much fewer in number than fermions. The only scalar boson (spin-0) is the Higgs boson which couples to all particles with mass. The vector bosons (spin-1) are the gluon which mediates the strong force between quarks, the photon which mediates the electromagnetic force between charged particles, and the $W^\pm$ and Z bosons which mediate the weak force.

2.1.1 Composite particles

Gluons and quarks share a property called *colour charge*. Colour charge is to the strong force as electric charge is to the electromagnetic force: it decides with whom gluons will interact; since quarks are the only fermions that carry colour charge, they are the only fermions that gluons interact with.

A very important property is *colour confinement*, which imposes the requirement that stable particles be colour-neutral. The result is that a bare quark, which carries one unit of colour (call it *red*= r), cannot exist in isolation. It can either be paired

with an anti-quark (carrying *anti-red*= $\bar{r}$), or grouped with two other quarks (carrying *blue*= b , and *green*= g colour charges), such that the overall particle is colour-neutral. These objects are held together by exchange of gluons, which themselves are colour-neutral, carrying for example $r\bar{r}$ colour. Collectively these composite particles are called *hadrons*, while two-quark hadrons are called *mesons*, and three-quark hadrons are called *baryons*. The only quark which does not form hadrons is the top; because its lifetime is so short it has no time to hadronize before it decays.

The most familiar baryon is the proton. A proton is composed of two up quarks and a down quark (uud), held together by a constant exchange of gluons. In fact, these quarks represent only what we call *valence quarks* in the proton: much like valence electrons determine the electromagnetic interactions of atoms, valence quarks determine the electroweak interactions of hadrons. “Underneath” the valence quarks is a sea of other $q\bar{q}$ pairs. A visualization of the substructure of a proton is shown in Figure 1.




Figure 1: Substructure of the proton, taken from [14].

The other familiar baryon is the neutron (udd), though many other baryons states are produced at the LHC. A myriad of mesons are also produced in collisions at the LHC. Some survive long enough to be observed directly (e.g. the most common is $\pi^\pm = u\bar{d}, \bar{u}d$), while others decay very quickly and are observed indirectly by their decay products. A sub-class of mesons worth distinguishing are “B mesons”, a term used to mean any meson with a valence bottom quark. While all mesons produce jets within ATLAS’ calorimeters (the calorimeters are discussed further in Section 3.2.3, while jet reconstruction is discussed further in Section 4.6), the topology jets from B mesons differ from those of the lighter mesons.

The quarks and gluons that make up a hadron are collectively called *partons*. The momentum fraction of the proton that is carried by each parton is distributed

according to their parton distribution function, a probability density function (PDF¹) on the range $[0, 1]$. Parton density functions are empirically determined, but as an approximate figure, a proton's momentum is divided roughly into 6 equal pieces to its three constituent quarks and three gluons. The important implication for the LHC is that when protons collide at 6.5 TeV, in reality it is partons with approximately ≈ 1 TeV energies which collide.

2.2 Interactions

It is the interactions between these particles which allow us to observe and understand them². Together the Strong and Electroweak interactions can describe virtually every interaction occurring inside ATLAS. The formal language of the SM describes these interactions using *symmetry groups*. Loosely speaking, the SM must be invariant under certain symmetry transformations, and this requirement precipitates the interactions we observe. Theories beyond the SM, discussed more in Section 2.4, try to impose new symmetry groups on the SM, and make predictions of processes that will be observed, if the postulated theory is true. The detailed mathematical structure of the SM is beyond the scope of this thesis, so this section aims for a mostly qualitative understanding: which particles interact with each other, and how these interactions can be used to describe processes observed in ATLAS.

Particle interactions can handily be illustrated using Feynman diagrams. These diagrams can be used to show the structure of any process, but are even more useful as calculation tools for evaluating the cross section of the process. For the purposes of this thesis, these calculations are left to phenomenologists, and the diagrams are used only graphically. In general any diagram which can be written down (using the basic pieces discussed in the following sections) does occur, but the rate of the interaction must be calculated, and could be negligible. The diagrams can also be rotated, flipped, or twisted in any orientation, and still be valid. The usual convention has time flowing from left-to-right each diagram. The greater the number of *vertices* in the diagram, the smaller its cross section. The form of a diagram with the least number of vertices is called the *leading order* (LO) diagram for a process, and gives a

¹Much of the literature will use 'PDF' to mean parton distribution function, but as this thesis is more often concerned with probability density functions, the acronym is defined accordingly.

²This statement holds true from a practical standpoint, that the particles must interact in some way in our particle colliders and detectors for them to be recorded, but it also holds true from an anthropic standpoint, since particles must be able to interact to form atoms and chemistry which allows the evolution of (eventually) intelligent life which can go on to build particle colliders and ask these questions.

useful visual of the initial and final-state particles. For precise predictions, evaluating higher order diagrams is necessary. As more vertices are added, these are called next-to-leading order (NLO), next-to-next-to-leading order (NNLO), and even next-to-next-to-next-to-leading order (NNNLO, or just N3LO). All the diagrams shown in this thesis are given at LO, though higher-order calculations are always used, or corrected for, in the simulated data. This is discussed further in Section 4.1.

2.2.1 Strong interaction

The Strong interaction is described by the theory of Quantum Chromodynamics (QCD) [15]. It is mediated by gluons, which interact (or “couple to”) colour charge. In practice this means they interact only with quarks and other gluons. In the mathematical formulation of the SM, the strong force is described by an $SU(3)_C$ symmetry group, where C is the colour-charge, and the gluons are actually an octet of 8 particles. Most of the LHC collisions begin with QCD, since the colliding particles are composed of quarks and gluons. The basic Feynman diagram pieces for QCD are shown in Figure 2. Any combination of these pieces can be “snapped” together to create a more complex process.




Figure 2: Feynman diagram components for strong interactions.

2.2.2 Electroweak interaction

The electroweak (EW) interaction is the unification of the electromagnetic and weak forces under one theory [15]. At low energy scales, far below m_W , they can be described separately, but at higher energies they are combined into the single EW force, mediated by photons and the vector bosons, $W^\pm, Z$. The symmetry governing these interactions is $SU(2)_L \times U(1)_Y$, where L denotes left-handed chirality, and Y is the weak hyper-charge (analogous to the electric charge). A truly self-consistent picture of the EW interactions is impossible without the Higgs mechanism which

will be discussed in Section 2.3.1, so for now its content must be taken at face value. Photons couple to electric charge, so interact only with charged particles. The $W^\pm, Z$ bosons couple to all left-handed fermions. When a $W^\pm$ interacts with a fermion, it always changes the particle's species (flavour) and charge, whereas neither the γ nor Z do so. The basic Feynman diagram pieces for EW interactions are shown in Figure 3.




Figure 3: Feynman diagram components for electroweak interactions.

These diagrams, together with Figure 2 are all that is needed to understand the source of the main backgrounds to the analysis in this thesis, and will return in Chapter 6. They are also essential in understanding the signal process under study in this thesis, $H \rightarrow ZZ^{(*)} \rightarrow 4\ell$.

2.2.3 Z boson decay

Of special importance to this thesis is the Z boson. The Z boson is an inherently unstable particle, with a large natural width Γ_Z . This width parameter determines, in part,

the shape of the production cross section of the Z boson, and therefore the distribution of Z boson masses. Although the “Z mass” is a single value, $m_Z = 91.1876 \text{ GeV}$ [16], individual Z bosons come with a range of masses - call them $m_{\ell\ell}$ for the di-lepton pair they may decay into. They are distributed, as for any unstable particle, according to a Relativistic Breit-Wigner (BW) distribution[15]:

$$BW(m_{\ell\ell}; m_Z, \Gamma_Z) = \frac{k}{(m_{\ell\ell}^2 - m_Z^2)^2 + m_Z^2 \Gamma_Z^2}, \quad (2)$$

where k is a normalization constant depending on m_Z and Γ_Z . This distribution is called the Z boson’s *lineshape*. It is shown in Figure 4. The width, Γ_Z , is the full width at half-maximum of the distribution.




Figure 4: Relativistic Breit-Wigner lineshape for Z boson decays.

At the LHC where the colliding quarks and gluons have a range of momentum, the production of the Z boson cannot be finely controlled, and so Z bosons are produced according to this lineshape. At an e^+e^- collider, where the colliding particles’ energy is finely controlled, the collider operator is in control of *where* on the Z lineshape the bosons are being produced. When the collider is tuned to the peak of the distribution, Z bosons are being produced “on-shell” (and at a higher rate), whereas away from the peak of the distribution, they are being produced “off-shell” (and at a lower rate). The same applies at the LHC, except that both on-shell and off-shell Z bosons will be produced simultaneously, distributed according to (2).

The decays $Z \rightarrow e^+e^-$ and $Z \rightarrow \mu^+\mu^-$ form the latter half of the $H \rightarrow ZZ^{(*)} \rightarrow 4\ell$

signal analyzed in this thesis (only $\ell = e, \mu$ are considered, τ 's are not part of the signal, by definition). They proceed exactly as seen in the top-center diagram of Figure 3. The branching ratio for these two decays [16] are

$$\begin{aligned}\text{BR}(Z \rightarrow \mu^+ \mu^-) &= 3.363 \pm 0.007\%, \\ \text{BR}(Z \rightarrow e^+ e^-) &= 3.363 \pm 0.004\%.\end{aligned}\tag{3}$$

The signal process studied in this thesis contains two Z bosons, which must both decay via one of these channels to be part of the signal.

2.3 The Higgs boson

The structure of the SM requires that its interactions be invariant under transformations of $SU(3)_C \times SU(2)_L \times U(1)_Y$ (combining the Strong and EW symmetry groups). The only way to achieve this is to require that all the fermions and bosons in the theory be *massless*, which stands in clear contention with experimental evidence. The reconciliation of this problem came in 1964 [1][17][18], with the introduction the (later-called) *Higgs mechanism*, which could explain the masses of the fermions and bosons by spontaneously breaking the $SU(2) \times U(1)$ symmetry in the theory. This mechanism is explained in Section 2.3.1. An important consequence of this mechanism is the prediction of another massive boson, the *Higgs boson*, whose production, decay, and use in this thesis are described in Section 2.3.2 and Section 2.3.3.

2.3.1 The Higgs mechanism

The Higgs mechanism works by introducing into the SM a doublet of complex scalar fields, called the Higgs doublet [15],

$$\Phi = \begin{pmatrix} \phi^+ \\ \phi^- \end{pmatrix},\tag{4}$$

which transforms under $SU(2)_L$. By introducing this doublet, the SM Lagrangian, which encodes all the particles and interactions in the theory, gets an additional contribution from

$$\mathcal{L}_{SM} \supset |D^\mu \Phi|^2 - V(\Phi),\tag{5}$$

where D^μ is the *covariant derivative*, used to represent the interactions of the doublet with other particles (technically, fields), and the $V(\Phi)$ is the Higgs potential, which can be parameterized as

$$V(\Phi) = \mu^2 \Phi^\dagger \Phi + \lambda (\Phi^\dagger \Phi)^2. \quad (6)$$

To guard against vacuum instability, we require $\lambda > 0$ (otherwise the potential falls to $-\infty$ for large values of $\Phi^\dagger \Phi$). If $\mu^2 < 0$, the potential takes on the form shown in Figure 5.




Figure 5: Shape of Higgs potential function leading to Spontaneous Electroweak Symmetry Breaking. Taken from [19].

This potential now has a non-zero minimum located in a ring of radius $\Phi^\dagger \Phi = v/\sqrt{2}$, where $v = \sqrt{-\mu^2/\lambda}$. This value v is called the *vacuum expectation value*, equal to 246 GeV. As the vacuum falls down to this non-zero minimum, the $SU(2)_L \times U(1)_Y$ symmetry describing the EW interactions is spontaneously broken. This process is called Electroweak Symmetry Breaking (EWSB).

Using a clever transformation to simplify (4), Φ can be expressed using a new field, H , which is a perturbation about the minimum of $V(\Phi)$,

$$\Phi = \frac{1}{\sqrt{2}} \begin{pmatrix} 0 \\ v + H \end{pmatrix}, \quad (7)$$

Using this expression, (6) becomes

$$\begin{aligned} V(\Phi) &= \frac{-\lambda v^2}{2} (v + H)^2 + \frac{\lambda}{4} (v + H)^4, \\ &= \frac{\lambda}{4} H^4 + \lambda v H^3 + \lambda v^2 H^2 + \text{constant}. \end{aligned} \quad (8)$$

The covariant derivative will include a term $\frac{1}{2}(\partial_\mu h)^2$, such that the terms in the Lagrangian involving directly and only the H field become

$$\mathcal{L}_{SM} \supset \frac{1}{2}(\partial_\mu H)^2 - \lambda v^2 H^2 - \lambda v H^3 - \frac{\lambda}{4} H^4. \quad (9)$$

Here we recognize the Lagrangian of a self-interacting scalar, the SM Higgs boson. The term $\frac{1}{2}(\partial_\mu H)^2$ is the kinetic term, the $-\lambda v^2 H^2$ represents a mass term for the H , where $\frac{1}{2}m_H^2 = \lambda v^2$, and so $m_H = \sqrt{2\lambda}v = \sqrt{-2\mu^2}$. The final two terms, $-\lambda v H^3$ and $-\frac{\lambda}{4}H^4$ are three-Higgs and four-Higgs self-interaction terms.

To see why the Higgs mechanism is so important to the SM, we turn to the covariant derivative, D^μ , which will show how the Higgs doublet's interactions with the EW sector are affected by EWSB. The derivative has the form

$$|D_\mu \Phi|^2 = \left| \left(\partial_\mu + \frac{ig_2 \tau_a W_\mu^a}{2} + \frac{ig_1 B_\mu}{2} \right) \Phi \right|^2. \quad (10)$$

In this expression, the $W_\mu^a = W_\mu^{1,2,3}$ are the massless fields associated to the $SU(2)_L$ symmetry, and the B_μ is the massless field associated to $U(1)_Y$. The τ are Pauli matrices, and g_1 (g_2) is the electromagnetic (weak) coupling constant. Inserting (7) and expanding the expression gives

$$|D_\mu \Phi|^2 = \frac{1}{2}(\partial_\mu H)^2 + \frac{1}{8}g_2^2(v+H)^2 |W_\mu^1 + iW_\mu^2|^2 + \frac{1}{8}(v+H)^2 |g_2 W_\mu^3 - g_1 B_\mu|^2, \quad (11)$$

where the first term was already extracted and used in (9). This expression can be simplified by a linear transformation to four new fields defined by,

$$\begin{aligned} W_\mu^\pm &= \frac{1}{\sqrt{2}}(W_\mu^1 \mp iW_\mu^2), \\ Z_\mu &= \cos(\theta_W)W_\mu^3 - \sin(\theta_W)B_\mu, \\ A_\mu &= \sin(\theta_W)W_\mu^3 + \cos(\theta_W)B_\mu, \end{aligned} \quad (12)$$

where $\tan \theta_W = g_1/g_2$ defines the weak mixing angle, θ_W . Using this transformation, several terms in (11) can be identified as

$$|D_\mu \Phi|^2 \supset \frac{g_2^2 v^2}{4} W_\mu^+ W^{-\mu} + \frac{g_2^2 v^2}{8 \cos^2(\theta_W)} Z_\mu Z^\mu + 0 \cdot A_\mu A^\mu. \quad (13)$$

As before, terms quadratic in a field represent masses of particles, from which we can

find the masses of these newly defined fields,

$$\begin{aligned} M_{W^\pm} &= \frac{1}{2}g_2v, \\ M_Z &= \frac{1}{2}\frac{g_2v}{\cos(\theta_W)}, \\ M_A &= 0. \end{aligned} \tag{14}$$

These fields are recognized (if the naming convention wasn't a giveaway already) as the physical $W^\pm, Z$ weak vector bosons, and the massless photon. The weak vector bosons acquired their mass through the introduction of the Higgs doublet, which breaks the $SU(2)_L \times U(1)_Y$ EW symmetry. The remaining symmetry, only on the massless A field, is called $U(1)_{EM}$, and is the more familiar symmetry group of the electromagnetic interaction. As a by-product of this symmetry breaking, a massive H field was introduced, which we recognize as the Higgs boson. This is the essence of the Higgs mechanism.

A similar mechanism works to generate the fermion mass terms. By introducing a coupling parameter for each fermion, the interaction with the Higgs doublet can be written as (for, e.g., the electron)

$$\mathcal{L}_{SM} \supset -\lambda_e \bar{E}_L \Phi e_R \tag{15}$$

where λ_e is a coupling parameter, and $\bar{E}_L, e_R$ were introduced in (1). By expanding Φ about the minimum using (7), we get a term

$$\mathcal{L}_{SM} \supset -\frac{1}{\sqrt{2}}\lambda_e v \bar{e}_L e_R, \tag{16}$$

which can be interpreted as a mass term for the electron,

$$m_e = \frac{1}{\sqrt{2}}\lambda_e v. \tag{17}$$

Note that since there is no right-handed neutrino, ν_R , neutrinos do not acquire a mass in this way. The SM is unable to explain the non-zero (but very small) neutrino masses observed experimentally.

For quarks the situation is a bit more complex as a linear transformation must be done to translate to the mass eigenstates. The final product however, appears similar to that of the fermions, where mass terms appear in the Lagrangian giving masses to

the quarks,

$$m_u = \frac{1}{\sqrt{2}}\lambda_u v, \quad m_d = \frac{1}{\sqrt{2}}\lambda_d v. \quad (18)$$

Also important for anyone wishing to observe a Higgs boson is to see how it interacts with particles, beyond just giving them mass. To see how the Higgs field couples with these new massive vectors bosons, we collect the terms in (11) with factors of H and H^2 , leaving the following in the Lagrangian,

$$\mathcal{L}_{SM} \supset \left[m_W^2 W_\mu^+ W^{-\mu} + \frac{1}{2} m_Z^2 Z_\mu Z^\mu \right] \left(\frac{2H}{v} + \frac{H^2}{v^2} \right). \quad (19)$$

This expression suggests interactions of the type $HVV, HHVV$, (where $V = W, Z$) with couplings

$$g_{HVV} = -2i \frac{m_V^2}{v}, \quad g_{HHVV} = -2i \frac{m_V^2}{v^2} \quad (20)$$

Similarly, the Lagrangian will contain terms for the fermions of the form

$$\mathcal{L}_{SM} \supset -m_f \bar{f} f \frac{h}{v}, \quad (21)$$

which suggest fermion couplings

$$g_{Hff} = i \frac{m_f}{v} \quad (22)$$

In all cases, the Higgs bosons' couplings are proportional to the mass of the particles with which it is interacting. The consequence is that the Higgs preferentially interacts with the heaviest particles in the SM: t quarks, Z , $W^\pm$, and b quarks. This coupling behaviour determines the ways in which Higgs bosons are produced and decay, discussed in the following section.

2.3.2 Higgs production and decay

The Higgs boson couples to any massive particles, which in the SM means it couples to all particles except the gluon and photon. The Feynman diagrams for Higgs interaction are shown in Figure 6. Only the top two diagrams contribute meaningfully to this thesis. With these two diagrams, and the preceding diagrams of the QCD and EW interactions, all of the Higgs production and decay modes considered in this thesis can be described.

In the bottom row, the three-Higgs and four-Higgs diagram are called the Higgs self-interaction diagrams, and the associated couplings are called the trilinear Higgs

coupling and quartic Higgs coupling, respectively. Measuring the strength of the trilinear Higgs coupling is one of the goals of the final stage of the LHC programme, called the High-Luminosity LHC (HL-LHC). Recent studies[20] project that it will be just in reach with 3000 fb^{-1} of data at $\sqrt{s} = 14 \text{ TeV}$, which the HL-LHC anticipates to collect by the end of its life. The quartic Higgs coupling is far out of reach of the LHC, and may only be accessible in a future collider. Measuring these couplings is valuable because it is equivalent to a measurement of each term in the Higgs Lagrangian, (9). Measurement of these terms would fully determine the shape of the Higgs potential, (6) (and shown in Figure 5), which has consequences as far reaching as the long-term stability of the universe.




Figure 6: Feynman diagram components for Higgs interactions. Only the upper diagrams contribute at rates relevant to this thesis.

2.3.2.1 Production

The mode of production of Higgs bosons depends on the collider at which they are produced, as well as its energy. At the LHC, the dominant production modes are necessarily either gluon or quark-initiated, whereas at an e^+e^- collider they would be different. The Higgs bosons analyzed in this thesis are produced through four production modes:

ggF The dominant production mode by far is *gluon gluon fusion* (ggF), wherein two gluons fuse via a loop of top quarks to produce a Higgs. The extremely high mass of the top quark is what makes this production mode so dominant, as its

coupling to the Higgs is far higher than any other particle. A contribution from a similar loop of bottom quarks is included in this production cross section, but it represents a very small fraction of ggF production, scaling like $(m_b/m_t)^2$ compared to the top loop. Other quarks technically enter into the loop as well, with even smaller contributions. The leading-order Feynman diagram for ggF production is shown in Figure 7.




Figure 7: Feynman diagram for ggF Higgs production.

VBF The second most important production mode is *vector boson fusion* (VBF), approximately an order of magnitude down from ggF. In VBF production, a pair of vector bosons (either ZZ or W^+W^-) radiate from quarks in the two colliding beams. These bosons fuse to produce the Higgs boson. The signature which distinguishes VBF from ggF is the pair of quarks appearing in the final state (observed as jets, to be explained in Section 4.6). The leading-order Feynman diagram for VBF production is shown in Figure 8.




Figure 8: Feynman diagram for VBF Higgs production.

VH The third leading Higgs production mode is *associated production* (VH), in which a Higgs boson is produced in association with a vector boson. In this mode, an

incoming quark-anti-quark pair fuse to produce a $W^\pm$ or Z boson, which then radiates a Higgs boson. The leading-order Feynman diagram for VH production at the LHC is shown in Figure 9.




Figure 9: Feynman diagram for VH Higgs production.

ttH, bbH The rarest Higgs production mode mentioned in this thesis is *top fusion*, or *bottom fusion* (ttH, or bbH). In this process, a Higgs boson is produced in association with a $t\bar{t}$ or $b\bar{b}$ pair. The ttH process is enhanced over bbH by the Higgs' strong coupling to top quarks due to their high mass, but suppressed by the enormous energy required to produce two tops (and a Higgs) in a single collision, as compared to that required to produce two bottoms (and a Higgs). Coincidentally, at $\sqrt{s} = 13$ TeV, these two competing factors balance out, and ttH and bbH have nearly identical cross sections. The leading diagram for both processes resembles that for ggF, but the quarks no longer form a loop, instead they survive as final state particles. It is shown in Figure 10.




Figure 10: Feynman diagram for ttH (bbH) Higgs production.

Figure 11 shows the cross section for each production mode, as a function of the Higgs boson mass. In this diagram, ggF is labelled as $pp \rightarrow H$, VBF is labelled as $pp \rightarrow qqH$, and VH is split into $pp \rightarrow WH$ and $pp \rightarrow ZH$. Table 3 summarizes

the cross sections and fraction of each production mode expected by the SM for a $m_H = 125 \text{ GeV}$ Higgs in $\sqrt{s} = 13 \text{ TeV}$ collisions at the LHC.




Figure 11: Plot of the SM predicted Higgs cross section for different production modes at the LHC at $\sqrt{s} = 13 \text{ TeV}$ as a function of the Higgs mass. Taken from [21]

Table 3: Production cross sections for a $m_H = 125 \text{ GeV}$ Higgs boson in $\sqrt{s} = 13 \text{ TeV}$ collisions. Values taken from [21], with uncertainties symmetrized and summed in quadrature, where necessary.

production mode	cross section [pb]	fraction [%]
ggF	48.58 ± 2.45	87.3
VBF	3.78 ± 0.08	6.8
VH	2.26 ± 0.04	4.1
ttH	0.51 ± 0.04	0.9
bbH	0.49 ± 0.11	0.9

2.3.2.2 Decay

Like many massive particles, the Higgs boson is inherently unstable. The SM predicts the lifetime of a 125 GeV Higgs to be approximately 10^{-22} s. The lifetime is related to the total width of the Higgs boson by $\tau_H = 1/\Gamma_H$, where Γ_H is the Higgs total width, or natural width. The decays of the Higgs boson occur by several modes, in proportion to their branching ratios (BR). The Higgs BR is predicted by the SM as a function of the Higgs mass, m_H , and is independent of the collider at which it is produced. An unstable particle decays preferentially to the particles it is most strongly coupled (assuming enough kinematic phase space to allow conservation of energy and momentum). For the Higgs, this means it prefers decays to heavy particles, though these often subsequently decay to light particles.

The BR of the SM Higgs boson as a function of its mass is shown in Figure 12. There is no decay to $t\bar{t}$, because the Higgs is lighter than the top quark. The decays to massless gg , $\gamma\gamma$, and (partly massless) $Z\gamma$ occur via a top loop (like the ggF production process, Figure 7, in reverse). The decays to W, Z boson pairs depend most strongly on m_H because the combined mass of the decay products exceed the mass of the Higgs. For this reason, one of the bosons must be produced off-shell, so is denoted as (e.g.) Z^* . The higher the m_H , the closer to on-shell the Z^* can be, so the $\text{BR}(H \rightarrow ZZ^*)$ grows very quickly: by over 20% between 124-126 GeV. By including the branching ratios for the subsequent Z decay, the total branching ratio to 4ℓ final states can be computed. The values of the branching ratios $\text{BR}(H \rightarrow ZZ^*)$ and $\text{BR}(H \rightarrow ZZ^{(*)} \rightarrow 4\ell)$ are given in Table 4 for a few m_H points.

Table 4: Branching ratios $\text{BR}(H \rightarrow ZZ^*)$, and $\text{BR}(H \rightarrow ZZ^{(*)} \rightarrow 4\ell)$ for different masses m_H . Values taken from [21]. The final column implicitly includes the middle column - they should not be multiplied.

m_H [GeV]	$\text{BR}(H \rightarrow ZZ^*)$ [%]	$\text{BR}(H \rightarrow ZZ^{(*)} \rightarrow 4\ell)$ [%]
124.0	2.38 ± 0.04	$(1.13 \pm 0.03) \times 10^{-2}$
124.5	2.50 ± 0.04	$(1.19 \pm 0.03) \times 10^{-2}$
125.0	2.62 ± 0.04	$(1.24 \pm 0.03) \times 10^{-2}$
125.5	2.74 ± 0.04	$(1.30 \pm 0.03) \times 10^{-2}$
126.0	2.87 ± 0.04	$(1.36 \pm 0.03) \times 10^{-2}$

Like the Z boson, the production cross section of the Higgs, and thus its invariant




Figure 12: Plot of the SM predicted Higgs branching ratio as a function of the Higgs mass. Taken from [21].

mass distribution, follows a BW. The lineshape for a SM $m_H = 125$ GeV Higgs is shown in Figure 13. Note that the SM Higgs width is nearly three orders of magnitude smaller than that of the Z boson. For $m_H = 125$ GeV, the width predicted by the SM is just 4.1 MeV. This lineshape will return in an important way in Chapter 7.




Figure 13: Relativistic Breit-Wigner lineshape for Higgs boson decays, drawn for a $m_H = 125$ GeV Higgs.

2.3.3 The $H \rightarrow ZZ^{(*)} \rightarrow 4\ell$ channel

This thesis presents measurements of the Higgs boson through its decay to four leptons: the $H \rightarrow ZZ^{(*)} \rightarrow 4\ell$ decay channel. By combining the diagrams in the preceding sections, the dominant way this process happens at the LHC is shown in Figure 14.




Figure 14: Feynman diagram for production and decay of $H \rightarrow ZZ^{(*)} \rightarrow 4\ell$.

This decay channel is a very powerful tool for measuring the Higgs' properties. The primary reason is that the final state is composed entirely of muons and electrons, which are both precisely reconstructed with very high efficiency. Secondly, the channel has a very high signal-to-background ratio (after the event selection described in

Chapter 5, it is $\mathcal{O}(1)$). For this reason, it was called the *golden channel* for Higgs discovery, and was one of the two channels essential to the discovery of the Higgs boson in 2012. In this thesis, it is used to make direct measurements of the Higgs mass, m_H , width, Γ_H , and total cross section times branching ratio, defined as

$$\sigma_{H \rightarrow 4\ell} \equiv \sigma_{p\text{-}p \rightarrow H} \times \text{BR} \left(H \rightarrow ZZ^{(*)} \rightarrow 4\ell \right), \quad (23)$$

where $\sigma_{p\text{-}p \rightarrow H}$ is the total cross section at $\sqrt{s} = 13 \text{ TeV}$ combining all production modes, and $\ell = e, \mu$. In the SM, m_H is a free parameter left to measure, while Γ_H and $\sigma_{H \rightarrow 4\ell}$ can be calculated once it is known. The SM predictions, when $m_H = 125 \text{ GeV}$, are

$$\begin{aligned} m_H &= 125.00 \text{ GeV}, \\ \Gamma_H &= 4.09 \pm 0.03 \text{ MeV}, \\ \sigma_{H \rightarrow 4\ell} &= 6.90 \pm 0.31 \text{ fb}. \end{aligned} \quad (24)$$

Measuring these parameters is the primary goal of this thesis. Although the SM offers no prediction of m_H , it has an innate value as a fundamental parameter of the theory. Additionally, precise knowledge of m_H improves the precision of other SM predictions, among them Γ_H and $\sigma_{H \rightarrow 4\ell}$. More importantly, precise m_H knowledge improves the predictive precision on the more rare Higgs production processes, which allow one to disentangle the Higgs' couplings to individual particles. The best existing measurement of m_H comes from an ATLAS+CMS combination of the Run 1 LHC data in the $H \rightarrow ZZ^{(*)} \rightarrow 4\ell$ and $H \rightarrow \gamma\gamma$ channels [8], where the mass was measured to be $m_H = 125.09 \pm 0.24 \text{ GeV}$. Although the measurement in this thesis cannot surpass the precision of this combined measurement, it exceeds the precision of any $H \rightarrow 4\ell$ mass measurement conducted so far. A review of the other existing measurements of m_H , $\sigma_{H \rightarrow 4\ell}$, and Γ_H will accompany the results of the measurements in Chapter 9 for ease of comparison.

2.4 Beyond the Standard Model

While the SM succeeds in describing nearly all observed phenomena with incredible precision, it is still an incomplete theory. Beyond Standard Model (BSM) physics refers to the phenomena observed in nature but not explained by the SM[22], and also to extensions of the SM attempting to explain these phenomena. Most BSM models will in some way affect the Higgs sector, which makes precision Higgs physics a powerful probe for testing them.

One of the most important questions for future experiments is the nature of Dark Matter (DM) - a class of particles making up 84% the matter content of our universe yet wholly unknown to the SM. The existence of Dark Matter is supported by an abundance of evidence from astrophysical and cosmological measurements [23], but has never been observed in a laboratory. One thing we do know is that *some of it* (and maybe most of it) is composed of non-SM particles, which we call Weakly Interacting Massive Particles (WIMPS). To explain the observed phenomena, WIMPS must have mass. The only known way particles can acquire mass is through the Higgs mechanism, which would mean that the WIMPS, even if they don't interact with any other SM particles, will interact with the Higgs boson. If this were the case, it is possible that the Higgs boson could decay to DM particles. Searches by ATLAS [24] have not detected any such decays, but even if the Higgs does not decay to DM particles, for example because they are too heavy for a 125 GeV Higgs to decay into, their coupling would still affect the production cross sections for the Higgs, and thus $\sigma_{H \rightarrow 4\ell}$.

Many BSM theories postulate an extended Higgs sector [25] to help resolve discrepancies in the SM. In these models, the observed Higgs near 125 GeV is the lightest of several Higgs bosons. These heavier counterparts have been searched for, including one search by the author [26], though no direct evidence of a heavy Higgs has been found. Direct detection is however not the only way an extended Higgs sector can reveal itself. The existence of a heavy Higgs also alters the couplings of the 125 GeV Higgs such that the cross sections, thus $\sigma_{H \rightarrow 4\ell}$, and the total width, Γ_H , are modified compared to the SM. As a result, precision measurement of these properties is itself an indirect test of a whole class of BSM theories.

The above examples are not meant to be complete, but rather to illustrate how precision measurement of an SM particle can be a powerful tool for finding clues, and guiding future searches and measurements, of BSM physics.

Chapter 3

ATLAS and the LHC

Though production of Higgs bosons can occur spontaneously in nature, the rate for this occurrence is prohibitively low for any observation[27]. For this reason, experimental searches and studies of the Higgs boson use high-energy particle colliders to actively produce Higgs bosons, while detectors observe the outcome of these collisions. By happenstance the Higgs boson has a mass near 125 GeV , and the collider currently best suited to its production is the Large Hadron Collider. Had the mass been lower, another collider such as the Tevatron or Large Electron-Positron collider would have preempted the LHC in the discovery and initial property measurement of the Higgs boson.

Section 3.1 will describe the Large Hadron Collider, used to produce the Higgs bosons under study. The products of these collisions are recorded by the ATLAS Detector, described in Section 3.2, and it is these data analyzed throughout the thesis. Section 3.3 will describe plans for a future hardware upgrade of the ATLAS Detector, with the author’s particular involvement described in Section 3.3.1.1.

3.1 The LHC

The Large Hadron Collider (LHC) [28] is the largest particle collider in the world. It is a circular collider with a circumference of approximately 27 km, located 100 m underground straddling the border between Switzerland and France, operated by the European Organization for Nuclear Research (CERN, an acronym based on its French name).

The primary use of the LHC is as a proton-proton (p - p) collider. Two adjacent beam pipes contain beams of protons circulating in opposite directions around the ring. The beams are clumped in 2808 bunches, each bunch containing approximately

10^{11} protons, each proton with an energy of 6.5 TeV. At this enormous energy, protons circle the LHC 11,000 times per second, travelling at 99.9999990% the speed of light. The opposing beams cross at four collision points along the ring. During each so-called “bunch crossing”, separated by intervals of 25 ns during normal operation, some of the protons in the bunches may collide, while the remainder of the bunch continues to circulate the LHC. The products of these collisions are observed by detectors at each of the four collision points: ATLAS, CMS, ALICE, and LHCb.

ATLAS (**A Toroidal LHC Apparatus**) [29] and CMS (**Compact Muon Solenoid**) [30] are both general-purpose detectors with similar physics goals, but distinct designs. ALICE (**A Large Ion Collider Experiment**) [31] is a more specialized detector designed to study quark-gluon plasma using the LHC’s other mode of operation: collision of lead ions ($^{208}\text{Pb}^{82+}$), called Pb-Pb collisions. LHCb (**LHC-beauty**) [32] is also a specialized detector, focusing on the study of bottom quarks and their interactions. The relative locations of the four detectors on the LHC ring is shown in Figure 15.

3.1.1 From canister to collision

Protons in the LHC begin their journey [33] in a small can of hydrogen gas. An electric field strips the electrons from the hydrogen, leaving only the nuclei which are bare protons. These protons pass into Linac2, a linear accelerator which uses Radio-Frequency Quadrupoles (RFQs) to focus the beam of protons and accelerate them to 50 MeV.

Linac2 is the only linear accelerator in the chain, before the protons pass into circular accelerators. The advantage of circular accelerators over linear accelerators is that the particles pass through the system many times, allowing the machine to give many small “kicks”, adding more energy with each orbit. All circular accelerators share the same main components: radio-frequency (RF) cavities are used to accelerate the particles, electromagnetic dipoles are used to bend the particle trajectories along the ring, while higher-order multipoles (quadrupole, sextupoles, octopoles) are used to shape and focus the beam. In order for the particles to be accelerated as they orbit, the dipole field strength, B , must ramp in sync (thus they are called “synchrotrons”) with the particles momenta, p , such that the bending radius given by

$$\rho = \frac{pc}{B} \tag{25}$$

remains constant, and equal to the physical radius of a circular accelerator. Here c

is the speed of light, and we assume particles with charge equal to one proton (or electron). The field strength must also compensate the incomplete coverage of the circumference ring by the dipole magnets (because of practical limitations), so must be proportionately stronger than for full coverage. As a result of the dynamic field strength of a synchrotron, all particles circulating must have the same momentum, so adding more particles during the acceleration is impossible. Typically the accelerator is filled at a lower starting energy, then all particles are accelerated to the target energy, and finally they are all extracted (or collided) to empty the accelerator and start over.

In the second stage of the accelerator chain, the protons enter the Proton Synchrotron Booster (PSB), a set of four stacked synchrotrons. The proton beam is split into these four synchrotrons, and then RF cavity harmonics with the orbiting beams are used to split these beams into bunches. The beams are merged (but not de-bunched) during extraction.

The third stage of the accelerator chain is the Proton Synchrotron (PS). The PS further splits the PSB bunches, and accelerates the protons up to 25 GeV.

The fourth stage of the accelerator chain is the Super Proton Synchrotron (SPS) which once operated as a proton-antiproton ($p\bar{p}$) collider, providing collisions for the Nobel-prize-winning [34] discovery of the W [35] and Z [36] bosons in 1984. Several injection cycles from the PS are required to fill the SPS with the full bunch structure of the LHC. Once full, the SPS accelerates the protons up to 450 GeV.

The final stage of the chain is the LHC itself. Bunches are extracted from the SPS at two points, leading to the two separate beam pipes in the LHC: one circulating clockwise and the other anticlockwise. Once the LHC is filled at 450 GeV, the particles are accelerated over the course of about 20 minutes to reach their maximum energy of 6.5 TeV per proton, when they are finally ready to be collided. The full accelerator chain is illustrated in Figure 15.

3.1.2 Collisions

A critical parameter in collider physics is the luminosity of the collisions, roughly speaking the amount of data collected. The amount of data being collected per unit time is the *instantaneous luminosity*, denoted by $\mathcal{L}$, whereas the total amount of data collected (in 2015 for example) is called the *integrated luminosity* (or commonly as in this thesis, just *luminosity*), written in shorthand as




Figure 15: Illustration of the LHC accelerator chain and the four detectors, based on [33]. Energies are per-proton. Sizes are not to scale.

$$\int \mathcal{L} = \int \mathcal{L}(t) dt \quad (26)$$

The importance of the luminosity becomes clear when we consider the measurement of (or search for) some process, $pp \rightarrow X$. If we observe some number of events, $N_{pp \rightarrow X}$, we may wish to compare to a predicted cross section, $\sigma_{pp \rightarrow X}$, from theory. To do this we use the integrated luminosity as follows:

$$\sigma_{pp \rightarrow X} = \frac{N_{pp \rightarrow X}}{\epsilon \int \mathcal{L}} \quad (27)$$

Where the quotient ϵ is introduced to represent an efficiency: what fraction of the produced events were expected to be observed. In Section 7.1, the $\sigma_{H \rightarrow 4\ell}$ parameter being measured in this thesis is parameterized in a similar way, with some additional caveats.

The instantaneous luminosity can be calculated directly from beam parameters [37]. For head-on beams, it is expressed as

$$\mathcal{L} = \frac{n_b n_1 n_2 f_r}{2\pi \Sigma_x \Sigma_y} \quad (28)$$

where n_b is the number of bunches per beam, n_i is the number of protons in beam i , f_r is the revolution frequency of the beams, $\Sigma_\mu = \sqrt{\sigma_{1\mu}^2 + \sigma_{2\mu}^2}$ for $\mu = x, y$ represents the x-y decomposition of the beam profiles. In reality the beams of the LHC cross at a small angle in order to control the point of collision, and in practice the equation must be modified slightly to take this into account. From (28) it is clear that $\mathcal{L}$ has units of $\text{s}^{-1}\text{m}^{-2}$, and thus $\int \mathcal{L}$ has units of m^{-2} . The units of inverse area are commonly expressed in pb^{-1} or fb^{-1} , which provides a convenient scale to work in (e.g. this thesis presents an analysis of 36.1 fb^{-1} of data).

To understand the collisions themselves it must be noted that, as introduced in Section 2.1, protons are not fundamental particles, they are made up of partons. It is the partons which interact directly, not the protons themselves¹. The two partons whose interaction makes up the bulk of the momentum transfer form what is called the hard-scatter (HS) interaction (located at the hard-scatter vertex). It is in these high-energy interactions where we search for Higgs bosons.

While the hard-scatter interaction is the most interesting for physics analyses, it is not the only interaction during a collision event. The remainder of the two protons form what is called the underlying event (UE) [39]. These leftover partons undergo multiple-parton-interaction (MPI), a collective interaction involving soft scatter and hadronization. The MPI process produces showers of particles, referred to as hadronic jets. UE is also taken to include beam-beam remnants, where other protons in the beam scatter elastically, producing additional jets near to the beam axis. Figure 16 illustrates a collision event with these components indicated.

A further source of activity in the detector is what we call pileup (PU). Pileup occurs in two forms: in-time pileup, and out-of-time pileup. In-time pileup refers to interactions from other protons within the same bunch crossing as the hard-scatter event. Out-of-time pileup refers to interactions from preceding or subsequent bunch crossings. In both cases, the products of the less interesting collisions are recorded by the detector indistinguishably from the hard-scatter, obfuscating the signal of interest. Pileup is sensitive to the instantaneous luminosity given in (28). The LHC beam

¹This is only a partial truth for two reasons:.. At low momenta, such that the de Broglie wavelength, $\lambda = h/p$ (where h is Planck's constant, and p is the particle momentum), is much larger than the effective radius of the proton, the substructure is irrelevant to the interactions, and we can consider direct proton-proton interactions. At high momenta, maybe 10 or 10^{10} times more than the LHC, perhaps it will be revealed that quarks themselves are not fundamental, and the same arguments can be made for them. The more complete truth is that the substructure of protons is essential to describing interactions at the LHC, whereas the [unknown, possibly non-existent] substructure of quarks and gluons is, so far, not [38].




Figure 16: Diagram showing distinction between underlying event and hard-scatter. The HS event in this case is a $H \rightarrow ZZ^{(*)} \rightarrow 4\ell$ event. Diagram inspired by [40].

parameters are tuned to maintain a level of pileup ATLAS and the other detectors can cope with, while still delivering the maximum amount of data.

3.2 The ATLAS Experiment

The ATLAS Experiment (or just “ATLAS”) is the largest detector on the LHC ring [29]. It is approximately cylindrical in shape, standing 25 m tall (and wide), and 45 m long. The detector weighs nearly 7000 tons, and is located along the LHC beam line, 100 m underground. The LHC beams collide at exactly (or nearly) the geometric center of the detector, and the debris flies outward in all directions through the detector from there.

The function of ATLAS is to record the particles that fly out of the collisions: what species of particle, what was its charge, momentum, where did it originate, etc. Knowing this information allows the study of the physical processes at hand. To this end, ATLAS is composed of many subdetectors specialized for different types of measurements. These subdetectors are illustrated in Figure 17 and will be described in detail in the following sections.




Figure 17: Diagram of the ATLAS Detector with various components illustrated. Image taken from [41].

3.2.1 The ATLAS coordinate system

In order to describe the detector, we must first define a coordinate system. The ATLAS detector uses a mix of Cartesian, cylindrical, and spherical coordinates. Depending on the application, one system of coordinates may lend itself more readily than another. All coordinate systems share the same origin, the geometric center of ATLAS, defined by $x = y = z = r = 0$.

The Cartesian coordinate system, (x, y, z) , used in ATLAS is defined as follows. The x axis points toward the center of the LHC. The y axis points up toward the sky (but is off vertical by 0.704° , because the LHC is not completely level). The z axis points along the beam line in the anti-clockwise direction. The Cartesian coordinate system is illustrated in Figure 18a.

The cylindrical coordinate system, (r, z, ϕ) , shares the same z coordinate as the Cartesian system. The transverse plane (x - y) is transformed into a radial distance r , and an azimuthal coordinate ϕ on the range $(-\pi, \pi]$, with $\phi = 0$ on the positive x -axis, and positive (negative) ϕ representing the top (bottom) half of the detector. The cylindrical coordinate system is illustrated in Figure 18b.

The spherical coordinate system introduces the polar angle θ , on the range $[0, \pi]$, defined as the angle from the positive z -axis. While normally a spherical coordinate

system would include a radial coordinate, this has little practical use in ATLAS. The spherical coordinate system is illustrated in Figure 18c.




Figure 18: Illustration of the three coordinate systems used in ATLAS.

3.2.1.1 Rapidity and pseudorapidity

Before moving on, consider the p - p collisions at the LHC. Both protons participating in the collision have the same momentum, 6.5 TeV, but the actual colliding partons have only a random fraction of this momentum, distributed according to their PDFs. This means that the center-of-momentum (c.o.m.) of the collision is generally not the rest frame of the ATLAS detector: the two have a large relative velocity along the z -axis.


How can we compare multiple data from one detector when each process occurred in a different c.o.m. frame? The straightforward idea would be to boost the data recorded in the ATLAS frame to the HS frame. Unfortunately this is generally not possible. To do so would require reconstructing the momentum of all particles in the detector, and selecting correctly only those which came from the HS vertex in order to balance the momenta and identify the HS frame. This is ruled out both by the production of neutrinos in many physical processes, and by the presence of products from UE and PU processes (especially near to the z -axis).

The solution to this problem is the introduction of rapidity [42], defined by

$$y = \frac{1}{2} \ln \left(\frac{E + p_z c}{E - p_z c} \right). \quad (29)$$

For a relativistic coordinate transformation (called a "boost") at speed $v = \beta c$ along the z -axis, the rapidity transforms as

$$y' = y - \tanh^{-1} \beta, \quad (30)$$

Figure 19: Diagram showing conversion between polar angle θ and pseudorapidity η .

which leads to a valuable property. Consider two particles measured in the ATLAS detector to be separated in rapidity by $\Delta y = y_1 - y_2$ in the ATLAS frame. What was their rapidity separation in their c.o.m. frame, which as discussed above is boosted along the z -axis?

$$\Delta y' = y'_1 - y'_2 = (y_1 - \tanh^{-1} \beta) - (y_2 - \tanh^{-1} \beta) = y_1 - y_2 = \Delta y. \quad (31)$$

Ergo, separations in rapidity are invariant under boosts along the z -axis. This very helpful property allows the comparison of *rapidity separations* from many different HS frames, as the separation in each HS frame is identical to what we can measure in the ATLAS frame.

In practice we can rarely measure both the momentum (or its component p_z) and the energy of a particle independently, and so it is hard to calculate the rapidity from (29). However if we examine the behaviour for highly relativistic particles, where $E, p \gg m$, we can reduce (29) to a simpler expression for a variable called pseudorapidity, defined by

$$\eta = -\ln \left(\tan \left(\frac{\theta}{2} \right) \right) \quad (32)$$

All particles directly observed (e.g. electrons, muons) in ATLAS will satisfy this condition. Instead of describing directions by (θ, ϕ) , they are described by (η, ϕ) . Figure 19 offers a fast visual conversion between θ and η . Heavier particles, such as h, Z, W bosons, must be described by the full rapidity, y , and Figure 19 cannot be used.

The pseudorapidity has another more qualitative advantage over polar angle, which is that particle production in a collider is roughly flat in η . This is useful information in detector design as it allows predictions for relative particle fluxes between different components of the detector - important to consider when designing for radiation hardness, granularity, detector live-time, etc. As a consequence, in the following sections many of the detector components are described by their location in η . Low (high) $|\eta|$ regions are also called *forward* (*central*) or *end-cap* (*barrel*), though the specific boundary in η varies by context.

3.2.2 Inner Detector

The ATLAS Inner Detector (ID) is the closest sub-detector to the collision point, and the first that particles encounter. It occupies the space $|r| < 1.15\text{ m}$, $|z| < 3.5\text{ m}$ with pseudorapidity coverage up to $|\eta| = 2.5$. The ID is composed of four main subsystems; from the beampipe outwards they are: the Insertable B-Layer (IBL), the Pixel Detector, the Semiconductor Tracker (SCT), and Transition Radiation Tracker (TRT). The entire ID is enclosed within a superconducting solenoid with inner radius $r = 1.2\text{ m}$ and magnetic field 2 T oriented along the z -axis. Due to the orientation of this field, charged particle tracks in the ID bend in the r - ϕ plane. As a result, momentum measurement in the ID depends more heavily on the precision in this plane, and not as much along the z -axis.

A 3D Render of these components is shown in Figure 20, excluding the solenoid. It should be noted that, in reality, additional detection planes would cover either end very closely and obscure any view of this sort - these are stripped away in the diagram. A cross section in $r - \phi$ is shown in Figure 21, including the specific envelopes of each ID component, excluding the IBL.

3.2.2.1 IBL

The IBL [45] is a layer of semiconductor pixel detectors immediately surrounding the beam pipe. It is the first active layer that particles produced in a collision will encounter. It is called the “B-Layer” because often B mesons will cross it before they decay, creating a distinct signature in the detector (a single track turning into 2 or more particles), and thus allowing them to be identified. In contrast, heavier unstable particles like Z bosons have smaller E/m , and so aren’t as boosted, and




Figure 20: Render of the barrel components (those within $|z| = 750\text{mm}$) of the ID as used in Run 2 at the LHC. Taken from [43].




Figure 21: Diagram showing all components of the ID projected in the r - z plane. The detector is symmetric in z , so shows only the positive z -axis. This diagram is taken from [44], and as such it describes the ID as it was in Run 1. The only change for Run 2 was the addition of the IBL between the beam pipe and the innermost Pixel layer.

thus decay much sooner, before reaching or even significantly approaching the B-Layer. While this function can be useful for someone interested in B mesons, in the case of this thesis it is used to reject background leptons from B decays. Furthermore the additional tracking layer slightly increases the lever arm used to measure track bending radius (discussed in Chapter 4), and so improves momentum resolution, in particular for electrons.

3.2.2.2 Pixel Detector

The Pixel Detector uses a similar technology to the IBL: layers of semi-conducting pixel detectors surround the beam pipe, which provide three-dimensional hit measurements along charged particle trajectories. Three concentric barrels surround the beampipe, while three coaxial disks cover the forward ($2.0 < |\eta| < 2.5$) regions. Coverage between regions is overlapping, and so any particles produced with $|\eta| < 2.5$ will cross at minimum three layers. Because of its proximity to the interaction point and high granularity, the Pixel Detector drives the precision of reconstruction of collision vertices.

3.2.2.3 SCT

The SCT uses layers of silicon microstrip detectors as its detection mechanism. In the barrel, the strips run nearly along the z direction, providing measurements in the r - ϕ plane. Each SCT layer indicated in Figure 21 is composed of 2 sets of strips, which are both at a small (and different) angle from the z -axis, providing a less precise measurement in the z coordinate as well.

3.2.2.4 TRT

The TRT uses a technology called straw detectors. The TRT barrel (endcap) is composed of 50,000 (320,000) individual straws oriented along the z (r) direction. The straws have a 4 mm diameter, and are filled with a gas mixture of Xe, CO₂, and O₂. As charged particles traverse the straws, they ionize the gas and this ionization is read out on a sense wire through the center of each tube. Between the straws is a radiator which induces passing electrons to emit transition radiation. This design feature gives electrons a unique signature in the TRT as compared to other charged particles, and is used as a quality criteria for identifying good electrons.

3.2.3 Calorimeters

The second sub-detector encountered by particles (and often the last) is the calorimeter. Located immediately outside the solenoid, it extends out to $r = 4.25$ m, and provides the most forward coverage in ATLAS, out to $|\eta| = 4.9$. It is designed to fully absorb electrons, photons, and hadrons, while allowing muons (and of course neutrinos) to pass through. Like in any calorimeter, the absorption happens by a dense material inducing the particle(s) to shower, and then measuring the total energy deposited in a shower. The ATLAS calorimeters are composed of two subsystems: the Electromagnetic Calorimeter, and the Hadronic Calorimeter. They are depicted in Figure 22.




Figure 22: Render of the ATLAS Calorimeter systems, taken from [46]

3.2.3.1 Electromagnetic Calorimeter

The Electromagnetic Calorimeter (ECAL) is designed to absorb the showers of photons and electrons. The ability of a calorimeter to absorb EM showers can be characterized by its depth in radiation lengths, X_0 . An X_0 is the length over which a high-energy electron loses an e -fold of its energy to bremsstrahlung radiation, and $7/9$ the length that pair-produced photon intensity falls by an e -fold. The ECAL has

a depth of $22X_0$ in the barrel, and $24X_0$ in the endcap. Both barrel and end-cap components are lead-liquid argon (LAr) calorimeters, which use alternating layers of lead as the absorbing material, and LAr as the detection medium to measure the showers. The geometry of the calorimeter is a nested accordion, which provides total symmetry in ϕ with no gaps in coverage.

3.2.3.2 Hadronic Calorimeter

Hadronic jets have a much deeper profile than electromagnetic showers, so as a result require a thicker calorimeter. The characteristic length for hadronic showers is the Nuclear Interaction Length, λ_I , the mean distance travelled before a hadron undergoes an inelastic collision, equivalently the distance over which particle intensity falls by an e -fold. The Hadronic Calorimeter (HCAL) has a depth of $9.7\lambda_I$ in the barrel, and $10\lambda_I$ in the most forward region. The HCAL is divided into three main subsystems: the Tile Calorimeter, the Hadronic End-Cap Calorimeter (HEC), and the LAr Forward Calorimeter (FCal) [29].

The Tile Calorimeter covers the outermost extent of the calorimeter system in the r direction, from $r = 2.28$ m to 4.25 m across its full length to $|z| = 6.0$ m (equivalently up to $|\eta| = 1.7$). It is a sampling calorimeter, using scintillating tiles as the detection medium and steel as the absorber.

The HEC bridges the space between the Tile Calorimeter and the FCAL, providing coverage in both endcaps from $1.5 < |\eta| < 3.2$. Like the ECAL, the HEC uses LAr as the detection medium, but instead uses thick copper plates as the absorber.

The FCal is the most forward main component of ATLAS, located from $3.1 < |\eta| < 4.9$ (note the overlap between the Tile Calorimeter, HEC, and FCal is deliberate to ensure gap-free coverage). Unlike the other hadronic calorimeters, the FCal is not overlapped by the ECAL, and so it is also responsible for absorbing EM showers. The FCal also uses LAr as the detection medium, and uses copper (first layer - for EM showers) and tungsten (subsequent layers - for hadronic jets) as the absorber. Because of the enormous particle flux at these values of pseudorapidity, the FCal uses a matrix of parallel tubes filled with the LAr to achieve a high granularity and prevent ion buildup that would otherwise occur in larger gas volumes.

3.2.4 Muon Spectrometer

The outermost set of detectors in ATLAS comprise the Muon Spectrometer (MS) [47]. The MS is a tracking detector immersed in an extensive magnetic field produced by

a series of toroid coils. The combined barrel and endcap toroids produce a complex magnetic field oriented roughly in the $\hat{\phi}$ direction, thus causing particle trajectories to bend in the r - z plane. The MS is primarily designed to measure muons, though some hadronic punch-through, where a jet is too energetic to be stopped by the calorimeters, does occur. The MS occupies the space beyond the calorimeter system, out to $r = 11$ m and $|z| = 22$ m. Most of that physical volume is empty however - the MS has a very long lever arm for measuring bending radius.

The MS is composed of four detector technologies. Monitored drift tubes (MDTs) and cathode strip chambers (CSCs) provide precision measurements in the bending plane, while resistive plate chambers (RPCs) and thin gap chambers (TGCs) are used for fast triggering (see more in Section 4.3) and measurement in the orthogonal coordinate.

The geometry of the MS is designed to provide at least three measurements of any track passing through it. Three concentric barrels at increasing radii surround the calorimeters. They are called respectively Inner, Middle, and Outer stations, and each contains overlapping chambers, called Small and Large according to their size in ϕ , to provide unbroken coverage up to $|\eta| < 1.05$. In Figure 24, which shows only the location of the Large chambers, they are labelled as BIL (meaning Barrel-Inner-Large), BML, BOL. Each station is composed of several layers of MDTs, while the Middle (Outer) station also contains two (one) layers of RPCs for triggering. In the forward region, three coaxial disks, called Inner, Middle, Outer, also composed of Small and Large components, provide coverage up to $|\eta| < 2.7$. Precision measurement is made by MDTs at all three stations, except for the Inner station and the region $2.0 < |\eta| < 2.7$, where the measurement is made by CSCs. Three layers of TGCs in the Middle station provide triggering, while the TGCs in the Inner station (also called the 'Small Wheel') are used only as a track measurement (but not triggering). In Figure 24, these chambers are labelled as EIL (ie Endcap-Inner-Large), EML, EOL. Though the combined Barrel and Endcap chambers provide near-complete coverage, some additional chambers are necessary to provide coverage in certain gaps caused by structural, service, or instrumentation requirements.




Figure 23: Render of ATLAS Muon Spectrometer showing various detector components. The disk closest to the front standing on its own is one of the Outer Wheels (EO*), while the equal-sized disk tightly capping the rest of the detector is the Large Wheel (EM*). Taken from [48]




Figure 24: Diagram of MS in r - $|z|$ plane, taken from [49]. Note that for each set of chambers labelled here ending in an "L" (for Large), there is a complementary "S" (for Small) chamber, at a slightly increased r (z) in the barrel (endcap). The Small and Large chambers have partial overlap, as can be seen in Figure 23, to provide nearly hermetic coverage for $|\eta| < 2.7$.

3.3 New Small Wheel

Over the course of the LHC run program, the instantaneous luminosity of collisions is expected to exceed the design value of $1 \times 10^{34} \text{ cm}^{-2} \text{ s}^{-2}$ by a factor of five [50]. The resultant increase in particle flux through the detector presents two key challenges for ATLAS: a degradation in tracking performance as a result of the increased activity in the detector, as well as an unsustainable Level-1 trigger rate (thus requiring the trigger thresholds to be raised). These problems are more present in the end-cap region ($|\eta| > 1.0$) of the MS which receives the largest particle flux.

Resolving these problems is critical for an effective physics program. Precision momentum measurement in the MS, as well as maintaining a high-efficiency trigger, particularly with low momentum threshold, is critical for any analyses using final-state muons. The solution undertaken by ATLAS is to replace the existing Small Wheel in ATLAS, with the (aptly named) New Small Wheel (NSW).

The NSW (actually two wheels - one for each end-cap) is composed of MicroMegas (MM) and small-strip thin gap chambers (sTGC). The MM will provide precision tracking, while the sTGCs provide fast triggering. The detectors are grouped into multi-layers 4 planes thick, and ordered as sTGC-MM-MM-sTGC. This configuration adds space between the two sTGC multi-layers, which improves angular resolution in triggering. The remainder of this section will discuss the sTGCs, as Carleton University is actively involved in their construction. The MM are certainly deserving of equal consideration, but are beyond the scope of this thesis. The location of the NSW, and an illustration of its importance in triggering is shown in Figure 25.

3.3.1 sTGC

The sTGC detection mechanism is the same as that of a normal TGC. A plane of parallel wires is held at high voltage inside a thin gas volume, which is sandwiched between two cathodes made up of a resistive graphite epoxy backed on one side by large rectangular copper pads, and the other by copper strips orthogonal to the wires. this design is pictured in Figure 26a. When a charged particle passes through the plane of the detector, it ionizes the gas. This ionization deposits charge on the wires, strips, and pads, allowing a measurement of the particle's trajectory in the x - y plane of the detector.

The ATLAS environment is highly demanding, where sTGCs can expect a particle flux of $20 \text{ kHz/cm}^2 \text{ s}$. The wires are held at 2900 V to ensure a small drift time,




Figure 25: Diagram showing location of the NSW inside one quadrant of the detector.

The importance of the NSW for triggering is also depicted: Using the Big Wheel alone, three tracks are observed, A, B, and C. When combining the information from the NSW, tracks B and C can be rejected because they did not originate from near the collision point.

accommodating these high rates. The graphite coating in the existing Small Wheel TGCs is too resistive to allow such operation. Furthermore, the strip spacing in the current TGC is inadequate to provide sufficient granularity and angular resolution for the high luminosity environment of Run 3 and beyond at the LHC. The sTGC uses a strip spacing of 3.2 mm, which will eventually provide a hit resolution of 100 μm in the r coordinate (recall the bending plane of the toroids is the r - z). This high resolution measurement will be used for triggering as well as a complementary measurement to the MM in the offline reconstruction.

Each gas volume is called a gap, or singlet. During assembly, two singlets are paired into a doublet, and then two doublets into a quadruplet. An existing doublet built at Carleton is pictured in Figure 26b. All quadruplets are trapezoidal in shape, with the strips oriented along the parallel sides. Several size variants of the quadruplets are needed, so that together with the MM detectors, they can be assembled into wedges (Figure 26c), and these wedges into wheels (Figure 26d).

The operating principle of the sTGCs as a fast trigger detector is quite simple. Inside one quadruplet, a 3-of-4 coincidence in the pads is used to identify tracks pointing approximately toward the interaction point. The pads are used to select nearby strips so only a small number of strips need to be read out. These strips provide a track measurement in the bending plane to be used for triggering. If the trigger fires, then all pads, strips, and wires can be read out for optimum track resolution in the offline reconstruction. The nature of the algorithms to combine this




Figure 26: Different stages of assembly of sTGC detectors into the final NSW. Figures 26a, 26b, 26d are taken from [50], while Figure 26b is a photo of an actual doublet assembled at Carleton.

information into coordinates for tracking is still under development.

3.3.1.1 sTGC alignment

An important aspect of sTGC assembly is *alignment*. High precision measurement within a detector component is worthless if you do not know where that component is. To do this, an optical locator device is placed on the exposed side of each quadruplet, which locates the quadruplet with respect to the ATLAS global coordinate system. Still unknown however are the relative locations of each singlet, and thus their strips, within the quadruplet.

The relative position of each singlet within the quadruplet is fixed (literally - by glue) at assembly time, in the assembly lab at Carleton. Any movement after this point would be the result of catastrophic failure (components breaking or warping due to mishandling) that would result in a chamber being discarded. Measurement of the alignment within a quad serves a dual purpose: it can be used for quality assurance and control (QA/QC) during assembling, and it can be used in offline reconstruction to correct the location of hits on a gap-by-gap basis. When used for QA/QC, the design specification requires that strips be aligned to their nominal location within 100 μm . Any misalignment larger than this requires special note and may result in a quadruplet being rejected.

The process of assembling boards with 100 μm alignment across dimensions larger than 1 m is non-trivial. During gluing of a doublet (quadruplet), one gap (doublet) is held in alignment with the other using an alignment jig. The jig is screwed into a flat granite table, assuring that its two precision pins sit vertically, as shown in Figure 27a. Brass inserts built into one of the angled sides on each strip board, shown in Figure 27b, are then pulled flush against the vertical pins, holding the boards in an aligned position during gluing. This whole system is held flat against the granite table under a plastic cover with light vacuum of 0.95 atm to maintain planarity over the large area of the detector.

By design this alignment system should hold the singlets tightly to their nominal respective alignment, but there are a number of reasons the system may fail to do so

- The alignment jig pins may be slightly out of parallel. Because of the roughly 2:1 proportions of the sTGCs, and the position of the alignment features on the short side of the boards, a small misalignment at the location of the pins could be twice as large on the opposite side.




Figure 27: During gluing of a doublet (quadruplet) sTGC singlets (doublets) are aligned using built-in brass inserts held firmly against the vertical pins of the alignment jig. Photos taken from Carleton assembly lab.

- Any small mote of dust between the pin and the insert may hold the board out of alignment, despite appearing flush to the pin. This is mitigated by checking for electrical continuity between the pin and the insert.
- The alignment may drift while glue is curing, a process which takes several hours. Because the entire assembly is hidden under a plastic bag (used to assure planarity), it is not possible to visually monitor the alignment during this time.

These factors all lead to a need to measure the relative alignment of the gaps within doublets and quadruplets, which the author developed. The method can be used while glue is wet and a problematic alignment can be addressed, or after glue is dry as a method of quality control.

A microscope assembly, pictured in Figure 28 was designed to sit on the granite table facing the exposed ends of the strips in the doublets or quadruplets (the other side is obscured by the alignment jig). A 45° reflecting prism allows a view very low to the surface of the table - an important feature since the bottom set of strips in a doublet is less than 5 mm from the surface of the table. This microscope assembly is used to record an image of the two sets of strips to measure. An example of such an image is Figure 33a.

This image is then opened using the alignment code written by the author. The user (assembly technician) needs only to mark off the corners of the strips in the image, separately for the top gap and the bottom gap. Typically an image can contain several (4, 5, 6, ...) pairs of strips. Once this user input is complete, the



Figure 28: Microscope assembly used for measuring alignment of sTGC doublets and quadruplets.

alignment algorithm takes over to compute the alignment.

An idealization of the alignment algorithm for a small number of strips is shown in Figure 29. A description of the algorithm is as follows:




Figure 29: Idealization of algorithm for sTGC alignment measurement. The same algorithm is applied on several pairs of adjacent strips, and the measured misalignments are averaged. Sizes not to scale.




Figure 30: Diagram showing calibration of pixel scale in each image. This process is performed separately and automatically on every image, the numbers printed here serve only as an example.

1. Find the image scale: The algorithm makes no assumption about the scale of the image, or the perspective of the camera relative to the strips. The scale is calibrated off the input image and user input (marked strip edges), using the known strip pitch of 3.2 mm. By finding the number of strips selected (based on the number of marked points), a conversion is made between pixel coordinates and real coordinates. This process is depicted in Figure 30.
2. Compute the strip-centers and centers-between-strips: Any pair of adjacent gaps in a doublet or quadruplet have their strips alternatively aligned, that is the strip-centers in one are designed to align with the centers-between-strips in the other (and vice-versa). Collectively these are just called the “centers”, and so henceforth this detail of alternatively aligned gaps is implied.
3. Compute the “flat” direction: A linear χ^2 fit is performed separately on the top set of centers and the bottom set of centers. Then the slopes of these two vectors is averaged to find the flat direction. This process is depicted in Figure 31.
4. Compute the perpendiculars: For each center in the bottom gap, a line intersecting this point and perpendicular to the flat direction is computed.
5. Compute the misalignments: Find the distance between each perpendicular and its corresponding center in the top gap, in the flat direction.
6. Compute the final misalignment: A final value and its error is computed by averaging the values from the previous step and finding their standard deviation, respectively.

A reasonable question that may arise, because of the trapezoidal shape of the sTGC chambers, is whether the microscope should be aligned to the cathode board




Figure 31: Diagram showing calibration of flat direction in each image. This process is performed separately and automatically on every image. Minor perspective effects, where the two rows of strips are not parallel, are mitigated by averaging the top and bottom slopes.





Figure 32: The two possible ways to orient the microscope while measuring misalignment. As explained in the text, configuration B is not possible because of the shallow depth of field of the microscope. Configuration A is the one used throughout the text here.


face, or look directly down the strips. These two configurations are depicted in Figure 32, labelled as A and B, respectively. In fact, the configuration B is ruled out because of the shallow depth of field of the microscope: in this configuration it is impossible to have more than one strip (or even one complete strip) in focus simultaneously. This leaves as the only option to face the edge of the cathode board, configuration A in the diagram. In this case, a further consideration is whether the self-calibration of the pixel scale is correct, given the observed strip pitch will be affected by the trapezoid angle θ , such that when facing the cathode board edge, $p' = p / \cos \theta$ (where p represents the true strip pitch, and p' represents the observed pitch). Fortunately, any misalignments of the boards perpendicular to the strip direction will *also* be scaled by this same amount, so when computing the misalignment the self-calibration nulls the effect of this term. The corollary is that the alignment algorithm is robust to *any* facing angle, it does not require the microscope to be precisely oriented to produce accurate alignment measurements.

The actual code itself produces a visualization of its work-flow, an example of which is shown in Figure 33b; the final alignment measured is printed on the image.

The measurement uncertainty arises from uniformity in strip etching, microscope focus, image resolution, and user error in marking strip corners. Each of these sources affects the value of the individual misalignments measured (step 5), and so the standard deviation of these values captures all sources of uncertainty together. The error as computed from the the standard deviation is confirmed by repeated measurement of the same set of strips producing values all within this error. The overall precision of this technique is around $25\text{ }\mu\text{m}$, significantly better than the design specification for the sTGCs of $100\text{ }\mu\text{m}$.



(a) Image from microscope



(b) Image after processing

Figure 33: Example of image from an early doublet built at Carleton showing significant misalignment. Shown at the image directly from the microscope (a), and the image after processing (b). Both images crop off useless portions at the top and bottom.

In order to describe the overall alignment of one quadruplet, multiple measurements are needed. Measuring the gap-to-gap misalignment near each of the four corners is sufficient to give a complete picture of the relative alignment of these two gaps, under the assumption that the relative alignments can be interpolated between these measurements. In the quadruplet, the relative alignment of each gap should be with respect to one single reference gap which will eventually have the precision optical alignment device to tie it into the ATLAS global coordinate system. Call these final misalignments $a_{12}^i, a_{13}^i, a_{14}^i$, where a_{xy}^i represents the alignment of gap y with respect to gap x , as measured at point i .

Because of geometrical constraints, the alignment microscope is only able to capture a field of view approximately 25 mm tall, limiting the measurement to adjacent strip boards inside a quadruplet. This necessitates calculating the desired misalignments from those that can actually be measured. Figure 34 indicates which values are calculated and measured, as well as the locations of the measurements. The final misalignments are calculated (except a_{12}^i , which requires no calculation) as follows

$$\begin{aligned} a_{12}^i &= a_{12}^i, \\ a_{13}^i &= a_{12}^i + a_{23}^i, \\ a_{14}^i &= a_{12}^i + a_{23}^i + a_{34}^i. \end{aligned} \tag{33}$$




Figure 34: Diagram indicating the four corners used for alignment measurements. At each corner, three measurements are made as indicated, and the remaining relative alignments are calculated from them.

Assuring and measuring the alignment of stGCs at assembly time is one of many details to be understood during their construction and testing. It is a small detail, but it can have resounding effects. Eventually when the NSW are installed in ATLAS,

precision alignment (or misalignment that can be corrected for) assures precision hit resolution. Hit resolution aids in efficient triggering and reconstruction of muons, discussed in the next chapter. Precision reconstruction in turn improves resolution of measurements such as the Higgs boson mass in final states containing muons. Though these particular alignment measurements in the sTGCs have no bearing on the analyses discussed in the remaining chapters of this thesis, they will have a consequence for analyses in 2019 and beyond.

Chapter 4

Reconstruction

Data from the ATLAS detector arrives in the form of hits in tracking detectors and energy deposits in various calorimeter cells. The interpretation of this information into useful “physics objects” to be used in analyses is called *reconstruction*, and is the subject of this chapter.

This chapter begins by differentiating simulated data and real data in Section 4.1, as simulated samples are used throughout the proceeding chapters. Next, Section 4.2 delves into the reconstruction, explaining important quantities we measure, and how they can be used to understand what happened in a collision. The sections that follow then describe the specific algorithms used to reconstruct the physics objects used in this thesis. More emphasis is placed on those prominent later in this thesis, and in particular to the author’s own contributions in Section 4.4.1 and Section 4.5.1.

4.1 Data and Simulation

Simulated data are broadly used in ATLAS as a tool to design components of the detector, develop software, and prepare physics analyses [51]. Using simulated data permits many things not possible in real data¹. The subset of these which are useful to this thesis are :

- Rare processes, such as Higgs events, can be simulated with far more statistics than can be collected in real data. Simulating these processes means *we know what to look for* in real data. The simulated dataset of a specific process is called a simulated *sample*.

¹The phrase “real data” is used in this section for emphasis, though most of this thesis will just use “data”.

- Processes can be simulated with different parameters of nature, e.g. values of m_H , which can be in turn used to measure the parameters in real data.
- In simulated data, the true particle kinematics are known, but not in real data. With this information in hand, we can compare reconstruction algorithms which attempt to measure these kinematics.

All simulated samples used in this thesis are generated by the Monte Carlo (MC) method [52]. Different *generators* can be used to simulate different physical processes. The plethora of generated particles, with their respective kinematics, are then passed through the digital model of the ATLAS detector using the **GEANT4** framework [53] to simulate their interaction with the detector. The detector response is then digitized, and the regular ATLAS reconstruction algorithms can be ran on these simulated data [51]. Regardless of how many events are simulated, the generated events can be re-weighted to the SM cross sections (or some other cross section, if a non-SM process is being simulated) and some integrated luminosity, so that the large number of simulated events will have a total weight equal to the number of events expected from that process in real data (e.g. 1,000,000 simulated events may be given a total weight of 8.53). Each of the simulated samples are further re-weighted in various ways to correct the simulations to better reflect real data, since neither the generators nor the model of ATLAS are perfect.

The details of the MC generators - why one is used for some samples but not others for instance - is beyond the scope of this thesis. They are included below for completeness only, and the reader is invited to consult the references inline.

The $H \rightarrow ZZ^{(*)} \rightarrow 4\ell$ signal samples are predominantly generated using the **POWHEG** MC event generator. **POWHEG** calculates ggH [54] and VBFH [55] production at NNLO, and a separate weighting is applied to make the Higgs p_T spectrum match the latest theoretical calculation as prescribed by [56]. This generator is interfaced to **PYTHIA8** [57] for simulation of showering and particle hadronization, and to **EVTGEN** [58] for the simulation of B-hadron decays. Events produced by VH production are generated using **PYTHIA8**, and ttH and bbH production are generated using **MadGraph5_aMC@NLO** [59], interfaced to **Herwig++** [60] and **PYTHIA8**, respectively. The cross section for ggF is calculated to N3LO QCD + NLO EW. For VBF and VH it is calculated to NNLO QCD + NLO EW. For ttH it is calculated to NLO QCD + NLO EW. For bbH it is calculated to NNLO or NLO, under different flavour schemes. The values of the cross sections used for $m_H = 125$ GeV were reported in Section 2.3.2, and sourced from [21].

Using the above generators, 500,000 ggF events, 300,000 VBF events, and 200,000 ttH and bbH events, where $H \rightarrow 4\ell (= e, \mu)$ were generated at m_H values of 120,122,124,125,126,128,130 GeV. For VH, 200,000 ZH and WH events (each), where $H \rightarrow 4\ell (= e, \mu, \tau)$ were generated at m_H values of 124,125,126 GeV.

The sources of background entering this thesis are described later in Chapter 6. The dominant source is from $q\bar{q} \rightarrow ZZ^{(*)}$, for which samples are generated using POWHEG interfaced to PYTHIA8 and EVTGEN as with the signal. For this background, 10,000,000 events with $4\ell = e, \mu$ were generated in the mass window 100-150 GeV, and 4,000,000 (not statistically independent) events were generated with no mass window cut. Background samples from $gg \rightarrow ZZ^{(*)}$ are generated in the same way, with 1,000,000 total events generated.

Backgrounds from tri-boson production - ZZZ, WZZ, and WWZ - are generated using SHERPA [61]. For ZZZ, 70,000 events with at least 4 real leptons were generated for each. For WZZ, and WWZ, 50,000 and 60,000 events were generated, respectively. For background from $t\bar{t}Z$, SHERPA was used to generate 1,500,000 events (this sample was used by other analysis groups - only a small fraction of these events decay via $t\bar{t} \rightarrow WbWb(WW \rightarrow \ell\nu\ell\nu)$ in order to be relevant).

The background from Z+Jets is generated using SHERPA. Many samples were generated separately for Z+1,2,3,4Jets, filtered into Z+b-jets, Z+c-jets, and Z+light jets according to the true origin of the quark-initiated jets. Alternative samples used for cross checks were generated using POWHEG and MADGRAPH, both interfaced to PYTHIA8 and EVTGEN. As Z+Jets samples are of broad use within ATLAS, huge numbers (10s of millions) of events were generated overall, though the vast majority do not enter the analysis in this thesis. Events from $t\bar{t}$ were generated using POWHEG with PYTHIA6 [62] for parton shower and hadronization, with PHOTOS [63] for QED processes, and with TAUOLA [64] for simulation of τ decays. Finally, WZ background events were generated using POWHEG interfaced to PYTHIA8 and EVTGEN.

A few other simulated samples are also used only in the current chapter for testing certain aspects of the reconstruction, and will be mentioned as needed.

4.2 Kinematic formulae

When a particle is produced in a collision (in data or simulation), it has some real kinematics, energy and momentum, that we call the *truth*, and that we attempt to measure through reconstruction. How well the *reco* matches the *truth* is the ultimate measure of how well reconstruction is performing, and is something that can be

quantified using simulated data, where the truth is known. The difference between reco and truth comes back to play a major role in Chapter 7. In this section, the formulae apply equally to either reco or truth quantities.

We can describe any particle's kinematics by its four-momentum,

$$\begin{aligned} p &= (E, \vec{p}), \\ &= (E, p_x, p_y, p_z), \end{aligned} \tag{34}$$

where E is the energy, and $\vec{p}$ is the classical three-momentum. A particle's mass can be calculated from the four-momentum²:

$$m^2 = p \cdot p = p^2 = E^2 - |\vec{p}|^2. \tag{35}$$

Note the minus sign appearing in this expression, but not, for example, in the classical momentum dot product: $\vec{p} \cdot \vec{p} = p_x^2 + p_y^2 + p_z^2$. The value of the four-momentum always depends on the reference frame of the observer, and this minus sign appears as a consequence of the transformation properties under changes of reference frame, called *Lorentz boosts*. To change to a frame moving with speed³ β , which corresponds to a Lorentz boost factor $\gamma = 1/\sqrt{1 - \beta^2}$, the four-momentum transforms as

$$\begin{aligned} E &\rightarrow \gamma(E + \beta p_z), \\ p_x &\rightarrow p_x, \\ p_y &\rightarrow p_y, \\ p_z &\rightarrow \gamma(p_z + \beta E), \end{aligned} \tag{36}$$

In contrast, the dot products of four-momenta, in any combination, *do not depend on reference frame*; for this reason they are called Lorentz Invariants, and the particle mass (m) is more precisely called the *invariant mass*, also the *rest mass* because it is always equivalent to the value in the rest frame. The corollary is that even though collisions do not occur in the rest frame of ATLAS, for reasons explained in Section 3.2.1.1, the invariant mass of any particle will be the same as if it did.

Conservation of energy and momentum dictate that four-momentum must be conserved through time. The useful consequence for ATLAS is that when an unstable particle decays and we measure only its decay products, we can access the kinematics of the initial particle without ever directly observing it.

²Note that this is the famous Einstein mass-energy equivalence relationship $E = mc^2$, except here expressed in a moving frame where $\vec{p} \neq 0$, and using a system of units where $c = 1$

³Note that speed ranges from 0 to 1, with a value 1 meaning the speed of light

For example, in a $Z \rightarrow \mu\mu$ decay,

$$\begin{aligned} p_Z &= p_{\mu_1\mu_2} = p_{\mu_1} + p_{\mu_2}, \\ m_{\mu_1\mu_2}^2 &= (p_{\mu_1} + p_{\mu_2})^2, \\ &= p_{\mu_1}^2 + p_{\mu_2}^2 + 2p_{\mu_1} \cdot p_{\mu_2}, \\ &= m_{\mu_1}^2 + m_{\mu_2}^2 + 2(E_{\mu_1}E_{\mu_2} - \vec{p}_{\mu_1} \cdot \vec{p}_{\mu_2}). \end{aligned} \tag{37}$$

At the high energies at the LHC, most final state particles have a relatively small invariant mass, e.g. $m_\mu \ll E_\mu$, and so the terms m_μ^2 can be neglected. If we are neglecting m_μ , then (35) tells us that $|\vec{p}| = E$, and so we find

$$m_{\mu_1\mu_2}^2 = 2E_{\mu_1}E_{\mu_2}(1 - \cos \theta_{12}), \tag{38}$$

where θ_{12} is the relative angle between the two muons. These mathematics extend easily to a more complex decay, such as $H \rightarrow ZZ^{(*)} \rightarrow 4\ell$, which will be exploited in Chapter 7.

Returning to the four-momenta of individual final state particles, we can establish a few other useful formulae. The *transverse momentum* of a particle, defined by $p_T = \sqrt{p_x^2 + p_y^2}$ is a useful quantity because HS events generally produce particles transverse to the beam. Additionally, the p_T is also invariant under boosts along the z -axis, so it is a useful quantity in which to compare experimental observation with theory (for example, in a differential cross section measurement like [65]). Similarly the *transverse energy* is defined by $E_T = \sqrt{p_T^2 + m^2}$.

A more intuitive way to express a particle's four-momentum is in cylindrical coordinates:

$$p = (E, y, \phi, m), \tag{39}$$

where y is the rapidity defined in Section 3.2.1.1. For light particles (including electrons and muons) using the pseudo-rapidity, η , is usually more convenient. Some expressions can be used to relate the quantities presented so far:

$$E = E_T \cosh y, \quad p_z = E_T \sinh y, \quad p_x = p_T \cos \phi, \quad p_y = p_T \sin \phi, \tag{40}$$

or, when $m \ll E$, the expressions become (p_x, p_y are unchanged)

$$E = |\vec{p}| = p_T \cosh \eta, \quad p_z = p_T \sinh \eta. \tag{41}$$

As explained in Section 3.2.1.1, separations in y are invariant under boosts along

the beamline. Using this, we can define a boost-invariant angular separation between two final state particles as

$$\Delta R = \sqrt{\Delta\phi^2 + \Delta\eta^2}. \quad (42)$$

Finally, some useful relations involving the speed β , the Lorentz boost factor γ , and the lifetime τ , of particles are

$$\gamma = E/m, \quad \beta = |\vec{p}|/E = \sqrt{1 - 1/\gamma^2}, \quad \tau = \gamma \cdot \tau_{rest}. \quad (43)$$

These expressions are useful in understanding the behaviour of particles: how fast are they travelling? will they decay before we observe them? will they pass through the whole detector before decaying? These questions can be answered quickly for a few example particles, shown in Table 5. Note that these simple expressions can be used to derive the canonical description of the LHC as “protons colliding at 99.9999988% light speed”.

Table 5: A few example particles, with their Lorentz boost (γ) and speed (β) calculated. Also calculated is the lifetime in the ATLAS detector frame, and the travel distance before decay, $c\beta\tau$. All particles use masses and decay widths from [16], except the Higgs boson for which a SM Higgs of $m_H = 125 \text{ GeV}$ is used.

Particle and momentum	γ	β	τ	$c\beta\tau$
50 GeV muon	473	0.9999978	1 ms	312 km
20 GeV B^+	3.91	0.967	6 ps	1.93 mm
20 GeV Z boson	1.02	0.214	3×10^{-7} as	79.9 nm
20 GeV Higgs boson	1.01	0.158	2×10^{-4} as	45.6 μm
6.5 TeV proton	6600	0.999999988	N.A.	N.A.

Though the above kinematic description is most useful for physics analyses, these are not always the quantities directly measured during reconstruction. A general outline to keep in mind is that calorimeters directly measure energy E , and coarse directional information θ, ϕ , while tracking detectors directly measure the parameters that define the particle trajectory: q/p (charge-to-momentum ratio), θ, ϕ, d_0 (transverse impact parameter), z_0 (longitudinal impact parameter). These measured quantities are used to derive the four-momenta discussed above. The detector signature is also used to identify the species of particle, and so its known mass (e.g. for electrons $m_e = 511 \text{ keV}$) can be used to construct the four-momentum.

4.3 Trigger

The first stage of ATLAS reconstruction is the Trigger. The Trigger decides which events to keep for further analysis, and which events to discard. Because of the enormously high event rate at the LHC, it is impossible to keep all events. Events which are kept are said to *fire* the trigger. Triggering happens *online*, in real-time as collisions happen, while the remainder of the reconstruction happens *offline*, and normally happens in the following days as data is transferred to the various computing sites which comprise the Worldwide LHC Computing Grid (WLCG) [66].

The trigger proceeds in two stages, called Level-1 (L1), and high-level trigger (HLT). The L1 stage is entirely implemented in hardware; it uses relatively simple vetoes to reject a huge number of uninteresting events. The HLT is implemented in software, and uses complex computations to identify the best remaining events. The approximate rate⁴ of events passing each trigger under nominal running conditions is shown in Table 6.

Table 6: Expected average trigger rates for Run 2, from [67].

Level	rate
no trigger	1.6 GHz
L1	100 kHz
HLT	1 kHz

The software running at the trigger level exists at a balance point between providing efficient trigger (capturing all the events with exciting physics), and fast triggering (being fast enough to process all events it receives). At the HLT level, the reconstruction algorithm is only slightly pared down from the offline reconstruction as it is described in the following sections.

This triggering procedure causes a certain bias in the data: triggers introduce specific search criteria, and so the data which passes the trigger already has a certain degree of model-dependence. To protect against this, ATLAS also records and analyses so-called *minimum-bias* events [68] - events triggered entirely at random with no requirements so as to be completely model-independent.

⁴Through the course of a collision run at the LHC, the two beams are depleted (protons are collided or lost by other means), and so the trigger rate slowly falls over time. If the rate falls a certain threshold, it becomes more worthwhile to dump the remaining beams and refill, rather than continue colliding at a low rate.

4.4 Electrons and photons

Electrons and photons are described together because their reconstruction [69][70] is very closely tied.

The reconstruction starts by building clusters out of energy deposits in the ECAL. Second, ID tracks are constructed using a combination of pattern-recognition tools, and performing a helical fit of the hits on each identified track, with the bending direction of the helix determining the track's charge. The fit assumes a charged pion hypothesis as a baseline for energy loss. As the fit extrapolates the particle trajectory layer-by-layer through the ID, it allows a Gaussian uncertainty on the track parameters, which is re-estimated at each step. This technique is called a Kalman filter.

The clusters in the ECAL are matched, if possible, with ID tracks. Electrons and photons are identified in the following ways

- The cluster matches no tracks: the particle is identified as an “unconverted photon”, meaning it was produced in the collision event as a photon, and remained so until it reached the ECAL.
- The cluster matches a pair of oppositely-charged tracks: the particle is identified as a “converted photon”, meaning it was produced in the collision event as a photon, but converted to an e^+e^- pair before (or while) it crossed the ID. A converted photon can also be identified by a single track if it has no hits in the innermost Pixel layers.
- The cluster matches a single track: this particle is identified as an electron (or positron).

4.4.1 Gaussian Sum Filter

As electrons traverse the material of the ID, they radiate energy (in the form of photons) by bremsstrahlung (translated as “braking radiation”, often simply called *brem*). The fraction of energy remaining, $z = E_f/E_i$, after traversing a thickness of material $t = x/X_0$ (normalized to radiation lengths) follows a Bethe-Heitler (BH) distribution [71],

$$p_{BH}(z) = \frac{(-\ln(z))^{\frac{t}{\ln 2}-1}}{\Gamma\left(\frac{t}{\ln 2}\right)}, \quad 0 < z < 1, \quad (44)$$

with some higher order corrections⁵.

Because the BH distribution is highly non-Gaussian, the approximation used by the Kalman Filter can fail to account for the extreme tails of energy loss. When the electron radiates a large fraction of its energy into a bremsstrahlung photon (an uncommon, but not entirely rare occurrence), its trajectory will exhibit a strong kink. Because the Kalman Filter uses a Gaussian description of the stochastic energy loss processes of electrons, it fails to account for these uncommon-but-not-entirely-rare events, and so electrons can be lost or else poorly reconstructed (e.g. bad momentum estimate).

To mitigate these problems, electrons are re-fit with a dedicated procedure called the Gaussian Sum Filter (GSF) [72][73]. The GSF approximates the BH as a weighted sum of Gaussian distributions,

$$p_{GSF}(z) = \sum_{i=0}^{N_G} w_i G(z; \mu_i, \sigma_i). \quad (45)$$

The particular sum of Gaussians used was optimized in [73] by minimizing the cumulative distribution function (CDF) distance, defined by

$$D_{CDF} = \int_{-\infty}^{+\infty} |p_{BH}(z) - p_{GSF}(z)| dz, \quad (46)$$

Figure 35 compares the six-Gaussian sum $p_{GSF}(z)$ with $p_{BH}(z)$ optimized by this metric.

With this Gaussian-sum approximation in hand, the track fitting can proceed as many weighted Kalman Filters in parallel. At each step, the track parameters' uncertainties are represented by a sum of multi-dimensional Gaussians. The dimension is equivalent to the number of track parameters, in this case five⁶. As the track is extrapolated to the next measurement, it passes through some material which causes an increase in the uncertainties according to the BH distribution. By describing this distribution as a sum of N_G Gaussians, the GSF allows the complex distribution of uncertainties to be propagated through material. At each propagation step, the number of Gaussians needed to describe the uncertainties increases by a factor N_G , as each initial Gaussian uncertainty is propagated through the Gaussian sum describing

⁵At the energies we are concerned with, 7 GeV and above, the most noteworthy contribution is from multiple scattering

⁶As mentioned at the top of Section 4.2 they are $q/p, \theta, \phi, d_0, z_0$




Figure 35: Comparison of Gaussian sum approximation (solid, blue) with $N_G = 6$ Gaussians to the Bethe-Heitler distribution (dashed, red), shown for $t = 0.02$. Taken from [73], Figure 5.2.

the BH distribution. As the propagation reaches each measurement, the track parameters can be updated and the uncertainties re-estimated. This process is illustrated in Figure 36.

The astute observer will see a potential danger with this approach: the number of Gaussians required explodes exponentially, as at step k , the number of Gaussians would be $(N_G)^k$. To prevent this catastrophe, after each propagation step the GSF reduces the number of Gaussians back to a manageable number. This process is a necessary approximation: the full details of the uncertainties are most accurately encoded in some huge number of Gaussians, but if most of that same information can be encoded in a smaller number of Gaussians (call it the “Number of Reduced Components“, or N_{red}) such that the uncertainties are still accurate, then the same result can be achieved in a more computationally feasible way. To understand how this reduction is done, we must establish some notation, which is adapted from [73].

The distribution of errors, ϵ , may be described in generic form as a sum of n multi-dimensional Gaussians:

$$p(\epsilon) = \sum_{i=1}^n w_i G(\epsilon; \mu_i, V_i), \quad (47)$$




Figure 36: Illustration of the GSF error propagation for a simplified example with $N_G = 2$ Gaussians. The initial track parameters, on Module 1, have their initial uncertainties described by a single Gaussian. As the track is propagated through the Material Layer, the sum of Gaussians describing the energy loss is propagated through the track parameter uncertainties, which are thereafter described by a sum of two weighted Gaussians. When the track reaches the measurement on Module 2, the same process is repeated for each of these two Gaussians, leaving four Gaussians. As a final step (not shown), the parameters would be re-estimated according to the actual measurement on this Module. Taken from [71], Figure 4.4

where $\boldsymbol{\mu}_i$, and $\mathbf{V}_i$ are the mean and covariance matrix of the d -dimensional Gaussian i , defined by

$$G(\boldsymbol{\epsilon}; \boldsymbol{\mu}, \mathbf{V}) = \frac{1}{(2\pi)^{d/2} \sqrt{\det(\mathbf{V})}} \exp \left(-\frac{1}{2} (\boldsymbol{\epsilon} - \boldsymbol{\mu})^T \mathbf{V}^{-1} (\boldsymbol{\epsilon} - \boldsymbol{\mu}) \right). \quad (48)$$

If the sum is too large, that is if $n > N_{red}$, then the goal is to reduce it to a manageable number. To do this, we identify “similar” Gaussians: if the information from two Gaussians is redundant or nearly redundant, it is obvious those should be the first to go. In practice, a scalar distance is computed between each pair of Gaussians in the sum, called the Kullback-Leibler distance [73][74]. For Gaussians $i = a, b$, this distance is

$$D_{KL}(a, b) = \text{tr} [(\mathbf{V}_a - \mathbf{V}_b) (\mathbf{V}_b^{-1} - \mathbf{V}_a^{-1})] + \left((\boldsymbol{\mu}_a - \boldsymbol{\mu}_b)^T (\mathbf{V}_a^{-1} + \mathbf{V}_b^{-1}) (\boldsymbol{\mu}_a - \boldsymbol{\mu}_b) \right). \quad (49)$$

When the pair of Gaussians with the smallest D_{KL} are found in the sum of (47), they are replaced by a single Gaussian with parameters

$$\begin{aligned} w &= w_a + w_b, \\ \boldsymbol{\mu} &= \frac{1}{w_a + w_b} (w_a \boldsymbol{\mu}_a + w_b \boldsymbol{\mu}_b), \\ \mathbf{V} &= \frac{1}{w_a + w_b} (w_a \mathbf{V}_a + w_b \mathbf{V}_b) + \frac{w_a w_b}{(w_a + w_b)^2} (\boldsymbol{\mu}_a - \boldsymbol{\mu}_b) (\boldsymbol{\mu}_a - \boldsymbol{\mu}_b)^T. \end{aligned} \quad (50)$$

This merged Gaussian preserves the total weight, the mean (1st moment), and the covariance (2nd moment) of their combined distribution. The terms involved are mostly obvious (e.g. the combined mean is the weighted sum of the individual means), except for the final term of $\mathbf{V}$, which is the contribution to the combined variance introduced by the difference between means of the two Gaussians. Once this combined Gaussian replaces Gaussians a, b in the sum, new D_{KL} are computed, and the process repeats iteratively until $n \leq N_{red}$.

Tuning of N_{red}

One study undertaken by the author was to optimize the configuration in the GSF of the parameter N_{red} . Too small a value results in poor description of the energy loss and thus poor momentum resolution, while too large a value results in a significant CPU time penalty. The initial value of $N_{red} = 12$ was established in 2006 [73], and this study sought to verify this value was still optimal, considering the volume of

changes to other aspects of the ID track fitting code since that time.

To evaluate the quality of track reconstruction, a sample of 10,000 simulated electrons with $E_T = 25$ GeV was used. These electrons were reconstructed by the GSF under different configurations of N_{red} . The relative residual distributions, defined by

$$1 - \frac{(q/p)_{reco}}{(q/p)_{truth}}, \quad (51)$$

are plotted for different N_{red} in Figure 37. These distributions are used to evaluate the performance according to three metrics:

- The σ of a Gaussian fit to the core of the distribution
- The Full-Width-at-Half-Max (FWHM) of the distribution
- The smallest width to contain 50% of the integral of the distribution




Figure 37: Distribution of (q/p) pulls for a sample of 10,000 electrons with $E_T = 25$ GeV reconstructed with the GSF under different configurations of the parameter N_{red} . The author learned how to produce much nicer plots since this study.

Figure 38 shows the evaluation time of the GSF algorithm, and the performance as a function of N_{red} , in steps of 2, 5, or 10 as N_{red} increases. The evaluation

time behaves approximately linearly with N_{red} , and will be discussed in a subsequent section. The resolution, as quantified by the three metrics above, exhibits a plateau of diminishing gains for higher N_{red} . The “smallest width” metric appears to be the most robust, providing a mostly smooth dependence as one expects, and the most distinct representation of the plateau. From this study, the value of $N_{red} = 12$ was confirmed to be optimal, and is the current value used by the GSF algorithm for all ATLAS electron reconstruction.




Figure 38: Tuning of GSF parameter N_{red} . Left shows the evaluation time of the fitter normalized to the $N_{red} = 12$ time. Right shows three metrics of performance as a function of N_{red} , evaluated on the histograms of Figure 37. All show a consistent picture of diminishing returns from an increase of N_{red} after approximately $N_{red} = 12$.

CPU optimizations

Another study by the author sought to improve CPU performance of the GSF. While CPU gains themselves don’t directly improve physics analysis, gaining (e.g.) 0.1 ms in one part of the code can leave “free time” to perform more expensive calculations to potentially improve resolution in another part.

The strategy used was to implement auto-vectorizable [75] code where possible within the GSF algorithm. Auto-vectorization is a compiler directive which allows a loop to be split onto multiple CPU cores and so allow iterations to be done in parallel in cases where one iteration of the loop does not affect another. A simple

code example where this could occur is in array multiplication (this example taken from [75]):

```
for (int i=0; i<12; ++i){
    A[i] = B[i] * C[i];
}
```

Here it is clear that each iteration i has no effect on the execution of some other iteration $j \neq i$. In this case the iterations can be split up onto N_{CPU} CPU cores and executed in parallel. Theoretically this can result in gains by a factor of as much as N_{CPU} for such sections of code. This optimization is schematically illustrated for the above loop in Figure 39.




Figure 39: Diagram showing theoretical example of auto-vectorization on a loop of size 12, and a CPU with 4 cores. The numbers in boxes represent the index of the loop.

From this example one can imagine potential opportunities for the use of auto-vectorization in the GSF code, which relies heavily on matrix manipulation, vector addition and subtraction, and in general numerical computation. The author implemented auto-vectorizable code into some such cases, and here are presented the results in terms of CPU performance. Figure 40 shows the same CPU scaling with N_{red} as seen in Figure 38, this time comparing before and after the code improvements. Note that since both lines, before and after, are normalized with respect to themselves,




Figure 40: Plot showing improved scaling of CPU time with N_{red} brought by auto-vectorization of GSF code. The blue line is prior to changes, and red line is after changes by the author. The vertical scale is normalized to $N_{red} = 12$ separately for both lines.

only the scaling can be compared, and not the absolute CPU consumption.

We can compare the CPU performance prior to changes (labelled *A*), and after the changes (labelled *B*) by fitting with a simple 2nd order polynomial⁷. The results of these two fits are

$$\begin{aligned} \left(\frac{t}{t_{12}}\right)_A &= 0.111 \left(\frac{N_{red}}{12}\right)^2 + 0.576 \left(\frac{N_{red}}{12}\right) + 0.221 \\ \left(\frac{t}{t_{12}}\right)_B &= 0.089 \left(\frac{N_{red}}{12}\right)^2 + 0.617 \left(\frac{N_{red}}{12}\right) + 0.212 \end{aligned} \quad (52)$$

where t_{12} is the CPU time for $N_{red} = 12$, separately for the two different code configurations. From this result it is clear the scaling with N_{red} is rendered more linear by the implementation of auto-vectorizable code. These improvements have made the penalty paid to increase N_{red} more affordable, should a replication of the preceding study find a higher value is more optimal to improve reconstruction resolution.

4.5 Muons

In contrast with the preceding objects which are all fully contained by the ATLAS detector, muons have much higher penetrating power (i.e. they deposit very little energy on average per unit length), and so they escape the detector entirely. Fortunately, as they do so, they deposit small amounts of energy through each detector they pass (ID, Calorimeters, then MS), which allows their trajectory to be reconstructed. Since they are charged particles, and most of the detector is immersed in strong magnetic fields from either the Solenoid or the Toroid Magnets, the bending of their trajectory allows measurement of their momentum.

Muon reconstruction initially proceeds separately in the ID and the MS [4]. In the ID, tracks are reconstructed exactly as for electrons, except the GSF re-fit is not used. ID tracks are labelled as *muon-associated* if they are matched to a calorimeter energy deposit consistent with a muon (low ionization, no shower), and if they satisfy the following requirements:

$$\begin{aligned} d_0 &\leq 50 \text{ mm}, \\ N_{pix} &\geq 1, \\ N_{SCT} &\geq 3, \\ N_{Si} &\geq 4, \end{aligned} \quad (53)$$

⁷for this fit I exclude the points below $N_{red} = 12$, because the appeal of ever lowering this parameter is very small

where $N_{pix(SCT)}$ is the number of hits in the Pixel (SCT) detector, and N_{Si} is the combined silicon hits (BLayer, Pixel, SCT).

In the MS, a separate algorithm designed around its large physical size is needed. First, in each MDT chamber and its associated TGC or RPC trigger chamber, a pattern-recognition algorithm, the Hough Transform [76], is used to identify sets of hits pointing in straight lines. These “track segments” are fit with a straight line, and required to point approximately toward the interaction region.

Full MS tracks are built by running a matching algorithm on the reconstructed track segments. When a possible match of segments is found, all the hits in the constituent segments are refit as one (curved) track. Tracks must have at least two segments, with a few geometrically-motivated exceptions. If the tracks contain hits found also on other tracks, they are removed from the track with the worse fit χ^2 . If a track passes close to hits which were not initially associated to one of its segments, and so were not included in the track fit, these hits are recovered. If the hit content of a track changes, it is refit. The collection of tracks produced by this algorithm are called MS tracks.

Muon objects to be used in physics analyses are finally built from ID tracks, calorimeter information, and MS tracks. Depending on the available information, these muons are classed into five *types* as follows, ordered roughly by performance (efficiency, purity, and momentum resolution).

- Combined (CB) Muons: built from a combined fit of matching ID and MS tracks. Tracks are matched using an outside-in pairing first, followed by inside-out pairing. Because ID tracks vastly outnumber MS tracks, the former accounts for most of the CB muons reconstructed. The hits from both tracks are refit together to create the combined track.
- Segment-tagged (ST) Muons: in some cases an ID track can only be paired with MS track segments in a particular MDT or CSC chamber, rather than full tracks. Since these tracks have lower performance than CB tracks, they are marked as a separate type.
- Silicon Associated Forward (SiAF) Muons: discussed separately in Section 4.5.1.
- Extrapolated (ME) Muons: In the region $2.5 < |\eta| < 2.7$, there is coverage by the MS but not the ID. In this region, MS tracks are refit with a loose constraint to have originated from the interaction point.

- Calorimeter-tagged (CT) Muons: In the region $|\eta| < 0.1$, the MS has very poor coverage due to service and instrumentation constraints (service and cabling for the ID and Calorimeters runs through this region). In this region muons are identified based only on the muon-associated ID track.

4.5.1 Silicon Associated Forward Muons

Silicon Associated Forward Muons are described in their own section as the author was responsible for implementing their combined refit into the reconstruction algorithm, and for the performance studies presented in this section.

On close examination of the inset of Figure 21, one can see that there is *some* ID coverage above $|\eta| > 2.5$ coming from the Pixel detectors. SiAF Muons attempt to take advantage of this coverage to improve muon performance in the narrow band around $|\eta| = 2.5$. For a real muon which passed only through these edges of Pixel detectors and then out into the MS, inclusion of these tracklets for ID-MS track matching (always as a second choice behind full ID tracks) can provide a more complete reconstruction, and therefore more precise momentum measurement, than the MS track alone.

In the rest of the ID, ie $|\eta| < 2.5$, tracks are required to meet the criteria of (53) to ensure a high level of purity. In the forward region, a separate set of ID "tracklets" can be reconstructed, almost exclusively from Pixel hits. For these tracklets, the SCT hit requirements of (53) are lifted, and the N_{Si} requirement is loosened to $N_{Si} \geq 3$. By comparison, Figure 41 shows the ID hit distributions for CB and SiAF muons (ME muons are not included because by definition have no ID hits) in a simulated sample. The difference in Pixel hit distribution is accounted for by requirements on N_{Si} : while CB muons are allowed fewer Pixel hits in practice, because they are able to make up for it in SCT hits, SiAF muon have usually no (or few) SCT hits to make up the difference. The particular shape of the N_{SCT} distribution for SiAF muons (e.g. the lack of $N_{SCT} = 1$ muons) is a product of the detector geometry in the endcap region. The combined Pixel and SCT hits are also shown.

For the SiAF muons, the matching of ID tracklets to MS tracks is done in the same way as for CB muons⁸. First the remaining unmatched MS tracks are extrapolated inwards to find a matching ID tracklet. As a secondary approach, inside-out matching is done where the ID tracklets are extrapolated outwards to be matched to MS tracks.

⁸but with lower priority: if an MS track can be paired to create a CB muon, it will take precedence over a SiAF pairing of the same MS track




Figure 41: Number of hits in Pixel (left) and SCT (center), and combined Pixel+SCT (right) systems for the ID track or tracklet of CB and SiAF muons. Number of hits includes hits plus dead sensors the track passed through. Certain loose quality cuts matching those of the $H \rightarrow 4\ell$ event selection, discussed later in Section 5.3.2, are applied. Distributions are normalized for comparison.

The possible matches are compared using a combined fit χ^2 , and the combination producing the smallest χ^2 is taken, unless no good match is found, in which case the MS track is left unpaired, and an ME muon is created. The net effect of the inclusion of this algorithm is a migration of some otherwise-ME muons to be SiAF muons.

To quantify the benefit of SiAF muons, we can look at simulated data containing $Z \rightarrow \mu\mu$ decays. Figure 42 shows the breakdown of muon types for muons reconstructed in these events⁹. Note that the rate of SiAF muons (2%) is approximately equal to that of ME muons, indicating that roughly half of the otherwise-ME muons were able to be reconstructed as SiAF muons. The ST and CT muons are reconstructed quite independently of this discussion, so are not included in proceeding plots.




Figure 42: Composition of reconstructed muons from simulated $Z \rightarrow \mu\mu$ events by muon type.

It can be useful to examine directly the reconstructed kinematic distributions for the various muon types. Figure 43 shows a comparison of some kinematic distributions of CB, SiAF, and ME muons reconstructed in the simulated $Z \rightarrow \mu\mu$ events. In particular the η distribution gives the most physical insight as to where these muons arise: values of $|\eta|$ near 2.5 skim the edge of the ID, exactly where one expects SiAF

⁹note that these muons already have certain quality criteria applied. The same selection is used for the $H \rightarrow 4\ell$ event selection, and discussed in Section 5.3.2




Figure 43: Kinematic distributions for muons reconstructed in simulated $Z \rightarrow \mu\mu$ events, separated by muon type. Note that the vertical scale and bin size of each histogram is arbitrary - the histograms are independently scaled and binned for convenient comparison.

muons. At higher values of $|\eta|$, approaching 2.7, the rate for SiAF muons falls while that of ME muons rises¹⁰.

We also expect improved reconstruction performance arising from the use of ID tracklets in a combined fit. Using simulated data permits the comparison of reconstructed kinematics to the truth kinematics, which we use to measure this benefit as explained in Section 4.2. Figure 44 shows the reconstruction performance for several kinematic variables - a narrower distribution indicates reconstructed values are closer (on average) to their truth values, and so reconstruction is performing better. For a balanced comparison, these plots require $|\eta| > 2.4$, an approximate pseudorapidity region¹¹ where all three muon types are reconstructed. This cut creates an apparent

¹⁰This plot extends to $|\eta| = 3$, even though muons used in analysis are required to have $|\eta| < 2.7$

¹¹More central pseudorapidities, for instance $|\eta| < 2.4$, have less pileup, and less material for muons




Figure 44: Resolution distributions for muons reconstructed in simulated $Z \rightarrow \mu\mu$ events, separated by muon type, plotted for $|\eta| > 2.4$. Note that the vertical scale and bin size of each histogram is arbitrary - they are normalized and binned for convenient comparison.

bias in the η estimation for CB muons which should be ignored (i.e. do not use this shift to draw conclusions about the muon types). From these figures we see the expected result: CB muons are always the best, but SiAF muons demonstrate a marked improvement over ME muons. To quantify what we can see by inspection, Table 7 compares the standard deviation of each of these distributions. SiAF muons are found to improve angular resolutions as compared to ME muons by a factor of 4.5 or better, and momentum resolution by a factor of 1.2.

This observed improvement in resolution arises because of the addition of Silicon hits to the final muon track. Delving deeper, we can compute the above resolutions to traverse (thus less scattering), making reconstruction “easier”, and giving an unfair advantage to CB muons.

Table 7: Performance of reconstruction for different muon types in kinematic variables η , ϕ , p_T , for $|\eta| > 2.4$. The table values are the sample standard deviations (denoted as σ) of the distributions shown in Figure 44, such that smaller numbers indicate better resolution. Note that while η and ϕ are shown as absolute resolutions (e.g. $x - x^{truth}$), the p_T is shown as a relative resolution $((p_T - p_T^{truth})/p_T^{truth})$. Statistical uncertainties are an order of magnitude below the final presented digit, though systematic uncertainties were not evaluated.

muon type	$\sigma_\eta(\times 10^{-3})$	$\sigma_\phi(\times 10^{-3})$	σ_{p_T}
CB	0.56	0.32	3.7 %
SiAF	1.1	0.71	4.3 %
ME	5.1	3.9	5.3 %

as a function of the number of silicon hits on the ID track (compared to Table 7 which shows inclusive numbers). Figure 45 shows exactly this result. ME muons are excluded from these plots as they have no ID hits. It should be noted that CB muons will in general also have TRT hits, which brings some benefit to resolution, but especially helps in rejection of electron tracks in the ID faking muons.

From these studies we can draw some conclusions. SiAF muons represent a distinct upgrade from ME muons they would otherwise be. This improvement arises from the addition of silicon hits to the track, in particular because they are so close to the interaction point. The more silicon hits the muon uses, the better its resolution will be, which can be seen from Figure 45. In some cases, SiAF muons can be nearly as well reconstructed as CB muons. Of course, the resolution of reconstructed quantities is not the only factor that should establish the quality of a certain algorithm. Also important is the *purity*: the fraction of reconstructed muons which were actually muons and not some other source of background. Because they lack TRT hits and have fewer silicon hits, it is expected that SiAF muons will have lower purity than CB muons, though it is left for future work to determine.




Figure 45: Resolution values for different muon types as a function of N_{Si} , for $|\eta| > 2.4$. The value plotted is the standard deviation of the histograms shown in Figure 44, split into bins of N_{Si} .

4.6 Jets

In contrast to photons, electrons, or muons, which are all both SM particles and a “physics object”, jets are a purely reconstruction-based concept. They are built from energy deposits in the HCal. Jets are the manifestation of a variety of strongly interacting final-state particles. Because of the hadronic nature of p - p collisions, many jets are produced from quarks and gluons directly, and also from baryons and mesons produced in collisions.

Jets are reconstructed using the anti- k_T algorithm with a cone size of $R = 0.4$ [77]. The algorithm iteratively groups nearby calorimeter clusters into jets, starting with the clusters with highest p_T and working downwards. The cone size, roughly speaking, defines the distance in which to search for nearby clusters, as defined in (42).

Jets are used only tangentially in the analysis in this thesis where their overlap with muons and electrons is concerned, described in Section 5.5.

4.7 Missing transverse energy

Missing Transverse Energy (MET, also denoted $\cancel{E}_T$) is not directly reconstructed so much as it is calculated. When neutrinos are produced in collision, for example from a $W \rightarrow \ell \nu$ decay, they stream freely through the detector (and in all likelihood several light-years of space) without interacting. The only window by which we observe neutrinos produced in the detector is by looking for the energy we *didn't see* - what they carried away from the collision.

As a consequence of conservation of momentum, the total momentum of all incoming particles in a collision is equal, unequivocally, to the total momentum of all outgoing particles from the collision. As mentioned in Section 3.2.1.1, the longitudinal momentum of partons participating in a collision is unknown, because the protons' very precisely known momentum is shared at random according to parton distribution functions. Only the momentum vector in the transverse (r - ϕ) plane is known, and it is zero. When summing the transverse momentum of all recorded particles (measured as described in the previous sections), any remaining imbalance called Missing Transverse Momentum, whose magnitude is the MET.

$$\begin{aligned}\vec{\cancel{p}}_T &= - \sum_{\text{particles}} \vec{p}_T \\ \cancel{E}_T &= |\vec{\cancel{p}}_T|\end{aligned}\tag{54}$$

This description of MET is a simplified version, as MET only enters tangentially to the analysis presented in this thesis. For a study which relies more heavily on MET, additional subtleties, such as when to use track-based or calo-based MET, must be considered and detailed.

4.8 Reconstruction summary

Figure 46 provides a summary of how different types of long-lived particles are reconstructed within ATLAS. Note that the distinction “long-lived” is important here, because short-lived particles (whether they are occasionally short-lived, or always short-lived) will have drastically different signatures: photons can convert to a e^+e^- pair, the neutral hadron π_0 always decays immediately to a $\gamma\gamma$ pair, among a plethora of other examples.




Figure 46: Diagram showing how different types of particles are reconstructed in ATLAS. Solid lines represent tracks, dashed lines represent small energy deposits, and dotted lines represent a complete absence of any signature. Electromagnetic and hadronic showering are indicated schematically to show in which component of the detector most of the energy is deposited. Sizes of any kind are not to scale.

Chapter 5

$H \rightarrow ZZ^{(*)} \rightarrow 4\ell$ Event Selection

Among the billions of collision events recorded by ATLAS, a very small minority contain $H \rightarrow ZZ^{(*)} \rightarrow 4\ell$ decays. This chapter describes how these events are identified for further analysis. While the full event selection described was implemented by the author, the development of the selection criteria and execution of the event selection was performed by ATLAS collaborators.

Events are selected using a series of requirements, or *cuts*. The cuts are designed around a trade-off between signal efficiency and background rejection: any cut in general could be modified to accept more signal at the cost of also accepting more background (called *looser*), or be adjusted to reject more background at the cost of also rejecting more signal (called *tighter*). Another pragmatic consideration affecting the event selection criteria is to only make cuts on variables that are well-described by MC, and only in a phase space which the ATLAS detector performance is well-understood, both factors which take higher priority over optimizing signal significance. The events surviving the event selection described in this chapter are called *candidate events*, and are used in the statistical analyses described in Chapter 7, 8, and 9.

Before any data are looked at, the analysis (which includes the event selection) is developed from end to end using MC simulated samples. This is done to prevent potential bias that could arise from modifying the analysis strategy after looking at the data¹ [78].

¹A worst-case scenario would be to choose analysis cuts which “cut around the data”: a dangerous bias is introduced if one decides to make a cut of $p_T > 4.9$ GeV instead of $p_T > 5.0$ GeV, just because it accepts several new data events (or vice-versa to remove these events). The cut at 4.9 GeV would have been acceptable in general only if it was decided upon before looking at the data.

5.1 Pre-Selection

A pre-selection of events is designed to remove data which are unsuitable to be used for analysis. One reason is a detector component may malfunction or fail during a data collection run (e.g. because of a high voltage trip). An offline component can create an effective “hole” in the detector from which particles can escape during collisions, and so it is safest to discard events during which this happens. Since ATLAS data is divided into runs - periods of LHC collisions usually lasting between hours to days - and further into lumi-blocks - 1 minute blocks of data - the first cut in the event selection is the Good Runs List (GRL), a list of runs and lumi-blocks that have been approved by the Data Quality (DQ) group for use in physics analyses. Malfunctioning detector components can often be rebooted by shifters in the ATLAS control room, so only a minimal number of lumi-blocks are lost. Further removal of events flagged for detector problems is done by checking for error states recorded in the SCT, the LAr, Tile, or for data corruption.

5.2 Event-level selection

The second set of cuts are designed to remove events which show no promise of having $H \rightarrow ZZ^{(*)} \rightarrow 4\ell$ decays based on event-level quantities.

The first is to require that events have a primary vertex with at least 2 tracks. The primary vertex is defined as the vertex with the highest $\sum p_T^2$ in the event. This definition is designed to identify hard scattering collision events, i.e. events most likely to contain interesting physics such as a Higgs boson.

Events are then required to have passed one of the triggers listed in Table 8. The requirement to have passed any one (or more) of these triggers is nearly 100% efficient for signal events. Note that the triggers presented here are those used in the final period of data taking in 2016 - earlier run periods use slightly different configurations according to their availability at the time.

Table 8: Trigger list for the $H \rightarrow ZZ^{(*)} \rightarrow 4\ell$ event selection. Rows are grouped by trigger type, as labelled.

trigger name	description
single electron triggers	single electron with $E \geq 24$ GeV (for e.g. <code>e24</code>) and of medium quality (for e.g. <code>lhmedium</code>). The string L1* represents L1 trigger requirements (V=varying thresholds with η , H=hadronic core isolation).
<code>e24_lhmedium_L1EM20VH</code>	
<code>e60_lhmedium</code>	
<code>e120_lhloose</code>	
di-electron triggers	two electrons, read as above.
<code>2e12_lhloose_L12EM10VH</code>	
tri-electron triggers	three electrons, read as above.
<code>e17_lhloose_2e9_lhloose</code>	
single muon triggers	one muon with $p_T \geq 20$ GeV (for e.g. <code>mu20</code>) of loose quality (for e.g. <code>ילוose</code>), and requiring the L1 15 GeV muon trigger (from <code>L1MU15</code>) or imposing $ \eta $ requirements (<code>0eta105</code> indicating $ \eta < 0.105$)
<code>mu20_ילוose_L1MU15</code>	
<code>mu40</code>	
<code>mu60_0eta105_msonly</code>	
di-muon triggers	two muons, read as above.
<code>2mu10</code>	
<code>mu18_mu8noL1</code>	
tri-muon triggers	three muons, read as above. Sometimes only an MS-track is required (indicated by <code>msonly</code>).
<code>3mu6</code>	
<code>3mu6_msonly</code>	
<code>mu18_2mu4noL1</code>	
electron-muon triggers	combination of electrons and muons, read as above. In the L1 trigger label, I=electromagnetic isolation requirements.
<code>e17_lhloose_mu14</code>	
<code>2e12_lhloose_mu10</code>	
<code>e12_lhloose_2mu10</code>	
<code>e24_medium_L1EM20VHI_mu8noL1</code>	
<code>e7_medium_mu24</code>	

5.3 Lepton selection

For mass measurement in the $H \rightarrow ZZ^{(*)} \rightarrow 4\ell$ channel, the lepton selection is critical: the invariant mass of the four leptons, $m_{4\ell}$, will be used as the observable sensitive ultimately sensitive to m_H , Γ_H , and $\sigma_{H \rightarrow 4\ell}$. This section describes the criteria for leptons to be used in the remainder of the event selection and analysis.

Note that “lepton” here and in the following refers only to electrons and muons. Neither taus nor neutrinos are included in this analysis as both result in MET and so require significantly different analysis strategies. They are also insensitive to the Higgs mass for the same reason.

5.3.1 Electron selection

In order to have precise four-momentum reconstruction, electrons need a good ID track. This is achieved by requiring a hit on the inner-most tracking layer, the B-Layer, as well as a minimum number of hits in the Pixel and SCT sub-detectors.

True electrons are distinguished from fake electrons using a multivariate likelihood discriminant (the predecessor method used during 8 TeV collisions is described in [79]). The discriminant is built from a set of variables with separating power between real electrons and fake electrons both reconstructed by the algorithms described in Section 4.4. It includes variables related to energy leakage into the hadronic calorimeter, fraction of energy deposited in the back layer of the EM calorimeter, variables related to shower width, and track-cluster matching. PDFs are built using the distributions of these variables in simulation for true electrons and fakes. These PDFs are combined in a likelihood to compute a single score for each electron. A cut is placed on this score, tuned on MC, so that together with the ID cuts a 95% efficiency for true electrons is achieved - this working point is called *Loose*, because it achieves high efficiency at a sacrifice of some purity. Section 6.3.1 provides a detailed measurement of the contamination from fake electrons passing this event selection.

Only electrons with $p_T > 7 \text{ GeV}$ and $|\eta| < 2.47$ (where ID tracking is available) are used.

5.3.2 Muon selection

Where they are available, CB muons are preferred for their high resolution and purity; they are found roughly in $0.1 < |\eta| < 2.5$. CT and ST muons are kept in the region $|\eta| < 0.1$ where no (or partial) MS coverage is available. SiAF and ME muons are

kept in the region $2.5 < |\eta| < 2.7$ where there is no (or trailing) ID coverage; of these, SiAF muons are preferred

A working point for muons, similarly called *Loose*, is built by placing cuts on ID and MS hits, specific to each muon type (one cannot ask ID hits of an SA track, nor ask MS hits of a CT track, for example). The goal of these requirements is to reject fake muons from pion and kaon decays. A measurement of the compatibility between ID and MS tracks, called *q/p significance* and calculated using the ID and MS track momenta and resolutions, also enters into the working point. These selections are described, excepting SiAF muons, in [4].

CB, ST, SiAF, and ME muons are used for $p_T > 5$ GeV, while CT muons are used for $p_T > 15$ GeV.

Muons are required to have $d_0 < 1$ mm, and $|z_0 \sin(\theta)| < 0.5$ mm, except for SA muons. These cuts are designed to reject muons which did not originate at the primary vertex, possibly originating from cosmics, displayed vertices, or pileup vertices. The definitions of d_0 and z_0 are shown in Figure 47.




Figure 47: Impact Parameter definitions. d_0 is the distance from the track to the z-axis at its point of closest approach (*perigee*). z_0 is the z-coordinate of the perigee, defined with respect to the primary vertex.

5.4 Jet and photon selection

Jets are required to pass a quality selection called *Bad Loose* in ATLAS [80]. Central, low-momentum jets (defined by $p_T < 60$ GeV and inside $|\eta| < 2.4$) are required to pass a jet-vertex tagger cut designed to eliminate contamination from pileup. Jets passing these requirements, as needed, are kept for $p_T > 30$ GeV and $|\eta| < 4.5$.

Photons are required to pass basic selection criteria based on shower properties as measured in the EMC.

5.5 Overlap removal

The reconstruction algorithms for each physics object run independently of other objects that have been reconstructed, leaving the possibility that two selected objects may share tracks or clusters in the detector. For example, the ATLAS reconstruction algorithm for electrons thinks a particular ID track belongs to an electron, and the reconstruction algorithm for muons thinks it belongs to a muon - both potentially reasonable conclusions, but at least one is incorrect. These cases are removed by the following cuts:

- For two electrons which share ID tracks, the electron with the highest track p_T is kept.
- For two electrons with nearby clusters, the electron with the highest cluster E_T is kept. “Nearby” is taken to mean $|\Delta\eta| < 0.075$ and $|\Delta\phi| < 0.125$.
- When an electron shares an ID track with a CT muon, only the electron is kept.
- When an electron shares an ID track with any other muon, the muon is kept.
- When a jet overlaps with an electron within $\Delta R < 0.2$, the jet is removed.

5.6 4ℓ selection

The goal of the event selection is to identify and re-assemble the $H \rightarrow ZZ^{(*)} \rightarrow 4\ell$ decay in the detector from the leptons passing the previously described selection. To do this, the following requirements are imposed in order.

Events are required to have either four muons, four electrons, or two electrons and two muons.

Next, all possible pairs of same flavour (electron-electron, or muon-muon), opposite sign (positive-negative charge) leptons are identified, called SFOS. These are the candidate Z boson decays. Events are required to have at least two independent SFOS pairs.

The SFOS pairs are themselves paired up to form all possible quadruplets, and the event selection proceeds by applying cuts to these quadruplets. If an event contains no quadruplet passing the selection, the event itself is cut.

The quadruplets must pass a kinematic cut which requires the leptons, ordered by p_T , to pass

$$\begin{aligned} p_{T_1} &> 20 \text{ GeV}, \\ p_{T_2} &> 15 \text{ GeV}, \\ p_{T_3} &> 10 \text{ GeV}. \end{aligned} \tag{55}$$

Quadruplets are also required to contain one (or more) of the leptons which caused the trigger to fire for this event.

Since the decay of the Higgs proceeds via two Z bosons, knowledge of this intermediate step can be used in the event selection. For m_H near 125 GeV, one of the Z s can be on-shell. In the quadruplets reaching this stage of the event selection, the primary di-lepton pair, called Z1, is identified as the one with an invariant mass, m_{Z1} , closest to $m_Z^{PDG} = 91.1876 \text{ GeV}$. Any quadruplets not containing this Z1 are removed. Among the remaining quadruplets, the one with a secondary di-lepton pair, called Z2, closest to m_Z^{PDG} is kept, while all others are discarded.

Note that this definition of Z1 and Z2 is what distinguishes the two channels called $2e2\mu$ and $2\mu2e$ - they are named according to which leptons form the primary and secondary di-lepton pairs (in that order). The channels $4e$ and 4μ use the same definitions of Z1 and Z2, though their naming introduce less grounds for confusion. The above cuts are applied on each channel separately, so at this stage of the cuts there remains at most one quadruplet per event and per channel.

Cuts are imposed on the di-lepton masses, m_{Z1} (also called m_{12}), and m_{Z2} (also called m_{34}) as shown in Figure 48. Note that the m_{Z2} cut depends on the four-lepton invariant mass, $m_{4\ell}$, because with increasing $m_{4\ell}$, m_{Z2} is increasingly able to be on-shell, so a stricter cut can be imposed. For m_H near 125 GeV, m_{Z2} from a signal event can never be on-shell.




Figure 48: Cuts imposed on the di-lepton masses m_{Z1} and m_{Z2} . Note that both sets of cuts extend flatly off the plotted range in $m_{4\ell}$, and that the m_{Z2} acceptance region is only a lower bound (it extends indefinitely off the plotted vertical range).

5.7 Additional Criteria

5.7.1 ΔR and cross-pairing

To reject boosted B-meson decays, quadruplets are required to have for all same-flavour leptons a separation of $\Delta R_{\ell_i \ell_j} > 0.10$, and for all different-flavour leptons a separation of $\Delta R_{\ell_i \ell_j} > 0.20$. Additional possible contamination from B-meson decays is rejected by cutting events where any cross-paired (one lepton from the Z1, the other from the Z2) SFOS di-lepton pair has $m_{\ell\ell} < 5$ GeV.

5.7.2 Isolation

To ensure high momentum resolution, and to reject backgrounds primarily from $t\bar{t}$ events, the leptons are required to be isolated from nearby particles in the detector. This isolation requirement is imposed in two steps: track isolation, and calo isolation, which are applied to all four leptons in the quadruplet.

Track isolation is a measure of nearby tracks to that of the lepton. The p_T of tracks within $\Delta R < 0.30$ of the lepton is summed (excluding the lepton itself), and expressed as a ratio over the lepton p_T . A cut is applied requiring $p_T^{\text{varcone30}}/p_T^{\text{lepton}} < 0.15$ for both electrons and muons. In plainer language, no more than 15% of the lepton's p_T should be found in the p_T of other tracks within $\Delta R < 0.30$ of the lepton.

Calo isolation applies a similar cut on the E_T of clusters within $\Delta R < 0.20$ of the lepton cluster. For electrons (muons) the cut is $E_T^{\text{topocone20}}/E_T^{\text{lepton}} < 0.20$ (0.30)). This is read in a similar way: no more than 20% (30%) of an electron's (muon's) E_T should be found in the E_T of other clusters within $\Delta R < 0.20$ of the lepton.

5.7.3 d_0 significance and vertexing

A cut is applied to reject displaced vertices by requiring $d_0/\sigma_{d_0} < 3$ for muons and $d_0/\sigma_{d_0} < 5$ for electrons, where d_0 is defined as in Figure 47a, and σ_{d_0} is the resolution on that quantity from the track reconstruction.

In addition to this, a simultaneous fit of the four leptons' ID tracks² is performed with a constraint to have come from a common vertex. A cut of $\chi^2 < 5$ is applied to the fit result.

5.8 FSR recovery

In a small percentage of events, one of the final-state leptons can emit final-state radiation (FSR), which carries away energy, leading to a loss of $m_{4\ell}$ resolution. An attempt is made to recover this energy by identifying candidate FSR photons to be re-added to the quadruplet. At most one such photon will be recovered.

Collinear ($\Delta R < 0.15$) photons are recovered only to muons coming from Z1. Electrons can also emit collinear radiation (bremsstrahlung), but this is typically caught in the same EM cluster as the electron, and so there is no need to recover it at this stage³. A collinear photon is accepted only if alters m_{Z1} from within 66-89 GeV to 66-100 GeV. Where multiple photons exist, the one with the highest E_T is used.

When no collinear FSR can be recovered, an attempt is made to recover far FSR ($\Delta R > 0.15$ for all leptons in the quadruplet). This is done for both flavours of lepton from both the selected Z1 and Z2. A photon can be recovered only if $m_{4\ell} < 190$ GeV, and if the recovery alters the di-lepton mass from within $m_{\ell\ell} < 81$ GeV to $m_{\ell\ell} < 100$ GeV. Where multiple photons exist, the one with the highest E_T is used. When separate candidate FSR photons meet these requirements both for Z1 and Z2, the photon which alters the di-lepton mass by the smallest amount is taken.

²When a quadruplet contains an ME muon, only three ID tracks are used.

³The electron reconstruction described in Section 4.4 would already have included co-linear FSR inside the cluster.

The invariant mass after the FSR recovery procedure (even if no FSR is recovered) is called $m_{4\ell}^{\text{fsr}}$, which has a slightly improved resolution compared to $m_{4\ell}^{\text{reco}}$, to be demonstrated in Section 5.10. This gain in $m_{4\ell}$ resolution translates directly to the final m_H resolution.

5.9 Z Mass Constraint

Knowledge of the decay chain $H \rightarrow ZZ^{(*)} \rightarrow 4\ell$ can be used to gain additional $m_{4\ell}$ resolution for the final state. For the same reason that the true invariant mass distribution of the Higgs' decay products will follow a Relativistic Breit-Wigner (with to-be-measured parameters m_H, Γ_H), so too does the decay of any on-shell Z boson in the decay chain. Unlike the Higgs, the parameters of the Z boson lineshape are very well-known: LEP measured [81] the mass $m_Z = 91.1875 \text{ GeV}$, and the width $\Gamma_Z = 2.4952 \text{ GeV}$ with uncertainty in the final two digits⁴. This known lineshape can be used to constrain the measurement uncertainties for leptons coming from an on-shell Z decay, a procedure call the Z Mass Constraint (ZMC) [82].

The mass of the Z boson to di-leptons (called ℓ_1, ℓ_2 with four-momenta p_1, p_2) before detector resolution is taken into account is called the *true* mass, and can be expressed as

$$(m_{12}^{\text{true}})^2 = (p_1^{\text{true}} + p_2^{\text{true}})^2, \quad (56)$$

$$= p_1^{\text{true}2} + p_2^{\text{true}2} + 2p_1^{\text{true}} \cdot p_2^{\text{true}}, \quad (57)$$

$$= 2E_1^{\text{true}} E_2^{\text{true}} (1 - \cos \theta_{12}^{\text{true}}), \quad (58)$$

where θ_{12} is the relative angle between the two leptons. The final equality applies only in the high energy environment of the LHC, where the invariant mass of an individual electron or muon, $p_i^2 = m_i^2$, is small enough to neglect, leaving $E_i \approx |\vec{p}_i|$. The quantity m_{12}^{true} is distributed as a BW distribution with parameters m_Z, Γ_Z , introduced in (2).

When reconstruction effects - a combination of many stochastic processes - are taken into account, the reco-level m_{12} no longer follows its true distribution. Each of the lepton energies, E_i^{true} are measured to be E_i^{reco} with some resolution σ_{E_i} (the same is true of the lepton direction, and thus θ_{12} , but this effect is far smaller), which effectively smears out the distribution of m_{12}^{true} into that of m_{12}^{reco} (also called m_{Z1}).

⁴This uncertainty could in principle be included in the statistical procedure to follow, but because it is so much smaller than individual lepton measurement uncertainties in ATLAS, it would have a vanishingly small effect.

The goal of the ZMC is to use our knowledge of the $P(m_{12}^{\text{true}})$ distribution to remove or reduce the effect of measurement uncertainties on the leptons. To do this, we write down a likelihood expression for m_{12}^{true} (a function of E_i^{true}) given the observed mass m_{12}^{reco} (a function of E_i^{reco}) and resolution σ_m (a function of E_i^{reco} , σ_{E_i}) as

$$\begin{aligned} P(m_{12}^{\text{true}}|m_{12}^{\text{reco}}, \sigma_m) &= \frac{P(m_{12}^{\text{reco}}|m_{12}^{\text{true}}, \sigma_m) P(m_{12}^{\text{true}})}{P(m_{12}^{\text{reco}})}, \\ &\propto P(m_{12}^{\text{reco}}|m_{12}^{\text{true}}, \sigma_m) P(m_{12}^{\text{true}}), \\ &\propto N(m_{12}^{\text{reco}}; m_{12}^{\text{true}}, \sigma_m) \text{BW}(m_{12}^{\text{true}}; m_Z, \Gamma_Z), \end{aligned} \quad (59)$$

where the first equality uses Bayes' Theorem [83], and the final proportionality assumes a Gaussian resolution distribution, and substitutes the actual BW distribution for m_{12}^{true} .

By maximizing this likelihood, the constrained Z1 mass is derived:

$$m_{12}^{\text{cons}} \equiv \underset{m_{12}^{\text{true}}}{\text{argmax}} P(m_{12}^{\text{true}}|m_{12}^{\text{reco}}, \sigma_m). \quad (60)$$

This mass should be interpreted as the “most likely true mass for this event, given reconstructed quantities”, and *should not* be interpreted as equivalent to m_{12}^{true} . A noteworthy consequence is that $P(m_{12}^{\text{cons}}) \neq P(m_{12}^{\text{true}})$, which can be seen by looking at the distributions of these quantities in simulated data, shown in Figure 49. Note that these distributions are presented in a chronological order: an event occurs (or is generated) with a di-lepton mass m_{12}^{true} , which then passes through the detector and is reconstructed with mass m_{12}^{reco} , then has FSR recovered to become m_{12}^{fsr} , and finally the ZMC is applied to give m_{12}^{cons} .

Note that in practice it is the lepton energies which are adjusted by this process, such that the invariant mass calculation remains consistent:

$$(m_{12}^{\text{cons}})^2 = (p_1^{\text{cons}} + p_2^{\text{cons}})^2. \quad (61)$$




Figure 49: Distributions of m_{12}^{true} , m_{12}^{reco} , m_{12}^{fsr} and m_{12}^{cons} on simulated signal sample inclusive of production modes with $m_H = 125$ GeV.

These same constrained leptons⁵ can then be used to calculate a constrained four-lepton invariant mass

$$(m_{4\ell}^{\text{cons}})^2 = (p_1^{\text{cons}} + p_2^{\text{cons}} + p_3^{\text{cons}} + p_4^{\text{cons}})^2, \quad (62)$$

which is used as the ultimate mass discriminant in this thesis. Note that the momenta $p_{3,4}^{\text{cons}}$ are unchanged from $p_{3,4}$, but the notation is kept consistent with leptons 1,2 to mean *after the ZMC*, even if no change was made in this case.

5.10 Invariant mass resolution

Simulated signal samples are used to verify the effects of the FSR recovery, and the ZMC on the final $m_{4\ell}$ resolution. To do this, samples with $m_H = 125$ GeV of all SM Higgs production modes are combined and weighted according to their respective cross sections. The $m_{4\ell}$ distributions in each decay channel for events

⁵The superscript “cons” should be interpreted to mean the leptons after the ZMC procedure, whether or not a lepton was affected by the constraint. For events around 125 GeV, the constraint only makes sense for the leptons from Z1, as only one Z boson can be on shell. At higher masses, approximately above 180 GeV, the constraint is applied to the Z1 and Z2 pairs, as they could have both come from on-shell Z boson decays.

passing the full selection are compared in the three cases: The $m_{4\ell}^{\text{reco}}$ distributions with no corrections are shown in Figure 50, the $m_{4\ell}^{\text{fsr}}$ distributions with FSR recovery are shown in Figure 51, and the $m_{4\ell}^{\text{cons}}$ distributions with FSR recovery + ZMC are shown in Figure 52. The inclusive distributions (all decay channels merged) are shown for the three cases in Figure 53.

In each plot, the core of the invariant mass distribution⁶ is fit with a Gaussian. Because the samples used have a very narrow width ($\Gamma_H = 4.1 \text{ MeV}$), the spread of events is determined almost entirely by detector resolution, and so this Gaussian can be used to quantify the detector resolution. As Chapter 7 will expound greatly on, a simple Gaussian is insufficient to describe the entire resolution, though for now it serves as an adequate first-order approximation. The detailed results of these fits are found in Table 9, and by this method we conclude the FSR recovery brings on average a 1.4% benefit to the $m_{4\ell}$ resolution, and the ZMC brings an additional 19% benefit, as compared by the Gaussian-fit σ of the inclusive distribution in each case.

Table 9: Summary table of four-lepton invariant mass resolution. Shown are the Gaussian-fit parameters, μ and σ , on the invariant mass distribution (i) uncorrected, (ii) with FSR recovery, and (iii) with FSR recovery and Z Mass Constraint applied, as determined on MC simulation at $m_H = 125 \text{ GeV}$.

	Uncorrected		With FSR		With FSR and ZMC	
channel	$\mu[\text{GeV}]$	$\sigma[\text{GeV}]$	$\mu[\text{GeV}]$	$\sigma[\text{GeV}]$	$\mu[\text{GeV}]$	$\sigma[\text{GeV}]$
4μ	124.48	1.92	124.55	1.88	124.80	1.55
$4e$	123.74	2.86	123.78	2.82	124.46	2.29
$2\mu 2e$	124.11	2.40	124.20	2.37	124.40	2.17
$2e 2\mu$	124.09	2.26	124.11	2.24	124.78	1.65
4ℓ	124.22	2.21	124.29	2.18	124.70	1.77

⁶The core is defined somewhat arbitrarily by the mean and standard deviation (s.d.) of the distribution to be the range [mean - s.d. , mean + s.d.]




Figure 50: Four-lepton invariant mass distribution and resolution from simulated $m_H = 125$ GeV signal in the four decay channels with no corrections.




Figure 51: Four-lepton invariant mass distribution and resolution from simulated $m_H = 125$ GeV signal in the four decay channels with FSR recovery.




Figure 52: Four-lepton invariant mass distribution and resolution from simulated $m_H = 125 \text{ GeV}$ signal in the four decay channels with FSR recovery and Z Mass Constraint.




Figure 53: Four-lepton invariant mass distribution and resolution from simulated $m_H = 125$ GeV signal inclusive of all decay channels, shown without corrections, with FSR recovery, and with FSR recovery + Z Mass Constraint

5.11 Analysis categories

In the following chapters, models for the signal and backgrounds will be described. To enhance sensitivity, the analysis is divided into four *categories* according to the decay channel: $4\mu, 4e, 2\mu 2e, 2e 2\mu$, where the mixed channels are defined as described in Section 5.6. Included in the definition of each category is a *mass window* surrounding the Higgs boson mass: $110 \text{ GeV} \leq m_{4\ell} \leq 135 \text{ GeV}$. The analysis will be executed using both $m_{4\ell}^{\text{fsr}}$ and $m_{4\ell}^{\text{cons}}$, and will use the same mass window in both cases. The mass window is chosen to capture any reasonable signal distribution (very conservatively, say $122 \text{ GeV} \leq m_H \leq 128 \text{ GeV}$) with high efficiency. The use of a mass window is purely pragmatic: using an arbitrarily large window puts a greater onus on the background modelling to be described in Chapter 6, without gaining much in return.

This is a good opportunity to distinguish between the two terms *channel*, and *category*, which at first may appear interchangeable. The channel is a physics-based concept, defined according to the decay mode of the two Z bosons in the decay chain. In reconstructed events this can't be known with perfect accuracy, but it is labelled according to the best of our knowledge. The category, however, is a statistical concept - nearly all analyses in ATLAS are separated into categories to increase sensitivity of the measurement. This improvement in sensitivity comes when categories are defined so as to separate classes of events with different resolutions, or different signal over background ratios. In this thesis, the analysis categories are defined according to the decay channel, because the decay channel is a reasonably good separator of event resolution as seen in Section 5.10. This concept is further explored in Chapter 7.

Chapter 6

Background Estimation and Modelling

The event selection described in Chapter 5 is designed to identify $H \rightarrow ZZ^{(*)} \rightarrow 4\ell$ events in ATLAS data. Any events which pass this selection, but did not come from Higgs boson decays, constitute the *backgrounds* to this analysis. Understanding and describing the backgrounds is essential to measuring the properties of the Higgs signal events.

This chapter is divided into four sections. Section 6.1 describes the backgrounds from SM (non-Higgs) ZZ^* production, called *irreducible* background because the four-lepton final state that gets reconstructed is nearly indistinguishable from signal events. Section 6.2 describes backgrounds from other SM processes producing four (or more) leptons - because of additional final state particles, or different topologies, these are more easily distinguished from signal. Section 6.3 describes backgrounds arising from other particles which *fake* electrons or muons in the detector; these backgrounds are called *reducible*, because generally they can be reduced by tighter quality cuts, primarily those in Section 5.3 and Section 5.7. Finally Section 6.4 summarizes the expected yields of each background.

For each background, two aspects are described:

- Yield** How many events from each source we expect to observe in the data, also called the *normalization*.
- PDF** The Probability Density Function (PDF); the expected distribution of $m_{4\ell}^{\text{fsr}}$ or $m_{4\ell}^{\text{cons}}$, defined in Section 5.8 and Section 5.9, respectively, for each source.

These two concepts are modelled independently, but in general they can be linked by other variables (e.g. a systematic uncertainty may affect the yield in an $m_{4\ell}$ -dependent way, so would be included in both the yield and the PDF). In all cases, the yield and the PDF will depend on the analysis category, as defined in Section 5.11.

6.1 Standard Model ZZ^*

The irreducible background from SM ZZ^* is produced by two main modes: $q\bar{q} \rightarrow ZZ^{(*)}$ (nicknamed qqZZ), and $gg \rightarrow ZZ^{(*)}$ (nicknamed ggZZ), where both Z bosons decay to di-lepton pairs. The leading Feynman diagrams for qqZZ and ggZZ are found in Figure 54 and Figure 55, respectively. Of the two, the qqZZ is dominant by more than an order of magnitude, because of the structure of the interaction. Both productions modes are non-resonant inside the mass window, so form continuums in $m_{4\ell}$.




Figure 54: Feynman diagram for $q\bar{q} \rightarrow ZZ^{(*)}$ background.




Figure 55: Feynman diagram for $gg \rightarrow ZZ^{(*)}$ background. The intermediate particle (the “box”), not labelled on the diagram, is a quark.

The yield for qqZZ is estimated using the 100-150 GeV filtered MC simulated sample, based on the number of events passing the full selection into each category, with the sample weighted according to the SM cross section. The yield for ggZZ is

estimated in the same way, using both the $2\ell 2\ell$ and $4\ell\ gg \rightarrow ZZ^{(*)}$ MC simulated samples.

The PDFs for qqZZ and ggZZ are built using the same samples as used for the yields. In both cases they are constructed using kernel smoothing (by the class `RooKeysPdf` [84]) of the simulated data. This smoothing algorithm uses the unbinned data as input, and creates a PDF by summing a Gaussian centered on each datum, with widths dynamically adjusted according to the local event density, and a smoothing parameter ρ (where higher ρ indicates more smooth, and lower ρ indicates less smooth). The algorithm is designed to capture fine structural details in the population distribution where many input data are available, without over-smoothing regions with fewer input data. Furthermore the algorithm includes an option of *mirroring* the data across the boundaries of the input dataset range¹. When the input distribution is non-zero at the boundary of the smoothing range, this option helps regulate the output PDF to prevent it from falling to zero².

The qqZZ background uses $\rho = 3.4$, with mirroring enabled. The smoothing range is set to $105\text{ GeV} \leq m_{4\ell} \leq 140\text{ GeV}$ (a 5 GeV extension on either side of the fit range) to ensure good behaviour near the boundary of the fit range. The qqZZ smoothing results for the constrained mass, $m_{4\ell}^{\text{cons}}$, are found in Figure 56. The ggZZ uses $\rho = 3.0$, and otherwise the same configuration as with qqZZ. The ggZZ smoothing results for the constrained mass are found in Figure 57. The corresponding plots for the unconstrained mass, $m_{4\ell}^{\text{fsr}}$, are found in Appendix A, as they are not substantially different.

At lower masses, outside of the fit window for the m_H analysis, $q\bar{q} \rightarrow ZZ^{(*)}$ can be resonant at the Z mass through the diagram shown in Figure 58 (the same diagram contributes inside the mass window, but to a smaller extent). This diagram enters the event selection when both bosons decay to di-lepton pairs. This diagram forms a peak in the $m_{4\ell}$ spectrum around $m_{4\ell} = m_Z$, and will be exploited later in Section 9.5.

¹An as example, in the case that the smoothing range ends at 140 GeV, the input data are mirrored such that the resulting PDF from $135 \rightarrow 140$ and from $145 \rightarrow 140$ are identical.

²Without input data on the other side of the boundary, events within the smoothing range “spill out” of the boundaries, but no events outside “spill in” to compensate.




Figure 56: Smoothed distribution of $m_{4\ell}^{\text{cons}}$ from simulated $q\bar{q} \rightarrow ZZ^{(*)}$ events. The black points represent the simulation, the solid blue line is the kernel smoothing result used as the PDF.




Figure 57: Smoothed distribution of $m_{4\ell}^{\text{cons}}$ from simulated $gg \rightarrow ZZ^{(*)}$ events. The black points represent the simulation, the solid blue line is the kernel smoothing result used as the PDF.




Figure 58: Feynman diagram for resonant $q\bar{q} \rightarrow ZZ^{(*)}$ background. This diagram is most important near $m_{4\ell} = m_Z$ where it forms a resonance.

6.2 $VVV, t\bar{t} + Z$

The background nicknamed VVV has two main components. The first is tri-boson production, VVV (meaning ZZZ, ZZW, ZWW). When some combination of the vector bosons decay leptonically, these events can enter the event selection. Examples of the leading diagrams for each are shown in Figure 59. The diagrams shown here are meant to be instructive but not complete - a plethora of diagrams with the bosons and leptons swapped into different configurations also contribute.

The second background grouped under the nickname VVV is $t\bar{t} + Z$. The leading diagrams for this process are shown in Figure 60.

The yield from both these processes is estimated together using MC simulated samples for each, weighted according to their SM cross sections. The PDF is built using kernel smoothing of the sample samples, with $\rho = 2.8$, mirroring enabled, and the smoothing range extended 5 GeV beyond the fit range, as was done for the SM ZZ^* backgrounds. The results when smoothing $m_{4\ell}^{\text{cons}}$ are shown in Figure 61, while $m_{4\ell}^{\text{fsr}}$ are found in Appendix A.




Figure 59: Feynman diagrams for tri-boson background.




Figure 60: Feynman diagram for $t\bar{t} + Z$ background. For both diagrams, the top pair decays as $t \rightarrow W^+b$, $\bar{t} \rightarrow W^-b$ (and then $W \rightarrow \ell\nu$), but it is left unlabelled to reduce clutter. In the analysis, this process is grouped together with VVV under the same label.




Figure 61: Smoothed distribution of $m_{4\ell}^{\text{cons}}$ from simulated $VVV, t\bar{t} + Z$ events. The black points represent the simulation, the solid blue line is the kernel smoothing result used as the PDF.

6.3 $Z + \text{Jets}, t\bar{t}, WZ$

The background $Z + \text{Jets}$, $t\bar{t}$ and WZ is collectively nicknamed $Z + \text{Jets}$, for its dominant component. It is the most challenging background entering the analysis, because it is composed of other particles faking prompt electrons and muons³. These processes are often difficult to simulate, and quite rare, such that the MC simulation-based approach used for the other backgrounds can not be used. For these reasons, $Z + \text{Jets}$ is the only background in this thesis whose yield is *estimated from data*. Because the estimation of fake electrons differs from that of fake muons, the two are treated separately: Section 6.3.1 describes the so-called $Z + ee$ yield estimation for events containing fake electrons (entering the $4e$ and $2\mu 2e$ categories), and Section 6.3.2 describes the same for fake muons called $Z + \mu\mu$ (entering the 4μ and $2e 2\mu$ categories). More detail is presented for the $Z + ee$ estimation than for $Z + \mu\mu$, as the author was directly involved in the former, but not the latter. The strategy and statistical tools for the $Z + ee$ estimation were developed by the author in collaboration with Ioannis Nomidis and William Leight, while the final execution (estimates, plots), were performed by Ioannis Nomidis. Section 6.3.3 describes how the PDF for the $Z + \text{Jets}$ backgrounds is built. Note that the yield estimation described in these sections is inclusive of $m_{4\ell}^{\text{cons}}$; the final $Z + \text{Jets}$ yield used in the analysis is derived by integrating the inclusive yield, distributed according to the PDF, from $110 \text{ GeV} - 135 \text{ GeV}$.

6.3.1 $Z + ee$

The electron background estimation is separated by the source of the fake electrons. While the final statistical analysis is unconcerned with the source of the fake electrons (only their number matters), understanding the composition of the background provides more confidence in its estimation, and is important for constructing the $Z + \text{Jets}$ PDF. The different sources of fake electrons are labelled as follows, ordered from most-to-least abundant:

- f Light jets. These are initiated by 1st or 2nd generation quarks, or gluons, where in certain cases the calorimeter shower profile can be reconstructed as an electron.
- γ Events where a photon converts via $\gamma \rightarrow e^+e^-$. In cases where the tracks are not separately reconstructed, and the EM showers overlap, these can be

³This includes non-prompt, but real-electrons and muons, which did not originate from the HS event.

mistakenly reconstructed as a single electron or positron.

- q** Heavy quarks, when they decay semi-leptonically, can produce real electrons that enter the event selection. These are still grouped together with the fake electrons because they are produced non-promptly (not an immediate product of the HS event) . Unlike the other two sources, the q background is estimated from MC.

In MC simulation, the true origin of any electron (real or fake) is known from its truth information, but in data the composition must be estimated⁴. To estimate the composition, a variable called n_{InnerPix} is used. It is defined as the number of pixel hits on the inner-most (closest to the beam-line) active pixel layer. In some cases the physically inner-most pixel modules are non-functioning, and the second layer is used. This variable is a powerful discriminant between f and γ , because each occupy mostly different bins of the n_{InnerPix} distribution.

To estimate the background in the Signal Region (SR) (i.e. the background which passes the full event selection), several Control Regions (CR) are used. Conceptually, the procedure begins “farther” from the signal region, and moves “closer” until an estimate in the SR is derived.

6.3.1.1 $Z + X$ Control Region

First, a CR called $Z + X$ is built in simulation and data. The selection for this CR proceeds similarly to the regular event selection, except,

- Only the single-electron and di-electron triggers of Table 8 are used, to mitigate selection bias on additional leptons in the event.
- Only a single di-lepton pair, called Z1, is required. It has all the same selection requirements as imposed in the main selection, including those of Section 5.7 modified for only two leptons. The only alteration is that a mass window cut of $76 \text{ GeV} < m_{Z1} < 106 \text{ GeV}$ is applied to be more confident a real Z boson was selected for the Z1.
- In addition to the Z1 pair an additional electron, called an X -electron, is required. This electron is not required to pass the likelihood criteria, or the ID hit requirements of Section 5.3.1. The only track quality cuts imposed are on

⁴The effective composition of individual electrons could be estimated too, but this is already done by the electron selection likelihood in Section 5.3.1. The working point used in that section deliberately allows some fake electrons to pass, as the necessary sacrifice to have high signal efficiency.

the total number of pixel and SCT hits, and the d_0 significance cut (described in Section 5.7.3). The isolation cuts of Section 5.7.2 are also removed for this electron.

The events passing the $Z+X$ selection in MC simulated are used to build templates of the n_{InnerPix} distribution for γ and f fake electrons. These expected distributions from simulation are shown in Figure 62.




Figure 62: Templates of the n_{InnerPix} distribution for γ and f -classified fake electrons built using the $Z + X$ CR in simulation. The template for real electrons (from ZZ^*) follows the exact same shape as for q . This plot was provided by Ioannis Nomidis.

In practice, the variable n_{InnerPix} is not perfectly modelled by simulation. To account for this, a data/MC comparison in the $Z + X$ CR is used to correct the template for f . Since f is the leading contributor (by far) to these events, only the f template is corrected by this procedure. Figure 63 shows the comparison of data and MC used to derive the correction. Included in this plot is the contribution from real electrons (combining non-prompt from q , and prompt electrons) from simulation. The truth information is used to classify the electrons in simulation.

Next, the $Z + X$ CR is divided into the f -enriched, and γ -enriched $Z + X$ CRs by requiring $n_{\text{InnerPix}} = 0$ and $n_{\text{InnerPix}} \geq 1$, respectively. After these selections, the purity of the X -electrons, defined here as N_γ/N_X in the γ -enriched CR, and N_f/N_X




Figure 63: Comparison of data and MC simulation of events passing the $Z + X$ CR selection. The distribution is simulation is separated by the origin of the X -electron as determined by the MC truth information. The difference between the data and simulation is used to correct the n_{InnerPix} template for f electrons. This plot was provided by Ioannis Nomidis.

in the f -enriched CR, is measured in simulation as 90% and 98% in the two CRs, respectively, indicating they are highly pure samples.

The data events in these enriched-CRs are used to determine the efficiency of the X -selection (defined, somewhat paradoxically, as the efficiency of the cuts that were *not* imposed on the X -electron) in data for f and γ . Like the template distributions, the yield must also be corrected for data/MC discrepancies. A scale-factor is defined for each species as:

$$N_{f(\gamma)}^{\text{data}} = s_{f(\gamma)} N_{f(\gamma)}^{\text{MC}}, \quad N_{f(\gamma)}^{\text{data, pass } X} = s_{f(\gamma)}^{\text{pass } X} N_{f(\gamma)}^{\text{MC, pass } X}, \quad (63)$$

where $N_{f(\gamma)}^{\text{data(MC)}}$ is the yield in data (MC) of $\gamma(f)$ electrons, $s_{f(\gamma)}$ denotes the data/MC scale factor to be determined for source $f(\gamma)$, and “pass X ” denotes electrons which pass the X -selection (like those in the SR). Of course, the yield of a given species of background can’t be determined exactly in data - it must be measured. This is done through an iterative procedure using the expected yields in MC simulation of each species of background. The two-step iteration starts by determining s_γ from the

γ -enriched $Z + X$ CR by subtracting from data the expected yields of $q + e$ and f ,

$$s_\gamma = \frac{N^{data} - N_{q+e}^{MC} - N_f^{MC}}{N_\gamma^{MC}}, \quad (64)$$

and then using this value of s_γ , determines s_f from the f -enriched $Z + X$ CR,

$$s_f = \frac{N^{data} - N_{q+e}^{MC} - s_\gamma N_\gamma^{MC}}{N_f^{MC}}, \quad (65)$$

and then re-evaluates s_γ , by using the updated s_f ,

$$s_\gamma = \frac{N^{data} - N_{q+e}^{MC} - s_f N_f^{MC}}{N_\gamma^{MC}}. \quad (66)$$

This process is repeated three times (terminating after the third calculation of s_γ), which is sufficient for the estimate to stabilize to fluctuations far below the systematic uncertainties assigned to the final estimate. The same iterative procedure is used to calculate $s_{f(\gamma)}^{\text{pass } X}$ from events passing the X -selection.

Finally the efficiency for the X -selection can be calculated as

$$\epsilon_{f(\gamma)} = \frac{N_{f(\gamma)}^{\text{data, pass } X}}{N_{f(\gamma)}^{\text{data}}} = \frac{s_{f(\gamma)}^{\text{pass } X} N_{f(\gamma)}^{\text{MC, pass } X}}{s_{f(\gamma)} N_{f(\gamma)}^{\text{MC}}} \quad (67)$$

In fact, this efficiency is calculated in three separate p_T bins, to account for the varying X -selection efficiency as a function of p_T . These efficiencies are shown in Figure 64 for two different MC generators. The resulting efficiencies, averaged over p_T , will be recorded in Table 10, though the following section is necessary to understand the remainder of the table.

6.3.1.2 $3\ell + X$ Control Region and Extrapolation to the Signal Region

A second CR, called $3\ell + X$, is built in simulation and data. This CR is more closely related to the SR selection described in Chapter 5, having all the same selection, except:

- The lowest- E_T electron of the Z2 pair has the exact same relaxed requirements as in the $Z + X$ control region. This electron is the X in $3\ell + X$.
- For the Z2 pair, only same-sign ($e^\pm e^\pm$) pairs are considered. This requirement




Figure 64: Final X selection efficiencies as a function of p_T of γ and f fakes, produced using two different generators of the MC simulated samples. These plots were provided by Ioannis Nomidis.

is designed to reject most of the contribution from ZZ^* , though even after this cut about 5% of the events in the CR are from ZZ^* . These surviving events occur when a fake lepton accompanies the ZZ^* event, and replaces one of the true leptons in the selection.

- To consider fakes occurring in any potential Z2 candidate, multiple quadruplets which share the same Z1 are allowed to pass the selection.

The key assumption of this background estimation technique is that the X in $3\ell + X$ and the X in $Z + X$ share the properties. This was tested by comparing the n_{InnerPix} distributions of $X = f, \gamma, q$, respectively, for the $Z + X$ selection and $3\ell + X$ selection using MC simulated samples, where the PDFs of n_{InnerPix} were found to agree.

The data distribution of n_{InnerPix} in the $3\ell + X$ CR was fit for the composition of N_f, N_γ using the data-corrected templates built using the $Z + X$ CR. This fit is shown in Figure 65. After the fit, the expected yields of ZZ^* and q as determined from the $3\ell + X$ CR in simulation are subtracted from the fit results. The remainder is the yield of f, γ in the $3\ell + X$ CR.

The final step is to extrapolate these yields to the SR using the efficiencies derived from the $Z + X$ CR. To account for the p_T -dependence of the efficiencies, the tool **sPlot** [85] is used to unfold the decomposed n_{InnerPix} distributions into decomposed p_T distributions, according to actual 2D distribution in data of n_{InnerPix} and p_T in




Figure 65: Fit on the n_{InnerPix} distribution in the $3\ell + X$ CR in data. The fit templates are built from the $Z + X$ CR.

Table 10: Summary table for reducible electron background estimate. The final yields in the SR are quoted with statistical uncertainty propagated from the fit $3\ell + X$ fit, and systematic uncertainty from the efficiency estimate. For the background from heavy flavour, the uncertainty is not separated.

type	data fit	$ZZ^* + q$	$N_i^{3\ell+X}$	efficiency [%]	N_i^{SR}
f	3075 ± 56	280 ± 6	2795 ± 56	0.22 ± 0.01	$5.68 \pm 0.36 \pm 1.19$
γ	208 ± 17	19.4 ± 0.5	179 ± 17	0.72 ± 0.03	$1.34 \pm 0.35 \pm 0.27$
q	estimated from simulation				5.7 ± 1.7

the $3\ell + X$ CR. This tool works by fitting the n_{InnerPix} distribution for the relative contributions from each source, and then projecting this fit result across the observed 2D data distribution of n_{InnerPix} and any other variable (in this case p_T). This decomposed p_T spectrum is multiplied by the p_T -dependent efficiencies to produce the final estimate for each background in the SR. Table 10 shows the yield estimates fit in the $3\ell + X$ CR, and extrapolated to the SR. The efficiencies quoted in the table are averaged over p_T .

6.3.2 $Z + \mu\mu$

Much like the electron background, the background from muons arises from multiple sources. The dominant source is from Z bosons accompanied by heavy flavour hadrons (e.g. B-mesons) which decay semi-leptonically, called Z +HF. Also contributing are muons coming from light jets via $\pi^\pm/K^\pm \rightarrow X + \mu^\pm$, as well as $t\bar{t}$, and WZ . The WZ contribution is estimated from simulation, while the others are estimated from data.

The estimation of $Z + \mu\mu$ proceeds as described in [5], with the dataset and all expectations from simulation updated to 36.1 fb^{-1} . The measurement defines several CRs each designed to be enriched in the different source of background. The selection for each CR is similar to the nominal event selection, with a few cuts inverted/relaxed in order to enhance the background contribution. The full details of the selections are omitted to put more emphasis on the cuts which do the most to enrich each background.

- **Inverted d0:** By inverting the d_0 significance cut, this CR is enriched in heavy flavour decays which tend to have large impact parameters (since the heavy flavour hadron has non-negligible $\beta c\tau$ before it decays)
- **$e\mu + \mu\mu$:** By requiring that the Z1 is formed from an $e^\pm\mu^\mp$ di-lepton pair, this CR is enriched in $t\bar{t}$, while the contribution from Z decays is highly suppressed.
- **Inverted Isolation:** By inverting isolation on the sub-leading muon of the Z2 pair, this CR is enriched in light jets.
- **Same-sign:** In this CR, the Z2 is required to have same-sign leptons ($\ell^\pm\ell^\pm$). All of the background sources enter into this CR.

In each CR, the expected distribution of each source is modelled using MC simulation. Using this model, the data in each CR enter a simultaneous fit to determine the yield of the different backgrounds in each CR. As an illustration, the fits from a preliminary result [5] using 14.8 fb^{-1} are shown in Figure 66. Finally these yields are extrapolated back to the signal region using transfer factors for the inverted/modified cuts, determined using simulation.

6.3.3 $Z + \text{Jets}, t\bar{t}, WZ$ PDF

Similar to the yield estimation, the PDF of the Z +Jets background is split between a $Z + ee$ shape (used for the $4e$ and $2\mu 2e$ categories), and a $Z + \mu\mu$ shape (used for




Figure 66: Distributions in the different CRs used in the simultaneous fit to estimate $Z + \mu\mu$ background composition, taken from [5]. These plots are shown for 14.8 fb^{-1} , while the estimations used in this thesis (performed by ATLAS collaborators) use the full 36.1 fb^{-1} .

the 4μ and $2e2\mu$ categories).

For the $Z + \mu\mu$ PDF, the $m_{4\ell}^{\text{cons}}$ distribution of simulated Z+HF events passing the nominal event selection is kernel smoothed to build a PDF. The same is done separately for $t\bar{t}$. Next, these two distributions are summed according to their relative expected yields⁵ from the data-driven estimates, to build the Z+Jets PDF. The relative contribution of each is varied in accordance with the uncertainty of the yield estimate to create systematic variations of the shape. The backgrounds from light flavour, and WZ are small enough that they are assumed to be absorbed within these systematic variations.

For the $Z + ee$ PDF, a partly data-driven approach is taken. In the same way that `sPlot` was used in the fit on n_{InnerPix} in the $3\ell + X$ CR to decompose the distribution of the X -lepton p_T , it can also be used to decompose the distribution of $m_{4\ell}^{\text{cons}}$ (or any other variable) in the CR. This distribution is then multiplied by the p_T -dependent efficiencies from the $Z + X$ CR to extract the $m_{4\ell}^{\text{cons}}$ distribution of f and γ backgrounds in the SR. The shape for q is taken from kernel smoothing of MC simulation. Finally, these three distributions are added together according to their relative expected yields to create the PDF for $Z + ee$. As done for the $Z + \mu\mu$, the yields are varied within their uncertainties to create systematically varied shapes to be used as an uncertainty.

The final shapes for Z+Jets, together with their systematic uncertainties are shown in Figure 67.

⁵By definition, the PDF is always normalized to unity.




Figure 67: PDF used to describe distribution of $m_{4\ell}^{\text{cons}}$ for $Z+\text{Jets}$, $t\bar{t}$, WZ backgrounds (black line). The red and blue lines represent systematic uncertainties both on the reducible background yield and shape (they are not normalized PDFs).

6.4 Backgrounds summary

The final results of the background yield estimates for the full 36.1 fb^{-1} dataset is found in Table 11. The categories are defined as in Section 5.11. Note that the source of uncertainties on these values (except the $Z+\text{Jets}$) has yet to be explained in this thesis; this will be covered in Chapter 8, together with all other systematic uncertainties. The central values are estimated as described above. The $Z+\text{Jets}$ yield in this table is smaller than in Table 10, because the latter was inclusive of $m_{4\ell}$, and here the analysis mass window of $110\text{--}135\text{ GeV}$ is applied on $m_{4\ell}^{\text{cons}}$. The yields using a window in $m_{4\ell}^{\text{fsr}}$ are nearly identical, as the ZMC has a minimal effect on the background distributions.

Table 11: Number of expected background events in each fit category for the full dataset. All sources of systematic uncertainty are included.

Category	$q\bar{q} \rightarrow ZZ^{(*)}$	$gg \rightarrow ZZ^{(*)}$	$VVV, t\bar{t} + Z$	$Z + \text{Jets}, t\bar{t}, WZ$
4μ	15.18 ± 1.49	0.65 ± 0.39	0.18 ± 0.01	1.56 ± 0.40
$4e$	6.39 ± 0.74	0.29 ± 0.18	0.13 ± 0.01	1.97 ± 0.47
$2\mu2e$	6.83 ± 0.77	0.20 ± 0.12	0.11 ± 0.01	2.05 ± 0.47
$2e2\mu$	10.59 ± 1.02	0.31 ± 0.19	0.18 ± 0.01	1.56 ± 0.40
combined	39.00 ± 2.10	1.45 ± 0.49	0.60 ± 0.03	7.14 ± 0.87

Chapter 7

Per-Event Response Signal Model

This chapter describes the signal model used for the mass measurement. Rather, it presents two signal models: the *template method* in Section 7.2 and the *per-event response method* in Section 7.3. The per-event response method is one of the primary contributions of this thesis - it is presented second for pedagogical and chronological reasons: the template method, in various manifestations, has been used to measure particle masses since the 1980's; it is a relatively straightforward technique, and a good introduction to the goals of a signal model. The per-event response model is a more ambitious strategy, which, well-executed, can deliver a more robust and precise measurement. The overarching goal of both models is to describe the expected distribution of $m_{4\ell}$ from $H \rightarrow ZZ^{(*)} \rightarrow 4\ell$ events as a function of m_H and Γ_H , which will allow both to be measured in data from the events passing the selection of Chapter 5. The template method does not include Γ_H dependence, and is forced to assume the SM value.

Before this, the chapter begins on much more humble footing. Section 7.1 describes how the expected number of signal events is modelled as a function of m_H , common to both methods. Also common to both methods are the background distributions described in Chapter 6, and the category definitions as described in Section 5.11.

7.1 Signal Normalization

The expected number of signal events in each category can be parameterized as

$$\begin{aligned} N_{H,c}^{\text{exp.}} &= \sigma_{p\bar{p} \rightarrow H} \times \text{BR}(H \rightarrow 4\ell) \times L \times A_c^{4\ell}(m_H) \times \left(\frac{4\text{BR}(H \rightarrow c)}{\text{BR}(H \rightarrow 4\ell)} \right), \\ &= \sigma_{H \rightarrow 4\ell} \times L \times A_c^{4\ell}(m_H) \times \left(\frac{4\text{BR}(H \rightarrow c)}{\text{BR}(H \rightarrow 4\ell)} \right), \end{aligned} \tag{68}$$

where c is the category (4μ , $4e$, $2\mu 2e$, $2e 2\mu$), $\sigma_{H \rightarrow 4\ell}$ is the total Higgs production cross section times branching ratio to four electrons or muons as it was defined in Section 2.3.3, and L is the integrated luminosity, 36.1 fb^{-1} . The factor $A_c^{4\ell}(m_H)$, is the detector acceptance times efficiency into category c as a function of m_H . For every 1 fb of $\sigma_{H \rightarrow 4\ell}$ and 1 fb^{-1} of data, this is how many events are expected to be reconstructed in category c . It is parameterized using a second order polynomial, fit on the number of events passing the event selection into each category in the MC simulated samples, combining production modes, at different mass points from $m_H = 120\text{-}130 \text{ GeV}$. This parameterization is shown in Figure 68. The final factor is an $\mathcal{O}(1)$ constant used to give the correct share of the 4ℓ branching ratio to each of the sub-channels. The branching ratios are normalized by

$$\text{BR}(H \rightarrow 4\ell) = \sum_c \text{BR}(H \rightarrow c), \quad (69)$$

but they are not identical. The $H \rightarrow eeee, \mu\mu\mu\mu$ channels have about 10% more total width than the $H \rightarrow 2e 2\mu, 2\mu 2e$ channels.

The signal yields expected from the SM, with detector acceptances time efficiency as parameterized by (68), are given for a few mass points in Table 12. The sources of uncertainty on these values are discussed later in Chapter 8.

Table 12: Expected number of $H \rightarrow 4\ell$ signal events, $N_{H,c}^{\text{exp.}}$, in 36.1 fb^{-1} of data in each category for select values of m_H .

category	$m_H = 123 \text{ GeV}$	$m_H = 125 \text{ GeV}$	$m_H = 127 \text{ GeV}$
4μ	19.84 ± 1.29	20.86 ± 1.35	21.80 ± 1.42
$4e$	10.45 ± 0.86	11.12 ± 0.92	11.70 ± 0.97
$2\mu 2e$	10.47 ± 0.81	11.31 ± 0.87	12.06 ± 0.93
$2e 2\mu$	13.98 ± 0.84	14.79 ± 0.89	15.48 ± 0.93
combined	54.74 ± 3.52	58.08 ± 3.73	61.04 ± 3.92




Figure 68: Parameterization of expected signal yield (per unit of cross section times branching ratio) as a function of m_H . The points are derived using simulated signal samples, and fit with a 2nd order polynomial.

7.2 Template Method

The *template method* serves as a benchmark against which the per-event response method can be tested by using a well-established method for measuring the mass of a heavy particles experiment. It was used in the first W boson mass measurement in 1987 [86] up to ATLAS' measurement of the Higgs boson mass from the Run 1 LHC data in 2013 [7]. Conceptually the method is straightforward:

1. Produce simulated data with a range of mass hypotheses.
2. Choose an observable whose distribution is sensitive to the hypothesis mass, and build a template for that distribution at each mass point.
3. Look at the observed distribution in data, and see which hypothesized mass point best matches; this is the best fit value.

In practice, these steps include more nuance and flexibility. One possibility is to use multiple observables (as in [7], among other examples), and correspondingly multi-dimensional template PDFs; this adds the complication of accounting for correlations

between the observables in the templates. Another point of flexibility is whether the fit will be binned (using χ^2 or binned likelihood fit), or unbinned (using an unbinned likelihood fit). If the fit will be unbinned, then one must decide how the templates at each mass point will be built: the most basic method would be a histogram from the simulated data [87], but a more physically motivated approach¹ is to create a smooth distribution using numerical methods, or an analytical parameterization ([7] provides examples of both in the signal model for the $H \rightarrow 4\ell$ channel, and $H \rightarrow \gamma\gamma$ channel, respectively). Finally, one must decide how to perform the final fit: in modern implementations, the templates are usually interpolated between simulated mass points to create a continuous parameterization, but discrete parameterizations together with some basic assumptions can also be functional².

The template method as presented in this thesis uses a 1D unbinned fit on observable $m_{4\ell}^{\text{cons}}$, using templates built from smoothed MC and interpolated using morphing to create a continuous parameterization.

7.2.1 Signal template smoothing

Template PDFs of $m_{4\ell}^{\text{cons}}$, called $P(m_{4\ell}^{\text{cons}})$, are built for each simulated mass point available from $m_H = 120\text{--}130\text{ GeV}$, separately for each analysis category. The templates are built by kernel smoothing the MC simulated data using the KEYS algorithm, described previously in Section 6.1, using $\rho = 2.2$. This value was chosen to produce a smooth, unimodal distribution, without any evidence of the result “chasing” the statistical fluctuations in the input data. Results of the smoothing in the 4μ category are shown in Figure 69, for the $4e$ category in Figure 70, for the $2\mu 2e$ category in Figure 71, and for $2e 2\mu$ in Figure 72, for each of the available simulated mass points.

7.2.2 Signal template morphing

The templates at each mass point are combined together to form a continuous parameterization in m_H using the class `RooParamKeysPdf` [88]. This class assumes that

¹Unless infinite MC statistics are available, then a histogram must have finite bins and the template will be “boxy”, which is not a physical PDF.

²One example, [87], assumes these discrete simulated mass points which form the likelihood curve will be parabolic, and so fits the minimum based on this assumption.




Figure 69: Smoothing result for template signal model. The $m_{4\ell}^{\text{cons}}$ distribution in category 4μ is shown for MC simulated mass points in order $m_H = 120, 122, 124, 125, 126, 128, 130$ GeV.

for some shift values Δ ,

$$P(m_{4\ell}|m_H + \Delta) \approx P(m_{4\ell} - \Delta|m_H), \quad (70)$$

e.g. $P(m_{4\ell}|125.2) \approx P(m_{4\ell} - 0.2|125).$

In this way, the signal templates at the simulated mass points can be interpolated to mass points values that were not simulated by what amounts to a small horizontal




Figure 70: Smoothing result for template signal model. The m_{4l}^{cons} distribution in category 4e is shown for MC simulated mass points in order $m_H = 120, 122, 124, 125, 126, 128, 130$ GeV.

shift. In practice, a sum of several nearby templates with masses m_H^i used,

$$P_{\text{morph}}(m_{4l}|m_H) = \sum_{m_H^i} w_{m_H^i} P_{m_H^i}(m_{4l}^{\text{cons}} - (m_H^i - m_H)),$$

$$\text{where } \sum_{\text{template } i} w_{m_H^i} = 1,$$
(71)




Figure 71: Smoothing result for template signal model. The m_{4l}^{cons} distribution in category $2\mu 2e$ is shown for MC simulated mass points in order $m_H = 120, 122, 124, 125, 126, 128, 130$ GeV.

where the weights are determined according to the distance between the template mass value and the hypothesis value, such that

$$w_{m_H^i} \text{ is largest when } (m_H^i - m_H) \rightarrow 0, \quad (72)$$

$$w_{m_H^i} \rightarrow 0 \text{ when } |m_H^i - m_H| \text{ is large.} \quad (73)$$

This process prevents discontinuities in morphing across the template values, and produces a smooth dependence in m_H . The result of this morphing is shown in Figure 73 for each of the analysis categories, in steps of 0.5 GeV from 120-130 GeV, demonstrating the smooth interpolation of the morphing. One consequence of the




Figure 72: Smoothing result for template signal model. The $m_{4\ell}^{\text{cons}}$ distribution in category $2e2\mu$ is shown for MC simulated mass points in order $m_H = 120, 122, 124, 125, 126, 128, 130$ GeV.

above weighting strategy is that although several of the input mass points are included in this plot (e.g. $m_H = 125$ GeV), the exact templates themselves are not: the input templates which surround 125 GeV are used to improve the continuity of the shape at this mass point, modelling it more reliably than the 125 GeV template alone. In this specific example, the weights take on the values as follows,

$$\begin{aligned}
 P_{\text{morph}}(m_{4\ell}|125) &= \dots + w_{124}P_{124}(m_{4\ell} - 1) + w_{125}P_{125}(m_{4\ell}) + w_{126}P_{126}(m_{4\ell} + 1) + \dots, \\
 &= \dots + 0.17P_{124}(m_{4\ell} - 1) + 0.64P_{125}(m_{4\ell}) + 0.17P_{126}(m_{4\ell} + 1) + \dots,
 \end{aligned}
 \tag{74}$$

where the even more distant templates carry the remainder of the weight. This morphing is stable even when individual mass samples are removed, a fact which is demonstrated in Appendix [B.1](#).




Figure 73: Demonstration of template morphing using RooParamKeysPdf in the four categories. The PDF $P(m_{4l}^{\text{cons}}|m_H)$ is drawn in 0.5 GeV steps from 120 GeV to 130 GeV.

7.3 Per-Event Response method

The *per-event response* (PER) method, and measurements performed with it, are the principal contribution presented in this thesis. It is a novel application of a Gaussian-mixture approach to describe the detector response (in this case, the $m_{4\ell}$ response) of each event. It is inspired by other uses in ATLAS, such as the Gaussian Sum Filter for electron tracking described in Section 4.4.1, where a Gaussian mixture is used to describe the Bethe-Heitler energy loss distribution. In the PER method, a Gaussian mixture is used to describe the energy response distribution of the four leptons in each selected event. These energy response distributions are propagated to calculate the $m_{4\ell}$ response for each event, which is consequently also described by a Gaussian mixture. By convolving this predicted $m_{4\ell}$ response distribution - unique to each event - with the common truth distribution for signal events with unknown parameters m_H , Γ_H the PER signal PDF is constructed.

The method bears similarities to a *per-event error* fit, so it is worthwhile to distinguish them. The principle in both cases is that every event has a unique detector resolution: some are better-measured, some are worse-measured than others. It is natural that these better-measured events should have more importance in the fit than worse-measured events. This is accomplished in practice for a $H \rightarrow 4\ell$ analysis by having the PDF of $m_{4\ell}$ depend explicitly on the measurement resolution. Where the two methods differ is that while per-event errors assumes the general form of the PDF [89] as a Gaussian, Crystal Ball (CB) [90] function, CB+Gaussian, two-sided-CB, etc., and includes a per-event resolution analytically in the width parameter, the PER method does not make such an assumption. Instead it predicts the entire PDF from first principles, based on the lepton kinematics in the event. The benefits of the PER method are most pronounced in channels with low statistics, or highly variable resolution between events, both of which are satisfied in the $H \rightarrow 4\ell$ channel. The PER method also more neatly disentangles the detector component of the $m_{4\ell}$ distribution (the response) from the physics component (the truth distribution). Both these methods differ from the template method, where the PDF is a fixed shape for all events in a given category, created by implicitly averaging the resolution of all events in that category when building each template. As a result, the template method is unable to give better-measured events within each category more importance in the fit than the poorly-measured events, whereas the PER method does exactly that.

It should be noted that in the earlier stages of development of the PER method, it was used as a cross-check of the ATLAS Run 1 Higgs mass measurement [7], and

as a direct measurement of the Higgs boson natural width. The development and execution of the method at this time was done by the author in close collaboration with others from CERN, Harvard University, and Max Planck Institute for Physics. The method at this stage was also presented in detail in [91]. As it is presented in this thesis, the method is executed at every stage by the author, and makes several improvements over the aforementioned version through the mathematical formulation, the treatment of FSR photons, the single muon response modelling, the use of tuning parameters to improve the model (to be explained in Section 7.3.4), and the treatment of systematics, as well as technical logistics throughout.

7.3.1 Per-Event Response PDF

In the PER model, the PDF for signal events is uniquely determined for each event based on the lepton kinematics, and the hypothesized Higgs parameters, taking the form $P(m_{4\ell}^{\text{fsr}} | m_H, \Gamma_H, \text{kinematics})$.

To reach this expression, note that any given lepton i has a true energy E_i^{true} , and an energy after reconstruction and FSR recovery³ $E_i^{\text{fsr}} \equiv E_i^{\text{reco}} + E^{\text{FSR}}$. The true energy used here is the *born energy* [92] - the lepton's energy at the Z boson decay vertex and, importantly, before any FSR. Using these quantities, we can define the lepton energy response as

$$x_i = \frac{E_i^{\text{fsr}} - E_i^{\text{true}}}{E_i^{\text{fsr}}}, \quad (75)$$

where x_i is distributed according to some general distribution $P(x_i)$ that will be described in Section 7.3.2. Some additional remarks about the definition of x_i are found in Appendix B.2.

Next, examine the four-lepton invariant mass at the born level:

$$\begin{aligned} (m_{4\ell}^{\text{true}})^2 &= (p_{4\ell}^{\text{true}})^2, \\ &= \left(\sum_{i=1, \dots, N_{\text{part}}} p_i^{\text{true}} \right)^2, \\ &= \sum_i (p_i^{\text{true}})^2 + \sum_{i,j \neq i} p_i^{\text{true}} \cdot p_j^{\text{true}}, \end{aligned} \quad (76)$$

where $N_{\text{part}} = 4$ for the four-lepton system⁴. Since the individual particles are muons

³The FSR photon recovered in Section 5.8 is matched to a particular lepton using the ΔR separation

⁴Except for the unimportant labelling of $m_{4\ell}$, the formulation is completely general to any number

and electrons, $(p_i^{\text{true}})^2 = m_e^2, m_\mu^2$ can be neglected compared to the much larger terms. Continuing,

$$\begin{aligned} (m_{4\ell}^{\text{true}})^2 &= 2 \sum_{i,j>i} p_i^{\text{true}} \cdot p_j^{\text{true}}, \\ &= 2 \sum_{i,j>i} E_i^{\text{true}} E_j^{\text{true}} (1 - c_{i,j}), \end{aligned} \quad (77)$$

where $c_{i,j}$ is the cosine of the relative angle between the leptons, assumed to be exactly measured in reconstruction⁵, such that $c_{i,j} \equiv c_{i,j}^{\text{fsr}} = c_{i,j}^{\text{true}}$. Next rearrange (75) to get $E_i^{\text{true}} = E_i^{\text{fsr}} (x_i - 1)$, which when substituted into (77) gives

$$\begin{aligned} (m_{4\ell}^{\text{true}})^2 &= 2 \sum_{i,j>i} E_i^{\text{fsr}} (x_i - 1) E_j^{\text{fsr}} (x_j - 1) (1 - c_{i,j}), \\ &= 2 \sum_{i,j>i} E_i^{\text{fsr}} E_j^{\text{fsr}} (1 - x_i - x_j + x_i x_j) (1 - c_{i,j}), \\ &= 2 \sum_{i,j>i} E_i^{\text{fsr}} E_j^{\text{fsr}} (1 - c_{i,j}) - 2 \sum_{i,j>i} E_i^{\text{fsr}} E_j^{\text{fsr}} (x_i + x_j) (1 - c_{i,j}), \\ (m_{4\ell}^{\text{fsr}})^2 - (m_{4\ell}^{\text{true}})^2 &= 2 \sum_{i,j>i} E_i^{\text{fsr}} E_j^{\text{fsr}} (1 - c_{i,j}) (x_i + x_j), \\ &= \sum_i \left(\sum_{j \neq i} (m_{ij}^{\text{fsr}})^2 \right) x_i, \end{aligned} \quad (78)$$

where in the third line we have dropped the term with a factor $x_i x_j$ which is negligible, and in the final line m_{ij} is the di-lepton mass of leptons i and j together. Finally if we divide both sides by $(m_{4\ell}^{\text{fsr}} + m_{4\ell}^{\text{true}}) \approx 2m_{4\ell}^{\text{fsr}}$ then⁶ we reach a neat expression

$$m_{4\ell}^{\text{fsr}} - m_{4\ell}^{\text{true}} = \frac{1}{2m_{4\ell}^{\text{fsr}}} \sum_i \left(\sum_{j \neq i} (m_{ij}^{\text{fsr}})^2 \right) x_i. \quad (79)$$

If we consider $(m_{4\ell}^{\text{fsr}} - m_{4\ell}^{\text{true}})$ to be a random variable, we are very interested in its PDF, which can be calculated using the properties of random variables to be

$$P(m_{4\ell}^{\text{fsr}} - m_{4\ell}^{\text{true}}) = \frac{1}{2m_{4\ell}^{\text{fsr}}} \prod_i \left(\sum_{j \neq i} (m_{ij}^{\text{fsr}})^2 \right) P(x_i), \quad (80)$$

of particles

⁵This is an imperfect but low-impact assumption: angular uncertainties contribute a much smaller effect to multi-lepton invariant masses than the uncertainties on individual lepton energies/momenta.

⁶To account for the difference between reconstruction and true, this factor is not exactly 2, but is determined on an event-by-event basis based on the mean of the predicted distribution $P((m_{4\ell}^{\text{fsr}})^2 - (m_{4\ell}^{\text{true}})^2)$.

where $\prod$ here represents *repeated convolution*, and not repeated multiplication. This is the response distribution of any particular event, if the $P(x_i)$ are known.

Considering the invariant mass again as a random variable, for each event the trivial expression $m_{4\ell}^{\text{fsr}} = (m_{4\ell}^{\text{fsr}} - m_{4\ell}^{\text{true}}) + m_{4\ell}^{\text{true}}$ leads to the expression

$$P(m_{4\ell}^{\text{fsr}}) = P(m_{4\ell}^{\text{fsr}} - m_{4\ell}^{\text{true}}) * P(m_{4\ell}^{\text{true}}). \quad (81)$$

By substituting (80), we have the expression

$$P(m_{4\ell}^{\text{fsr}} | \text{kinematics}) = \left[\frac{1}{2m_{4\ell}^{\text{fsr}}} \prod_i \left(\sum_{j \neq i} (m_{ij}^{\text{fsr}})^2 \right) P(x_i) \right] * P(m_{4\ell}^{\text{true}}) \quad (82)$$

This derivation shows that if the lepton energy responses are accurately known, *no matter the form of their distributions*, then the correct PDF for any given event can be calculated exactly, for any given signal hypothesis of $P(m_{4\ell}^{\text{true}})$ and its corresponding parameters. In this thesis, the parameters are m_H and Γ_H , as will be seen shortly. The following section describes how we model $P(x_i)$ as a Gaussian-mixture for any given lepton kinematics.

7.3.2 Describing Single Lepton Energy Response

The PDF of each lepton's energy response, $P(x_i)$, is described using a weighted sum of three Gaussian distributions. The number of Gaussians is chosen as the smallest number sufficient to describe the energy response x_i in any kinematic region. Note that using Gaussians is not obligatory - any parameterization of $P(x_i)$ which accurately describes the distribution would do fine. The choice of a Gaussian mixture for the PDF is convenient because

- The Gaussians can usually (but not necessarily) be associated with particular effects: a Gaussian resolution core, and radiative/bremsstrahlung tails.
- The convolutions necessitated by (80) are always analytical.
- The response distribution of (80) will also be a sum of weighted Gaussians.
- The response distribution can be simplified in an efficient way - a procedure to be described in Section 7.3.3.
- The propagation of systematic uncertainties is more easily understood, and technically implemented.

The response PDF for each lepton is then parameterized using 8 free parameters:

$$P(x_i) = \sum_{g=1,2,3} w_g G(x_i; \mu_g, \sigma_g), \quad (83)$$

where the weights are normalized $\sum_{g=1,2,3} w_g = 1$ (reducing the free parameters from 9 to 8). These Gaussian parameters (w_g, μ_g, σ_g) describing the lepton energy response in general depend on the species and kinematics of the lepton.

To construct a model of the lepton energy response, all available signal and $q\bar{q} \rightarrow ZZ^{(*)}$ MC simulated samples are used. Only leptons from events passing the full selection and entering the $m_{4\ell}$ fit range are used. These leptons are separated by species (muon, or electron), and divided into bins of (Z1 or Z2) x (p_T bins) x ($|\eta|$ bins). The p_T and $|\eta|$ bins are chosen based on available MC statistics (finer granularity would always be better than coarser), expected resolution behaviour as quantified by ATLAS Combined Performance (CP) groups using Data and MC comparisons [79][4] (rapidly varying resolution requires more bins, slowly varying resolution requires fewer), and detector geometry (transitions in detector layout as a

function of $|\eta|$ normally define bin boundaries). The bins used are as follows:

- muons
 - Z bins: 1,2
 - p_T bin boundaries (GeV): 5. , 6.5, 7.5 , 8.5, 9.5, 10.5, 11.6, 12.7, 13.8, 15.0, 16.2, 17.4, 18.6, 19.8, 21., 22.3, 24 , 26 , 28.1, 30., 31.5, 33, 34.5, 36, 38, 40, 42, 44, 46, 49, 52, 56, 60, 64, 68, 72, 77, 82, 87, 92, 98, 104, 110, 118, 126, 134, 144, 154, 166, 178, 180, 200, 225, 250, 400, 1000
 - $|\eta|$ bin boundaries: 0., 0.1, 0.22, 0.34, 0.46, 0.58, 0.70, 0.82, 0.94, 1.06, 1.18, 1.30, 1.42, 1.54, 1.68, 1.80, 1.92, 2.04, 2.19, 2.34, 2.50, 2.70
- electrons
 - Z bins: 1,2
 - p_T bin boundaries (GeV): 7., 8.6 , 10.0 , 11.0 , 12.2 , 13.4 , 14.6 , 15.8, 17, 18.4, 19.8, 21.2, 22.6 , 24 , 25.6 , 27.2 , 28.8 , 30.6 , 32.3 , 34.2 , 36 , 38.2 , 40.4 , 42.6 , 44.8 , 47, 49.4 , 51.8, 55, 58, 61, 65, 69, 74, 79, 85, 91, 97, 103, 110, 118, 126, 135, 144, 154, 164, 176, 188, 200, 225, 250, 1000
 - $|\eta|$ bin boundaries: 0., 0.1, 0.15, 0.2, 0.4, 0.6, 0.8, 1.00, 1.20, 1.37, 1.52, 1.6, 1.75, 1.90, 2.0, 2.1, 2.3, 2.4, 2.5

As a result, the muons are subdivided into $2 \times 55 \times 21 = 2310$ bins. The electrons are subdivided into $2 \times 51 \times 18 = 1836$ bins. In each of these bins the distribution of x_i as defined in (75) is fit using an unbinned maximum likelihood fit⁷ with the PDF of (83). The parameters are stored in a database for later retrieval. It's not practical to include the 4146 fits (in fact, double this number, as Section 7.3.5 will explain) in this text, but a few samples bins are included in Figure 74 for muons, and in Figure 75 for electrons. Additional examples are included in Appendix B.3.

Note that the response distribution for muons is nearly Gaussian (meaning w_1 is large), and the additional Gaussians are only needed to account for small non-Gaussian behaviour in the tails, whereas the electrons have more significant asymmetries which requires larger w_2 , w_3 to describe. To gain an overall picture of the performance of the fits, the distribution of χ^2/n , defined as in [93], from each of the

⁷The fit is re-executed as many as 30 times with randomly varying seeds in order to escape local minima and find the best fit possible in this bin




Figure 74: Example fits of unconstrained muon energy response (1/2 - see Appendix B.3 for more). The distributions are built from MC simulation (black points) fit using a sum of three Gaussians (blue curve). Each plot shows the same fit in linear and log scale. The kinematic bins and the fit parameters are printed on each plot.

fits is plotted in Figure 76. There is no analytical prediction against which to measure this distribution, because it combines the χ^2/n values from many fits with different values of n ⁸, so it can only be examined qualitatively. Lower values indicate a closer

⁸The χ^2 is only defined for non-empty bins, so the number of bins used to calculate it can vary from fit to fit




Figure 75: Example fits of unconstrained electron energy response (1/2 - see Appendix B.3 for more). The distributions are built from MC simulation (black points) fit using a sum of three Gaussians (blue curve). Each plot shows the same fit in linear and log scale. The kinematic bins and the fit parameters are printed on each plot.

agreement of model and data, and a value of 1 is expected from statistical fluctuation alone, on average.

It can be interesting to examine trends in the model, in particular by looking at σ_1 , the width of the leading Gaussian, or most simply thought of as the scalar energy




Figure 76: Distribution of χ^2/n from unconstrained lepton response model fits.

resolution of the lepton. The dependence of σ_1 on p_T is plotted⁹ in all Z and $|\eta|$ bins in Figure 77. These plots exhibit the expected behaviour for both species of lepton: for electrons, calorimeter energy resolution is proportional to $\sqrt{E}$ [94], so the fractional resolution *falls* as $1/\sqrt{E}$, which is nicely replicated in the plot. For muons, since the momentum is measured according to the bending of the muon track, the fractional resolution *grows* with E as the track becomes more straight. Some deviations from this behaviour exist because the leading (sub-leading and sub-sub-leading) Gaussian(s) is (are) not forced to represent the core (tails) of the resolution distribution - they are free to find the best representation of $P(x_i)$ given the MC simulated distribution in that bin. The behaviour of the other parameters of the Gaussian sum is presented in Appendix B.4, as they are less instructive than σ_1 .

As further validation, the lepton resolution as described by the PER model was compared with that from the ATLAS CP groups. Note that this can never be a perfect apples-to-apples comparison: the PER model is by definition a *response*, not a resolution. To make a comparison, the PER model was reduced to a single number by using only the leading-Gaussian width. Then the two models of resolution were compared by examining the ratio of σ_1/σ_{CP} for many leptons from simulated signal events. This comparison is shown in Figure 78 (note: on the plot σ_1 is labelled as σ_{PERg1}). As expected, the two methods do not agree perfectly, but they do show indications of providing complimentary descriptions of the lepton resolution: At higher

⁹Note that these plots don't include the full p_T range covered by the model; they include the range where most $H \rightarrow 4\ell$ leptons exist, and where the p_T dependence is most interesting.




Figure 77: Behaviour of parameter σ_1 in lepton response model as a function of p_T in the $|\eta|$ bins for muons and electrons, split by Z1 and Z2. The parameters take on discrete values, but a piece-wise linear interpolation is drawn to emphasize the trend.

values of p_T , where the electron response becomes more Gaussian-like, the two resolution models converge. For muons the same convergence exists, but there is more $|\eta|$ -dependence, so it is less distinct.

7.3.3 PDF simplification

As the lepton response x_i is dimensionless by definition, it is important to see how the factors of mass found in (82) enter the expression to create a $m_{4\ell}$ response. Using the properties of random variables we can see that




Figure 78: Comparison of lepton energy resolution as modelled in PER method, with the resolution as modelled by CP groups. The markers show the average ratio of the PER leading-Gaussian resolution (labelled σ_{PERg1}), with the resolution as parameterized by the ATLAS CP group (labelled σ_{CP}) as a function of p_T , in slices of $|\eta|$. The vertical errors bars represent the standard deviation of the same ratio, for all leptons in the same bin. The leptons come from ggF and VBF Higgs signal simulated samples at various mass points.

$$\begin{aligned}
 \text{if } P(x_i) &= \sum_g w_g G(x_i; \mu_g, \sigma_g), \\
 \text{then } P(Ax_i) &= \sum_g w_g G(Ax_i; A\mu_g, A\sigma_g),
 \end{aligned} \tag{84}$$

where A can be any scalar number, but is often the factor $(\sum_{j \neq i} (m_{ij}^{\text{fsr}})^2)$ from (79). In this expression now the parameters of the Gaussians carry the same dimensions as A (e.g. mass). We can also see what happens when these distributions are convolved, as instructed by (80):

$$\begin{aligned}
 P(Ax_i + Bx_j) &= \left[\sum_g w_g^i G(Ax_i; A\mu_g^i, A\sigma_g^i) \right] * \left[\sum_{g'} w_{g'}^j G(Bx_j; B\mu_{g'}^j, B\sigma_{g'}^j) \right], \\
 &= \sum_{g,g'} w_g^i w_{g'}^j [G(Ax_i; A\mu_g^i, A\sigma_g^i) * G(Bx_j; B\mu_{g'}^j, B\sigma_{g'}^j)], \\
 &= \sum_{g,g'} w_g^i w_{g'}^j G\left(Ax_i + Bx_j; A\mu_g^i + B\mu_{g'}^j, \sqrt{(A\sigma_g^i)^2 + (B\sigma_{g'}^j)^2}\right).
 \end{aligned} \tag{85}$$

From this expression we see that when the input lepton energy response model is a Gaussian mixture, so too will be the final invariant mass response.

A consequence of (85) is that when 3 weighted Gaussians are used to describe $P(x_i)$, then a sum of $3^{N_{\text{part}}} = 81$ weighted Gaussians will be needed to describe the $m_{4\ell}$ response. This number is evidently quite unwieldy and must be reduced. The reduction follows the same procedure used in electron reconstruction using a Gaussian Sum Filter [72]. The 81 Gaussians are compared using the Kullback-Leibler distance [73] defined for a pair of Gaussians (call them a, b) by

$$D_{\text{KL}} = \frac{(\sigma_a^2 - \sigma_b^2)^2 + (\mu_a - \mu_b)^2 (\sigma_a^2 + \sigma_b^2)}{\sigma_a^2 \sigma_b^2}, \quad (86)$$

and the closest pair (smallest D_{KL}) is identified. This pair is replaced in the sum by a single Gaussian with parameters

$$\begin{aligned} w' &= w_a + w_b, \\ \mu' &= \frac{w_a \mu_a + w_b \mu_b}{w_a + w_b}, \\ \sigma'^2 &= \frac{w_a \sigma_a^2 + w_b \sigma_b^2}{w_a + w_b} + \frac{w_a w_b}{(w_a + w_b)^2} (\mu_a - \mu_b)^2. \end{aligned} \quad (87)$$

The algorithm is repeated until the number of remaining Gaussians reaches the target, $N_G = 4$. This procedure by design preserves the 1st moment (mean) and 2nd moment (variance) of the original distribution. Higher order moments are not preserved, but by retaining $N_G = 4$ Gaussians, most of the unique shape of any individual event's $m_{4\ell}$ response is preserved. An example of this procedure for a single MC event is shown in Figure 79, where the Gaussian sum representing the mass response for this event is shown at various stages of the collapse. In practice the procedure stops at $N_G = 4$, though here it is shown down to $N_G = 1$.

Returning to (82), the PER signal PDF can be written down as the convolution of this Gaussian mixture (describing the response), with the true distribution for any given signal hypothesis,

$$P(m_{4\ell}^{\text{fsr}} | \text{kinematics}) = \left[\sum_{g=1}^{N_G} w_g G(m_{4\ell}^{\text{fsr}} - m_{4\ell}^{\text{true}}; \mu_g; \sigma_g) \right] * P(m_{4\ell}^{\text{true}}), \quad (88)$$

where the weighted sum in square brackets and all the Gaussians parameters incorporate the invariant mass factors implicitly (the same ones explicitly shown in (85)).




Figure 79: Demonstration of collapsing procedure on an example MC signal event. The initial 81-Gaussian sum is shown at various stages of collapse down to a single Gaussian. The x-axis is $m_{4\ell}^{\text{fsr}} - m_{4\ell}^{\text{true}}$ in units of GeV, left unlabelled to minimize clutter. Some Gaussians have weights and sigmas such that they are below the plotted y-axis range.

The signal hypothesis used in this thesis is of a Higgs boson with mass m_H , and natural width Γ_H . Its true mass will be distributed, as described in Section 2.3.2.2, by a BW distribution,

$$P(m_{4\ell}^{\text{true}}) = \text{BW}_\epsilon(m_{4\ell}^{\text{true}}; m_H, \Gamma_H), \quad (89)$$

where the BW_ϵ indicates the BW is corrected by an efficiency function $\epsilon(m_{4\ell}^{\text{true}})$ to account for non-uniform detector efficiency across the width of true lineshape when Γ_H is large. For small widths, as in the SM, the efficiency is sufficiently flat that this

factor has no effect (even up to $\Gamma_H = 100\Gamma_H^{\text{SM}}$, the effect is very small). Conveniently, this factor is exactly equivalent to the acceptance polynomial function of Section 7.1,

$$\begin{aligned} \text{BW}_\epsilon(m_{4\ell}^{\text{true}}; m_H, \Gamma_H) &= \epsilon(m_{4\ell}^{\text{true}}) \text{BW}(m_{4\ell}^{\text{true}}; m_H, \Gamma_H), \\ &= A_c(m_{4\ell}^{\text{true}}) \text{BW}(m_{4\ell}^{\text{true}}; m_H, \Gamma_H). \end{aligned} \quad (90)$$

The description of $P(m_{4\ell}^{\text{true}})$ as a BW distribution assumes by definition that we select the correct four leptons in an $H \rightarrow 4\ell$ signal event. For ggF and VBF production, events passing the event selection and the mass window used in this analysis have the correct four leptons 99.97% of the time. For VH and ttH production modes with additional leptons, the correct four leptons are selected in 97.59% of the events in the same mass window (most events with a wrong lepton selected end up outside the mass window). When these production modes are weighted by their SM cross section the non-resonant contribution is extremely small.

The PER signal PDF as it is used in this measurement can be written down by combining (88) and (89),

$$P(m_{4\ell}^{\text{fsr}} | m_H, \Gamma_H, \text{kinematics}) = \left[\sum_{g=1}^{N_G} w_g G(m_{4\ell}^{\text{fsr}} - m_{4\ell}^{\text{true}}; \mu_g; \sigma_g) \right] * \text{BW}_\epsilon(m_{4\ell}^{\text{true}}; m_H, \Gamma_H). \quad (91)$$

This final expression contains one single observable ($m_{4\ell}^{\text{fsr}}$), several conditional observables (kinematics, which are transformed into w_g, μ_g, σ_g , combined as $\vec{g}$) whose distribution is not specified in advance but must be observed in data, and two parameters of interest (m_H, Γ_H) that can be fit.

7.3.4 Gaussian sum parameter transformation

While the method described above very closely predicts the lineshape of each event, the final result can be improved by tuning on the $m_{4\ell}$ distribution from signal MC. This transformation accounts for effects introduced by the choice of binning in the lepton response model, for MC statistics availability in each of these bins and their effect on the Gaussian parameters, and for any small shape detail lost during the collapsing procedure. This transformation uses one tuning parameter for each parameter of the Gaussian sum in (91), which are transformed as follows:

$$\begin{aligned}
w_i &\rightarrow n_i w_i, \\
\sigma_i &\rightarrow r_i \sigma_i, \\
\mu_i &\rightarrow \mu_i + s_i + (1 - r_i)(\mu_i - \mu_1),
\end{aligned} \tag{92}$$

where r denotes a modification on the resolution, s denotes a modification on the scale, n denotes a modification on the relative contributions of each of the Gaussians. Note that the r_i are used both to scale the resolution of the Gaussians, and to shift them closer together (reducing the RMS of their combined distribution). The parameters are fit simultaneously for all events in $m_H = 125$ GeV ggF MC sample, separately for each category, and recorded in Table 13. No improvement was found for the $2e2\mu$ category, and so no transformation was applied. Note that values of $s_i = 0, r_i = 1, n_i = 1$ indicate no effect; the transformation overall has a very small (but none-the-less valuable) effect.

Table 13: Transformation parameters used to account for slight discrepancies in the invariant mass response, as used in the unconstrained PER model. The parameters are derived by a simultaneous fit on simulated ggF event with $m_H = 125$ GeV.

parameter	4μ	$4e$	$2\mu 2e$	$2e2\mu$
r_1	0.963	0.997	0.948	1
r_2	0.953	0.984	1.001	1
r_3	0.994	0.933	1.051	1
r_4	1.110	1.995	1.154	1
s_1	0.060	0.148	0.101	0
s_2	-0.252	0.043	-0.040	0
s_3	-1.214	-0.611	-0.836	0
s_4	-1.140	-2.997	-2.967	0
n_1	0.998	0.983	1.012	1
n_2	1.001	1.000	0.981	1
n_3	1.026	1.117	0.988	1
n_4	1.155	1.177	0.794	1

Figure 80 shows the effect of this transformation on the $m_H = 125$ GeV ggF MC sample (the same one used to derive the parameters). The transformation yields a

distinct improvement over the uncorrected case. Values of χ^2/n are printed on each plot for comparison. For a robust fit result, these parameters must be independent of m_H and of the Higgs production mode - by definition they exist only to correct minor imperfections in the response model, something wholly detector-based. To verify this, the same set of transformation parameters are tested on the $m_H = 124$ GeV and $m_H = 126$ GeV ggF MC samples, shown in Figure 81 and Figure 82 respectively, as well as the $m_H = 125$ GeV VBF MC sample in Figure 83. In each case, the same set of transformation parameters perform very well at improving model agreement.

7.3.5 Response after Z Mass Constraint

As demonstrated in Section 5.10, the introduction of the Z Mass Constraint reduces the resolution of four-lepton invariant mass by about 14%, and consequently makes $m_{4\ell}^{\text{cons}}$ more sensitive to m_H than $m_{4\ell}^{\text{fsr}}$ by an equal amount. Inclusion of this benefit is evidently worthwhile for a precision measurement of m_H .

If the constrained four-momenta are considered as an alternative estimate of the lepton kinematics¹⁰, the model described in Section 7.3.2, and Section 7.3.3 can be modified to describe these *new* responses simply by substituting $E_i^{\text{fsr}} \rightarrow E_i^{\text{cons}}$, $m_{4\ell}^{\text{fsr}} \rightarrow m_{4\ell}^{\text{cons}}$. With this substitution, the mathematics which predict the detector response of $m_{4\ell}^{\text{cons}}$ are unchanged as compared to $m_{4\ell}^{\text{fsr}}$. The fits of the lepton response (those shown in Figures 74,75) are altered by this new definition, though visually they are not substantially different; for this reason they are recorded in Appendix B.3. Both versions of the response model serve valuable purposes: the unconstrained PER is used as a proof-of-principle and cross check, and the constrained PER is used in the final m_H measurement.

An additional caveat when using the ZMC is that the constraint introduces a relationship between the two leptons which form the Z1 pair. As shown in Section 5.9, the most likely values m_{Z1} on an event-by-event basis form a distribution significantly narrower than the Z lineshape. As a result, the energies E_1^{cons} , E_2^{cons} are correlated in order to make a di-lepton invariant mass very close to $m_Z^{\text{PDG}} = 91.19$ GeV. This introduces a potential pitfall in the error propagation of the PER method, particularly (80), which assumes the lepton energy responses are *uncorrelated*; using this

¹⁰Some revolutionary calibration procedure which improves the invariant mass resolution by 18% would be treated the same way.




Figure 80: Demonstration of effect of transformation of Gaussian parameters on the unconstrained PER model on the $m_H = 125$ GeV ggF MC sample, separated by category. The transformation parameters, defined in (92), are derived by a simultaneous fit on this sample. The black points show the MC simulation, the blue line shows the prediction from the untransformed per-event response model, and the red line shows the prediction from the transformed per-event response model. At the bottom the ratio of the MC simulation over each prediction is shown.




Figure 81: Demonstration of effect of transformation of Gaussian parameters on the unconstrained PER model on the $m_H = 124$ GeV ggF MC sample.




Figure 82: Demonstration of effect of transformation of Gaussian parameters on the unconstrained PER model on the $m_H = 126$ GeV ggF MC sample.




Figure 83: Demonstration of effect of transformation of Gaussian parameters on the unconstrained PER model on the $m_H = 125$ GeV VBF MC sample.

relationship with *correlated* random variables is incorrect¹¹. To account for this effect in the PER method, the transformation parameters in Section 7.3.4 are re-interpreted to not only account for small discrepancies in the model, but also to absorb the effect of the relationship between x_1, x_2 in the error propagation. The approach is to say that while the co-dependent uncertainties can't be propagated analytically, their effect can be observed numerically in simulation, and incorporated parametrically. In this way we can still benefit from the increased mass resolution of the ZMC, together with the PER method. As in the unconstrained case, the transformation parameters are derived by a simultaneous fit on simulated ggF event with $m_H = 125$ GeV. The parameters are given in Table 14.

Table 14: Transformation parameters used to account for the effect of the ZMC on the invariant mass response, as used in the constrained PER model. The parameters are derived by a simultaneous fit on simulated ggF event with $m_H = 125$ GeV.

parameter	4μ	$4e$	$2\mu 2e$	$2e 2\mu$
r_1	0.850	0.917	0.896	0.826
r_2	0.949	0.978	0.955	1.023
r_3	0.956	0.942	0.952	0.964
r_4	1.011	1.998	1.078	1.070
s_1	0.079	0.185	0.115	-0.010
s_2	-0.081	0.059	0.068	0.131
s_3	-0.815	-1.021	-0.223	-0.453
s_4	-0.592	-2.990	-0.795	-1.974
n_1	0.906	0.996	1.039	1.017
n_2	1.047	1.041	1.071	0.964
n_3	1.051	1.016	1.299	1.023
n_4	1.013	0.896	1.105	1.201

Note in particular the parameter r_1 , which transforms the core width of the resolution distribution¹². The values for each category (e.g. a 15% reduction in 4μ)

¹¹An obvious solution may appear to be to introduce a correlation ρ_{12} between x_1 and x_2 , and use simple correlated Gaussian error propagation: $\sigma'^2 = \sqrt{\sigma_1^2 + \sigma_2^2 - 2\rho\sigma_1\sigma_2}$ in (85), but the correlation is not a simple one, and this expression is inadequate to describe the relationship between x_1 and x_2 .

¹²The Gaussian-parameter indices are ordered by largest weight. Index 1 corresponds to the

are consistent with those found by fitting the core of the distributions alone, summarized in Table 9. Results of this transformation on the $m_H = 125$ GeV ggF MC sample can be seen in Figure 84. The same transformation parameters applied to the $m_H = 124$ GeV ggF, $m_H = 126$ GeV ggF, and $m_H = 125$ GeV VBF MC samples can be seen in Figure 85, Figure 86, and Figure 87, respectively, justifying that these parameters correct only the response model, independent of m_H or the production mode.

7.3.6 Distribution of resolutions

While the macro properties of the PER model shown in the preceding sections provide a good measure of its overall performance through agreement with simulation, it can also be interesting to examine individual events. Figure 88 shows the PDF predicted by the constrained PER model for a few events randomly chosen from ggF simulation. The distributions are plotted assuming $m_H = 125$ GeV, with $\Gamma_H = 4.4$ MeV.

Included on the distributions is an effective resolution [95], σ_{eff} , defined as half the smallest width that contains 68.3% of the distribution; this is a simple metric that enables comparison of the generally complex resolution distributions. For a Gaussian distribution, it is equivalent to the Gaussian σ . This quantity can also be used to visualize the spread of resolutions within each category; this is plotted in Figure 89.

The distributions of σ_{eff} exhibit some expected features: The 4μ effective resolution has significantly less variance than the categories containing electrons. This is not surprising when Figure 77 is considered: energy resolution varies much more rapidly for electrons than for muons, leading to larger variance in the final $m_{4\ell}$ resolution (both with and without constraint). The mean of each histogram, printed on the same plots, represents the average resolution in each category. Events with σ_{eff} smaller than this value are better-than-average events, and those with larger σ_{eff} are worse-than-average; consequently, these events become either more or less important in the final m_H fit, respectively.

largest weight, making up the core of the resolution distribution, while index 4 corresponds to the far far tails.




Figure 84: Demonstration of effect of transformation of Gaussian parameters on the constrained PER model on the $m_H = 125$ GeV ggF MC sample, separated by category. The transformation parameters, defined in (92), are derived by a simultaneous fit on this sample. The black points show the MC simulation, the blue line shows the prediction from the untransformed per-event response model, and the red line shows the prediction from the transformed per-event response model. At the bottom the ratio of the MC simulation over each prediction is shown.




Figure 85: Demonstration of effect of transformation of Gaussian parameters on the constrained PER model on the $m_H = 124$ GeV ggF MC sample.




Figure 86: Demonstration of effect of transformation of Gaussian parameters on the constrained PER model on the $m_H = 126$ GeV ggF MC sample.




Figure 87: Demonstration of effect of transformation of Gaussian parameters on the constrained PER model on the $m_H = 125$ GeV VBF MC sample.




Figure 88: Predicted $P(m_{4\ell}^{\text{cons}})$ from the constrained PER model for a sample of simulated ggF events in each category. Printed on the plots are the $m_{4\ell}^{\text{cons}}$ from each event, and the effective resolution, σ_{eff} .




Figure 89: Distribution of σ_{eff} in the constrained PER model from simulated ggF events in each category. Better-measured events have lower values of σ_{eff} , and more importance in the fit. Printed on the plots are the mean and standard deviation of the distributions.

7.4 Choosing a signal model

Both signal models presented in this chapter - the per-event response method, and the template method - are valid techniques for measuring the Higgs boson mass, m_H . It is also quite clear that the PER method is significantly more complex, requires more ingredients, and proportionately more validation. Perhaps less obvious is that the CPU consumption of the PER method is drastically higher than for the template; a consequence of each event carrying its own signal PDF, each of which requires a numerical convolution¹³. Naturally this begs the question: why use the PER method? Or phrased another way, what are the advantages of the PER method which outweigh these potential drawbacks?

There are several advantages to the PER method which clearly justify the complexity of the method, and the CPU requirements:

- The PER provides a first-principles derivation, starting from the lepton kinematics and propagating uncertainties upwards to the $m_{4\ell}$ lineshape. This method makes a clear delineation between detector and physics, something impossible through template-style methods.
- Only the PER method is capable of fitting the Higgs width, Γ_H , since the truth lineshape is included analytically within the PDF.
- Only the PER method is capable of incorporating systematic uncertainties on the detector resolution into the model. Chapter 8 will present the treatment of systematic uncertainties, and elaborate on this point.
- The PER method is as model-independent as possible. While the PER and template methods both use the same simulated samples as input, the template method by definition assumes the kinematic distributions of the sample to be correct (so, implicitly, the SM is correct), while the PER does not. Should the model used in the simulation be incorrect, because either the SM is flawed, or new higher-order corrections alter the model, then the lineshape predicted by the template method will be incorrect, and so the m_H value it measures will

¹³This point is not to be understated: along with the non-inclusion of the ZMC, the CPU consumption was a limiting factor in the implementation of a PER method during Run 1. This drawback has been radically improved by the implementation presented in this thesis: where the the Run 1 implementation requires two numerical convolutions per-event (resolution * FSR * truth), the implementation in this thesis requires only one ((resolution + FSR) * truth). This change, among others, has drastically reduced CPU consideration as a limiting factor from the method, though it is still important.

be subject to change. In contrast, since the PER method predicts the PDF of the observed events based only on the observed lepton kinematics and the ATLAS detector, the m_H value measured is not subject to change based on a different model hypothesis. The only model assumption embedded in the PER is the very safe claim that the true distribution of $H \rightarrow 4\ell$ events follows a Breit-Wigner lineshape, which is a basic property of any unstable particle (see Section 2.3.2.2). A demonstration of this model independence is found in Appendix B.5.

- The statistical power of each event in determining m_H is more closely related to its measurement uncertainty when using the PER method. As Section 7.3.6 demonstrates, the resolution of 4ℓ events reconstructed in ATLAS can vary quite significantly. While the template method separates events by category to control this effect, it is clear that within each category there are still large differences. Because the template method effectively assumes the average resolution within each category, better-than-average events are wasted (in a statistical sense), and worse-than-average events are given too much impact. In the PER method, each event has a statistical impact proportionate to its resolution.

It is these advantages which make the PER method the most powerful tool available for measurement of the Higgs boson resonance properties. A clear demonstration of the power of the method is found in Chapter 9.

Chapter 8

Systematic Uncertainties

Every measurement has an uncertainty composed of statistical and systematic components. The statistical component decreases as more data are added, giving better precision. The systematic component however is an inherent limitation in the measurement strategy, and (mostly) doesn't decrease as more data are added. Systematic uncertainties arise from many sources: the theory inputs used to generate the MC simulation, the lepton energy/momentum calibrations, the scale factors used to correct the simulation to match the data, the available MC statistics¹, as well as the measurement implementation, such as the choice of signal and background models. These uncertainties are the subject of this chapter.

Section 8.1 begins by describing how systematics uncertainties are parameterized in the statistical model. Section 8.2 then describes all the different sources of systematic uncertainty entering the analysis. The following sections then divide the systematics based on where their effect is felt: Section 8.3 describes uncertainties affecting the yield of the signal or backgrounds, while Section 8.4 describes those affecting the energy scale or resolution. Depending on whether the measurement is of m_H , $\sigma_{H \rightarrow 4\ell}$, or Γ_H , these systematics alter the expected resolution of the measurement, the central value of the measurement, or both.

The strategies used to estimate the systematics which follow were developed by a large number of individuals, both within and without the ATLAS $H \rightarrow 4\ell$ analysis group. The systematics are then made available to analyzers like the author through tools and documentation. Some systematics - e.g. the theory systematics - can be used directly in the statistical analysis, while others need a further step of calculation. In the latter case, it may be that ATLAS detector performance groups have provided systematic variations of the (e.g.) energy of an individual lepton. For any

¹This is an example of a “statistical systematic” - systematic uncertainties which could be improved by increased data or MC statistics, although their effect on the measurement is indirect.

such systematic which impacts the event selection², the full selection was repeated with the systematically varied quantities set to their nominal and $\pm 1\sigma$ values. The event selection was run on all the MC simulated signal and background samples to evaluate the impact of the systematic uncertainty associated to each sample. The author’s direct involvement appears at the final step: evaluating the impact on the $m_{4\ell}$ distribution of candidate events, and thus the impact on m_H , $\sigma_{H \rightarrow 4\ell}$, and Γ_H measurements.

8.1 Parameterization of systematic uncertainties

The fluctuation of each systematic uncertainty is described using a nuisance parameter (NP), and the set of NPs is denoted as $\boldsymbol{\theta}$. These NPs effectively normalize the effect of every systematic to the same scale, so the effect of two systematics, e.g. A with a $\pm 1\%$ effect and B with a $\pm 2\%$ effect, are both mapped to the same NP value: $\theta_{A(B)} = \pm 1$. For the same reason, this definition normalizes asymmetric systematics, so that the effect of C , $^{+3\%}_{-4\%}$, is also mapped to $\theta_C = \pm 1$. This definition lets the full set of NPs, $\boldsymbol{\theta}$, be constrained by symmetric Normal distributions to their nominal values, and allows the fluctuations of each to be more easily interpreted and compared. Two systematics which fit to the same value are being “pulled” equally hard by the data, although they may have wildly different effects. How the nuisance parameters and their constraints enter the statistical model used for Higgs property measurement will be shown in Chapter 9.

In practice, the systematics are all implemented using the class `FlexibleInterpVar` (FIV). This class takes as input a nominal and $\pm 1\sigma$ value of a systematic A , and produces a continuous function $f(\theta_A)$ to represent the effect of the systematic. Inside (outside) the given range of $\pm 1\sigma$, the class interpolates (extrapolates) piece-wise-linearly to determine the effect of the systematic. Figure 90 shows an example of two functions built in this way. Better still, the class accepts as input the nominal and $\pm 1\sigma$ values of multiple systematics simultaneously, and so builds a function $f(\boldsymbol{\theta})$ incorporating the effect of many systematics in an efficient manner.

Note that although two different effects are shown in this figure, they could (but don’t have to) correspond to the *same* NP: some given source of uncertainty, A , could

²In this example, if an energy scale uncertainty variation moves a lepton across a p_T cut, the entire selection algorithm for this event will be altered.




Figure 90: Example of the continuous function created by `FlexibleInterpVar` to model the effect of a given systematic which is specified at its nominal and $\pm 1\sigma$ values. Two examples are shown: the blue line shows a $\begin{smallmatrix} +2\% \\ -1\% \end{smallmatrix}$ effect, while the red line shows a much smaller $\begin{smallmatrix} -0.2\% \\ +0.5\% \end{smallmatrix}$ effect.

have two separate effects on two different aspects of the analysis. Furthermore, as illustrated in this figure, $\theta = +1$ does not always correspond to a positive systematic effect, or vice-versa with $\theta = -1$. An example where both these aspects come into play is in the electron resolution: a $+1\sigma$ variation of the electron resolution systematic, corresponding to $\theta_{\text{el. res.}} = +1$, will *increase* the detector resolution, and simultaneously *decrease* the expected yield in the $4e$ category, because poor resolution means fewer events pass the selection and mass window cut. Effects of this nature will be seen later in this chapter.

8.2 Sources of systematic uncertainty

This section is divided according to the source of systematic uncertainty. Section 8.2.1 deals with the theory systematics, Section 8.2.2 deals with the experimental systematics associated to detector performance, and Section 8.2.3 deals with analysis systematics associated to the measurement strategy. In each section several sources of systematic uncertainty are presented according to their NP names. More detail is given to describing the systematic uncertainties which have larger impacts.

8.2.1 Theory systematics

The theory uncertainties in this thesis primarily pertain to the rate of processes, and thus their respective yields. In reality they may also affect the shape of distributions, but these effects are considered to be negligible in the context of the other uncertainties entering the analysis.

The theoretical calculation of the rate of any process, via its cross section, inherently has some theory uncertainty. Because cross sections are calculated perturbatively, if the MC generator includes some process calculated only at LO, then an uncertainty exists because of the missing NLO components; if it is calculated at NLO, then uncertainty still exists from the NNLO component, and so on. Furthermore, these cross section calculations require knowledge of the parton-density function within the colliding protons - it is evidently important to know the energies of colliding gluons in ggF-Higgs production, or of quarks in $q\bar{q} \rightarrow ZZ^{(*)}$ production. On top of these, there are uncertainties associated both to the energy scale of the process from running couplings in QCD and QED, and from the scale at which partons shower. These uncertainties will appear both for the signal [21], and the backgrounds [96].

The full list of theory uncertainties included in the analysis is as follows.

pdf_ggH Parton-density function uncertainties for ggH signal events from [21]. As ggH events are expected to dominate, theory systematics on other production modes are neglected.

br_H4l Uncertainty on calculation of $\text{BR}(H \rightarrow ZZ^{(*)} \rightarrow 4\ell)$.

Signal_comp Uncertainty on signal composition between different production modes as it pertains to varying detector acceptance for $H \rightarrow 4\ell$ events.

QCDscale_VV QCD scale uncertainties for $q\bar{q} \rightarrow ZZ^{(*)}$.

pdf_qq Parton-density function uncertainties for $q\bar{q} \rightarrow ZZ^{(*)}$.

HOEW_QCD_syst High-order electroweak and QCD corrections for $q\bar{q} \rightarrow ZZ^{(*)}$.

HOEW_syst High-order electroweak corrections for $q\bar{q} \rightarrow ZZ^{(*)}$.

QCDscale_ggVV QCD scale uncertainties for $gg \rightarrow ZZ^{(*)}$.

pdf_gg Parton-density function uncertainties for $gg \rightarrow ZZ^{(*)}$.

scale_VVV Parton-density function and scale uncertainties for $VVV, t\bar{t} + Z$.

8.2.2 Experimental systematics

In addition to the theoretical systematics, several experimental systematics enter into the analysis. Most of these systematics come from aspects of the muon and electron reconstruction and performance - their effects on each lepton are provided by the ATLAS detector performance groups (for muons and electrons, see [4] and [79] respectively), and propagated to the $H \rightarrow 4\ell$ system by the method described above. Many of the systematics are described in groups, labelled here using a wildcard (e.g. `basename_*`), as the details of each are beyond the scope of this section.

EG_SCALE_* Electron energy scale uncertainty. The calibration of the EM clusters used to measure electron energy is derived using $Z \rightarrow e^\pm e^\mp$ events, and precision knowledge of the Z pole from the LEP experiments. This uncertainty, under the name `EG_SCALE_ZEESYST`, dominates the electron energy scale uncertainties. The other sources come from material in the ATLAS detector, individual component gain and calibration, detector pedestal correction, and various other minute details.

EG_RESOLUTION_* Electron energy resolution uncertainty. The energy resolution in MC simulation is smeared to match the detector resolution measured in data around the Z pole, and so the dominant component to this uncertainty is called `EG_RESOLUTION_ZSMEARING`. Other contributions come from material and calibration effects.

MUON_SCALE Muon momentum scale; the uncertainty on the momentum scale for all muon types. This scale is calibrated using $Z \rightarrow \mu^\pm \mu^\mp$ data events, and applied to the MC simulation.

MUON_ID Muon momentum resolution as measured in the ID. The muon reconstruction in the ID has an associated resolution in data. The MC simulation often misses the small geometric distortions of the detector, and must be corrected by scale factors to match the data. This is the systematic uncertainty associated to these scale factors.

MUON_MS Muon momentum resolution as measured in the MS. All the same concepts as for `MUON_ID`, but in the MS.

MUON_SAGITTA_* A weak rotational distortion in the ATLAS ID led to a correction of the measured muon momentum in data. Two NPs account for the uncertainty on this correction.

EL_EFF_* Several systematics associated with the electron efficiency. These arise from isolation (***ISO***) and reconstruction (***RECO***) efficiencies, which can in general vary by detector region.

MUON_EFF_* Muon efficiency-related uncertainties.

MUON_ISO_* Muon isolation-related uncertainties.

MUON_TTVA_* Muon track-to-vertex-association related uncertainties. This is used in matching muon tracks to vertices for the combined 4ℓ vertex fit.

PRW_DATASF Pile-up re-weighting uncertainty. The MC simulation is generated with a certain pileup distribution, often before the data are collected. The simulated events are then reweighted from their generated pileup values to match the data distribution. This is the uncertainty in that procedure.

JET_* Systematics related to jet reconstruction, which can have tangential effects on the $H \rightarrow 4\ell$ analysis through particle overlap and isolation.

8.2.3 Analysis Systematics

Several analysis systematics impact the measurements in this thesis. The luminosity uncertainty is central to ATLAS, while the others are all specific to the analysis in this thesis.

lumi Uncertainty only integrated luminosity recorded by ATLAS, 36.1 fb^{-1} . The uncertainty was estimated as 3.2%, and applied to all MC-based yield estimates.

reducible_llee_* Uncertainty associated to data-driven $Z + ee$ background estimate.

reducible_lluu_* Uncertainty associated to data-driven $Z + \mu\mu$ background estimate.

H4lMCStat Uncertainty associated to the available statistics in the MC simulated signal samples. This uncertainty incorporates the systematic uncertainty due to the choice and implementation of the signal model.

8.3 Systematic affecting event yields

This section presents the observed effects of each systematic affecting the signal or background yields. The effects are calculated by the number of weighted events in

the corresponding MC simulated sample passing the full selection with the systematic set to its $\pm 1\sigma$ value, as compared to when it is set at the nominal value,

$$f(\theta = \pm 1) = \frac{N_{\text{sys at } \pm 1\sigma}^{\text{pass,selection}}}{N_{\text{sys at nominal}}^{\text{pass selection}}} . \quad (93)$$

Systematics which affect the event yields are incorporated as FIVs multiplying the expected yield of each source in each category. For example, the $q\bar{q} \rightarrow ZZ^{(*)}$ normalization in the $2\mu 2e$ category with systematics incorporated is

$$N_{\text{qqZZ},2\mu 2e}^{\text{with sys}} = N_{\text{qqZZ},2\mu 2e}^{\text{no sys}} \times f_{\text{qqZZ},2\mu 2e}^{\text{norm.}}(\boldsymbol{\theta}), \quad (94)$$

where the FIV, $f_{\text{qqZZ},2\mu 2e}^{\text{norm.}}$, is labelled as such because a separate function is needed for every parameter affected by systematic uncertainties.

Experimental and analysis systematics have separate effects in each category, while theory systematics have the same effect in each category. The effect of the theory systematics on the normalizations are shown in Figure 91. The effects of experimental and analysis systematics on the $H \rightarrow 4\ell$ signal at $m_H = 125 \text{ GeV}$ are shown in Figure 92, and for the qqZZ background in Figure 93. The ZJets systematics are shown in Figure 94; only analysis systematics are included because they dominate so significantly. For ggZZ and VVV, experimental systematics are negligible compared the theory systematics so they are omitted. The dominant systematics affecting the signal yield in each category are summarized in Table 15.

An additional consideration is that the signal uncertainties may in general depend on the Higgs mass. To account for this, the systematic uncertainties are estimated as described above separately for each MC simulated mass point. At each mass point, a FIV is created as described, and then these FIV are interpolated to intermediate values of m_H , using the BSpline-based interpolation scheme described in Section 7.2.2, where it was used for morphing Template distributions. Two examples of this morphing are shown in Figure 95.




Figure 91: The effect of theory systematics on the signal and backgrounds entering the analysis.




Figure 92: Experimental and analysis systematics affecting the $H \rightarrow 4\ell$ signal yield for $m_H = 125$ GeV. Systematics with effects smaller than 0.2% are omitted.




Figure 93: Experimental and analysis systematics affecting the $q\bar{q} \rightarrow ZZ^{(*)}$ background yield. Systematics with effects smaller than 0.2% are omitted.




Figure 94: Analysis systematics affecting the $Z + \text{Jets}, t\bar{t}, WZ$ background yield. Only analysis systematics, arising from data-driven background estimates of Section 6.3, are shown as experimental systematic have relatively negligible effects.




Figure 95: Example of morphing of systematic uncertainties on the signal yield. The effects are calculated at discrete mass points, and their effect is interpolated to intermediate values of m_H . Shown here are a few of the dominant experimental systematics for muons and electrons in the 4μ and $4e$ categories, respectively. Their effect is demonstrated on the signal acceptance functions in each category.

Table 15: Summary of leading systematics affecting signal yield. Systematics are labelled according to the scheme in Section 8.2. Where multiple systematics contribute under one name, the values from Figure 92 were added in quadrature.

$H \rightarrow 4\ell$ yield systematic [%]				
Systematic	4μ	$4e$	$2\mu 2e$	$2e 2\mu$
EL_EFF_*	0.0	5.8	4.9	1.0
pdf_ggH	3.9	3.9	3.9	3.9
lumi	3.2	3.2	3.2	3.2
MUON_EFF_*	2.3	0.0	1.4	1.0
br_H4l	2.2	2.2	2.2	2.2
PRW_DATASF	0.9	1.7	1.1	0.8
MUON_ISO_*	1.6	0.0	0.4	1.2
Signal_comp	1.3	1.3	1.3	1.3
MUON_TTVA_*	1.0	0.0	0.3	0.7
EG_SCALE_*	0.0	0.5	0.4	0.1
EG_RESOLUTION_*	0.0	0.1	0.1	0.1
MUON_SCALE	0.1	0.0	0.1	0.1

8.4 Systematics affecting scale and resolution

Systematics which affect the shape of the $m_{4\ell}$ distribution, beyond just its normalization, must be treated differently. Systematics such as these are also handled in a different way for signal and background processes.

For the signal, the effect of each source of systematic uncertainty is decomposed into an effect on the invariant mass *scale*, and an effect on the invariant mass *resolution*. To evaluate the effect on the scale, the systematically varied MC samples, after the full event selection, are used. The scale effect is quantified as the change in mean $m_{4\ell}$ value from all events passing the selection in the systematically varied sample, as compared to the nominal sample,

$$f_{\text{scale}}(\theta = \pm 1) = \frac{\langle m_{4\ell}^{\text{sys at } \pm 1\sigma} \rangle}{\langle m_{4\ell}^{\text{sys at nominal}} \rangle}. \quad (95)$$

The same strategy is used to evaluate the effect of each systematic uncertainty on the mass resolution, except the standard deviation (denoted here as SD) of the distribution is used to measure the effect,

$$f_{\text{resolution}}(\theta = \pm 1) = \frac{\text{SD} [m_{4\ell}^{\text{sys at } \pm 1\sigma}]}{\text{SD} [m_{4\ell}^{\text{sys at nominal}}]}. \quad (96)$$

One exception to this strategy comes from the `H4l_MCStat` uncertainty. This uncertainty is only considered to affect the mass scale, where it is estimated by the standard deviation of the mean $m_{4\ell}$ value,

$$\begin{aligned} \sigma_{\bar{x}} &= \sigma_x / \sqrt{N}, \\ \sigma_{MC\ stat} &= \text{SD}[m_{4\ell}] / \sqrt{N_{\text{MC}}}, \end{aligned} \quad (97)$$

for all events passing the event selection in the $m_H = 125 \text{ GeV}$ MC simulated signal samples, with all production modes and all channels combined. By this calculation, the actual value is $\pm 8.1 \text{ MeV}$, which ends up being about a factor of five less than the dominant experimental uncertainties. This point is elaborated on later in Section 9.2.2.

Another exception applies only to the Template method. As brought up in Section 7.4, the construction of the smoothed templates by definition assumes the kinematic distribution of the MC simulated samples is correct. Theoretical uncertainties in general give rise to different kinematic distributions, and therefore different $m_{4\ell}$

lineshapes. This change in the actual $m_{4\ell}$ distribution with respect to the original template introduces a systematic uncertainty. This uncertainty was estimated by varying the relative fraction in MC simulation of ggF, VBF, VH, ttH within experimental uncertainties, and observing how the mean of the $m_{4\ell}$ distribution shifted. By this technique, the Template method is assigned an additional 6 MeV scale systematic not present in the PER method (and so not included in the plots that follow).

The sources of systematic uncertainty which impact the scale are shown in Figure 96 for a $m_H = 125$ GeV signal. Those which impact the resolution are shown in Figure 97. Similarly to the signal normalization, this evaluation is done at all simulated mass points, and interpolated to intermediate m_H values.

The leading systematics affecting energy scale are summarized in Table 16, where they are expressed in MeV assuming $m_{4\ell} \approx 125$ GeV. The leading resolution systematics are summarized in Table 17.

These results exhibit some of the expected features. The dominant mass scale systematics are those coming from the muon (electron) momentum (energy) scale uncertainty. These effects should translate nearly directly to the m_H measurement, which is, in a simplistic sense, just a measurement of the mean $m_{4\ell}$ of signal events. The systematics which dominate the resolution effects are exactly those designed to describe resolution uncertainties, and also the pile-up re-weighting uncertainty. In both cases, the systematics from electrons are larger than for muons, which has always been the case for $H \rightarrow 4\ell$ mass measurement in ATLAS [7].

Table 16: Summary of leading systematics affecting signal mass scale. Systematics are labelled according to the scheme in Section 8.2. Where multiple systematics contribute under one name, the values from Figure 96 were added in quadrature.

$H \rightarrow 4\ell$ scale systematic [MeV]				
Systematic	4μ	$4e$	$2\mu 2e$	$2e 2\mu$
EG_SCALE_*	0.9	81.4	66.1	24.0
MUON_SCALE	71.7	0.0	41.3	27.6
EL_EFF_*	0.0	11.1	8.8	3.4
PRW_DATASF	1.0	9.5	4.2	15.4
H4lMCStat	8.1	8.1	8.1	8.1
EG_RESOLUTION_*	0.9	6.3	7.3	5.3




Figure 97: Systematics affecting the $H \rightarrow 4\ell$ signal mass resolution, shown for $m_H = 125$ GeV. Systematics with effects smaller than 0.05% are omitted.

Table 17: Summary of leading systematics affecting signal mass resolution. Systematics are labelled according to the scheme in Section 8.2. Where multiple systematics contribute under one name, the values from Figure 97 were added in quadrature.

$H \rightarrow 4\ell$ resolution systematic [%]				
Systematic	4μ	$4e$	$2\mu 2e$	$2e2\mu$
EG_RESOLUTION_*	0.02	0.50	0.31	0.93
PRW_DATASF	0.51	0.21	1.01	0.12
MUON_ID	0.50	0.00	0.36	0.06
EG_SCALE_*	0.27	0.20	0.19	0.42
MUON_MS	0.42	0.00	0.28	0.13
EL_EFF_*	0.00	0.35	0.21	0.35
MUON_SAGITTA	0.22	0.00	0.29	0.02

The scale systematics are incorporated into the signal model in a somewhat non-intuitive way. These systematics are properties of the detector performance and analysis strategy, not fundamental physics. Accordingly, they belong in the signal model as part of the response distribution, and not the physics distribution. However, for technical reasons, they are implemented in exactly the opposite way: the scale systematics are made to transform the m_H parameter inside the BW truth distribution of (91),

$$m_H^{2e2\mu} = m_H \times f_{\text{scale}}^{2e2\mu}(\boldsymbol{\theta}). \quad (98)$$

After this transformation, the parameter to determine is still m_H , but the mean of the BW distribution in each category is a function thereof. This is just a choice of convenience; the final signal PDF is unaffected by whether the scale systematic was applied on the response or truth component of the distribution. The resolution systematics are incorporated in a more natural way, directly on the response functions where they belong. The Gaussian parameters of the final mass response in (91) are scaled by the relevant systematic. For example,

$$\sigma_{i,\text{with sys.}}^{2e2\mu} = \sigma_{i,\text{no sys.}}^{2e2\mu} \times f_{\text{resolution}}^{2e2\mu}(\boldsymbol{\theta}). \quad (99)$$


Each Gaussian sigma (indexed by i) in the sum describing the per-event response is multiplied by this same factor, and all responses in a given category are assumed to have the same systematic uncertainty. Resolution uncertainties are not included in

the Template method, as there is no analytical resolution parameter that could be modified, so the technical implementation is much more difficult.

When using the PER signal model, the implementation of signal systematics presents an opportunity for future study. For all the same reasons that each event has a unique mass response, it will also have unique experimental systematics. This fact is neglected in this thesis, where the effects on $m_{4\ell}$ were calculated as the average effect of individual lepton systematics, much like the Template signal model uses the average response of many individual events. This averaging of systematic uncertainties comes with the same caveats as averaging detector resolution: events with below-average systematic uncertainty are assigned unrealistically large uncertainties, and those with above-average systematic uncertainties are given undue importance in the fit. Developing a strategy for implementing per-event systematics into the PER method has potential to yield a reduction in the overall systematic uncertainty on m_H measurement, though the magnitude of this difference is expected to be small, and the uncertainty in the $H \rightarrow 4\ell$ channel will remain by far statistically dominated until at least Run 3 at the LHC.

For the background processes, a different strategy is used. In this case, a histogram of the $m_{4\ell}$ distribution is built from events passing the full event selection with and without the systematic variations applied at $\pm 1\sigma$. These histograms use a very coarse 5 GeV binning to minimize effects of statistical fluctuations. The systematically varied histograms are then linearly interpolated from bin-to-bin, and this function is used as the $m_{4\ell}$ -dependent systematic effect on the signal. The class `RooStarMomentMorph` is used to interpolate (extrapolate) the effect on the PDF to variations smaller (larger) than $\pm 1\sigma$ at which the systematic was evaluated. One exception to this prescription applies to the Z+Jets background, for which the systematically varied distributions are produced as described in Section 6.3.3.

Two examples of the shape variations created in this way are shown in Figure 98 for the muon momentum scale and the dominant electron resolution systematic. The variations are shown at $\pm 5\sigma$ to emphasize the effect; when shown at $\pm 1\sigma$ the difference is invisible. This highlights that the experimental shape uncertainties for the background are less important than those of the signal. This is a result of the non-resonant background shape (as compared to the resonant signal shape) - an energy scale systematic which in general could displace events upward in $m_{4\ell}$ simultaneously “replaces” them with events moving upwards from still lower values of $m_{4\ell}$, leaving the shape only marginally affected.



(a) Effect on the 4μ category from muon momentum scale uncertainty

(b) Effect on the $4e$ category from dominant electron resolution uncertainty

Figure 98: Demonstration of systematic uncertainties affecting the $q\bar{q} \rightarrow ZZ^{(*)}$ PDF. The effects are shown for the muon momentum scale (MUON_SCALE) and the primary electron resolution systematic (EG_RESOLUTION_ZSMEARING) on the 4μ and $4e$ categories, respectively. The nominal PDF of $m_{4\ell}^{\text{cons}}$ is shown in black, while the $+5\sigma(-5\sigma)$ systematic variation is shown in red (blue).

How to assign import to the different systematic uncertainties presented in this chapter is not always evident. The ultimate question, of course, is what will be their impact on the m_H , $\sigma_{H \rightarrow 4\ell}$, and Γ_H measurements? Is a 0.01% scale systematic more or less important than a 5% normalization systematic? And where does a 1% resolution systematic fit in that comparison? The following chapter will present the statistical methods used for the property measurement in this thesis, and various validation steps and cross-checks which lead up to the final measurement. Among these checks, Section 9.3 addresses the preceding questions, and estimates the impact of each systematic uncertainty on the final measurements.

Chapter 9

Statistics and Property Measurement

This chapter presents the final statistical methods used for the property measurement in this thesis. All the components of the preceding chapters, especially Chapter 6 and onwards, come together to perform the measurement. The studies and results of all three methods are shown in this chapter: the constrained PER method as the primary measurement strategy, as well as the template method, and unconstrained PER method which serve as cross-checks. The extra checks are shown both as a point of comparison, and to inspire confidence in the constrained PER method as a measurement strategy. The constrained PER method is the exclusive source of the final results of this thesis, summarized in Chapter 10.

This chapter is laid out as follows: Section 9.1 presents the statistical formalism and background needed to understand the results in the following sections. Section 9.2 presents checks of the measurement using an effective dataset built from simulation. Section 9.3 presents a break-down of the systematic uncertainties by their effect on the final measurements. Section 9.4 presents checks using pseudo-datasets built with realistic statistics from what is expected in the data. Section 9.5 presents a validation of the PER method using the $Z \rightarrow 4\ell$ resonance in simulation and data. Finally Section 9.6 presents the results of the $H \rightarrow 4\ell$ property measurements in the data.

9.1 The method of maximum likelihood

The main statistical tool used to make inferences in this thesis is the method of maximum likelihood [83]. In this method, a likelihood function is constructed to convert the complex dataset (many numbers, $\boldsymbol{x}$) into a likelihood value (a single number $L(\boldsymbol{x})$), so as to make inferences from that value. Generally, the likelihood function will depend on some parameters of interest (POIs, denoted $\boldsymbol{\mu}$), so we can find

the best fit value of these parameters by finding where the likelihood is at its maximum value, i.e. the most likely values of the parameters. Confounding this strategy are the NPs associated to systematic uncertainties, $\boldsymbol{\theta}$, which enter the likelihood as additional free parameters.

The observables, $\mathbf{x}$, depend on the model being used:

$$\begin{aligned} \text{constrained PER: } \mathbf{x} &= \{m_{4\ell}^{\text{cons}}, \vec{g}\}, \\ \text{Template method: } \mathbf{x} &= \{m_{4\ell}^{\text{cons}}\}, \\ \text{unconstrained PER: } \mathbf{x} &= \{m_{4\ell}^{\text{fsr}}, \vec{g}\}, \end{aligned} \tag{100}$$

where the curly braces, e.g. $\{m\}$, denote “the set of all m in data”, and the $\vec{g}$ are the Gaussian-sum parameters representing the mass response, defined in (91) and the text that follows. Note that $\vec{g}$ are *conditional observables* - their distribution is not specified by the model; it is determined by the dataset.

The POIs, $\boldsymbol{\mu}$, are some subset of the Higgs mass, m_H , cross section times branching ratio to four leptons, $\sigma_{H \rightarrow 4\ell}$, and the natural width, Γ_H , depending on the goal of a given fit. For any given fit, the POI is just one of these parameters, while the others are treated as NPs, except for Γ_H which is fixed at 4.4 MeV (unless it is the POI). As only one parameter is considered the POI at a time, we can simply denote $\boldsymbol{\mu} = \mu$. Note that the POIs and NPs are both free in the fit; the distinction between the two is that we are interested in the value of the POIs, but not that of the NPs. Applying this rule gives the following:

$$\begin{aligned} \text{fitting mass: } \mu &= m_H, & \boldsymbol{\theta} &= (\sigma_{H \rightarrow 4\ell}, \theta_i, \theta_j, \dots), \\ \text{fitting cross section: } \mu &= \sigma_{H \rightarrow 4\ell}, & \boldsymbol{\theta} &= (m_H, \theta_i, \theta_j, \dots), \\ \text{fitting width: } \mu &= \Gamma_H, & \boldsymbol{\theta} &= (m_H, \sigma_{H \rightarrow 4\ell}, \theta_i, \theta_j, \dots). \end{aligned} \tag{101}$$

The above may seem an unnecessarily complex way to re-label the observables and parameters, but it allows the formulation which follows to be more universal to all signal models and types of fit. With these parameters defined, we are prepared to construct the likelihood.

For a given set of parameters $\mu, \boldsymbol{\theta}$, the *likelihood* of each event x_i can be expressed as

$$P(x_i | \mu, \boldsymbol{\theta}) = \sum_s^{\text{sources}} N_{s,c}^{\text{exp.}}(\mu, \boldsymbol{\theta}) P_{s,c}(x_i, \mu, \boldsymbol{\theta}). \tag{102}$$

Recall that when the PER model is used, the signal component of this function differs

on an event-by-event basis, whereas for the Template method it is fixed within a given category. Each event can be given a weight, w_i , which appears as an exponent on this value. Data events always have $w_i = 1$, while MC simulated events have non-unit weights in general, as discussed in Section 4.1. If all events are treated as independent observations, then their combined likelihood can be expressed as

$$\prod_i^{\text{events}} \left[\sum_s^{\text{sources}} N_{s,c}^{\text{exp.}}(\mu, \boldsymbol{\theta}) P_{s,c}(x_i, \mu, \boldsymbol{\theta}) \right]^{w_i} \quad (103)$$

To add some further complication, these events must be divided into categories, and within each category we must include an *extended term* accounting for the Poisson probability of observing N^{obs} events when N^{exp} are expected,

$$\prod_c^{\text{categories}} \text{Pois}(N_c^{\text{obs.}}; N_c^{\text{exp.}}(\mu, \boldsymbol{\theta})) \quad (104)$$

where c is an index over each category. This term effectively introduces a penalty in the likelihood when the overall number of events differs from the expected number (for a given set of POIs and NPs), which adds statistical power in constraining the parameters which affect normalization. The overall expected number of events per category is defined by

$$N_c^{\text{exp.}}(\mu, \boldsymbol{\theta}) = \sum_s^{\text{sources}} N_{s,c}^{\text{exp.}}(\mu, \boldsymbol{\theta}). \quad (105)$$

Also included in the likelihood are constraint terms for each of the systematic uncertainties entering the model,

$$\prod_{\alpha}^{\text{systematics}} G(\theta_{\alpha}; 0, 1). \quad (106)$$

When all these pieces are put together, the *extended likelihood* can be written down, according to the predicted distribution of the dataset $\mathbf{x}$, with parameter set $\mu, \boldsymbol{\theta}$, as

$$\begin{aligned} L(\mathbf{x}|\mu, \boldsymbol{\theta}) = & \prod_c^{\text{categories}} \left[\text{Pois}(N_c^{\text{obs.}}; N_c^{\text{exp.}}(\mu, \boldsymbol{\theta})) \times \prod_i^{\text{events}} \left[\sum_s^{\text{sources}} N_{s,c}^{\text{exp.}}(\mu, \boldsymbol{\theta}) P_{s,c}(x_i, \mu, \boldsymbol{\theta}) \right]^{w_i} \right] \\ & \times \prod_{\alpha}^{\text{systematics}} G(\theta_{\alpha}; 0, 1). \end{aligned} \quad (107)$$

The parameters $\mu, \boldsymbol{\theta}$ which maximize this value are called *maximum likelihood estimators*, and are denoted as $\hat{\mu}, \hat{\boldsymbol{\theta}}$. More colloquially, they are just the best fit values.

To get the most use of this likelihood function, consider the *profile likelihood ratio* [97], defined by

$$\lambda(\mu) = \frac{L(\mathbf{x}, \mu, \hat{\boldsymbol{\theta}}(\mu))}{L(\mathbf{x}, \hat{\mu}, \hat{\boldsymbol{\theta}})}, \quad (108)$$

where $\hat{\boldsymbol{\theta}}(\mu)$ indicates the best fit values of $\boldsymbol{\theta}$ for the given POI μ . Using the profile likelihood ratio, three test statistics can be defined that support all measurements in this thesis.

The first and most important test statistic used for measurement of m_H and $\sigma_{H \rightarrow 4\ell}$ is defined as

$$t_\mu = -2 \ln \lambda(\mu), \quad (109)$$

The use of a natural logarithm is computationally convenient for a few reasons. The most basic is that $t_\mu = 0$ at the best fit value, and is larger everywhere else. Additionally, the quotient in (108) becomes a simple subtraction of a constant (only the numerator inside the log needs to be re-evaluated with each μ), the repeated multiplications of (103) and (104) become repeated additions, and, instead of multiplying very small numbers repeatedly, relatively modest numbers are added, so that floating-point arithmetic is no longer an important concern.

To perform a measurement using t_μ , we utilize the property that it is asymptotically distributed as a chi-square distribution with one degree of freedom [97],

$$p(t_\mu|\mu) = \frac{1}{\sqrt{2\pi}} \frac{1}{\sqrt{t_\mu}} e^{-t_\mu/2}. \quad (110)$$

Using this property, a two-sided confidence region in parameter μ can be set where $t_\mu^{\text{data}} < t_\mu^{\text{crit}}$. The critical value is defined where

$$\int_{t_\mu^{\text{crit}}}^{\infty} p(t_\mu|\mu) dt_\mu = 1 - \text{CL}, \quad (111)$$

where CL is the desired confidence level of parameter μ . Using this expression, we find that

$$\begin{aligned} \int_{1.00}^{\infty} p(t_\mu|\mu) dt_\mu &= 1 - 68.3\%, \\ \int_{3.84}^{\infty} p(t_\mu|\mu) dt_\mu &= 1 - 95.0\%, \end{aligned} \quad (112)$$

which indicates the $\pm 1\sigma$ (68.3% CL) confidence region is set within the bounds where $t_\mu < 1.00$. The corresponding $\pm 2\sigma$ (95% CL) confidence region is set within $t_\mu < 3.84$.

With this information, doing a fit is as simple as scanning t_μ across values of μ , finding where it is zero to determine the best fit value, and where it crosses these critical values¹ to find the uncertainty - this procedure is called a *likelihood scan*.

The other two test statistics are used only for Γ_H . These new test statistics must be used, instead of t_μ , because the existing dataset does not have sufficient statistics to perform a two-sided direct measurement of Γ_H . In the following, the POI $\mu = \Gamma_H$ in every equation, even though it is written as μ for generality.

First, the consistency of the observed data with the hypothesis $\Gamma_H = \Gamma_H^{\text{SM}}$ is checked using the test statistic defined by

$$q_0 = \begin{cases} -2 \ln \lambda(0), & \text{when } \hat{\mu} > \mu^{\text{null}}, \\ 0, & \text{when } \hat{\mu} < \mu^{\text{null}}, \end{cases} \quad (113)$$

where $\mu^{\text{null}} = \Gamma_H^{\text{SM}} = 4.1 \text{ MeV}$. The conditional equation is designed such that downwards fluctuations in the observed width are not interpreted as being inconsistent with the null hypothesis of $\Gamma_H = \Gamma_H^{\text{SM}}$. The asymptotic PDF of q_0 under the null hypothesis is given by

$$p(q_0) = \frac{1}{2} \delta(q_0) + \frac{1}{2\sqrt{2\pi}} \frac{1}{\sqrt{q_0}} e^{-q_0/2}, \quad (114)$$

where $\delta(q_0)$ is the Dirac delta function. This distribution arises under the mathematical assumption that half of datasets will have $\hat{\mu} > \mu^{\text{null}}$, and the other half $\hat{\mu} < \mu^{\text{null}}$ (i.e. the maximum likelihood estimator is symmetric about the true value). The former will be distributed as in (110), while the latter will be assigned a value of zero, resulting in the $\delta(q_0)$ function. Consistency with the SM hypothesis is typically quoted as a pvalue, defined by

$$p_0 = \int_{q_0^{\text{data}}}^{\infty} p(q_0) dq_0, \quad (115)$$

which specifies the probability to have observed a deviation in data as large (or larger) than that observed, if the SM hypothesis is correct. In particle physics, *evidence* for a deviation from the SM is usually claimed at $p_0 < 1.35 \times 10^{-3}$, while a *discovery* (such as that of the Higgs boson) is claimed when $p_0 < 2.86 \times 10^{-7}$, corresponding to 3σ and 5σ Gaussian significance, respectively.

¹To work out any other CL or critical value, the integral in (110) is equal to $\text{erfc}(\sqrt{t_\mu^{\text{crit}}/2})$, where erfc is the complementary error function.

If no significant deviation from the SM is observed, then we turn to the final test statistic, which can be used to set an upper limit (UL) on Γ_H . It is defined by

$$q_\mu = \begin{cases} -2 \ln \lambda(\mu), & \text{when } \hat{\mu} \leq \mu, \\ 0, & \text{when } \hat{\mu} > \mu. \end{cases} \quad (116)$$

The effect of this conditional expression is to remove the low- Γ_H rejection region caused by upwards fluctuations of Γ_H (labelled in (116) as μ), and therefore set an upper limit on the width. Similar to the other test statistics, q_μ is asymptotically distributed as

$$p(q_\mu) = \frac{1}{2} \delta(q_\mu) + \frac{1}{2\sqrt{2\pi}} \frac{1}{\sqrt{q_\mu}} e^{-q_\mu/2}, \quad (117)$$

where the two conditions in (116) lead to the two terms in this expression. By computing the integral in the right-hand tail, we find

$$\begin{aligned} \int_{0.22}^{\infty} p(q_\mu|\mu) dq_\mu &= 1 - 68.3\%, \\ \int_{2.71}^{\infty} p(q_\mu|\mu) dq_\mu &= 1 - 95.0\%, \end{aligned} \quad (118)$$

so that a one-sided 68.3% (95%) frequentist confidence level UL can be set where q_μ rises to 0.22 (2.71).

The differences in the definitions of t_μ and q_μ are mostly immaterial for this thesis, but are presented in the standard way for consistency with the literature. For analyses performing a search (for a new resonance for example), the prospect of a downward or upward fluctuation causing misinterpretation of the results is a much more real effect. In this analysis, the two test statistics are only differentiated by where on the likelihood scan the critical values (and thus the exclusion regions) are drawn. For measurements of m_H and $\sigma_{H \rightarrow 4\ell}$, $\pm 1\sigma$ ($\pm 2\sigma$) exclusion regions are set where $t_\mu > 1.00(3.84)$. For measurements of Γ_H , a one-sided 68.3% (95%) exclusion region is set where $q_\mu > 0.22(2.71)$.

9.2 Tests using an effective dataset

The first set of tests used to validate the measurement strategy are based on an *effective dataset*. This effective dataset is similar to an *Asimov dataset* [97], where a dataset is created matching exactly the expected distributions and yields with no statistical fluctuation. The purpose of such a dataset is to test the statistical power

(significance, discovery potential, expected precision) of a measurement without requiring huge numbers of pseudo-experiments. When actual statistics are sufficiently high, the precision from fitting this Asimov dataset will reflect the median precision that would have been found from pseudo-experiments.

The effective dataset used in this section serves both this same purpose, and another purpose inaccessible to an Asimov dataset created as described above. The effective dataset is constructed by combining all the MC simulated samples for the different processes which enter the analysis: the signal (all production modes), qqZZ, ggZZ, VVV, Z+Jets. All samples are combined according to their SM cross sections, except for Z+Jets which uses the data-driven estimates to renormalize the samples. By this process, the sum of weights, $\sum_i w_i$, of each sample exactly matches the expected number of events in the likelihood, $N_{s,c}^{\text{exp}}$. For the background samples only, the $m_{4\ell}^{\text{cons}}$ and $m_{4\ell}^{\text{fsr}}$ values in the simulated samples are replaced by evaluating the PDF for each, $P_{s,c}(x_i, \boldsymbol{\theta} = \boldsymbol{\theta}_{\text{nominal}})$, in steps of 0.25 GeV, such that background exhibits no statistical fluctuation (similar to an Asimov dataset). For the signal samples, the events are used directly from the MC simulation without alteration, and when the PER method is being used, the response parameter $\vec{g}$ are also used without alteration. This permits an important check: while Chapter 7 validated the signal model(s) on signal samples alone, using this strategy permits the validation of the signal model(s) *in the presence of backgrounds*. This check is not possible using the Asimov dataset, which assumes the full signal and background model is perfect, and generates events from this combined PDF. Conversely, this could be seen as a disadvantage of the PER method: because the $\vec{g}$ are not specified by the model, the PER model is not capable of generating events directly from the PDF, unless a proto-dataset with values of $\vec{g}$ is provided (which would itself have statistical fluctuations), so generating a true Asimov dataset is not possible using the PER method.

9.2.1 Tests of expected resolution

For the first round of tests using the effective dataset, $m_H = 125$ GeV MC simulated samples are used for the signal. The NLL scans of m_H are shown for the unconstrained PER, the Template method, and the constrained PER in Figure 99. Each plot includes the likelihood scans from each category individually, and that from the combined model, with and without systematics. The systematic uncertainties affect the mass scale are extremely small, and are barely visible on these plots.

The same set of scans for $\sigma_{H \rightarrow 4\ell}$ are shown in Figure 100, and for Γ_H in Figure 101.




Figure 99: Scans of the NLL on parameter m_H using an effective dataset with $m_H = 125$ GeV. Subfigures show the scans using the different signal models as labelled. The NLL scans including (not including) systematics are shown with solid (dashed - barely visible) lines.

Note that for scans on Γ_H , only a one sided 95% CL upper limit is set. Also be reminded that the Template method does not include the Γ_H parameter, so is not able to serve as a cross check of this measurement.

These scans are used to extract the expected resolution of each measurement, assuming a SM $m_H = 125$ GeV Higgs boson. Since the SM production cross sections and the detector resolution vary as a function of m_H , the expected measurement resolutions will vary as a function of m_H , but by a relatively small amount. The results of these scans are recorded in Table 18, including all statistical and systematic




Figure 100: Scans of the NLL on parameter $\sigma_{H \rightarrow 4\ell}$ using an effective dataset with $m_H = 125 \text{ GeV}$. Subfigures show the scans using the different signal models as labelled. The NLL scans including (not including) systematics are shown with solid (dashed) lines.

uncertainties. These results indicate the constrained PER provides an improvement of approximately 2.6% on the m_H resolution, and 1.7% on the $\sigma_{H \rightarrow 4\ell}$ resolution, compared to the template method. The unconstrained PER method, while a satisfactory cross-check, does not attain the expected resolution of the other two methods on m_H measurement, but is competitive with the template method in the $\sigma_{H \rightarrow 4\ell}$ measurement.




Figure 101: Scans of the NLL on parameter Γ_H using an effective dataset with $m_H = 125$ GeV. Subfigures show the scans using the different signal models as labelled. The NLL scans including (not including) systematics are shown with solid (dashed) lines.

Table 18: Summary table of effective dataset results on m_H , $\sigma_{H \rightarrow 4\ell}$, Γ_H for the three signal models. The effective dataset is the same in all cases, with the true values from the input dataset as shown. Values inside the table use the same units as the reference values. Upper limits for Γ_H are given at 95% CL.

	constrained PER	Template	unconstrained PER
m_H	Input dataset value $m_H = 125.00$ GeV		
4μ	$124.99^{+0.53}_{-0.52}$	$124.98^{+0.52}_{-0.52}$	$125.00^{+0.61}_{-0.60}$
$4e$	$125.00^{+1.17}_{-1.15}$	$125.00^{+1.25}_{-1.20}$	$125.02^{+1.30}_{-1.27}$
$2\mu 2e$	$125.00^{+0.98}_{-0.97}$	$124.99^{+1.06}_{-1.03}$	$125.01^{+1.07}_{-1.06}$
$2e 2\mu$	$124.98^{+0.75}_{-0.73}$	$124.98^{+0.72}_{-0.72}$	$125.00^{+0.96}_{-0.93}$
combined	$124.99^{+0.37}_{-0.37}$	$124.99^{+0.38}_{-0.38}$	$125.00^{+0.43}_{-0.43}$
$\sigma_{H \rightarrow 4\ell}$	Input dataset value $\sigma_{H \rightarrow 4\ell} = 6.90$ fb		
4μ	$6.85^{+1.95}_{-1.69}$	$6.87^{+1.96}_{-1.70}$	$6.85^{+1.97}_{-1.71}$
$4e$	$6.94^{+2.86}_{-2.37}$	$6.90^{+2.90}_{-2.40}$	$6.94^{+2.89}_{-2.40}$
$2\mu 2e$	$6.94^{+2.77}_{-2.29}$	$6.93^{+2.78}_{-2.30}$	$6.94^{+2.77}_{-2.30}$
$2e 2\mu$	$6.91^{+2.36}_{-2.01}$	$6.90^{+2.38}_{-2.03}$	$6.92^{+2.40}_{-2.02}$
combined	$6.90^{+1.21}_{-1.10}$	$6.90^{+1.23}_{-1.12}$	$6.90^{+1.23}_{-1.12}$
Γ_H	Input dataset value $\Gamma_H = 0.0041$ GeV		
4μ	< 2.90	n.a.	< 3.47
$4e$	< 8.68	n.a.	< 10.23
$2\mu 2e$	< 6.97	n.a.	< 7.83
$2e 2\mu$	< 4.31	n.a.	< 6.08
combined	< 1.67	n.a.	< 2.13

9.2.2 Tests of model closure

The second set of tests using the effective dataset are a validation of the signal model(s) in the presence of background. They seek to better understand why the best fit values of m_H appearing in Table 18 often show a small deviation from the input mass $m_H = 125$ GeV, where they appear as 124.99 GeV or 125.01 GeV, and so forth. The test is only particularly meaningful for m_H measurement, and not the $\sigma_{H \rightarrow 4\ell}$ or Γ_H measurements.

As alluded to in Section 8.4, the best fit value in a given signal dataset is expected to fluctuate in accordance with MC statistics by an amount given in (97) as

$$\sigma_{MC\ stat} = \text{SD}[m_{4\ell}] / \sqrt{N_{MC}}$$

While it is the overall fluctuation that gets assigned as a systematic uncertainty, it is the per-category fluctuations that matter when interpreting the values in Table 18. To this point, Table 19 compiles the statistical variation in the mean $m_{4\ell}$, and thus best fit m_H , expected from each fit.

Table 19: Standard deviation of the mean $m_{4\ell}$ value, translating to the variations in the best fit m_H value, for independent signal datasets with the same actual statistics as the MC simulated signal samples used in this thesis. The unconstrained combined value was used as a systematic uncertainty in Section 8.4.

$H \rightarrow 4\ell$ scale systematic [MeV]		
category	unconstrained $\sigma_{MC\ stat}$ [MeV]	constrained $\sigma_{MC\ stat}$ [MeV]
4μ	13.8	13.4
$4e$	20.2	19.6
$2\mu 2e$	18.7	18.3
$2e 2\mu$	15.7	15.0
combined	8.1	7.8

To get a better indication of whether the small shifts observed in Table 18 arise purely from MC statistical fluctuation, three independent effective datasets are created with signals at different MC simulated mass points $m_H = 124, 125, 126$ GeV. For each dataset, NLL scans were performed, and the best fit mass, $\hat{m}_H$, was compared with the simulated mass, $m_H^{\text{simulated}}$, to calculate the *residual bias*,

$$\Delta_{m_H} \equiv \hat{m}_H - m_H^{\text{simulated}}. \quad (119)$$

Table 20: Residual bias Δ_{m_H} calculated using NLL scans on effective datasets at different mass values for the three signal models.

	$m_H = 124 \text{ GeV}$	$m_H = 125 \text{ GeV}$	$m_H = 126 \text{ GeV}$
constrained PER, $\Delta_{m_H} [\text{MeV}]$			
4μ	+8.7	-7.0	-11.0
$4e$	-27.9	+0.7	-45.1
$2\mu 2e$	+34.9	+6.6	-21.6
$2e 2\mu$	-0.9	-13.9	-23.2
combined	+6.2	-5.8	-19.1
Template, $\Delta_{m_H} [\text{MeV}]$			
4μ	-9.0	-10.6	-13.0
$4e$	-29.4	+1.6	-13.7
$2\mu 2e$	+3.4	+22.4	-13.5
$2e 2\mu$	-6.1	-16.0	-11.6
combined	-8.5	-13.3	-12.5
unconstrained PER, $\Delta_{m_H} [\text{MeV}]$			
4μ	+2.8	+1.7	-10.4
$4e$	-20.3	+22.5	-24.2
$2\mu 2e$	+26.2	+11.5	-21.3
$2e 2\mu$	-29.5	+12.8	+12.4
combined	+9.9	+8.7	-8.7

If the signal and background model(s) are flawless, then in a mathematical sense these values were randomly sampled from a Gaussian distribution with zero mean, and widths as specified in Table 19, $G(\Delta_{m_H}; 0, \sigma_{MC \text{ stat}})$. Using the three independent datasets gives an indication of the truth of this postulate. The residual biases for the different datasets, models, and categories are collected in Table 20.

At first glance, the observed residual biases are of comparable magnitude to that expected from statistical fluctuation alone, suggesting that the signal models introduce no additional uncertainty. The number of samples available is too limited to do a formal test on a per-category basis, so instead all categories are combined. Each of the twelve values per-method (the combined values are derivatives of the others,

so can't be used here) are standardized into twelve values of $z = \Delta_{m_H}/\sigma_{MC\ stat}$. The sample mean, standard deviation, and their respective errors are calculated [98] for each set of values to be

$$\begin{aligned}\bar{z}^{\text{constrained PER}} &= -0.49 \pm 0.34, & s_z^{\text{constrained PER}} &= 1.14 \pm 0.55, \\ \bar{z}^{\text{Template}} &= -0.51 \pm 0.21, & s_z^{\text{Template}} &= 0.71 \pm 0.22, \\ \bar{z}^{\text{unconstrained PER}} &= -0.08 \pm 0.32, & s_z^{\text{unconstrained PER}} &= 1.07 \pm 0.49,\end{aligned}\tag{120}$$

which are consistent with the proposal of $\bar{z} = 0.00$, $s_z = 1.00$, i.e. that the observed residual biases come from MC statistical fluctuations only.

The tests in this section leave the encouraging conclusion that the different methods under study - primarily the constrained PER method, but also the two cross-checks methods - are robust across a range of possible Higgs masses, and will serve as effective tools for measurement of m_H .

9.3 Effects of systematics

The likelihood scans in the preceding section are sufficient to see the combined effect of all systematics on the measurement; this can be read directly from the difference between the likelihood scans with systematics and those without. It is still very useful to decompose the effect of the systematics on each measurement, which allows the systematics to be ranked according to their impact on the measurement. To this end, the following quantities are calculated, for each source of systematic uncertainty α :

$\hat{\theta}_\alpha$ The best fit value and uncertainty of NP θ_α when all other parameters are fixed at their best fit values, called the *pull* of α .

$\hat{\mu}(\theta_\alpha = \pm 1)$ The best fit value of POI μ (or upper limit of Γ_H) when all other parameters are fixed at their best fit values, except θ_α which is set to ± 1 . These values are compared against the global best fit (or upper limit) of μ (when all other parameters, including θ_α , are at their best fit values) to see the *effect* of α .

Since the implementation of systematics is common to all methods, this study was performed using only the constrained PER method. The effects were studied on the effective dataset of Section 9.2 with $m_H = 125$ GeV, and in the data. The decomposition of systematic uncertainties for the m_H measurement is shown in Figure 102, for




Figure 102: Effects of different sources of systematic uncertainty on the m_H measurement, as seen in the effective dataset from simulation with $m_H = 125$ GeV described in Section 9.2, and in the data. Shown are the pulls (black points and error bars) of each NP, and its effect on m_H when set to ± 1 . Systematics with effects less than ± 4 MeV are omitted.

the $\sigma_{H \rightarrow 4\ell}$ measurement in Figure 103, and for the Γ_H measurement in Figure 104.

In the effective dataset, the pulls should indicate $\theta_\alpha = 0.00 \pm 1.00$ since all systematics are implicitly at their nominal values. In the data, some statistical fluctuation is expected (and observed), since the global best-fit value of the NPs is a-priori unknown.

In interpreting the effects in these plots, note that the systematics are ordered from top to bottom in decreasing order of their effect². These plots can therefore be used to determine the most important systematic uncertainties for each measurement, something not possible from examining the fit inputs in Chapter 8 alone. Table 21 summarizes the dominant systematics observed for each measurement. Note that the values do not exactly match the plots, as the table adds (in quadrature) some systematics which were omitted from the plots.

²Specifically, they are sorted by $\left| \hat{\mu}(\theta_\alpha = +1) - \hat{\mu}(\theta_\alpha = -1) \right|$




Figure 103: Effects of different sources of systematic uncertainty on the $\sigma_{H \rightarrow 4l}$ measurement, as seen in the effective dataset from simulation with $m_H = 125$ GeV described in Section 9.2, and in the data. Shown are the pulls (black points and error bars) of each NP, and its effect on $\sigma_{H \rightarrow 4l}$ when set to ± 1 . Systematics with effects less than ± 60 ab are omitted.

One noteworthy feature can be seen by looking at systematics on the m_H measurement. While the effect of the electron energy scale uncertainties on the $4e$ category are about 12% larger than the effect of the muon momentum scale on the 4μ category (refer back to Section 8.4), in the combined fit the effect of the muon momentum scale is about 130% larger. This occurs because the 4μ category dominates the measurement statistically, so systematics in this category have a larger impact on the combination.

The ranking of the sources of systematic uncertainty in this section is useful in directing future measurement strategy. Although statistical uncertainties are at this stage by far the dominant source of uncertainty, that will not always be the case. The systematic uncertainties which dominate the measurements now are those on which the greatest focus should be to improve in future. For all measurements, the theoretical calculation of the dominant $q\bar{q} \rightarrow ZZ^{(*)}$ background is one of the leading effects. For the m_H measurement specifically, the electron energy and muon momentum scales




Figure 104: Effects of different sources of systematic uncertainty on the Γ_H measurement, as seen in the effective dataset from simulation with $m_H = 125$ GeV described in Section 9.2, and in the data. Shown are the pulls (black points and error bars) of each NP, and its effect on the upper limit of Γ_H when set to ± 1 . Systematics with effects less than ± 20 MeV are omitted.

are both very important. For the $\sigma_{H \rightarrow 4\ell}$ measurement, the electron efficiency systematics also play a very important role, followed closely by pile-up uncertainties, and muon efficiency and isolation. For the Γ_H measurement, the dominant sources of systematic uncertainty are all related to the event yields, and (somewhat surprisingly) not the resolution. This arises because as the natural width increases, the contribution to the overall lineshape width from resolution systematics falls off. Also as the natural width increases, the signal looks more and more background-like, and so the background normalizations play a large role in where the Γ_H upper limit can be set. Focusing time and effort in reducing these sources of systematic uncertainties will pay the most benefit in the resolution of future measurements.

Table 21: Summary table of dominant systematics affecting measurement of $m_H, \sigma_{H \rightarrow 4\ell}$, and upper limits of Γ_H . The value presented are the observed effect of each systematic, coming from Figure 102b, Figure 103b, Figure 104b. The naming scheme introduced in Section 8.2 is used to group systematics where necessary.

Dominant systematics for m_H [MeV]			
MUON_SCALE	46	QCDscale_VV	20
EG_SCALE_*	17	pdf_qq	14
lumi	11	H41MCStat	8
reducible_llee_*	7	EL_EFF_*	7
total	58		
Dominant systematics $\sigma_{H \rightarrow 4\ell}$ [ab]			
lumi	370	EL_EFF_*	230
QCDscale_VV	190	PRW_DATASF	150
pdf_qq	140	MUON_EFF_*	120
Signal_comp	110	MUON_ISO_*	100
total	570		
Dominant systematic for Γ_H upper limit [MeV]			
QCDscale_VV	130	pdf_qq	92
lumi	62	reducible_llee_*	59
QCDscale_ggVV	43	EL_EFF_*	43
reducible_lluu_*	30	MUON_ISO_*	23
total	190		

9.4 Tests using pseudo-experiments

The effective dataset approach used in Section 9.2 is very powerful for determining properties of the fit. In particular, it is very useful for determining the median expected fit result [97]. However, the median behaviour is not always a satisfactory metric, especially in the low-statistics regime of the current measurement. One reason is that by using the median behaviour of a huge number of events, the highly variable mass responses around which the PER method is built are averaged out and appear more like the Template method. So the effective dataset in practice hides the aspects of the PER which distinguish it from the Template method. This section will

demonstrate that in the low-statistics regime expected in the data, the constrained PER method outperforms the Template method, specifically because of its use of per-event resolution information. As the methods all use a common implementation and values for systematic uncertainties (with the exception of uncertainties on the signal resolution being omitted for the Template method), the inclusion of systematics is not useful in differentiating these methods³. For this reason, no systematic uncertainties are included in these studies.

To test the measurement at low statistics, a large number of pseudo-experiments were performed, also called *toys*. The datasets in each toy, called *pseudo-datasets*, were built using events drawn from MC simulated samples with the event selection and analysis mass window applied. By using directly MC simulation, no assumption is placed on the validity of the signal or background PDFs. The events were drawn in random order with probabilities proportional to their respective event weights. This procedure forces the toys to have the same composition on average as the effective dataset of Section 9.2. The number of events drawn from each source in each category was a Poisson-distributed random number based on the expected number of events,

$$N_{qqZZ,2e2\mu}^{\text{toy } i} = \text{random.Poisson}(N_{qqZZ,2e2\mu}^{\text{exp.}}). \quad (121)$$

For the signal, $m_H = 125 \text{ GeV}$ samples were used for each production mode, and the total number of signal events to be sampled was split up among production modes according to their relative contributions inside the mass range using SM cross sections and the simulated samples⁴:

$$\begin{aligned} N_{ggF,2e2\mu}^{\text{exp.}} &= 88.2\% N_{\text{signal},2e2\mu}^{\text{exp.}}, \\ N_{VBF,2e2\mu}^{\text{exp.}} &= 7.6\% N_{\text{signal},2e2\mu}^{\text{exp.}}, \\ N_{VH,2e2\mu}^{\text{exp.}} &= 3.5\% N_{\text{signal},2e2\mu}^{\text{exp.}}, \\ N_{ttH+bbH,2e2\mu}^{\text{exp.}} &= 0.7\% N_{\text{signal},2e2\mu}^{\text{exp.}}. \end{aligned} \quad (122)$$

Using these fractions, a random number of events, following the same prescription as for the backgrounds, was drawn to be used in each toy. This procedure was used to construct a large number of toys (15,000 or more, depending on which measurement),

³Systematics would be useful in differentiating these measurements from ATLAS' m_H measurement in Run 1 [7], which used different analysis method with larger systematic uncertainties.

⁴Note that these are different than the SM production fractions in Table 3 because of the mass window applied - fewer WH,ZH,ttH events survive when the mass window is applied, because the four-lepton candidate can often contain a lepton that did not originate from the Higgs decay. As a result, the relative fractions of ggF and VBF are increased.

each representing what *could* be seen in the data, if $m_H = 125$ GeV. These toys can be used to validate and compare the fit methods in a complementary fashion to the effective dataset approach.

9.4.1 Fit results in pseudo-experiments

The first set of tests using the toy experiments comes from fitting each model (constrained PER, Template, unconstrained PER) on the toy datasets; the results of the fits are used to make comparisons between the methods. The three metrics by which the toys are compared are:

$\hat{\mu}$ The best fit value of the parameter of interest μ (e.g. m_H). The distribution of best fit values is expected to be Gaussian, centered on the true value, with width equal to the resolution of the measurement. For Γ_H , it is expected to be zero, with a right-hand tail (negative Γ_H is not possible, by definition).

p_μ The *pull* of μ , defined by the residual of the best fit value, divided by the estimated error. The pull is, by definition, expected to be distribution as a dimensionless Gaussian with zero mean, and unit width. Deviations of the mean indicate a measurement bias, while deviations of the width indicate mis-estimation of uncertainties in the fit. The pull is ill-defined by Γ_H , and is not shown.

95% UL The 95% CL upper limit, applicable to Γ_H toys only.

The pull can be calculated in the case of asymmetric errors[99] by

$$p_\mu = \begin{cases} \frac{\hat{\mu} - \mu^{\text{true}}}{|\sigma_\mu^+|} & \text{when } \hat{\mu} < \mu^{\text{true}}, \\ \frac{\hat{\mu} - \mu^{\text{true}}}{|\sigma_\mu^-|} & \text{when } \hat{\mu} > \mu^{\text{true}}, \end{cases} \quad (123)$$

where $\sigma_\mu^\pm$ are the positive and negative errors of the fit. This definition is constructed such that 68.3% of pulls will fall within ± 1 , if the errors are accurately estimated in the fit.

The distributions of the best fit values and pulls from the three methods for m_H are found in Figure 105, and for $\sigma_{H \rightarrow 4\ell}$ in Figure 106. In each case, the distribution is fit with a Gaussian. The mean and standard deviation of the histogram itself are recorded on the plot as $\text{mean}_{\text{hist}}$ and RMS_{hist} , respectively. The mean and width of the Gaussian fit are recorded as μ_{fit} and σ_{fit} , respectively. The equivalent plots, separated

by analysis category, are shown in Appendix C.3. The values of these Gaussian fits are recorded in Table 22 for m_H , and Table 23 for $\sigma_{H \rightarrow 4\ell}$. Recall that the uncertainty on μ_{fit} is approximately $\sigma_{\text{fit}}/\sqrt{N_{\text{toys}}}$, falling with the number of toys ⁵.

The distributions of the best fit values and 95% CL upper limits on Γ_H from the PER method with and without ZMC are found in Figure 107. The mean, median, and standard deviation of the histogram is recorded on both plots. The equivalent plots, separated by analysis category, are shown in Appendix C.3. The median fit value is exactly zero⁶ for both methods (and in all individual categories). The median and mean 95% CL upper limits from the toy experiments are recorded in Table 24.


Table 22: Summary table of toy experiment fits of m_H . The parameters of the Gaussian fits in Figure 105 are recorded here. Units in the table match the header.

	constrained PER		Template		unconstrained PER	
	μ_{fit}	σ_{fit}	μ_{fit}	σ_{fit}	μ_{fit}	σ_{fit}
$\hat{m}_H$, $m_H^{\text{true}} = 125 \text{ GeV}$						
4μ	125.00	0.52	124.99	0.52	125.00	0.61
$4e$	125.01	1.08	124.97	1.12	125.04	1.20
$2\mu 2e$	125.02	0.93	124.99	0.98	125.02	1.01
$2e 2\mu$	124.94	0.72	124.95	0.72	124.95	0.88
combined	124.99	0.36	124.98	0.37	125.00	0.43
p_{m_H}						
4μ	-0.01	1.00	-0.02	1.01	-0.01	1.01
$4e$	0.01	1.02	-0.03	1.05	0.02	1.03
$2\mu 2e$	0.01	1.01	-0.02	1.04	0.01	1.01
$2e 2\mu$	-0.09	1.01	-0.07	1.03	-0.06	1.03
combined	-0.03	0.98	-0.05	1.00	-0.01	0.99


The upper section of Table 22 can be used to compare the constrained PER model with the Template method in a context closely related to the actual data. The widths

⁵The specific number of toys entering each number is printed on the corresponding plot


⁶Technically, it is the lower limit of the fit parameter, which isn't allowed to go below 10^{-5} so that the BW remains finite.



(a) constrained PER




(b) Template




(c) unconstrained PER


Figure 105: Results of m_H fits from three methods on toy datasets. Shown are the best fit values $\hat{m}_H$ (left), and the pulls p_{m_H} (right) from each method. The raw distribution of results from the toys is shown in blue, and a Gaussian fit is shown with a red line.



(a) constrained PER



(b) Template



(c) unconstrained PER

Figure 106: Results of $\sigma_{H \rightarrow 4\ell}$ fits from three methods on toy datasets. Shown are the best fit values $\hat{\sigma}_{H \rightarrow 4\ell}$ (left), and the pulls $p_{\sigma_{H \rightarrow 4\ell}}$ (right) from each method. The raw distribution of results from the toys is shown in blue, and a Gaussian fit is shown with a red line.




Figure 107: Results of Γ_H fits from Per-event response method, with and without Z Mass Constraint, on toy datasets. Shown are the best fit values $\hat{\Gamma}_H$ (left of each subfigure), and the 95% CL upper limit (right) from each method. The raw distribution of results from the toys is shown in blue. Note the log axis and large first bin in the left-hand plots.

of the $\hat{m}_H$ distribution indicate the expected resolution of the measurement: the constrained PER method shows no difference in the 4μ category, a 3.6% improvement in the $4e$ category, a 5.1% improvement in the $2\mu 2e$ category, no change in the $2e 2\mu$ category, and a 2.6% improvement overall, demonstrating that it is more *precise*. These improvements are consistent with what was found using the effective dataset in Section 9.2.1. While the effective dataset can test the bias of the models at high statistics, where all models were found to be unbiased in Section 9.2.2, it does not test the bias expected in data, where limited statistics are available. The means of the $\hat{m}_H$ distribution provide a measure of the bias of the methods in this low statistics regime. In the combined fit toys, the constrained PER method shows less bias compared to

the Template method (mean of $\hat{m}_H$ is 124.99 versus 124.98), indicating it is also more *accurate*. These are two of the key reasons why the PER method is such a powerful tool for measurement of resonance mass.

Table 23: Summary table of toy experiment fits of $\sigma_{H \rightarrow 4\ell}$. The parameters of the Gaussian fits in Figure 106 are recorded here. Units in the table match the header.

	constrained PER		Template		unconstrained PER	
	μ_{fit}	σ_{fit}	μ_{fit}	σ_{fit}	μ_{fit}	σ_{fit}
$\hat{\sigma}_{H \rightarrow 4\ell}$, $\sigma_{H \rightarrow 4\ell}^{\text{true}} = 6.90 \text{ fb}$						
4μ	6.99	1.71	6.98	1.72	7.01	1.74
$4e$	7.14	2.38	7.07	2.37	7.09	2.41
$2\mu 2e$	7.11	2.33	7.12	2.31	7.07	2.35
$2e 2\mu$	6.95	2.03	6.94	2.06	6.96	2.04
combined	6.95	1.09	6.94	1.10	6.96	1.10
$p_{\sigma_{H \rightarrow 4\ell}}$						
4μ	0.01	0.97	0.01	0.97	0.02	0.97
$4e$	0.05	0.95	0.02	0.95	0.04	0.95
$2\mu 2e$	0.05	0.95	0.04	0.95	0.03	0.94
$2e 2\mu$	-0.02	0.96	-0.03	0.96	-0.02	0.95
combined	0.02	1.01	0.00	1.00	0.03	1.00

Examining Table 23 reveals an interesting feature. All methods in (nearly) all categories on average fit a $\sigma_{H \rightarrow 4\ell}$ value *higher* than the input to the toys. Curiously, the average fit value in the individual categories is (almost) always higher than the combined average fit value. How can this be? The cause is the *Greedy Bump Bias* [100], wherein the floating m_H in the toy fits allows the signal shape to “chase after” upward statistical fluctuations of the background. With more statistics in the combined fit, m_H is more restricted from chasing statistics, and the fit is “less greedy”. The lower section of Table 23 demonstrates that this bias overall is $\mathcal{O}(2\%)$ of the statistical uncertainty, so is negligible next to this and to the other systematic uncertainties in the fit result. The overall resolution from the constrained PER method has a slight advantage, approximately 1% improvement, over the Template method. This is also consistent with the improvement found from the effective dataset.

Table 24: Summary table of toy experiment fits of Γ_H . The median and mean 95% CL upper limits from toy experiments are recorded here, as found in Figure 107. Units in the table match the header.

	constrained PER		unconstrained PER	
	median 95% UL	mean 95% UL	median 95% UL	mean 95% UL
$\Gamma_H^{\text{true}} = 0.0041 \text{ GeV}$				
4μ	2.51	3.74	3.00	4.61
$4e$	7.58	15.10	8.90	17.01
$2\mu 2e$	5.81	11.64	6.68	13.23
$2e 2\mu$	4.25	8.75	5.75	11.16
combined	1.56	1.94	1.92	2.39

Comparing the median upper limits from toys in Table 24 with the upper limits in the effective dataset in Table 18 suggests that the effective dataset strategy overestimates the upper limits that can be set by about 10%, depending on the category. This difference arises because the upper-limit setting on the width is extremely sensitive to statistical fluctuation, far more than the m_H or $\sigma_{H \rightarrow 4\ell}$ measurements, so the effective dataset is less able to replicate the median expected result that a true Asimov dataset would have done.

A few other interesting points arise from comparing the fit results with toys (Table 22, 23) to those with the effective dataset (Table 18). Both methods are intended to provide an estimate of the expected resolution of the fit on data, but under different assumptions. By focusing on the constrained PER toys fitting m_H as an example, the results indicate that the 4μ uncertainty is well-estimated by the effective dataset approach (both methods suggest 0.52 GeV), as well as the $2e 2\mu$ (about 0.72 GeV from both methods). Contrarily, in the $4e$ category the effective dataset overestimates the resolution (toys show 1.08 GeV versus 1.16 GeV from the effective dataset), as it also does in the $2\mu 2e$ category (0.93 GeV versus 0.97 GeV). The combined fit uncertainty agrees from both methods, indicating 0.36 GeV expected resolution on m_H (without systematics). The same comparison is more difficult for $\sigma_{H \rightarrow 4\ell}$, where the uncertainties from the effective dataset are significantly asymmetric for the individual categories. For the combined fit, the effective dataset and toys are in agreement on the expected resolution on $\sigma_{H \rightarrow 4\ell}$.

Complementary to this feature are the widths of the pull distributions, which are equal to 1.00 when the uncertainties are perfectly estimated. Deviations from

this value would indicate uncertainties are not being accurately estimated, and so the *asymptotic approximation*, used in Section 9.1 to define rejection regions of the test statistics may not be satisfied. However, for all three methods, the width of the combined fit pulls of m_H and $\sigma_{H \rightarrow 4\ell}$ are within 0.98-1.01, indicating the uncertainties are being accurately estimated even at low statistics. In the individual categories, the deviations are slightly larger, but all pull widths are within 0.94-1.05, suggesting even the individual categories have mostly sufficient statistics to satisfy the asymptotic approximation and set confidence regions as described.

9.4.2 Toy-by-toy comparison

Comparing the properties of an ensemble of toy experiments yields important, yet not complete, information. Some properties become more illustrative when they can be compared on a toy-by-toy basis. This is more akin to the situation in data, where there is only a single dataset, and the results between the methods will be compared only on that dataset. Fortunately, the process by which toys are created allows them to be uniquely “barcoded”, so that the result of fitting one particular dataset with the constrained PER method can be compared with the fit results when fitting the Template method on the exact same set of events (but without using the per-event response information). This barcoding strategy does not allow the unconstrained PER method to be compared on a toy-by-toy basis, because the application of the ZMC will move events in or out of the mass window, and so the population of MC simulated events from which the toys are drawn is not equivalent between the unconstrained and constrained toy generation. This is not a great loss, as the preceding checks have demonstrated the unconstrained PER is a valid, but lower-resolution method. This section will therefore make toy-by-toy comparisons only of the constrained PER and Template methods.

Figure 108 contains the two-dimensional correlations of $\hat{m}_H$ and $\hat{\sigma}_{H \rightarrow 4\ell}$, respectively, from the constrained PER method and the Template method. These plots indicate that the fit results are highly correlated. The correlation coefficient ρ is printed on each plot, and is found to be ordered by the expected number of events per-category (higher statistics yields higher correlation, not surprisingly). The same plots separated by category can be found in Appendix C.3.

While the correlations are visually interesting, other distributions can be more useful in quantifying the models. Two interesting things to explore are *how much the models agree with each other*, and *how much the models agree with the truth*. The




Figure 108: Correlation plots of $\hat{m}_H$ and $\hat{\sigma}_{H \rightarrow 4\ell}$, respectively, between the constrained PER model, and the Template model, from pseudo-experiments. The diagonal line indicates where both methods agree exactly.

former, which contains all the same information as Figure 108 but in one dimension, is shown in Figure 109. This quantifies how much the two models are expected to agree (or disagree) when fit on the data. The latter can be understood, in a somewhat less intuitive fashion, by comparing the absolute difference between the best fit values and the truth value from the input to the toys ($m_H = 125$ GeV, and $\sigma_{H \rightarrow 4\ell} = 6.90$ fb). This comparison is shown in Figure 110. Entries to the left of zero are cases where the constrained PER model fit a more accurate value (closer to the truth) than the Template model, and entries to the right of zero are cases where the Template model returned a more accurate result. Consequently, the mean value shows how much *closer to the true value* the constrained PER method is, on average. The area under the distribution from negative infinity to zero gives the fraction of toy datasets in which the constrained PER model was more accurate. The most impressive difference comes from m_H measurement in the $4e$ category, where the constrained PER method is on average 72 MeV closer to the true m_H , and more accurate in 54.5% of datasets. The differences in the other categories are more subdued, so that the combined result is that the PER method fits more accurately the true m_H ($\sigma_{H \rightarrow 4\ell}$) value in 51.0% (50.2%) of datasets, being on average 4 MeV (< 1 ab) closer to the true value. These quantities are summarized in Table 25, and the corresponding plots for each category




Figure 109: Difference in best fit values of m_H and $\sigma_{H \rightarrow 4\ell}$, respectively, between the constrained PER model, and the Template model, from pseudo-experiments. The standard deviation, recorded on the plots as RMS_{hist} , gives an indication of how much the two models may agree or disagree when fit on the data.




Figure 110: Comparison of distance from the best fit value to the true value, of m_H and $\sigma_{H \rightarrow 4\ell}$, one a toy-by-toy basis. In toys to the left of zero, the constrained PER model produces a more accurate fit result than the Template model. In toys to the right of zero, the Template method was more accurate.

are found in Appendix C.3.

Table 25: Expected improvement in fit results in constrained PER method compared to Template method. The second column shows the RMS deviation between the best fit values of m_H ($\sigma_{H \rightarrow 4\ell}$) from the two methods. The third column shows the average improvement; that is, on average, how much closer to the true value of m_H ($\sigma_{H \rightarrow 4\ell}$) is the constrained PER method compared to the Template method. The fourth column shows in what percentage of datasets the constrained PER method is expected to produce a fit result closer to the true value of m_H ($\sigma_{H \rightarrow 4\ell}$) than the Template method.

category	RMS deviation	Average improvement	More accurate results
m_H			
4μ	182 MeV	1 MeV	50.0%
$4e$	530 MeV	72 MeV	54.5%
$2\mu 2e$	381 MeV	54 MeV	55.1%
$2e 2\mu$	373 MeV	4 MeV	50.0%
combined	138 MeV	4 MeV	51.0%
$\sigma_{H \rightarrow 4\ell}$			
4μ	222 ab	11 ab	51.2%
$4e$	350 ab	-11 ab	49.2%
$2\mu 2e$	290 ab	4 ab	51.4%
$2e 2\mu$	315 ab	5 ab	50.9%
combined	227 ab	< 1 ab	50.2%

9.5 Validation using the $Z \rightarrow 4\ell$ resonance

As a final check before looking at the data in the signal region, the PER methods can be tested on another resonance with well-known properties: the $Z \rightarrow 4\ell$ resonance described in Section 6.1. Measuring these known quantities in data gives an additional handle on the performance of the PER technique. Section 9.5.1 describes some of the additional considerations necessary to describe the lineshape of the Z , and Section 9.5.2 reports measurements of the resonance properties. As this section is intended as a validation of the PER method, rather than a standalone result, the tests in this section are all reported with statistical uncertainties only, while no systematics are included. The result could otherwise stand on its own as a valid measurement, but because the statistics are so limited compared to a $Z \rightarrow \ell\ell$ measurement, the resolution is several orders of magnitude worse than existing measurements. Another

contributing factor is that the LHC is a p - p collider, and so the collisions are far messier than at LEP, the e^+e^- collider where the best existing measurements of the Z boson properties were performed.

The category definitions used for the $Z \rightarrow 4\ell$ analysis are the same as in the primary analysis, except a mass window of $78 \text{ GeV} \leq m_{4\ell} \leq 104 \text{ GeV}$ is used.

9.5.1 Per-event response and the $Z \rightarrow 4\ell$ lineshape

The lepton response model developed to describe the $m_{4\ell}$ response of the Higgs resonance, described in Section 7.3.2, can be used for the Z resonance. One caveat to this arises because m_Z is significantly lower than m_H . The lower resonance mass means that a lepton with a given energy has a different response distribution when originating from a Z , than if it had come from a Higgs⁷. For this reason a separate set of tuning parameters, originally described in Section 7.3.4, is required when fitting the $Z \rightarrow 4\ell$ peak.

An additional consideration is the interpretation of $\epsilon(m_{4\ell}^{\text{true}})$, first defined in Section 7.3.1. This function is necessary to properly describe the true distribution of events before reconstruction,

$$\text{BW}_\epsilon(m_{4\ell}^{\text{true}}; m_Z, \Gamma_Z) = \epsilon(m_{4\ell}^{\text{true}}) \text{BW}(m_{4\ell}^{\text{true}}; m_Z, \Gamma_Z). \quad (124)$$

The efficiency function $\epsilon(m_{4\ell}^{\text{true}})$ was initially defined to account for non-uniform detector efficiency as a function of $m_{4\ell}^{\text{true}}$. Now it must be re-interpreted to absorb two other effects: Many of the selected Z 1 di-lepton pairs are produced off-shell creating a significant low-side tail, and, even more important, on the high side there is a continuous transition between $q\bar{q} \rightarrow Z \rightarrow 4\ell$ into $q\bar{q} \rightarrow ZZ^* \rightarrow 4\ell$ (described in Section 6.1), creating a high-side tail. The function $\epsilon(m_{4\ell}^{\text{true}})$ as it is used in this section incorporates both these effects. It is described by a third-order polynomial, fit on the ratio of the $m_{4\ell}^{\text{true}}$ distribution from MC simulation and a pure $\text{BW}(m_{4\ell}^{\text{true}}; m_Z, \Gamma_Z)$ distribution. These fits are shown in Figure 111. Note that the efficiency functions are fit on a wider range than the category definitions to encourage smooth behaviour. The discrepancies at the extreme ends are therefore unimportant for the purposes of this cross-check.

⁷If this seems bizarre, consider a specific example: a lepton reconstructed with $E_i^{\text{reco}} = 20 \text{ GeV}$. If asked to estimate the PDF of its true energy $P(E_i^{\text{true}})$, would the answer be the same with the knowledge this lepton came from a Z , as it would knowing that it came from a H ? Surely it would not be the case. While this effect is present within the m_H fit range itself, it is considered to be




Figure 111: $Z \rightarrow 4\ell$ efficiency function, $\epsilon(m_{4\ell}^{\text{true}})$, modelling the difference between the MC simulated distribution of $m_{4\ell}^{\text{true}}$ and a pure BW distribution, in each category. The upper plot shows the MC simulated distribution (blue line), and pure BW distribution (red line). The lower plots shows their ratio (black points and line), and the 3rd order polynomial used for $\epsilon(m_{4\ell}^{\text{true}})$ (red curve).

Another point of note is that the likelihood of (107) is more simple for this measurement because the vast majority of events are from the $Z \rightarrow 4\ell$ signal itself; the measurement is nearly background-free. The only background contribution is from Z +Jets, which is nearly two orders of magnitude smaller than the signal inside the mass range used. One could have also considered the right-hand $q\bar{q} \rightarrow ZZ^{(*)} \rightarrow 4\ell$ tail as a separate background, but this is more difficult to disentangle in MC simulation, with no particular benefit. The signal normalization in the likelihood is parameterized with respect to the SM as,

$$N_{Z \rightarrow 4\ell, c}^{\text{exp.}} = \mu_Z N_{Z \rightarrow 4\ell, c}^{\text{SM}} \quad (125)$$

where μ_Z is a free parameter called the *signal strength*, equal to 1.00 in the SM.

Also, because Γ_Z is orders of magnitude larger than Γ_H , ATLAS has sufficient experimental resolution to make a two-sided width measurement of the Z , and so the t_μ test statistic is used to perform the measurement in the same way as for m_Z

absorbed by the range of samples used in building the lepton model in Section 7.3.2.

and μ_Z .

The final important difference between fitting the $Z \rightarrow 4\ell$ peak, as compared to the $H \rightarrow 4\ell$ peak is that the unconstrained PER and the constrained PER are expected to give nearly identical results. The reason is that the ZMC is designed to constrain the Z1 di-lepton pair to have an invariant mass near m_Z . But this is of course not possible if the four-lepton invariant mass is already near m_Z itself. So in practice, the $m_{4\ell}$ values of events in the mass window used here shift upwards by a small amount (≈ 0.1 GeV), but the resolution of the Z peak is not improved, so neither is the expected resolution from the constrained PER as compared to the unconstrained (in data, subject to statistical fluctuations, they may differ). For this reason, only the constrained PER is used in what follows, while the unconstrained results (expected and observed), are kept in Appendix C.4.

For this study, the transformation parameters of the constrained PER model were derived on simulated $q\bar{q} \rightarrow Z \rightarrow 4\ell$ and $q\bar{q} \rightarrow ZZ^{(*)} \rightarrow 4\ell$ events. These fits are shown in Figure 112, demonstrating that the model can work very effectively at describing the $Z \rightarrow 4\ell$ resonance.

9.5.2 Property measurement of the $Z \rightarrow 4\ell$ resonance

To validate the method, maximum likelihood fits were performed on an effective dataset containing the same events shown in Figure 112, with a small additional component from Z+Jets. The $Z \rightarrow 4\ell$ component of the dataset is weighted to the SM cross section, while the Z+Jets component is weighted to its data-driven estimate. One drawback to this dataset is that it has limited statistics compared to the $H \rightarrow 4\ell$ simulation used in Section 9.2, as far fewer events were simulated for this resonance.

The configurations of parameters were handled in an analogous way to the $H \rightarrow 4\ell$ resonance: for fits on m_Z , μ_Z was left to float; for fits on μ_Z , m_Z was left to float; for fits on Γ_Z , both μ_Z and m_Z were left to float. As no systematics were applied for the $Z \rightarrow 4\ell$ study, there are no nuisance parameters included in the fit. Likelihood scans of m_Z , μ_Z , and Γ_Z on this dataset are found in Figure 113.

Scans of m_Z exhibit residual biases smaller than 5 MeV. This value is smaller than that of the $H \rightarrow 4\ell$ analysis, both because this measurement is nearly background-free, and because the tuning parameters are derived on the very same sample as used for this test, so the closure arises more automatically. For the same reason, μ_Z is guaranteed to show no residual bias. The width, Γ_Z , is the most difficult of




Figure 112: Agreement between MC simulation and constrained PER model on $m_{4\ell}^{\text{cons}}$ distribution of the $Z \rightarrow 4\ell$ resonance, in each of the analysis categories. The resonance parameters m_Z , Γ_Z are set at their best measured values from [16].

the three, as it is highly sensitive to the dataset statistics. The limited $Z \rightarrow 4\ell$ MC statistics cause some fluctuation in this case, as large as 130 MeV in the $2\mu 2e$ category, which has the fewest simulated events. The combined measurement of Γ_Z on the simulated data converges within 15 MeV to the reference value, approximately 2.5% the statistical resolution of the measurement.

The checks on the MC simulation demonstrate the PER model is an effective fit strategy, and is prepared to be used on $Z \rightarrow 4\ell$ data. Table 26 summarizes the numbers of events observed in each category, compared to the SM expectation. Although no systematic uncertainties were evaluated or used in the $Z \rightarrow 4\ell$ measurements, the table includes a conservative 10% uncertainty on the expected number of $Z \rightarrow 4\ell$ events. This is roughly consistent with the total uncertainty on the $q\bar{q} \rightarrow ZZ^{(*)}$




Figure 113: Likelihood scans of m_Z , μ_Z , Γ_Z on the effective dataset using the constrained PER model. Dashed lines indicate the scans include only statistical uncertainty; no systematics are included.

background yield in the $H \rightarrow 4\ell$ analysis, though in the case presented here the uncertainty is not decomposed into all its sources. The uncertainty on the (small) Z+Jets contribution comes from the data-driven estimate. The likelihood scans of m_Z , μ_Z , and Γ_Z on the data are found in Figure 114.

The distribution of data events is shown in Figure 115 together with the constrained PER model drawn at the best fit values of m_Z , μ_Z , and with Γ_Z set to its SM value.

The results of all $Z \rightarrow 4\ell$ likelihood scans are summarized in Table 27. All measurements are within 2σ of the reference values, while most are within 1σ . In particular, the measured m_Z shows a $+1.2\sigma$ deviation, μ_Z shows a -0.2σ deviation,

Table 26: Expected and observed yields in the mass window $78 \text{ GeV} \leq m_{4\ell}^{\text{cons}} \leq 104 \text{ GeV}$ for the $Z \rightarrow 4\ell$ validation study.

category	$Z \rightarrow 4\ell$	Z+Jets	observed
4μ	48.8 ± 4.8	0.20 ± 0.05	45
$4e$	17.0 ± 1.7	0.28 ± 0.07	18
$2\mu 2e$	10.6 ± 1.0	0.29 ± 0.07	8
$2e 2\mu$	18.2 ± 1.9	0.20 ± 0.05	23
combined	94.4 ± 9.4	0.97 ± 0.25	94




Figure 114: Likelihood scans of m_Z , μ_Z , Γ_Z on the data using the constrained PER model. Dashed lines indicate the scans include only statistical uncertainty; no systematics are included.

and Γ_Z shows a $+1.3\sigma$ deviation from their respective reference values. Recall that the quoted uncertainties are conservative as they include only statistical uncertainty,




Figure 115: Plot of the $m_{4\ell}^{\text{cons}}$ distribution of the $Z \rightarrow 4\ell$ resonance in data, in each of the analysis categories. The constrained PER model is plotted after the fit, with parameters m_Z , μ_Z set at their best-fit values.

and so these quoted deviations would be reduced in a more complete analysis. The measurements of the $Z \rightarrow 4\ell$ resonance in this section serve to verify the readiness of the constrained PER model for measuring the $H \rightarrow 4\ell$ resonance in data.

Table 27: Summary table of results from constrained PER $Z \rightarrow 4\ell$ likelihood scans of m_Z , μ_Z , and Γ_Z , in MC simulated data, and in the real data. The “SM value” of m_Z , Γ_Z indicates the current best-measured value, used as input in other SM calculations, not a fundamental prediction of the SM. Units in the table match those of the reference values.

Category	MC simulated dataset	observed data
m_Z	SM value $m_Z = 91.1876$ GeV	
4μ	$91.19^{+0.41}_{-0.41}$	$91.46^{+0.42}_{-0.41}$
$4e$	$91.18^{+1.03}_{-1.03}$	$91.75^{+1.08}_{-1.06}$
$2\mu 2e$	$91.18^{+1.11}_{-1.11}$	$91.31^{+1.62}_{-1.33}$
$2e 2\mu$	$91.18^{+0.90}_{-0.90}$	$92.49^{+0.91}_{-0.94}$
combined	$91.19^{+0.34}_{-0.34}$	$91.62^{+0.35}_{-0.35}$
μ_Z	SM value $\mu_Z = 1.00$	
4μ	$1.00^{+0.15}_{-0.14}$	$0.92^{+0.14}_{-0.13}$
$4e$	$1.00^{+0.26}_{-0.22}$	$1.03^{+0.27}_{-0.23}$
$2\mu 2e$	$1.00^{+0.34}_{-0.28}$	$0.69^{+0.30}_{-0.23}$
$2e 2\mu$	$1.00^{+0.25}_{-0.22}$	$1.25^{+0.28}_{-0.24}$
combined	$1.00^{+0.11}_{-0.10}$	$0.98^{+0.11}_{-0.10}$
Γ_Z	SM value $\Gamma_Z = 2.4952$ GeV	
4μ	$2.42^{+0.95}_{-0.75}$	$2.57^{+1.10}_{-0.84}$
$4e$	$2.63^{+2.70}_{-1.63}$	$4.69^{+3.92}_{-2.28}$
$2\mu 2e$	$2.35^{+2.97}_{-1.70}$	$5.83^{+7.65}_{-3.47}$
$2e 2\mu$	$2.35^{+2.20}_{-1.39}$	$5.05^{+3.68}_{-2.35}$
combined	$2.43^{+0.74}_{-0.62}$	$3.61^{+1.04}_{-0.84}$

9.6 Results with data

The preceding sections are all intended to engender confidence in the measurement strategy. Section 9.2 demonstrated that the constrained PER model introduces no apparent bias within the available MC simulated statistics, and provided expected resolutions, assuming the SM, for each measurement to be performed. Section 9.3 broke down all sources of systematic uncertainty entering the analysis, ranking them according to their importance for each measurement. Section 9.4 tested the measurement strategy with realistic statistics like those expected in the data, demonstrating

that the method is robust in the existing regime of low statistics. As a final validation step, Section 9.5 showed a measurement of m_Z , and Γ_Z using the constrained PER model on the $Z \rightarrow 4\ell$ resonance in MC simulation and data, demonstrating the high degree of flexibility of the method. In each case where appropriate, the Template method and unconstrained PER method were validated alongside the constrained PER method both as cross-checks of the final measurement, and as benchmarks of performance. With this confidence in hand, we are prepared to do the measurement on the data.

The expected and observed yields in each category are recorded in Table 28. The expected signal yields are given for a $m_H = 125$ GeV SM Higgs boson.

Table 28: Table of expected and observed event yields. For the $H \rightarrow 4\ell$ signal yield, the values are given for a $m_H = 125$ GeV Higgs.

category	qqZZ	ggZZ	VVV	Z+Jets	$H \rightarrow 4\ell$
4μ	15.18 ± 1.49	0.65 ± 0.39	0.18 ± 0.01	1.56 ± 0.40	20.86 ± 1.35
$4e$	6.39 ± 0.74	0.29 ± 0.18	0.13 ± 0.01	1.97 ± 0.47	11.12 ± 0.92
$2\mu 2e$	6.83 ± 0.77	0.20 ± 0.12	0.11 ± 0.01	2.05 ± 0.47	11.31 ± 0.87
$2e 2\mu$	10.59 ± 1.02	0.31 ± 0.19	0.18 ± 0.01	1.56 ± 0.40	14.79 ± 0.89
combined	39.00 ± 3.84	1.45 ± 0.88	0.60 ± 0.05	7.14 ± 1.23	58.08 ± 3.73
		total expected	observed		
4μ			38.43 ± 2.40	38	
$4e$			19.91 ± 1.56	24	
$2\mu 2e$			20.49 ± 1.51	26	
$2e 2\mu$			27.43 ± 1.60	34	
combined			106.27 ± 6.39	122	

The distributions in data of $m_{4\ell}^{\text{cons}}$ are shown in Figure 116, overlaid on top of the expected distributions from the difference sources of signal and background. For the signal, the constrained PER PDF is drawn with m_H and $\sigma_{H \rightarrow 4\ell}$ is set to their best fit values (to be presented shortly), and Γ_H to its SM value. The corresponding distributions of $m_{4\ell}^{\text{fsr}}$ can be found in Appendix C.1. These plots deliberately bin the model to allow comparison bin-by-bin; the unbinned post-fit signal+background PDFs plotted on top of the same data can be found in Appendix C.2, for both the constrained and unconstrained PER model.

Perhaps more interesting for the PER model is the 2D distribution of $m_{4\ell}^{\text{cons}}$ vs. σ_{eff} (defined in Section 7.3.6). This distribution, colour-coded by category, is shown in Figure 117. Note that a narrower window of $m_{4\ell}^{\text{cons}}$ is used here than in the fit, so that the plot puts more emphasis on the signal events. To interpret this plot, consider that each event will effectively “pull” m_H toward its $m_{4\ell}^{\text{cons}}$ value (actually, slightly above its $m_{4\ell}^{\text{cons}}$ value, because the detector response is not symmetric). The strength of this pull is dictated by the resolution, so that events with smaller values of σ_{eff} have stronger pulls, and events with larger values have weaker pulls. Marked on the plot are the average σ_{eff} from MC simulation, and its standard deviation. In this case, the observed $4e$ and $2\mu 2e$ events are worse than the average expected from simulation, while the other two categories are approximately equal. Using this plot is very useful in interpreting the difference between the PER method, which uses this information, and the Template method, which assumes every event in a given category has the same resolution. Of course, the effective resolution is not a complete picture when describing the detector response: each event has a unique, generally complex and asymmetric response distribution. The full response distributions of each event in the data, as modelled using the constrained PER method, are collected in Appendix C.5. These distributions provide everything needed to interpret the results in data.

The likelihood scans for m_H in the full 36.1 fb^{-1} dataset are found in Figure 118, and the results are summarized in Table 29. The combined best fit value of m_H of the two cross-checks agree with the constrained PER method to within 140 MeV, a value consistent with the level of disagreement that could be expected from the same dataset based on the toy experiments (second column of Table 25). Some of differences between the constrained PER and Template method in individual channels are worth highlighting, as they serve well to differentiate the two methods.

- In the 4μ category, the events observed have quite average resolutions, according to Figure 117a, and so the two methods agree very closely. However, as seen in Figure 116a, the events are spread more broadly in $m_{4\ell}^{\text{cons}}$ than expected from MC simulation, which results in a worse-than-expected m_H resolution in this category.
- In the $4e$ category, the constrained PER returns a higher m_H value than the Template method. This is a result of the distribution of events in Figure 117b




Figure 116: Distribution in data of m_{4l}^{cons} in the four analysis categories drawn with the distribution predicted by the signalbackground model, using the constrained PER model with m_H and $\sigma_{H \rightarrow 4l}$ set to their best fit values. The plotted axis range is reduced from the fit range by 0.5 GeV on either side to allow tidier binning.




Figure 117: Distribution of $m_{4\ell}^{\text{cons}}$ vs σ_{eff} in data. Each datum is marked by a round point. Triangles instead of points indicate the actual value is off the plotted range. The horizontal line represents the average σ_{eff} of the events in simulation, and the hashed area represents the standard deviation of those values. Events below the line have a stronger-than-average pull on the m_H value, while events above the line have a weaker-than-average pull on m_H .

- both methods are drawn by the clump of events near 122 GeV, but the constrained PER method gives extra importance to the well-measured event near 124 GeV, which pulls the best fit value upwards.

- Similarly, in the $2\mu 2e$ category, the constrained PER is pulled to a lower fit

value, and given a larger negative uncertainty, by a cluster of well-measured events between 123-124 GeV, seen in Figure 117c. The Template method cannot “see” the resolution of these events, so is only pulled by the $m_{4\ell}^{\text{cons}}$ values, of which the majority fall in a slightly higher range.

- In the $2e2\mu$ category, a cluster of well-measured events around 126-128 GeV in Figure 117d pull the m_H value upwards in the constrained PER method, compared to the Template method.

These results demonstrate one of the key design principles of the PER method, that better-measured events are given more importance in the fit result than poorly measured events. Whether this results in a more accurate measurement is often up to good luck in the events that are observed, but as Table 25 demonstrated, the odds are better than even.

The likelihood scans for $\sigma_{H \rightarrow 4\ell}$ are found in Figure 119, with the results summarized in Table 29. The constrained PER and Template methods agree exactly on the final best fit value, with slight differences in the uncertainty, and unconstrained PER method returns only a slightly higher value. The fit results in each category also show excellent agreement between all methods.

The likelihood scans for Γ_H are found in Figure 120. The value of the test statistic q_0 , defined in (113), can be read as the left-most value of the combined (blue) curve, at $\Gamma_H = 4.1\text{MeV}$. For the constrained (unconstrained) PER model, it is $q_0 = 4.14$ (2.09). Using (115), the probability of SM fluctuations producing as large or larger an observed width is $p_0 = 0.0089$ (0.041). These values can be checked by considering the distribution of $\hat{\Gamma}_H$ from toys, in Figure 107 which suggest a similar probability for such an excess. As these p_0 values do not meet the criteria for identifying an excess, an upper limit is set on Γ_H , as prescribed in Section 9.1. The q_μ test statistic is equal to the plotted likelihood curve everywhere to the right of $\hat{\Gamma}_H$. Most the observed upper limit-setting power comes from the 4μ and $2e2\mu$ categories, whereas in the $4e$ and $2\mu 2e$ categories, only extremely loose limits are able to be set.

The fit results from the constrained PER model in Table 29 represent the ultimate product of this thesis. The uncertainties in the table can be decomposed into their statistic and systematic components trivially using Table 21, and $\sigma_{\text{stat+sys}}^2 = \sigma_{\text{stat}}^2 + \sigma_{\text{sys}}^2$, which is how the final results of Chapter 10 are presented.




Figure 118: Scans of the NLL on parameter m_H in data. Subfigures show the scans using the different signal models as labelled. The NLL scans including (not including) systematics are shown with solid (dashed) lines.

The results can be compared with existing measurement and the SM predictions using

$$z[\sigma] = \frac{x_1 - x_2}{\sqrt{\sigma_{x_1}^2 + \sigma_{x_2}^2}}, \quad (126)$$

where for asymmetric uncertainties, the σ_{x_i} is taken in the direction of the other measurement. The comparisons are signed to express the direction of compatibility, and expressed in units of Gaussian σ . These comparisons are shown only for the constrained PER results, although the results of the cross checks have similar compatibility. Some of the measurements being compared have correlated systematic uncertainties, although this is not taken into account in the simple consistency check as the measurements in this thesis are highly statistically dominated.




Figure 119: Scans of the NLL on parameter $\sigma_{H \rightarrow 4\ell}$ in data. Subfigures show the scans using the different signal models as labelled. The NLL scans including (not including) systematics are shown with solid (dashed) lines.

The results of the constrained PER measurements are highly compatible with existing measurements. The Higgs boson mass measured in this thesis, $m_H = 124.61 \pm 0.44 \text{ GeV}$, is compatible with:

- The ATLAS and CMS Run 1 combined m_H measurement [8],
 $m_H = 125.09 \pm 0.24 \text{ GeV}$, with compatibility -0.96σ .
- The ATLAS Run 1 $H \rightarrow ZZ^{(*)} \rightarrow 4\ell$ m_H measurement [7],
 $m_H = 124.51 \pm 0.52 \text{ GeV}$, with compatibility $+0.15\sigma$.
- A more recent CMS $H \rightarrow ZZ^{(*)} \rightarrow 4\ell$ m_H measurement [101],
 $m_H = 124.50^{+0.48}_{-0.46} \text{ GeV}$, with compatibility $+0.17\sigma$.




Figure 120: Scans of the NLL on parameter Γ_H in data. Subfigures show the scans using the different signal models as labelled. The NLL scans including (not including) systematics are shown with solid (dashed) lines.

The mass measurement presented in this thesis exceeds the precision of the two $H \rightarrow 4\ell$ measurements entering into [8], and that in [101], making it the most precise single-experiment Higgs boson mass measurement in the $H \rightarrow 4\ell$ channel to date.

The measurement of $\sigma_{H \rightarrow 4\ell}$ requires slightly more care in its comparison with existing measurements, as the existing measurements can be framed in different ways and must first be converted to a common reference, in this case $\sigma_{H \rightarrow 4\ell}$ as it is defined in this thesis. The measurement can also only be compared with Run 2 measurements taken at $\sqrt{s} = 13$ TeV collisions, as the cross section depends on the collision energy. The measurement of this thesis, $\sigma_{H \rightarrow 4\ell} = 8.55^{+1.38}_{-1.26}$ fb, is compatible with:

- The SM prediction for $\sigma_{H \rightarrow 4\ell}$ at the observed Higgs boson mass [21],
 $\sigma_{H \rightarrow 4\ell}^{SM} = 6.63 \pm 0.30$ fb, to within $+1.57\sigma$.
- The most recent ATLAS $H \rightarrow ZZ^{(*)} \rightarrow 4\ell$ measurement [5], which measured $\sigma_{H,tot} = 81^{+18}_{-16}$ pb. Using the branching ratio $\text{BR}(H \rightarrow 4\ell)$ as in this conference note to translate this to a $\sigma_{H \rightarrow 4\ell}$ value, the approximate corresponding measurement is $\sigma_{H \rightarrow 4\ell} = 10 \pm 2$ fb, with which the measurement in this thesis is compatible within -0.61σ . Note that these two measurements use an overlapping dataset, so this is not a wholly independent comparison.
- The most recent CMS measurement [101], which measured the quantity
 $\mu_{H \rightarrow 4\ell} = N_{H \rightarrow 4\ell}^{obs} / N_{H \rightarrow 4\ell}^{SM} = 0.99^{+0.33}_{-0.26}$. This can be translated, using the SM

Table 29: Summary table of the final data fit results on m_H , $\sigma_{H \rightarrow 4\ell}$, Γ_H for the three signal models. Compare these numbers with expected resolutions from simulation in Table 18. Upper limits for Γ_H are given at 95% CL.

	constrained PER	Template	unconstrained PER
m_H [GeV]			
4μ	$124.15^{+0.64}_{-0.63}$	$124.22^{+0.61}_{-0.60}$	$123.95^{+0.69}_{-0.73}$
$4e$	$123.90^{+1.20}_{-1.13}$	$123.24^{+1.74}_{-1.47}$	$123.86^{+1.24}_{-1.27}$
$2\mu 2e$	$126.01^{+1.34}_{-2.13}$	$126.68^{+1.05}_{-1.20}$	$125.37^{+1.02}_{-1.07}$
$2e 2\mu$	$125.33^{+0.62}_{-0.65}$	$124.65^{+0.72}_{-0.64}$	$125.81^{+0.67}_{-0.68}$
combined	$124.61^{+0.44}_{-0.44}$	$124.47^{+0.46}_{-0.44}$	$124.76^{+0.44}_{-0.45}$
$\sigma_{H \rightarrow 4\ell}$ [fb]			
4μ	$7.39^{+2.05}_{-1.78}$	$7.58^{+2.08}_{-1.82}$	$7.51^{+2.09}_{-1.83}$
$4e$	$8.08^{+3.13}_{-2.60}$	$7.93^{+3.21}_{-2.65}$	$7.95^{+3.11}_{-2.58}$
$2\mu 2e$	$8.49^{+3.13}_{-2.63}$	$7.92^{+2.99}_{-2.50}$	$9.35^{+3.15}_{-2.67}$
$2e 2\mu$	$10.96^{+2.83}_{-2.48}$	$11.00^{+2.87}_{-2.51}$	$11.18^{+2.84}_{-2.49}$
combined	$8.55^{+1.38}_{-1.26}$	$8.54^{+1.38}_{-1.27}$	$8.72^{+1.38}_{-1.27}$
Γ_H [GeV]			
4μ	< 3.73	n.a.	< 4.53
$4e$	< 30.58	n.a.	< 23.55
$2\mu 2e$	< 18.23	n.a.	< 13.75
$2e 2\mu$	< 3.83	n.a.	< 3.87
combined	< 3.65	n.a.	< 3.71

predictions at $m_H = 125.09$ GeV to $\sigma_{H \rightarrow 4\ell} = 6.8^{+2.3}_{-1.8}$ fb, with which the measurement in this thesis is compatible within $+0.68\sigma$.

The measurement of $\sigma_{H \rightarrow 4\ell}$ at $\sqrt{s} = 13$ TeV is also the most precise measurement of this quantity to date. The above-mentioned measurements are the first performed at $\sqrt{s} = 13$ TeV, with a smaller dataset than the one used in this thesis. Measurement strategies focused on the cross section, rather than designed around measuring m_H , may be able to exceed to precision of this thesis using the same dataset.

The observed Γ_H upper limit, $\Gamma_H < 3.65$ GeV at 95% CL, is not the most stringent set so far. As demonstrated using pseudo-experiments in Section 9.4, the PER technique has the potential to set a very stringent limit (as far as direct limits

go) on Γ_H , potentially even below 1 GeV. However, by the same demonstration, in some datasets only very weak limits can be set. The observed data happen to the latter, not the former, and so only a weak limit is set on Γ_H . For comparison, other direct Γ_H measurements are:

- By ATLAS using the Run 1 $H \rightarrow ZZ^{(*)} \rightarrow 4\ell$ data and the PER technique in its early stages [7], an upper limit of $\Gamma_H < 2.6$ GeV was set.
- A CMS combination of $H \rightarrow ZZ^{(*)} \rightarrow 4\ell$ and $H \rightarrow \gamma\gamma$ using the Run 1 data [89] set an upper limit $\Gamma_H < 2.6$ GeV (in $H \rightarrow ZZ^{(*)} \rightarrow 4\ell$ and $H \rightarrow \gamma\gamma$ alone it was 2.4 GeV and 3.4 GeV, respectively).

All these limits pale in comparison to the strength of *indirect* limits on Γ_H , of which the strongest [102] is $\Gamma_H < 13$ MeV. These limits are generally accompanied by significantly more model assumptions than the *direct* limit set by this thesis, and so the two can still be complementary.

The measurements in this chapter serve to refine our knowledge of one of nature's unknown parameters, m_H , and put the SM to the test with a precision measurement of $\sigma_{H \rightarrow 4\ell}$, and upper limit set on Γ_H . No significant deviations from the SM, or from existing measurements, were observed.

Chapter 10

Conclusion

In July of 2012, the ATLAS and CMS collaborations jointly announced the discovery of a new particle consistent with the Standard Model Higgs boson. This elusive particle holds a critical role in the Standard Model, where it is responsible for giving fundamental particles their mass. To understand whether the newly discovered particle was indeed the predicted Higgs boson, a wide range of property measurements have been conducted by ATLAS and CMS using the full Run 1 LHC dataset, all of which are so far consistent with Standard Model predictions within their uncertainties.

In the higher energy collisions of Run 2 at the LHC, Higgs bosons are being produced at more than twice the rate of Run 1, offering a significant opportunity to measure the Higgs more precisely than ever before. One of the most powerful tools for Higgs property measurement is the $H \rightarrow ZZ^{(*)} \rightarrow 4\ell$ decay channel because all the final state particles from the Higgs decay can be fully and precisely reconstructed. For this reason it was one of the discovery channels for the Higgs, and continues to be one of the most powerful for Higgs boson property measurements. Precision measurement of the properties of the Higgs boson continue to be an effective tool for testing the Standard Model to its limits, and garnering clues for what may lay beyond it.

This thesis presented a measurement of the Higgs boson mass, m_H , its cross section times branching ratio to four electrons or muons, $\sigma_H \times \text{BR}(H \rightarrow 4\ell = e, \mu)$, and an upper limit on its natural width, Γ_H . While the mass is an unknown free parameter of the Standard Model, the other two are predicted once the mass is known. The mass measurement has ramifications in a huge range of Higgs analyses from ATLAS and CMS, as all the Higgs coupling and branching ratio predictions depend on it.

The measurements were conducted using an unbinned maximum likelihood fit on the four-lepton invariant mass, $m_{4\ell}$, providing statistical models for the signal and background distributions. The signal model used was the *per event response* method, which decouples the detector resolution component of the $m_{4\ell}$ distribution from the

physics component. It works by building a model of lepton resolutions dependent on their kinematics, and propagates this analytically as a Gaussian mixture to the final detector $m_{4\ell}$ response. As the response is predicted on an event-by-event basis, well-measured events are given more importance in the fit than poorly-measured events. The measurements in this thesis were all conducted using the full Run 2 dataset collected by ATLAS so far, 36.1 fb^{-1} of p - p collisions at $\sqrt{s} = 13 \text{ TeV}$. Using this technique and dataset, the Higgs boson properties were measured to be

$$\begin{aligned} m_H &= 124.61^{+0.44}_{-0.44}(\text{stat.})^{+0.06}_{-0.06}(\text{syst.}) \text{ GeV}, \\ \sigma_H \times \text{BR}(H \rightarrow 4\ell) &= 8.55^{+1.23}_{-1.15}(\text{stat.})^{+0.63}_{-0.52}(\text{syst.}) \text{ fb}, \\ \Gamma_H &< 3.65(\text{stat.}) \oplus 0.19(\text{syst.}) \text{ GeV at } 95\% \text{ CL}. \end{aligned}$$

The measured values are consistent within their uncertainties with all existing measurements and the Standard Model predictions. The m_H measurement exceeds the precision of any single-experiment m_H measurement in the $H \rightarrow 4\ell$ channel by ATLAS or CMS conducted to this point. The measurement of $\sigma_{H \rightarrow 4\ell}$ also exceeds the precision of other such measurements at $\sqrt{s} = 13 \text{ TeV}$. The Γ_H upper limit is not the most stringent set, but complements existing limits. All three are the first measurements of their kind using the full 2015 and 2016 $\sqrt{s} = 13 \text{ TeV}$ LHC collision datasets.

The results in this thesis are just a small piece of the Higgs boson property measurements being performed by ATLAS and CMS. Despite this, they represent a valuable contribution to the field as they, in particular the m_H measurement, help refine the SM predictions against which other Higgs property measurements are compared. Tests such as these could eventually uncover whether the Higgs boson near 125 GeV is indeed the Higgs boson of the Standard Model, or is a manifestation of an underlying theory.

Appendix A

Backgrounds supplementary plots

This appendix supplements Chapter 6, providing plots of the smoothing used to create PDFs of the unconstrained mass, $m_{4\ell}^{\text{fsr}}$. The unconstrained smoothing results for qqZZ, ggZZ, VVV , and are found in Figure 121, Figure 122, Figure 123, respectively. Due to the large uncertainty on the Z+Jets background yield and shape, the $m_{4\ell}^{\text{cons}}$ and $m_{4\ell}^{\text{fsr}}$ distributions are considered to be interchangeable.




Figure 121: Smoothed distribution of $m_{4\ell}^{\text{fsr}}$ from simulated $q\bar{q} \rightarrow ZZ^{(*)}$ events. The black points represent the simulation, the solid blue line is the kernel smoothing result used as the PDF.




Figure 122: Smoothed distribution of $m_{4\ell}^{\text{fsr}}$ from simulated $gg \rightarrow ZZ^{(*)}$ events. The black points represent the simulation, the solid blue line is the kernel smoothing result used as the PDF.




Figure 123: Smoothed distribution of $m_{4\ell}^{\text{fsr}}$ from simulated $VVV, t\bar{t} + Z$ events. The black points represent the simulation, the solid blue line is the kernel smoothing result used as the PDF.

Appendix B

Extra Discussion of the Signal Model

This appendix contains supplementary information concerning the signal models. Appendix B.1 verifies the stability of the Template method’s morphing function. Appendix B.2 provides an explanation of the non-conventional definition of the lepton response. Appendix B.3 presents supplementary information on the individual lepton response fits. Appendix B.4 presents the overall behaviour of the fit parameters in the response model. Appendix B.5 presents a test of the kinematic model independence of the PER model.

B.1 Stability of Template interpolation

One potential drawback to the Template method is that the model could become over-tuned to the particular input samples used to build it. To demonstrate that this *is not* the case, individual signal mass points were removed from the morphing, and the prediction $P(m_{4\ell}^{\text{cons}}|m_H)$ with and without those mass points were compared. This comparison is shown in Figure 124, where the Template drawn is in each category at $m_H = 124.5 \text{ GeV}$ for three distinct cases: when all simulated mass points are used to build the morphed shape, when all except the $m_H = 124 \text{ GeV}$ sample are used to build the morphed shape, and when all except the $m_H = 125 \text{ GeV}$ sample are used to build the morphed shape. From this test, we can see that the interpolation is highly stable, even when individual mass points are removed. The actual uncertainty on the interpolation procedure is smaller than the variations from these plots, as they by definition use less input information.




Figure 124: Demonstration of the stability of the Template interpolation (morphing). Each plot shows the Template PDF drawn at $m_H = 124.5 \text{ GeV}$, when the input to the morphing is altered along the following lines: all samples used (black curve), all except the 124 GeV sample (blue curve), and all except the 125 GeV sample (red curve). At the bottom, the ratio of each with respect to the black curve is shown.

B.2 On the definition of x_i

In the derivation of Section 7.3.1, the reader may notice that the definition of the lepton response, replicated from (75) as

$$x_i = \frac{E_i^{\text{fsr}} - E_i^{\text{true}}}{E_i^{\text{fsr}}},$$

does not match the typical definition, of

$$x_i^{\text{typ}} = \frac{E_i^{\text{fsr}} - E_i^{\text{true}}}{E_i^{\text{true}}} \quad (127)$$

or similar. It may even appear that the choice was arbitrary. However, the definition made in (75) is quite deliberate. If the same derivation is continued using x_i^{typ} instead of x_i , then (79), which is replicated here

$$m_{4\ell}^{\text{fsr}} - m_{4\ell}^{\text{true}} = \frac{1}{2m_{4\ell}^{\text{fsr}}} \sum_i \left(\sum_{j \neq i} (m_{ij}^{\text{fsr}})^2 \right) x_i,$$

becomes instead

$$m_{4\ell}^{\text{fsr}} - m_{4\ell}^{\text{true}} = \frac{1}{2m_{4\ell}^{\text{fsr}}} \sum_i \left(\sum_{j \neq i} (m_{ij}^{\text{true}})^2 \right) x_i^{\text{typ}}. \quad (128)$$

While this new expression is mathematically valid, it is unfortunately of no use in predicting the response. To calculate the response of a reconstructed event using the right hand side of this expression, we would need to know the true lepton kinematics to calculate m_{ij}^{true} . While this could serve as a validation tool in simulation, where the true kinematics are known, it would be useless on data, where only reconstructed kinematics are available. It is the atypical definition of lepton energy response in (75) which allows the predicted mass response to be calculated.

B.3 Response model fits

This section contains supplementary plots of the lepton response model fits in various kinematic bins. The plots shown here represent only a small subset of the greater-than 8000 fits used to build the models. Table 30 summarizes where to find the plots for the per-event response model with and without Z Mass Constraint. Also find

Figure 131 which shows the χ^2/n distributions for the constrained lepton response model.

Table 30: Reference table for response model example fits

muon response	electron response
unconstrained	(more in Section 7.3.2)
Figure 125	Figure 126
constrained	
Figure 127	Figure 129
Figure 128	Figure 130




Figure 125: Example fits of unconstrained muon energy response (2/2). The distributions are built from MC simulation (black points) fit using a sum of three Gaussians (blue curve). Each plot shows the same fit in linear and log scale. The kinematic bins and the fit parameters are printed on each plot. See Table 30 for more plots.




Figure 126: Example fits of unconstrained electron energy response (2/2). The distributions are built from MC simulation (black points) fit using a sum of three Gaussians (blue curve). Each plot shows the same fit in linear and log scale. The kinematic bins and the fit parameters are printed on each plot. See Table 30 for more plots.




Figure 127: Example fits of constrained muon energy response (1/2). The distributions are built from MC simulation (black points) fit using a sum of three Gaussians (blue curve). Each plot shows the same fit in linear and log scale. The kinematic bins and the fit parameters are printed on each plot. See Table 30 for more plots.




Figure 128: Example fits of constrained muon energy response (2/2). The distributions are built from MC simulation (black points) fit using a sum of three Gaussians (blue curve). Each plot shows the same fit in linear and log scale. The kinematic bins and the fit parameters are printed on each plot. See Table 30 for more plots.




Figure 129: Example fits of constrained electron energy response (1/2). The distributions are built from MC simulation (black points) fit using a sum of three Gaussians (blue curve). Each plot shows the same fit in linear and log scale. The kinematic bins and the fit parameters are printed on each plot. See Table 30 for more plots.




Figure 130: Example fits of constrained electron energy response (2/2). The distributions are built from MC simulation (black points) fit using a sum of three Gaussians (blue curve). Each plot shows the same fit in linear and log scale. The kinematic bins and the fit parameters are printed on each plot. See Table 30 for more plots.




Figure 131: Distribution of χ^2/n from constrained lepton response model fits.

B.4 Response model parameters

This section presents how the parameters of $P(x_i)$ (defined in (83)) describing the lepton response PDF depend on selected Z , $|\eta|$, and p_T . Parameter σ_1 is the most easily recognized as the lepton energy resolution, while the others characterize additional properties of the response distribution and can't be so easily interpreted in isolation. Table 31 summarizes where to find each of the parameters for two models - with Z Mass Constraint (“constrained”) and without (“unconstrained”).

Note that the parameter w_2 contained in these plots is a recursive weight, while in (83) it is absolute (meaning $w_2^{\text{plotted}} = w_2^{(83)} / (1 - w_1)$). The parameter w_1 is unaffected by this definition. The p_T range used in these plots is reduced as compared to the full model coverage to emphasize the region where most $H \rightarrow 4\ell$ leptons are reconstructed.

Table 31: Reference table for response model plots for each of the parameters of (83) for model with and without Z Mass Constraint.

parameter	unconstrained model	constrained model
σ_1	Figure 132	Figure 140
σ_2	Figure 133	Figure 141
σ_3	Figure 134	Figure 142
μ_1	Figure 135	Figure 143
μ_2	Figure 136	Figure 144
μ_3	Figure 137	Figure 145
w_1	Figure 138	Figure 146
w_2	Figure 139	Figure 147




Figure 132: Behaviour of parameter σ_1 in unconstrained lepton response models. See Table 31 for other parameters.




Figure 133: Behaviour of parameter σ_2 in unconstrained lepton response models. See Table 31 for other parameters.




Figure 134: Behaviour of parameter σ_3 in unconstrained lepton response models. See Table 31 for other parameters.




Figure 135: Behaviour of parameter μ_1 in unconstrained lepton response models. See Table 31 for other parameters.




Figure 136: Behaviour of parameter μ_2 in unconstrained lepton response models. See Table 31 for other parameters.




Figure 137: Behaviour of parameter μ_3 in unconstrained lepton response models. See Table 31 for other parameters.




Figure 138: Behaviour of parameter w_1 in unconstrained lepton response models. See Table 31 for other parameters.




Figure 139: Behaviour of parameter w_2 in unconstrained lepton response models. See Table 31 for other parameters.




Figure 140: Behaviour of parameter σ_1 in constrained lepton response models. See Table 31 for other parameters.




Figure 141: Behaviour of parameter σ_2 in constrained lepton response models. See Table 31 for other parameters.




Figure 142: Behaviour of parameter σ_3 in constrained lepton response models. See Table 31 for other parameters.




Figure 143: Behaviour of parameter μ_1 in constrained lepton response models. See Table 31 for other parameters.




Figure 144: Behaviour of parameter μ_2 in constrained lepton response models. See Table 31 for other parameters.




Figure 145: Behaviour of parameter μ_3 in constrained lepton response models. See Table 31 for other parameters.




Figure 146: Behaviour of parameter w_1 in constrained lepton response models. See Table 31 for other parameters.




Figure 147: Behaviour of parameter w_2 in constrained lepton response models. See Table 31 for other parameters.

B.5 Model Independence of the PER signal model

Theoretical uncertainties in general give rise to kinematic differences of the $H \rightarrow 4\ell$ events, which will result in an altered invariant mass distribution. These can come from parton-density-function uncertainties, or more directly from calculation of the Higgs boson cross section. If the signal model does not account for these differences, a bias can be introduced. This appendix is meant to demonstrate that the Per-Event Response model is *independent* of any kinematic hypothesis for the signal model, and contrast it with the Template method which by definition assumes that the kinematic distributions of the events in the MC simulated samples used to built the templates are correct, and is therefore sensitive to the aforementioned theoretical uncertainties. To do so, simulated ggF samples are used in three configurations:

1. The nominal with all weights according to the SM,
2. a very hard Higgs, which is the nominal sample re-weighted by $w' = w \times 0.03 (p_T^{4\ell}/GeV)$
3. a very forward Higgs, which is the nominal re-weighted by $w' = w \times (y^{4\ell})$

The overall weight is arbitrary for this test, and so all samples are re-weighted to the same total cross section. The PER model, and Template model are drawn on top of these three MC configurations in Figure 148, using the $4e$ channel as an example. Note that in all cases, the models are drawn “as nominal”, and no fit is done. As can be seen from the plots, the Template PDF is identical in all cases; for a given m_H , its shape is fixed. However, the PER method *adapts dynamically* to the kinematics in each event, and succeeds at describing the shape. So any alteration of event kinematics from a new theory, or higher-order calculation, etc., will not affect the best fit mass value produced by the PER method. For this reason, it is called “model independent”, depending only on accurate description of detector effects.

The variations presented in this appendix are not meant to be physically reasonable variations. They are extreme examples to demonstrate (i) how the reconstructed distribution depends on the kinematics of the $H \rightarrow 4\ell$ events, and (ii) how the PER signal model automatically adapts to these kinematics to be able to describe arbitrarily altered kinematic models. This demonstration shows that the PER technique is as model-independent as possible.




Figure 148: Per-event Response model (red line) and Template model (blue line) drawn on top of ggF $m_H = 125$ GeV Higgs MC (black points) in the 4e category. The bottom of each plot shows the ratio of the MC and Template model to the PER model. The three plots show the MC simulation as compared to each model, in each of the different variations described in the text. The PER model adapts dynamically to the changing kinematic distributions, while the Template model is static and fails to describe the altered distributions.

Appendix C

Statistics and Measurement Supplement

This appendix supplements several aspects of Chapter 9. Appendix C.1 shows data distributions using $m_{4\ell}^{\text{fsr}}$. Appendix C.2 shows the data distributions with the fitted S+B model on top. Appendix C.3 shows supplementary plots of the toy experiments, presenting many of the same plots from the main body of the thesis separated by analysis category. Appendix C.4 shows the $Z \rightarrow 4\ell$ analysis conducted using $m_{4\ell}^{\text{fsr}}$ instead of $m_{4\ell}^{\text{cons}}$. Appendix C.5 shows the response distribution of $m_{4\ell}^{\text{cons}}$ for every recorded data event in the m_H fit range.

C.1 Data distributions of $m_{4\ell}^{\text{fsr}}$

The distributions in data of $m_{4\ell}^{\text{fsr}}$ are found in Figure 149, together with expectations from each source. For the signal, the unconstrained PER model is drawn with m_H and $\sigma_{H \rightarrow 4\ell}$ set to their best fit values. The 2D distributions of $m_{4\ell}^{\text{fsr}}$ and σ_{eff} are found in Figure 150.




Figure 149: Distribution in data of $m_{4\ell}^{\text{fsr}}$ in the four analysis categories drawn with the distribution predicted by the signalbackground model, using the constrained PER model with m_H and $\sigma_{H \rightarrow 4\ell}$ set to their best fit values. The plotted axis range is reduced from the fit range by 0.5 GeV on either side to allow tidier binning.




Figure 150: Distribution of $m_{4\ell}^{\text{fsr}}$ vs σ_{eff} in data. Each datum is marked by a round point. Triangles instead of points indicate the actual value is off the plotted range. The horizontal line represents the average σ_{eff} of the events in simulation, and the hashed area represents the standard deviation of those values. Events below the line have a stronger-than-average pull on the m_H value, while events above the line have a weaker-than-average pull on m_H .

C.2 Post-fit PDF

The data distributions of $m_{4\ell}^{\text{cons}}$ with the combined signal+background PDF of the constrained PER model, with m_H , $\sigma_{H\rightarrow 4\ell}$ set to their final best fit values, and Γ_H to its SM value, are shown in Figure 151. The corresponding plots using $m_{4\ell}^{\text{fsr}}$ and the unconstrained PER model are shown in Figure 152.




Figure 151: Distribution in data of m_{4l}^{cons} in the four analysis categories drawn with the full signal+background model, using the constrained PER model with m_H and $\sigma_{H \rightarrow 4l}$ set to their best fit values. The plotted axis range is reduced from the fit range by 0.5 GeV on either side to allow tidier binning.




Figure 152: Distribution in data of $m_{4\ell}^{\text{fsr}}$ in the four analysis categories drawn with the full signal+background model, using the unconstrained PER model with m_H and $\sigma_{H \rightarrow 4\ell}$ set to their best fit values. The plotted axis range is reduced from the fit range by 0.5 GeV on either side to allow tidier binning.

C.3 More results from pseudo-experiments

This section contains supplementary plots to Section 9.4, where tests of the three prepared methods were performed using pseudo-experiments.

To complement Figure 105, the fit results and pulls of m_H in the toy datasets separated by category are shown for the constrained PER, Template, and unconstrained PER models in Figure 153, Figure 154, and Figure 155, respectively. Correspondingly, the per-category results of toy fits of $\sigma_{H \rightarrow 4\ell}$, shown for all category combined in Figure 106, are shown for the same three methods in Figure 156, Figure 157, and Figure 158, respectively. The per-category results for Γ_H , which can only be fit using the PER technique, are shown for the constrained and unconstrained PER models in Figure 159 and Figure 160, respectively.




Figure 153: Results of m_H fits using the constrained PER model, per category. Shown are the best fit values $\hat{m}_H$ (left of each subfigure), and the pulls p_{m_H} (right) in each category. The raw distribution of results from the toys is shown in blue, and a Gaussian fit is shown with a red line.




Figure 154: Results of m_H fits using the Template model, per category. Shown are the best fit values $\hat{m}_H$ (left of each subfigure), and the pulls p_{m_H} (right) in each category. The raw distribution of results from the toys is shown in blue, and a Gaussian fit is shown with a red line.




Figure 155: Results of m_H fits using the unconstrained PER model, per category. Shown are the best fit values $\hat{m}_H$ (left of each subfigure), and the pulls p_{m_H} (right) in each category. The raw distribution of results from the toys is shown in blue, and a Gaussian fit is shown with a red line.




Figure 156: Results of $\sigma_{H \rightarrow 4\ell}$ fits using the constrained PER model, per category. Shown are the best fit values $\hat{\sigma}_{H \rightarrow 4\ell}$ (left of each subfigure), and the pulls $p_{\sigma_{H \rightarrow 4\ell}}$ (right) in each category. The raw distribution of results from the toys is shown in blue, and a Gaussian fit is shown with a red line.




Figure 157: Results of $\sigma_{H \rightarrow 4\ell}$ fits using the Template model, per category. Shown are the best fit values $\hat{\sigma}_{H \rightarrow 4\ell}$ (left of each subfigure), and the pulls $p_{\sigma_{H \rightarrow 4\ell}}$ (right) in each category. The raw distribution of results from the toys is shown in blue, and a Gaussian fit is shown with a red line.




Figure 158: Results of $\sigma_{H \rightarrow 4\ell}$ fits using the unconstrained PER model, per category. Shown are the best fit values $\sigma_{H \rightarrow 4\ell}$ (left of each subfigure), and the pulls $p_{\sigma_{H \rightarrow 4\ell}}$ (right) in each category. The raw distribution of results from the toys is shown in blue, and a Gaussian fit is shown with a red line.




Figure 159: Results of Γ_H fits using the constrained PER model, per category. Shown are the best fit values $\hat{\Gamma}_H$ (left of each subfigure), and the 95% CL upper limit (right) in each category. The raw distribution of results from the toys is shown in blue. Note the log axis and large first bin in the left-hand plots.




Figure 160: Results of Γ_H fits using the unconstrained PER model, per category. Shown are the best fit values $\hat{\Gamma}_H$ (left of each subfigure), and the 95% CL upper limit (right) in each category. The raw distribution of results from the toys is shown in blue. Note the log axis and large first bin in the left-hand plots.

To complement Figure 108, the 2D correlation of $\hat{m}_H$ between the constrained PER and Template models, separated by category, are shown in Figure 161. Those for $\sigma_{H \rightarrow 4\ell}$ are shown in Figure 162.




Figure 161: Correlation plots of $\hat{m}_H$ between the constrained PER model, and the Template model, from pseudo-experiments in each analysis category. The diagonal line indicates where both methods agree exactly.

To complement Figure 109, the difference in best fit m_H between the constrained PER and Template models, separated by category, is shown in Figure 163. Those for $\sigma_{H \rightarrow 4\ell}$ are shown in Figure 164.

To complement Figure 110, the actual fit accuracy (based on the distance to the true value) of the constrained PER and Template models are compared for fits of m_H in Figure 165, and of $\sigma_{H \rightarrow 4\ell}$ Figure 166.




Figure 162: Correlation plots of $\hat{\sigma}_{H \rightarrow 4\ell}$ between the constrained PER model, and the Template model, from pseudo-experiments in each analysis category. The diagonal line indicates where both methods agree exactly.




Figure 163: Difference in best fit values of m_H between the constrained PER model, and the Template model, from pseudo-experiments in each analysis category. The standard deviation, recorded on the plots as RMS_{hist} , gives an indication of how much the two models may agree or disagree when fit on the data.




Figure 164: Difference in best fit values of $\sigma_{H \rightarrow 4\ell}$ between the constrained PER model, and the Template model, from pseudo-experiments in each analysis category. The standard deviation, recorded on the plots as RMS_{hist} , gives an indication of how much the two models may agree or disagree when fit on the data.




Figure 165: Comparison of distance from the best fit value to the true value of m_H one a toy-by-toy basis. In toys to the left of zero, the constrained PER model produces a more accurate fit result than the Template model. In toys to the right of zero, the Template method was more accurate.




Figure 166: Comparison of distance from the best fit value to the true value of $\sigma_{H \rightarrow 4\ell}$ one a toy-by-toy basis. In toys to the left of zero, the constrained PER model produces a more accurate fit result than the Template model. In toys to the right of zero, the Template method was more accurate.

C.4 $Z \rightarrow 4\ell$ analysis using $m_{4\ell}^{\text{fsr}}$

The $Z \rightarrow 4\ell$ analysis using $m_{4\ell}^{\text{fsr}}$ proceeds exactly as for $m_{4\ell}^{\text{cons}}$ detailed in Section 9.5, using the same $m_{4\ell}^{\text{fsr}}$ window, using all the same samples for the tuning fit (separate tuning parameters) and for the effective dataset, and the same configuration of free and fixed variables during the fits. With these points considered, the likelihood scans of m_Z , μ_Z , and Γ_Z on this dataset are found in Figure 167. The expected and observed event yields for $m_{4\ell}^{\text{fsr}}$ are found in Table 32. The likelihood scans on data are found in Figure 168. The data distribution of $m_{4\ell}^{\text{fsr}}$ is plotted in Figure 169 with the unconstrained PER model plotted at the best fit m_Z , μ_Z , with Γ_Z set to its SM value. A summary of the expected and observed results are found in Table 33, which confirms that the ZMC has a relatively minor effect on the $Z \rightarrow 4\ell$ analysis, and does not improve the resolution of this measurement.

Table 32: Expected and observed yields in the mass window $78 \text{ GeV} \leq m_{4\ell}^{\text{fsr}} \leq 104 \text{ GeV}$ for the $Z \rightarrow 4\ell$ unconstrained validation study.

category	$Z \rightarrow 4\ell$	Z+Jets	observed
4μ	48.8 ± 4.8	0.20 ± 0.05	45
$4e$	17.1 ± 1.7	0.28 ± 0.07	18
$2\mu 2e$	10.6 ± 1.0	0.29 ± 0.07	8
$2e 2\mu$	18.4 ± 1.9	0.20 ± 0.05	23
combined	94.4 ± 9.4	0.97 ± 0.25	94




Figure 167: Likelihood scans of m_Z , μ_Z , Γ_Z on the effective dataset using the unconstrained PER model. Dashed lines indicate the scans include only statistical uncertainty; no systematics are included.




Figure 168: Likelihood scans of m_Z , μ_Z , Γ_Z on the data using the unconstrained PER model. Dashed lines indicate the scans include only statistical uncertainty; no systematics are included.




Figure 169: Plot of the $m_{4\ell}^{\text{fsr}}$ distribution of the $Z \rightarrow 4\ell$ resonance in data, in each of the analysis categories. The unconstrained PER model is plotted after the fit, with parameters m_Z , μ_Z set at their best-fit values.

Table 33: Summary table of results from unconstrained PER $Z \rightarrow 4\ell$ likelihood scans of m_Z , μ_Z , and Γ_Z , in MC simulated data, and in the real data. The “SM value” of m_Z, Γ_Z indicates the current best-measured value, used as input in other SM calculations, not a fundamental prediction of the SM. Units in the table match those of the reference values.

Category	MC simulated dataset	observed data
m_Z	SM value $m_Z = 91.1876 \text{ GeV}$	
4μ	$91.19^{+0.41}_{-0.41}$	$91.44^{+0.41}_{-0.41}$
$4e$	$91.18^{+1.03}_{-1.03}$	$91.75^{+1.05}_{-1.03}$
$2\mu 2e$	$91.20^{+1.09}_{-1.10}$	$91.45^{+1.46}_{-1.35}$
$2e 2\mu$	$91.18^{+0.89}_{-0.89}$	$92.36^{+0.88}_{-0.91}$
combined	$91.19^{+0.33}_{-0.33}$	$91.61^{+0.35}_{-0.34}$
μ_Z	SM value $\mu_Z = 1.00$	
4μ	$1.00^{+0.15}_{-0.14}$	$0.91^{+0.14}_{-0.13}$
$4e$	$1.00^{+0.26}_{-0.22}$	$1.03^{+0.27}_{-0.23}$
$2\mu 2e$	$1.00^{+0.34}_{-0.28}$	$0.69^{+0.30}_{-0.23}$
$2e 2\mu$	$1.00^{+0.25}_{-0.22}$	$1.24^{+0.28}_{-0.24}$
combined	$1.00^{+0.11}_{-0.10}$	$0.97^{+0.11}_{-0.10}$
Γ_Z	SM value $\Gamma_Z = 2.4952 \text{ GeV}$	
4μ	$2.46^{+0.95}_{-0.75}$	$2.46^{+1.07}_{-0.81}$
$4e$	$2.70^{+2.64}_{-1.61}$	$4.04^{+3.42}_{-2.01}$
$2\mu 2e$	$2.38^{+2.96}_{-1.69}$	$5.45^{+7.50}_{-3.41}$
$2e 2\mu$	$2.60^{+2.21}_{-1.41}$	$5.07^{+3.46}_{-2.25}$
combined	$2.50^{+0.74}_{-0.62}$	$3.43^{+0.99}_{-0.80}$

C.5 Response distribution of data events

The rarity of $H \rightarrow 4\ell$ events may limit the statistical precision of the measurement, but at least one consolation to this fact is that the number of events in data is small enough to examine one at a time. This allows one to see the dramatic effect of the Per-event response technique: each event has a uniquely predicted signal distribution, for a given mass hypothesis. This appendix contains the per-event signal PDF, drawn for each event drawn with m_H set to its combined best fit value. Table 34 indicates where to find the PDFs, which are organized by category of the events. On each plot the actual value of $m_{4\ell}^{\text{cons}}$ is marked by a red line. The height of the red line (the value of the PDF at that $m_{4\ell}$ value) indicates roughly how likely such an event is, if it came from a signal decay.

Table 34: Summary of figures containing Per-event signal PDFs for the data events. The constrained PER signal model is drawn with m_H set to the best fit value, 124.61 GeV.

category	figures references
4μ	Figure 170 and Figure 171
$4e$	Figure 172 and Figure 173
$2\mu 2e$	Figure 174 and Figure 175
$2\mu 2e$	Figure 176 and Figure 177




Figure 170: Per-event signal PDF of 4μ events in data according to the constrained PER model (1/2). Printed on each plot is the reconstructed mass of the event after ZMC and FSR recovery, and the effective resolution σ_{eff} as it is defined in Section 7.3.6. The dashed red line is drawn at the observed value of $m_{4\ell}^{\text{cons}}$ for this event.




Figure 171: Per-event signal PDF of 4μ events in data according to the constrained PER model (2/2). Printed on each plot is the reconstructed mass of the event after ZMC and FSR recovery, and the effective resolution σ_{eff} as it is defined in Section 7.3.6. The dashed red line is drawn at the observed value of $m_{4\ell}^{\text{cons}}$ for this event.




Figure 172: Per-event signal PDF of 4e events in data according to the constrained PER model (1/2). Printed on each plot is the reconstructed mass of the event after ZMC and FSR recovery, and the effective resolution σ_{eff} as it is defined in Section 7.3.6. The dashed red line is drawn at the observed value of $m_{4\ell}^{\text{cons}}$ for this event.




Figure 173: Per-event signal PDF of 4e events in data according to the constrained PER model (2/2). Printed on each plot is the reconstructed mass of the event after ZMC and FSR recovery, and the effective resolution σ_{eff} as it is defined in Section 7.3.6. The dashed red line is drawn at the observed value of $m_{4\ell}^{\text{cons}}$ for this event.




Figure 174: Per-event signal PDF of $2\mu 2e$ events in data according to the constrained PER model (1/2). Printed on each plot is the reconstructed mass of the event after ZMC and FSR recovery, and the effective resolution σ_{eff} as it is defined in Section 7.3.6. The dashed red line is drawn at the observed value of m_{4l}^{cons} for this event.




Figure 175: Per-event signal PDF of $2\mu 2e$ events in data according to the constrained PER model (2/2). Printed on each plot is the reconstructed mass of the event after ZMC and FSR recovery, and the effective resolution σ_{eff} as it is defined in Section 7.3.6. The dashed red line is drawn at the observed value of $m_{4\ell}^{\text{cons}}$ for this event.




Figure 176: Per-event signal PDF of $2e2\mu$ events in data according to the constrained PER model (1/2). Printed on each plot is the reconstructed mass of the event after ZMC and FSR recovery, and the effective resolution σ_{eff} as it is defined in Section 7.3.6. The dashed red line is drawn at the observed value of $m_{4\ell}^{\text{cons}}$ for this event.




Figure 177: Per-event signal PDF of $2e2\mu$ events in data according to the constrained PER model (2/2). Printed on each plot is the reconstructed mass of the event after ZMC and FSR recovery, and the effective resolution σ_{eff} as it is defined in Section 7.3.6. The dashed red line is drawn at the observed value of $m_{4\ell}^{\text{cons}}$ for this event.

List of References

- [1] P. W. Higgs, “Broken symmetries and the masses of gauge bosons,” *Phys. Rev. Lett.*, vol. 13, pp. 508–509, 16 1964.
- [2] ATLAS Collaboration, “Observation of a new particle in the search for the Standard Model Higgs boson with the ATLAS detector at the LHC,” *Physics Letters B*, vol. 716, no. 1, 2012.
- [3] CMS Collaboration, “Observation of a new boson at a mass of 125 GeV with the CMS experiment at the LHC,” *Phys. Lett. B*, vol. 716, no. 1, 2012.
- [4] ATLAS Collaboration, “Muon reconstruction performance of the ATLAS detector in protonproton collision data at $\sqrt{s}=13$ TeV,” *Eur. Phys. J.*, vol. C76, no. 5, p. 292, 2016. arXiv: [1603.05598 \[hep-ex\]](#).
- [5] ATLAS Collaboration, “Study of the Higgs boson properties and search for high-mass scalar resonances in the $H \rightarrow ZZ^* \rightarrow 4\ell$ decay channel at $\sqrt{s} = 13$ TeV with the ATLAS detector,” CERN, Geneva, Tech. Rep. ATLAS-CONF-2016-079, 2016.
- [6] “Measurements of the Higgs boson production cross section at 7, 8 and 13 TeV centre-of-mass energies and search for new physics at 13 TeV in the $H \rightarrow ZZ^* \rightarrow \ell^+ \ell \ell'^+ \ell'$ final state with the ATLAS detector,” CERN, Geneva, Tech. Rep. ATLAS-CONF-2015-059, 2015.
- [7] ATLAS Collaboration, “Measurement of the Higgs boson mass from the $H \rightarrow \gamma\gamma$ and $H \rightarrow ZZ^* \rightarrow 4\ell$ channels with the ATLAS detector using 25 fb $^{-1}$ of pp collision data,” *Phys. Rev.*, vol. D90, no. 5, p. 052 004, 2014. arXiv: [1406.3827 \[hep-ex\]](#).
- [8] ATLAS Collaboration and CMS Collaboration, “Combined Measurement of the Higgs Boson Mass in pp Collisions at $\sqrt{s} = 7$ and 8 TeV with the ATLAS and CMS Experiments,” *Phys. Rev. Lett.*, vol. 114, p. 191 803, 19 2015.

- [9] G. Cree, “Search for a high mass neutral Higgs boson using the ATLAS detector,” 2015.
- [10] D. Griffiths, *Introduction to Elementary Particles*. Wiley-VCH, 2008.
- [11] CMS Collaboration, “Measurement of the top quark mass using proton-proton data at $\sqrt{s} = 7$ and 8 TeV,” 2015. arXiv: [1509.04044 \[hep-ex\]](#).
- [12] SNO Collaboration, “Measurement of the Rate of $\nu_e + d \rightarrow p + p + e^-$ Interactions Produced by 8B Solar Neutrinos at the Sudbury Neutrino Observatory,” *Phys. Rev. Lett.*, vol. 87, p. 071 301, 7 2001.
- [13] Super-Kamiokande Collaboration, “Evidence for oscillation of atmospheric neutrinos,” *Phys. Rev. Lett.*, vol. 81, pp. 1562–1567, 8 1998.
- [14] B. N. Laboratory. (). Facility upgrades: rhic-ii and erhic, [Online]. Available: <https://www.bnl.gov/rhic/RHIC2.asp>.
- [15] M. E. Peskin and D. V. Schroeder, *An Introduction to Quantum Field Theory*. Westview Press, 1995.
- [16] Particle Data Group, “Review of Particle Physics,” *Chin. Phys.*, vol. C38, p. 090 001, 2014.
- [17] F. Englert and R. Brout, “Broken Symmetry and the Mass of Gauge Vector Mesons,” *Phys. Rev. Lett.*, vol. 13, pp. 321–323, 1964.
- [18] G. S. Guralnik, C. R. Hagen, and T. W. B. Kibble, “Global Conservation Laws and Massless Particles,” *Phys. Rev. Lett.*, vol. 13, pp. 585–587, 1964.
- [19] L. Alvarez-Gaume and J. Ellis, “Eyes on a prize particle,” *Nature Physics*, pp. 2–3, 2011.
- [20] C. ATLAS, “Studies of the ATLAS potential for Higgs self-coupling measurements at a High Luminosity LHC,” CERN, Geneva, Tech. Rep. ATL-PHYS-PUB-2013-001, 2013.
- [21] D. de Florian *et al.*, “Handbook of LHC Higgs Cross Sections: 4. Deciphering the Nature of the Higgs Sector,” 2016. arXiv: [1610.07922 \[hep-ph\]](#).
- [22] J. D. Lykken, “Beyond the Standard Model,” in *CERN Yellow Report CERN-2010-002, 101-109*, 2010. arXiv: [1005.1676 \[hep-ph\]](#).
- [23] M. Roos, “Dark Matter: The evidence from astronomy, astrophysics and cosmology,” 2010. arXiv: [1001.0316 \[astro-ph.CO\]](#).

- [24] ATLAS Collaboration, “Search for invisible decays of a Higgs boson using vector-boson fusion in pp collisions at $\sqrt{s} = 8$ TeV with the ATLAS detector,” *JHEP*, vol. 01, p. 172, 2016. arXiv: [1508.07869 \[hep-ex\]](#).
- [25] J. F. Gunion, “Extended Higgs sectors,” in *Supersymmetry and unification of fundamental interactions. Proceedings, 10th International Conference, SUSY’02, Hamburg, Germany, June 17-23, 2002*, 2002, pp. 80–103. arXiv: [hep-ph/0212150 \[hep-ph\]](#).
- [26] ATLAS Collaboration, “Study of the Higgs boson properties and search for high-mass scalar resonances in the $H \rightarrow ZZ^* \rightarrow 4\ell$ decay channel at $\sqrt{s} = 13$ TeV with the ATLAS detector,” CERN, Geneva, Tech. Rep. ATLAS-CONF-2016-079, 2016.
- [27] J. Unger, M. Kobel, and D. Stoeckinger, “Higgs Boson Production in the Atmosphere of the Earth,” Presented 20 Nov 2014, PhD thesis, Dresden, Tech. U., 2014.
- [28] O. S. Brning, P. Collier, P. Lebrun, S. Myers, R. Ostojic, J. Poole, and P. Proudlock, *LHC Design Report*. Geneva: CERN, 2004.
- [29] ATLAS Collaboration, “The ATLAS Experiment at the CERN Large Hadron Collider,” *J. Instrum.*, vol. 3, S08003. 437 p, 2008, Also published by CERN Geneva in 2010.
- [30] CMS Collaboration, “The CMS Experiment at the CERN LHC,” *JINST*, vol. 3, S08004, 2008.
- [31] ALICE Collaboration, “The ALICE experiment at the CERN LHC,” *JINST*, vol. 3, S08002, 2008.
- [32] LHCb Collaboration, “The LHCb Detector at the LHC,” *JINST*, vol. 3, S08005, 2008.
- [33] L. Evans and P. Bryant, “Lhc machine,” *Journal of Instrumentation*, vol. 3, no. 08, S08001, 2008.
- [34] Nobel Media, *The Nobel Prize in Physics 1984*.
- [35] UA1 Collaboration, “Further evidence for charged intermediate vector bosons at the {SPS} collider,” *Physics Letters B*, vol. 129, no. 34, pp. 273 –282, 1983.
- [36] UA1 Collaboration, “Observation of muonic Z^0 -decay at the pp collider,” *Physics Letters B*, vol. 147, no. 1, pp. 241 –248, 1984.

- [37] ATLAS Collaboration, “Improved luminosity determination in pp collisions at $\sqrt{s} = 7$ TeV using the ATLAS detector at the LHC,” *Eur. Phys. J. C*, vol. 73, no. arXiv:1302.4393. CERN-PH-EP-2013-026, 2518. 40 p, 2013.
- [38] CMS Collaboration, “Search for quark compositeness in dijet angular distributions from pp collisions at $\sqrt{s} = 7$ TeV,” *JHEP*, vol. 05, p. 055, 2012. arXiv: [1202.5535 \[hep-ex\]](#).
- [39] S. Chatrchyan *et al.*, “Measurement of the Underlying Event Activity at the LHC with $\sqrt{s} = 7$ TeV and Comparison with $\sqrt{s} = 0.9$ TeV,” *JHEP*, vol. 09, p. 109, 2011. arXiv: [1107.0330 \[hep-ex\]](#).
- [40] J. M. Butterworth, G. Dissertori, and G. P. Salam, “Hard Processes in Proton-Proton Collisions at the Large Hadron Collider,” *Ann. Rev. Nucl. Part. Sci.*, vol. 62, pp. 387–405, 2012. arXiv: [1202.0583 \[hep-ex\]](#).
- [41] J. Pequeno, “Computer generated image of the whole ATLAS detector,” 2008.
- [42] E. Daw, *Lecture 7 - rapidity and pseudorapidity*, 2012.
- [43] ATLAS Collaboration, “Track Reconstruction Performance of the ATLAS Inner Detector at $\sqrt{s} = 13$ TeV,” CERN, Geneva, Tech. Rep. ATL-PHYS-PUB-2015-018, 2015.
- [44] ATLAS Collaboration, “A study of the material in the ATLAS inner detector using secondary hadronic interactions,” *JINST*, vol. 7, P01013, 2012. arXiv: [1110.6191 \[hep-ex\]](#).
- [45] M Capeans, G Darbo, K Einsweiler, M Elsing, T Flick, M Garcia-Sciveres, C Gemme, H Pernegger, O Rohne, and R Vuillermet, “ATLAS Insertable B-Layer Technical Design Report,” CERN, Geneva, Tech. Rep. CERN-LHCC-2010-013. ATLAS-TDR-19, 2010.
- [46] J. Pequeno, “Computer Generated image of the ATLAS calorimeter,” 2008.
- [47] S. Palestini, “The muon spectrometer of the atlas experiment,” *Nuclear Physics B-Proceedings Supplements*, vol. 125, pp. 337–345, 2003.
- [48] J. Pequeno, “Computer generated image of the ATLAS Muons subsystem,” 2008.
- [49] ATLAS Collaboration, “Standalone vertex finding in the ATLAS muon spectrometer,” *JINST*, vol. 9, P02001, 2014. arXiv: [1311.7070 \[physics.ins-det\]](#).

- [50] T Kawamoto, S Vlachos, L Pontecorvo, J Dubbert, G Mikenberg, P Iengo, C Dallapiccola, C Amelung, L Levinson, R Richter, and D Lellouch, “New Small Wheel Technical Design Report,” CERN, Geneva, Tech. Rep. CERN-LHCC-2013-006. ATLAS-TDR-020, 2013, ATLAS New Small Wheel Technical Design Report.
- [51] ATLAS Collaboration, “Expected Performance of the ATLAS Experiment - Detector, Trigger and Physics,” 2009. arXiv: [0901.0512 \[hep-ex\]](#).
- [52] S. Weinzierl, “Introduction to Monte Carlo methods,” 2000. arXiv: [hep-ph/0006269 \[hep-ph\]](#).
- [53] S. A. et al., “Geant4a simulation toolkit,” *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment*, vol. 506, no. 3, pp. 250–303, 2003.
- [54] S. Alioli, P. Nason, C. Oleari, and E. Re, “NLO Higgs boson production via gluon fusion matched with shower in POWHEG,” *JHEP*, vol. 04, p. 002, 2009. arXiv: [0812.0578 \[hep-ph\]](#).
- [55] P. Nason and C. Oleari, “NLO Higgs boson production via vector-boson fusion matched with shower in POWHEG,” *JHEP*, vol. 02, p. 037, 2010. arXiv: [0911.5299 \[hep-ph\]](#).
- [56] D. de Florian, G. Ferrera, M. Grazzini, and D. Tommasini, “Transverse-momentum resummation: Higgs boson production at the Tevatron and the LHC,” *JHEP*, vol. 11, p. 064, 2011. arXiv: [1109.2109 \[hep-ph\]](#).
- [57] T. Sjostrand, S. Mrenna, and P. Z. Skands, “A Brief Introduction to PYTHIA 8.1,” *Comput. Phys. Commun.*, vol. 178, pp. 852–867, 2008. arXiv: [0710.3820 \[hep-ph\]](#).
- [58] D. J. Lange, “The EvtGen particle decay simulation package,” *Nucl. Instrum. Meth.*, vol. A462, pp. 152–155, 2001.
- [59] J. Alwall, R. Frederix, S. Frixione, V. Hirschi, F. Maltoni, O. Mattelaer, H. S. Shao, T. Stelzer, P. Torrielli, and M. Zaro, “The automated computation of tree-level and next-to-leading order differential cross sections, and their matching to parton shower simulations,” *JHEP*, vol. 07, p. 079, 2014. arXiv: [1405.0301 \[hep-ph\]](#).
- [60] M. Bahr *et al.*, “Herwig++ Physics and Manual,” *Eur. Phys. J.*, vol. C58, pp. 639–707, 2008. arXiv: [0803.0883 \[hep-ph\]](#).

- [61] T. Gleisberg, S. Hoeche, F. Krauss, M. Schonherr, S. Schumann, F. Siegert, and J. Winter, “Event generation with SHERPA 1.1,” *JHEP*, vol. 02, p. 007, 2009. arXiv: [0811.4622 \[hep-ph\]](#).
- [62] T. Sjostrand, S. Mrenna, and P. Z. Skands, “PYTHIA 6.4 Physics and Manual,” *JHEP*, vol. 05, p. 026, 2006. arXiv: [hep-ph/0603175 \[hep-ph\]](#).
- [63] P. Golonka and Z. Was, “PHOTOS Monte Carlo: A Precision tool for QED corrections in Z and W decays,” *Eur. Phys. J.*, vol. C45, pp. 97–107, 2006. arXiv: [hep-ph/0506026 \[hep-ph\]](#).
- [64] S. Jadach, Z. Was, R. Decker, and J. H. Kuhn, “The tau decay library TAUOLA: Version 2.4,” *Comput. Phys. Commun.*, vol. 76, pp. 361–380, 1993.
- [65] ATLAS Collaboration, “Measurements of the Total and Differential Higgs Boson Production Cross Sections Combining the $H \rightarrow \gamma\gamma$ and $H \rightarrow ZZ^{(*)} \rightarrow 4\ell$ Decay Channels at $\sqrt{s} = 8\text{TeV}$ with the ATLAS Detector,” *Phys. Rev. Lett.*, vol. 115, p. 091 801, 9 2015.
- [66] I Bird, P Buncic, F Carminati, M Cattaneo, P Clarke, I Fisk, M Girone, J Harvey, B Kersevan, P Mato, R Mount, and B Panzer-Steindl, “Update of the Computing Models of the WLCG and the LHC Experiments,” Tech. Rep. CERN-LHCC-2014-014. LCG-TDR-002, 2014.
- [67] Y. Nakahama, “The atlas trigger system: ready for run-2,” *Journal of Physics: Conference Series*, vol. 664, no. 8, p. 082 037, 2015.
- [68] ATLAS Collaboration, “Charged-particle distributions in $\sqrt{s}=13\text{ TeV}$ pp interactions measured with the ATLAS detector at the LHC,” *Phys. Lett.*, vol. B758, pp. 67–88, 2016. arXiv: [1602.01633 \[hep-ex\]](#).
- [69] ATLAS Collaboration, “Electron and photon energy calibration with the ATLAS detector using LHC Run 1 data,” *Eur. Phys. J.*, vol. C74, no. 10, p. 3071, 2014. arXiv: [1407.5063 \[hep-ex\]](#).
- [70] ATLAS Collaboration, “Electron and photon energy calibration with the ATLAS detector using data collected in 2015 at $\sqrt{s} = 13\text{ TeV}$,” CERN, Geneva, Tech. Rep. ATL-PHYS-PUB-2016-015, 2016.
- [71] A. Keith Morley, “Electron Bremsstrahlung Studies and Track Based Alignment of the ATLAS Detector,” PhD thesis, Melbourne U., 2010.
- [72] ATLAS Collaboration, “Improved electron reconstruction in ATLAS using the Gaussian Sum Filter-based model for bremsstrahlung,” CERN, Geneva, Tech. Rep. ATLAS-CONF-2012-047, 2012.

- [73] T. M. Atkinson, “Electron reconstruction with the ATLAS inner detector,” Presented 2006, PhD thesis, Melbourne U., 2006.
- [74] S. Kullback and R. A. Leibler, “On Information and Sufficiency,” *Ann. Math. Statist.*, vol. 22, no. 1, pp. 79–86, Mar. 1951.
- [75] “*auto-parallelization and auto-vectorization*”, <https://msdn.microsoft.com/en-us/library/hh872235.aspx>, Accessed: 2016-10-17.
- [76] J. Illingworth and J. Kittler, “A survey of the hough transform,” *Computer Vision, Graphics, and Image Processing*, vol. 44, no. 1, pp. 87–116, 1988.
- [77] ATLAS Collaboration, “Jet energy measurement and its systematic uncertainty in proton–proton collisions at $\sqrt{s} = 7\text{TeV}$ with the ATLAS detector,” *The European Physical Journal C*, vol. 75, no. 1, pp. 1–101, 2015.
- [78] A. Roodman, “Blind analysis in particle physics,” *eConf*, vol. C030908, TUIT001, 2003, [,166(2003)]. arXiv: [physics/0312102](https://arxiv.org/abs/physics/0312102) [[physics.data-an](#)].
- [79] ATLAS Collaboration, “Electron efficiency measurements with the ATLAS detector using 2012 LHC proton-proton collision data,” 2016. arXiv: [1612.01456](https://arxiv.org/abs/1612.01456) [[hep-ex](#)].
- [80] ATLAS Collaboration, “Selection of jets produced in 13TeV proton-proton collisions with the ATLAS detector,” CERN, Geneva, Tech. Rep. ATLAS-CONF-2015-029, 2015.
- [81] S. Schael *et al.*, “Precision electroweak measurements on the Z resonance,” *Phys. Rept.*, vol. 427, pp. 257–454, 2006. arXiv: [hep-ex/0509008](https://arxiv.org/abs/hep-ex/0509008) [[hep-ex](#)].
- [82] V. Ippolito, “Measurement of the properties of the new particle observed within the search for the Higgs boson in the $H \rightarrow ZZ^{(*)} \rightarrow 4\ell$ channel at ATLAS,” PhD thesis, Rome U., 2013-12-06.
- [83] G. Cowan, *Statistical Data Analysis*. Oxford University Press, Oxford, 1998.
- [84] K. S. Cranmer, “Kernel estimation in high-energy physics,” *Comput. Phys. Commun.*, vol. 136, pp. 198–207, 2001. arXiv: [hep-ex/0011057](https://arxiv.org/abs/hep-ex/0011057) [[hep-ex](#)].
- [85] M. Pivk, “sPLOT: A Quick introduction,” in *Statistical Problems in Particle Physics, Astrophysics and Cosmology (PHYSTAT 05): Proceedings, Oxford, UK, September 12-15, 2005*, 2006, pp. 173–177. arXiv: [physics/0602023](https://arxiv.org/abs/physics/0602023) [[physics](#)].

- [86] UA1 Collaboration, “Intermediate-vector-boson properties at the CERN Super Proton Synchrotron collider,” *Europhys. Lett.*, vol. 1, no. CERN-EP-85-185, 327–345. 23 p, 1986.
- [87] CMS Collaboration, “W-like measurement of the Z boson mass using dimuon events collected in pp collisions at $\sqrt{s} = 7$ TeV,” CERN, Geneva, Tech. Rep. CMS-PAS-SMP-14-007, 2016.
- [88] K. C. Sven Kreiss, $H \rightarrow ZZ^* \rightarrow 4\ell$ *Likelihood in ATLAS*, Likelihoods for the LHC Searches, CERN Workshop, 2012.
- [89] CMS Collaboration, “Precise determination of the mass of the Higgs boson and tests of compatibility of its couplings with the standard model predictions using proton collisions at 7 and 8 TeV,” *Eur. Phys. J.*, vol. C75, no. 5, p. 212, 2015. arXiv: [1412.8662 \[hep-ex\]](#).
- [90] M. Oreglia, “A Study of the Reactions $\psi' \rightarrow \gamma\gamma\psi$,” PhD thesis, SLAC, 1980.
- [91] W. Spearman, “Measurement of the mass and natural width of the Higgs boson in the H to ZZ to 4ℓ decay channel with the ATLAS detector,” PhD thesis, Harvard U. (main), 2014-09-11.
- [92] “Proposal for truth particle observable definitions in physics measurements,” CERN, Geneva, Tech. Rep. ATL-PHYS-PUB-2015-013, 2015.
- [93] N. D. Gagunashvili, “Comparison of weighted and unweighted histograms,” *ArXiv Physics e-prints*, May 2006. eprint: [physics/0605123](#).
- [94] R. Fernow, *Introduction to experimental particle physics*. Cambridge University Press, 1986.
- [95] CMS Collaboration, “Measurement of the properties of a Higgs boson in the four-lepton final state,” *Phys. Rev.*, vol. D89, no. 9, p. 092 007, 2014. arXiv: [1312.5353 \[hep-ex\]](#).
- [96] M. Grazzini, S. Kallweit, and D. Rathlev, “ZZ production at the LHC: fiducial cross sections and distributions in NNLO QCD,” *Phys. Lett.*, vol. B750, pp. 407–410, 2015. arXiv: [1507.06257 \[hep-ph\]](#).
- [97] G. Cowan, K. Cranmer, E. Gross, and O. Vitells, “Asymptotic formulae for likelihood-based tests of new physics,” *Eur. Phys. J.*, vol. C71, p. 1554, 2011, [Erratum: *Eur. Phys. J.*C73,2501(2013)]. arXiv: [1007.1727 \[physics.data-an\]](#).

- [98] S. Ahn and J. A. Fessler, *Standard Errors of Mean, Variance, and Standard Deviation Estimators*, 2003.
- [99] L. Demortier and L. Lyons, “Everything you always wanted to know about pulls,” The Rockefeller University, University of Oxford, Tech. Rep., 2002.
- [100] Thomas Dorigo and Michael Schmitt, “On the Significance of the Dimuon Mass Bump and the Greedy Bump Bias,” Tech. Rep., 2000.
- [101] “Measurements of properties of the Higgs boson and search for an additional resonance in the four-lepton final state at $\sqrt{s} = 13$ TeV,” CERN, Geneva, Tech. Rep. CMS-PAS-HIG-16-033, 2016.
- [102] CMS Collaboration, “Search for Higgs boson off-shell production in proton-proton collisions at 7 and 8 TeV and derivation of constraints on its total decay width,” *JHEP*, vol. 09, p. 051, 2016. arXiv: [1605.02329 \[hep-ex\]](#).