



University
of Glasgow | Experimental
Particle Physics

School of Physics and Astronomy
University of Glasgow
Glasgow
G12 8QQ

PhD Thesis

Higgs boson studies: associated production with a vector boson and decay into b-quarks using the ATLAS Run-2 dataset.

Dwayne Isaac Patrick Laurence SPITERI

January 5, 2021

Submitted in fulfilment of the requirements for the
Degree of Doctor of Philosophy



Abstract

On the 4th July 2012, the Standard Model of Particle Physics received further validation with the discovery of the Higgs boson; ushering in a new age of Higgs physics. This thesis presents some of my contributions to the current research in this field as a member of the ATLAS experiment at CERN. It explains how the ATLAS experiment fits within the CERN accelerator complex and the structure of the ATLAS detector, leading to a description of some of the work that I did towards the upgrade of its hardware, and my studies on the reconstruction of tracks in the detector.

The thesis goes on to present the results of two analyses I worked on: The $VH, H \rightarrow b\bar{b}$ (read: VHbb) Resolved analysis and the $VH, H \rightarrow b\bar{b}$ Boosted analysis. Specific attention will be drawn to my fit studies in the Resolved analysis, and the new trigger strategy I designed for the Boosted analysis.

The analyses make use of 139 fb^{-1} of proton-proton collision data at the centre-of-mass energy of $\sqrt{s} = 13 \text{ TeV}$, collected by the ATLAS detector between 2015 and 2018, to observe the Higgs boson decay to b -quarks via associated Vector Boson production ($VH, H \rightarrow b\bar{b}$), and go on to provide differential cross-section measurements in bins of the transverse momentum of the vector bosons.

To my rambunctious immediate family: Josette, Lionel, Elene, Stephane and Nikita;
and for Rowan.

Acknowledgements

First and foremost, I would like to thank my supervisors: Aidan Robson and Craig Buttar. Aidan has been a constant source of support and guidance throughout the entirety of the PhD and has helped me to become a better science communicator. In turn I have always respected his advice. Craig always asked to see how things were going and always made time to listen.

Secondly I must extend a general thanks to the members of Glasgow's Particle Physics Experimental (PPE) Group. I have enjoyed and been enriched by many a conversation about physics, politics or life itself especially over lunch or Friday doughnuts, and it is here I shall specifically thank Andy Blue, Lars Eklund, and all the PPE IT support staff.

Thirdly, this PhD would not be possible without the funding of the Science and Technology Facilities Council (STFC) and the University of Glasgow (UoG). I would also like to offer my sincerest gratitudes to: the UoG Graduate School, The PPE Secretaries (Valerie Flood, Karen Hegyi and Jill Borland), and the UK Liaison Office (Karine Boutemy-Amdal, Petra Heemskerk and Andrea Diallo) whose administrative and housing support made my life a lot easier.

Since ATLAS is a collaboration of thousands of people, there have been many people who have been of great help to the physics content of this thesis, especially while I was based at CERN. From my time in Tracking CP Group I have to thank Pierfrancesco Butti, Nick Styles and Matthias Danninger for their invaluable help in my qualification

task and its extension. On the physics analysis front, while the VH , $H \rightarrow b\bar{b}$ family has been a delight to work alongside, I have to thank Giuseppe Callea, Andy Buckley, Thomas Calvet, Hannah Arnold, Stephen Jiggins, Valerio Dao and Tatsuya Masubuchi in particular for their patience in responding to all my questions.

For the inspiration of the structure of my thesis I am in debt to the theses of: Carlos Miguel Vergel Infante, Danilo Enoque Ferreira de Lima and Michael Fenton.

Finally, to the cornucopia of friends that I made during the PhD whether we met in: the Kelvin building; the Glasgow University Catholic Chaplaincy; the Burn Estate; the UK-wide 2016 cohort of particle physicists; or while on my long term attachment to CERN. I have encountered some amazing people from around the world that have genuinely made the last four years of my life a fantastic experience. I shall list some of the more important ones to me here with their given nicknames.

- The two best people I know currently on CMS: DJ Vim and Shiny Scizor.
- My lifelong CoSE PhD buddies: The Two Ronnies ('Sour-doughy' Smaranda and "Stuff you steal is free" Hatton), SomeRandomBloke and Quinnine.
- The members of *Beans on Roast*: The Moth, Steamy, Slayer, Tablet, Lorry, Don Simon, The Conor, Binman (Log), Joannah, TripleJ, Batessey and Professor Pixel.
- The non-members of *Beans on Roast*: Dr Philosotea and Ethan2ol.
- The Lyon Contingent: Sixty, Egg, (I and Leg of) Naoron, Ororo, Ven4, Carl.I.Am and R2rr. Théonnaise. And Tiny Tee.
- The people I met at the CLASHEP and MCNet conferences with special mentions for Doc Floc and the rest of the crew in the Back Seat Voices.
- For the board, badinage, baking and bonding with long-suffering housemates: Jojoba, Cardamom and Special K.

- To the Ostriches (especially continued moral support from Gorleb Roosh and Preecision).
- To the Moroccan, the Turk and the United Statians in 1-R-13 in 2018.
- And to the definition of maverick: WD30.

I acknowledge that this chosen path is not the easiest to tread. It is fraught with mental tolls and I have met many excellent people that have had to take time off or dropped out the doctoral programme altogether. To those who: are struggling but continue resolutely out of sheer stubbornness; on the verge of giving up; going through a dark patch; or just need to vent, if by some strange set of circumstances you ever read this and are in a bad way, I will always be available for a chat.

Contents

Abstract	i
Acknowledgements	iii
List of Figures	xii
List of Tables	xix
Declaration	xx
I High Energy Physics Background	1
1 Introduction	2
1.1 Opening Remarks	2
1.2 The Forces of Nature	3
1.3 The Standard Model	4
2 Standard Model theory	8
2.1 Theory	8
2.2 The theoretical foundations of particle physics	9
2.2.1 Constructing a theory	12
2.2.2 Feynman diagrams	14
2.3 The Standard Model Lagrangian	17
2.3.1 $\mathcal{L}_{\text{Fermions}}$: Nature's building blocks	17
2.3.2 $\mathcal{L}_{\text{Gauge}}$: Nature's forces	18
2.3.3 $\mathcal{L}_{\text{Higgs}}$: Solving a massive problem	20
2.3.4 $\mathcal{L}_{\text{Yukawa}}$: Giving mass to fermions	24
2.4 Probing beyond the Standard Model ($\mathcal{L}_{\text{BSM}}$)	27
2.4.1 Parametrising beyond-the-Standard-Model (BSM) physics	28
2.5 Theoretical conclusion	29
3 CERN and the LHC	30
3.1 CERN	30
3.2 The Large Hadron Collider	33
3.2.1 Why and How: Producing Higgs Bosons at the LHC	37
3.2.1.1 Instantaneous and Integrated Luminosity	38

3.2.1.2	Cross-sections	39
3.2.2	A brief look at what the LHC has done	40
3.2.3	Limitations of the LHC	42
3.2.4	A short glance at what the LHC is going to do	43
II	The ATLAS Experiment	45
4	The ATLAS detector	46
4.1	The ATLAS experiment	46
4.2	The ATLAS detector	47
4.2.1	The ATLAS co-ordinate system	48
4.2.2	Muon Spectrometer	49
4.2.2.1	Tracking	50
4.2.2.2	Triggering	53
4.2.3	Calorimeters	53
4.2.3.1	The Hadronic Liquid Argon Calorimeter	57
4.2.3.2	The Hadronic Tile Calorimeter	57
4.2.3.3	The Forward Calorimeter	57
4.2.3.4	The Electromagnetic Calorimeter	58
4.2.4	Magnet system	59
4.2.4.1	Toroid magnets	60
4.2.4.2	The central solenoid	61
4.2.5	The Inner Tracker	62
4.2.5.1	Transition Radiation Tracker	64
4.2.5.2	Semiconductor Tracker	66
4.2.5.3	Pixel System	69
4.3	How ATLAS identifies particles	72
4.4	The ATLAS trigger	75
4.4.1	Level 1 Trigger	76
4.4.2	High Level Trigger	76
4.5	Tracker upgrade	78
4.5.1	Layout	79
4.5.2	ABC130 chips	80
4.5.3	Data acquisition	81
4.5.4	ABC130 identification	82
4.6	How to make the collected data useful	84
5	Monte Carlo event generation	85
5.1	Monte Carlo integration	85
5.2	Monte Carlo simulation	87
5.3	Generating LHC events	88
5.3.1	Parton distribution functions	90
5.3.2	Matrix element	90
5.3.3	Parton shower	91
5.3.4	Hadronisation	92
5.3.5	Underlying event and multi-parton interactions	92

5.3.6	Pile-up	93
5.3.7	Detector simulation	93
5.3.8	Samples	94
5.4	Calibration of physics objects	94
6	Fake tracks in the ATLAS detector	95
6.1	(Re)Constructing Track Objects	95
6.1.1	Inside-Out Track Reconstruction	96
6.1.1.1	TRT Extensions	99
6.1.2	Outside-In Track Reconstruction	99
6.1.3	Track Contents	99
6.2	Origin of Fake Tracks	101
6.3	Modelling of Fake Tracks	101
6.4	Mis-modelling Recommendations in 2017	104
6.4.1	Extracting Uncertainties	105
6.5	Mis-modelling Recommendations in 2018	109
6.5.1	Additional Cuts on d_0	110
6.5.2	Extracting Uncertainties	112
6.6	Monte Carlo Simulated Tracks	116
6.6.1	Truth Linking	117
6.6.2	Truth Matching Probability (P_{match})	119
6.6.3	Combining Truth Linking and P_{match}	120
6.6.4	Track Barcodes	122
6.6.5	Conclusions	125
7	Object definitions	127
7.1	Low-Level Physics Objects	127
7.1.1	Vertices	128
7.1.2	Calorimeter Topo-clusters and Towers	128
7.1.2.1	Towers	129
7.1.2.2	Topoclusters	129
7.2	Decision Trees and Classifiers	130
7.2.1	Boosted Decision Trees	132
7.3	Leptons	133
7.3.1	Electrons	134
7.3.2	Muons	135
7.3.3	Taus	136
7.4	Missing Transverse Energy: E_T^{miss}	137
7.5	Jets	139
7.5.1	What is a jet	139
7.5.2	Choosing an algorithm	139
7.5.3	Types of Jet	142
7.5.3.1	Track jets	143
7.5.3.2	Calorimeter Jets	144
7.5.4	Jet Grooming	145
7.5.4.1	Pruning	146

	7.5.4.2	Trimming	146
	7.5.5	Jet Calibrations	147
7.6		B-Tagging Jets	148
	7.6.1	Direct B-Tagging	150
	7.6.2	Truth B-Tagging	150
	7.6.3	Hybrid B-Tagging	151
	7.6.4	B-jet Corrections	151
7.7		How will these objects be used	152

III Physics Analysis 153

8	An overview of the VHbb analyses		154
8.1	Why VH, $H \rightarrow b\bar{b}$?		154
8.2	Past analyses		156
	8.2.1	Run-1 analysis - 2014/2015	158
	8.2.2	The first Run-2 analysis - 2017	158
	8.2.3	The second Run-2 analysis - 2018	159
		8.2.3.1 Event selection and categorisation	159
		8.2.3.2 Control regions	161
		8.2.3.3 Backgrounds	162
		8.2.3.4 Analysis techniques	163
		8.2.3.5 The fit	164
		8.2.3.6 Systematic uncertainties	165
		8.2.3.7 Cross-check analyses	165
		8.2.3.8 Results	166
	8.2.4	Run-2 STXS paper - 2019	170
	8.2.5	Parametrising BSM effects	170
		8.2.5.1 Simple template cross-sections (STXS)	172
		8.2.5.2 STXS methodology	173
		8.2.5.3 The 2019 STXS binning	175
		8.2.5.4 Results	178
8.3	Common themes between the current analyses		178
	8.3.1	Common data samples and MC simulations	180
	8.3.2	Modelling	181
		8.3.2.1 Multi-jet background estimation	182
	8.3.3	Common backgrounds	182
		8.3.3.1 V+jets	183
		8.3.3.2 Top and anti-Top ($t\bar{t}$)	184
		8.3.3.3 Single top	186
		8.3.3.4 Multi-jet	187
		8.3.3.5 Diboson	187
	8.3.4	Common workflow	189
8.4	The VHbb Resolved analysis - 2020		192
	8.4.1	Event selection and categorisation	193
		8.4.1.1 0-lepton selection	194

8.4.1.2	1-lepton selection	194
8.4.1.3	2-lepton selection	195
8.4.1.4	Signal regions	195
8.4.1.5	Control regions	196
8.4.2	The analysis BDT	197
8.4.2.1	BDT_R	199
8.4.3	Additional points on background modelling	200
8.4.3.1	Data-driven $t\bar{t}$ estimation	200
8.4.4	Systematic uncertainties	201
8.4.4.1	Background Systematic uncertainties	201
8.4.4.2	Experimental uncertainties	201
8.4.4.3	Signal uncertainties	204
8.4.5	Analysis methodology	205
8.4.6	Cross check analyses	205
8.4.6.1	Diboson Analysis	205
8.4.6.2	Cut-Based (Dijet mass) analysis: CBA	206
8.5	The VHbb Boosted analysis - 2020	206
8.5.1	Event selection and categorisation	207
8.5.1.1	0-lepton selection	207
8.5.1.2	1-lepton selection	208
8.5.1.3	2-lepton selection	209
8.5.1.4	Signal and Control Regions	209
8.5.2	Systematic Uncertainties	211
8.5.2.1	Experimental Uncertainties	211
8.5.3	Analysis methodology	213
8.5.4	Cross-check analyses	213
8.6	The significance of statistics	213
9	The VHbb fits and post-fit modelling investigations	214
9.1	Statistics, Fits and Significances	214
9.1.1	The Null and Alternative Hypothesis	216
9.1.2	The Likelihood and the Test Statistic	217
9.1.3	Measurement and Results	220
9.1.4	Significance and the Asimov Dataset	221
9.2	Likelihood Fit Model from first principles	222
9.2.1	Adding additional systematics	236
9.2.2	Adding additional bins	243
9.3	Systematic Manipulations	247
9.3.1	Eigenvector Decompositions	247
9.3.2	Symmetrising	249
9.3.3	Smoothing	249
9.3.4	Pruning	251
9.4	Post-fit Modelling Investigations	251
9.5	The 0-Lepton Resolved VHbb Fit	252
9.5.1	The 0L Flavour-tagging Fit Study	254

9.5.2	0L Pruning Study	266
9.5.3	Diboson p_T^V Fit Study	273
9.6	Future Changes and Summary	277
10	E_T^{miss} triggers in the 2-Lepton channels	279
10.1	Motivation	279
10.2	Possible options	281
10.2.1	Di-muon trigger addition	281
10.2.2	E_T^{miss} trigger alternative	282
10.2.3	Choice of E_T^{miss} triggers	284
10.3	Implementation	285
10.4	Preliminary trials in the Resolved analysis	287
10.4.1	Results for E_T^{miss} study in the Resolved analysis	290
10.5	Preliminary trials in the Boosted analysis	294
10.5.1	E_T^{miss} Trigger Scale Factors	297
10.5.2	Missing E_T^{miss} Trigger	297
10.6	Final Results in the Boosted Analysis	302
10.6.1	Final Results from Running the Boosted WorkSpace Maker	302
11	Results and Conclusion	306
11.1	The Boosted analysis results	307
11.2	The Resolved analysis results	311
11.3	Future plans for the analysis	314
11.4	Conclusion	316
	Appendices	318
A	From Newtonian to Lagrangian Mechanics	319
B	Axioms of Quantum Mechanics and their consequences	324
C	Basics of Group Theory	327
C.1	Tenets of Group Theory	327
C.2	Types of Group Relevant to Physics	328
C.2.1	Lie Groups	328
C.2.2	Unitary Groups ($U(n)$)	328
C.2.3	Special Unitary Groups ($SU(n)$)	328
D	Impact of 2-Lepton significance improvements on the total significance	329

List of Figures

1.1	Depiction of the Standard Model showing its seventeen particles (five bosons and twelve fermions), as spheres on plinths that represent the fundamental forces.	6
1.2	The particle data booklets that I currently own.	7
2.1	Picture showing the rough scales at which certain theories are valid. . .	12
2.2	Images showing the three different kinds of 2-2 processes as depicted by basic Feynman diagrams	15
2.3	Feynman diagrams showing the three different kinds of corrections to a 2-2 annihilation processes that can occur.	16
2.4	Image showing the three Standard Model Feynman vertices allowed by the $\mathcal{L}_{Fermions}$ term in the Standard Model Lagrangian.	18
2.5	Image showing the self interaction vertices for the gauge bosons post symmetry- breaking.	21
2.6	Image showing a rough sketch of a cross-section the Higgs potential with a fixed λ and a running m^2 for four different values of m^2	22
2.7	Image showing the Feynman vertices introduced to the Standard Model by the addition of a Higgs potential.	23
2.8	Image showing the one Feynman interaction between the Higgs Boson and fermions: the ffH vertex.	25
2.9	Pie chart showing the decay ratios of the Higgs boson to other particles.	26
3.1	Image showing the CERN accelerator complex complete with the particles that traverse each of the sections.	31
3.2	View of the Swiss/French countryside above the LHC tunnel.	35
3.3	Labelled diagram of all the parts of the CERN accelerator complex that are important for the running of the LHC.	36
3.4	CERN's official projection of the timeline and future development of the LHC Machine.	44
4.1	Cumulative luminosity delivered by the LHC and recorded by ATLAS detector for the entire length of Run-1 (a) and Run-2 (b).	47
4.2	Schematic of the ATLAS detector.	48
4.3	Schematic of the ATLAS detector with particular emphasis on the components that make up the Muon Spectrometer.	50
4.4	Labelled cross-section schematic of the Muon Spectrometer in the barrel region.	52

4.5	Labelled quarter-section schematic of the Muon Spectrometer.	54
4.6	Schematic showing the Calorimeter System in the Barrel region of ATLAS.	55
4.7	Sketch of the EM Barrel LAr Calorimeter.	59
4.8	Schematic showing the Magnet System at ATLAS.	60
4.9	Plot showing the bending power as a function of $ \eta $	62
4.10	Schematic showing a close-up cross-section of the Barrel and Endcap Inner Tracker System at ATLAS	63
4.11	Schematic showing a blown-up cross-section of the Barrel Inner Tracker System at ATLAS, and the distance each layer of the sub-system is from the beamline.	64
4.12	Schematic showing a blown-up cross-section of the Endcap Inner Tracker System at ATLAS, and the distance each layer of the sub-system is from the beamline.	65
4.13	Cross-section of a typical semiconductor module.	67
4.14	Picture of a finished SCT barrel module.	69
4.15	Cross-section of the Pixel sub-detector without the IBL.	71
4.16	Cross-section of a 3D Pixel	72
4.17	Image showing the typical energy deposits of various particles traversing the ATLAS detector.	73
4.18	Average number of pile-up interactions per bunch crossing for each year in Run-2.	76
4.19	Schematic showing the Inner Tracker System at ATLAS.	78
4.20	The RAL DLLH3 set-up at the University of Glasgow.	83
5.1	Diagram showing how the Monte Carlo Integration method is used to evaluate the integral of $p(x)$ between the points 'a' and 'b'.	86
5.2	Image showing the different factorised stages in MC simulations.	89
6.1	Image showing the different objects that are formed while creating a tracks at the ATLAS experiment.	98
6.2	Pictorial definition of the parameters that make up a track at the ATLAS detector.	100
6.3	Plot of the average number of tracks as a function of pile-up for the Loose and TightPrimary tracks in Run 2 2016 and 2017 data.	104
6.4	Plot of the average number of tracks as a function of pile-up for the Loose and TightPrimary tracks in Run-2 (2016) data compared to the equivalent MC simulations events.	105
6.5	Plot of the average number of fake tracks as a function of pile-up for the Loose and TightPrimary tracks in Run 2 data.	106
6.6	Histograms for the residuals of the ratio between data and MC simulations for the average number of tracks against μ	109
6.7	Final .txt output file produced by the tracking code I created which is the basis of the recommendations provided to the group.	110
6.8	Plots of the average number of tracks as a function of pile-up (for the Loose and TightPrimary tracks in 2016+2017 data.	111
6.9	Number of tracks versus μ for the Loose track selection data (filled) and MC simulated events (blank) with additional cuts on d_0 applied.	113

6.10	Histograms for the residuals of the ratio between data and MC for the average number of tracks against μ after two rounds of outlier removal.	114
6.11	Final .txt output files produced by the tracking code I created which is the basis of the recommendations provided to the group.	115
6.12	Graphic showing how MC truth track and reconstructed track objects are made.	116
6.13	Number of reconstructed tracks that have truth links vs μ plotted alongside the total number of tracks for the Loose selection (a) and the Tight-Primary selection (b).	118
6.14	Loose and TightPrimary that have $P_{match} < 0.5$	120
6.15	Image showing the percentage of tracks of a given selection without truth links but have matching probabilities of greater than 0.5.	121
6.16	Plot showing the number of Loose reco tracks that have $P_{match} > 0.5$ compared to one showing the number of loose reco tracks that have $P_{match} < 0.5$ and have a truth link.	121
6.17	Plots showing the number of reconstructed tracks vs μ for different barcode groups.	124
6.18	Number of stable reconstructed tracks vs μ	125
7.1	A schematic showing an example decision tree separation of a training set into signal and background leaves.	131
7.2	A graphic giving a brief overview of the types of jet algorithm that exist.	141
7.3	Three sequential combination algorithms run on a sample event with many soft particles added.	142
7.4	The MV2c10 output for the three jet types: b -jets, c -jets, and light-flavour jets in simulated $t\bar{t}$ events.	149
8.1	Feynman diagram of the Higgs production process with the highest cross-section: gluon-gluon fusion.	155
8.2	Feynman diagram of the production of a Higgs particle via vector boson fusion. V here can either represent a W-boson or a Z-boson.	155
8.3	Feynman diagram of associated production of a Higgs particle with a vector boson.	155
8.4	Feynman diagram of production of a Higgs boson produced alongside a pair of top quarks.	155
8.5	Feynman diagram showing the full Higgsstrahlung process of Higgs production at the LHC, and its decay into the final states of interest to the analysis.	157
8.6	BDT_{VH} post-fitting output distributions for the 0-lepton (top), 1-lepton (middle) and 2-lepton (bottom) channels, in the 2-jet (left) and 3-jet (or ≥ 3 jets for the 2-lepton case) (right) categories in the $p_T^V < 150$ GeV region.	168
8.7	Event yield as a function of $\log_{10}(S/B)$ for data, background, and a 125 GeV Higgs boson signal.	169
8.8	Proposed Stage 1 (top) and Stage 1.1 (bottom) Simple Template Cross-Section Regions.	175
8.9	Proposed Analysis regions for the 2019 STXS VH , $H \rightarrow b\bar{b}$ analysis. . . .	176

8.10	Proposed 5 POI (top) and 3 POI (bottom) Simple Template Cross-Section Regions for the analysis.	177
8.11	Measurements of cross-section $\times$ branching fractions of $H \rightarrow b\bar{b}$ and $V \rightarrow ll$ for each of the 5 POI's.	179
8.12	Feynman diagram showing how the V +jets process enters the signal region.	183
8.13	Feynman diagram showing how the $t\bar{t}$ process enters the signal region.	185
8.14	Image showing the four tree-level single top Feynman diagrams.	186
8.15	Multi-jet background example Feynman diagram.	188
8.16	Feynman diagram of the two leading diboson contributions to the background.	188
8.17	Picture showing the rough workflow of the analysis.	191
8.18	Distribution of the signal yield under the radial difference between the two leading jets for (a) 2-jet and (b) 3-jet 1-lepton channel.	196
8.19	Feynman diagram showing the full Higgsstrahlung process of Higgs production, and its subsequent decay into the final states of interest to the Boosted analysis.	207
9.1	Example formulation of the test statistic for the null (q_{H_0}) and alternative hypothesis (q_{H_1}).	219
9.2	Upon measurement of q_{obs} , the alternative hypothesis is rejected as the q_{obs} is larger than the p-value set.	220
9.3	Variation of the likelihood for different values of the number of expected events ($\mathcal{N}^{exp}$).	224
9.4	Plot of the likelihood value at different values for the number of sigma the systematic θ_j is shifted.	225
9.5	Plot of the likelihood value at different values for the number of sigma the systematic θ_j is shifted.	226
9.6	Maximum likelihood variations for fixed $\mathcal{N}_{data}$ and (top) fixed pre-fit impact of θ_j but varying $\mathcal{N}_{MC}$; (bottom) fixed $\mathcal{N}_{MC}$ but varying the pre-fit impact of θ_j	228
9.7	Plot showing the variations of the NLL when the systematic θ_j is varied around nominal pull value.	230
9.8	Example 'pull plot' for θ_j	231
9.9	1σ shifts from the minimum NLL for the $\theta + j$ (a) and the subsequent change in the number of events in the bin (b).	234
9.10	Pre-fit (9.10a) and post-fit (9.10b) plots for the example fit.	236
9.11	3D plot showing the change of the likelihood with variation of two systematics: θ_i and θ_j	237
9.12	Pull plot showing the pull and variance for θ_i and θ_j	238
9.13	1σ shifts from the minimum NLL for the two systematics θ_i (top left) and θ_j (top right).	239
9.14	Ranking plot for two sets of systematics.	240
9.15	Correlation matrix for θ_i and θ_j	242
9.16	Comparison of the post-fit plots between the single and double systematic fits.	243

9.17	Illustration of the possible effect a systematic can have when it is present in multiple bins.	245
9.18	Image showing the normalisation effect of a systematic SysJET_CR_JET_JER_EffectiveNP_1.	246
9.19	Example post-fit plots from the 0L standalone analysis.	253
9.20	Pull plot for the systematics related to b-tagging in the Combined 012L fit.	255
9.21	Fit study flow chart.	256
9.22	B-tagging pulls comparing the fits of the nominal fit with the fit where the Light_0 pull is fully decorrelated.	258
9.23	Light_0 flavour pulls in three different fits.	258
9.24	Selection of flavour tagging pulls from fits where the control and signal regions are separated, and where they are merged.	259
9.25	Light_0 (and Light_2) flavour pull(s) for two different fits.	260
9.26	Image showing the shape effect of the systematic SysFT_EFF_Eigen_Light_0 in three samples.	261
9.27	Image showing the normalisation effect of the systematic SysFT_EFF_Eigen_Light_0.	262
9.28	Light_0 (and Light_2) flavour pull(s) for two different fits.	263
9.29	Comparison between the NormVal plots of the single bin fit and the V+jets decorrelated single bin fit.	263
9.30	Relevant pulls that have changed in the fit where the normalisation factors have doubled with respect to the nominal fit.	264
9.31	Pull plot where the nominal fit is compared to the fit decorrelated in both n-jet and V+light categories and the normalisation systematics are doubled.	265
9.32	Pulls from the Nominal 0L fit compared to a fit where an additional systematic is added.	266
9.33	Normalisation factors from all the samples separately for each region in the 0L fit.	268
9.34	Selection of flavour tagging pulls from the Nominal 0L fit compared to a fit where no systematic are pruned.	269
9.35	Selection of non-flavour tagging pulls from the Nominal 0L fit compared to a fit where no systematic are pruned.	270
9.36	Comparison of shape and normalisation effects for the systematic eigenvalues Light_1, Light_2, and Light_5.	272
9.37	NormVal plots for the two systematics VVPTVME and VVPTVPSUE with and without pruning for the 12 regions in the 0L standalone fit.	273
9.38	Figure showing the shapes for all the uncertainties for VVPTVME that appear in the fit when the pruning is removed.	274
9.39	Figure showing the shapes for all the uncertainties for VVPTVPSUE that appear in the fit when the pruning is removed.	275
9.40	Plots of $\frac{1_{up} - nominal}{nominal}$ for the WZ (left) and the ZZ diboson sample (right).	276
9.41	Plots of the number of events that correspond to the 1σ upward fluctuation of the systematic ZZPTVPSUE in the previous and current analysis.	277

10.1	Graphic showing possible event yield with the introduction of certain trigger information into the event selection.	283
10.2	Graphic showing differences between the new and the older trigger regime.	286
10.3	Preliminary plots showing the effect of the E_T^{miss} study in the Resolved 2-lepton analysis using the logical OR of the available triggers.	288
10.4	Preliminary plots showing the effect of the E_T^{miss} trigger study in the Resolved 2-lepton analysis using only E_T^{miss} triggers to trigger muons. . . .	289
10.5	Plot showing the difference between the event selection yields when using MET_{dimuon} ('MET') as the trigger proxy, and MET. Here the event selection uses both muon and E_T^{miss} information to filter events.	290
10.6	Preliminary results for the impact on the decided E_T^{miss} trigger regime on the Resolved analysis. MET + 'MET' refers to the final E_T^{miss} proxy used by this study which is a vector addition of MET_{dimuon} and MET.	290
10.7	Simplified graphic showing differences between the new and the older trigger regime.	291
10.8	Final results for the impact on the decided E_T^{miss} trigger regime on signal and background samples present in the Resolved analysis.	292
10.9	Maximal percentage increases in the event signal yield from changing trigger strategy in the Boosted analysis, testing the difference between the logical OR and the E_T^{miss} -trigger-only triggering regimes.	295
10.10	Preliminary results for the impact on the decided E_T^{miss} trigger regime on signal and background samples present in the Boosted analysis on MC16a signal events.	296
10.11	Trigger scale factor values versus p_T^V for the five E_T^{miss} triggers used in the Boosted analysis.	298
10.12	Differences between the impacts of the two E_T^{miss} trigger regimes on signal, background and data samples present in the Boosted analysis. . . .	300
10.13	p_T^Z distributions in the $250 \text{ GeV} \leq p_T^V < 400 \text{ GeV}$ and $p_T^V \geq 400 \text{ GeV}$ regions in the 2-lepton channel for signal, all backgrounds, and the data separately.	303
10.14	Plot showing the pre-fit m_J distributions and the data-MC ratio for both p_T^V regions in the Boosted analysis 2-lepton signal region, before and after the new trigger regime is implemented.	305
11.1	Mass distributions of the large- R jet in the 0-lepton, 1-lepton, and the 2-lepton 2- b -tagged jet signal regions.	308
11.2	m_J distributions in all of the lepton channels and p_T^V regions combined for data and MC simulations after the removal of all non-diboson backgrounds in the Boosted analysis.	309
11.3	Measured and predicted values of the STXS's for WH and ZH in the Boosted analysis's STXS framework in each p_T^V region.	310
11.4	Output of the multivariate BDT_{VH} distributions in the (a,b) 0-lepton, 1-lepton, and the 2-lepton 2-jet, 2- b -tagged jet signal region of the Resolved analysis.	312

11.5	Final m_{bb} distributions in all of the lepton channels and p_T^V regions combined for data and MC simulations after the removal of all non-diboson backgrounds in the Resolved analysis.	313
11.6	Values of μ_{ZH}^{bb} and μ_{WH}^{bb} for the Resolved analysis 2-POI fit, and their combined 1-POI value and the Resolved analysis diboson cross-check fit.	314
11.7	Measured and predicted values of the STXS's for the two WH and three ZH sub-channels in the Resolved analysis's STXS framework.	315

List of Tables

1.1	Table showing the strength of all the fundamental forces relative to the strongest force; the strong force.	4
4.1	Table showing the sub-detectors various particles deposit energy into. . .	72
6.1	Table showing the difference in selection requirement made on Loose and TightPrimary tracks	103
8.1	STXS categories for the 2019 VH , $H \rightarrow b\bar{b}$ analysis paper showing which analysis regions contribute the most to which STXS category.	178
8.2	Generators used for the simulation of the signal and background processes.	180
8.3	Table showing the summary of the event selection and categorisation in the Resolved analysis 0-, 1- and 2-lepton channels.	193
8.4	Table showing the variables used for the multivariate BDT discriminant in each of the lepton channels.	198
8.5	Summary of the systematic uncertainties on the main backgrounds in the analysis.	202
8.6	Tabulated summary of the systematic uncertainties on of the diboson background.	203
8.7	Summary of the systematic uncertainties on the signal.	204
8.8	Table showing the summary of the event selection and categorisation in the Boosted analysis 0-, 1- and 2-lepton channels.	208
8.9	Summary of the Boosted analysis regions.	211
8.10	Breakdown of the systematic contributions to μ_{VH}^{bb}	212
9.1	Table showing the significances of the baseline standalone 0-, 1- and 2-lepton fits and the significance when the pruning is fully removed. . . .	271
10.1	Table showing the E_T^{miss} triggers used during the Run-II (2015-2018) data collection period.	284
10.2	Table showing the event selection criteria for the 2-lepton Resolved analysis at the time a decision was made about the study.	293
10.3	Yields for signal and background inclusive in p_T^V in the signal regions and fit significances.	304
D.1	Evaluated significances of the channels in the Boosted analysis correct as of 20/11/19.	330

Declaration

The material presented in this thesis is the result of my own research work in the Particle Physics Experimental group of the School of Physics and Astronomy at the University of Glasgow. It has not been submitted for any other degree at the University of Glasgow or any other institution. The list below specifies which parts of the thesis contain my own original work.

Chapter 4 In the final part of the chapter (section 4.5), I talk about the tracker upgrade and my contributions to the chip testing software.

Chapter 6 This chapter is dedicated to the studies I performed for the Tracking Combined Performance Group, parametrising the differences between the Monte Carlo modelling of tracks and those reconstructed from data, and deriving recommendations for the collaboration.

Chapter 9 The second part of the chapter (from section 9.4 onwards) shows the work I did in the Resolved analysis to understand and improve the fit model in the 0-lepton Standalone fit.


Chapter 10 This chapter summarises the work done to increase the significance of the Boosted analysis 2-lepton channel.

Chapter 11 This chapter summarises the overall results for the analyses discussed in this thesis.

This work contributed to the following journal publications:

- [1] Observation of $H \rightarrow b\bar{b}$ decays and VH production with the ATLAS detector.
- [2] Measurement of $H \rightarrow b\bar{b}$ as a function of the vector-boson transverse momentum in 13 TeV pp collisions with the ATLAS detector.
- [3] Measurements of WH and ZH production in the $H \rightarrow b\bar{b}$ decay channel in pp collisions at 13 TeV with the ATLAS detector.
- [4] Measurement of the associated production of a Higgs boson decaying into b -quarks with a vector boson at high transverse momentum in pp collisions at $\sqrt{s} = 13$ TeV with the ATLAS detector.


Part I



High Energy Physics Background

“Trying to construct a broad understanding of the world from a memory of your own experience would be like looking at the ceiling of the Sistine chapel through a long, thin cardboard tube: you can try to remember the individual portions you have spotted here and there, but without a system or a model, you are never going to appreciate the whole picture. Let’s begin.”

- Ben Goldacre, *Bad Science*



Banner image for part headings from <https://atlasvirtualvisit.web.cern.ch/>

*

Introduction

1.1 Opening Remarks



Autumn is the season. It is a dry September afternoon and you find yourself on a train. Temporarily setting aside your book, your phone, or your laptop, your attention turns to the picturesque countryside that streams past.

What do you see?

Do you see the solar photons undergoing Rayleigh scattering in the upper atmosphere to make the sky appear blue¹? Or the refraction of that light in water droplets to produce a rainbow? Do you look at the trunks of the trees racing by and wonder exactly how you see a brown colour as there was no brown in the rainbow you have just seen²? Do you see the birds flying past as being in an eternal battle with the forces of gravity, air resistance and upthrust keeping them aloft, all the while oblivious to the billions of neutrinos passing through them? Chances are while, like myself, you are a massive

¹The sun gives off lower intensities of violet light than blue light; this coupled with the fact that our eyes are more sensitive to blue light than violet light meaning that we see the sky as blue.

²Colour is made from three components: hue, brightness, and saturation. Brown is just an orange that does not reflect much light back to your eyes.

nerd because you are reading this, you do not see the world in this way (unless, of course, you are under the influence of some sort of narcotics). But **Physics**TM is the science of trying to describe how the natural world works and the properties of matter and energy. Particle physics attempts to do this by deconstructing the universe (and by extension, the world) into a set of fundamental building blocks with rules that govern their interactions. This is referred to as a model. Currently the best model that is consistent with two centuries of experimental data, is a set of quantum field theories known as the Standard Model.

1.2 The Forces of Nature

A *force* is the name given to any interaction that changes the motion of an object if not equally opposed by another force. Physics asserts that every force in nature that acts on a body is reducible to one of four fundamental forces. Two forces govern long-range interactions between objects: electromagnetism and gravity; and two forces are only relevant on atomic-scales: the weak and strong forces.

The strong force is the strongest of these forces and is responsible for holding atomic nuclei together despite the fact the constituent protons all have the same sign electric charge and therefore repel each other. The electromagnetic force is the next strongest force and governs the interactions between particles that are electrically charged. The weak force is responsible for radioactive decay and is the main process of nuclear fusion in stars. The last force is gravity, the only force relevant across the large distances of the cosmos, but by far the weakest at the atomic scale (see Table 1.1 below). Gravity is an attractive force that acts between all objects that have mass.

These four forces are conservative, which means the relative strength of these forces is only dependent on the distance between two objects that can undergo the same inter-

Table 1.1: A table showing the strength of all the fundamental forces relative to the strongest force: the strong force [5]. The force-distances mentioned here are governed by the lifetime of the particles that mediate the forces. Since these relative strengths are energy dependent, these values are approximations made at 200 MeV.

Force	Relative strength	Range (10^{-15} m)
Strong	1	~ 1
Electromagnetic	$\sim 1/137$	∞
Weak force	$\sim 10^{-6}$	$\sim 10^{-3}$ [6]
Gravitational	$\sim 10^{-39}$	∞

action. A description that permeates all of space able to describe the way these forces propagate and interact with objects is called a ‘field’, and it allows us to come up with a mathematical formalism to describe these forces. For sub-atomic particles that experience the fundamental forces, in order to accurately depict their behaviour, a single model interweaving these forces and the fundamental particles that interact via them is required. This mathematical description is the Standard Model.

1.3 The Standard Model

The Standard Model (SM) describes the observable universe on a sub-atomic level by the interactions of fundamental particles.

There are twelve *fermionic* or ‘matter’ particles: six quarks that interact via the strong, electromagnetic and weak forces; three electrically charged leptons that interact with the electromagnetic and weak forces; and three electrically neutral leptons called neutrinos that interact only with the weak force. These twelve particles also have distinct anti-particles. Anti-matter particles, however, are seen as trivial extensions of matter particles and are thus not depicted as individually separate from their matter counterparts in the Standard Model³. These matter particles can combine in many different

³An anti-particle has the same mass as its particle counterpart, but all its quantum numbers are reversed. For example the up quark has a unit of colour charge [red, blue, green] and a electric charge relative to the charge of an electron of $\frac{2}{3}$, and its antiparticle counterpart (the anti-up quark) has a unit of anti-colour charge [anti-red, anti-blue, anti-green] and an relative electric charge of $-\frac{2}{3}$.

ways to create other particles.

There are four force-carrying *gauge bosons* that interact with particles to mediate the fundamental forces: the gluon, which mediates the strong force; the photon, which mediates the electromagnetic force; and two weak bosons, the W-boson and the Z-boson which mediate the weak force. Since the W-boson has electric charge, two W-bosons exist: the W^+ and the W^- , but they are anti-particles of each other. The gluon (and quarks also) carry colour change, so technically all possible colours that the particles can have count as separate particles, but again this ‘colour index’ is a trivial extension of the number of particles and are thus discounted.

The last particle in the Standard Model description is the Higgs boson. The Higgs boson is a physical manifestation of the Higgs field. Particles that interact with this Higgs field gain their ‘rest’ mass. A pictorial depiction of the particles in the Standard Model is shown in Figure 1.1.

This means that there are seventeen fundamental particles, and of these, the Higgs particle, discovered only in July of 2012 by the ATLAS (A Toroidal LHC ApparatuS) [7] and CMS (Compact Muon Solenoid) [8] experiments, is the least well known in terms of its characteristics. This PhD thesis will describe my contributions to increasing our understanding of the Higgs boson, and specifically the discovery of the decay mode of a Higgs particle into a bottom and anti-bottom quark ($H \rightarrow b\bar{b}$).

The astute readers among you (who didn’t already know) will have spotted one of the failings of the Standard Model in this brief description. It has no mechanism through which particles interact gravitationally with each other. While a gravitational attraction between particles is expected to be extremely weak, such a mechanism is possible. However, there are both experimental and theoretical challenges in uniting gravity




Figure 1.1: Depiction of the Standard Model. There are seventeen particles, antiparticles notwithstanding, shown as spheres consisting of five bosons and twelve fermions. The particles known as fermions are all on the central plinth. Quarks are depicted in blue, charged leptons in yellow/gold and the neutrinos in red. The quarks are separated into two rows; the top row has a fundamental electric charge of $\frac{2}{3}$ while the bottom has a fundamental electric charge of $-\frac{1}{3}$. The force-carrying bosons have interactions with all particles of the colour of the bars that come off the central plinth. All particles on plinths have interactions with all other particles on all levels higher than it, i.e. the gluon will interact only with the quarks but the photon will interact with both the quarks and the charged leptons but not the gluon. The exception to this is the W-boson as it has electromagnetic charge, so can interact with a photon. The fermion plinth is divided into three groups of particles called generations. Between each generation all the quantum numbers barring flavour are the same but the higher the generation number the more massive the particles called generations. All particles in white text are massive, and the sizes of the spheres are indicators of the relative masses of the particles with the exception of the electron and neutrinos which would be too small so have been scaled up. All particles in black text are massless. The Higgs particle interacts with all massive particles. If the white-gradient shading on the ball is to the bottom right then these particles interact with the Higgs via the Higgs Mechanism, and if the black-gradient shading on the particle ball is to the bottom left then the interactions with the Higgs are via the Yukawa interaction.

with the Standard Model. From a theoretical point of view in constructing the Standard Model, it has not proved possible thus far to weave Einstein's General Relativity into the Standard Model. In addition, experimentally, the particle that would be responsible for mediating the interaction, the graviton (g^0), has not been detected.

The field of particle physics is ever-evolving and minor alterations are made to the Standard Model 'particle zoo' all the time such as the addition of newly discovered particles and the recording of properties to greater precisions [9–19]. These are regularly done and one such way of communicating these changes are the Particle Data Group Booklets, some of which are present in Figure 1.2. The first step on this journey is a step into the theoretical side to ask why the Standard Model needs a Higgs Boson.



Figure 1.2: The particle data booklets that I currently own starting from the year I was born. If you have physical copies of the 1996 or 2004 versions lying around, I would be interested in taking them off your hands.

CHAPTER 2

Standard Model theory

2.1 Theory

Tackling the interesting and large-scale questions about matter and forces requires the ability to make sense of any observations that are made. Data taken in experiments has to be linked to the underlying fundamental properties of the system. In order to do this the interpretation of the results needs to be grounded into a theoretical framework which gives us the means to interpret the results with some level of predictive power¹.

To describe particle physics, the Standard Model must be able to describe highly-energetic

¹A footnote by Kenny Wraight: To build such a cosmological narrative, we try to wrestle with as much of the maelstrom of phenomena as we can manage. Peering through the prism of theory we distinguish, filter, edit and understand the observed world. At the same time, we use a pragmatic sense to adapt the framework, honing the lens and sharpening its focus. Predictive power and empirical accuracy are soft metrics to follow and the process is iterative and unending. The prize of this endeavour is an interpretative structure which, with as few blocks as possible can support the weights of current observations and future measurements. An addition by me: We also must acknowledge our potential personal bias in all of this. Often there are multiple ways to describe the same phenomena but our cultural and societal biases (not only of the prevailing ideas in a given field, but more generally as well) guides one way of thinking over another. At some level the basis we use to frame the universe also tells us a lot about ourselves. The theory becomes then a picture of the heavens less so described and more wished for.

particles interacting with each other on small distance scales. The Standard Model as described in Chapter 1 is a surface-level explanation of the aims of the theory, the next section delves into what this means from a theoretical standpoint, and shows how the Standard Model formulation as it is known today came about.

2.2 The theoretical foundations of particle physics

Until the 18th century, Classical Mechanics was the framework for describing the world as it was known: large distance scales and low energies. For this we had Newton's equations but they were limited in scope. To be able to solve a problem using Newtonian Mechanics one has to find all the forces acting on a system and the direction they are acting in, and this can be non-trivial. Newtonian Mechanics also unfortunately happens to be constructed within and work best in Cartesian co-ordinates. The Lagrange formulation of Classical Mechanics, or Classical Lagrangian mechanics (CLM), is a reformulation that adds nothing intrinsically new in terms of physics. It allows the mechanics of a particle to be represented by a single object called the *Lagrangian* (see appendix A) independently of the co-ordinate system and to solve equations of motion, one only needs to identify the degrees of freedom of a system.

CLM, however, applies only to the dynamics of particles that have finite degrees of freedom, and like classical mechanics breaks down for relativistic (high energy) particles. Particles have small masses, so those that have large energies will have large velocities, which is crucial for the regime of particle physics to be tested. The Lagrange formalism however is applicable to any theory and therefore retains its mathematical usefulness.

Another problem with CLM is that there are certain important phenomena (for particle physics anyway) that it fails to describe. One of the more famous early examples is the Young's Slit Experiment [20]. In this experiment photons (and later electrons) are

demonstrated to not only exhibit both wave and particle-like properties individually, but able to interfere with themselves to behave as a particle and a wave at the same time!

A potential solution to the latter problem is to describe particles abstractly as discrete energy-state packets called quanta, akin to particles. These *quantum states* are also described by *wavefunctions* which obey wave mechanics. These single particle quantum states are able under non-observation to exist in all possible paths, and to collapse in a probabilistic manner to one option when observed. This is Quantum (Lagrangian) Mechanics, and to obtain this from CLM one has to impose five axioms (see appendix B). *Quantisation* introduces a constant $\hbar$, which both governs the quanta of energy a system can have and the amount of information you can know about the system at the same time (also appendix B).

Ultimately QLM fails when we try to solve the former problem as well. Turning QLM to Relativistic QLM solves the issue of relativity but in doing so the probabilistic wavefunction has solutions that are no longer positive definite. Essentially this either means that the energy states available to single particles can be negative, or that the probability density of the single particles can be negative. The only way to get around this issue is to create particles in particle-hole pairs² [21, p60-61] but then this breaks the single-particle picture that underpins quantum mechanics. A field theory description can solve this problem, as it has the degrees of freedom to define the negative energy states as particles themselves, so lets take a step back and look at what a classical field looks like.

In both Newtonian mechanics and its Lagrangian equivalent counterpart, the combination of two states meant that their momenta were simply added together to resolve the

²Here a ‘hole’ is a solution capable of removing a particle with is a positive quanta of energy and hence acts like a particle with a negative quanta of energy

overall force of this new state. In the relativistic regime, this can no longer be done due to there being a ‘speed limit’ on particles: the speed of light (c). A Classical Field Theory (CFT), such as Electrodynamics, moves from being a description of isolated particles to one of fields, which have infinite degrees of freedom. CFT’s solve the issue of relativity by being compatible with Einstein’s Special Relativity which introduces the idea of Lorentz transformations. Particles travelling with speeds comparable to the speed of light (‘Lorentz boosted’) that interact will have a non-linear velocity-dependent combination of their velocities such as to respect this light-speed limit [22, p48]. A CFT, however, can only describe the mechanics of fields and hence ignores the parts of the theory where discrete particles are involved. In the case of Electromagnetism, it can deal with calculations of electric fields through matter but fails to describe phenomena where photons (particles of light) are present (e.g the photoelectric effect).

This problem can be solved by ‘quantising’ CFT in a similar manner as was discussed previously with CLM. Quantum Field Theory or QFT is able to accurately describe relativistic particles on small distance scales, the fields they interact with, and the probabilistic nature of these interactions. The focus then shifts on not only trying to describe directly the particles themselves, but the fields that these particles operate in as well. In this picture we describe particles in these fields via their interactions with the fields and the particles that mediate the forces as excitations of these fields.

Figure 2.1 shows a summary of the discussion above showing rough regions in distance-energy space where certain theoretical approaches are valid. QFT is now the stage on which we set the scene. Thus we can start to craft a model of particle physics by finding or imposing *symmetries* upon our particle system and using Group Theory (see Appendix C) to analyse the consequences of these symmetries.




Figure 2.1: Picture showing the rough scales at which certain theories are valid. At small distances the description moves from particle having any energy to energy coming in small packets called quanta. This 'quantisation' comes with a constant. At large energies, particles travel closer to the cosmic speed limit - the speed of light and theories need to be compatible with Einstein's Special Relativity.

2.2.1 Constructing a theory

Now we have established a theoretical framework capable of describing some objects, the next step is to place it within the constraints of our universe. This sounds overly-grandiose and obvious, but these theories are so general they can be applied to a variety of fictional universes. For the theory to be able to make statements about our universe, we can impose some general observations or symmetries about the universe that are experimentally validated. From the remaining structure of the theory, various operations can be formed to make predictions.

The structure of the universe has some symmetries that our QFT theory should observe. One such symmetry is that largely we do not see a variation in the laws of physics at

differing points in space³. This is the same thing as saying that the theory is invariant under *translations*. Another such symmetry is that the space seems to be isotropic, so with no preferential direction, and is the same between two observers irrespective of their relative speeds to each other (the atoms that make up your body don't fall apart when you get on a train for example). This is Lorentz invariance. The group that combines Lorentz transforms and translations is the Poincaré group, $\mathcal{P}$. Transformations of states in QFT therefore take the form of

$$|\psi_i\rangle \rightarrow \mathcal{P}_{ij} |\psi_j\rangle. \quad (2.1)$$

In the quantum mechanical picture, the interactions between two objects or *dynamics* of the system are described by the inner product of two states (again see appendix B), often referred to as the *matrix element*, $\mathcal{M}$ and is given by

$$\mathcal{M} = \langle \psi_1 | \psi_2 \rangle. \quad (2.2)$$

The inner product therefore transforms under the Poincaré group as

$$\langle \psi_1 | \psi_2 \rangle \rightarrow \langle \psi_1 | \mathcal{P}^\dagger \mathcal{P} | \psi_2 \rangle. \quad (2.3)$$

For this matrix element to be normalised and Poincaré invariant, $\mathcal{P}^\dagger \mathcal{P} = 1$. This is the definition of unitarity. A *particle* is therefore defined as an object whose free state and interaction dynamics transform under a 'global' unitary representation of the Poincaré group [23, p110].

Particle states that are solutions to this (representations of the Poincaré group) have many *internal* properties⁴ (also known as a *quantum numbers*) to uniquely define them. Most of these (charge, flavour, etc) are unaffected by Poincaré Group transforms but

³This is not the case when the space itself is warped immensely by a large ultra-dense object like a black hole, or in the centre of a neutron star, but these are exceptional cases.

⁴Parameters that describe the particle.

mass and ‘spin’ (or more accurately the component of spin projected on a given axis) are subject to change and therefore define the types of representation. Generally there are two types of particle that transform in the required way. Those with integer spin-states are called *bosons* and those with non-integer spins are called *fermions*.

The fact that there exist properties of these particles that are unaffected by Poincaré Group transforms means that in the theory fermions and bosons have additional ‘internal’ degrees of freedom. These internal degrees of freedom are interesting, because while they do not affect the underlying theory, there can be many choices that can be made to fix the properties of particles on the sub-theory level (or better known as the ‘local’ as opposed to ‘global’ level⁵) in order to tackle it. These choices are called *gauges*. Since the overall theory has already been validated, the choice of gauge should not matter. The freedom to choose any gauge to solve problems is referred to as gauge invariance.

Like their global cousins, local invariances imply local symmetries. These symmetries give the overall theory a sub-structure, as each set of internal transforms will form their own group. These are called gauge symmetries. We map these internal symmetries to the fundamental forces.

2.2.2 Feynman diagrams

In order to do interesting calculations with the theory, we need a way of describing particle interactions. This is done in the form of the matrix element. This matrix describes two incoming particles interacting in some way and producing two outgoing particles, known as a 2-2 scattering process. Richard Feynman pioneered a visual representation of this 2-2 scattering process.

⁵The main difference here being that global symmetries are observed independent of the position of the particle and local ones are observed at each point in spacetime.




Figure 2.2: Images showing the three different kinds of 2-2 processes as depicted by basic Feynman diagrams. Fermions are depicted as straight continuous lines, and all other types of lines (wavy, curly or dashed) depict a boson of some description. If the y-axis represents spatial separation and the x-axis represents time, vertices occur when particles interact via collisions ($2 \rightarrow 1$ lines) or decays ($1 \rightarrow 2$ lines). The left annihilation diagram shows two fermion particles merging into a boson which then decays into two fermions. The central scattering diagram shows two particles approaching one another, exchanging a boson and then drifting apart. The right absorption diagram depicts a fermion absorbing a boson and then emitting another boson later. The combinations of fermions and bosons (fff , ffb , fBB , BBB) that can meet at vertices are limited by the theory the diagram is trying to describe.

Each line (squiggly or otherwise) in Figures 2.2a, 2.2b and 2.2c represents a particle, and particles that have both ends connected to a vertex are called *propagators*. Propagators are interesting because they are ‘virtual’ particles, since they never interact with the outside world, they can have funny properties like negative energy. All the other particles are initial or final state particles and are real and have real and physical energies and momenta.⁶

Each place where three or four particles meet is called a vertex and either represents the emission of a particle or the creation of a particle from the collision of two particles.

Due to this, the vertices allowed by a theory depend on which particles can or cannot

⁶Imagine the particles are three vectors that all meet at a single point (vertex). Let’s say that vector towards the point is positive and away from the point is negative (it will be different for each vertex). For all of the vectors to be positive, they must all be oriented towards the point. In this case, the points and particles act like field lines of a magnet going to or from a pole. If the diagrams have no flow then this is a valid description, however this is not the case. We want to describe particles interacting and changing which indicates a flow; the flow of time. For there to be a sense of ‘flow’ the point cannot act as a pole. To do that some of the vectors will point towards the vertex and others will point away. If these vectors correspond to particles then those vectors could have properties like the energy and momentum of the particle, and this direction will therefore indicate whether that particle is real or virtual. To keep the theory able to describe experimental findings, all the final state particles have to be real, which can be achieved by adding in additional vertices and confining the virtual particles between them.

interact with one another. Each vertex and leg are mathematical objects and the matrix element is the product of all of the components present in the diagram. The probability of such a process occurring is related to the square of this matrix element.

The basic processes shown in Figure 2.2 have low numbers of vertices. The diagrams with only two vertices are referred to as *tree-level* diagrams. These simple diagrams can be modified with more particles emanating from legs or from the propagator to create diagrams with more vertices. Some examples of these *higher order* corrections to the 2-2 annihilation process in Figure 2.2a can be seen in Figure 2.3.




Figure 2.3: Feynman diagrams showing the three different kinds of corrections to a 2-2 annihilation processes that can occur. All of these processes have four vertices as opposed to 3. All the correction diagrams that have three vertices would change the number of input or output particles and describe a different process.

The diagrams in Figures 2.2 and 2.3 only have one type of vertex (ffB) in them.

For the probabilistic nature of the matrix element to be realised, each vertex must contribute a factor between 0 and 1. Hence the more vertices a diagram has, the rarer the process. Diagrams with more vertices are often referred to as being of higher order. Higher order diagrams with the same initial and final states as lower ones therefore confer corrections on the lower order process.

The *cross-sections* of processes (how likely they are to occur) is given by the integral of

the square of the sum of the matrix elements of all the diagrams that enter the process with respect to energy of the particles in the process.⁷

2.3 The Standard Model Lagrangian

The presence of gauge symmetries in our model imply that the Lagrangian of the theory can be represented as many additive parts, with internal symmetries manifesting themselves in the gauge transformations of multiplicative sections. The additive structure of the Standard Model Lagrangian is shown in Equation 2.4.

$$\mathcal{L}_{SM} = \mathcal{L}_{Gauge} + \mathcal{L}_{Fermions} + \mathcal{L}_{Yukawa} + \mathcal{L}_{Higgs} \quad (2.4)$$

This Lagrangian has terms in each section which relate to the interactions between various fermionic and bosonic fields. The full terms can be expressed as Equation 2.5.

$$\mathcal{L}_{SM} = -\frac{1}{4}F_{\mu\nu}F^{\mu\nu} + i\bar{\psi}\not{D}\psi + \bar{\psi}_L\hat{Y}\psi_R\Phi + \bar{\psi}_R\hat{Y}\psi_L\bar{\Phi} + D_\mu\Phi^\dagger D^\mu\Phi - \mathcal{V}(\Phi) \quad (2.5)$$

where $F_{\mu\nu}$ and D represent bosonic fields and their couplings, respectively; ψ represents the fermion fields and Φ and Y represent the Higgs field and its coupling.

In this description, fermions are the building blocks of the universe and interact with each other via the exchange of bosons. The rest masses of all of these particles come from their interaction with the Higgs field. The rest of this chapter will briefly go over each part of this Lagrangian.

2.3.1 $\mathcal{L}_{Fermions}$: Nature's building blocks

Given that the structure of a Lagrangian generally is in the form of $\mathcal{T} - \mathcal{V}$ (see Appendix A), the first three terms of Equation 2.4 are the 'kinetic terms' of the Standard Model

⁷also known as the 'phase space'.

but are more often referred to as the system ‘dynamics’. Starting with the second term of the Lagrangian we have the fermion interaction term. Sandwiched between two spinor fermion fields is the covariant derivative $\not{D}$, which encodes the interactions of the fermions with the boson mediators of the forces. The Feynman vertices this term corresponds to are depicted in Figure 2.4.




Figure 2.4: Image showing the three Standard Model Feynman vertices allowed by the $\mathcal{L}_{Fermions}$ term in the Standard Model Lagrangian. The f represents quarks, and both charged and neutral leptons. For fermion species interacting with a W-boson the two incoming and outgoing fermions must be different but for fermions interacting with a Z-boson or a photon, the species must be the same. For the strong qqg vertex, while the flavour of q_1 and q_2 are the same they will have different colour charges as the gluon also carries colour charge.

2.3.2 $\mathcal{L}_{Gauge}$: Nature’s forces

The gauge invariances mentioned earlier are expressed entirely in the first part of the Lagrangian. The symmetries that make up these gauge groups correspond to the fundamental forces in the Standard Model.

The elements in a particular gauge group consist of fermionic particles (or objects made up of several similar particles) that can feel the corresponding forces, and the generators of the group (that can create all of the elements) are taken to be the bosons that

mediate the forces. These bosons are therefore referred to as *gauge bosons* and in the unbroken form of the Standard Model Lagrangian, both the fermion group members and the gauge bosons are massless.

The unbroken gauge group structure of $\mathcal{L}_{Gauge}$ and therefore of the entire Standard Model Lagrangian is a union of the group that represents the strong force, also known as Quantum Chromodynamics or *QCD*, and the group that represents the combination of the electromagnetic and weak forces, known as the electroweak interaction. In group terms this is expressed as $SU(3)_C$ and $SU(2)_L \times U(1)_Y$, respectively. The subscripts on the groups refer to the property that unites all the elements of that group. The subscript C , refers to the ‘colour’ charge of group elements in QCD mediated by gluons with 8 different types of colour charge. The subscripts L and Y refer to the left-handed weak isospin and the weak hypercharge of the members respectively. In the unbroken SM the electromagnetic forces and the weak force are not separate forces. Essentially this means that the boson mediators (the $W_\mu^{1,2,3}$ and the B_μ), which have different isospins and hypercharges, transform under both $SU(2)$ and $U(1)$ symmetries.

The gauge group structure of the Standard Model has to be broken, because masses for the weak bosons and the fermions that interact in groups have been experimentally observed. For the theory to be consistent with measurements, it needs another mechanism that can give mass to both the gauge bosons and fermions.

To give mass to the gauge bosons, the symmetry of the $SU(2)_L \times U(1)_Y$ group needs to be ‘broken’. Each generator in this group that breaks gives rise to a mediating gauge boson known as a Goldstone boson [24, p198-p201]. It is these Goldstone bosons that have to interact with this new mechanism to give rise to the massive gauge bosons the W-boson and the Z-boson. This mechanism that simultaneously breaks the electroweak symmetry and gives mass to the Goldstone Bosons is represented in the $\mathcal{L}_{Higgs}$ part of

the Standard Model Lagrangian.

The masses for the fermions come from an addition to the interaction term between fermions and bosons, $\mathcal{L}_{Fermions}$ given by $\mathcal{L}_{Yukawa}$.

The $\mathcal{L}_{Gauge}$ itself is responsible for the self-interaction of the gauge bosons with themselves. This only happens for forces that correspond to the *non-abelian* groups $SU(2)_L$ and $SU(3)_C$. The four three-point vertices and two four point vertices that are allowed in the SM post symmetry breaking are shown in Figure 2.5.

2.3.3 $\mathcal{L}_{Higgs}$: Solving a massive problem

As mentioned in the previous section, the presence of observed masses for the weak gauge bosons in an otherwise massless theory means that the gauge field $SU(2) \times U(1)_Y$ is ‘broken’ down into $U(1)_{EM}$. The strength of the pre-breaking sub-forces relative to each other is represented by the coupling strength of the forces. For $SU(2)$, it is denoted by g and for the $U(1)$ group, g' . The mixing of the forces is defined by the Weinberg angle θ_w defined in Equation 2.6:

$$\tan \theta_w = \frac{g'}{g}. \quad (2.6)$$

In order to break the symmetry of the $SU(2)$ group without throwing away the rest of the gauge invariances of the model⁸, an additional scalar field (Φ) must be added that can transform under $SU(2)$. This field is known as the Higgs field.

This is done via the ABEGHHKL (Anderson-Brout-Englert-Guralnik-Hagen-Higgs-Kibble-Landau) mechanism, more commonly known as the BEH (Brout-Englert-Higgs) mech-

⁸If the theory breaks, why bend over back to fix it? When this mechanism was formulated, it was already known from the study of beta decays that the weak bosons had to be massive. Electromagnetism on the other hand worked well under this theory, had successfully already come up with a description that was accurate and replicated experimental data extremely well.




Figure 2.5: Image showing the self interaction vertices for the gauge bosons post symmetry-breaking. Since $F_{\mu\nu} = \partial_\mu A_\nu - \partial_\nu A_\mu + gf^* A_\mu A_\nu$ [23, p509], for a group to be non-Abelian, f^* must be non-zero. Since the term in the Standard Model Lagrangian is essentially the square of this, and each A_μ represents a bosonic field, in addition to the three point vertices, there are also the four point vertices: gggg and WWZZ.

anism or simply the Higgs mechanism [25–27]. The Higgs Mechanism adds a potential to the Standard Model that particles can interact with and it takes the form $\mathcal{V}(\Phi) = m^2\Phi^\dagger\Phi + \lambda(\Phi^\dagger\Phi)^2$ where m and λ are additional degrees of freedom, also known as free parameters [23, p584].

This potential has a quartic and a quadratic component to it, and varies with energy. For large values of energy, the potential has one stable minimum. As the energy gets smaller, the interplay between the quartic and the quadratic terms changes the shape of the potential as shown in the toy cross-section example of such a potential in Figure 2.6.




Figure 2.6: Image showing a rough sketch of the Higgs potential complete with a Higgs Boson for a fixed λ (the coefficient of the x^4 term) and a running m^2 (the coefficient of the x^2 term) for four different values of m^2 . The y-axis represents energy, and the x-axis represents the magnitude of the complex field, ϕ , assuming that the potential is symmetric in the phase. E_B signifies the point where the potential becomes non-trivial and this is taken to be $O(100 \text{ GeV})$. The potential is spherically symmetric such that there is a continuous ring of minimum point values. Image generated using [28].

If λ is assumed to be positive definite, then if at any point m^2 becomes negative, instead of one stable minimum, the potential becomes an unstable local maximum, with an annulus of points of lower potential. The fact that any potential can be selected is a

symmetry, and when that potential is selected, the symmetry is broken. The selection of this minimum causes a vacuum expectation value or VEV of the field⁹.




Figure 2.7: Image showing the Feynman vertices introduced to the Standard Model by the addition of a Higgs potential. Here, V generalises both the W -boson and the Z -boson as Vector bosons. In each diagram with a pair of V 's, both species must be the same. The self interaction diagrams are the reasons that the Higgs Boson has mass.

When the symmetry breaks, the massless $W_\mu^{1,2,3}$ and B_μ bosons form linear combinations of themselves to be infused with the VEV and produce the massive W^+ , W^- and Z^0 bosons (more commonly referred to as the W -boson and the Z -boson as both charges of the W -boson have the same mass) and a massless γ (photon) that mediates the remnant $U(1)_{EM}$ group. The VEV (or ν) of the Higgs potential is given by the relation of

⁹Before the symmetry breaking, this field does indeed already have a value of the potential, but we take this to be zero. The main thing that is of interest is the difference between the old and the new values of the potential and it is this difference that is called the VEV.

the two free parameters of the Higgs potential.

$$\nu = \frac{m}{\sqrt{\lambda}}. \quad (2.7)$$

The masses of all the bosons that transform through SU(2) are given by the combinations of Equations 2.6 and 2.7 given in 2.8. The latest measured masses of the Higgs Boson, the W-bosons and the Z-boson are 125.10 ± 0.14 GeV, 80.379 ± 0.001 GeV and 91.187 ± 0.002 GeV [29].

$$m_H = 2m^2 \quad | \quad m_W = g \frac{\nu}{2} \quad | \quad m_Z = \frac{\nu}{2} \sqrt{g^2 + g'^2} = \frac{m_W}{\cos \theta_w}. \quad (2.8)$$

$\mathcal{L}_{Higgs}$ adds both a Higgs dynamics term for the field and a potential. This term of the Standard Model Lagrangian details how it interacts with the gauge bosons of the broken SU(2) and the gauge U(1). This is shown in Figure 2.7.

2.3.4 $\mathcal{L}_{Yukawa}$: Giving mass to fermions

The masses for the fermions are given by the Higgs VEV and the Higgs fermionic coupling known as a Yukawa coupling (y_f). In the Feynman picture this is the inclusion of one more vertex shown in Figure 2.8. The stronger the interaction between a particular fermion and the Higgs, the larger its rest mass (m_f) will be. The relation is given by Equation 2.9:

$$m_f = y_f \frac{\nu}{\sqrt{2}} \quad (2.9)$$

This equation is important for this thesis as the coupling strengths y_f are all free parameters. Given this flexibility, the best way of integrating possible new physics models into the Standard Model would be via these Yukawa terms. The true values of these will only ever be known, therefore, by taking experimental measurements. Chapters 8, 9 and 10 will detail some of the work I did looking into one of these vertices: the $b\bar{b} - H$




Figure 2.8: Image showing the one Feynman interaction between the Higgs Boson and fermions: the ffH vertex. It is this interaction that gives the fermions their mass and the interaction probability is proportional to the mass of the fermion.

vertex. Equivalently, for convenience later on, it can be written as the $H \rightarrow b\bar{b}$ vertex or $H \rightarrow b\bar{b}$.

At particle colliders, the energy deposits and measurements that can be accessed are those from final-state particles. These final particles are going to be descendants of the initial particles that come from the vertex of choice.

The best and easiest way of measuring these properties is to design an experiment to separate the final states in the detector dependent on processes that have the vertex of interest from those that do not. A collection of events that contain a $H \rightarrow b\bar{b}$ vertex can be used to extract the relevant factors, often with ratios to cancel out common terms, and check consistency of measurements with the Standard Model.

One of the things that would be most useful to measure from the $H \rightarrow b\bar{b}$ vertex is the *branching fraction* of the Higgs boson to b -quarks: the fraction of the Higgs boson decays that are to these particular daughter particles. Since the b -quark is the heaviest particle in the Standard Model that the Higgs boson can decay into kinematically [9], and the

Higgs boson couples proportionally to the squared mass of that particle, the $H \rightarrow b\bar{b}$ decay is the most prevalent. Until 2018 [1], this decay had not yet been observed. Figure 2.9 below shows our current understanding of the branching fractions of the Standard Model interactions with the Higgs.




Figure 2.9: Pie chart showing the decay ratios of the Higgs boson to other particles. Even though pairs of W - and Z -bosons are heavier than the Higgs rest mass, the masses of the virtual particles can differ from the nominal masses such that the decay is still kinematically possible. This is known as being produced ‘off-shell’ and the asterisk shows that at least one of the particles is produced in this manner. The Higgs boson also can decay into gluons via an intermediate (top) quark loop. These branching fractions are dependent on the energy of the Higgs particles. These ratios are taken from a Higgs with its nominal mass.

The utility of this Yukawa term is such that the question can be flipped on its head. All particles that have mass whether predicted in the Standard Model or not should have a Yukawa coupling term¹⁰. By measuring the partial decay fractions for all the particles that are known, it can be inferred how much of the Higgs branching fraction

¹⁰While this is true for neutrinos if they are Dirac particles, the way neutrinos interact with the Higgs is different in the case that they are Majorana particles.

is made up of particles that we do not know. This requires measuring tiny couplings and is not the most favourable approach. In the case of $H \rightarrow b\bar{b}$ as the largest branching fraction, measuring the difference between the predicted and observed values of the number of events present in a regions of interest (called the signal strength parameter, μ) which is dependent on the branching fraction, can already constrain contributions from additional non-SM processes.

2.4 Probing beyond the Standard Model ($\mathcal{L}_{BSM}$)

As mentioned in the Introduction, the Standard Model does not explain all of the phenomena that are observed. Free parameters notwithstanding, there are some traits of particles that are not understood within the current framework of the Standard Model itself ¹¹. However ‘Beyond the Standard Model’ physics or BSM can mean more than the direct detection of new particles. BSM vertices can contribute to higher order Standard Model processes via virtual particle loops. These processes if they exist will augment existing known Standard Model processes so therefore these known processes need to be accurately predicted and measured.

Standard Model Quantum Field Theory or SMQFT is a *mass-order-4* QFT theory. This means that the sum of the mass-dimensions of all the particles in allowed Standard Model particle vertices is ≤ 4 . Fermions have a mass-dimension of $3/2$, and bosons have mass-dimension 1. Tensor fields, which are a linear combination of bosonic fields, have a mass-dimension of 2 ¹².

Since the Standard Model works well at ‘low energies’, one of the most popular ways of extending the Standard Model is to assume that it is a low energy approximation of

¹¹*squints eyes* I’m looking at you neutrinos!

¹²A Z-boson decaying into two quarks is an interaction between two fermion fields and a bosonic field. It is an allowed process in the Standard Model because the sum of the mass dimensions of these particles ($3/2 + 3/2 + 1$) is 4.

a more fundamental theory. This means the addition of higher-order terms in mass-dimension to the model. This is Standard Model Effective Field Theory or SMEFT ($\mathcal{L}_{SMEFT} = \mathcal{L}_{SM} + \mathcal{L}_{BSM}$).

By constructing BSM effects that interact with Standard Model observables, it is possible to determine how extra higher order terms in the Lagrangian affect the theory.

2.4.1 Parametrising beyond-the-Standard-Model (BSM) physics

The generic form of the SMEFT Lagrangian looks like this:

$$\mathcal{L}_{SMEFT} = \mathcal{L}_{SM} + \sum_{n=5}^{n=\infty} \sum_i \frac{c_i}{\Lambda^{n-4}} \mathcal{O}_i^{(n)} \quad (2.10)$$

where n is the mass dimension, Λ is a relatively arbitrary choice of cut-off energy scale (chosen here to be 1 TeV) and the c_i are the coefficients of the Wilson Operators $\mathcal{O}_i$. At each mass dimension there is a different number, i of these Wilson Operators (for example in the most generic basis; at $n=5$, there is 1 Wilson Coefficient; and at $n=6$, there are 2499).

The cut-off energy scale is the point at which the higher order terms have such a small effect that *SMEFT* simplifies to the Standard Model. Adding a dependence on this energy scale to the mass-order ensures that the theory is renormalisable (each successive term's contribution has a smaller effect).

Odd-numbered terms in the bracketed exponent of the Wilson coefficients are cut because they violate lepton and baryon number conservation [30],[31, p2]. This is not normally seen in interactions within the Standard Model and its effect is mostly orthogonal to those that are desired, so for the analyses discussed later in this thesis, they are ignored. In addition, terms of order-10 or greater in Λ are ignored due to the

negligibility of their effects. With these assumptions Equation 2.10 becomes Equation 2.11.

$$\mathcal{L}_{SMEFT} = \mathcal{L}_{SM} + \sum_i \frac{c_i}{\Lambda^6} \mathcal{O}_i^{(6)} + \sum_j \frac{c_j}{\Lambda^8} \mathcal{O}_j^{(8)} \quad (2.11)$$

While direct order-8 terms are theoretically important to include, they are second-order effects (with respect to those brought about by the order-6 operators) and there are so many of them, most generators cannot produce enough simulations to make this addition worthwhile. In order to make a first attempt at this process, the second-order effects from the dimension-8 operators are ignored.

In the most generic SMEFT regime, there are some 2499 operators. To reduce this to a manageable number a *basis* is chosen to organise the particle interactions into operators and some assumptions on the particles are made¹³. The basis selected is the ‘Warsaw basis’ chosen because it is *complete*, so is the most general parametrisation of BSM physics and therefore the operators themselves have the minimum model dependencies. With the selection of the Warsaw basis and these assumptions, the number of operators is reduced to 59. This number is reduced further as not all of the operators will affect the VH , $H \rightarrow b\bar{b}$ process. Of these 59 operators, only 15 were relevant to the analysis (detailed in Chapter 8) and only 4 had any reasonable impact.

2.5 Theoretical conclusion

To explore the parameter space able to test these ‘SMQFT’ or ‘SMEFT’ theories, high-energy particles need to be created such that collisions of these particles have enough energy to access new energy frontiers. For the data that pertains to studies in this thesis, how this is done is described in the following chapter.

¹³These assumptions are that: all the masses of non-b-quarks and leptons are 0; the CKM matrix is completely diagonal; and the SM follows a $U(3)^5$ group structure.

CHAPTER 3

CERN and the LHC

3.1 CERN

 rganisation Européenne pour la Recherche Nucléaire, formerly Conseil Européen pour la Recherche Nucléaire, from which the popular acronym 'CERN' originates, was founded in the 1950's as a pioneering European project. Its aim was to unite and attract scientific talent to stay in post-war Europe and collaborate on a large scale, bridging the borders of nations such that scientific progress could be achieved for the benefit of mankind independently of the politics between nations. For its famous political neutrality, Switzerland was chosen as the country to host the facility. Another advantage of CERN was the ability to share the burden of rising costs of cutting edge nuclear research among several nations.

On the 11th May 1957 CERN's first accelerator, the Synchrocyclotron (SC), started up and with it came the particle beams for CERN's first particle physics and nuclear experiments [32]. Over the past 65 years CERN has grown to have 31 member states and affiliate nations and now hosts 25 experimental facilities of varying sizes including the

Anti-Matter Spectrometer [33] (AMS), which is currently on the International Space Station; ISOLDE, which was CERN's first facility based on ion research and which is still in operation today, the fixed target experiment NA62 [34]; and the collaborative giants of the ATLAS [35] and CMS [36] experiments. Some of these can be seen in figure 3.1 below.




Figure 3.1: Image showing the CERN accelerator machine complex complete with the particles that traverse each of the sections [37].

Arguably, the most famous thing about CERN is its collider physics programme. Various record-breaking colliders have been built at CERN, starting with CERN's first synchrotron, the Proton Synchrotron (PS). Built in 1959 and with a circumference of 628m [38] the PS accelerated particles¹ up to 25 GeV and used a complex of room tempera-

¹To this date the PS has accelerated electrons, positrons, protons, anti-protons and the nuclei of helium, oxygen, and sulphur.

ture electro-magnets to bend the particle beams into a circle². The PS's main focus was to provide beams for many experiments across three sites including a number of bubble chambers and some fixed-target experiments. Over the years successive colliders have been built with larger circumferences and have accelerated particles to higher and higher energies.

In 1976 the Super Proton Synchrotron (SPS) was completed. This 7 km machine allowed a single beam of protons with beam energy of 400 GeV to be collided against fixed targets [39]. In 1981, after a proposal by Carlo Rubbia³, the SPS accelerator was upgraded to deliver collisions between beams of protons and anti-protons, which allowed for a $\sqrt{s}$ of up to 900 GeV [39]. The SPS was now dubbed the $S\bar{p}\bar{p}S$ and four experiments were built around the accelerator in colliding-beam mode: Underground Area 5 (UA5) designed to measure Centauro events⁴; UA1 and UA2, the general-purpose detection experiments designed to try and detect the W and Z bosons predicted by Sheldon Glashow, Steven Weinberg, and Abdus Salam in 1968 [40, 41] (which won them the Nobel Prize in 1979 [42]); and UA4, designed to measure the proton-antiproton ($p\bar{p}$) cross-section. In 1983, the W and Z bosons were discovered at UA1 and UA2, which earned the head of the UA1 collaboration Carlo Rubbia and the accelerator physicist Simon van der Meer⁵ the 1984 Nobel prize [43].

After the successful run of the SPS, an even more ambitious project was undertaken: the Large Electron Positron collider (LEP). Lying in a tunnel of 27 km in circumference

²This is because a charged particle travelling in a direction perpendicular to a (dipole) magnetic field has a force exerted upon it perpendicular to both these directions. This allows the direction of the charged particles to be manipulated. The strength of the magnetic field and the momentum of the particles are the two main factors in determining how large the circular collider has to be to contain the particle beam at the turning circle of its optimal energy.

³Carlo Rubbia suggested the adaptation of the SPS to collide protons with antiprotons and recognised the need for stochastic cooling of the anti-protons the storage beams in order to get the energies and intensities needed to produce W and Z bosons at SPS.

⁴a form of anomalous events in cosmic-ray experiments first seen in 1972.

⁵Simon Van Der Meer came up with and developed the idea of stochastic cooling which corrects the energy spread of beams in circular colliders to make them more stable.

which took three years to bore, on average 100m below the ground, LEP collided electrons and positrons at four collision points corresponding to the detector caverns of four experiments on the LEP ring: ALEPH, DELPHI, L3 and OPAL. The main shared aim for all these experiments was to take data to provide a detailed study of the electroweak interaction. Initially LEP operated for seven years around the pole-mass of the Z-boson at an $\sqrt{s}$ energy of 100 GeV to maximise the production of the newly-discovered particles.

Observing in more detail the creation and decay of the Z-boson was a critical test of the Standard Model as it was the heaviest particle known at the time, and the one therefore, with the largest phase space for new physics. In 1995 LEP was upgraded for a ‘phase two’ of operation, with the aim to run at double the energy so that the resulting collisions could produce pairs of W-bosons [44]. LEP climbed in energy and would eventually reach $\sqrt{s} = 209$ GeV in 2000. Measurements performed at LEP during this time on the decay-width of the Z-boson (which is proportional to the number of particles it can decay to) also proved that there are only three generations of light particles of matter⁶ [45, p61].

The LEP collider was closed in November of the year 2000, and in the same tunnels a new circular hadron⁷ collider was built; The Large Hadron Collider or LHC.

3.2 The Large Hadron Collider

The LHC is currently the world’s largest and most powerful particle accelerator. While the previous colliders built at CERN used normal-conducting electromagnets to bend and focus beams, the LHC was the first one (at CERN) to be built with superconduct-

⁶There could be more generations of particles, but they must be heavy enough such that the Z-boson cannot kinematically decay to them, and these have yet to be seen.

⁷The name given to particles made up of multiple quarks - like protons

ing magnets from the outset⁸. An array of magnets that guide charged particle beams is called a magnet lattice. The LHC magnet lattice consists of various kinds of superconducting magnets cooled to 1.9K using liquid helium. 1232 main dipole magnets apply a force going into the collider centre to bend the beams and keep them going in a circular trajectory [46, 47]. 392 quadrupole magnets focus the beams in the plane perpendicular to the direction of travel making the beam profile narrower. This process is energy dependent and causes the beam to spread along the beamline [46, 47]; 688 lattice sextupole magnets and 168 octupole magnets introduce an energy-dependent force along the direction of the beamline to restore the particle bunch integrity along the direction of the beamline⁹ (by introducing a force inversely proportional to the distance from the centre of the magnet which corrects the spread of energy of the particles in the beams [46]). Figure 3.2 shows an aerial view of the LHC.

The LHC houses seven experiments in total. The four large experiments are: ATLAS (**A** Toroidal LHC Apparatu**S**¹⁰) and CMS (**C**ompact **M**uon **S**olenoid), the general-purpose discovery experiments; LHCb [49] (**L**arge **H**adron **C**ollider **b**eauty¹¹) designed to record composite particles with b-quarks in them; and ALICE [50] (**A** **L**arge **I**on **C**ollider **E**xperiment) whose main aim is to replicate the conditions of the early universe with their own experimental caverns and collision points. The three small experiments are: TOTEM [51] (**T**OTAL Elastic and diffractive cross-section **M**easurement) designed to measure basic properties of the colliding beams; LHCf [52] (**L**arge **H**adron **C**ollider-forward) an astrophysics experiment measuring properties of π^0 particles in order to better understand ultra-high-energy cosmic rays; and MoEDAL [53] (**M**onopole and **E**xotics **D**etector **A**t the **L**HC), whose aim is to search for highly-ionising stable and

⁸In LEP phase two, LEP upgraded their magnets to superconducting ones.

⁹With all these accelerating cavities attached to each other the LHC is less of a circle and more of a dischilitetrakosiogdoëkontagon (2480 sided shape).

¹⁰This commits so many acronym crimes it isn't funny. More like A Terribly Lamé Acronym Sighted.

¹¹LHCb was to be called BEAST - Beauty Experiment At Small Theta, which I think we could all agree is a better name but alas the name was taken. Though in my opinion LHCb's logo far and away makes up for the lacklustre name.




Figure 3.2: View of the Swiss/French countryside above the LHC tunnel. There are several places at which the tunnels have service access shafts. These are referred to as points. The ATLAS detector cavern is stationed at Point 1. The ALICE, CMS and LHCb detector caverns are housed at points 2,5, and 8, respectively. Image from [48].

pseudo-stable massive particles and magnetic monopoles.

The smaller experiments are designed to share an interaction point with a larger experiment and are roughly tens of meters away from their respective interaction points. All these experiments straddle the beamline and take data from proton-proton (pp) collisions, as well as collisions: of protons and lead nuclei (p -Pb); and between lead nuclei (Pb-Pb) and other heavy ions thanks to the LEIR (Low Energy Ion Ring) machine.

The LHC was built with the aim of colliding beams of various particles together at energies of up to 14 TeV. It is not efficient for these particles to be accelerated from rest to these energies in a single machine, so the particles reach these energies in stages by using the older colliders as injectors to further rings. This forms the CERN accelerator complex. For the LHC, the accelerator chain is depicted in Figure 3.3.




Figure 3.3: Labelled diagram of all the parts of the CERN accelerator complex that are important for the running of the LHC. Image from [54].

The particles in the beams that end up circling the LHC roughly 11,000 times a second (99.999999% the speed of light) come from unostentatious origins. Let us take the example of proton-proton collisions as the data from those forms the basis for the analyses and results later on in this thesis.

The protons start as molecular hydrogen (H_2) in a small hydrogen tank. This H_2 is then passed into a duoplasmatron where it is ionised to form $^1H^+$; which is a proton. These protons are then accelerated by LINAC2¹² from rest to 50 MeV, where they enter the Proton Synchrotron Booster Ring (PS Booster). The PSBooster is made from 4 stacked synchrotron rings that accelerate the protons from 50 MeV to 1.4 GeV before injection into the PS, which in turn further accelerates the protons to 25 GeV for injection into

¹²LINAC2 replaced LINAC1 in 1978 and will be replaced by LINAC4 later on this year. LINAC3 is used to accelerate heavy ions into LEIR.

the SPS. This 25 GeV beam is also delivered to experiments that are based in the East Area. The SPS boosts the beam energy from 25 GeV to 450 GeV. As of today the SPS has offshoot beamlines that service the NA62 and COMPASS experiments (among others). This 450 GeV bunch of protons is inserted into the LHC circulating in one direction and the entire process is repeated for the second LHC beam circulating in the other direction. The beams are cleaned, ramped up to the required collision energy (6.5 TeV), and stabilised before being collided together every 25 nanoseconds in the 4 locations of the large experiments [55].

Why accelerate the particles in this way with many stages? In order to be useful particle beams have to be collimated, chromatically stable, and organised in a bunch structure. Using only one machine to do this makes the job of attaining the perfect beams much harder, as the superconducting magnets have to be able to cope with beams from a much larger range of energies, which means they need to be able to operate under a wider array of modes which can cause issues like quenching. Using existing architecture to accelerate particles is cheaper as it allows each stage to be designed/retrofitted to deal with particle beams of a much narrower set of parameters for controlling the beam. This makes it more likely to achieve stable running conditions for the beams overall. In addition to this, the use of multi-tiered infrastructure allows for more flexibility in terms of machine variety that can be cost-effectively operated at once. At each stage or circular ring, an experiment can be run with a different target beam energy and not affect the other experiments that lie earlier or later on in the chain (i.e the SPS can run the NA62 and COMPASS [56] experiments while the LHC provides data to ATLAS, LHCb, etc.).

3.2.1 Why and How: Producing Higgs Bosons at the LHC

To study any particle reliably we need to be able either to observe a natural process that produces the particle, or to produce them in a controlled manner. For the Higgs,

only the latter option is available. Higgs particles (denoted as H) are produced at the LHC by interactions that occur when particles are collided. The probability of a particular interaction occurring as a result of two colliding particles (in this case $pp \rightarrow H$) is proportional to the reaction's production cross-section ($\sigma_{pp \rightarrow H}$). The rate of expected Higgs events R_H produced at a pp collider is defined as this production cross-section multiplied by the instantaneous luminosity of the collider shown in Equation 3.1.

$$R_H = \mathcal{L}_{inst} \times \sigma_{pp \rightarrow H}. \quad (3.1)$$

3.2.1.1 Instantaneous and Integrated Luminosity

The *instantaneous luminosity* ($\mathcal{L}_{inst}$) of a collider is a measure of how many particles are able to be squeezed through a given area in a given time. It is not equal to an *event rate*, $R = N/s$ as not all the particles forced into this area in this time will interact, but the larger the $\mathcal{L}_{inst}$ of a collider, the more likely an interaction is to occur and the more events of a particular process are likely to be created. Since the particles are collided in bunches, the figure of merit is the *bunch luminosity* or $\mathcal{L}_b$ where the total instantaneous luminosity is given by

$$\mathcal{L}_{inst} = \Sigma_b \mathcal{L}_b = g \mathcal{L}_b, \quad (3.2)$$

where g is the number of bunches in a beam if all the bunches are identical. In a circular collider where the beams travel through each other several times in a given time window, and there are two identical beams of particles colliding head-on, the bunch luminosity is given by

$$\mathcal{L}_b = \frac{f_r n^2}{4\pi\sigma_h\sigma_w} \quad (3.3)$$

[57], where the *revolution frequency*, f_r , is the number of times a given bunch crossing will pass through the same interaction point in a given period of time, n is the number

of particles in a bunch, and σ_h and σ_w are the heights and widths of the particle beams, respectively.

Integrating out the time dependence on the instantaneous luminosity $L = \int \mathcal{L}_{inst} dt$ gives the number of particle collision events per target cross-sectional area in units of inverse femtobarns (fb^{-1}), where the barn is a unit of area equal to 10^{-24} cm^2 . The unit is used for time-integrated luminosity, since the instantaneous luminosity of a collider is usually given in $\text{cm}^{-2} \text{ s}^{-1}$. The design (or nominal) luminosity of the LHC, for example is $1 \times 10^{34} \text{ cm}^{-2} \text{ s}^{-1}$ [58]. When the LHC operates the total amount of data taken is measured in inverse femtobarns.

3.2.1.2 Cross-sections

In Particle Physics a *cross-section* is a measure of the probability of a specific process occurring when two particles collide. When this happens, a large variety of processes could take place; some more likely than others. For example when two protons interact, they can simply bounce off one another, both protons remaining intact; or they could create a Higgs boson. The larger the cross-section of a process, the larger the chance of the corresponding interaction when two particles collide. The units of cross-sections are given in area and are normalised to the area of the colliding particles. The term ‘cross-sections’ originated in nuclear physics where the colliding particles were atomic nuclei with a diameter of approximately 10^{-14} m . Hence, the maximum area for two nuclei colliding head-on is therefore 10^{-28} m^2 also known as a *barn* (b), which is also equal to 10^{-24} cm^2 .

For many rare processes such as $pp \rightarrow H$, a barn is many orders of magnitude larger than is typical, so nanobarns (nb), picobarns (pb), or femtobarns (fb) are used instead. The cross-section for Higgs boson production, for example, is low (≈ 1 in every 10^{11} 13 TeV proton-proton collisions will produce a Higgs particle) so to be able to effec-

tively probe its properties, we need to be able to produce and detect large numbers of them such that we can measure their properties with high precision. The only particle collider in operation currently capable of doing this is the LHC at CERN, which collides particles with high enough centre of mass energies and luminosities to produce plenty of events with Higgs candidates in them. For example, the number of Higgs events produced (N_H) at the LHC is given by,

$$N_H = L \times \sigma_{pp \rightarrow H}. \quad (3.4)$$

Thus if a detector has accumulated 2 fb^{-1} of integrated luminosity over a period of time, a process with a total production cross-section ($\sigma_{pp \rightarrow H}$) of 55 pb ($= 55000 \text{ fb}$) [59, p8-15] is expected to produce 110000 Higgs events.

3.2.2 A brief look at what the LHC has done

The LHC construction was finished in 2008 with the aim of colliding particles at $\sqrt{s} = 14 \text{ TeV}$. However, while powering up the magnets on the LHC, a failure in the electrical bus between two dipole magnets occurred which caused an electrical arc in the helium enclosure, venting several tonnes of liquid helium into the tunnel [60]. The resulting liquid rapidly vaporised with explosive force and damaged about 30 magnets, but more had to be removed for cleaning. This had three effects on the operation of the LHC. Firstly, pp physics data taking did not commence until 2009; secondly, the beams were made to collide at a lower instantaneous luminosity than the intended design; and thirdly, thus far the colliding beams have not yet been circulated at the design energy.

Consecutive years of operation and data collection are called Runs. The LHC Run-1 started in September of 2009 and started with a $\sqrt{s}$ of 900 GeV , climbing up to $\sqrt{s} = 7 \text{ TeV}$ in 2010, and finally $\sqrt{s} = 8 \text{ TeV}$ in 2012. The collider instantaneous luminosity was steadily increased across the run from about 50% to plateau at 75% of the design

luminosity in 2011. Some of the highlights of Run-1 include but are not limited to: ALICE confirming in 2010 that it had made the hottest ever quark-gluon plasma¹³ [61, p2] in Pb-Pb collisions; the LHC discovering its first new particle, the $\chi_b(3P)$ in 2011¹⁴; and LHCb discovering an exotic five-quark bound state: the ‘pentaquark’ [62]. The most important discovery in terms of importance to this thesis was made in 2012, when the data taken in 2011 and earlier was analysed and the Higgs Boson was discovered by the ATLAS [7] and CMS [8] collaborations (which earned the theorists François Englert and Peter Higgs the 2013 Nobel Prize [63]). Due to this discovery Run-1 was extended to February 2013 and ended up delivering roughly 29 fb^{-1} of data to the four main experiments.

After a 2 year break for machine maintenance, upgrades and repairs called Long Shutdown 1, Run-2 started in 2015 operating at a $\sqrt{s}$ of 13 TeV. Like in the last run, the luminosity was increased as the run went on reaching the design luminosity in 2016 and further increasing until it reached double the design luminosity by the end of the run. The 3rd of December 2018 marked the end of Run-2 for the ATLAS detector and 147 fb^{-1} of data was taken during Run-2 and physicists are now processing and analysing this data.

While the Higgs itself was discovered in Run-1, there was not enough data to make precise measurements on all of its most important properties. This is one of the main motivations that drives further runs, along with looking for phenomena not currently explained by the Standard Model such as dark matter. To be able to ensure that the Higgs particle that was found is the Standard Model Higgs boson, and is not hiding any new physics, comparisons need to be made between the decay branching ratios of the Higgs boson found using experimental data, and theoretical predictions. As

¹³A new exotic state of matter discovered in 2000. Hadronic particles ‘melt’ and the constituent quarks are free to move.

¹⁴a meson consisting of a b-quark and an anti-b-quark

discussed in Section 2.3.4, the Higgs boson couples to a particle with a strength proportional to its mass squared. The largest theoretical decay branching fraction is therefore to the heaviest fundamental particle-antiparticle pair whose mass is less than that of the Higgs boson: a b -quark and an anti- b -quark. This thesis will discuss two such $H \rightarrow b\bar{b}$ analyses of the Run-2 dataset (Resolved VH , $H \rightarrow b\bar{b}$ and Boosted VH , $H \rightarrow b\bar{b}$) later. The Higgs boson was not discovered in this channel, however, because of the difficulty in measuring hadronic final states at the LHC.

3.2.3 Limitations of the LHC

Colliding protons together to produce Higgs (and other) particles causes some further complications because protons are not fundamental particles. The energy available at the collision between two protons is not given by the energies of the protons, but by the energies of the constituent colliding partons¹⁵, which take some fraction of the protons energy. This parton energy is not well defined, so the energy available at the center-of-momentum is drawn from an energy distribution, and is usually significantly lower than the $\sqrt{s}$ that the collider reaches.

Given the current magnetic lattice at the LHC, the energy of the beams cannot be increased easily. To obtain meaningful statistics at higher energies without increasing the beam energy, more data is required in the tails of these energy distributions than is currently available. One way of achieving this is to increase the luminosity of the beam and have more particles colliding per unit time, effectively extending the energy scales

¹⁵Protons are made up of three quarks (also known as the valence quarks) held together by the strong force. The strong force is categorised by a particle having ‘colour charge’ and coloured particles like quarks exchange gluons between one another to interact. However the gluon also carries colour charge, so it can self-interact and produce more gluons. This means that in high energy regimes, the perceived colour charge of a particle changes with relative distance and the content of a particle comprised of quarks depends on the energy scheme being probed. The theory that governs interactions between coloured objects is called Quantum Chromodynamics, or QCD. At low energies, the picture of the proton being made up of three valence quarks is usually satisfactory, but in high energy regimes, the valence quarks radiate gluons, which self-interact and form more gluons, or pair produce to form quark pairs, and the proton effectively becomes a sack of colour-carrying objects. It is these colour-carrying sacks that are collided at the LHC.

that can be studied.

3.2.4 A short glance at what the LHC is going to do

By going to higher energies and luminosities than Run-1, Run-2 was able to produce more Higgs-candidate events to study. This thesis will look at the entirety of the Run-2 dataset which exceeded its target of 120 fb^{-1} to deliver 156 fb^{-1} compared to the 30 fb^{-1} that was present in the Run-1 dataset. In addition to this, since the full Run-2 dataset is now available, the resulting six-year gap between runs will mean that newer and more efficient techniques developed in that time-frame will be applied to Run-2 analyses to make them more effective at extracting information from the experimental data.

The LHC is planned to continue operation all the way up to 2040. Figure 3.4 shows the most up-to-date plan for the LHC machine in this time. Following Run-2, the LHC has entered Long Shutdown 2 to prepare the machine and to allow the collaborations to optimise their detectors for Run-3 and beyond. Run-3 aims to operate at double the nominal luminosity of the collider at an energy of 14 TeV for 3 years and aims to add another 160 fb^{-1} to the total LHC dataset.

After Run-3 has ended, CERN will enter Long Shutdown 3, and the final upgrades for the High-Luminosity Large Hadron Collider (HL-LHC) will be installed. Opening in about 2025¹⁶, the HL-LHC takes the idea of extending the luminosity much further. It is expected to start collisions with a fixed integrated luminosity of 5 times the nominal luminosity and increase over the lifetime of the HL-LHC to 7.5 times the nominal (and maybe even beyond). All the LHC detectors have various upgrade programmes to deal with the additional particle flux that this will generate.

The HL-LHC is due to start colliding particles in mid 2027 [58] and will accommodate

¹⁶Hopefully this is not set back due to the coronavirus outbreak.




Figure 3.4: CERN's official projection of the timeline and future development of the LHC Machine [58].

at least two further runs: Run-4 and Run-5. The aim is to deliver 2500-3500 fb⁻¹ of additional data to the LHC experiments over the course of the ten years of its operation. By increasing $\mathcal{L}_{inst}$ by these factors, and by running for so long, the maximum possible energies that can be probed for new particles can be extended by between 300 GeV and 400 GeV higher than the current LHC. Another effect is that with the full 3000-4000 fb⁻¹ of data, certain measurements will go from an observation (3σ tolerance) to a discovery (5σ tolerance). The HL-LHC will also be able to increase the precision on current Standard Model measurements.

Of the four larger experiments on the LHC, the ATLAS experiment is the largest. It is at this facility that the research reported in this thesis takes place. The following chapter looks at the detector in more detail.


Part II



The ATLAS Experiment

“Where a computer on the ENIAC is equipped with 18,000 vacuum tubes and weighs 30 tons, machines in the future may have only 1,000 vacuum tubes and weigh only 1.5 tons”.

- Popular Mechanics, March 1949



**

The ATLAS detector

4.1 The ATLAS experiment

The ATLAS experiment is one of seven experiments on the LHC's 27 km ring. It is the largest collaboration of active scientists ever to have existed, currently with over 3,000 scientific authors [64]. ATLAS is a 'general purpose' discovery detector, designed to take full advantage of the largest man-made energy scales to make discoveries and observe phenomena that involve highly massive particles; processes that could not be observed using the previous generations of lower-energy accelerators.

Figure 4.1 shows the cumulative integrated luminosity of data delivered to the ATLAS experiment over the course of the experimental runs: Run-1 and Run-2.

Run-2 provides most of the data for the results in this thesis. Of the 156 fb^{-1} of data delivered to the ATLAS experiment over the course of Run-2, 147 fb^{-1} of it was recorded. After removing data with with detector problems, 139 fb^{-1} of it was post-processed and




Figure 4.1: Cumulative luminosity delivered by the LHC and recorded by ATLAS detector for the entire lengths of Run-I (a) [65] and Run-II (b) [66]. The total physics-quality data recorded over Run-1 and Run-2 is 163.87 fb^{-1} .

disseminated into various smaller formats for physicists on which to perform analyses.

4.2 The ATLAS detector

The ATLAS detector is cylindrical in shape measuring 25 metres in diameter by 46 metres in length and weighing roughly seven kilotonnes. The detector is symmetric in the plane parallel to the beam direction at the interaction point (IP). About this IP, it has a solid angle coverage of nearly 4π steradians. The ATLAS detector is made up of three sections: a central barrel region and two ‘endcap’ regions. In each section, the ATLAS detector is made up of five parts: the Magnet System, the Inner Detector, the Hadronic Calorimeter, the Electromagnetic Calorimeter, and the Muon Spectrometer. Figure 4.2 shows a graphical depiction of this.

The sub-detectors can be considered as one of two types: those that identify the path of charged particles (trackers); and those are designed to absorb the particle and try to determine its energy irrespective of its charge (calorimeters).




Figure 4.2: Schematic of the ATLAS detector: Red = Inner Tracker, Orange = EM Calorimeter, Green = Hadronic Calorimeter, Yellow = Magnet System, Blue = Muon Spectrometer. *Tyrannosaurus Rex* for scale. Image from [54].

4.2.1 The ATLAS co-ordinate system

A Cartesian co-ordinate system can be used to denote the placement of different parts of the detector. The origin is the nominal interaction point at the centre of the detector. The x -axis points towards the centre of the LHC ring, the y -axis points upwards towards the Earth's surface, and the z -axis points along the direction of the beamline. Given that the detector is roughly cylindrical in shape, and most of the information that is recoverable about the particles passing will be perpendicular to the beamline, cylindrical polar co-ordinates are more often used. In this scheme, the x - y plane, also known as the transverse plane is described by the variables ' r ' and ' ϕ ' denoting the radial distance to the beam and the azimuthal angle about the z -axis respectively. In the x - z plane, the angle subtended by a particle and the beamline is given by θ , but as the particles travel in the z -plane, these will have momentum boosts along that axis. This means that depending

on the reference frame, this angle will change. This is called a non-Lorentz invariant quantity and is not ideal for experiments. The quantity $\eta = -\ln[\tan(\frac{\theta}{2})]$, also referred to as the pseudo-rapidity, is a function of this angle θ , is used instead. This is done because pseudo-rapidity differences are Lorentz invariant¹. The ATLAS barrel region encompasses $|\eta| < 1.05$, and its endcaps $1.05 < |\eta| < 4.9$.

4.2.2 Muon Spectrometer

Starting from the detector components furthest away from the beamline, we have the Muon Spectrometer (MS). With a lifetime of 2.2×10^{-6} s [9, p16], muons are long-lived particles in particle physics terms. A 1 GeV muon will have an average lab-frame free decay range of ~ 6.3 km² and this number gets larger, the more energetic the muons are. Even though the presence of material effects effectively reduces this distance considerably, muons will still traverse the entirety of the detector without decaying and is the only known charged particle to be able to do so consistently. The radially furthest parts of the detector, therefore, are there to report any hits from charged particles whose trajectory has been bent by the surrounding toroidal magnet (more on that later), and to trigger on them. When used alongside the Inner Detector (more on that later as well), the MS identifies and measures the charge sign and momentum of muons.

The Muon Spectrometer reconstructs tracks via information from four different types of component: Monitored Drift Tubes Chambers (MDTC's), Cathode Strip Chambers (CSC's), Resistive Plate Chambers (RPC's) and Thin Gap Chambers (TGP's). It only has solid angle coverage of 2π and up to $|\eta| < 2.7$ in pseudo-rapidity. Labelled parts of the MS with respect to the ATLAS detector can be seen in Figure 4.3.

¹The pseudo-rapidity η is a low-velocity approximation of the rapidity y , which is Lorentz invariant. For massless particles, η and y are equivalent; and for particles whose masses are small compared to their energy, η and y are approximately equivalent.

²Calculated using $x = \gamma c\tau$. γ is given roughly by the energy of the muon divided by its mass. The muon is assumed travelling close to 3×10^8 m/s (c), which is a fair assumption for all subatomic particles produced at large energy collisions at the LHC.




Figure 4.3: Schematic of the ATLAS detector with particular emphasis on the components that make up the Muon Spectrometer. Image from [67].

4.2.2.1 Tracking

Since the MS is far away from the beamline, it needs to cover a larger area to get the same solid angle coverage, but the granularity of this system is not as important to determine the path of a particle, and the components need not be ‘radiation-hard’. As drift tubes are a cheap(er) way to cover a lot of detector-area, the bulk of the active detector material in the MS is made up of them. ATLAS’ MDTCs are aluminium tubes 3 cm in diameter filled with a pressurised mixture of 93% Argon and 7% Carbon Dioxide (and a few parts per million of water vapour), with a $50\text{ }\mu\text{m}$ diameter anode wire in the centre [68, p3].

As a charged particle passes through a drift tube, it ionises the atoms of the gas in the tube. The resulting electrons then ‘drift’ in an electric field to the centre positively-charged wire where the signal is collected.

The MDTCs perform precise tracking in the particles’ r - z plane, also referred to as the ‘bending plane’ with the momentum measurement being based on the curvature of the track in this plane.

The MDTCs in the barrel are arranged in three interlaced cylindrical layers called stations. A station has several units. Each unit contains two *multilayers*³ of MDTCs separated by a spacer. Figure 4.4 shows a clearer picture of how this looks.

In the endcaps, the MDTC multilayers are arranged in three ‘main’ radial discs. A fourth auxiliary annulus of endcap MDTCs covers the limited pseudo-rapidity region of $1.05 < |\eta| < 1.4$ spanning from the barrel to the endcap known as the ‘Transition Region’.

While MDTCs provide most of the precision tracking coverage in the detector, at larger pseudo-rapidities ($2 < |\eta| < 2.7$), where the detector components are closer to the beamline, the rate of particles incident on this area is larger. Thus the innermost region of the first endcap ring is composed of CSCs instead of MDTs. The CSCs are Multi-Wire Proportional chambers with additional cathodes set into the wire complex. This provides higher detector granularity [69, p13].

To ensure a high performance of the MS over the large areas involved, the accuracy to which the muon chambers are aligned both with respect to each other and with respect

³A multilayer consists of either 3 or 4 planes of between 12 and 72 MDT chambers 1 to 6 metres long. The inner stations have shorter MDT chambers, and the outer stations have longer ones. The number of MDTCs in each plane and the number of planes in a multilayer also depend on how close the multilayer is to the IP. The same is true for the endcap regions.




Figure 4.4: Schematic of the Muon Spectrometer shown here in indigo and violet. The dashed yellow lines indicate the presence of the magnetic toroids and the grey hashed lines represent non-tracking infrastructure. This image shows how the layers in the barrel interlace in a cylindrical fashion and how they surround the barrel toroid coil. Peter Parker and accompanying spider for scale. Image from [54].

to the overall detector is of the utmost importance. Roughly 12000 alignment sensors monitor the internal multilayer deformations and relative positions between MDTCs. The algorithms for doing this are all based on monitoring the deviations on straight lines through the MS. The precision of the muon momentum measurement requires a $30\text{ }\mu\text{m}$ spatial precision on the alignment of these chambers [69, p13].

MDTCs and CSCs, however, relying on low-field charge collection are not fast enough to be able to send information to the trigger to make live decisions. The speed at which these detector components can be read out would cause a backlog. So another system is needed in addition to the MDTCs to provide time-sensitive trigger information.

4.2.2.2 Triggering

The Resistive Plate Chambers (RPCs) and Thin Gap Chambers (TGCs) attached to MS components in the barrel and endcap regions respectively are tasked with providing information about the particles fast enough for rapid trigger decisions to be made. They do this by providing three pieces of information:

- which bunch-crossing the particle was detected in.
- a well-defined threshold on transverse momentum of the particle.
- information about the particle's position in the r - ϕ 'non-bending' plane orthogonal to the bending plane.

A cross-section representation of the MS is shown in Figure 4.5.

4.2.3 Calorimeters

The calorimeter systems are the next thing we encounter on our journey inwards. With the obvious exception of the muons discussed in the previous section, the calorimeters are designed to stop all particles, barring neutrinos, within its volume and measure




Figure 4.5: Schematic of the Muon Spectrometer showing how the MS was designed such that particles passing through it will interact with at least three chambers. Image from [70].

their energies. The calorimeter system at ATLAS is made up of three sub-systems: the Hadronic Calorimeter (HCal), the Electromagnetic Calorimeter (ECal) and the Forward Calorimeter (FCal), and has a η coverage of $-4.9 < |\eta| < 4.9$. These can be seen in Figure 4.6.

It has been known for a long time [72] that heavy ($m \gg m_e$) charged and neutral particles traversing through matter lose energy in a variety of ways⁴. For heavy charged particles the effectiveness of materials in causing this energy loss is referred to as the *stopping power* and is given by $S[E] = -\frac{dE}{dx}$ [73, p23]. For neutral particles which lose energy via interactions with nucleons and nuclear interactions only, the average distance between nuclear interaction or collision length, λ is the more appropriate figure of merit.

⁴Including but not limited to: Elastic nuclear scattering, Inelastic collisions of atomic electrons, Cherenkov radiation, Transition radiation, Bremsstrahlung and nuclear reactions.




Figure 4.6: Schematic showing the Barrel Calorimeter System at ATLAS. Image from [71].

For a charged particle, the total range a particle with an energy E will travel in a particular medium is therefore given by;

$$\mathcal{R} = \int_0^E \frac{1}{S[E]} dE. \quad (4.1)$$

The ATLAS hadronic and electromagnetic calorimeters are *sampling calorimeters* [69, p4] which means that they aim to determine the energies of particles through the usage of two types of material used in alternation: an ‘absorber material’ designed to turn the passing particles into lower energy daughter particles and a ‘sampler material’ designed to measure the energies of these daughter particles. This results in series of collimated particle deposits made in the calorimeters referred to as *showers*.

The quantity, S , and therefore $\mathcal{R}$, are obviously different for different materials. To ensure that the entire energy of the particles is deposited within the volume of the calorimeters, a comparative figure of merit is needed to compare the stopping powers of different materials. The same is true, of course, for λ , but this quantity can already be directly compared to other materials.

Instead of trying to characterise materials by their stopping power, one can come up with a quantity that measures the distance a particle has to travel in a material to reduce its energy by a certain fraction (to go from E to pE). Assuming the identity shown in Equation 4.2 [73, p26],

$$S[E] = \frac{E}{\Delta x}. \quad (4.2)$$

It can be shown that the partial range $\mathcal{R}_p$ of the particle is equal to

$$\begin{aligned} \mathcal{R}_p &= \int_{pE}^E \frac{\Delta x}{E} dE = \Delta x \int_{pE}^E \frac{1}{E} dE \\ &= \Delta x [\ln E]_{pE}^E \\ &= \Delta x [\ln E - \ln pE] \\ &= \Delta x \ln \left(\frac{E}{pE} \right) \\ &= \Delta x \ln \left(\frac{1}{p} \right) \end{aligned} \quad (4.3)$$

If the fraction of energy lost p is chosen to be $\frac{1}{e}$, then the partial range given by Equation 4.3 is only given by Δx . It is this length that is called the radiation length X_0 .

Hadronic showers are usually longer and wider than electromagnetic showers as there is a strong interaction component as well as an EM one in the showering process. Hence the λ of materials tends to be much larger than X_0 .

The total thickness of the calorimeter system in the barrel is $22 X_0$ (ECal) plus 10λ (HCal) and in the endcaps is $24 X_0$ and 10λ [69, p8]. Measurements have shown this is enough material to limit the energy deposits to the calorimeters and not to cause ‘punch through’ to the muon spectrometer.

4.2.3.1 The Hadronic Liquid Argon Calorimeter

In the endcap region of the detector, the HCal uses copper plate absorbers and liquid Argon (LAr) a scintillating medium⁵ as the sampler material. It consists of two independent wheels divided into two segments, for a total of four layers per end-cap. The outer wheels are constructed from 50 mm copper plates while the inner ones are constructed from copper plates half as thick. The endcap HCal covers an η range of between 1.5 and 3.2 overlapping with the Tile Extended Barrel Calorimeter and the Forward Calorimeter respectively.

4.2.3.2 The Hadronic Tile Calorimeter

The next subsystem radially inwards is the Hadronic Calorimeter. In the barrel region of the detector, the HCal has a radial thickness of 1.97 m and uses steel as the absorbing medium and scintillating tiles as the sampler material.

The Barrel HCal has two main sections made of the same materials: one radially adjacent to the central LAr calorimeter called the Tile Barrel covering an η of $|\eta| < 1.0$; and one radially adjacent to the endcap calorimeters called the Tile Extended Barrel servicing a η range of $0.8 < |\eta| < 1.7$ [69, p10].

4.2.3.3 The Forward Calorimeter

The FCal is a hybrid calorimeter that only exists in the endcap region and acts as both a Hadronic and electromagnetic calorimeter within the same volume. It allows the de-

⁵A material that produces photons that can be converted into an electric signal when a charged particle passes through.

termination of the energies of particles travelling between $3.1 < |\eta| < 4.9$ [69, p10].

Given the larger particle flux incident on the detector surface at large $|\eta|$, effects that were negligible at lower η are more of a concern. One of these is the electromagnetic radiation reflected back into the Inner Detector from the calorimeter system by colliding neutrons (also known as neutron albedo). To reduce this effect the FCal is set back with respect to the front face of the ECal by 1.2 m. Since the same requirement of interaction lengths need to be made (FCal is about 10λ), this calorimeter system is more densely designed.

The FCal has three layers in each endcap. All three use liquid argon as a scintillating material. The outer two layers use tungsten as an absorbing layer to primarily target the hadronic interactions and the inner layer uses copper layers to tune showering to that of electromagnetic interactions.

4.2.3.4 The Electromagnetic Calorimeter

Like the HCal, the ECal is split into barrel and endcap sections but unlike its hadronic counterpart, all parts of the ECal are made entirely of the same materials: lead absorbing plates and liquid argon (LAr) as a sampling material. Because of this, the ECal is often referred to as the LAr Calorimeter.

The Barrel ECal ($|\eta| < 1.475$) is made up of three layers as labelled in Figure 4.7. Layer 1 has the highest granularity and is designed to discriminate between different shower seeds that are close together. Layer 2 is the radially largest section and should contain most of the energy for the majority of electromagnetic showers. Layer 3 is designed as redundancy to help absorb energies of higher energy EM showers.

Like the HCal, the end cap ECal ($1.375 < |\eta| < 3.2$) is made up of two sets of two




Figure 4.7: Sketch of the EM Barrel LAr Calorimeter. Image from [74].

coaxial wheels in each endcap section. They are situated directly in front of the endcap hadronic calorimeter wheels. In the endcap region devoted to precision physics ($|\eta| < 2.5$) (we will see more of this later), the EM calorimeter is segmented in a similar fashion to the barrel with three sections in depth. For the rest of the endcap volume, the calorimeter is segmented in two sections in depth and has a coarser granularity in the η - ϕ plane than the rest of the acceptance [69, p9].

For optimum efficiency, the ECal needs to be operated at a cold temperature and in an environment free from dust. To this end, the ECal is housed in the LAr vacuum vessel, and is cooled in the LAr cryostat.

4.2.4 Magnet system

The Magnet System comes in two parts: a solenoid that surrounds the Inner Detector, and as previously mentioned, a variable field toroid that lies outwith the calorimeters

to provide a magnetic field for the muon spectrometer. The magnetic field from these magnets bends passing charged particles such that information about their charge and momentum can be obtained from their tracks. An image showing the layout of the Magnet system with respect to the ATLAS detector can be seen in Figure 4.8.




Figure 4.8: Schematic showing the Magnet System at ATLAS. Image from [54].

4.2.4.1 Toroid magnets

Like its circumferent muon spectrometer subsystem, the toroidal magnet consists of a barrel and two endcap sections. The air-core toroidal magnetic field is generated by eight rounded rectangular coils for each section that are arranged symmetrically about the central region of the detector [75, p1].

The two end-cap toroids are inserted in the barrel toroid at each end to line up with the central solenoid. Each of the eight endcap coils are shifted by 22.5° with respect to the

barrel toroid coils to optimise the bending power of the system in the transition region [76]. The bending power is given by the field integral shown in Equation 4.4,

$$\mathcal{I} = \int_{L_1}^{L_2} B_{\perp} dl, \quad (4.4)$$

where $B_{\perp}$ is the component of the magnetic field perpendicular to the chosen integral path [77, p22] which is the path of a particle with infinite energy (unbent by the magnetic field) through the planes of the inner and outer muon chambers. This equation is further complicated by the misalignment of the magnetic coils and the presence of other magnetic material within the detector volume, like the iron present in the hadronic calorimeter. Because of this the bending power will vary as a function of $|\eta|$. Between the pseudo-rapidity range $0 < |\eta| < 1.4$, the barrel toroid provides 1.5 to 5.5 Tm of bending power. In the transition regions where the two magnets overlap ($1.4 < |\eta| < 1.6$), and in the region ($1.6 < |\eta| < 2.7$), the end-cap toroids approximately 1 to 7.5 Tm [69, p12]. This is shown pictorially in Figure 4.9.

4.2.4.2 The central solenoid

Situated between the Electromagnetic Calorimeter and the Inner Tracker, the central solenoid, constructed from an alloy of niobium and titanium, is designed to provide a 2 T axial field to the Inner Tracker [69, p20]. To ensure that desired calorimeter performance is reached, interactions between the traversing particles and the materials within the detector⁶ must be kept to a minimum. For the central solenoid, this means making the coil as thin as possible and optimising the layout of the cooling and vacuum systems to reduce the material thickness in front of the calorimeters. To this end the central solenoid is housed within the LAr cryostat, and the LAr vacuum vessel is combined with that of the solenoid, thus necessitating two fewer walls to be constructed [76, p30].

⁶resulting in significant energy losses before the particles reach the calorimeters.




Figure 4.9: Plot showing the bending power as a function of $|\eta|$. Image from [75, p15].

4.2.5 The Inner Tracker

The Inner Detector (ID) is the part of the detector closest to the beamline. Similarly to the Muon Spectrometer, a charged particle that passes through the ID will leave a series of signals that when fitted with algorithms form tracks to give the trajectory of travel of a particular particle. Every 25 ns some tens of hundreds of particles are incident on this detector volume. Within the magnetic field of the central solenoid, tracks from particles with higher energies and therefore greater momenta will be less curved than those of lower energies.

To provide the resolution required on these tracks and their origins, the detector needs

to be of a high granularity. It fulfils these requirements by the use of three detector technologies: a ‘straw-tube’ gaseous tracker called the Transition Radiation Tracker (TRT) , silicon micro-strip semiconducting sensors that make up the Semiconductor Tracker (SCT), and the Pixel Tracker which is made up of semiconducting silicon pixel sensors.

Like the other sub-systems described earlier on in this chapter, the ID has a barrel and endcap region. This can be seen depicted in Figure 4.10. The η region devoted to precision physics ($|\eta| < 2.5$) is covered by the SCT and the Pixel detectors.




Figure 4.10: Schematic showing a close-up cross-section of the Barrel and Endcap Inner Tracker System at ATLAS. Image from [78].

Figure 4.11 shows a detailed look at the barrel region of the Inner Detector, blowing up each of the layers of the sub-detectors and displaying their distances from the beam-line. Going radially inwards, the barrel region layers of the inner detector consist of 73 straw-tube gaseous tracker planes of the TRT followed by 4 layers of silicon micro-strip

sensors that make up the SCT, which in turn surround 4 layers of high granularity silicon pixel sensors. The endcap region of the ID, depicted in Figure 4.12, has 160 TRT planes, 9 disks of silicon sensors followed by 3 disks of pixel sensors.




Figure 4.11: Schematic showing a blown up cross-section of the Barrel Inner Tracker System at ATLAS, and the distance each layer of the sub-system is from the beamline. Image from [78].

4.2.5.1 Transition Radiation Tracker

The Transition Radiation Tracker or TRT covers an η range of $|\eta| < 2$ and is made up of arranged units called straws. These straws work in the same way as the ones in the




Figure 4.12: Schematic showing a blown up cross-section of the Endcap Inner Tracker System at ATLAS, and the distance each layer of the sub-system is from the beamline. The IBL Image from [69, p56].

muon spectrometer, but each straw is made from carbon-fibre reinforced Kapton⁷ and is 4 mm thick [80, p1]. In the centre of this straw is a gold-plated tungsten wire, and the gas used to fill the tubes is composed of 70% xenon, 27% CO₂ and 3% oxygen.

The barrel tubes are positioned parallel to the beam axis and are 144 cm long. The central wire inside is split at $\eta = 0$, hence only r - ϕ information can be extracted from tracks as the information about the longitudinal component is lost. The tubes in the endcap are 37 cm long and arranged radially in the wheels in 160 planes.

In addition to the straws, the spaces between the straws in the barrel (endcap) are filled with polymer fibres (foils) to create a medium with a different refractive index. When a charged particle passes between two homogeneous media with different refractive indices, it will emit a low energy (or soft) x-ray photon. Known as *transition radiation*, these soft x-rays are absorbed by the xenon in the tubes exciting the gas and causing larger energies to be deposited along the central wire. This leads to higher energy read-outs. The intensity of this kind of particle radiation scales with the relativistic factor

⁷Known for its good electrical insulation, temperature resistance and robustness to radiation damage, but is usually a film so requires another material to maintain rigidity for larger structures [79].

$\gamma = \frac{E}{m}$ of the particle [81, p1-3]. Hence this is predominantly a feature of passing electrons due to their low mass, and hence the presence of transition radiation can be used to aid in particle identification.

Due to the large area taken up by a single straw, and those straws only able to give out one reading, the spatial resolution per point is lower than the silicon sensors closer to the beamline, however this is mitigated by the average number of hits per track (36), and the longer track lengths measured. Because of this tracks with TRT deposits have better momentum measurements.

4.2.5.2 Semiconductor Tracker

From this point going radially inwards, the tracking volume is made up of semiconducting silicon sensors. These gossamer detecting sheets are used to provide radiation-resistant position measurements without adding too much material budget to the detector, which would reduce the accuracy of particle calorimetry that takes place at larger radii.

A semiconducting sensor works in a similar manner to the ionisation chambers mentioned in earlier sections. A passing electrically-charged particle will excite atomic electrons in the medium it passes through to make a charge pair. One set of this pair (usually the more mobile electrons) is accelerated to an electrode, amplified and turned into an electric signal. The difference here is that instead of collecting charge from the ionisation of gaseous atoms, the passing of a charged particle through or nearby the active silicon (whether it be a strip or a pixel) induces an excitation of electrons from the valence band to the conductive band of a solid lattice, and the electrodes that collect charge are segmented to provide more accurate position measurements [82, p9]. The semiconductor sensors on the ATLAS experiment are designed in a similar fashion to that seen in Figure 4.13.




Figure 4.13: Cross-section of a typical semiconductor module. Image from [83].

The problem with silicon is that the resistivity ρ ($= R_l^A$) of pure silicon is low [82, p12]. This means the power dissipated to the sensor when collecting charges is extremely high⁸. The electrical properties of the silicon lattice can be altered by inserting different elements (impurities) into the lattice. This is called *doping*.

A silicon lattice has four neighbours, and in pure silicon, an electron is shared among each one of them. Doping atoms that are in Group-3 of the periodic table (they have three valence electrons) would mean that only three bonds are formed with 3 nearest neighbours, leaving the introduced impurity with one fewer pair. This leaves a 'hole' in the lattice that would attract electrons from neighbouring atoms causing this hole to appear to move. This type of doping is called *n-type doping*. Conversely *p-type doping* is when atoms that are in Group-5 are introduced into the lattice causing an additional free electron to move around⁹.

⁸Pure silicon has a resistivity of $10^4 \Omega\text{cm}$. So with the some trial sensor dimensions (thickness $300\mu\text{m}$ and area $= 2\text{cm}^2$) would get a resistance of 600Ω . Using a field of 30V to accelerate electrons will induce a current of 200mA and a power dissipation into the sensor of 12W !

⁹This 'spare' electron is still bound electrically to the atoms and does not have the energy to escape to the conduction band, so sits between the valence and conduction bands

Now that silicon with the desired conductivity can be created by alternating the concentration of dopants in the silicon substrate, the issue becomes how one removes the residually-bound electrons and holes. By introducing the n-type substrate with p-type detecting strips embedded, and applying an electric field, the electrons from the n-type silicon will drift to the p-type and fill the holes. This is referred to as ‘depletion’. Upon full depletion, the silicon is said to be ‘active’. The field used to deplete the silicon is the same one used to accelerate charges that enter the conductance band [84, p21].

The external electric field across the sensor accelerates these charged particles to a collection point which in this case is a strip, creating a current. A silicon sensor in the SCT will typically have 768 strips. To read out the current and gain information about the position of the charge, each strip needs to be dealt with separately. Hence each strip is connected to an input channel on an readout chip called a ABCD3TA chip. This current is amplified in front-end electronics of the ABCD3TAs and turned into a binary readout signal, and transmitted into software logic which will store the component that read out the current, and the time window this occurred in. This is a ‘silicon hit’.

Since these ABCD3TA’s have 128 channels, six chips are needed to read out one sensor. As the strips are 12 cm long, the position resolution can be increased by placing two modules back to back with a relative rotation of 40mrad with respect to each other [85, p4]. Six ABCD3TA’s are mounted on a kapton circuit to form a *hybrid*. The hybrids are attached to two sensors via wire bonds to form a *module*, as shown in Figure 4.14.

In the barrel of the SCT, 12 rectangular modules¹⁰ are mounted onto a *ladder* to provide additional support, rigidity and common cooling and optical services [86, p51]. Going inwards from the outermost micro-strip sensor layers, the four concentric cylindrical layers in the barrel are formed from 56, 48, 40 and 32 full length ladders respectively, which leads to a total of 2112 modules (4224 sensors). In each endcap of the SCT, there

¹⁰made from two square ones ‘daisy-chained’ together




Figure 4.14: Picture of a SCT barrel module. The middle electronics board is placed over the junction of two sensors. Image from [69, p65].

are 9 rings of 52 trapezoidal-shaped modules for a total of 936 modules (1872 sensors).

One of the limiting factors in the efficiency of data collection on ATLAS is the architecture of the read-out chips. This determines how fast data is read out and how long information can be stored before a decision on whether to keep or discard it has to be made. To cope with the increase of particle flux expected in the upgrade to the HL-LHC, the ABCD3TA chips are going to be replaced. Some of the work I did in this regard is discussed in Section 4.5.

4.2.5.3 Pixel System

The Pixel sub-system is the closest sub-system to the beamline. As the sub-detector with the highest granularity, the function of the Pixel System is to provide highly accurate position resolution on passing tracks. Its proximity to the beamline means that the position it records are usually the first ones in a track. These are the most important hits

when trying to find the origin of various particles also known as the *production vertex*. The process of identifying this production vertex is called *vertexing*.

To increase the granularity, the pixel sensors are different to the SCT sensors in two ways. Firstly the active silicon contact is a small rectangle instead of a long strip, and said rectangle is read out by electronics bump-bonded directly to the surface of the sensor rather than at the end of the row to form ‘planar’ modules.

The basic unit of this detector is called a ‘pixel’. These pixels encompass an area in $\Delta\phi - z$ of 50×400 square micrometres. Each pixel module in this sub-detector has the same external dimensions ($19 \text{ mm} \times 63 \text{ mm}$) and is arranged in 18 columns of 164 pixels, but spatial requirements mean that the four pixels closest to the read-out are effectively merged. Each array is read out by a front-end electronics chip with 2880 channels, and there are 16 arrays per module, providing 47232 pixels but only 46080 readout channels per module [69, p57].

Like the SCT in the barrel region, the pixel modules are mounted onto staves for common support and cooling. There are 13 modules on a pixel staff, and each of the layers going inwards (commonly referred to Layer-2, Layer-1 and the B-Layer) contains 676, 494 and 286 staves respectively. Each endcap disk houses the modules in 8 sectors with each sector containing 6 modules, for a total of 144 modules in each endcap and 1744 pixel modules in total [69, p60]. The arrangement of these modules can be seen in Figure 4.15.

The inner-most layer of pixels at a radius of 33 mm from the beam pipe is the Insertable B-Layer (IBL). Installed in Long Shutdown 1, it covers a range of $|\eta| < 3$. It provides an additional high resolution track point closer to the beamline, which increases track reconstruction and flavour tagging performance [87, p8,13]. Due to the more extreme




Figure 4.15: Cross-section of the entire pixel detector without the Inner Barrel Layer. Image from [69, p86].

conditions in terms of radiation damage this close to the beamline, the IBL has the ability to be removed from the IP without affecting the alignment of the rest of the detector, such that it can be moved into position only when data is required to be recorded.

The pixels in the IBL are in arrays of 336×80 on a module, with 20 modules per stave. The IBL consists of 14 staves. 12 of these modules are planar modules similar to the ones discussed before, but the eight end modules (4 on each side) are ‘3D’ sensors.

In a 3D sensor the electrodes used to read out the charge are immersed within the semiconductor substrate rather than being placed on the surface of the substrate, and the depletion electric field is applied parallel to the surface. While these pixel modules are more complicated to make, the distance between penetrating electrodes can be small, thus reducing the field required to deplete the volume, and the power diffused to the sensor. This means that the detectors do not need to be cooled as much. A cross section of this detector fabricated by CNM¹¹ is shown in 4.16.

¹¹Centro Nacional de Microelectronica, Barcelona (Spain).




Figure 4.16: Cross-section of a 3D Pixel fabricated at CNM. Image from [87, p18].

4.3 How ATLAS identifies particles

Combining the energy deposits in the various parts of the detector described previously allows the reconstruction and identification of particles through their characteristic signatures. Table 4.1 shows the sub-systems that are typically activated when a particle travels through the detector. Figure 4.17 shows a visual representation of the information in Table 4.1.

Table 4.1: Table showing the sub-detectors various particles deposit energy into. The crosses show cases where there could be some small non-negligible deposits. The pluses next to ticks indicate where the subsystem produces additional information for certain types of particle.

Signal Deposit	SCT/Pixel	TRT	EMCal	HCal	MS
Electron (e)	✓	✓+	✓		
Muon (μ)	✓	✓	×	×	✓
Neutrino (ν_l)					
Photon (γ)			✓		
Neutral Hadron				✓	
Charged Hadron	✓	✓	✓	✓	




Figure 4.17: Image showing the typical energy deposits of various particles traversing through a simplified cross-section of the ATLAS detector. Image from [88].

Photons, while not electrically charged, are still able to be detected at calorimeters. Because they interact via the electromagnetic interaction, energetic photons passing through matter are able to ‘convert’ into an electron-positron pairs, which will produce deposits in the EMCalo.

Since neutrinos are not able to be directly detected in the detector, their existence has to be inferred by the presence of other particles and an imbalance in the total momentum present in the r - ϕ plane transverse to the direction of the particle beams. This *missing transverse energy* or E_T^{miss} is assumed to be the energy taken away by a neutrino. It is calculated by determining the vector of momentum required for the vector sum of the

momentum in the r - ϕ plane to equal zero as there is no initial particle motion in this plane.

Elementary particles not listed in Table 4.1 are unstable. This means that they will decay into multiple particles present in the list (a W^+ boson may decay into a μ and a ν_μ) and hence can be identified by the presence of multiple signal deposits in the set of these particles and the identification of a production vertex some distance from the beamline.

If a particle decays into leptons, it is said to have decayed *leptonically*. Likewise if a particle decays into quarks, it is said to have decayed *hadronically*. Because taus are similar in mass to a proton, they are the only lepton able to decay both leptonically and hadronically [89, p1]. This means that the tau does not have a unique energy signature and can look like an electron or muon with some E_T^{miss} or several hadronic particles which means that particles that decay leptonically into taus can be misidentified. This demonstrates that some signals in the detector volume are less useful for physics analyses.

Given that the detector has some 90M channels¹², if they were all read out every 25 ns, a data rate of 2.25 exabits of data per second needs to be recorded¹³. It is currently a physical and financial impossibility to handle these levels of data flow, processing and storage, so a decision about what events are to be kept, and those that are to be discarded, needs to occur live while the data is being taken. This is the function of the ATLAS trigger system.

¹²Number of channels in various subsystems: ID = 80.4M+6.3M+0.35M, Calorimetry = 164k+92k, MS = 354k+31k+373k+318k

¹³If the download speed of the average person's internet is 60 Mbit/s, you would need 3.75^{10} computers to get all this data.

4.4 The ATLAS trigger

Every single channel in every module is connected via wires to an output chip. When those channels are activated by the collection of charge, the signals must be processed quickly such that those channels are available again to take data¹⁴. When the detector is operational under beams from the LHC, bunches of protons are collided every 25 ns at the ATLAS experiment, and this amounts to a data rate of 40MHz.

Every 25 ns, it is expected that more than one collision between protons will occur. This means that on average the detector will have to be able to handle several event signatures overlapping in the detector. The mean number of interactions per bunch crossing, $\langle\mu\rangle_{bunch}$, also referred to as *pile-up*, for each year in Run-2 is shown in Figure 4.18.

This is important as MC simulated events cannot perfectly simulate the distribution of the $\langle\mu\rangle_{bunch}$ extracted from data, or the μ -profile. When simulated data is generated, a correction between the assumed MC simulations μ -profile and the real μ -profile from the data will have to be applied.

The ATLAS Trigger and the Data Acquisition (TDAQ) [90] uses a three stage trigger to reduce the data rate: the Level 1 Trigger (L1), the Level 2 Trigger (L2), and the Event Filter (EF). The latter two are more commonly known together as the High Level Trigger (HLT). The TDAQ has a modular design with a configurable infrastructure such that it can be changed to suit the requirements of various physics analyses.

¹⁴The information that a particular channel has not been hit in a given timeframe is also important




Figure 4.18: Average number of pile-up interactions per bunch crossing for each year in Run-2. Image from [66].

4.4.1 Level 1 Trigger

The L1 trigger is a hardware trigger situated close to the beamline designed to reduce the data rate from 40 MHz to 100 kHz [90, p3]. It does this by using rough granularity to scan the triggering parts of the muon spectrometer and the calorimeter system to find a particle candidate or signature (electrons, photons, jets, E_T^{miss} , muons, hadronically decaying taus). This is called a Region of Interest (RoI) and if the information is sent to the level 2 trigger, a 'L1 Accept'.

4.4.2 High Level Trigger

The L2 and EF are software triggers designed to use a larger fraction and the full granularity of the detector with respect to the L1 trigger respectively. The L2 is designed

to analyse detector data within the RoI to see if the particle deposits correspond to a relevant signal and to reduce the event rate from 100 kHz to about 5 kHz. Events that make it past this filter are referred to as ‘L2 Accepts’. The EF will analyse the whole detector instance with L2 Accepts and with the full detector granularity and reduce the rate to $\mathcal{O}(100 \text{ Hz})$ [90, p4].

The trigger system is also more nuanced than this as it is possible to ‘steer’ the event selection. The HLT starts from the L1 Accepts and applies ‘trigger decisions’ in a series of sequential steps, each refining the existing information by acquiring additional data at higher granularity from increasingly number of ATLAS sub-systems.

A *trigger chain* is a list of required physics signatures, implemented at the HLT level. This means that multiple requirements can be combined at the L2 and EF stages such that events are selected with more specific objects in it. An example of one of these triggers used in the analyses presented in Chapters 7-9 of this thesis would be the trigger ‘HLT_xe90_mht_L1XE50’. Here the L1 Accept passed to the HLT an event with 50 GeV of E_T^{miss} and events that will have passed this trigger will have fulfilled the additional trigger requirement of having 90 GeV of E_T^{miss} .

There are also several types of data, called *data streams*, collected for ATLAS for which there are different trigger menus. The results in this thesis come from two streams: the ZeroBias data stream (which randomly triggers on events), and the ‘physicsmain’ data stream (which is the nominal stream for data that will be used in the physics analyses in Chapters 8 and 10).

4.5 Tracker upgrade

To operate in the environment of the HL-LHC, the ID is planned to be upgraded. The design and construction of such a detector takes many years and I have been involved in testing prototype chips that will replace the current ABCD3TA's. This section briefly describes the upgrade and my work.

While the trajectories of charged particles in the current inner tracker rely on hits being reconstructed in gaseous detectors and silicon semiconductor technologies, the planned upgrade aims to solely use semiconductor silicon sensor technologies. The current make up of the detector has been explored in depth in the previous section but a summary can be seen in Figure 4.19.




Figure 4.19: Schematic showing the Inner Tracker System at ATLAS. Here the TRT is in Red, the SCT is in orange and the Pixel is in yellow. Image from [54].

The Inner Detector (ID) of ATLAS was designed to operate for a decade at an incident luminosity of $1 \times 10^{34} \text{ cm}^{-2} \text{ s}^{-1}$ and have an expected pile-up of 23 events every 25 ns bunch crossing [91, p57]. The proposed LHC upgrade (see Section 3.2.4) means that the ID would have to cope with a pile-up of about 200, which means that data acquisition systems have to deal with more raw collision data. As well as requiring larger band-

widths to deal with the data increases, the increased amount of particles in the detector mean that the detector experiences higher levels of radiation damage. In addition to this, if the experiment wants to maintain its track resolution, it also has to increase in granularity which means smaller strips/pixels are needed.

It is not possible for the ATLAS machine with its current specification to maintain its efficiency over time (especially near the beampipe) under these conditions and requirements on its performance. The solution is to replace the ID in its entirety with a similar detector called the Inner Tracker or ITk, but this requires research and development into radiation-harder technologies.

4.5.1 Layout

In the barrel region of the new detector there will be [92];

- 5 cylindrical layers of pixel sensors that will extend out to larger radii, followed by
- 2 cylindrical layers of micro-strip sensors with small strips (23.8 mm) and
- 2 cylindrical layers of micro-strip sensors with long strips (47.8 mm) to replace the current TRT.

The semiconductor microstrip sensors proposed to be used in the upgrade are a p-type bulk with 1280 n-type strips covering their 95.14 cm^2 surface [91, p85-86]. The sensor substrate is made from $300 \mu\text{m}$ thick silicon. The most probable number of electron-hole pairs released when the type of charged particle produced at ATLAS passes through a thin layer of matter can be estimated using a Landau distribution [93]. For $300 \mu\text{m}$ of silicon, this number is 22,000, which is equivalent to 3.52 fC of charge. The threshold charge traditionally required for a hit to be registered is set at 1.5 fC.

Once the signal reaches the chip it is then amplified to the order of a few milli-volts. After binary discrimination the signal is sent to the memory, awaiting a decision from the triggering process and opening up the channel to take reliable data again.

The current ABCD3TA chips have a smallest-feature¹⁵ size of 250 nm and use a binary readout architecture [94, p1844 (p2)] to read the 128 silicon strip sensors. If after the length of a readout bin (25 ns) the channel has received a signal that is larger than a given threshold set inside the chip beforehand, it outputs a “1”, if not, then it outputs a “0”. The advantage of binary chips over analogue ones is that there is less information to read out and process which is an advantage in the high-information generating environment that is the LHC.

4.5.2 ABC130 chips

The ABC(N)130 chips are the latest in the ‘ABCN series’ [95, p11-12]. These are Application-Specific Integrated Circuits (or ASIC’s) designed to replace the ABCD3TA chips. The first iteration of chip was the ABCN250 (or ABCN-25 in some literature). It shared the same lithographic size of 250 nm as the ABCD3TA chip and had a similar architecture and internal structure but with an additional power-management block [96, p374 (p2)].

The ABC130 is the intermediary chip between the ABCN250 and the ABC production chip, which is the chip currently planned to be used in the ATLAS detector during the HL-LHC programme. Both the ABC130 and the ABC have a 130 nm smallest feature size to the ABCN250’s 250 nm one. This means that the ABC130 can have twice as many channels as the ABCN250, which in turn means that half as many chips can be used to read out the same number of strips.

Smaller feature sizes also require less power to operate. Lower operating power means

¹⁵The size of any one element on a chip, such as a transistor, or block. It is also called the lithography.

less material to cool and less non-tracking infrastructure such as cables and pipes is needed nearby to keep the chips functioning. This reduces the chances of background scattering in the detector. Due to smaller lithographic feature sizes of the ABCN series, the proposed chips will also be more radiation hard and operate at the necessary bandwidth suitable for operation in the HL-LHC.

Both of these factors reduce the amount of non-detecting material inside the detector which has all the advantages mentioned above. However, the ABC130 is a new prototype, and there are not many that are stable enough to undergo testing. Hence much work is needed to understand the capabilities of the ABC130 chips to the same standard as the ABCN250.

4.5.3 Data acquisition

ITSDAQ is the Data AcQuisition program used to characterise ASICs, hybrids and modules for the ATLAS Strips ITk and is based on the SCTDAQ used to read out the current SCT. ITSDAQ utilises both C++ and python source code, and is open for modification to allow for improvements to testing parameters and procedures. In order to centralise the testing of wafers, the ITSDAQ code is centrally managed by the Rutherford Appleton Lab.

For the previous ABCN250 chip sets, a FUSEid feature¹⁶[97] was implemented, allowing at the wafer probing stage a unique ID to be assigned to all ASICs. This ID was then electrically read out at the module test stage allowing the ASIC wafer test results to be paired to the modules test results for subsequent storage in the Production Database. However, for the ABC130 chip set, there is no such FUSEid feature.

¹⁶Essentially it is a series of electronic fuses that can be selectively blown to modify the chip architecture while the chip is in operation

4.5.4 ABC130 identification

My contribution to the testing infrastructure was to automate the process of assigning identities to ABC130 chips sent to the lab at Glasgow in the absence of a FUSEid marking.

While each chip is made to the same specifications, they have unique responses when individually tested using the ITSDAQ. Assuming that the chips are not damaged in transit, and these unique responses are recorded, the initial tests done on the chip should be reproduced in the lab. This can be used to identify the chip.

The test set-up at the Glasgow Detector Development Group consists of a 100-channel mini-sensor attached to a single ABC130 chip mounted onto a hybrid. This is shown in Figure 4.20.

The production round of ABC130 chips being tested came from 17 wafers. These eight-inch wafers each have 143 ASICs and are then tested on a wafer machine. This set of tests that the chips undergo carry information about the general status of the chip and parameters that optimise the chips usage (each wafer will have this information in a separate file). The method I investigated compared these ASIC wafer tests to the ITSDAQ module tests in the lab to try and identify the chip from the wafers.

These wafer tests are referred to as threshold scans and are designed to test the characteristics of the front end of the chips. The threshold scan injects a known charge into each of the channels of the chips. For a given threshold, the channels occupancy (response of '1' or '0') is sampled 500 times. The resulting curve of occupancy against threshold is called an s-curve. Threshold scans performed with different charges injected at different times produce thousands of numbers to characterise a chip.




Figure 4.20: The RAL_DLLH3 set-up at the University of Glasgow. The test hybrid set-up has slots for 10 ABC130 chips. Only one of them is used, and that chip has a mini-sensor put across some of the pins to add a capacitance load across the chip. The buffer board ensures that the correct voltages are being applied to the chip from the power supply.

The centralised ITSDAQ code was not designed to be run multiple times to check the identities of chips. I modified the base ITSDAQ code to make it more efficient when looping through information from multiple wafers, and added an option to toggle the chip identification (chipID) on or off. These changes were successfully run on the ITSDAQ framework at Glasgow.

Each wafer is tested sequentially such that the best candidate chip on each wafer is found. The chipID uses a sum-of-squares minimisation (SOSmin) method of my own design to compare the numbers from the threshold scan of the ASIC on the module to the scans made on all the ASICs on each of the 17 wafers. Of the 17 candidates, the one

with the lowest SOSmin score is determined to be the overall best match.

The SOSmin method squares the differences between the numbers obtained in the lab ITSDAQ test and their corresponding numbers for each wafer in the file. These numbers are then added and the lowest value of these summed squares is taken as the chip candidate for that wafer.

Following this I optimised the code to reduce the computing time spent in this loop, and to reduce the number of outputs produced when this software is run. These more streamlined outputs will be made compatible with the forthcoming ITk production code-base. This code was validated and later incorporated into the general release for use in testing at all other institutes working on modules for the ITk.

4.6 How to make the collected data useful

In order to use the data collected from the sub-detectors to make predictions, it needs to be compared to what is expected from the Standard Model. To do so, the Standard Model ‘response’ needs to be known and this is created from simulated data generated via Monte Carlo processes. These are going to be discussed in the next chapter.

Monte Carlo event generation

Relating the collider data collected by the detector described in Chapter 4, to the theoretical framework described in Chapter 2, requires Standard Model processes to be generated and passed through a simulation of the detector to obtain the expected ‘response’. A ‘measurement’ of a Standard Model process is then a comparison between the collider data and the simulated expectation. Monte Carlo numerical techniques are used both in computing the matrix elements for all the relevant processes, and in simulating the detector response.

5.1 Monte Carlo integration

MC simulations are a tool used to evaluate integrals numerically. With this method, to integrate some function $p(x)$ between points a and b , it is enclosed within a ‘volume’ known to contain the entire integrable phase-space of the solution (i.e the function is finite within the chosen volume and is only discontinuous at the integration boundaries). Random points are sampled within this volume and some will fall inside the area between $p(x)$ and the x -axis (called ‘hits’), and some will fall outside (called ‘misses’). This is seen in Figure 5.1.




Figure 5.1: Diagram showing how the Monte Carlo Integration method is used to evaluate the integral of $p(x)$ between the points 'a' and 'b'. The integrating volume is given by the area between the points a and b on the x-axis extended up to the line $y = c$. Points are randomly sampled in $[x, y]$ and those that are less than $p(x)$ are called 'hits' while those that are larger than $p(x)$ are called 'misses'.

The integral of $p(x)$ is then estimated by the ratio of hits to the total number of points sampled multiplied by the integrating volume (see Equation 5.1).

$$\int_a^b p(x) dx \approx \frac{\# hits}{\# hits + \# misses} \times c(b - a). \quad (5.1)$$

The main advantage of this method is that no prior knowledge about the solution is needed. Since the function is contained within the volume, the more points sampled, the more accurate the estimate [98, p1].

In the case above the integral is simple, so the convergence ($\frac{1}{\sqrt{N}}$, where N is the total

number of sampled points) to an accurate answer is slow compared to other methods. Where this method outperforms those other methods is that this convergence is independent of dimension, so becomes the more favoured method for higher-dimensional problems [99, p6-7].

5.2 Monte Carlo simulation

The integration example has a direct correspondence to the generation of particle events. Each dimension in the integrating volume corresponds to a property or distribution of a particle (or set of particles) such that the integrating volume itself becomes the allowed phase-space of interactions.

The randomly sampled points are then ‘events’ in this space. A *sample*, then, is a file that contains information about the integrating volume and the points sampled within that volume. Each physics process (or groups of similar ones) will usually generate its own sample set. Physics analyses harvest events from samples by creating an event selection (a $p(x)$) and taking all the events that fall within it.

The entire phase-space of allowed interactions does not need to be sampled across to get accurate physics results. This is because it is known that the selections will target kinematically interesting points of the phase-space. To reduce the time to generate the samples, and the size of the samples, the points sampled can be biased by assigning event weights [100, p90-93] which gives more ‘weight’ to events generated in niche corners of the phase-space¹.

¹At some point the number of dimensions catches up to you and weighting samples is not enough. ‘Markov Chain’ MC simulations (MCMC) sample points in the phase-space to form a Markov Chain where each point generated in MCMC depends only on the previous point [101, p13-15]. This point-to-point sampling is done according to some distribution and the likelihood (read:probability) of the particle parameter set.

5.3 Generating LHC events

Protons are composite QCD objects. They are made of ‘valence’ quarks (uud), that give the proton its properties, and a slew of ‘sea’ quarks and gluons that are present due to QCD interactions between these valence quarks². Due to the omnipresent effects of the strong nuclear force, the initial hard collisions between incoming protons produce strongly-interacting sprays of particles. These sprays form large backgrounds that cover up the desired signals from interesting physics processes.

The process from high-energy hard scattering to low-energy detector observable has many stages.

Hard Scattering A high-energy parton from each of the colliding proton interacts creating new partons.

Parton Showering The high energy regime, where the principal partons from an event split and produce more quarks and gluons ($q \rightarrow qg, g \rightarrow gg, g \rightarrow q\bar{q}$) of lower energies

Hadronisation The low energy regime, where quarks bind to produce colour charge neutral states like mesons ($q\bar{q}$) or baryons (qqq)

Decay The mesons and baryons decay hadronically or leptonically into lower-energy hadrons and leptons that interact with the detector volume.

Because the decay products strongly interact with themselves at each stage, there are so many variables, energy scales, and integrable dimensions that analytically solving this problem becomes impossible. MC simulations are therefore required to predict the distributions of both the backgrounds and the signal as measured by the detector in the generation of LHC events. Figure 5.2 shows what a fully MC simulated event looks like.

²Imagine like a sack of marbles with three big marbles and lots of smaller different sized ones in it.




Figure 5.2: Image (from [102]) showing the different factorised stages in MC simulations. The first stage are the incoming hadrons, and these are represented by dark green lines with arrows pointing towards the IP. The next stage is the hard scattering process depicted in red and interactions with the rest of the proton (underlying event (UE) or secondary hard scattering events (Multiple Particle Interactions or MPI)) are shown in purple. The parton shower that follows the main hard scatter process is shown in blue. The subsequent hadronisation is represented by light green, with hadron decays and final state particles shown in dark green. The wavy yellow lines represent accompanying electromagnetic radiation from the final state particles.

In the event sequence above, the hard-scattering cross-section is the most difficult part to compute. This is because the strength of the strong coupling constant (α_s) decreases with increasing energy, and the interaction between two protons to create other partons requires the inter-meshing of low energy (short-distance) and high energy (long-distance) interactions.

The short-distance QCD interactions take place between partons inside the proton. At

this short a distance, the α_s is so large, that when the process is expanded via perturbation, the higher-order terms do not converge. This is called *non-perturbative* QCD. The long-distance QCD interactions take place between the colliding partons from each proton. At this range, α_s is small, and effects converge. This part is the *perturbative* QCD part.

The *Factorisation theorem* says that if you denote an energy scale at which effects go from being perturbative to non-perturbative, then these effects can be treated separately and independently [99, p15].

5.3.1 Parton distribution functions

The non-perturbative QCD effects that occur inside the proton are represented by what is known as the parton distribution functions or PDF's. The PDF's describe how the energy of the proton is distributed between all the valence and sea quarks that constitute it. For a given parton, i , a PDF gives the probability density $P_i(x)$ to carry a particular fraction, x , of the momentum of the proton. The interaction scale of these interactions is about 100-1000 GeV. As this is not calculable, this is parametrised from data taken at experiments. These PDF's evolve with energy, and this evolution is given by the DGLAP [103–105] equations.

5.3.2 Matrix element

The calculable part of the QCD hard-scatter process goes into the *Matrix Element*, or ME. These matrix elements are the same as the one referred to in Sections 2.2.1 and 2.2.2. Each matrix element is a sum of all of the Feynman diagrams that enter a given process. The S-Matrix, $\mathcal{M}$, which comprises these elements, is then a unitary operator that contains all the information about the scattering of the two partons. The total cross-section is proportional to the integral of $\mathcal{M}^2$.

While the energy scale of these interactions is about the TeV level, the calculation is still not analytically feasible, so the matrix element is calculated from perturbation theory in increasing powers of α_s . Since each strong vertex in a Feynman diagram contributes a power of α_s , the *leading-order* (LO) contributions are going to be from Feynman diagrams with the fewest number of vertices. The *next-to-leading-order* or NLO contributions are going to be from Feynman diagrams with the same final state but have more vertices (usually 1 or 2). The same is to be said from the *next-to-next-to-leading-order* or NNLO calculations.

The order at which this calculation is terminated is a source of error or *uncertainty* for sample generation. ‘Matrix Element corrections’ (from Chapter 8) match the hard scatter processes to possible soft and collinear radiations that come from particles in the event.

The ME PDF’s used for the MC simulations in this thesis are the ones available to all LHC experiments: NNPDF3.0NLO and NNPDF3.0NNLO³

5.3.3 Parton shower

The next step for particles that emerge after the hard-scattering process is to form high-energy splitting from the couple of particles at the TeV level to several around the 10-100 GeV level.

In this regime, each parton has a finite probability of decaying into two particles ($g \rightarrow q\bar{q}$) or creating particles out of the vacuum ($q \rightarrow gq$ or $q \rightarrow qq\bar{q}$). The parton shower generators calculate these probabilities at each energy, and approximates physical pro-

³The perturbative effects in the PDF come from fit parameters defined at scales ~ 1 GeV, so it has to be perturbatively evolved up to the 100-1000 GeV interaction scale.

cesses down to about 1 GeV.

5.3.4 Hadronisation

After the showering has finished, the parton energy scale is now low enough that quark confinement kicks in and the partons start to bind together and form hadrons. This is called hadronisation.

This low-energy regime cannot be analytically solved and can therefore only be modelled, and then fine tuned with corrections from data. For example the PYTHIA8 [106] event generator uses a Lund String model [107] to complete this step.

In the Lund String model, two quarks form a colour ‘string’ between them⁴. These strings have a tension as the strong force wants to pull the quarks together. As the quarks move apart, the energy stored in the string increases, increasing the colour field of the vacuum. At some point this high-colour field induces a $q\bar{q}$ pair to appear from the vacuum, which breaks the string into two smaller strings. The $q\bar{q}$ pairs ‘attach’ themselves to these strings and accelerate from each other before they can annihilate. These two quark-pair systems each have less energy than the initial pairing and hence the energy of the collision will continue to dissipate until the final hadrons that interact with the detector are formed. The energy scales of this interaction are calculated from the GeV to the MeV level.

5.3.5 Underlying event and multi-parton interactions

At the same time as the hard-scatter interaction is taking place, the other partons in the protons are interacting with each other. Lower-energy collisions between sea-quarks are referred to as the *underlying event*, and the higher-energy interactions between

⁴gluons are modelled as ‘kinks’ in the strings

other partons are referred to as *multi-parton interactions*.

The partons in the underlying event are no longer in colour-bound states so some hadronisation occurs here as well. The underlying event contributions can be modelled from other processes.

Parton Shower corrections are those that come from differences between shower schemes, and the reliability of underlying event and hadronisation tunes. All event generators come with default tunings but in some cases, these can be modified after fits to data. The ATLAS 2014 tuning set (A14) is based on measurement of $t\bar{t}$ and Z-boson observables [108] and the AZNLO [109] come from W-boson and γ^* momentum measurements in pp collisions.

5.3.6 Pile-up

In addition to the hard-scatter event of choice, each sample event has additional hard-scatter events from pp collisions overlaid to simulate pile-up. These are generated using PYTHIA8. These simulated events then undergo reconstruction with the same algorithms used to process data. Since the pile-up distributions between data and MC simulations are not guaranteed to be the same, the MC simulated events are given weights so that they do match. This is called a pile-up re-weighting.

5.3.7 Detector simulation

The last step in this process is to pass these MC simulations through a simulation of the detector to model the effects between the final-state particles and the detector hardware.

Since the ATLAS detector is a large and complicated machine, the full simulation of the ATLAS detector via GEANT4 [110] is usually the dominant contribution to the total

CPU time taken to generate samples for the ATLAS experiment.

5.3.8 Samples

In terms of sample generation, physicists are generally interested in the matrix elements as they contain the vertices that give rise to possible interesting physics. It is therefore useful to separate out processes that have different vertices from each other.

A sample is the name given to a group of MC simulated events that come from the same set of matrix element contributions. Events in a $t\bar{t}$ sample, for example, will have all been generated via the propagation of all the Feynman diagrams that have two top quarks (and nothing else) in the final state.

5.4 Calibration of physics objects

In Chapter 4, data collection techniques at the ATLAS detector were reviewed. In this chapter the methods used to model the data in MC simulations were discussed. Before these can be compared, the MC simulations need to be calibrated such that the objects created in MC simulations are the same as those in data.

The job of the various Combined Performance (CP) groups in the ATLAS experiment is to measure properties of the various basic objects that are collected in MC simulations and produce corrections with respect to data for various groups using said objects, usually referred to as *recommendations*. You will see an example of the kind of studies and recommendations provided from one of these CP groups from the study I conducted in the next chapter for the Tracking CP group.

Fake tracks in the ATLAS detector

Particle tracks are important for particle identification on the ATLAS experiment. It is the job of the Tracking CP group to ensure that the detector software is able to build track objects accurately. This chapter will detail the study into misidentified tracks that I performed for the Tracking CP group. Firstly I will tell you how these tracks are built.

6.1 (Re)Constructing Track Objects

When protons collide at the interaction point (IP) of the ATLAS experiment, the resulting stable particles produced traverse the ATLAS detector. As described in Chapter 4, the ATLAS detector is made up of 4 main sub-systems and the innermost one, the Inner Tracker, is responsible for turning the multitude of ‘hits’ that are generated from each passing charged particle into ‘track objects’; thus aiming to reconstruct the trajectory of a passing charged particle through the Inner Tracker.

A *hit* is the boolean output¹ of a semiconductor detector (or TRT) component that in-

¹While generally this output can be analogue or digital, In the ATLAS experiment to reduce the band-

teracts with a charged particle. Given the many layers of the Inner Detector, a passing charged particle should cause several hits to be registered.

Reconstructing a track from several of these hits at the ATLAS experiment is a multi-stage process using a series of algorithms that are contained within the *ATHENA software framework* [111, p27-32], a common tool for all analysers on the ATLAS experiment used to manipulate data taken from the ATLAS detector. These algorithms can use information from the different parts of the Inner Tracker's sub-detectors to produce different kinds of tracks. The tracks are created in two passes. The first pass is where the majority of tracks are created. Tracks are seeded from the inner silicon layers and extended outwards, in so-called *Inside-Out* track reconstruction. The second pass, *Outside-In* track reconstruction, is the reverse of this. Figure 6.1 shows an example of different types of objects used in the reconstruction of track candidates.

6.1.1 Inside-Out Track Reconstruction

In Inside-Out track reconstruction, hits in close spatial proximity to each other in the pixel and SCT sub-systems (silicon hits) are clustered into *space points* which are shown in Figure 6.1 as yellow dots. In the pixel sub-systems, a space point is a single cluster. In the SCT, where entire strips are hit, these space points come from clusters either side of a double-sided module. This allows for the point to have some information about its position along the strip direction².

Then, all combinations of three space points in the same sub-system are collected to form *track seeds* (SCT and Pixel track seeds are referred to as SSS seeds and PPP seeds respectively). As a typical charged particle will be responsible for more than 3 space points, the total possible number of track seeds will far outnumber the number of parti-

width of data that has to be dealt with in a given time-frame the electronic output is digitised.

²For the barrel the z -direction is parallel to the strip, and in the end-caps it's the radial direction

cle tracks expected. Early removal of un-physical track seeds, therefore, helps to make the tracking pattern recognition less computationally expensive.

The trajectories of the SSS seeds are then extrapolated to the beamline and the intersections are collected to try and identify their production vertex. SSS seeds that do not converge to one of a small number of vertices are usually rejected. The maximum and minimum z -positions of these vertices are then used to reject PPP seeds whose beamline extrapolations fall outside this z -region. Seeds still available at this point are said to have been *accepted*.

The next step is to turn accepted track seeds into *track candidates*. The accepted track seeds' initial parameters are then used to define a *search road* to locate possible track hits as the track is extended outwards through the magnetic field. If a space point within the seed has not been used to make another track candidate previously and sufficient hits lie on the search road, a Kalman Filter [112] is used to identify the hits compatible with the track, which then become part of the track candidate. Each seed therefore only ever results in a single track candidate. The track candidates parameters are then updated using a Runge-Kutta method [113].

In order for these track candidates to *survive*, they must pass two further requirements: to have at least seven silicon hits, and to have a p_T larger than 500 MeV. This is to reduce non-genuine tracks that will have a small number of hits and low-energy tracks which are usually not of interest to the high-energy physics programme.

Track candidates that have survived then go into the 'ambiguity solver' to further reduce non-genuine tracks. The ambiguity solver assigns a score to each track based on the number of sub-detector hits and holes³, and the χ^2 of the track fit calculated using

³If the interpolation of two space points of a track intersects an active silicon module that has not registered a hit, the intersection point is called a *hole*.




Figure 6.1: Image showing the different objects that are formed while creating tracks at the ATLAS experiment. Charged particle interactions with detector components induce electrical signals which in turn inform the position of this interaction. This information is contained within space points, shown as yellow points. These points are combined to form track ‘seeds’ which are cyan if they are formed in the Semiconductor Tracker (SSS seeds); and navy blue if they are formed in the Pixel system (PPP seeds). The momentum information from the seeds is then extended to find other points that lie on this to form silicon track candidates (shown in dashed green). The third track going clockwise shows the tracks are then extended into the TRT to see if they align with TRT deposits. Track candidates that fulfil further quality requirements become tracks, shown in red. The lighter red track has one pixel ‘hole’ and one SCT ‘hole’ so passes loose track requirements, and the darker red track has hits in every layer, so passes the tight track requirement. The two track candidates left in the diagram do not become actual tracks because they either miss the nominal interaction point when extended, or do not have any unique space points attributed to the track in the pixel system.

the residuals between the space points and the track. Track candidates are ranked by score and space points shared by multiple tracks are given to the track with the highest score. It is at this stage that track candidates can be merged if they share a lot of hits, or rejected if the score is deemed too low.

6.1.1.1 TRT Extensions

The straws that make up the TRT are also called drift tubes. Due to the number of drift tubes that are present in the detector, charged particles that have large relativistic factors travelling through the TRT can sometimes cause a lot of hits. These hits form a large number of points on the arc of a circle, and therefore this TRT track object is known as a drift circle. The silicon-based tracks are extended to the TRT, and any drift circles that are consistent with the track are added to it.

6.1.2 Outside-In Track Reconstruction

Outside-In track reconstruction works in the same way as Inside-Out track construction, but as a starting point takes unassigned TRT drift circles and tries to extend them inwards to match them to silicon hits that are currently unclaimed. Fewer than about 2% of accepted tracks are made this way.

6.1.3 Track Contents

As mentioned in Section 6.1.1, the track parameters are determined by the track seeds and are updated using a Runge-Kutta method as additional points are added to the track. Owing to the constant magnetic field provided by the surrounding 2T solenoid magnet, all charged particles that pass through will travel on a helical path through the detector. This helical path can be parametrised by 5 variables determined at the shortest spatial distance the track is from the beamline axis and with respect to a surface perpendicular to the direction of the beam at a particular point, called a perigee surface. These track parameters are as follows: the distance of closest approach to the

beamline or the transverse impact parameter, d_0 ; the distance along the axis of the beam (longitude) that separates d_0 and the perigee surface (the longitudinal impact parameter), z_0 ; the track curvature parameter, $\frac{q}{|\vec{p}|}$, where q is the charge of the particle and $\vec{p}$ is the momentum of the particle at the spacial point $[d_0, z_0]$; the angle subtended by the momentum of the particle in the (x, z) direction at $[d_0, z_0]$ and the longitudinal axis (the polar angle), θ_0 ; and the angle subtended by the momentum of the particle in the (x, y) direction at $[d_0, z_0]$ (also known as the transverse momentum $\vec{p}_T$ or p_T) and the x -axis (the azimuthal angle) ϕ_0 . A pictorial representation of these parameters can be found in Figure 6.2.




Figure 6.2: Pictorial definition of the parameters that make up a track at the ATLAS detector. All of the parameters are defined with respect to a perigee surface, shown here in blue. The transverse momentum ($\vec{p}_T$) is the 2-D projection of the total momentum of the track ($\vec{p}$) at the distance of closest approach onto the perigee surface. The B subscript on the axis directions refers to the beam, which travels in the z -direction.

This selection of parameters are chosen as they are a representation of the track state at its theoretical origin and is thus generally the closest approximation to the description of the track near its production vertex.

6.2 Origin of Fake Tracks

To increase the number of collisions, protons are collided in bunches. This results in multiple interactions per bunch-crossing or ‘pile-up’ (see Section 4.4) and means more hits are registered in the inner detector in a given time frame. The resulting jump in the complexity of the combinatorics makes it harder for these algorithms to correctly create tracks that correspond to the trajectory of the charged particles as they pass through the detector. Combinations of hits that are from different tracks and look like hits from actual tracks, are referred to as *fake tracks* [114]. While the process above is good at mitigating the levels of fake tracks that are classified as genuine, a fair few fake tracks are still labelled as real.

Since the reconstruction of fake tracks is inevitable given the conditions of the proton collisions, being able to model them correctly is essential to reducing the uncertainties that they cause. This, combined with the fact that the luminosity increase scheduled to take place for HL-LHC will only increase the pile-up further (to $\langle\mu\rangle_{\text{bunch}} = 200$), means there is a strong motivation for tackling and calibrating such effects now while the pile-up is relatively low.

6.3 Modelling of Fake Tracks

To be able to model the number of combinatorial fake tracks as a function of pile-up, a baseline model of the number of genuine (i.e non-fake) tracks is required. To do this an environment free from processes/decisions that inflate or bias the number of tracks is required. Real physics data from the ATLAS experiment (using the ‘physics-main’ data

stream) is biased towards events with lots of tracks in them. The triggers used to select events prefer those that contain high- p_T charged objects that will leave good-quality tracks in the detector. Hence the events obtained within this trigger are more likely to have a higher number of tracks. To remove this physics bias, data from a different data stream called 'ZeroBias' data is used. The triggers for this stream are randomly seeded there are no requirements on the tracks created.

For the data from the ZeroBias data stream, the assumption is made that the collider conditions are stable for long enough, since the particles that collide at the IP do not change. This means that the average number of genuine tracks is the same for each pile up interaction. A first-order approximation can therefore be made that the average number of genuine tracks is directly proportional to the number of pile-up interactions. Then any deviation from the linear trend of average number of tracks with pile up is assumed to be solely attributed to fake tracks.

To generate MC simulated events that are akin to this data selection, a base event is selected, and pile-up processes are added. Since a base event that has some tracks will artificially increase the number of tracks, the event used as a base is a single neutrino event.

The contribution of fake tracks to the average number of tracks also depends on the kind of tracks created. Tracking algorithms can create several types of tracks (that all meet the baseline requirements) by using different sets of requirements to cluster hits. These sets of requirements are called selections, and there are two main types of track selection: Loose and TightPrimary.

The Loose and TightPrimary track selections are only applied to tracks that lie within the barrel region of the ATLAS detector and pass a momentum cut of 400 MeV. These

track selections also stipulate requirements on how many total hits are needed to be registered in both of Pixel and SCT sub-systems ('silicon hits') as well as them separately. The Loose selection, for example, requires at least 7 silicon hits but the TightPrimary selection requires at least 9 hits, and one of them must be in the two innermost pixel layers [115]. More details about the differences between the track selections can be seen in Table 6.1. Given that the the TightPrimary selection has more requirements on the tracks than the Loose selection, it is expected that there is a lower fraction of fake tracks that pass the selection.

Table 6.1: Table showing the difference in selection requirement made on Loose tracks and TightPrimary tracks. Modules shared correspond to tracks made from hits that are traced to the same module.

Selection Criteria	Loose	TightPrimary
p_T Cut	$p_T > 0.4 \text{ GeV}$	$p_T > 0.4 \text{ GeV}$
η Range	$ \eta < 2.5$	$ \eta < 2.5$
Minimum Number of Silicon Hits (MNSH)	≤ 7	≤ 9 (≤ 11 if $ \eta \leq 1.65$)
Required Pixel Hits	0	1 (in IBL or B-Layer)
Maximum Number of Silicon Modules Shared	1	0
Maximum Number of Silicon Holes	2	2
Maximum Number of Pixel Holes	1	0

Figure 6.3 shows the average number of tracks against the average pile up per bunch crossing, $\langle \mu \rangle_{\text{bunch}}$ (for the sake of brevity this will be simply referred to as μ from now on) for both the Loose and TightPrimary selections between a range of $6 \leq \mu \leq 60$ in the ZeroBias data stream. The distance between the points and the fitted line indicate possible fake tracks. A deviation from the expected linear trend was observed at high μ . For the TightPrimary selection, the points sit close to the linear fit indicating low levels of fake tracks. The parameters that define the linear trend are seeded from the region $9 \leq \mu \leq 16$ where it is expected that the number of fake tracks is negligible. It is thought that effects other than fake tracks could contribute to a deviation from linearity but were expected to be sub-dominant and hence the first iteration of the study conducted included none of those possible effects.




Figure 6.3: Plot of the average number of tracks as a function of pile-up for the Loose and TightPrimary tracks in Run 2 data. Here the p_T cut-off at 1 GeV rather than 400 MeV reflects the fact that the majority of analyses that use tracks do not find tracks of that low energy to be useful. The black and red lines are linear fits made between $9 \leq \langle\mu\rangle_{\text{bunch}} \leq 16$ and extended out to higher $\langle\mu\rangle_{\text{bunch}}$.

6.4 Mis-modelling Recommendations in 2017

The first iteration of the study aimed to replicate the efforts made by previous analysers and used their methods to create recommendations on a newer dataset.

The raw number of fake tracks present in each tracking selection, while important, is not the main figure of merit for this study. Since most analyses use MC simulated events as a basis to test the Standard Model by comparing it to data, understanding the differences between the reconstruction of track objects in MC simulated events and data is important. Ideally a track made from MC simulated events and a track made from real data should have the same definition. If there are differences, the Tracking CP group should come up with a correction to be applied such that the overall number of tracks is consistent between the two.

The aim of the recommendations, therefore, is to provide a measure of the uncertainty on the number of MC simulated tracks present in each selection. Figure 6.4 shows the average number of tracks versus μ for the study in 2017 which only uses 2016 data. Since the runs selected to produce the points in Figure 6.4 had small amounts of data at high values of μ , the data points above $\mu = 45$ were deemed unreliable and the trends were capped at this value.




Figure 6.4: Plot of the average number of tracks as a function of pile-up for the Loose (in black circles) and TightPrimary (in magenta squares) tracks in Run 2 2016 data (filled) compared to the equivalent MC simulations events (empty). The solid lines represent extrapolations of a linear fit to MC simulations seeded between $9 \leq \mu \leq 16$.

6.4.1 Extracting Uncertainties

The study of fake tracks has two main purposes.

- To check the number of fakes in generated MC simulations matches that of the data.
- To try to minimise the number of fake tracks in the track selections.

For both of these aims, a better quantity to look at is the linear fit deviation which is given by

$$\frac{\langle N_{Tracks} \rangle - f(\mu)}{\langle N_{Tracks} \rangle}, \begin{cases} Loose & f(\mu) = m_L \mu + c_L \\ TightPrimary & f(\mu) = m_{TP} \mu + c_{TP} \end{cases}$$

where $f(\mu)$ is a linear fit seeded in the region of low μ where it is assumed there are negligible amounts of fake tracks, and extrapolated out to high μ . Figure 6.5 shows the difference between the number of tracks and the seeded linear fit seen in Figure 6.4 at each value of μ and the ratio between the average number of events given by data to that of generated MC. It is these residuals from the line Data/MC = 1 line that will provide uncertainties on the fake rate.




Figure 6.5: Plot of the average number of fake tracks as a function of pile-up for the Loose (in black circles) and TightPrimary (in magenta squares) tracks in Run-2 (2016) data (filled) compared to the equivalent MC simulations events (empty). The ‘fake rate’ is the difference between the points and the line of the respective colour in Figure 6.4 divided by the total number of tracks for that selection.

The ratio of the fake rate in data and MC at a given value of μ is calculated across the entire range of $0 \leq \mu \leq 45$. Then for each track selection, the residual of the ratio from MC/Data = 1 line is taken for each value of μ . These points are then put into a

histogram and then a Gaussian is fitted to calculate the un-weighted mean ($\frac{\sum_{i=1}^N x_i}{N}$) and RMS ($\sqrt{\frac{\sum_{i=1}^N (x_i)^2}{N}}$) of these residuals for all points with a $\mu > 20$. Since there are 40 points entering each histogram and they are unweighted, outliers can have a large effect on the resulting distribution used to make recommendations.

Outliers occur in the distribution of residuals for two main reasons. The first is that at a particular value of μ there is a very large discrepancy between the modelling of tracks in data and MC simulated events that is not present at surrounding values of μ . This is not expected to happen as it implies a spontaneous mismodelling effect independent on μ , not seen for any other values of μ . The second reason is that you divide two numbers that are close to zero by each other, and obtain large values (seen in Figure 6.5). This is not a characteristic of mismodelling per-se, but is an indicator of statistical effects which are largely irrelevant. Hence it is assumed that any outliers seen can be removed.

To remove the outliers, all points that are greater than two RMS deviations from the mean are removed and the mean and RMS are recalculated. This ‘trimming’ is done twice. The final mean and RMS of the distributions of trimmed residuals provide the basis of the recommendations.

The histograms generated on the fake study framework for the trimmed ratio plot residuals are consulted. Depending on the distribution of trimmed residuals one of either the mean or RMS is used to characterise this distribution. Since most of the mean and RMS outputs are between -1 and 1, these can be easily converted into percentages, which will be the recommendation. The criteria for selecting the mean or RMS is outlined below.

- If the trimmed bin-by-bin residuals are centred around 0, this implies that the MC simulations reproduce the data well overall. It is expected in this case that the

RMS will be small, but the mean will be smaller still. A reasonable conservative estimate would be to utilise the RMS of the points as the uncertainty. However, if the RMS-based uncertainty is significantly larger than expected from a 'by-eye' inspection of the residuals, then there could be a problem with the linear fit assumption and it is better to use the mean.

- If the residuals are not centred at 0, but the RMS is small, this implies that the distribution is clustered around a different value. The uncertainty should therefore be based on the deviation of the mean. In this case it would be good to consider a one-sided uncertainty, as the dominant effect seems to be the systematic shift is larger than expected. The systematic shift of the points is therefore more important to characterise than their spread, so the mean should be used, as in this case it represents the difference from zero.
- If the residuals are not based at 0, and the RMS is large, with no obvious clustering at a particular value then neither the mean or the RMS give a good method of characterising what is going on. This is an indication that the fitting procedure is not working well and an alternative measure needs to be used.

The greatest cause of systematic error in this process comes from the relatively arbitrary choice of the μ range used to seed the linear fit. To prove that this process is stable, the above steps are repeated for a variety of different μ ranges. Depending on the distributions of the final means and the RMS under different μ ranges, either the mean or the RMS will be used as the basis of the recommendation. Figure 6.6 shows some of these mean and RMS histograms and Figure 6.7 shows the presentation of the data in the format used for the recommendations.

Since the residuals for the loose selection are tightly centred on 0, the RMS spread shown in the top table of Figure 6.7 was used to generate recommendations. With a spread of between 0.49 and 0.52, an uncertainty of 50% on the statistical precision of




Figure 6.6: Histograms for the residuals of the ratio between data and MC simulations for the average number of tracks against μ . On the top these are shown for the Loose and the TightPrimary (referred to in these figures as simply 'Tight') selections and on the bottom after two rounds of outlier removal, for the TightPrimary (left) and Loose selections (right) respectively.

the number of fake tracks was suggested. For the TightPrimary as the distribution is centred about -1, the mean spreads shown in the bottom table of Figure 6.7 are used which were used to justify an uncertainty of 100%. Since the number of assumed fake tracks for the TightPrimary selection is low, it is expected that the uncertainty on any fake tracks found will be high.

6.5 Mis-modelling Recommendations in 2018

By 2018, more of the 2017 dataset had been processed, and runs with a larger average μ could be incorporated into the study and thus extend the range of the track recommendations from $\mu = 45$ to $\mu = 60$.

LOOSE RECOMMENDATION		
Linear Fit Range	Mean	RMS
6 < μ < 13	-0.216089	0.498792
7 < μ < 14	-0.233022	0.517859
8 < μ < 15	-0.246893	0.517656
10 < μ < 17	-0.251301	0.493711
11 < μ < 18	-0.257637	0.497104
12 < μ < 19	-0.230041	0.505758
9 < μ < 16	-0.229709	0.492658
TIGHTPRIMARY RECOMMENDATION		
Linear Fit Range	Mean	RMS
6 < μ < 13	-0.838512	0.874920
7 < μ < 14	-0.963887	1.008796
8 < μ < 15	-0.951618	0.990783
10 < μ < 17	-0.950041	1.005986
11 < μ < 18	-0.939612	0.993488
12 < μ < 19	-1.005434	1.277743
9 < μ < 16	-0.987692	1.059458


Figure 6.7: Final .txt output file produced by the tracking code I created which is the basis of the recommendations provided to the group. The last entry in both is the nominal $9 \leq \mu \leq 16$ linear seed, but the recommendation provided is generally taken to be the rough midpoint of all these values rounded upwards to the nearest 5%-10%.

Figure 6.8a and Figure 6.8b shows the updated plot of the number of tracks versus μ and the difference between the number of tracks and the seeded linear fit, respectively.


If the assumption that the number of fake tracks is linear in μ was correct, the points in the deviations from linearity plot seen in Figure 6.8b should also follow a linear trend. Given that the Loose selection points for both the data and the MC simulations do not follow a linear trend, this indicates a higher-order contribution to the number of fake tracks.

6.5.1 Additional Cuts on d_0

One way of producing better recommendations is to try to profile the phase space that is most likely to result in a fake track. Since tracks have several parameters that describe



(a)



(b)

Figure 6.8: Plots of the average number of tracks as a function of pile-up (μ) for the Loose and TightPrimary tracks in Run-2 (2016+2017) data for $\mu < 60$ (a), and the difference between the number of tracks and the linear fit seeded from $9 \leq \mu \leq 16$ (b).

them, the distribution of some track properties may help distinguish real tracks from fake ones. Given the definition of d_0 in Section 6.1.3, the distribution of d_0 is expected to be peaked around small values for real tracks. Since fake tracks are created from random combinations of clustered hits, it is assumed there should be no preferential distances from the beamline for fake tracks, and hence their distribution should be uniform across the entire range of d_0 . To test this hypothesis, several slices of d_0 were added to both of the track selections and the region at which the tracks start to diverge from linearity is attempted to be isolated.


Figure 6.9a shows that there are more fake tracks that deviate from the linear fit assumption at low μ , the larger the cut on d_0 applied to the track selection. Figure 6.9b shows that with an additional cut on d_0 , the deviations from the linear fit for the Loose selection are greatly reduced compared to the plot without this cut shown previously in Figure 6.4.

6.5.2 Extracting Uncertainties


Recent studies by an adjacent CP group had shown that there was a reduction in the efficiency of TightPrimary track selection with respect to μ . To improve the study and the recommendations, a μ -dependent efficiency loss term was included in the fit. Given the improvement in tracking quality that including an additional d_0 cut had on the number of tracks versus mu plot, recommendations were produced for both the selection with a d_0 cut of 5 mm and those without this cut, (called the inclusive selection, and most compatible with the previous study).

The updated linear fit deviation is given by

$$\frac{\langle N_{Tracks} \rangle - f(\mu)}{\langle N_{Tracks} \rangle}, \begin{cases} Loose & f(\mu) = m_L \mu + c_L \\ TightPrimary (2018) & f(\mu) = m_{TP} \mu * (1 - b(\mu)) + c_{TP} \end{cases}$$



(a)



(b)

Figure 6.9: (a) Number of tracks versus μ for the Loose track selection data (filled) and MC simulated events (blank) with three additional cuts on d_0 applied. In ascending order. The yellow points have a selection of $0.5\text{ mm} \leq d_0 \leq 2\text{ mm}$, the blue points have a selection of $2\text{ mm} \leq d_0 \leq 5\text{ mm}$ and the green points have a selection of $d_0 \geq 5\text{ mm}$. (b) Number of tracks versus μ for the Loose (black) and TightPrimary (red) selections with an additional $< 0.5\text{ mm}$ cut on d_0 , also referred to as innermost d_0 .

where, as in the previous iteration of the study, $f(\mu)$ is a linear fit seeded in the region of low μ where it is assumed there are negligible fake tracks, extrapolated out to high μ , and b is an efficiency gradient for the TightPrimary selection.




Figure 6.10: Histograms for the residuals of the ratio between data and MC for the average number of tracks against μ after two rounds of outlier removal. Here the TightPrimary selection is referred to simply as ‘Tight’ Top left and top right, for the inner- d_0 TightPrimary and Loose selections respectively. The track selections for their respective inclusive tracks are below them.

Using the criteria listed in the previous selection on the distribution of the data/MC residuals, Figure 6.10 shows that for the mis-modelling recommendations on both the inclusive and inner- d_0 TightPrimary track selection, the mean should be used, and for the recommendations on the inclusive and inner- d_0 Loose track selection, the RMS should be used.

From Figure 6.11, and the conclusions following the histograms in Figure 6.10, the recommendations for uncertainties on the tracking fake mis-modelling presented internally to the Tracking combined performance group were:

LOOSE RECOMMENDATION			
Linear Fit Range	Mean	RMS	
6 < μ < 13	-0.036739	0.050882	
7 < μ < 14	0.014215	0.043818	
8 < μ < 15	0.004710	0.038911	
10 < μ < 17	-0.038976	0.056869	
11 < μ < 18	-0.047892	0.065157	
12 < μ < 19	-0.060820	0.082002	
9 < μ < 16	-0.030386	0.046996	
TIGHTPRIMARY RECOMMENDATION			
Linear Fit Range	Mean	RMS	
6 < μ < 13	-0.766952	0.804112	
7 < μ < 14	-0.977999	1.020847	
8 < μ < 15	-0.959904	0.972539	
10 < μ < 17	-0.860097	0.892967	
11 < μ < 18	-0.869742	0.920246	
12 < μ < 19	-0.808628	0.941430	
9 < μ < 16	-0.905481	0.969513	

(a)

LOOSE RECOMMENDATION			
Linear Fit Range	Mean	RMS	
6 < μ < 13	0.645740	0.878288	
7 < μ < 14	0.722007	2.212326	
8 < μ < 15	0.732433	1.812710	
10 < μ < 17	0.373180	0.519231	
11 < μ < 18	0.289021	0.454227	
12 < μ < 19	0.157182	0.366509	
9 < μ < 16	0.529948	0.735999	
TIGHTPRIMARY RECOMMENDATION			
Linear Fit Range	Mean	RMS	
6 < μ < 13	-0.825191	0.894543	
7 < μ < 14	-0.899189	0.928469	
8 < μ < 15	-0.946295	0.971105	
10 < μ < 17	-0.819933	0.848748	
11 < μ < 18	-0.807826	0.844708	
12 < μ < 19	-0.719495	0.806533	
9 < μ < 16	-0.862689	0.903464	

(b)

Figure 6.11: Final .txt output file produced by the tracking code I created which is the basis of the recommendations provided to the group. (a) shows the means and RMS' for inclusive track selections and (b) shows them for the track selections with an additional requirement of a $d_0 < 0.5$ mm. The last entry in both is the nominal $9 \leq \mu \leq 16$ linear seed, but the recommendation provided is generally taken to be the rough midpoint of all these values rounded upwards to the nearest 5%.

Inclusive Loose Selection: 5%

Inclusive TightPrimary Selection: 85%

Inner- d_0 Loose Selection: 50%

Inner- d_0 TightPrimary Selection: 85%

To be conservative, these numbers were rounded up. The final uncertainty recommendations presented in the summer of 2018 for the track selections are as follows:

Inclusive Loose Selection: 10%

Inclusive TightPrimary Selection: 100%

Inner- d_0 Loose Selection: 50%

Inner- d_0 TightPrimary Selection: 100%

Later on in the year, these results helped motivate the decision to increase the silicon hits required for the Loose selection from 7 to 8 to reduce the raw number of fake tracks.

6.6 Monte Carlo Simulated Tracks

The uncertainty recommendations that are generated are based on a ratio between the fake tracks found in data and fake tracks found in MC simulations. Since it is understood that the assumption of linearity with μ for fake tracks is incorrect to some degree, improvements to the recommendations can come from having a better understanding as to what is labelled as a fake track in MC simulations. In order to do this we should follow step-by-step the processes that lead to track objects being created in MC simulated events. This can be seen below in Figure 6.12. After the tracks are obtained, to reduce the size of the files, non-interesting tracks can be removed.




Figure 6.12: Graphic showing how MC truth track and reconstructed track objects are made.

With this, it is possible to compare the performance of the tracking algorithm in MC simulations to the event record of what happened, referred to as the 'truth' or 'truth information'. After slimming, the reconstructed (reco) tracks and truth tracks are stored

in separate ‘containers’ and can be compared to try and understand how the tracking algorithm performs on MC simulated events. Comparing the trend of certain truth-level features with μ could therefore help identify some unexplained non-linearities in the general $\langle N_{Tracks} \rangle$ vs μ plots.


6.6.1 Truth Linking

One such comparison between reco tracks and truth tracks is the presence of a *truth link* in the reconstructed track container. When GEANT4 is run to make a HITS dataset, the MC simulations truth record is also augmented⁴ and links are produced between each hit and the responsible particle. Multiple particles, however, can produce the same digital signal. When you apply a digital reconstruction to go from hits to tracks, because each track has one truth link, the relationship between hit-truth and track-truth information is a many-to-one mapping. The full truth information is then thinned on the conversion to various filtered xAOD analysis formats.⁵ After the xAOD’s are formed, the production of reduced sets of data for analyses purposes means that information about the associated truth particle and the link themselves are no longer propagated, but information derived from these quantities, like truth track particle ID and whether it was a signal event or a pile-up event is propagated.


The truth link takes the form of a boolean and its presence indicates that early on in the process of producing a reco track, there was a truth track with certain properties that are exactly equal or close to the equivalent properties of the reco track. This is a property that is assigned to a reco track before the slimming of the truth events takes place. If everything was perfectly efficient, it could be assumed that to first approximation all tracks that do not have a truth link are fakes.

⁴to contain any long-lived particle decays and any conversions due to detector material interactions

⁵When a ‘link’ is made what is physically happening is that a reference to an object in another container in the same file is being stored, like a symbolic link. If this container is discarded, then the link will be broken and that information will be no longer accessible. This is done so that when building the xAOD’s there is as much information as possible.



(a)



(b)

Figure 6.13: Number of reconstructed tracks that have truth links vs μ plotted alongside the total number of tracks for the Loose selection (a) and the TightPrimary selection (b).

In both the figures 6.13a and 6.13b the trends of tracks with truth tracks (the blue lines) increases linearly with μ . This means that on its own, the truth track information is not enough to provide any insight into the non-linearity of $\langle N_{Tracks} \rangle$ such that recommendations can be altered.

6.6.2 Truth Matching Probability (P_{match})

Another such comparison between reco and truth tracks is the Truth matching probability, P_{match} . This quantity is attributed to the track after slimming takes place. It is a number between 0 and 1 that describes the likelihood that a track is a fake. P_{match} is defined by the ratio of the number of space points that are shared (or common) to both the reco and its best-matched truth track, and the number of space points that make up the reco track. Since the tracker is made up of three segments, the space points on each track will fall into one of three types: pixel clusters, SCT clusters or TRT clusters. The pixel clusters are closer to the beamline and are most important in the reconstruction of the track, followed by the SCT and then the TRT clusters. P_{match} assigns weights to clusters from each sub-detector to reflect this. Clusters from pixel are given a weight of 10, those from the SCT get a weight of 5, and those from the TRT, a weight of 1 leading to Equation 6.1.

$$P_{match} = \frac{10 \times N_{Pixel}^{Common} + 5 \times N_{SCT}^{Common} + N_{TRT}^{Common}}{10 \times N_{Pixel}^{Track} + 5 \times N_{SCT}^{Track} + N_{TRT}^{Track}}. \quad (6.1)$$

TightPrimary tracks are partially defined by a cut of $P_{match} > 0.5$, so looking at the tracks that fail this for both selections serves as a test to whether the reconstruction algorithms are behaving as expected. This can be seen in Figure 6.14.

A track having a $P_{match} < 0.5$ is a good indication that the track is a fake. There are two reasons that a track will pass this requirement. The first is that the reconstructed track has lots of hits, but the selection of these hits is unlike any truth track. The chance




Figure 6.14: Loose and TightPrimary that have $P_{\text{match}} < 0.5$. It is important to note that the failure rate for the Loose selection increases non-linearly. At high μ there are roughly 20 tracks failing this requirement. This is about 8% of the 250 tracks seen in Figure 6.13a.

of this becomes increasingly unlikely with the increased number of hits in an event, especially as hits can be shared between multiple track candidates. The second, and more likely reason, is that the reco track does not have many hits attributed to it, this is an indication of a fake track as real tracks leave lots of hits in the detector.

6.6.3 Combining Truth Linking and P_{match}

Since P_{match} and ‘Truth Linking’ happen at different stages of the data preparation process, by combining these two criteria and seeing what tracks, if any, have or do not have such information, it can be ascertained what happens during the slimming step.

Figures 6.15 and 6.16 demonstrate that some information is lost in the slimming process. The truth link is established around the same time as the P_{match} . Every particle that has a high P_{match} should have a truth link, but Figure 6.16 shows that this is not the




Figure 6.15: Image showing the percentage of tracks of a given selection without truth links but have matching probabilities of greater than 0.5.




Figure 6.16: Plot showing the number of Loose reco tracks that have $P_{\text{match}} > 0.5$ compared to one showing the number of loose reco tracks that have $P_{\text{match}} < 0.5$ and have a truth link.

case. For a reco track, the associated truth particle and the link themselves stop being stored, but the information derived from them continues to be propagated. Figure 6.15 shows that roughly 80% of reconstructed Loose tracks (which is the selection with the higher fake contribution) at low μ are best matched to a particle that is no longer stored, but given that the fake contribution is predominantly present at high μ for the Loose selection and at these values of μ it drops to 0%, the lost information is probably not due to fakes.

The best explanation of the results in figures 6.15 and 6.16 is that the truth record is cleaned after the calculation of the P_{match} . So to solve ambiguities in the MC simulated events reconstruction, one has to go into the slimming step of the track algorithm and try to correct the loss of information. This was beyond the scope of the study.

6.6.4 Track Barcodes

Different types of clustered hits are assigned *barcodes* depending on whether they are constructed from hits produced from the original particle as generated by an event generator modelling a signal process (a *primary vertex*), or if they were generated via interactions with the detector material in a GEANT4 detector simulation (a *secondary vertex*⁶).

Tracks created from hits from the original generator have low barcodes, and those generated later in GEANT4 have higher barcodes⁷. The number of primary and secondary vertices should rise linearly with μ . All real tracks therefore should have a barcode but

⁶Not to be confused with its namesake created in the event generator event record, which means that the particle produced has a measurable lifetime so will decay into other products at a different vertex to the one where it was produced. This particle and daughter hits will all be given the same type of event-generator barcodes.


⁷This is roughly correct if P_{match} is sufficiently high. If the P_{match} is low, selecting the track origin via the barcode only says that more of the hits that constitute a track come from a track that is primary or secondary than any other track. However most of those hits are still not from that track, so the track is neither primary or secondary - it is a fake.

because the track is assigned a barcode at the start of the reconstruction process, fakes can have non-zero barcodes. However, all tracks that have barcodes of 0 will be fake. Figure 6.17 shows plots of the stored barcodes plotted against μ .


Of all of the track objects found, none of them had a barcode of zero. Figures 6.17b and 6.17c have many more track objects present than any of the plots shown in this chapter. This is because not all track objects that have barcodes are ‘trackable’. Not all charged particles generated are stable to the generator, and as such there exists a status flag to show a particles generator stability. Stability is the distance measure of how long a particle exists before it decays. Often this is given to be true if the product of the particles lifetime and the speed of light⁸ is greater than 10 mm, but this is the average distance and not an actual distance a particle travels before it decays.

The generator produces a lot of particles, some of which are for book-keeping but all of them are real and physical. Their lifetime, however, means that they will not be able to leave enough hits in the detector to be reconstructed. Figure 6.18 shows the barcode information shown previously but only for particles that are stable to the generator. Figures 6.17a and 6.18a show that the tracks with low and high barcodes rise linearly with μ as expected. It is possible that the tracks that the stable tracks that share barcodes (as shown in figure 6.18b) are fakes, but these are numbers generated over several hundred events, and since the barcodes are randomly generated for each event, and it was unknown to what precision a barcode, it is not unexpected that a small number of tracks between events happen to have the same barcode. Further study had to be conducted into how and why tracks could share barcodes.


⁸which is a good enough approximation for the speed at which all particles traverse the detector



(a)



(b)



(c)

Figure 6.17: Number of reconstructed tracks vs μ for all track objects that have high barcodes ($\geq 2 \times 10^5$) (a); all track objects that have low barcodes ($0 < \text{barcode} < 2 \times 10^5$) (b); and for all track objects that share the same barcode (c). The lines plotted in these plots are lines to guide the eye, and such have little significance attached to them.




Figure 6.18: Number of stable reconstructed tracks vs μ for all track objects that have low barcodes ($0 < \text{barcode} < 2 \times 10^5$) (a); and for all track objects that share the same barcode (b).

6.6.5 Conclusions

An alternative hypothesis was suggested to explain these results. It is possible that the tracks that have been assumed to be fake because they do not follow a linear trend with μ are real tracks that at some point were processed such that they stopped fulfilling truth criteria. One example being an interaction with matter causes a pion to fragment to more pions. The daughter pions lose the truth link to the original pion, and the P_{match} ascribed to the original pion falls below 0.5 and the track is labelled as fake even though the track is a true one. If this is true then at the point of the splitting, the barcode information is also lost and a reduction of secondary tracks is seen. A counterpoint to this argument would be that these tracks would have barcodes of 0, and none were seen, but it was argued that an event weight could be present squashing this 0-barcode contribution to zero.

This could be a question of definitions. One could ask if one needs a precise definition of a fake track, but the answer to this is not straight-forward. For a detector-based truth level study, then the answer is most likely yes, but to produce Combined Performance recommendations, where what matters is only how an analysis downstream of the re-

construction would classify such objects, then the answer is probably no.

The main barrier to further progress, and/or a definitive answer to either of these questions is that the truth definitions that we are trying to diagnose have no information about them at the analysis level that this study is conducted at. The best way of continuing would be to trace a single event through the entire reconstruction process and manually figure out what occurs to the tracks at each stage, but this is a large undertaking and the scope and difficulty of such a task is unknown. Since it is unclear if any more diagnostic plots could be produced in the framework that I created to tackle these questions, and no new recommendations could be created based on this new information, this study was left for future work.

CHAPTER 7

Object definitions

Tracks, which were discussed in the previous chapter, are just one type of *physics object*. A physics object is the term for a collection of signals from sub-detectors that are consistent with the passage of one or more particles. In Chapter 4.3, the process of identifying certain particles¹ via their deposits in various sub-detectors was discussed. However, the precise definitions of these particles is done in terms of physics objects.

The main topic of the thesis is based around the VH , $H \rightarrow b\bar{b}$ analyses presented in the subsequent Chapters 8, 9 and 10. Only objects that are integral to these analyses will be discussed here, and as such this excludes photons.

7.1 Low-Level Physics Objects

‘Low-level’ physics objects are those that come from direct measurements from within the detector. These include *tracks* from the ID and muon spectrometer, vertices inferred from ID tracks, and *towers* and *clusters* from the electronic and hadronic calorimeters.

¹and other ‘high-level’ physics objects but what these are will be explained soon.

The information from low-level objects is combined to form ‘high-level objects’ like particles or jets.

Mis-modellings, bad assumptions, and mistakes that can occur in an experiment of this scale are quantified through the treatment of *systematic uncertainties*. These mistakes can be present in the construction of low-level physics objects but systematic uncertainties are not derived on low-level objects for the analyses in this thesis. Systematic uncertainties will only be defined directly on the resultant high-level objects.

7.1.1 Vertices

After tracks are constructed in the Inner Detector, the information contained within them is extrapolated backwards along the track to identify the point of origin. As mentioned in Chapter 6, the point at which multiple tracks meet is referred to as a vertex.

In each event the vertex with the largest root sum square of the transverse momenta of contributing tracks ($\sqrt{\sum p_T^2}$) attributed to it is called the *primary vertex* [116, p10-12]. It is this vertex which is assumed to be from a hard-scattered proton-proton collision. All the other vertices in the event will be labelled *secondary vertices* and these are thought to originate from pile-up interactions, multi-parton interactions and the decay of particles created at the primary vertex.

7.1.2 Calorimeter Topo-clusters and Towers

Particles that pass through the detector can leave energy deposits in the calorimeters; the structure of which are subject to particle type. The energy deposits are grouped together to form topological clusters (topo-clusters) or towers.

7.1.2.1 Towers

Within given η ranges, the $\eta - \phi$ space of the calorimeters are divided into grids of N_η by N_ϕ elements of size $\Delta \eta \times \Delta \phi$ [117, p5]. The energy deposited in these cells are summed in longitudinal layers, and then summed into towers. If the cell is present in multiple towers, then the energy of the cell is shared according to the fractional area taken up by the cell in each of the towers. Towers are used in the construction of jets.

7.1.2.2 Topoclusters

Unlike the towers, a *topocluster* is a collection of variable size that starts from a ‘seed’ deposit. The ‘topological’ element comes from the fact that these topo-clusters are designed to include neighbouring satellite deposits caused by bremsstrahlung photon emission [118, p2].

The primary seeds are identified using the variable ζ_{cell}^{EM} , computed as

$$\zeta_{cell}^{EM} = \left| \frac{E_{cell}^{EM}}{\sigma_{noise,cell}^{EM}} \right| \quad (7.1)$$

where $|E_{cell}^{EM}|$ is the absolute energy detected in the cell and $|\sigma_{noise,cell}^{EM}|$ is the expected noise. The superscript ‘EM’ indicates that these measures are made at the electromagnetic energy scale² [118, p4]. To start the clustering process, a cell is required to have a $\zeta_{cell}^{EM} \geq 4$. A search is then carried out around this *proto-cluster* for neighbouring cells with $\zeta_{cell}^{EM} \geq 2$ to merge with it. Proto-clusters that share $\zeta_{cell}^{EM} \geq 2$ are then merged. Finally the set of neighbouring cells with $\zeta_{cell}^{EM} \geq 0$ are merged with the proto-cluster to form a ‘4-2-0’ topocluster. Topo-clusters that are then calibrated using test measurements in the ECal, are hence referred to as EMTopoClusters. These form the basis of the subsequent candidate particle reconstruction.

²This is the scale set by typical deposits made by electromagnetically interacting particles, such as electrons

7.2 Decision Trees and Classifiers

Often when there are multiple variables contributing to an object's definition, the object is formed by requiring these variables to be above certain cut-off values. This is called the single-cut technique. This can result in a low yield in the number of objects formed if a complicated object has harsh requirements. In order to try and increase the yield of objects formed, another technique can be used: a multi-variate analysis or MVA.

An MVA is designed to take several variables as an input, and output a *classifier* to identify whether the object wanting to be created exists in the event or not. One of the most common MVA's uses a Boosted Decision Tree or BDT to do this.

With an 'unboosted' Decision Tree (DT), when trying to divide events into signal and background, they are divided into two parts; a training set and a testing set: the training set is designed to be used to train the Decision Tree in the creation of a *classifier* from the input variables, and the testing set is used to test the classifier.

To train the BDT, the events are ordered by their values in each variable. Then consecutively for each variable the sample will be split into two groups: one where the variable value is less than a cut-off (or splitting value); and one where the event variable is greater than the splitting value. The splitting value is chosen to give the best 'separation'. In this case separation means events with values larger than the cut-off are more likely to have the object in (signal), and events with values smaller than the cut-off will be less likely to the object in (background). It is important to note that at this stage, no events are discarded outright.

After the events have been sorted into whether they look like signal or background in a given variable, these events are separated and treated like different samples (but are

never again split in the same variable). Each successive split in a variable creates two more branches until certain criteria (discussed later) are met.

Figure 7.1 shows an example splitting of a training sample which is known to have 200 signal and 200 background events. Each letter from 'T' to 'Z' represents a cut-off value in a variable that splits the event in two. In the training set, it is known which events are signal and which are background.




Figure 7.1: A schematic showing an example decision tree separation of a training set into signal (red) and background (blue) leaves. The black represents branches without any such designation and the letters T-Z are values in different variables that perform the best optimisation of signal and background at each splitting stage. Since Y and Z are variable splittings in a different branch path to T,U,V and X, Y and Z can be (but don't have to be) any one of the variables that correspond to the splittings T,U,V,X but they both cannot be the same.

Depending on how the MC simulations events are generated, they may have been given event weights. Each splitting decision that the training set makes in the branch can be ranked according to the ratio of signal to background that is present in the branch and their event weights. This variable is called purity (P) and is given by Equation 7.2 [119, p2].

$$P = \frac{\sum_s W_s}{\sum_s W_s + \sum_b W_b} \quad (7.2)$$

where $\Sigma_s W_s$ and $\Sigma_b W_b$ are the sum of all the event weights in the signal and background events in a branch respectively. This is defined such that $P(1 - P) = 0$ for a branch that is either pure signal or pure background.

The initial sample keeps being split into two branches, and this process will eventually stop. The branch end is referred to as a leaf and these leaves are either designated to be signal or background. A branch is designated as a leaf if one of four conditions occur:

- A branch has greater than the desired purity in signal or background
- A branch has too few events to continue
- A branch cannot be split in a new variable
- The maximum number of leaves, a variable which can be set beforehand, would be obtained

Figure 7.1 shows an example output for this process. Events that fall on signal leaves (shown in red) are classified as ‘signal-like’ events, and those that fall on background leaves (shown in blue) are treated as background events.

Figure 7.1 shows an example that is not perfect. There are many events that are known to be signal that incorrectly fall on background leaves. In addition to this, Decision Trees are unstable. Small changes in the training sample can lead to large changes to the criteria used to form branches which affects the structure of the entire tree and the number of final leaves produced. This process can be improved upon by *boosting* events to form a *Boosted Decision Tree* or BDT.

7.2.1 Boosted Decision Trees

A BDT constructs trees in the same manner as the unboosted version starting off with unweighted events but then introducing a ‘misclassification’ weight for each event. The

misclassification of an event occurs when a signal event ends up in a background leaf or a background event ends up in a signal leaf. At the point where leaves are formed, misclassified events will have their weight increased or *boosted*.

The process of building trees from the starting sample is repeated using these new boosted events which will produce a different tree with yet different event weights. This is repeated thousands of times until the changes from one tree generation to another is stable and the end result is the formation of a Boosted Decision Tree.

Each event is then scored. During the training, every time an event lands on a signal leaf the event is assigned a score of +1, and every time an event lands on a background leaf it is given a score of -1. The final event score of the event is then the renormalised sum of all these scores with any associated boosted weightings. The higher the score, the larger the probability that a chosen event is a signal one. By cutting on a particular score, the desired ratio of signal to background in a sample can be selected.

The same logic can be used to describe and form classifiers that are used to describe a region in a continuous fashion. Previously, to identify events that belong in a signal region, a signal-to-background optimisation would be performed on each variable. This would select a cut-off value in each variable and events that have the variable below the cut-off value are completely removed. Using BDT's a classifier can be formed to discriminate signal from background while throwing away no events outright.

7.3 Leptons

While leptons are technically the name given to particles that interact with the electroweak force and not the strong force; in the chapters to come, the use of the word lepton refers to only electrons and muons.

This is because it is useful to have a term to separate electrons and muons from the other leptons³ as these are objects that have high reconstruction efficiencies.

Neutrinos are only detected by momentum imbalances seen within the detector; and taus are heavy enough to decay hadronically meaning that their signature usually requires its own analysis outright.

Similar to the track selections for ID tracks, electrons and muons can be designated *loose* or *tight*. In both cases the particle candidates are required to have a p_T of > 7 GeV, a small impact parameter, and to be constructed within the region of the detector devoted to precision physics ($|\eta| < 2.47$ ⁴).

7.3.1 Electrons

Electrons are constructed from topo-clusters that are matched to ID tracks [118, p8]. Instead of simply building electrons from these deposits from a cut-based method, the constituent objects are fed into a multivariate analysis (MVA) which constructs the *likelihood* that the input features resemble something that is known to be an electron. In this there are additional possible requirements on this probability in the form of ‘LooseLH’, ‘MediumLH’ and ‘TightLH’ [120, p11] computed from ‘shower shape’ and track quality variables.

The analyses make use of both *loose* and *tight* quality electrons. The tight requirement will be used when only a single electron track is expected in an event and the loose requirement will be used when more than one electron track (usually two) is expected.

³Yet more proof that in general physicists are terrible at naming things. I say why not use ‘emu’ instead! It sounds funny because it’s the name of bird and best of all it’s technically correct.

⁴Not quite 100% coverage

In addition to the properties already mentioned above, *loose* electrons are required to have tracks that are isolated from those that go into other objects. The quantity used to determine the track isolation is the scalar sum of the transverse momentum (p_T) of the tracks with $p_T > 0.4$ GeV in a cone of ΔR around the electron⁵, excluding the track of the electron itself. The tracks to be considered must be of good quality and originate from same the primary vertex as the electron track [120, p12]. and fulfil a LooseLH quality criteria. *Tight* electrons are selected to TightLH criteria and have stricter isolation but based on calorimeter deposits.

7.3.2 Muons

Since the muons have their own sub-detector, muon reconstruction can happen independently in both the ID and the MS. The information from the two sub-systems is then combined to form muon tracks. The ID muons are constructed much like their electron counterparts, but the requirements in the MS are slightly different. While the p_T still has to be above the same threshold of ≥ 7 GeV, the tolerated acceptance is increased to $|\eta| < 2.7$ to cover the entire pseudo-rapidity of the MS [1, p4].

Loose and *Tight* muons are used in the analyses, and similarly to their electron counterparts the tight requirement is used when only a single muon track is expected and the loose requirement is used when more than one track is expected.

Also similarly to electrons there are four different types of muon quality selections that can be built: ‘Loose’, ‘Medium’, ‘Tight’ and ‘High- p_T ’ [121, p6]⁶ with ‘Medium’ being the ATLAS default.

⁵ ΔR measures the size of a resulting object in the (rapidity- ϕ plane)

⁶Each criterion is designed to maximise a certain aspect of muon reconstruction: ‘Loose’ is designed to maximise efficiency; ‘Medium’ for the determination of statistical uncertainties; ‘Tight’ for muon purity at the cost of some efficiency; and ‘High- p_T ’ for momentum resolution.

Loose muons are selected for the analyses using the ‘Loose’ [121, p7] quality criterion and are required to have a similar level of track isolation to that of electrons. *Tight* muons fulfil ‘Medium’ [121, p7] quality criteria and have a stricter track isolation requirement.

7.3.3 Taus

Taus are the heaviest known lepton (in the original sense) and while the analysis will not use taus to look for signal, they are massive enough to decay into quarks as well as into other leptons. Leptonically decaying taus will decay into muons and electrons, but these can be distinguished from the ones that are produced from the decay of heavier objects, as they typically will come from displaced vertices and have large impact parameters. Hadronically decaying taus decay into quarks which form jets in the detector. These jets, however, are not easily distinguished from deposits from other hadronically decaying particles.

Taus, therefore, feature in the analysis in the form of a ‘veto’. Events that reconstruct tau candidates are removed from the analysis to avoid taus being misidentified as jets. The information from these tau candidates also goes into the missing transverse energy calculation (see Section 7.4).

When taus decay they mostly do so hadronically ($\sim 65\%$) [29, p17] into pions (and kaons). Pions can have charges of $+1$, 0 and -1 and have masses of about 135 MeV (kaons are heavier at 493 MeV). Neutral pions do not leave tracks in the detector so the ones that are mostly useful for the construction of taus are the charged ones. Since taus have a charge of ± 1 , and are many times heavier than pions there are two main types of decay mode that are interesting for forming the tau veto. These are classified with the number of charged tracks in the final state: 1-prong ($\tau^- \rightarrow \pi^- \nu_\tau$) and 3-prong

$$(\tau^- \rightarrow \pi^- \pi^- \pi^+ \nu_\tau)^7.$$

Tau-leptons constructed within the analysis must have $p_T \geq 20$ GeV and $|\eta| < 2.5$ but lie outside of the transition region between the barrel and end-cap ECal, $1.37 < \eta < 1.52$ [1, p4].

Like the lighter charged leptons, tau candidates are reconstructed using MVA techniques. The quality labels ‘loose’, ‘medium’ and ‘tight’⁸ are based on the target reconstruction efficiency and hence are often referred to as *working points*.

The working points for tau candidate reconstruction in the ‘loose’ : ‘medium’ : ‘tight’ scheme are 0.60 : 0.50 : 0.45 for 1-prong and 0.50 : 0.40 : 0.30 for 3-prong [122, p8]. The analyses use ‘medium’ quality taus [1, p4].

7.4 Missing Transverse Energy: E_T^{miss}

When two beams collide head-on at the ATLAS experiment, their energy is dispersed throughout the detector. Since the particles have no initial motion in the r - ϕ plane perpendicular to the beams thanks to quadrupole magnets, the total energy present in the r - ϕ plane in any given event must sum vectorially to 0⁹. If this is not the case, then due to the conservation of energy, there must be energy that is “missing” (having either been carried off by an undetected particle, encountered a dead region of the detector, or been assigned to the wrong time-frame). The energy vector ($\vec{E}_T^{\text{miss}}$) and its magnitude (E_T^{miss}) required to balance the visible energy deposits are referred to as missing transverse energy. In practice, when it comes to recording important events, only the magnitude of $\vec{E}_T^{\text{miss}}$ is important.

⁷In the case of pions, there are possible 5-prong modes, but these occur less than 1% of the time.

⁸ yay for originality

⁹In practice, this is an approximation as the beams are collided at a small angle to increase the luminosity.

There are two main types of objects that are added together to get the $\overrightarrow{E_T^{\text{miss}}}$ of an event. First is the negative vector sum of the transverse momenta of all the high-energy physics objects associated with the hard scatter vertex; leptons, photons, hadronically decaying taus and jets [123, p6]. These are called ‘hard’ objects. The second is the negative vector sum of all energies inferred from charged particle tracks not associated with any hard object (unused tracks) [123, p5]. This term is also called soft E_T^{miss} .

In order to correctly evaluate the amount of E_T^{miss} in an event, the objects that go into the calculation must come from detector signals that are mutually exclusive. An overlap removal procedure is therefore applied to remove double-counting between the reconstructed leptons - here, taus and jets.

The way this E_T^{miss} object is constructed while the experiment is running (online) is that it only takes into consideration the ID tracks and the calorimeter deposits, but not the information from the MS. So while it does take into account of the energy of the muon deposited in the calorimeter, since muons are minimum ionising particles at typical collider energies, this is only a tiny fraction of the muon’s true energy. Not correcting for this muon energy means that most of the energy of the muons appears as E_T^{miss} . This can be corrected some time after data taking (offline), where later triggers can recalculate E_T^{miss} using the full detector information.

This means that the algorithms used to calculate E_T^{miss} in the Level-1 trigger and the High Level Trigger do not use muon-corrected L1 calorimeter information. This will become important in Chapter 10.

7.5 Jets

7.5.1 What is a jet

The analyses in later chapters focus on the particular decay of the Higgs boson into quarks, specifically b -quarks. Due to quarks having colour charge and interacting via the strong force (QCD), they are not able to be observed in isolation. A principle of QCD is that only colour neutral particles can be observed [124] and thus these quarks form hadrons made of two (mesons) or three (baryons) quarks that have an overall neutral colour charge.

As mentioned in Section 5.3, this process is called hadronisation and the initial high-energy hadrons are unstable. They decay on short timescales into multiple lower-energy hadrons that are picked up by the detector.

These final-state hadrons are combined using algorithms to form ‘jets’. The jet definition comes from what hadrons are attributed to a given jet, and how the momenta of those hadrons are combined. It is these jets that are then used to infer properties of the initial quark/gluon, such as energy, flavour and momentum. There are many ways of doing this, as no particular way is intrinsically correct, hence there is no unique jet definition.

7.5.2 Choosing an algorithm

A jet algorithm is a set of rules that constructs jets from groups of candidate particles. Since there are many algorithms that exist to construct jets, it is useful to have some rules about what is wanted from the definition of a jet.

The “Snowmass Requirements” [125, p3] are a series of conditions placed on jet algo-

rithms. They require that the jet algorithm is simple to implement in both the experiment and in theoretical calculations; and that physical values of cross-sections yield physical answers at any order of perturbation theory and somewhat independent of the hadronisation process ¹⁰.

Together these requirements ask that the jet observables (most often the number of jets in a given event) from the algorithm are stable under small fluctuations in circumstances. Such fluctuations include the emission of a soft (low energy) particle (*infrared safety*) and the emission of particles at angles close to the direction of travel (*collinear safety*).

Generally speaking there are two types of jet algorithm [126, Section 2.1 & 2.2]: those that sweep out areas around seed particles, called cone algorithms; and those that create jets by adding particles one at a time, called sequential recombination algorithms. Figure 7.2 shows a few varieties of these two types of jet algorithm, some notes on how each algorithm creates jets, and for the conical algorithms, experimental examples of each type.

Sequential recombination algorithms are the ones now mostly used at colliders, as they are reasonably fast and are both infrared and collinear safe. As shown by Figure 7.2, they are controlled by two parameters: a and R . R is the same parameter mentioned previously that determines the size of the jet in the rapidity- ϕ space, and a changes the order at which particles are clustered into jets.

¹⁰At the time of the 1990 conference, it wasn't clear what sort of algorithm could fulfil these requirements, and it took the best parts of two decades to create algorithms that did.




Figure 7.2: A graphic giving a brief overview of the types of jet algorithm that exist. y_i , ϕ_i and $p_{T,i}$ are the rapidity, azimuthal angle, and the transverse momentum of the candidate object to be added to the jet; B is the beampipe; R is the radius of the jet cone, and ΔR_{ij} is the radial distance between the candidate and 'seed' jet object.




Figure 7.3: Three sequential combination algorithms run on a sample event with many soft particles added. This is designed to illustrate one of the properties of the jet: the catchment areas of the hard jets. The shapes in the k_T and Cambridge-Aachen algorithms are partially determined by the specific set of soft particles added, making anti- k_T the ideal algorithm. Images taken from [127, p4].

A positive a -value clusters the particles from softest to hardest. The k_T algorithm is an example of this where a is 1. An a -value of 0, is an energy independent clustering called the Cambridge/Aachen algorithm. A negative a -value clusters the objects from hardest to softest. The anti- k_T algorithm is an example of this where a is -1, which has the additional bonus of creating near-perfect conic jets as shown in Figure 7.3 which compares the ‘conic’ directrices of the three algorithms.

7.5.3 Types of Jet

Jets can be built from calorimeter energy deposits or from tracks, and as well as having multiple types of algorithm that one can use for making jets. There are many types of jets used in the analyses described in later chapters.

7.5.3.1 Track jets

Track jets only use information from the ID to create jets. These are primarily used for identifying smaller jet objects within larger ones, and therefore they are primarily used in the Boosted analysis where large jets with substructure will be constructed (more on this in the Large- R Jets section). The track objects that are considered candidates for track jets are required to at least fulfil loose track criteria (as seen in Table 6.1 in Section 6.3).

Fixed-Radius (FR) track jets use the anti- k_T algorithm with a constant radius of 0.2. At larger transverse energies, the decay products from the Higgs boson are closer together, and the performance can degrade.

Variable-Radius (VR) track jets instead use an ‘effective radius’ that changes as a function of p_T given by Equation 7.3 [128, p6].

$$\Delta R_{eff}(p_T) = \frac{\rho}{p_T} \quad (7.3)$$

The parameter ρ determines how quickly the effective jet size decreases with p_T . In order to ensure that the jet radius remains sensible, the R_{eff} usually have an upper and lower bound, given by a R_{min} and a R_{max} . In the Boosted analysis, ρ is set to 30 GeV and produces jets bound between R-values of 0.02 and 0.4.

The upper bound is chosen to be $R=0.4$, as this is the value used in the Resolved analysis. Events in the lower- p_T limits of the Boosted analysis have some overlap with those considered for the Resolved analysis.

Track jets are required to have a $p_T > 10$ GeV, $|\eta| < 2.5$, and have at least two track constituents [4, p6].

7.5.3.2 Calorimeter Jets

Calorimeter Jets in ATLAS are built using the anti- k_T algorithm on 4-2-0 EMTopoClusters from the ECal and HCal. There are two types of calorimeter jets that are used in the analyses described in the next chapter: Small- R and Large- R jets.

The ΔR between the decay products from the Higgs is given approximately by the following formula 7.4 [129]:

$$\Delta R_{ij} \approx \frac{m_H}{p_T^H} \frac{1}{\sqrt{z(1+z)}} \quad (7.4)$$

where m_H and p_T^H are the mass and the transverse momentum of the Higgs respectively and z is the momentum fraction taken by the jet with the most energy (the leading jet - b_1). It is this rough calculation that decides what size a jet used in the analysis should be.

Small- R Jets are reconstructed using the anti- k_T algorithm with $R = 0.4$. Outside the acceptance of the tracker ($2.5 < \eta < 4.5$), jets are required to have a $p_T \geq 30$ GeV, and within, a $p_T \geq 20$ GeV. ‘Cleaning criteria’ are required to separate out and remove events with jets that are a result of calorimeter noise or non-collision backgrounds¹¹. Since hard-scatter events are of principle concern, a jet vertex *tagger* removes jets associated with all secondary vertices not associated to the primary vertex for jets of $p_T < 120$ GeV and $|\eta| < 2.5$.

Large- R Jets are used in higher energy regimes. The momentum fraction taken by the leading jet in comparison to the sub-leading jet approaches $\frac{1}{2}$ which means Equation 7.4 can be simplified to Equation 7.5.

$$\Delta R_{ij} \approx \frac{2m_H}{p_T^H} \quad (7.5)$$

¹¹These are events such as cosmic rays interacting with the detector, or neutron albedo

Hence in the region of $p_T^H > 250 \text{ GeV}$, the ΔR between the two jets will be less than 1, and both particles will be able to be attributed to a single jet of $\Delta R=1$.

An $R=1$ jet is built and then *groomed* (grooming will be covered in Section 7.5.4). After this a VR track jet algorithm is run on the tracks associated to the event. These jets have to be matched to the large- R jet, but this cannot be done using only geometric quantities like a ΔR separation due to irregular jet boundaries, and the dense hadronic environment. The process that is used is called ghost association.

Ghost Association

A *ghost* is a track-jet object whose 4-vectors are made to have infinitesimal p_T such that essentially only part of the vector that remains is the part that indicates the tracks direction. Ghost association is where ghosts of track jets made from only the directional information of a track are matched to a re-clustered large- R jet.

The EMTopoClusters are once again clustered into a large- R jet but containing the ghost objects. The removal of the ghosts energies is done to ensure that the jet reconstruction is not affected by the presence of the ghosts. Ghosts that fall within the cone of the jet are said to be associated to that jet.

7.5.4 Jet Grooming

Since jets are sensitive QCD objects, there is a delicate balance to be had in capturing all the energy deposits associated with a jet while minimising contaminations from other hadrons such as: initial state radiation, multiple interactions, and event pileup. With larger jet radii, one can be more certain that all the daughter emissions that come from the primary quark of interest have been captured in the jet cone, but the resolution/reconstruction of the said particle decreases as more contaminants are included within the jet.

The choice of an algorithm may not be enough to ensure that the final observables are unaffected by the myriad background processes that can contaminate jets. After being formed, jets have to be groomed. Grooming is defined as any procedure that outputs a jet when a jet is input, oftentimes with a different energy (mass) [130, p15]. This is done by a selective combination of pruning and trimming which require re-running jet algorithms on stable jet objects (called seed jets) and removing objects that fail cuts.

7.5.4.1 Pruning

When re-evaluating the deposits that make up jets, it is useful to create criteria that are desired for sub-jets. This can be done by defining two variables: θ_{12} and $z_{12,j}$. These are mathematically defined in Equation 7.6 [131, p2].

$$\theta_{12} \equiv \Delta R_{12}, \quad z_{12,j} \equiv \frac{\min(p_{T,1}, p_{T,2})}{p_{T,j}} \quad (7.6)$$

θ_{12} is the characteristic radial (R) separation between two sub-components of the jet and $z_{12,j}$ is the fraction of the transverse momentum of the jet that ends up the smallest of the two chosen sub-jets for a jet seed, j .

For any two chosen sub-components of the jet, if their angular distance θ_{12} is too large and their $z_{12,j}$ falls below a given cut value, then the softer jet sub-component is discarded and the algorithm is repeated with a new sub-component of the jet. If one or both of the conditions are not satisfied, both objects are recombined into the jet j . Once all the sub-components have either been merged into j or discarded, the jet is said to have been pruned. Pruning is good at removing soft, large angle emissions from a jet.

7.5.4.2 Trimming

When a jet is trimmed, a seed jet is reclustered with a smaller radius. Each sub-jet within the seed jet that falls below an energy cut-off is removed from the seed jet [132,

p8]¹². Trimming is useful for removing unimportant sub-structure from a jet.

When forming Large- R jets for the analysis, the track jets are ghost associated to the re-clustered ungroomed Large- R jet. Since the hardest objects in the re-clustered ungroomed Large- R jets are identical to those in the trimmed Large- R jets with the addition of associated ghost objects, the kinematics of Large- R jets are determined by the trimmed jets.

7.5.5 Jet Calibrations

The algorithms that create jets (both real and simulated) are not perfect, and there are some important effects that are not considered when forming them. Jets are calibrated in order to correct for some of these effects. A couple of these such calibrations are listed below.

EMTopo Jets are built out of energy deposits in the electromagnetic calorimeters, and these objects are “extended” into the hadronic calorimeters to pick up emissions that fall within the cone, but as discussed in Section 4.2.3, the response of the ECal and HCal to charged particles is different. EMTopo Jets are calibrated from the electromagnetic (EM) scale to the hadronic (jet energy) scale through a series of steps [133, p6-9].

- Firstly the jet axis is realigned such that it points to the primary vertex.
- Secondly jet pile-up effects are corrected for by subtracting the average pile-up calculated during simulations.
- Thirdly the jet energy scale correction (JES) is applied to MC simulated events jets seeded from ECal deposits to reach the hadronic particle scale.
- Lastly “in-situ” corrections are applied to accommodate for differences between jets created from actual data and those formed in simulations. In the (2-lepton

¹²There is a nice image here [132, p9]

channels of the) analyses this is done by the Kinematic Likelihood Fitter (KLFit-ter). The KL Fitter for $Z(\gamma)$ +jet events uses the fact that the entirety of the energy of the system is contained within “fully reconstructable” objects to balance out the energy of the decay products of the Z-boson with any jets formed.

In addition, the jets made from MC simulated events are often built without pileup. These differences need to be compensated for, and in ATLAS it is done at the jet calibration stage.

In order to investigate and add effects from pile-up, a *nominal* sample of events with two jets (also called dijet events) is simulated and distributed according to the average number of additional collisions. These *minimum-bias* dijet events are added to the hard-scattered events. These events end up with an additional energy density over them and therefore inside any jets formed. The average pile-up energy density is calculated and the typical energy of a hadronic jet (jet energy scale) is estimated.

7.6 B-Tagging Jets

Identifying the flavour of the original parton which leads to the formation of a jet is important as this differentiates possible signal jets from the other background jets. This is done by jet-flavour tagging jets, and jets that are identified as coming from b -quarks are said to be b -tagged.

Simulated jets undergo the process of *flavour tagging*. In this process, jets are then given the label of either b -, c - or light¹³-jets. This happens in two stages. In the first step ‘low-level’ tagged objects are assigned as b -, c - or light-jets according to the respective flavour of the hadrons with $p_T \geq 5 \text{ GeV}$ found within a $\Delta R = 0.3$ cone around the jet axis [134, p7].

¹³ u -, d -, s -quark and gluon

For jets inside the tracker acceptance, the second step is to tag jets as containing b -hadrons using a multivariate discriminant, MV2 (the 2 comes from Run-2), which combines the outputs of the ‘low-level’ tagged objects as well as the kinematic properties of the jets such as p_T and $|\eta|$ [134, p8-9]. The MV2 algorithm is trained on a $t\bar{t}$ sample with a 7% fraction of c -jets (which guarantees that the sample has a good amount of all three types of jet in it). The analysis used a modified version of this algorithm called MV2c10 which increases the c -jet background fraction to 10%.

The selection of MV2c10 is tuned for an average b -tagging efficiency of 70% for b -jets (in the simulated $t\bar{t}$ events), which corresponds to a light-flavour misidentification efficiency of 0.3% and a c -jet misidentification efficiency of 12.5% [3, p7]. Figure 7.4 shows the MV2c10 output distributions for different jet types in an event.




Figure 7.4: The MV2c10 output for b -jets (solid blue line), c -jets (dashed green line), and light-flavour jets (dotted red line) in simulated $t\bar{t}$ events. Image from [135, p6].

One of the most important background samples in the analysis is where a vector boson decays into two jets. This is commonly referred to as the V+jets sample. Simulated V+jets events are further categorised according to the two b -tagged jets in the event. If both the jets are designated light-flavour jets, it is given the designation $V + ll$. If one jet is tagged as a c -jet and the other is a light jet the event is classified as $V + cl$. In the rest of the cases ($V + bl$, $V + cc$, $V + bc$, $V + bb$), the events are designated V+HF (heavy flavour), but after the b -tagging selection it is expected that this mainly consists of $V + bb$ events.

7.6.1 Direct B-Tagging

The normal way to tag a jet directly is to extract the MV2c10 value of the jet and if it exceeds a certain threshold, the jet is b -tagged. This is called *direct* tagging. In practice this approach is highly efficient for backgrounds that are dominated by c - or light jets and results in these samples having large MC simulations uncertainties due to low statistics as it discards jets that fail these cuts. To avoid this, these events need to be tagged another way.

7.6.2 Truth B-Tagging

In *truth* tagging, instead of cutting on the MV2c10 score directly, all events are allocated into all three 0(b)-tag, 1-tag and 2-tag categories regardless of whether the jets in those events contain b -hadrons or not. The event is then weighted on the probability that each jet is a b -jet (it passes a b -jet cut) [136, p8]. This way all events enter all regions but the event is weighted in each category accordingly. A-priori a decision is made, based on the low-level tagging information, on which jets are b -tagged and which are not, and the total permutation weight is the product of the b -tagging efficiencies for tagged jets multiplied by the product of the inefficiencies for untagged jets. These efficiencies are provided as a function of η , p_T and truth flavour. In this method a series of tagged scenarios and probabilities are created. One is then selected at random.

This method significantly increases the statistics present after the full event selection. However, adopting only the truth tagging approach for all the jets in an event results in some closure problems in events with two actual b -jets in them.

7.6.3 Hybrid B-Tagging

The solution is to adopt both approaches. Truth tagging is only used on events that do not contain a truth ' b ' and in the cases where a truth ' b ' is present, direct tagging is used. This method is called *hybrid* tagging.

7.6.4 B-jet Corrections

In addition to the standard calibrations for a jets energy scale, b -tagged jets receive additional corrections to improve their energy measurement.

μ -in-jet Adds the four-momentum of the closest muon in the vicinity¹⁴ of the b -jet to the b -jet.

p_T reco A residual correction designed to recover energy from neutrinos caused by decays of daughter leptonic or hadronic decays of heavy-flavour hadrons within the jet.

In the case of b -tagging Large- R jets, after they have been trimmed, the two leading ghost-associated VR track jets, have the b -tagging algorithm applied to them. A final scan of the jet will remove events where one of the VR track jets is contained within another and remove the track jet object with smaller p_T when the distance between both of the track axes is less than the radius of the Large- R jet¹⁵.

¹⁴a variable cone of $\Delta R(jet, \mu) < \min(0.4, 0.04 + 10/p_T^\mu)$ where p_T^μ is the transverse momentum of the muon in GeV [3, p7]


¹⁵This happens about 1% of the time

7.7 How will these objects be used

These VH , $H \rightarrow b\bar{b}$ analyses described in the following chapter will make use of the object definitions to define physics objects in the signal and background regions. The physics analyses will be contingent on the ATLAS detector detecting events in which combinations of objects are found in coincidence¹⁶.

¹⁶The specific combination of physics objects in an event, and the way they are distributed with respect to one another is called an event topology

Part III



Physics Analysis

"There's an old Sysan saying that the soup of life is salty enough without adding tears to it."

- Iain M.Banks, *Look to Windward*



An overview of the VHbb analyses

Discussed previously in Chapters 3 and 7, the main task of this thesis will be measuring the cross section of $H \rightarrow b\bar{b}$ at the ATLAS experiment. It was also previously mentioned that this channel is important as it is the largest Higgs branching ratio within the Standard Model. In addition, deviations in the measured cross-section times branching ratio compared to the theoretical one could point to new physics beyond the Standard Model, as it shows that some physics effect is contributing to the mechanics of this process in ways we do not yet understand or are not accounted for in current theoretical models.

This chapter will cover the evolution of the $H \rightarrow b\bar{b}$ analysis over the course of the three years of my involvement, resulting in the publications of [1–4].

8.1 Why VH, $H \rightarrow b\bar{b}$?

There are four processes that produce Higgs particles at the LHC: gluon-gluon fusion (88%), vector-boson fusion (7%), associated production (3%) and $t\bar{t}H$ production (1%). Feynman diagrams for these processes can be seen in Figures 8.1 to 8.4 respectively.




Figure 8.1: Feynman diagram of the Higgs production process with the highest cross-section: gluon-gluon fusion. Theoretically, the loop in the diagram is possible for all quarks, but due to the substantially larger mass of the top quark than the other quarks and the Higgs couples preferentially to heavier particles, the loop mainly consists of top quarks. This process has the highest cross section because at high energies the parton distribution function of the proton is roughly 50% gluons [137, p117] making gluon-gluon interactions the most likely occurrence from colliding protons.




Figure 8.2: Feynman diagram of the production of a Higgs particle via vector boson fusion. V here can either represent a W -boson or a Z -boson.




Figure 8.3: Feynman diagram of associated production of a Higgs particle with a vector boson. V here can either represent a W -boson or a Z -boson. The cross-section represented by this diagram is lower than the other diagrams because this interaction requires that a 'sea' anti-quark from one proton interacts with a valence quark of the other.




Figure 8.4: Feynman diagram of production of a Higgs boson produced alongside a pair of top quarks. Theoretically, this diagram is possible for all quarks, but due to the substantially larger mass of the top quark relative to the other quarks, this specific process has the only relevant cross-section.

There are several difficulties in any analyses such as this, and the majority of them relate to the detection of the final state particles. The channel $gg \rightarrow H, H \rightarrow b\bar{b}$ has a large background at the LHC as the identification of two b -quarks requires the formation and identification of b -tagged jet objects. This process is made much harder when there are numerous other jets created in the QCD ‘sea’ produced by the collision of protons. This reduces the likelihood of forming the correct jets that come from the b -quarks produced by the Higgs boson.

Any analysis that wants to construct this signal needs a way of reducing this ‘multi-jet’ background. One such solution to this problem is to look for events that have both b -jets and other peripheral non-jet objects to tag simultaneously, like leptons with large momenta or E_T^{miss} . This vastly reduces the likelihood that a background fluctuation can reproduce the signal. The disadvantage of this approach is that it lowers the cross-section of production, as more processes are required that produce particles alongside the Higgs boson. Figure 8.5 expands on the Feynman diagram shown in Figure 8.3, showing possible relevant decay modes for the particles produced alongside the Higgs boson. This process, now to be referred to simply as $VH, H \rightarrow b\bar{b}$ is the process specifically considered in this thesis. The process of $VH, H \rightarrow b\bar{b}$ is analysed in two complementary approaches here referred to as VHbb Resolved and VHbb Boosted¹. The idea is to create a signal event topology that is vastly different to that of the background. A *signal region* that is too niche will typically not have enough events to control the background to a high level of precision.

8.2 Past analyses

After the discovery [7, 8] of the Higgs boson in 2012, successive studies into its nature have tried to identify the extent its expected decays into various particles match those

¹The Boosted Analysis was launched separately to the resolved analysis after the discovery paper of 2018 but they share common goals and analyses frameworks and tools.




Figure 8.5: Feynman diagram showing the full Higgsstrahlung process of Higgs production at the LHC, and its decay into the final states of interest to the analysis. Gluon initiated (rather than quark) diagrams do exist but they require two additional vertices and are thus rarer.

predicted by theory. In this respect, the $H \rightarrow b\bar{b}$ analyses (with VH , $H \rightarrow b\bar{b}$ being of particular relevance) have existed in many forms.

The basic idea of an analysis like this is to compare data taken from measurements to what is expected from simulations of the Standard Model and extract various figures of merit. The Standard Model simulations are Monte Carlo samples (discussed in Chapter 5) and subsequent modelling studies are performed to make the best estimate of what the signal and background would look like in the regions of interest (ROI) specific to the analyses.

In the case of all the VH , $H \rightarrow b\bar{b}$ analyses this means: using versions of the event generators such as POWHEG-BOX [138], PYTHIA8 [106] and SHERPA [139] to create events; using techniques like template fitting data to MC simulations in regions close to the signal region but with few intended signal events (sideband regions or sidebands); and studies on specific backgrounds to create a better ROI model. The aim of this is to extract the ratio of the observed yield of the Higgs boson decaying to two b-quarks against the expected yield in the signal regions (signal strength), and find the statistical

significance of this ratio.

The main figure-of-merit, the signal strength, is interpreted as a measurement of the VH production cross-section multiplied by the $H \rightarrow b\bar{b}$ branching ratio (the $VH, H \rightarrow b\bar{b}$ process). It is a ratio of the observed $VH, H \rightarrow b\bar{b}$ process with respect to the Standard Model expectation.

The signal strength is usually shown as having two separate uncertainties, showing whether the lack of data (or MC simulated events) in analysis ROI's (statistical effects) or modelling uncertainties (systematic effects) are the larger driving force behind the precision in results.

8.2.1 Run-1 analysis - 2014/2015

The first iteration of the $VH, H \rightarrow b\bar{b}$ analysis in its modern form (i.e with the knowledge of the mass of the Higgs boson) attempted to study the Standard Model expectation of Higgs decay into bottom quarks using Run 1 data [140] which consisted of two datasets: one with a centre-of-momentum collision energy of 7 TeV; and one with that of 8 TeV, totalling 25 fb^{-1} .

The Run-1 $VH, H \rightarrow b\bar{b}$ analysis found the observed (expected) deviation of events with a Higgs boson decaying to two b-quarks from background events to be 1.4σ (2.6σ), and a signal strength of $0.52 \pm 0.32(\text{statistical}) \pm 0.24(\text{systematic})$ [140, p54].

8.2.2 The first Run-2 analysis - 2017

The first analysis to run on the Run-2 dataset builds largely on the one used to conduct the Run-1 measurement. As discussed in Chapter 3, the Run-2 dataset has a higher centre-of-momentum energy ($\sqrt{s}$), which increases the cross-section of $VH, H \rightarrow b\bar{b}$, resulting in more expected signal events in a given dataset. In addition to this, the

analysis performed using the data taken in 2015 and 2016 which is a larger dataset - 36.1 fb^{-1} as opposed to the Run-1 Analysis' 25 fb^{-1} .

The nominal 2017 analysis observed (expected) a deviation with respect to the background-only model corresponding to a significance of 3.5σ (3.0σ). This exceeds the 3σ requirement for *Evidence* of a process and this paper (and the analysis) is now referred to as the evidence paper (analysis). The fitted value for the signal strength parameter (μ_{bb}^{VH}) was found to be $1.20^{+0.24}_{-0.23}(\text{stat.})^{+0.34}_{-0.28}(\text{syst.})$ [136, p31].

In order to try and increase the significance of the nominal result, the results from the Run-1 and the 2017 analyses are combined. This combination makes the background excess correspond to an observed (expected) significance of 3.6σ (4.0σ), which is a slight improvement. The combined μ_{bb}^{VH} was found to be $0.90 \pm 0.18(\text{stat.})^{+0.21}_{-0.19}(\text{syst.})$.

8.2.3 The second Run-2 analysis - 2018

The second analysis to run on the Run-2 dataset was done a year later, this time incorporating the data taken in the previous year. The size of this dataset was 79.8 fb^{-1} .

It was at this stage where I joined the team and started to contribute to the analysis. This analysis forms the basis of my studies and the future results and so is described in more detail in the following sections.

8.2.3.1 Event selection and categorisation

Events that constitute background are Standard Model processes without the Higgs producing a b -quark and an anti- b -quark (two b -quarks), and the signal being the Higgs decaying into two b -quarks. $VH, H \rightarrow b\bar{b}$ events were categorised into three different channels depending on how many leptons were observed in the final state alongside the Higgs b -jets: 0-lepton, 1-lepton, and 2-lepton. These leptons, subject to further re-

quirements, are assumed to have come from the associated vector boson and are used to calculate the transverse momentum of the associated vector boson (p_T^V).

In the 2-lepton channel, either a pair of electrons or muons are present in the event from the decay of a Z-boson. In the 1-lepton channel, it is assumed that a W-boson has decayed into a single electron or muon and a neutrino, the latter of which manifests as an amount of E_T^{miss} . In the 0-lepton channel, all the energy of an accompanying Z-boson is turned into neutrinos and hence, the total E_T^{miss} is used as a proxy for the transverse momentum of the Z-boson.

In this iteration of the analysis only E_T^{miss} (for the 0-lepton and 1-lepton channels) and single-lepton triggers (for the 1-lepton and 2-lepton channels) are used to trigger leptons in events. While di-lepton triggers exist for possible use in the 2-lepton channel, they have low efficiency and reduce signal yield considerably compared to triggering on only one of the two leptons (this will be covered more in Chapter 10).

In addition to this there are two types of signal region in each lepton channel: one for events with two jets both of which are b -tagged, and one for those with three jets of which at least two are b -tagged jets². In the 0-lepton and 1-lepton channels, events containing additional jets with $p_T \geq 30 \text{ GeV}$ and $|\eta| < 2.5$ are discarded to reduce the amount of background from $t\bar{t}$ processes. The number of jets in an event is referred to as the jet *multiplicity*.

All ‘selected’ (those that pass the requirements shown in the previous chapter) jets in events are run through b -tagging algorithms. In each jet-region, two objects must be b -tagged and at least one of those tagged b -jets must have $p_T \geq 45 \text{ GeV}$. In the 3-jet region, in addition to this, if the jet with the lowest p_T is one of the two b -tagged jets, then

²Since jets are QCD objects, high energy radiations are not uncommon. Hence in order to recover some events where one of the b -jets radiates a hard particle, a three jet (3-jet) region is included.

the event is discarded. In the 0-lepton and 1-lepton channels, to reduce the impact of $t\bar{t}$ background, the jet regions are split into those with exactly two jets, and those with exactly three. In the 2-lepton channel, because there is a *control region* (or CR) to specifically target $t\bar{t}$ contributions, the exactly-three-jet requirement can be relaxed to include events with higher multiplicities.

Since there are energy-dependent corrections that have to be made, and the signal region covers a large range of possible values for the p_T^V , the channels are further divided in p_T^V . Due to the large $\sqrt{s}$, the number of miscellaneous events with p_T^V large enough to enter the analysis increases. Since the 0-lepton and 1-lepton channels are dependent on accurate E_T^{miss} calculations, a greater number of background contaminants can enter these regions. To mitigate against this, the minimum p_T^V cut-off is increased to 150 GeV where the E_T^{miss} trigger becomes reliably efficient. Events with p_T^V below this are discarded. In the 2-lepton channel, there is no such E_T^{miss} expected in reconstructed events, and thus a p_T^V region below this cut-off can be introduced: the $75 \text{ GeV} \leq p_T^V < 150 \text{ GeV}$ region. Events that have p_T^V below 75 GeV are likely to not have a lot of signal so are not included.

In total the 2017 analysis has 8 signal regions: 2-jet and 3-jet $p_T^V > 150 \text{ GeV}$ regions for both the 0-lepton and 1-lepton channels, and a 2-jet and 3-jet for both the $p_T^V > 150 \text{ GeV}$ and $75 \text{ GeV} \leq p_T^V < 150 \text{ GeV}$ regions in the 2-lepton channel.

8.2.3.2 Control regions

Control regions (CR's) are signal regions with inverted kinematic/flavour requirements designed to isolate specific background contributions. A region dominated by one particular background means that background can be better modelled. Events that fall into any CR's are removed from any SR's.

In the 2-lepton channel, the two leptons are expected to come from the decay of a Z-boson, and as such are expected to be of the same flavour (antiparticles have the same flavour as their particle counterparts)³. If a new region is created with this condition inverted such that different-flavour leptons are requested, then the background most likely to contribute to this is $t\bar{t}$. This is because each of the two leptons in the final state come from a different top quark, and as such have no flavour restrictions. This control region is referred to as the $e\mu$ CR.

The 1-lepton channel also has a control region designed to enhance the contribution of backgrounds with heavy-flavour jets and a W-boson candidate. This control region is referred to as the W+HF CR. In the 1-lepton signal regions, there is either the requirement that the $m_{bb} \geq 75 \text{ GeV}$ or that the ‘top-quark’ mass $m_{top} \leq 225 \text{ GeV}$ ⁴. In the 1-lepton CR, there is a requirement that events have both $m_{bb} \leq 75 \text{ GeV}$ AND $m_{top} \geq 225 \text{ GeV}$. The m_{top} requirement decreases the effect of $t\bar{t}$ background.

In each case, there is a control region for each jet category and p_T^V category, making four CR’s in the 2-lepton channel and two CR’s in the 1-lepton channel.

8.2.3.3 Backgrounds

Each of the backgrounds are modelled individually in the analysis. There are 5 main backgrounds in the analysis: W+jets, Z+jets, diboson, $t\bar{t}$, and single top. How each of these backgrounds enter various lepton channels will be expanded on in Section 8.3.3.

A less prevalent contribution, but nonetheless an important one are backgrounds with

³For muons, since they leave longer tracks, it is easier to identify which way the tracks curve in a magnetic field. Hence there is an additional requirement that both muons observed have charges that are opposite-sign.

⁴A top quark more than 99% of the time decays to a W-boson and a b-quark, and in the 1-lepton a W-boson decays into a lepton and a neutrino. The quantity m_{top} can be evaluated by estimating first the mass of the W-boson in the event from the reconstructed E_T^{miss} and lepton, and then combining this with the mass of the b-tagged jets that gives the answer closest to the known pole mass of the top quark, 172.5 GeV.

high jet multiplicities. These are referred to as multi-jet backgrounds.

Multi-jet background

The analysis signal regions assume that leptons triggered in events come from leptonic decays of the weak bosons via electroweak (EW) processes. However, in rare cases jets can be mis-identified as leptons⁵. Multi-jet backgrounds are produced with large cross-sections so events that have multiple jets in them (multi-jet), can have a non-negligible contribution to the signal region. This background is modelled separately in each of the three channels, but is found to be negligibly small in the 0-lepton and 2-lepton channels. In the 1-lepton channel, the transverse W-boson mass (m_T^W) of a candidate W-boson is used to discriminate between the multi-jet and EW contributions. This was found to be about 1.9% (2.8%) of the total background contribution in the 2-jet electron (muon) sub-channel, and 0.2% (0.4%) for the 3-jet sub-channel [1, p11].

8.2.3.4 Analysis techniques

In addition to the construction of final objects in these channels given largely by requirements on objects similar to those stated in the previous chapter, additional topological and kinematic constraints were applied. These constraints were made on variables constructed from event-level objects and include things like; the distance in R between the two b-jets $\Delta R(jet_1, jet_2)$; the amount of the mass of the W-boson resolvable in the transverse direction (the transverse mass) m_T^W ; and the mass of the dilepton system m_{ll} . Some of the constraints, like $\Delta R(jet_1, jet_2)$, are common to all the channels, but most are specific to a particular channel (m_T^W in 1-lepton and m_{ll} in 2-lepton). However instead of cutting hard on these constraints, the variables in each region are fed into a boosted decision tree to create a multivariate discriminant, the output of which is de-

⁵Contributions to the multi-jet background can include detected electrons and muons from semi-leptonic decays of heavy-flavour hadrons or neutral pions that decay into two photons, where the photon *converts* into a e^+e^- pair. In the case of the latter, the electrons would normally fail isolation requirements put on leptons that go into the analysis, but on rare occasions they could pass.

signed to separate signal from background.

Each individual background and the signal is modelled separately with the aim of obtaining the normalisations (required changes in the expected number of events for a particular sample), shapes (changes in the expected number of events with relation to some variable), acceptances (changes in the expected number of events between regions expected to behave the same due to detector effects) and the signal yield. A likelihood fit is used to extract all of this in all regions simultaneously (an explanation of fitting in general is given in Chapter 9).

8.2.3.5 The fit

The fit takes all the final selections from the analysis and transforms this into our figures-of-merit of interest. This essentially means combining all the information gleaned in every bin of the analyses in all the signal and control regions and turning that into one (oftentimes a few) final number(s). Any change in the analysis model will affect the result from the fit in some way, and hence any changes to the fit model requires a re-running of some or all of the analysis chain.

A more detailed description on how the VH , $H \rightarrow b\bar{b}$ fitting scheme works will be given in Chapter 9, so only an overview is given here.

The MVA is used for all the nominal signal regions. It performs a BDT signal-to-background optimisation in each of the eight SR's to create 8 BDT_{VH} discriminants. These will form the primary inputs into the global likelihood fit. This likelihood fit also takes input from the CR's with event yields coming from all of the W+HF CR's and the 2-jet $p_T^V < 150 \text{ GeV } e\mu$ CR, and m_{bb} distributions coming from the rest of the $e\mu$ CR's. The global likelihood fit is carried out on all eight SR's and six CR's regions simultaneously.

8.2.3.6 Systematic uncertainties

In an analysis like this, the model used to describe the physics and the assumptions that go into the model have associated uncertainties. All those that are not directly statistical in nature are referred to as *systematic uncertainties* as mentioned in Section 7.1.

In this analysis the systematic uncertainties are split into 4 types: those of experimental origin; those associated with the modelling of the simulated backgrounds generally, with a specific focus on those related to the multijet background estimation; and those concerning the Higgs boson signal [1, p11]. While individual analyses may have different uncertainties, they are roughly handled in the same way for all analyses, so these will be covered in a later section.

8.2.3.7 Cross-check analyses

Since an analysis is made up of many intricate parts, it can be hard to determine whether the larger-scale analysis decisions are biasing the result as a whole. The idea behind the cross-check is to look at the data mostly with the same parameters, background templates, systematics, but with a different approach to extracting the signal. This is designed to test whether the analysis as a whole is biased towards seeing an excess of events or to see if the treatment of common systematics is done fairly.

There are two types of cross-check analysis: one where the ‘signal’ is a background process with a similar topology as the main analysis signal; and one where events are processed into signal and background regions differently (for example having a cut-based analysis (CBA) instead of a MVA). The 2018 analysis had two cross-checks to this effect.

The dijet-mass analysis is a CBA scanning the spectrum of the invariant mass of the two-b-quark system; m_{bb} designed to test the usefulness of the MVA.

The diboson analysis uses the same MVA as the main analysis but the ‘signal’ is $VZ Z \rightarrow b\bar{b}$ designed to test the treatment of the signal uncertainties in the main analysis.

Since the main (nominal) analysis and the cross-check are computed with the same dataset, in order to compare the two, the cross-check is *bootstrapped*⁶ to the nominal analysis. When a measurement is taken at an experiment, it is assumed that the dataset follows an unknown distribution. Bootstrapping is the process where random events are drawn and replaced (so it’s possible to draw the same event more than once) from a generated sample to replace this distribution. This is a quick and accurate way of assessing the relative differences between two approaches on a dataset [141, p336-339]. In this way, the relative correlation between the two approaches and the statistical consistency (how stable the result is to the size of the sample dataset) of any observed results can be seen.

8.2.3.8 Results

Figure 8.6 shows the final BDT_{VH} discriminants for six of the eight signal regions. When all the lepton channels are combined, the probability p_0 of explaining at least the observed amount of data only using the background alone is 5.3×10^{-7} . This is perhaps made more clear with Figure 8.7 which is a combination of all eight signal regions and the six control regions into one plot. The expected value of p_0 was 7.3×10^{-6} . The observation corresponds to an excess with an observed (expected) significance of 4.9σ (4.3σ). This was unfortunately⁷ shy of the benchmark set for an ‘observation’ of $VH, H \rightarrow b\bar{b}$. The fitted value of $\mu_{VH}^{b\bar{b}}$ was $1.16_{-0.25}^{+0.27} = 1.16 \pm 0.16(\text{stat.})_{-0.19}^{+0.21}(\text{syst.})$. In order to attempt to push the significance over the 5σ threshold, the analysis once again had to combine the result with other analyses.

⁶Factoid taken from [141, p336]: The name *Bootstrap* goes back to the famous book of Erich Raspe in which he narrates the adventures of the lying Baron von Münchhausen. Münchhausen had pretended to have saved himself out of a swamp by pulling himself up with his own bootstraps (but in the original version he pulls himself out with his own hair.). In statistics, the expression bootstrap is used because from a small sample the quasi complete distribution is generated. There is not quite as much lying as in Münchhausen’s stories.

⁷not for me!

The first combination was done with the analysis result from Run-1 to get more sensitivity in the entire $VH, H \rightarrow b\bar{b}$ result. The significance, however, stayed at 4.9σ even though 5.1σ was expected.

Failing to get an observation over the entire process, the analysers tried to observe it in parts: VH production and $H \rightarrow b\bar{b}$ decays.

To make a measurement of VH production, the $VH, H \rightarrow b\bar{b}$ analysis was combined with others that have the same production mode, but the Higgs boson decays into different particles. The two analysis chosen were the $VH, H \rightarrow \gamma\gamma$ (referred to as $H \rightarrow \gamma\gamma$) and the $VH, H \rightarrow ZZ^* \rightarrow 4$ leptons (referred to as $H \rightarrow 4l$). A combination of $VH, H \rightarrow b\bar{b}$ with $H \rightarrow 4l$ and $H \rightarrow \gamma\gamma$ was only expected to produce a result that was significant to the level of 4.8σ , but a result of 5.3σ was observed.

To make a measurement of $H \rightarrow b\bar{b}$ decay, the $VH, H \rightarrow b\bar{b}$ analysis was combined with other analyses targeting $H \rightarrow b\bar{b}$ decay vertices but with differing production modes. These were: the $t\bar{t}H, H \rightarrow b\bar{b}$ which at leading order uses the production vertex identified in Figure 8.4 referred to as ' $t\bar{t}H$ '; and the vector boson fusion (VBF) analysis whose main production mode can be seen in Figure 8.2 but has a significant contribution from gluon-gluon fusion (Figure 8.1) events, referred to as ' $VBF + ggH$ '. The combination of VH with $t\bar{t}H$ and $VBF + ggH$ resulted in a 5.4σ result, when 5.5σ was expected.

While the main analysis was not able to get an observation of its own, or in combination with Run-1 data, splitting it into production and decay modes and combining it with other analyses resulted in individual observations of the production and the decay. The motivation for the next iteration of the analysis is that the addition of more data will be enough to push the measurement of the entire $VH, H \rightarrow b\bar{b}$ process over 5σ . This is




Figure 8.6: BDT_{VH} post-fitting output distributions for the 0-lepton (top), 1-lepton (middle) and 2-lepton (bottom) channels, in the 2-jet (left) and 3-jet (or ≥ 3 jets for the 2-lepton case) (right) categories in the $p_T^V < 150$ GeV region. The background contributions are shown as coloured filled histograms which each contribution having its own colour. The 125 GeV Higgs boson signal is shown as both a red filled histogram on top of the fitted backgrounds but is normalised to the μ extracted from the data (1.16), and unstacked as a red line, scaled by a factor, which is indicated in the legend. The dashed line shows the total level of the pre-fit background. The hashed region in each bin indicates the total size of the statistical and systematic uncertainties for the fitted signal and background. The ratio of the data to the sum of the fitted signal and background is shown in the lower panel.




Figure 8.7: Event yield as a function of $\log_{10}(S/B)$ for data, background, and a 125 GeV Higgs boson signal. The final BDT_{VH} discriminants from all the analysis regions are combined into bins of $\log_{10}(S/B)$ where S and B are the fitted signal and background yields respectively. The 125 GeV Higgs boson signal is shown as a red filled histogram on top of the fitted backgrounds but is normalised to the μ extracted from the data (1.16). The lower panel shows the statistical significance (pull) of the difference between the data and the fitted background-only model. The points show the observed pulls and the red line is the expected pulls. The vertical error bars on the points are statistical uncertainties only.

not guaranteed though, as the combination with Run-1 proves. The analysis itself will need to be modified to make this more likely. However, as the required boost is small, there was a lack of urgency such there was time to be a bit more ambitious and try new things.

My contributions during this analysis were to test and validate the samples going into

the fit, and to unify selections in the CBA cross-check analysis. The analysis improvements I was working towards were moved to the next iteration of the analysis.

8.2.4 Run-2 STXS paper - 2019

The 2019 paper is an interpretation of the result presented in the previous section. This result, while expected, is a setback to the hopes of seeing direct interaction of BSM particles to Higgs bosons in the near future, since the result of $VH, H \rightarrow b\bar{b}$ produces an amount of $H \rightarrow b\bar{b}$ that is consistent with the SM.

This means that direct $H \rightarrow BSM$ is an ultra-rare process unlikely to be seen⁸. This, however, as discussed in Chapter 2, is only one way that BSM physics can interact with the SM.

Another way that BSM particles/processes can interact with the SM is by increasing or decreasing the rate of $H \rightarrow b\bar{b}$ present in the data via secondary means. In this sense the ‘heavy’ BSM contributions are decoupled from the ‘low energy’ SM regime.

With the observation of the $VH, H \rightarrow b\bar{b}$ process on the horizon, the next stage is to turn μ_{bb}^{VH} into something more useful. One of the ways this can be done is to use the result as an anchor point to reduce some systematic uncertainties, specifically the theoretical uncertainties that affect the signal processes [142, p437]. In this way we can see what effect these BSM processes would have on this figure.

8.2.5 Parametrising BSM effects

When experiments measure signal strength, it is approximately equivalent⁹ to the measurement of the cross-section of VH production multiplied by the branching fraction of

⁸This does not mean that there aren’t analyses looking for this signature, they are just not in the scope of this thesis.

⁹The narrow-width approximation.

the decay of the Higgs boson to b-quarks, normalised by the SM prediction (as seen in Equation 8.1).

$$\mu_{bb}^{VH} = \frac{\sigma_{VH} \times BR(H \rightarrow b\bar{b})}{[\sigma_{VH} \times BR(H \rightarrow b\bar{b})]^{SM}} \quad (8.1)$$

where the branching ratio $BR(H \rightarrow b\bar{b})$, is the fraction of all possible direct Higgs boson decays that produce b-quarks. This could be viewed as the multiplication of three parts, each of which can be individually parametrised.

$$\frac{\sigma_{VH} \times BR(H \rightarrow b\bar{b})}{[\sigma_{VH} \times BR(H \rightarrow b\bar{b})]^{SM}} = \frac{\sigma_{VH}}{[\sigma_{VH}]^{SM}} \times \frac{\frac{\Gamma(H \rightarrow b\bar{b})}{[\Gamma(H \rightarrow b\bar{b})]^{SM}}}{\frac{\Gamma(H \rightarrow all)}{[\Gamma(H \rightarrow all)]^{SM}}} \quad (8.2)$$

The branching ratio is expressed as a ratio of *decay widths*, Γ such that it is easier to parametrise. Unstable particles that have very short lifetimes have limits on how well their energy can be known thanks to the energy-time form of the uncertainty principle ($\Delta E \Delta t > \frac{\hbar}{2}$). A *full* decay width ($\Gamma(H \rightarrow all)$) is equal to $2\Delta E = \frac{\hbar}{\tau}$ where τ is the lifetime of the particle when it decays in this manner. It is also equal to the sum of all of the *partial* decay modes that arise when one focusses on the decay of the Higgs boson into a particular final state.

The part of Equation 8.2 in red would be affected by the parametrisation of the Higgsstrahlung production; the part in green by the parametrisation of the Higgs to b-quarks decay; and the part in blue by the total decay width of the Higgs.

The singular figure-of-merit μ_{bb}^{VH} , therefore, is not sufficient enough to see the detailed effect different values of Wilson coefficients (see Section 2.4.1) have on the analysis result. Since the effects are likely to be subtle, the signal strength needs to be calculated in separate kinematic bins.

SMEFT operators (or combinations thereof) interact with the Standard Model processes by affecting these cross-sections and decay amplitudes. Mostly though, their effects are more visible in the way they change parameter distributions or event yields in particular sub-channels.

This means multiple measurements of the signal cross-section in each of these kinematic bins, leading to many measurements of the signal strength. These results should be also independent on experiment and for this simple template cross-sections (STXS's) are used.

8.2.5.1 Simple template cross-sections (STXS)

STXS's are used to define an exclusive cross-section according to the properties of the Higgs production and/or Higgs decay. They specifically target categories where BSM contributions are expected to be enhanced. It is a dedicated assessment of theoretical systematics, so no other systematics are changed with respect to the previous analysis.

The analysis result is spread across many bins and each experiment that attempts to measure the same process will have different detector effects. *Unfolding* is the process of removing effects due to the interaction with detector hardware (deconvolution, unsmearing) that allows data from different experiments to be fairly compared. The cross-sections that come about as a result of unfolding are called fiducial cross-sections. Unfolding, however, is complicated and requires a lot of statistics.

A template is a histogram that contains an estimate of the number of events in each bin, prepared from a high-statistics region. A *template* of each sample is made to enter into all analysis regions.

These templates are simple because with respect to their fiducial cross-section counter-

parts fewer considerations are taken into account when determining which bins affect the cross-section.

To constrain BSM parameters and contribute to (SM)EFT theories, these STXS bins are parametrised with the effective Lagrangian shown in Equation 8.3 and the analysis will set constraints on coefficients of operators that have an effect on VH processes and/or decays¹⁰.

$$\mathcal{L}_{SMEFT} = \mathcal{L}_{SM} + \sum_i \frac{c_i}{\Lambda^6} \mathcal{O}_i^{(6)} \quad (8.3)$$

It is assumed that in each STXS bin the VH cross-sectional shapes and $H \rightarrow b\bar{b}$ branching ratios are exactly SM-like and the rate or cross sections of these can be scaled up or down by factors.

When running an STXS fit, instead of using a single signal template, each STXS bin has its own signal template. A template is a histogram that contains all BDT bins in all analysis signal regions - each STXS bin enters all the analysis regions. Each STXS template contains only the VH events that lie in that corresponding STXS bin.

When the fit is done, an independent normalisation is assigned to each of the signal templates, which the fit will determine. The result is a signal strength of each STXS bin separately.

8.2.5.2 STXS methodology

The STXS methodology is designed to be staged, so as more data is recorded, the granularity of the measurement increases. It is an agreement by the LHC community [59, 143] to standardise these kinds of results for use outside of the experiments the analysis was conducted in.


¹⁰Though in the 2019 paper, only the operators that affected the production processes were looked at.

‘Stage 0’ corresponds to the separate measurements of the four main production-mode cross-sections. For the VH production, the cross-sections are further divided into $q\bar{q} \rightarrow WH$, $q\bar{q} \rightarrow ZH$ and $gg \rightarrow ZH$. These measurements are done in the rapidity region < 2.5 for the Higgs Boson, as the detector acceptances for the majority of the reconstructed objects such as leptons, photons and b-jets is largest inside this region.


In ‘Stage 1’ these regions are split into 31 subregions by: the multiplicity of non-Higgs-boson jets with transverse momentum $p_T > 30 \text{ GeV}$ (referred to as njets); and the transverse momenta of either the Higgs boson or the weak gauge boson for $VH, V \rightarrow ll$ production [142, p446]. Of the 31 regions, 11 are associated with the $VH, V \rightarrow ll$ production mode and these can be seen in Figure 8.8a. Each of these regions is a STXS bin. If there is enough data, every sub-region of p_T^V and njets space shown in Figure 8.8a can get its own STXS bin, and a $p_T^V > 250 \text{ GeV}$ region is added for the $gg \rightarrow ZH$ sub-channel. This is known as the ‘Stage 1.1’ binning and is shown in Figure 8.8b.

‘Stage 2’ has all the same region splits as in the STXS Stage 1.1 but separates $q\bar{q} \rightarrow ZH$ into the events where the Z-boson decays into $\nu\nu$ and ll .

The Stage 1.1 binning is combined in various ways to align with the analysis regions. The aim is to come up with a STXS regime for the analysis as similar to the Stage 1 binning as possible while ensuring each bin is well populated with events.



(a) Stage 1 STXS Binning



(b) Stage 1.1 STXS Binning

Figure 8.8: Proposed Stage 1 Simple Template Cross-Section Regions from [142, p446] (top). This can be further divided into p_T^V regions and multiplicity to form the Stage 1.1 binning (bottom). Image is a modified version of the one from [143, p10].

8.2.5.3 The 2019 STXS binning

The STXS binning chosen by the analysis is shown in Figure 8.9. With respect to the Stage 1 binning, it differs in three major ways.

- It splits the $p_T^V < 150$ GeV into a $0 \text{ GeV} \leq p_T^V < 75$ GeV and a $75 \text{ GeV} \leq p_T^V < 150$ GeV and removes all of the former from the analysis as this phase space is not accessible.

- The $75 \text{ GeV} \leq p_T^V < 150 \text{ GeV}$ region for the $q\bar{q} \rightarrow WH$ region is removed.
- A $p_T^V \geq 250 \text{ GeV}$ region is added to the $gg \rightarrow ZH$ channel.






Figure 8.9: Modified Stage 1 STXS Binning Proposed Analysis regions for the 2019 STXS VH , $H \rightarrow b\bar{b}$ analysis. Image is a modified version of the one from [143, p10].

These 11 regions are also presented in Table 8.1. These regions are then combined together to form two sets of ‘reduced stage 1’ binnings called the 5 parameters of interest (5-POI) binning, and the 3-POI binning respectively. How the regions are combined into these is shown in Table 8.1, along with the analysis regions with the most dominant event contribution to each STXS bin.

The 5(and 3)-POI binning regime merges the $q\bar{q} \rightarrow ZH$ and $gg \rightarrow ZH$ regions because there is not enough statistics to separate these regions even though they are kinematically different. As this was done, because the 2-lepton region has a 3-or-more jet region as opposed to the 0-lepton and 1-lepton’s exactly 3 jet region, to reduce uncertainty the 0-jet and 1-or-more-jet STXS regions are also merged. This is pictorially represented in Figure 8.10.



(a) Stage 1 STXS Binning.



(b) Stage 1.1 STXS Binning.

Figure 8.10: Proposed '5 POI' Simple Template Cross-Section Regions for the analysis obtained by merging the region present in Figure 8.9. This can be further merged in p_T^V regions to form the '3 POI' binning (bottom). Image is a modified version of the one from [143, p10].

Table 8.1: STXS categories for the 2019 VH , $H \rightarrow b\bar{b}$ Analysis Paper showing which analysis regions contribute the most to which STXS category. It also shows how these are combined to form the final quoted numbers: the 5-POI STXS regime. Table from [2, p6]

Merged region 3-POI scheme	Merged region 5-POI scheme	Stage 1 (modified) STXS region	Reconstructed-event categories with largest sensitivity		
			N_{lep}	$p_{\text{T}}^{V,\text{r}}$ interval	N_{jet}
$WH, p_{\text{T}}^{\text{W}} > 150 \text{ GeV}$	$WH, 150 < p_{\text{T}}^{\text{W}} < 250 \text{ GeV}$	$q\bar{q} \rightarrow WH, 150 < p_{\text{T}}^{\text{W}} < 250 \text{ GeV}, 0\text{-jet}$ $q\bar{q} \rightarrow WH, 150 < p_{\text{T}}^{\text{W}} < 250 \text{ GeV}, \geq 1\text{-jet}$	1	> 150 GeV	2, 3
	$WH, p_{\text{T}}^{\text{W}} > 250 \text{ GeV}$	$q\bar{q} \rightarrow WH, p_{\text{T}}^{\text{W}} > 250 \text{ GeV}$			
$ZH, 75 < p_{\text{T}}^{\text{Z}} < 150 \text{ GeV}$	$ZH, 75 < p_{\text{T}}^{\text{Z}} < 150 \text{ GeV}$	$q\bar{q} \rightarrow ZH, 75 < p_{\text{T}}^{\text{Z}} < 150 \text{ GeV}$ $gg \rightarrow ZH, 75 < p_{\text{T}}^{\text{Z}} < 150 \text{ GeV}$	2	75–150 GeV	$2, \geq 3$
$ZH, p_{\text{T}}^{\text{Z}} > 150 \text{ GeV}$	$ZH, 150 < p_{\text{T}}^{\text{Z}} < 250 \text{ GeV}$	$q\bar{q} \rightarrow ZH, 150 < p_{\text{T}}^{\text{Z}} < 250 \text{ GeV}, 0\text{-jet}$	0	> 150 GeV	2, 3
		$gg \rightarrow ZH, 150 < p_{\text{T}}^{\text{Z}} < 250 \text{ GeV}, 0\text{-jet}$			
		$q\bar{q} \rightarrow ZH, 150 < p_{\text{T}}^{\text{Z}} < 250 \text{ GeV}, \geq 1\text{-jet}$	2	> 150 GeV	$2, \geq 3$
		$gg \rightarrow ZH, 150 < p_{\text{T}}^{\text{Z}} < 250 \text{ GeV}, \geq 1\text{-jet}$			
	$ZH, p_{\text{T}}^{\text{Z}} > 250 \text{ GeV}$	$q\bar{q} \rightarrow ZH, p_{\text{T}}^{\text{Z}} > 250 \text{ GeV}$			
		$gg \rightarrow ZH, p_{\text{T}}^{\text{Z}} > 250 \text{ GeV}$			

The 3-POI scheme produces cross section measurements with lower uncertainties but the 5-POI regime is more sensitive to possible high- p_{T}^V enhancements from BSM processes.

8.2.5.4 Results

Figure 8.11 shows the breakdown of the individual signal strengths in each of the 5 STXS bins.

While the experimental data is consistent with the theoretical bounds, the uncertainties in each bin are large. These will be reduced in the next round of iteration of the analysis which leads us to the most up-to-date analyses.

8.3 Common themes between the current analyses

After the evidence paper, it was decided to split the analysis into two interlinked, but distinct directions: one where the decay products of the Higgs boson are easily resolved

called the VH , $H \rightarrow b\bar{b}$ Resolved analysis (or the Resolved analysis); and one specifically looking at higher energy regimes where the Higgs decay products start to overlap: the VH , $H \rightarrow b\bar{b}$ Boosted analysis (or the Boosted analysis).




Figure 8.11: Measurements of cross-section $\times$ branching fractions of $H \rightarrow b\bar{b}$ and $V \rightarrow ll$ for each of the 5 POI's. The pad below shows the ratio of these partial signal strengths to the SM values. Image from [143, p11].

These are the analyses that formed the greater part of my contributions for this thesis, so will be expanded on in the following section. The results of the Resolved and the Boosted Analyses will not be displayed in this chapter, but can be found later in Chapter 11.

While these analyses are distinct ways of analysing the Run-2 dataset, there are many similarities, definitions and problems common to both analyses. These will be discussed first before going into each analysis separately.

8.3.1 Common data samples and MC simulations

The data used in both Resolved and the Boosted analysis was the entirety of the Run-2 dataset which corresponds to 139 fb^{-1} .

MC simulated events that contribute to both the nominal signal and backgrounds are generated via one of several “event simulators” (see Chapter 4). A summary of this can be found in Table 8.2. The exception to this is the multijet production, as this contribution is measured using ‘data-driven’ techniques elaborated on in Section 8.3.3.

Table 8.2: Generators used for the simulation of the signal and background processes. The acronyms ‘ME’, ‘PS’ and ‘UE’ stand for: matrix element, parton shower and underlying event, respectively. The cross-section order refers to how many orders of the strong coupling constant α_s (unless specified) have been taken into consideration (i.e how many QED vertices are present in contributing Feynman diagrams). Table from [1, p6].

Process	ME generator	ME PDF	PS and Hadronisation	UE model tune	Cross-section order
Signal, mass set to 125 GeV and $b\bar{b}$ branching fraction to 58%					
$qq \rightarrow WH \rightarrow \ell\nu b\bar{b}$	POWHEG-Box v2 + GoSAM + MiNLO	NNPDF3.0NLO ^(*)	PYTHIA 8.212	AZNLO	NNLO(QCD)+NLO(EW)
$qq \rightarrow ZH \rightarrow \nu\nu b\bar{b}/\ell\ell b\bar{b}$	POWHEG-Box v2 + GoSAM + MiNLO	NNPDF3.0NLO ^(*)	PYTHIA 8.212	AZNLO	NNLO(QCD) ^(†) +NLO(EW)
$gg \rightarrow ZH \rightarrow \nu\nu b\bar{b}/\ell\ell b\bar{b}$	POWHEG-Box v2	NNPDF3.0NLO ^(*)	PYTHIA 8.212	AZNLO	NLO+NLL
Top quark, mass set to 172.5 GeV					
$t\bar{t}$	POWHEG-Box v2	NNPDF3.0NLO	PYTHIA 8.230	A14	NNLO+NNLL
s -channel	POWHEG-Box v2	NNPDF3.0NLO	PYTHIA 8.230	A14	NLO
t -channel	POWHEG-Box v2	NNPDF3.0NLO	PYTHIA 8.230	A14	NLO
Wt	POWHEG-Box v2	NNPDF3.0NLO	PYTHIA 8.230	A14	Approximate NNLO
Vector boson + jets					
$W \rightarrow \ell\nu$	SHERPA 2.2.1	NNPDF3.0NNLO	SHERPA 2.2.1	Default	NNLO
$Z/\gamma^* \rightarrow \ell\ell$	SHERPA 2.2.1	NNPDF3.0NNLO	SHERPA 2.2.1	Default	NNLO
$Z \rightarrow \nu\nu$	SHERPA 2.2.1	NNPDF3.0NNLO	SHERPA 2.2.1	Default	NNLO
Diboson					
$qq \rightarrow WW$	SHERPA 2.2.1	NNPDF3.0NNLO	SHERPA 2.2.1	Default	NLO
$qq \rightarrow WZ$	SHERPA 2.2.1	NNPDF3.0NNLO	SHERPA 2.2.1	Default	NLO
$qq \rightarrow ZZ$	SHERPA 2.2.1	NNPDF3.0NNLO	SHERPA 2.2.1	Default	NLO
$gg \rightarrow VV$	SHERPA 2.2.2	NNPDF3.0NNLO	SHERPA 2.2.2	Default	NLO

The total luminosity is measured by the LUCID-2 forward detector on ATLAS, and the uncertainty in the value is 1.7%. The samples generated typically underestimated the average number of interactions in a given bunch crossing, so this is rescaled by 1.03 [144] to give a better agreement to the data. A luminosity uncertainty equal to this

correction is added to the analysis.

8.3.2 Modelling

In multivariate analyses, the background modelling is especially tricky. Since regions are connected, both underestimating uncertainties and being overly cautious with them are problematic. Overestimations of background uncertainties can lead to reduced sensitivities to the signal (as $\frac{s}{\sqrt{b}}$ reduces) and underestimations can lead to poor closure with data in background regions of the analysis¹¹.

It is hard to know for certain if the MC simulated events are giving the correct description of the desired physical effects, but the potential differences between prescriptions can be quantified via shape and normalisation effects. Some backgrounds can be so complicated to generate that the MC simulations have large corrections. These can be hard to use in an analysis so often data-driven methods are used to estimate backgrounds instead.

When the background spectrum is easy to model (e.g it is smoothly falling), and the signal is a resonant peak on top of this (like the signal for the $H \rightarrow \gamma\gamma$ analysis in which the Higgs was first discovered), the estimation of the background in the signal region can be done by parametrising a data sideband with an empirical function and interpolating into the SR. If you do this, then MC simulations are not needed as the estimate of the background contribution in the SR comes from the sideband. This method of determining the background contribution can be used for the entire background or particular samples: either way, these are data-driven techniques.

The background in these analyses is not able to be simply parametrised by a functional

¹¹This example is more pertinent to systematics with normalisation effects (so those that affect the overall event yield). It also depends on the direction of the uncertainty as the inverse could be equally true.

form. In this case we require the use of templates of each sample that goes into the various analysis regions.

Because the background is non-trivial, the signal expected to be seen also needs to be modelled. In a single- μ fit, usually there is only one template for the signal, but in a STXS analysis, there is a signal template for each STXS bin.

Similar to a functional form, a template has parameters to be varied, and these can be done either before or after the fit to the data in all regions. The fit is run multiple times on the control regions to identify the optimal values of these parameters before bins most sensitive to the signal are un-blinded.

8.3.2.1 Multi-jet background estimation

Both analyses use data-driven techniques to estimate the multi-jet background contribution. This is because the MC simulations are statistically limited and are unlikely to reproduce an accurate representation of the multi-jet contribution in this corner of the phase space. Both analyses therefore use a template fit in a dedicated CR to model this.

Both find that the multi-jet contribution is negligible in their respective 0-lepton and 2-lepton channels, and is of the order of a percent for the 1-lepton channel.

8.3.3 Common backgrounds

Here each of the backgrounds will be briefly discussed with a small discussion on uncertainties allowed to float when the global fit is conducted.

As the Resolved analysis and the Boosted analysis effectively share the same final state, the background processes that contribute are the same. Since both analyses have a 0-lepton, a 1-lepton, and a 2-lepton channel, the processes that are dominant in a given

channel in the Resolved analysis will also be dominant in the same channel of the Boosted analysis.

8.3.3.1 V+jets

Figure 8.12 shows a tree-level Feynman diagram of how the process can produce two jets that can pass b-tagging requirements and some leptons.




Figure 8.12: Feynman diagram showing how the V+jets process enters the signal region.

Since the choice of boson on the upper part of Figure 8.12 is free, final states with zero, one or two charged leptons can be seen. This means that this process will enter all the signal regions of the analyses.

The V+jets background is the largest contribution to the background in the entire analysis and is therefore by far the most computationally expensive background to generate with enough statistics to extract a template with the required precision. This is because the production cross-section of W+jets (Z+jets) in the 1-lepton (0-lepton and 2-lepton) is about 5 (4) orders of magnitude larger than the signal.

As explained in Section 7.6, the V+jet backgrounds are subdivided into 6 categories depending on what the jets in the event are tagged as: $V(+)\bar{b}b$, Vbc , Vbl , Vcc , Vcl and

$Vl(l)$. The first four of which are jointly treated as two single backgrounds: $W + HF$ and $Z + HF$ (collectively as V+HF).

The $V + HF$ backgrounds have separate floating normalisations for the 2-jet and 3-jet categories. In the Resolved analysis, the $Z + HF$ is also floated independently in each of the lowest p_T^V categories. Since the $W + HF$ and $Z + HF$ have multiple compositions of jet flavours, there are also sets of flavour uncertainties for each. They are the differences in the ratio of the yield of the relative composition of these flavours with respect to the dominant bb (i.e. $(Z + \text{or } W+)bc/(Z + \text{or } W+)bb$, bl/bb and cc/bb). This is done separately in each lepton channel.

Since the remaining flavour components Vcl and Vl only contribute to $\sim 1\%$ in the Resolved analysis and $\sim 10\%$ in the Boosted analysis they are treated as separate normalisation-only uncertainties.

8.3.3.2 Top and anti-Top ($t\bar{t}$)

99% of the time a top quark is created it decays into a W-boson and a b -quark. When an ATLAS collision event creates two of them, then a signature with two b -quarks and some amount of leptons from a semi-leptonic decaying W-boson are observed. An example of this process at tree-level is shown in Figure 8.13.

$t\bar{t}$ production enters all of the analysis sub-channels. For it to contribute to the 1-lepton channel, one of the W-bosons must decay into a tau lepton which then decays hadronically and registers as E_T^{miss} . In the 0-lepton channel, both W-bosons decay in this manner. The efficiency of the detector keeps the fraction of $t\bar{t}$ events entering the the 1- and 2-lepton channels low, however with a cross-section roughly 3000 times that of the signal, its presence in the 1-lepton channel is considerable.




Figure 8.13: Feynman diagram showing how the $t\bar{t}$ process enters the signal region.

The $t\bar{t}$ process has separate normalisations for the 2-jet and 3-jet categories, but these are constrained by CR's in their respective analyses. Uncertainties are derived by comparing the nominal event generated sample to an alternative one.

Often-times as the entire $t\bar{t}$ decay chain is not contained within the detector, the relative number of events between the exactly-two jet, and the three-or-more jet regimes can change the shape of the template. Both analyses have an additional acceptance uncertainty to account for this. The Resolved analysis has it as an event migration between the 0-lepton and 1-lepton channels, and the Boosted analysis has it between regions with exactly two VR track jets associated to the large-R jet, and the ones with three-or-more associations (this event selection will be detailed in the Section 8.5.1.). In the Boosted analysis, this is because the second b -quark from the second top often lies outside of the large-R jet.

8.3.3.3 Single top

Events that only contain a single top can also enter the analysis. Mostly this is due to hadronic decays of quarks or W-bosons in the final state being misidentified as b -quarks. The four leading order single top diagrams are shown in Figure 8.14.




Figure 8.14: Image showing the four tree-level single top Feynman diagrams. The top diagrams are weak processes referred to as s -channel (top left) and t -channel (top right). The bottom two are irreducible (can't be distinguished from each other) 'Wt-channel' diagrams. Particles misidentified as b -jets are labelled as " b ".

Single top diagrams enter the analyses via the 0-lepton and 1-lepton channels. Since the 's-channel' and 't-channel' diagrams are completely weak-force processes, their vertices are more suppressed and thus are less likely to happen. In the boosted analysis, for example, the s - and t -channel single top diagrams only contribute to 5% of the total single top background.

The dominant Wt-channel (and t -channel for the Resolved analysis) backgrounds have uncertainties for the normalisation, acceptance and shape for the p_T^V and m_{bb} distribu-

tions. The s-channel only has normalisation uncertainties since its contribution is small.

The generators that produce single top (stop) events do so not only using tree-level diagrams like those in Figure 8.14, but at higher orders as well (NLO). The addition of another vertex in the stop diagrams means there is an overlap between diagrams for $t\bar{t}$ and stop. This is usually known as interference and can affect the yield of stop in the analysis.

The nominal samples attempt to solve this using a technique called diagram removal, which removes stop diagrams at NLO that are that enter $t\bar{t}$ at LO. This method can be problematic as the process is no longer gauge invariant. To try and combat this, there is an additional stop uncertainty. This uncertainty compares the nominal sample with an alternative sample but uses a different technique of generating stop events: diagram subtraction. This process is more complicated but in a nutshell, it removes $t\bar{t}$ events at the cross-section level [145, p3].

8.3.3.4 Multi-jet

Both analyses have a non-negligible contribution from events with multiple QCD jets. In both analyses, the estimation on how many events impact the analysis is done by creating a template using data in a multi-jet sideband. Figure 8.15 shows an example of a Feynman process that could produce an analysis signal.

8.3.3.5 Diboson

The last background of note in the analysis is the diboson background. Here there are three separate process being considered: $ZZ(\rightarrow b\bar{b} + leptons)$, WZ and WW . The last of these is the least likely to create two b -jets and is thus the weakest contribution. Figure 8.16 shows the Feynman diagram for the two dominant diboson processes.

The only difference between this background and the signal is the replacement of the




Figure 8.15: Multi-jet background example Feynman diagram. Here a gluon emission from a jet has been misidentified by the detector as a lepton.




Figure 8.16: Feynman diagram of the two leading diboson contributions to the background.

Higgs boson with a Z-boson. When this ‘virtual’ Z-boson is about the same mass as the Higgs boson, this background is irreducible. It is important to understand this background and hence both analyses have a cross-check where they replace the signal with diboson events to see if the results and the systematics are not biased.

The WZ and ZZ contributions have overall normalisation systematics, comparisons between the p_T^V and m_{bb} shapes, and an acceptance uncertainty between regions. The WW contribution only has a normalisation uncertainty.

8.3.4 Common workflow

Since the two analyses share a lot of systematics and backgrounds, it is useful for them to generate their results on a similar *framework*.

An analysis will typically have two main dataflows: one for theory and one for measurement. The theory dataflow loosely takes the most current understanding of the physics and generates MC simulated events and looks a bit like this: Theory $\rightarrow$ Event Generation $\rightarrow$ Simulation $\rightarrow$ Trigger $\rightarrow$ Reconstruction $\rightarrow$ Analysis. The measurement dataflow takes the data from experiments and turns it into a format where questions can be asked of it. It looks something like this: Detector $\rightarrow$ Trigger $\rightarrow$ Reconstruction $\rightarrow$ Analysis.

An analysis framework is a series of programs, packages and code that processes pre-prepared data and MC simulated events simultaneously through the final stages of their dataflows to produce results.

The data from the ATLAS detector (or Monte Carlo simulated data) is processed centrally with the latest reconstruction and calibrations to create objects called Analysis

Object Data (xAOD's)¹². These xAOD's are typically large and have lots of processes that are not of interest to particular set of analyses. These xAOD's are then run through the Derivation Framework to reduce their size and produce 'derivations' or DxAOD's. These DxAOD's are designed to be shared by several analyses. It is at this derivation stage that individual analyses will diverge in the way that these derivations are treated.

The main workflow on how an analysis is done in the VHbb Analysis Group using ATLAS data is represented in Figure 8.17.

The analysis uses two main packages to create results: the $VH, H \rightarrow b\bar{b}$ CxAOD framework CxAODFramework_VHbb; and the $VH, H \rightarrow b\bar{b}$ workspace maker WSMaker_VHbb. The CxAODFramework_VHbb (CxAODFramework) starts by running a package known as the CxAODMaker (*Maker*) on the DxAOD's to create smaller objects called CxAOD's, which is what an individual analysis will be run on. The Maker is designed to trim the derivations such that the resulting CxAOD's include some cuts common to all sub-analyses in the group. Studies that are performed within the analysis will run various configurations of another program in the framework known as the CxAODReader (*Reader*). The Reader turns CxAODs into ROOT Ntuples and histograms. The Reader is a steered C++ encoded file which contains more specific cuts and object definitions, suited for the specific analysis designed to run on.

As each lepton channel has different problems to solve, it is useful to split the analyses into three groups, depending on the decay of the vector boson: 0-lepton, 1-lepton and 2-lepton analyses. In practice this means that each of these subgroups use different derivations but are interested in the same final state, and each can run the framework independently on one another, effectively creating three mini-analyses that can be individually optimised.

¹²AOD's are not ROOT readable, but xAOD's are ROOT readable. They contain the same information.




Figure 8.17: Picture showing the rough workflow of the analysis. Black lines represent data, blue lines represent MC simulations, red lines represent systematics. Green curved lines are packages in the analysis. Detector-based systematics are propagated to Monte Carlo Simulations into the xAOD format and end up being stored in CxAOD's as 'shallow copies'.

The majority of the work presented in this thesis comes from work undergone in the 0-lepton and 2-lepton analyses which means primarily, I am interested in the Higgs candidate being produced alongside a Z-boson.

The WSMaker_VHbb (WSMaker) is the tool that is in charge of creating the main and auxiliary fit diagnosis plots by performing a statistical analysis of the histograms it gets fed. This allows the yields and kinematics of both the signal and the background processes to be estimated and later changed if required.

The ultimate aim of the fit is to provide the signal strengths for each STXS bin. From this, limits on the STXS coefficients on the Lagrangian operators that affect both the VH production and the $H \rightarrow b\bar{b}$ decay can be found.

The functionality of the WSMaker and a walk through its outputs will be covered in the next chapter which covers statistics.

A large part of working on the analysis is defined by working on the framework that will produce inputs for the main analysis run. Most of this comes down to editing either the Maker or the Reader for Analysis Framework to produce specific CxAOD sets on which to perform physics studies. Systematic uncertainties and physics object definition improvements are looked at and the results feed back into the framework in a cycle until this round of analysis is complete.

8.4 The VHbb Resolved analysis - 2020

The Resolved analysis is a direct continuation of the 2019 STXS analysis. Like its predecessors, it uses boosted decision tree multivariate discriminants built from the event-

level variables to maximise the sensitivity to the Higgs boson signal.

The CxAODFramework creates output BDT distributions which are used as the inputs into the WSMaker which creates the binned maximum-likelihood fit, referred to as the global likelihood fit.

Apart from running on a larger dataset, several improvements have been made to the analysis itself, like: the formation of new control regions; the addition of new information into the BDT, and the new event categorisation to better align to STXS regions. These will be touched upon in later sections.

8.4.1 Event selection and categorisation

The event selection used in this analysis is summarised in Table 8.3.

Table 8.3: Table showing the summary of the event selection and categorisation in the Resolved analysis 0-, 1- and 2-lepton channels. Table from [3, p9].

Selection	0-lepton	1-lepton		2-lepton
		e sub-channel	μ sub-channel	
Trigger	$E_{\text{T}}^{\text{miss}}$	Single lepton	$E_{\text{T}}^{\text{miss}}$	Single lepton
Leptons	0 <i>loose</i> leptons	Exactly 1 <i>tight</i> electron 0 additional <i>loose</i> leptons $p_{\text{T}} > 27$ GeV	Exactly 1 <i>tight</i> muon 0 additional <i>loose</i> leptons $p_{\text{T}} > 25$ GeV	Exactly 2 <i>loose</i> leptons $p_{\text{T}} > 27$ GeV Same-flavour Opposite-sign charges ($\mu\mu$)
$E_{\text{T}}^{\text{miss}}$	> 150 GeV	> 30 GeV	–	–
$m_{\ell\ell}$	–	–	–	$81 \text{ GeV} < m_{\ell\ell} < 101 \text{ GeV}$
Jet p_{T}		> 20 GeV for $ \eta < 2.5$ > 30 GeV for $2.5 < \eta < 4.5$		
b -jets		Exactly 2 b -tagged jets		
Leading b -tagged jet p_{T}		> 45 GeV		
Jet categories	Exactly 2 / Exactly 3 jets	Exactly 2 / Exactly 3 jets		Exactly 2 / ≥ 3 jets
H_{T}	> 120 GeV (2 jets), > 150 GeV (3 jets)	–		–
$\min[\Delta\phi(E_{\text{T}}^{\text{miss}}, \mathbf{jets})]$	$> 20^\circ$ (2 jets), $> 30^\circ$ (3 jets)	–		–
$\Delta\phi(E_{\text{T}}^{\text{miss}}, \mathbf{b}\mathbf{b})$	$> 120^\circ$	–		–
$\Delta\phi(\mathbf{b}_1, \mathbf{b}_2)$	$< 140^\circ$	–		–
$\Delta\phi(E_{\text{T}}^{\text{miss}}, \mathbf{p}_{\text{T}}^{\text{miss}})$	$< 90^\circ$	–		–
	–	–		$75 \text{ GeV} < p_{\text{T}}^V < 150 \text{ GeV}$
p_{T}^V regions	$150 \text{ GeV} < p_{\text{T}}^V < 250 \text{ GeV}$	$150 \text{ GeV} < p_{\text{T}}^V < 250 \text{ GeV}$		$150 \text{ GeV} < p_{\text{T}}^V < 250 \text{ GeV}$
	$p_{\text{T}}^V > 250 \text{ GeV}$	$p_{\text{T}}^V > 250 \text{ GeV}$		$p_{\text{T}}^V > 250 \text{ GeV}$
Signal regions	$\Delta R(\mathbf{b}_1, \mathbf{b}_2)$ signal selection			
Control regions	High and low $\Delta R(\mathbf{b}_1, \mathbf{b}_2)$ side-bands			

8.4.1.1 0-lepton selection

As the 0-lepton event signature requires the Z-boson to decay into two neutrinos, events that would fall in this category are triggered with E_T^{miss} triggers. The minimum energy threshold for the E_T^{miss} varies with the data-taking period, starting off at 70 GeV in 2015 and reaching 110 GeV in 2018.

The efficiency of these triggers is tested in $t\bar{t}$ and $V + jets$ events alternatively triggered with single-muon triggers because the energy of the muon(s) are not featured into the calculation of E_T^{miss} ¹³. The resulting correction factor is about 0.95 for events with 150 GeV rising to a plateau of unity at around 200 GeV.

Multi-jet contributions enter the 0-lepton channel when they have large E_T^{miss} . This happens when jets in calorimeters fail to be measured. These events are mostly removed via additional requirements on E_T^{miss} , jets and p_T^{miss} .

Since the only detectable objects in these events are jets, it is desirable for the trigger efficiency to be independent of the number of jets in the event. This phase space where this is not possible is small, and is removed by putting a requirement on the transverse momenta of the Higgs jets (H_T).

0-lepton events with any loose leptons are discarded.

8.4.1.2 1-lepton selection

The 1-lepton channel is split into two sub-channels: an electron (e) sub-channel and a muon (μ) sub-channel. This is done because the two channels use different triggers.

The e sub-channel uses single-lepton triggers, with electrons required to pass one of a

¹³a design feature that is exploited in Chapter 10

list of possible triggers. The minimum energy threshold placed on electrons is 24 GeV in 2015 data rising to 26 GeV from 2016. Since electrons are more commonly created as background, the e sub-channel has an additional $E_T^{\text{miss}} > 30$ GeV requirement to reduce the background from multi-jets.

The μ sub-channel uses the same E_T^{miss} triggers and corrections as the 0-lepton channel. This is because the single lepton triggers in the analysis are less efficient in this phase space as the coverage of the muon system in the central region of the detector is limited.

As this channel aims to detect one lepton from a W-boson, a tight selection is required of said lepton, and events with additional loose (or tight) leptons are rejected.

8.4.1.3 2-lepton selection

The 2-lepton channel uses the same single-lepton trigger regime used in the 1-lepton e sub-channel for both ee and $\mu\mu$ events.

Exactly two loosely tagged leptons of the same flavour are required for events to be kept and because leptons are usually well-measured, the two leptons are required to have a combined invariant mass close to a Z-boson. Events with two muons have the additional stipulation of having to have opposite-sign charges. This is done because the particles live long enough for track curvature direction to be accurately assigned and the opposing charges will bend in different directions under the barrel and toroidal magnetic fields.

8.4.1.4 Signal regions

To align with the STXS regions, the analysis decided to split the $p_T^V < 150$ GeV region in each of the channels into a $150 \text{ GeV} \leq p_T^V < 250 \text{ GeV}$ region and a $p_T^V \geq 250 \text{ GeV}$ region for each jet category. Adding to the existing $75 \text{ GeV} \leq p_T^V < 150 \text{ GeV}$ in 2-lepton, there

are a total of 14 signal regions.

8.4.1.5 Control regions

Also new to this iteration of the analysis, each of the 14 signal regions have additional p_T^V -dependent requirement on the ΔR_{bb} , creating a lower and upper ΔR_{bb} CR sideband for every signal region.

Figure 8.18 shows the distribution of the event yield when these sidebands are created. These sidebands are rich with either $V + HF$ or $t\bar{t}$ events. The same ΔR_{bb} selections are made for all three channels and 93% of the signal is kept in all 2-jet signal regions, 81% in all 3-jet signal regions, and 68% in the 3-or-more-jet signal region.




Figure 8.18: Distribution of the signal yield under the radial difference between the two leading jets for (a) 2-jet and (b) 3-jet 1-lepton channel. The dark lines show the distinction between the High ΔR CR, SR and the Low ΔR CR. Image from [3, p11].

In addition to this there are the CR's that were present in the previous analysis. These are made in a particular channel by having largely the same event selection as the nominal one for that channel, but having a few key requirements inverted.

Top $e\mu$ -CR In the 2-lepton channel, the Top (stop + $t\bar{t}$) CR is made by 'reversing' the requirement that the leptons have the same flavour (to require an electron and a

muon) and by removing the requirement that the charges have to have the opposing sign.

MJ CR In the 1-lepton channel, the MJ (multi-jet) CR is made by inverting the energy requirement on the tight lepton and reducing the number of required b-jets from 2 to 1.

These control regions are not put into the fit, but BDT's from the 14 SR regions and the event yields from the other 28 ΔR_{bb} CR's are used as inputs in the global likelihood fit. BDT's that are generated in the main analysis are referred to as BDT_{VH} .

8.4.2 The analysis BDT

The variables that go into the BDT are different for each lepton channel. Table 8.4 shows a comprehensive list of all the variables and the lepton channels they are included in.

- m_{bb} is the invariant mass of the $b\bar{b}$ system.
- ΔR_{bb} is the angular distance between the two leading b -tagged jets.
- p_T^V is the transverse momentum of the weak boson.
- $\Delta\eta(\mathbf{b}_1, \mathbf{b}_2)$ is the angle subtended by the leading b -jets in the plane marked by transverse plane and the beamline (psuedorapidity separation).
- $\Delta\eta(\mathbf{V}, \mathbf{bb})$, $\Delta\phi(\mathbf{V}, \mathbf{bb})$ and $\Delta y(\mathbf{V}, \mathbf{bb})$ are respectively the the psuedorapidity separation, azimuthal angle and the rapidity difference between the vector boson and the two b -tagged jets which are used as a proxy for the Higgs boson system.
- E_T^{miss} is the negative scalar sum of all the detectable objects in the event. In the 0-lepton channel this is exactly equivalent to p_T^V .
- m_{eff} is the scalar sum of transverse momenta of all the jets (H_T) and the E_T^{miss} .
- $\min[\Delta\phi(\mathbf{l}, \mathbf{b})]$ is the angle between the lepton and the closest b -tagged jet.
- m_T^W is the transverse mass of the W-boson.
- m_{top} is the reconstructed top quark mass.

Variable	0-lepton	1-lepton	2-lepton
m_{bb}	×	×	×
$\Delta R(\mathbf{b}_1, \mathbf{b}_2)$	×	×	×
$p_T^{b_1}$	×	×	×
$p_T^{b_2}$	×	×	×
$p_T^V \equiv E_T^{\text{miss}}$	×	×	×
$\Delta\phi(\mathbf{V}, \mathbf{bb})$	×	×	×
$MV2(b_1)$	×	×	
$MV2(b_2)$	×	×	
$ \Delta\eta(\mathbf{b}_1, \mathbf{b}_2) $	×		
m_{eff}	×		
$p_T^{\text{miss, st}}$	×		
E_T^{miss}	×	×	
$\min[\Delta\phi(\ell, \mathbf{b})]$		×	
m_T^W		×	
$ \Delta y(\mathbf{V}, \mathbf{bb}) $		×	
m_{top}		×	
$ \Delta\eta(\mathbf{V}, \mathbf{bb}) $			×
$E_T^{\text{miss}}/\sqrt{S_T}$			×
$m_{\ell\ell}$			×
$\cos\theta(\ell^-, \mathbf{Z})$			×
Only in 3-jet events			
$p_T^{\text{jet}_3}$	×	×	×
m_{bbj}	×	×	×

Table 8.4: Table showing the variables used for the multivariate BDT discriminant in each of the lepton channels. Table from [3, p15].

- $E_T^{\text{miss}}/\sqrt{S_T}$ is a 2-lepton-specific variable. $\sqrt{S_T}$ is the scalar sum of the transverse energies of all the charged leptons and jets in the event.
- $m_{\ell\ell}$ is the invariant mass of the dilepton system.
- $p_T^{\text{jet}_3}$ is the transverse momentum of the third jet in the system (if there is one).
- m_{bbj} is the invariant mass of the $b\bar{b}j$ system in the event that there are three jets and the third jet is labelled as ‘j’.

All of the BDT inputs above, have been used in the previous iterations of the analysis, the following new variables are also input to the BDTs.

- **MV2(b1)** and **MV2(b2)** is the MV2c10 score for the leading (b1) and sub-leading (b2) b -tagged jets in the event. This inclusion provides additional rejection in the 0-lepton and 1-lepton channels where c -jets and light-jets are more likely to be misidentified as b -jets. This is especially the case in $W \rightarrow cq$ decays in $t\bar{t}$ and stop- Wt backgrounds. The sensitivity in the 1-lepton (0-lepton) channel is improved by $\sim 10\%$ ($\sim 7\%$). No additional sensitivity is gained from using this variable in the 2-lepton channel, as the backgrounds are mostly dominated by processes containing two real b -jets.

- $p_T^{\text{miss,st}}$ is the magnitude of the E_T^{miss} soft-term (as mentioned in Section 7.4), $p_T^{\text{miss,st}}$. For the 0-lepton channel this provides additional rejection against $t\bar{t}$, which may contain leptons or b -jets that go unreconstructed, due to kinematic and detector acceptance. This typically means that a $t\bar{t}$ event will have a larger p_T^{miss} than signal events. This improves the sensitivity in the 0-lepton channel by $\sim 2\% - 3\%$.

- $\cos\theta(l^-, Z)$ is known as the Z-boson polarisation, and is the cosine of the polar angle between the Z-boson's direction of travel in the lab frame, and the plane of the leptonic decay products in the Z-boson's rest frame. The polarisations of Z-bosons from the signal process are expected to be different to those from the dominant Z+jets background, thus providing additional background rejection in the 2-lepton channel. The inclusion of this variable in the BDT improves the 2-lepton channel sensitivity by $\sim 7\%$.

8.4.2.1 BDT_R

Also new to this analysis is the process of BDT re-weighting or BDT_R . This is where the BDT's are created using variables from both the nominal samples and the alternative samples. Since the event generators often use different models in the hadroniza-

tion/showering stages of event generation, it is good to have samples that generate events differently to the one used to generate the nominal prediction. The prescriptions can be compared and contrasted to see how the model copes with slightly different definitions of the events that make up the samples.

The ratio of the BDT scores in each bin provides a re-weighting function which can be used to correct the nominal sample to match the alternative sample or vice versa.

The BDT_R technique is used to produce shape uncertainties for the p_T^V distribution in 1-lepton and 0-lepton $t\bar{t}$ and $W + HF$. For the former, the matrix element and the parton shower are varied separately.

8.4.3 Additional points on background modelling

Since $t\bar{t}$ has multiple compositions of jet flavours, there are also sets of flavour uncertainties for each. They are the differences in the ratio of the relative composition of these flavours with respect to the dominant bb , akin to the treatment of the uncertainties under $V + HF$ (i.e. $(t\bar{t} \rightarrow)bc/(t\bar{t} \rightarrow)bb$. But bl/bb and cc/bb are combined to make $other/bb$). The difference here is that the ratio comparison is done between the nominal samples and the alternative matrix element and parton shower generator samples separately.

8.4.3.1 Data-driven $t\bar{t}$ estimation

In the 2-lepton channel, the $t\bar{t}$ was modelled previously using simulated event samples. However the selection that creates the top $e\mu$ -CR in the 2-lepton analysis creates a region with over 99% contribution from $t\bar{t}$ and stop- Wt events. The differing-flavour Top events present in the top $e\mu$ -CR are thus used to model the same-flavour lepton top background that enters into the nominal signal selection. Both shape and normalisation effects are modelled.

Biases with this technique are determined by a yield comparison between top background events in the SR with those in the top $e\mu$ -CR. With said ratio of the yields being 1.00 ± 0.01 (1.01 ± 0.01) for the 2-jet (≥ 3 -jet) region, no such biases are observed. The uncertainty in the ratio is purely a statistical uncertainty resulting from the simulated samples.

With this method, all the experimental and theoretical uncertainties are eliminated as none of them creates any biases in the ratio beyond the statistical one. This means that the data statistics in the top $e\mu$ -CR become the dominant source of uncertainty for the 2-lepton top background estimate.

8.4.4 Systematic uncertainties

Systematic uncertainties are determined mostly by performing fitting studies in the framework. This section will briefly cover the ones that have not already been mentioned and show summary tables of their final values.

8.4.4.1 Background Systematic uncertainties

Tables 8.5 and 8.6 shows a summary of the final background uncertainties for the Resolved analysis, and each contribution has been discussed previously.

8.4.4.2 Experimental uncertainties

The experimental uncertainties that have the largest impact on the analysis are b -tagging correction factors, jet energy scale calibrations and the jet energy resolution modelling. The latter two are determined centrally by the ATLAS collaboration [146] (Jet Combined Performance Groups) for use by the relevant ATLAS analyses. From these, the latest recommendations¹⁴ are incorporated into the analysis.

¹⁴Those that reach the analysis before a certain date at least.

Table 8.5: Summary of the systematic uncertainties on the main backgrounds in the analysis. If a systematic has a ‘S’ then it is a shape-only uncertainty in the signal region. ‘S+M’ represents a shape uncertainty that allows a migration of events to other regions. Table from [3, p19].

Z + jets	
Z + ll normalisation	18%
Z + cl normalisation	23%
Z + HF normalisation	Floating (2-jet, 3-jet) $\times$ (75 GeV < p_T^V < 150 GeV, p_T^V > 150 GeV)
Z + bc -to-Z + bb ratio	30%–40%
Z + cc -to-Z + bb ratio	13%–16%
Z + bl -to-Z + bb ratio	20%–28%
SR-to-low ΔR CR ratio	3.8%–9.9% (75 GeV < p_T^V < 150 GeV, p_T^V > 150 GeV)
SR-to-high ΔR CR	2.7%–4.1% (75 GeV < p_T^V < 150 GeV, p_T^V > 150 GeV)
0-to-2 lepton ratio	7%
p_T^V	M+S (75 GeV < p_T^V < 150 GeV, p_T^V > 150 GeV)
m_{bb}	S (75 GeV < p_T^V < 150 GeV, p_T^V > 150 GeV)
W + jets	
W + ll normalisation	32%
W + cl normalisation	37%
W + HF normalisation	Floating (2-jet, 3-jet)
W + bc -to-W + bb ratio	15% (0-lepton) and 30% (1-lepton)
W + cc -to-W + bb ratio	10% (0-lepton) and 30% (1-lepton)
W + bl -to-W + bb ratio	26% (0-lepton) and 23% (1-lepton)
SR-to-CR ratio	3.6%–15%
0-to-1 lepton ratio	5%
p_T^V	M+S (2-jet, 3-jet)
R _{BDT}	S
$t\bar{t}$ (0+1-lepton channels only)	
$t\bar{t}$ normalisation	Floating (2-jet, 3-jet)
0-to-1 lepton ratio	8%
$t\bar{t}$ (flavour composition) bc -to- bb ratio (ME)	7.6%–8.2% (0-lepton), 1.3%–3.8% (1-lepton)
$t\bar{t}$ (flavour composition) bc -to- bb ratio (PS)	2.1%–3.2% (0-lepton), 1.5%–7.1% (1-lepton)
$t\bar{t}$ (flavour composition) other-to- bb ratio (ME)	2.8%–6.4% (0-lepton), 3.3%–5.7% (1-lepton)
$t\bar{t}$ (flavour composition) other-to- bb ratio (PS)	5.6%–13% (0-lepton), 0.3%–2.1% (1-lepton)
p_T^V	M+S (2-jet, 3-jet)
R _{BDT} ME variation	M+S (2-jet, 3-jet)
R _{BDT} PS variation	M+S (0-lepton, 1-lepton)
Single top quark	
Cross-section	4.6% (s -channel), 4.4% (t -channel), 6.2% (Wt)
Acceptance 2-jet	17% (t -channel), 55% ($Wt(bb)$), 24% ($Wt(other)$)
Acceptance 3-jet	20% (t -channel), 51% ($Wt(bb)$), 21% ($Wt(other)$)
m_{bb}	M+S (t -channel, $Wt(bb)$, $Wt(other)$)
p_T^V	M+S (t -channel, $Wt(bb)$, $Wt(other)$)
Multi-jet (1-lepton)	
Normalisation	30%–200% (2-jet), 100% (3-jet)
BDT template	M+S

Table 8.6: Tabulated summary of the systematic uncertainties on of the diboson background. If a systematic has a ‘M+S’, it represents a shape uncertainty that allows a migration of events to other regions. Parentheses for any uncertainty denotes where an uncertainty is independently varied. Table from [3, p20].

<i>ZZ</i>	
Normalisation	20%
0-to-2 lepton ratio	6%
Acceptance from scale variations	10%–18%
Acceptance from PS/UE variations for 2 or more jets	6%
Acceptance from PS/UE variations for 3 jets	7% (0-lepton), 3% (2-lepton)
m_{bb} from scale variations	M+S (correlated with WZ uncertainties)
p_T^V from scale variations	M+S (correlated with WZ uncertainties)
m_{bb} from PS/UE variations	M+S (correlated with WZ uncertainties)
p_T^V from PS/UE variations	M+S (correlated with WZ uncertainties)
m_{bb} from matrix-element variations	M+S (correlated with WZ uncertainties)
<i>WZ</i>	
Normalisation	26%
0-to-1 lepton ratio	11%
Acceptance from scale variations	13%–21%
Acceptance from PS/UE variations for 2 or more jets	4%
Acceptance from PS/UE variations for 3 jets	11%
m_{bb} from scale variations	M+S (correlated with ZZ uncertainties)
p_T^V from scale variations	M+S (correlated with ZZ uncertainties)
m_{bb} from PS/UE variations	M+S (correlated with ZZ uncertainties)
p_T^V from PS/UE variations	M+S (correlated with ZZ uncertainties)
m_{bb} from matrix-element variations	M+S (correlated with ZZ uncertainties)
<i>WW</i>	
Normalisation	25%

The b -tagging correction factors are the differences between the efficiency of the b -tagging algorithm in data and simulation. They are evaluated separately for b -jets, c -jets and light-jets in five bins of the value of the MV2 discriminant.

These correction factors have estimated uncertainties from multiple sources, and are decomposed¹⁵ into many uncorrelated independent components. Components that have a negligible impact are then removed. The analysis then has 29 uncertainties for b -jets, 18 for c -jets and 10 for light-jets.

8.4.4.3 Signal uncertainties

The values of the signal uncertainties are summarised in Table 8.7.

Similar to the estimation of the background, the uncertainties in the m_{bb} and p_T^V distributions are determined from variations in an alternate sample. Unlike the backgrounds however, there are additional theoretical systematics from the proton PDF and the scale/strength of the strong force α_s (combined to form PDF+ α_s), and the higher-order electro-weak corrections to the signal process.

Table 8.7: Summary of the systematic uncertainties on the signal. An ‘M+S’ represents a shape uncertainty that allows a migration of events to other regions. Parentheses for any uncertainty denotes where an uncertainty is independently varied. The $H \rightarrow b\bar{b}$ branching fraction error is the error on the Standard Model prediction. Table from [3, p21].

Signal	
Cross-section (scale)	0.7% (qq), 25% (gg)
$H \rightarrow b\bar{b}$ branching fraction	1.7%
Scale variations in STXS bins	3.0%–3.9% ($qq \rightarrow WH$), 6.7%–12% ($qq \rightarrow ZH$), 37%–100% ($gg \rightarrow ZH$)
PS/UE variations in STXS bins	1%–5% for $qq \rightarrow VH$, 5%–20% for $gg \rightarrow ZH$
PDF+ α_s variations in STXS bins	1.8%–2.2% ($qq \rightarrow WH$), 1.4%–1.7% ($qq \rightarrow ZH$), 2.9%–3.3% ($gg \rightarrow ZH$)
m_{bb} from scale variations	M+S ($qq \rightarrow VH$, $gg \rightarrow ZH$)
m_{bb} from PS/UE variations	M+S
m_{bb} from PDF+ α_s variations	M+S
p_T^V from NLO EW correction	M+S

¹⁵This is elaborated on in the following chapter.

In each STXS bin, acceptance uncertainties are evaluated to account for possible migration of events and correlations between STXS regions. These acceptance uncertainties are evaluated for energy scale variations, alternative parton shower or underlying event (PS/UE) models and PDF+ α_s .

In addition to these, there are theoretical uncertainties taken from the recommendations of the LHC Higgs Cross Section Working Group [59, 147, 148] for the calculations of the VH production cross-section and the $H \rightarrow b\bar{b}$ decay.

8.4.5 Analysis methodology

The fit is fed the 14 SR BDT_{VH} created by the framework and from the previous STXS binning shown in Figure 8.9, three different STXS ‘POI’ configurations are studied.

- 1) The single-POI fit is done to make a new measurement of μ_{VH}^{bb} which informs the VH production cross-section and the $H \rightarrow b\bar{b}$ branching fraction.
- 2) A two-POI fit to measures separate ZH , $H \rightarrow b\bar{b}$ and WH , $H \rightarrow b\bar{b}$ components.
- 3) A 5-POI fit as seen previously in Figure 8.10a is performed to measure the signal cross-section ($\times$ the $H \rightarrow b\bar{b}$ and $V \rightarrow ll$ branching fractions) in each STXS bin.

8.4.6 Cross check analyses

Like its predecessors, the 2020 Resolved analysis has two cross-checks to validate the methodology of the main analysis. They use the same framework but different distributions are put into the fit.

8.4.6.1 Diboson Analysis

The diboson analysis modifies the nominal multivariate analysis to extract the VZ , $Z \rightarrow b\bar{b}$ diboson process. To this end BDT_{VZ} objects are used in place of BDT_{VH} .

Normalisation uncertainties on the diboson process are removed, and the overall normalisations are floated. The VH , $H \rightarrow b\bar{b}$ signal enters this version of the analysis as a background and is normalised to the production rate predicted by the Standard Model. Here only two POI configurations are evaluated.

- 1) A single-POI fit is done to make a measurement of μ_{VZ}^{bb} which informs the VZ production cross-section and the $Z \rightarrow b\bar{b}$ branching fraction.
- 2) A two-POI fit to measure separate ZZ , $Z \rightarrow b\bar{b}$ and WZ , $Z \rightarrow b\bar{b}$ components.

8.4.6.2 Cut-Based (Dijet mass) analysis: CBA

The Higgs boson signal measurement is also cross-checked with a dijet-mass analysis, where the signal yield is measured using the mass of the dijet system (m_{bb}) as the main observable entered into the fit instead of the multivariate discriminant BDT_{VH} .

Here only the single-POI fit is done to make a measurement of μ_{VH}^{bb} .

8.5 The VHbb Boosted analysis - 2020

The Boosted analysis mainly deals with events where the Higgs is so energetic that the two b-quarks from the decay are not able to be resolved efficiently as two separate jets but form a single ‘fat’ jet (or a large- R jet). The increase in energy of a particle is referred to as ‘boosting’. Figure 8.19 shows this in the Feynman diagram representation.

Requiring larger energies of the Higgs Boson means that the number of events that fill this requirement are lower, meaning the Boosted analysis has lower statistics than the Resolved analysis. As a new proof-of-concept analysis, the Boosted analysis will also not feature a BDT and all the complexities that it brings. The main analysis will be a cut-based selection based on a fit of the mass of the large- R jet: m_J .




Figure 8.19: Feynman diagram showing the full Higgsstrahlung process of Higgs production at the LHC, and its decay into the final states of interest to the Boosted analysis. The purple cone is a representation of what a boosted signal event looks like. One large-radius fat jet with two objects tagged as b -jets inside it.

8.5.1 Event selection and categorisation

The event selection used in this analysis is summarised in Table 8.8. It shares some common features with the Resolved analysis, such as all the requirements on leptons that enter each sub-channel, and for the most part, the triggers that are used on these leptons.

The Boosted analysis requires both the transverse momentum of the vector boson to be larger than 150 GeV and at least one large- R jet with $p_T > 250$ GeV and $|\eta| < 2$ to be present in all events. The large- R jet with the largest p_T is referred to as the ‘Higgs-jet candidate’ or the H_{cand} . This large- R jet needs to have at least two track-jets associated to it, with the most energetic two needing to be b -tagged. Once the large- R jet has been established, the jet algorithms are run again with $R=0.4$ to produce small- R jets in the event which are checked to see if they are topologically consistent with the large- R jet.

8.5.1.1 0-lepton selection

Like the Resolved analysis, events that fall into the 0-lepton selection are triggered with E_T^{miss} triggers with the same year-dependent minimum energy thresholds. However be-

Table 8.8: Table showing the summary of the event selection and categorisation in the Boosted analysis 0-, 1- and 2-lepton channels. Baseline electrons are ‘loose’ quality electrons with additional requirements on track parameters and are the same as the Loose leptons in the Resolved analysis. Signal leptons are a subset of baseline electrons with tighter requirements and are the same as the Tight leptons from the Resolved analysis. Table from [4, p8].

Selection	0 lepton channel	1 lepton channel		2 leptons channel	
		<i>e</i> sub-channel	μ sub-channel	<i>e</i> sub-channel	μ sub-channel
Trigger	$E_{\text{T}}^{\text{miss}}$	Single electron	$E_{\text{T}}^{\text{miss}}$	Single electron	$E_{\text{T}}^{\text{miss}}$
Leptons	0 <i>baseline</i> leptons	1 <i>signal</i> lepton $p_{\text{T}} > 27 \text{ GeV}$ $p_{\text{T}} > 25 \text{ GeV}$ no second <i>baseline</i> lepton		2 <i>baseline</i> leptons among which ≥ 1 <i>signal</i> lepton, $p_{\text{T}} > 27 \text{ GeV}$ both leptons of the same flavour - opposite sign muons	
$E_{\text{T}}^{\text{miss}}$	> 250 GeV	> 50 GeV	-	-	
p_{T}^V	$p_{\text{T}}^V > 250 \text{ GeV}$				
Large- <i>R</i> jets	at least one large- <i>R</i> jet, $p_{\text{T}} > 250 \text{ GeV}$, $ \eta < 2.0$				
Track-jets	at least two track-jets, $p_{\text{T}} > 10 \text{ GeV}$, $ \eta < 2.5$, associated to the leading large- <i>R</i> jet				
<i>b</i> -jets	leading two track-jets associated to the leading large- <i>R</i> must be <i>b</i> -tagged (MV2c10, 70%)				
m_{J}	> 50 GeV				
$\min[\Delta\phi(E_{\text{T}}^{\text{miss}}, \text{small-}R \text{ jets})]$	> 30°	-			
$\Delta\phi(E_{\text{T}}^{\text{miss}}, H_{\text{cand}})$	> 120°	-			
$\Delta\phi(E_{\text{T}}^{\text{miss}}, E_{\text{T, trk}}^{\text{miss}})$	< 90°	-			
$\Delta y(V, H_{\text{cand}})$	-	$ \Delta y(V, H_{\text{cand}}) < 1.4$			
$m_{\ell\ell}$	-			66 GeV < $m_{\ell\ell}$ < 116 GeV	
Lepton p_{T} imbalance	-			$(p_{\text{T}}^{\ell_1} - p_{\text{T}}^{\ell_2})/p_{\text{T}}^Z < 0.8$	

cause the efficiency of these triggers was found to plateau to unity at around 200 GeV, and the lowest p_T^V region to enter the analysis starts at 150 GeV, the channel does not implement any correction factor.

Similar to the Resolved analysis 0-lepton selection, possible multi-jet contributions are mitigated by further cuts on the boosted equivalent of the Resolved analysis objects (E_T^{miss} , jets, p_T^{miss}).

8.5.1.2 1-lepton selection

The Boosted analysis 1-lepton trigger selection is identical to the Resolved analysis one. It has a *e* sub-channel which is triggered by single-lepton triggers and a μ sub-channel triggered by E_T^{miss} triggers.

$t\bar{t}$ and W+jets backgrounds are further reduced by applying a requirement that the rapidity difference between the vector boson and the H_{cand} is less than 1.4.

8.5.1.3 2-lepton selection

The Boosted analysis trigger strategy in the 2-lepton channel is exactly identical to the 1-lepton one. The reason behind the move from exclusively using single-lepton triggers to a combination of single lepton and E_T^{miss} triggers is the basis of my contribution to the Boosted analysis. This study is detailed in Chapter 10 of this thesis.

Like the 1-lepton selection, the 2-lepton selection has a maximum rapidity difference requirement between the vector boson and the H_{cand} , but the backgrounds it is designed to reduce are Z+jets ones.

The event topology of the Z+jet backgrounds is also expected to be different to that of the signal due to the polarisation of the Z-boson. In the Resolved analysis this is exploited by the introduction of the variable $\cos\theta(l^-, Z)$ into the BDT. In the Boosted analysis this difference is exploited by the inclusion of a requirement on the *lepton imbalance* $((p_T^{l1} - p_T^{l2})/p_T^Z)$ to the event selection. Like the Resolved analysis, the 2-lepton channel also has a requirement that the invariant mass of the two leptons be around that of the Z-boson.

8.5.1.4 Signal and Control Regions

As the official STXS regions (see Figure 8.9) have a split at 400 GeV, the Boosted analysis has two p_T^V categories: $250 \text{ GeV} \leq p_T^V < 400 \text{ GeV}$ and $p_T^V \geq 400 \text{ GeV}$.

The analysis splits the p_T^V regions further in the 0- and 1-lepton channels by sorting the regions into how many additional small- R track jets are associated with the H_{cand} . Events with no additional b -tagged track-jets (b -track-jets) associated with the H_{cand} are

designated as signal, and those that have one or more are designated as background.

This is because in Boosted analysis $t\bar{t}$ events, due to the topology of the decaying top pair, the b -track-jets that are associated with the large- R jet are often a b -labelled jet, and a c -labelled jet from the decay of one of the W-bosons. This means that the second actual b -jet in the event from the other top is not part of the H_{cand} . Events, therefore, with additional b -track-jets not associated to the H_{cand} with respect to the nominal selection, are likely to be background. The requirement on additional b -track-jets forms a CR which is enriched with $t\bar{t}$ events.

In the Resolved analysis there are two jet categories: 2-jet and 3-jet. In the Boosted analysis, since signal emissions resulting in a third jet are more likely to be emitted between the b -quark and travelling in the same direction, this emission is likely to be picked up within the large- R jet. Occasionally, however, this emission is not in the cone of the travelling b -jets and is not then associated with the H_{cand} . While it is important to catch signal events like this, it is expected that there are also background events (namely W +jets and $t\bar{t}$) that can produce similar signatures. This is referred to as a ‘low-purity’ signal.

The signal regions in the 0- and 1-lepton channels are therefore further categorised depending on the number of small- R jets not matched to the H_{cand} . Events without small- R jets associated outside the H_{cand} are referred to as ‘high purity’ (HP) signal, and those with small- R jets associated outside the H_{cand} are designated as ‘low purity’ (LP) signal.

The 2-lepton analysis has good $t\bar{t}$ and W +jets background rejection and thus events with additional small- R jets or b -track-jets not associated to the H_{cand} need not be categorised separately from signal events.

For each p_T^V category, the 0- and 1-lepton channels have a HP-SR, a LP-SR and a CR, and the 2-lepton only has a SR. These ten SRs and four CRs are summarised in Table 8.9.

Table 8.9: Summary of the Boosted analysis regions. Regions with a lot of signal are designated SR. However there are regions with comparatively lower and higher purity of signal events and these are labelled ‘low-purity’ (LP) SRs and ‘high-purity’ (HP) SR’s respectively. The abbreviation ‘add’ designates additional small- R jets that are not matched to the large- R jet. Table from [4, p10].

Channel	Categories					
	$250 < p_T^V < 400 \text{ GeV}$			$p_T^V \geq 400 \text{ GeV}$		
	0 add. b -track-jets		≥ 1 add. b -track-jets	0 add. b -track-jets		≥ 1 add. b -track-jets
	0 add. small- R jets	≥ 1 add. small- R jets		0 add. small- R jets	≥ 1 add. small- R jets	
0-lepton	HP SR	LP SR	CR	HP SR	LP SR	CR
1-lepton	HP SR	LP SR	CR	HP SR	LP SR	CR
2-lepton	SR			SR		

8.5.2 Systematic Uncertainties

Since the signal and backgrounds are shared with the Resolved analysis, the Boosted analysis has the same uncertainties for the most part. The main differences come down to the modelling of new objects, so the calibration of the track-jets and large- R jets. Table 8.10 shows the average impact of all the systematics on μ_{VH}^{bb} in the final fit.

8.5.2.1 Experimental Uncertainties

The uncertainties and corrections for small- R jets are treated in the same manner as jets from the Resolved analysis as they are essentially the same objects.

Like the Resolved analysis there are also a number of b-tagging uncertainties which are decomposed into many analysis systematics. These uncertainties, however, come from calibrations with p_T^V limits. To extrapolate these uncertainties beyond this data limit for

Table 8.10: Breakdown of the systematic contributions to μ_{VH}^{bb} . The effect of uncertainties in both p_T^V regions is combined. The average impact is the mean of the upwards and downward systematic fluctuations and the total impact is different from the quadrature sum of all the contributions in the table because not all the systematics are independent. Table from [4, p16].

Source of uncertainty		Avg. impact
Total		0.372
Statistical		0.283
Systematic		0.240
Experimental uncertainties		
small-R jets		0.038
large-R jets		0.133
E_T^{miss}		0.007
Leptons		0.010
b -tagging	b -jets	0.016
	c -jets	0.011
	light-flavour jets	0.008
	extrapolation	0.004
Pile-up		0.001
Luminosity		0.013
Theoretical and modelling uncertainties		
Signal		0.038
Backgrounds		0.100
$\hookrightarrow Z + \text{jets}$		0.048
$\hookrightarrow W + \text{jets}$		0.058
$\hookrightarrow t\bar{t}$		0.035
$\hookrightarrow$ Single top quark		0.027
$\hookrightarrow$ Diboson		0.032
$\hookrightarrow$ Multijet		0.009
MC statistical		0.092

b -, c - and light-jets requires an additional extrapolation uncertainty. This uncertainty is not a large effect for the Resolved analysis due to the average lower energies of objects.

Uncertainties in the energy and mass scales of large- R jets are determined by comparisons between the ratio of calorimeter-based to track-based measurements in dijet data and in simulation. These uncertainties are used to smear jet observables and under different calibration scales, the variation in the large- R jet mass are observed.

8.5.3 Analysis methodology

Results are obtained via a simultaneous fit to the data of the m_J distribution in order to extract μ_{VH}^{bb} and μ_{VZ}^{bb} at the same time. For the measurements of the signal strength in the STXS framework, the analysis created a 4-POI scheme by splitting the signal into WH and ZH production in each of the p_T^V bins.

8.5.4 Cross-check analyses

As the methodology is new, and a simple cut-based selection input to a fit there are no analysis cross-checks.

8.6 The significance of statistics

Aside from the discussion of the results, which is done in Chapter 11, one of the most important things about the analyses that has intentionally been skipped over is the statistical process of fitting. This includes extracting signal strengths from the data and determining the effects of systematics in order to interpret the final event selection in terms of the underlying physics. This is the focus of the next chapter which details the statistical workings of the fit, and my work on the optimisation of the fit.

The VHbb fits and post-fit modelling investigations

9.1 Statistics, Fits and Significances

After generating the analysis CxAOD inputs from the Reader, the information contained within these CxAOD's (background samples, number of jets identified, region, various systematic fluctuations etc) needs to be condensed into a measure of signal strength to estimate the confidence with which the data matches a set of theoretical predictions.

In the VH , $H \rightarrow b\bar{b}$ analyses, the input distributions are translated into a signal strength measurement by a *fit* via the WSMaker package which performs a comparison of the MC simulations plus all its uncertainties to the data. To monitor the fit process, many auxiliary plots are produced alongside the final plots designed to inform the analysers about how variables and assumptions interact with each other.

Fundamentally, the fit is a mathematical formula that is maximised or minimised and at a simple level asks how many data points $\mathcal{N}_i^{data}$ appear in each bin given that there are $\mathcal{N}_i^{exp}$ MC predicted events for all bins, i. A fit needs parameters to vary, and since the data is the observed counts, only the predictions can change and these will vary within their own respective uncertainties. These uncertainties are introduced as parameters (θ) for $\mathcal{N}_i^{exp}$.

Producing a fit in this manner is done based on the fundamental tenets of hypothesis testing. These are fourfold.

- The null (H_0) and alternative hypotheses (H_1) are defined
- Define the test statistic q under each of these hypotheses: q_{H_0} and q_{H_1}
- Compute from observations the observed value: q_{obs}
- Decide based on q_{obs} whether to reject H_0 or H_1

The degree of confidence the data is described by one hypothesis or another is called the *significance* (Z). This significance is the number of standard deviations (σ) under the null hypothesis that separate q_{obs} and the central value of q_{H_0} . This significance can also be described as how much the background model would have to fluctuate to look like the data. A ($Z =$) 1σ effect can be caused by a background fluctuation roughly 35% of times, and a 2σ fluctuation occurs roughly 5% of the time. When looking for possible upwards (or downward) fluctuations over many bins, more significant fluctuations are expected to be seen (this is called the look-elsewhere effect [149, p1-2]). Using this kind of thinking there are two standards used to describe possible new effects: *Evidence* and *Discovery*. Evidence is the name given to a 3σ effect, which can be caused by a background fluctuation with a probability of roughly 0.3%, and a discovery is a 5σ effect that is caused by a background fluctuation less than one in a million times¹.

¹roughly 1 in every 1.7 million trials

9.1.1 The Null and Alternative Hypothesis

In analyses like the VH , $H \rightarrow b\bar{b}$ analyses, the aim is to discover a particle in a new decay mode. An observation of this particle therefore will manifest itself as additional events in the bins of $H \rightarrow b\bar{b}$.

Firstly the null and alternative hypotheses are constructed. The null hypothesis is taken to be the case that no such decay exists, also known as the background-only model. The alternative hypothesis is the case where the Higgs boson decays into b -quarks, known as the signal+background model.

A robust model of the background is needed to form the null hypothesis such that it can test how likely possible upward fluctuations are in the signal region. Here it is important to note that while both upward and downward fluctuations from the background can occur, any downward fluctuations in the background do not serve as evidence against the background-only model in comparison to the signal+background model.

In addition to defining the null hypothesis, at this stage the p-value is selected. The p-value is a probability at which it is decided that the null hypothesis will be rejected. For example, if the p-value is set to 5%, if an observation is made that corresponds to a background fluctuation of more than 2σ , then the background-only hypothesis will be rejected.

To test that the background model is accurate, it can be compared to data in a region of phase space that is not relevant to the signal. This is one where the physics model should still be able to predict the event rates and distributions, and the signal makes little or no contribution to the event count. This is why analyses have control regions. The simultaneous test of the background model in these control regions and the signal

regions is where the likelihood comes in.

9.1.2 The Likelihood and the Test Statistic

Whenever an experiment is built to measure an outcome, one of the most important features an experiment can have is that it is reliable. This is to say that the same (or a similar) result can be obtained if the experiment is repeated. This demand is made because it is assumed that with the same conditions, the same ‘true’ value of the object in question is measured. Each independent dataset ($\bar{x}$ that is obtained measuring a value of ‘x’) is distributed around the hypothesised true value: μ . This can be represented as Equation 9.1.

$$\mathcal{L} = p(\bar{x}|\mu) = \mathcal{L}(\mu, \vec{\theta}, \vec{k}, \mathcal{N}^{data}) \quad (9.1)$$

$\vec{\theta}$ and $\vec{k}$ are vectors of parameters (called *systematics* and *normalisations* respectively) that allow the background model to be manipulated. Their relevance will be expanded upon later on in this section and in Section 9.2.

The function in Equation 9.1 is usually referred to as the likelihood, and is denoted by $\mathcal{L}$. The likelihood in the base form is a function of how likely a dataset is generated given a particular hypothesis [150, p24-34]. It is not the probability of a hypothesis being true given the data, which while being the ultimate aim, has to be extracted from the likelihood using Bayes Theorem [151, p15] (see Equation 9.2).

$$p(\mu|\bar{x}) = \frac{p(\bar{x}|\mu) \times p(\mu)}{p(x)} \quad (9.2)$$

The likelihood is a way of mathematically encoding the null and alternative hypotheses. It is a probability distribution made up of a prediction which comes from MC simulations created under a ‘nominal’ physics model. Probabilistic variance from this central value comes from changing values of Nuisance Parameters (NP) in the model.

A NP is a given parameter in the physics model that is not fully understood or subject to change. Each NP has a nominal value, and its own distribution of how likely deviations from this nominal value are.

Variations in NP's will vary processes in the model and therefore increase or decrease the number of counts seen in the bins where that NP is relevant. Some NP's will have a larger effect on the number of counts in a bin than others.

The *test statistic* [152, p5] used to conduct the hypothesis test is a one-number summary of the physics model and is therefore a function of the likelihood. The Neyman-Pearson lemma says that the best possible test statistic to choose, $q(\mu)$, is the logarithm of the ratio between the likelihoods that come from the null and alternative hypothesis [152, p19], as expressed by Equation 9.3 below.

$$q(\mu, H_0) = -2 \ln \frac{\mathcal{L}(H_1|x_{\sim H_0})}{\mathcal{L}(H_0|x_{\sim H_0})} \quad | \quad q(\mu, H_1) = -2 \ln \frac{\mathcal{L}(H_1|x_{\sim H_1})}{\mathcal{L}(H_0|x_{\sim H_1})} \quad (9.3)$$

The important difference between them is that the dataset 'x' is generated differently in each case to correspond to the hypotheses being tested. The test statistic for the null hypothesis ($q(\mu, H_0)$) is evaluated as the log likelihood ratio of the null and alternative hypothesis where the dataset is generated under the *null* hypothesis ($x_{\sim H_0}$). Likewise the test statistic for the alternative hypothesis ($q(\mu, H_1)$) is evaluated as the log likelihood ratio of the null and alternative hypothesis where the dataset is generated under the *alternative* hypothesis ($x_{\sim H_1}$).

These $q(\mu)$'s are probability densities that integrate to 1 and typically overlap. The smaller the overlap, the more powerful any one result is in distinguishing which of the distributions a the trial test statistic from the data would belong to. An example of what this looks like can be seen in Figure 9.1.




Figure 9.1: Example formulation of the test statistic for the null (q_{H_0}) and alternative hypothesis (q_{H_1}). α , given by the purple hatched area, corresponds to the probability that the null hypothesis is rejected given that the null hypothesis is true, and β , given by the orange hashed area, is the probability that the alternative hypothesis is rejected given that the null hypothesis is not true. Smaller q values are in general more signal like. The value of q that corresponds to α equalling the p -value is denoted by $q_{p\text{-value}}$. Before any measurement is taken, the strength or power of the statistical test can be measured from the relative areas of α and β given a observed $q(q_{obs})$ equalling this $q_{p\text{-value}}$. Image generated using [28].

It is also important to measure the contribution of a single systematic in the model. This can actually be done by creating a test statistic for each systematic $q_{\text{NLL}}(\mu, \theta_j)$ so that each systematic is separated out from the likelihood to determine its optimal value and its variance. The maximised likelihood with optimised values of the entire set systematics except θ_j , is divided by the maximised likelihood with the optimised value θ_j . Similarly to the $q(\mu)$ above, the negative natural logarithm of this likelihood ratio is taken (see Equation 9.4).

$$q_{\text{NLL}}(\mu, \theta_j) = -2 \ln \frac{\mathcal{L}(\mu, \hat{\theta}_j)}{\mathcal{L}(\mu, \hat{\theta})} \quad (9.4)$$

According to Wilks theorem [153, p62], $q_{\text{NLL}}(\mu, \theta_j)$, referred to as the negative log likelihood or NLL is distributed as a χ^2 with one degree of freedom. This allows the variation of the NLL by a certain number to give rise to the uncertainty in the systematic. Details of this process accompanied by example plots will be shown in Section 9.2.

9.1.3 Measurement and Results

The process of running the ‘fit’ on data consists of optimising the values of a set of NP’s by varying them from their nominal values and testing the goodness of fit.

The parameter set is modified such that the number of counts in the MC simulations replicates the data as closely as possible while minimising the systematic uncertainties in order to achieve this. Once the most likely alternative model is identified, a measurement is conducted and the test statistic for the data, q_{obs} , is calculated. As one set of values were observed, q_{obs} is a delta function given by Equation 9.5,

$$-2 \ln \frac{\mathcal{L}(H_1|x_{obs})}{\mathcal{L}(H_0|x_{obs})}. \quad (9.5)$$

If q_{obs} is larger than the given p-value, then the null hypothesis is rejected. This can be seen in Figure 9.2.




Figure 9.2: Upon measurement of q_{obs} , the alternative hypothesis is rejected as the q_{obs} is smaller than the p-value set. The significance of the result is given by Z , the number of standard deviations by which the central values of q_{H_0} and q_{obs} are separated. Image generated using [28].

As discussed previously one of the final figures of merit is the significance. The *significance* of this q_{obs} is calculated by finding out how many deviations from the nominal value of q_{H_0} , this q_{obs} lies (given by Z in Figure 9.2). The *expected significance* is given by the number of standard deviations the centre value of q_{H_1} is from q_{H_0} , as this is the most likely value for the test statistic to have in the case that the signal model describes the data.

9.1.4 Significance and the Asimov Dataset

It is also useful to have an idea of whether the analysis has any chance of obtaining its goals before looking at the data or ‘unblinding’.

As explained before this is done by generating MC simulations for the alternate hypothesis, finding the distribution of the test statistic under the alternative hypothesis and the median value of the test statistic (q_{median}), then finding the p-value of the null hypothesis that corresponds to this q_{median} . This is the sensitivity.

However, MC simulations do not need to be generated to get the median of the alternative hypothesis. Given that the alternative analysis is often a modification of the null hypothesis characterised by differing values of certain parameters, if the distribution of your alternative hypothesis is known, the median can be obtained via mathematics. A single representative point to represent the population is used to determine this value, and this point is provided by the *Asimov Dataset* [152, p10].

In the Asimov dataset, all the parameter values represent the values that would be estimated from the Monte Carlo model with high statistics. It is computationally cheap as only the maximal point is determined as opposed to an entire distribution.

9.2 Likelihood Fit Model from first principles

This is a small practical introduction into likelihood fitting, which is a common tool in many analyses in High Energy Physics. By looking at the relatively trivial case of a single bin fit, the aim of this section is twofold. Firstly, to demonstrate what rough processes happen in a fit that occurs over several hundred bins; and secondly it is to introduce some of the plots that are used to display VH , $H \rightarrow b\bar{b}$ fit information that will be shown later on in this chapter and to describe how they originate.

The likelihood in its most generic form is given by the Equation 9.6 below

$$\mathcal{L}(\mu, \vec{\theta}, \vec{k}, \mathcal{N}^{data}) = \prod_{i \in bins} \text{Pois} \left(\mathcal{N}_i^{data} | \mu s_i(\vec{\theta}) + k_i b_i(\vec{\theta}) \right) \times \prod_{\theta \in \vec{\theta}} \frac{1}{\sqrt{2\pi}} e^{-\frac{\theta^2}{2}} \quad (9.6)$$

$$\times \prod_{i \in bins} \frac{(k_i \beta_i)^{\beta_i} e^{-k_i \beta_i}}{\Gamma(\beta_i + 1)}$$

- μ is the signal strength parameter. It quantifies how much signal is in this bin.
- $\vec{\theta}$ is the set of Nuisance Parameters (NP) in the fit normalised such that $\theta = 1$ is a 1σ shift in that NP.
- $\vec{k}_{(i)}$ is the set of scale/normalisation factors for the backgrounds present (in bin i).
- $\mathcal{N}_{(i)}^{data}$ is the number of data events observed (in bin i)
- s_i is the number of signal events expected in bin i
- b_i is the number of background events expected in bin i
- β_i is the weights associated with events generated in bin i

The first part of the likelihood, in orange, is the model Poisson probability. For each bin in the model, the Poisson probability is calculated by obtaining the number of data events in that particular bin given the number of events that were expected to be in that bin. The model Poisson probability is the product of all these Poisson probabilities.

The part in green is the Gaussian prior, which penalises against large NP pulls. This is an ‘Occam’s Razor’-like term that punishes bad or rare models/deviations with small multiplication values.

The part in purple is a Monte Carlo Statistics (MCStat) term, which arises due to the fact that in a complex counting experiment sometimes the occurrence of rare events is boosted to increase the statistics in certain regions.

In a simplified case: MC simulations have been generated without event boosting in any regions, so the term on the end disappears. There is only one bin and one uncertainty, θ_j , and the number of events that constitutes to the signal does not matter, just the total background count. Equation 9.6 simplifies to

$$\mathcal{L}(\mu, \theta_j, \vec{k}, \mathcal{N}^{data}) = \frac{(\mathcal{N}^{exp}(\mu, \theta_j, \vec{k}))^{\mathcal{N}^{data}} e^{-\mathcal{N}^{exp}(\mu, \theta_j, \vec{k})}}{\mathcal{N}^{data}!} \times \frac{1}{\sqrt{2\pi}} e^{\frac{-\theta_j^2}{2}} \quad (9.7)$$

In the case where there are no prior estimations, and there are no scale or normalisation factors to be implemented, Equation 9.7 can be further simplified to

$$\mathcal{L}(\mu, \theta_j, \mathcal{N}^{data}) = \frac{(\mathcal{N}^{exp}(\mu, \theta_j))^{\mathcal{N}^{data}} e^{-\mathcal{N}^{exp}(\mu, \theta_j)}}{\mathcal{N}^{data}!} \times \frac{1}{\sqrt{2\pi}} e^{\frac{-\theta_j^2}{2}} \quad (9.8)$$

$$\mathcal{L}_{no\ prior}(\mu, \theta_j, \mathcal{N}^{data}) = \frac{(\mathcal{N}^{exp}(\mu, \theta_j))^{\mathcal{N}^{data}} e^{-\mathcal{N}^{exp}(\mu, \theta_j)}}{\mathcal{N}^{data}!} \quad (9.9)$$

So let us take a toy example. For $\mathcal{N}^{data} = 10$ with no Gaussian prior or systematics, the value of $\mathcal{N}^{exp}$ which maximises $\mathcal{L}_{no\ prior}(\mu, \theta_j)$ should be the one that exactly predicts the data, so $\mathcal{N}^{exp} = 10$. This can be seen in Figure 9.3.

In the case where there are no systematic uncertainties, event weights or normalisations, the expected number of events is given by the number of MC simulated events.




Figure 9.3: Variation of the likelihood for different values of the number of expected events ($\mathcal{N}^{exp}$). The maximum possible likelihood is reached when $\mathcal{N}^{exp} = \mathcal{N}^{data}$.

In the case where there are systematics, however, the number of expected events becomes

$$\mathcal{N}^{exp} = \mathcal{N}^{MC} + \sum_n (1\sigma_{\theta_n} \times \theta_n), \quad (9.10)$$

where $1\sigma_{\theta_n}$ is the number of events that a 1σ pull in that systematic corresponds to, and θ_n is how many standard deviations the systematic is pulled before the fit is run. Putting some small trial numbers into Equation 9.10, if the number of events generated via MC simulations is 8, and a $+1\sigma$ shift in some arbitrary systematic called θ_j corresponds to 0.5 more events, then $\mathcal{N}^{exp}(\theta_j = 0) = 8$ and $\mathcal{N}^{exp}(1) = 8.5$. Figure 9.4 displays the value of the number of shifts in this systematic that maximises this likelihood when the number of events present in the data in this bin remains at 10.

Figure 9.4 shows that the optimised fit only needs a small shift in the systematic θ_j (i.e




Figure 9.4: Plot of the likelihood value at different values for the number of sigma the systematic θ_j is shifted. For the simple fit, the systematic is only slightly pulled from no shift at all.

the central value of 0.121 is close to 0). A ‘pull’ will be quoted as this central value. Systematics that are not pulled at all when plotted against the likelihood, will be gaussian with width 1 and a central value of 0. The constraint on this pull which will serve as a 1 sigma shift is not the width of this distribution, that will come later.

If the problem is scaled up, a situation can be engineered where a pull is required. If the effect of the systematic θ_j is made to be twenty times as large, and the number of both the data events and the MC simulated events are raised by an order of magnitude, Figure 9.5 shows that the pull of this new θ_j is 1.06.

It is also important to note here that in both cases, the systematic was not expected to have a pull before it entered the fit. While this is mostly the case, if an effect is known not to be accounted for in the model, it is possible to insert a pre-fit pull on the systematic before it goes into the fit. More accurately then, a pull is the difference between the




Figure 9.5: Plot of the likelihood value at different values for the number of sigma the systematic θ_j is shifted. When $N^{data}=100$ and $N^{exp}=80$, the systematic is pulled up by roughly 1σ .

post-fit and the pre-fit value of the pull (which is normally 0).

In these two examples, the maximum of the likelihood distributions do not correspond to a pull of the systematic such that $\mathcal{N}^{exp}$ equals $\mathcal{N}^{data}$. This is because there is an interplay between the statistical uncertainty terms, and the penalty terms from nuisance parameter constraints. Pulling the systematic introduces a Gaussian prior term, which has a negative impact on the likelihood, and gets exponentially larger the larger the required shift. To this end, one should remember that the error on the number of counts in a given bin is $\sqrt{N}$. While the non-prior maximum is $\mathcal{N}^{exp} = \mathcal{N}^{data}$, in practice, a fit with a small number of free parameters, will only pull systematics to change the number of expected events to be close to, or equal to, the bottom of this uncertainty range $(\mathcal{N}^{data} - \sqrt{N})^2$.

²Due to the symmetrical shape of the likelihood around its peak, this same is true for the reverse case. If $\mathcal{N}^{exp} > \mathcal{N}^{data}$ and the fit has low flexibility, systematics will be pulled to move the number of expected events to the top of the uncertainty range: $\mathcal{N}^{data} + \sqrt{N}$


This is seen in the examples that were shown previously. In the first example, the error on the number of counts in the bin is $10 \pm \sqrt{10} = 10 \pm 3.2$. The lower bound on this error is $10 - 3.2 = 7.8$ which encompasses the $\mathcal{N}^{MC}$ of 8, so there is little need to incur the penalty imposed by pulling θ_j the full 4σ required to give an exact match. In the second example 100 ± 10 does not cover the MC_{pred} of 80, but a 1σ upwards pull in the systematic results in $\mathcal{N}^{MC} = 90$.

It is reasonably useful to find the variance in the optimised θ_j chosen when the value of the 1σ varies, and when $\mathcal{N}^{MC}$ varies. This is displayed in Figures 9.6a, 9.6b and 9.6c which show that the maximal likelihood is dependent on the physics model.


As seen in Figure 9.4, Figure 9.6a shows that the maximum possible likelihood under a given set of systematics is when the $\mathcal{N}^{MC}$ matches that of the data. So independent of the systematics, the closer the MC simulations are to that of the data, the larger the maximum likelihood. As well as optimising systematics, some amount of time must be invested in ensuring simulated process are behaving as expected.

Figures 9.6b and 9.6c show some of the dangers of systematics. The larger the systematics are, the less the fit has to pull them to get the desired effect. Large systematics are usually unphysical as corrections on such scales should be included within the physics model. However, if introduced, their pulls could dwarf all other effects, or risk the number of events in a bin going negative, which is also unphysical. Since there is no upper bound on the likelihoods improvements, it is important to understand how each systematic is relevant and justify each one and its size before including it in the final model.


It has been established with the previous examples that for a fixed MC simulations



(a)



(b)



(c)

Figure 9.6: Maximum likelihood variations for fixed N_{data} and (top) fixed pre-fit impact of θ_j but varying N_{MC} ; (bottom) fixed N_{MC} but varying the pre-fit impact of θ_j . The difference between the bottom left and bottom right plots is that the bottom left plot shows the prefit impacts and the bottom right shows the postfit impacts.

prediction, the likelihood can be maximised globally by choosing the right set of parameters. Some of these parameters will be independent and some will not, but for full control over the physics model, there needs to be a way of tracking and adapting the behaviour of each systematic individually as one requires a way of quantifying which systematics are more impactful to the fit.

To do this, a test statistic for the systematic can be created. As mentioned in the previous section, this test statistic is the NLL. The two likelihoods that are divided against each other are the global maximised likelihood given by Equation 9.11,

$$\mathcal{L}(\mu, \hat{\theta}_j, \mathcal{N}^{data}) = \frac{(\mathcal{N}^{exp}(\mu, \hat{\theta}_j))^{\mathcal{N}^{data}} e^{-\mathcal{N}^{exp}(\mu, \hat{\theta}_j)}}{\mathcal{N}^{data}!} \times \frac{1}{\sqrt{2\pi}} e^{-\frac{\hat{\theta}_j^2}{2}} \quad (9.11)$$

and the maximised likelihood for varying systematic θ_j , which is given by Equation 9.8.

$$\begin{aligned} NLL &= -\ln \left(\frac{\mathcal{L}(\mu, \theta_j)}{\mathcal{L}(\mu, \hat{\theta}_j)} \right) \\ &= -\ln \left[\frac{\frac{(\mathcal{N}^{exp}(\mu, \theta_j))^{\mathcal{N}^{data}} e^{-\mathcal{N}^{exp}(\mu, \theta_j)}}{\mathcal{N}^{data}!} \times \frac{1}{\sqrt{2\pi}} e^{-\frac{\theta_j^2}{2}}}{\frac{(\mathcal{N}^{exp}(\mu, \hat{\theta}_j))^{\mathcal{N}^{data}} e^{-\mathcal{N}^{exp}(\mu, \hat{\theta}_j)}}{\mathcal{N}^{data}!} \times \frac{1}{\sqrt{2\pi}} e^{-\frac{\hat{\theta}_j^2}{2}}} \right] \\ &= -\ln \left[\frac{e^{-\frac{\theta_j^2}{2}}}{e^{-\frac{\hat{\theta}_j^2}{2}}} \right] = -\ln \left(e^{-\frac{1}{2}(\theta_j^2 - \hat{\theta}_j^2)} \right) \\ NLL &= +\frac{1}{2}(\theta_j^2 - \hat{\theta}_j^2) \end{aligned} \quad (9.12)$$

Since θ_j and $\hat{\theta}_j$ will be equal at one point, according to Equation 9.12, the minimum of the NLL is 0. The value of θ_j at this minimum is the nominal post-fit pull value. To evaluate the post-fit error bands on θ_j , the $\pm 1\sigma$ variation of the θ_j (68% confidence limits) around the minimum of the NLL are taken.

Since θ_j is formulated as the number of σ shifts in a systematic, at the minimum $\hat{\theta} = 0$ and a 1σ shift of $\hat{\theta} = 1$. This 1σ expansion around the minimum is equivalent to $\theta_j \rightarrow \hat{\theta}_j \pm 1\sigma = 0 \pm 1$.

$$\begin{aligned}
 NLL_{1\sigma} &= +\frac{1}{2}((\hat{\theta}_j \pm 1\sigma)^2 - \hat{\theta}_j^2) \\
 &= +\frac{1}{2}((0 \pm 1)^2 - (0)^2) \\
 &= +\frac{1}{2}((\pm 1)^2) \\
 &= +\frac{1}{2}(1) = +\frac{1}{2}
 \end{aligned}$$

Hence the 68% confidence limits on the pull of the systematic θ_j correspond to the values of θ_j that have a NLL value of 0.5. These values can be seen for θ_j calculated from the NLL in Figure 9.7.




Figure 9.7: Plot showing the variations of the NLL when the systematic θ_j is varied around nominal pull value. When the NLL changes by +0.5 from the minimum, the pull has varied by $\pm 1\sigma$. A change of +2 corresponds to a $\pm 2\sigma$ shift.

The fit can generate this kind of plot, but the more generally used output is known as a ‘pull plot’ shown in Figure 9.8. This plot takes the values of a systematic that corresponds to the minimum and 1σ variations and displays them in a way multiple pulls can be viewed at the same time.




Figure 9.8: Example ‘pull plot’ for θ_j . Multiple pulls can be shown on such a plot, and the fit produces several of these with pulls common to various processes displayed on the same pull plot.

If a systematic has been ignored by the fit, it will be displayed on the pull plot, as a pull of 0 with 1σ variance of ± 1 . If a systematic is pulled beyond expectation (which can sometimes come with small error bars) or is pulled when no pull was expected, then that systematic is said to be *overconstrained*. If a systematic has a variance larger than ± 1 it is said to be *underconstrained*. It is important to note that in the definition of this methodology, pulls with variances greater than 1σ are not allowed so this indicates an error in the coding of the model or a numerical issue.

Underconstraining can happen for parameters whose 1σ effects change the number of

expected events by a very small amount. Variations in the pulls of those systematics do not look gaussian like that of Figure 9.5 but seem flat. This makes the numerical estimation of where the $\pm 1\sigma$ points are less accurate.

It can also occur when there are systematics with up/down fluctuations in the same direction (e.g. 1σ up = +5 events and 1σ down = +2 events). This definition of a systematic is not well incorporated into the definition of the likelihood function, so the calculation may fail, or act as if it failed.

Knowing the size and variance of the pulls is important, but some systematics have more of an effect than others as they correspond to larger changes in the number of events. A large pull in the model may be excused, if its change to the overall model borders on insignificant. To evaluate this, the *impact* of the systematic is required.

The impact of a systematic is calculated by taking the values of the systematic that correspond to the $\pm 1\sigma$ fluctuations and inserting them into the relevant figure of merit with all the other systematics at their likelihood-maximising values. In the case of the analyses to be discussed, this parameter is signal strength parameter μ ³, which is essentially the ratio between N^{exp} and N^{data} (the ratio between cross-sections). The normalised difference between the values of μ at the $\theta_j \pm 1\sigma$ systematic fluctuations and the maximal μ is the impact.

The impact of the systematic θ_j on μ is therefore given by Equation 9.13,

$$\Delta\mu^\pm = \frac{\hat{\mu}_{1\sigma^\pm} - \hat{\mu}}{\hat{\mu}}, \quad (9.13)$$

where $\hat{\mu}$ is the number of events predicted in the one bin after fit systematic variations

³This is not the same μ that is in the equations above. This is equivalent to the μ from the previous chapter


and $\hat{\mu}_{1\sigma^\pm}$ are the $\pm 1\sigma$ variations from $\hat{\mu}$.

Figure 9.9 shows how μ changes within the 68% confidence interval of θ_j . It also shows that the number of expected events in the bin is 90.5 and a $+1\sigma$ or -1σ variance in the systematic θ_j corresponds to this changing by +6.83 and -6.63 events respectively. Since the number of observed events is 100, the impact of this systematic on $\hat{\mu}$ is +0.068 and -0.066.


In Figure 9.9, the equation used to calculate the impact of θ_j on μ was linear. Often when the prediction is close to zero and the systematics are large enough, downward fluctuations in the systematic can lead to the number of events in a given bin becoming negative. A negative event count is not a physical occurrence, and the problems it can cause can be exacerbated by scale factors that multiplicatively augment the number of events and can make the negative bin larger. Since the number of events must stay the same, in a standard fit with more than one bin, the fit will also try to compensate for negative bins by adding events to other bins.

To avoid this, a systematic being varied in the fit should never be able to reduce the number of events in a given bin to less than 0. The equation that forms the line in the impact plot can be augmented to do this by adding exponentials to the function when it is extrapolated. This extrapolation occurs when the systematic $|\theta_j| > 1\sigma$ because the behaviour of the systematic is defined between its 1σ fluctuations. Having a plot which is two exponentials joined by a linear function forms a more stable description of how μ varies with θ_j .

On the lower end this ensures that μ can never be negative, but on the upper end, this approach makes the systematic have a wilder effect on μ if it is ever pulled up that far, thus more easily drawing attention to systematics with reasonably large upward pulls.



(a)



(b)

Figure 9.9: 1σ shifts from the minimum NLL for the $\theta + j$ (a) and the subsequent change in the number of events in the bin (b). Since the number of data points is 100, to get the impact, the y-axis in the bottom plot is divided by 100.

When more systematics are introduced, the correlations between these systematics become important. Since correlations between systematics are derived from double derivatives of a function of the likelihood [150, p158-161], systematics need to be continuous and not discrete. The solution to the systematic making the number of predicted events not be less than zero in a continuous way is to connect the two extrapolation exponentials with a piecewise linear function. This means the systematic is linear in the region it needs to be linear while being fixed at the edges with defined values as to not have discontinuities with the exponential functions on its boundaries.

The last set of diagnosis plots to be introduced in this section are pre-fit and post-fit plots. These are plots that show histograms with the number of (weighted) sample events and the data-to-Monte-Carlo-generated event ratio (data/MC) in each bin of a particular region. Each sample is typically denoted by different colour, and these colours will be consistent across the different analysis channels. This plot is the amalgamation of all the intricacies of the fit, so is a good indication of the overall outcome of the fit (the post-fit plot is the first plot that gets looked at when a fit is created). If there are any problems or trends in the data/MC ratio, one would look to see if there is the same trend in the pre-fit plots. Trends here would merit a detailed look at the fit to find the cause.

Being an overall look at the fit the final result, if consistent with the data, may in actual fact mask large pulls in systematics required to achieve the ‘goodness’ of the fit. Figure 9.10 shows the prefit and post-fit plots for the 1-bin, 1-systematic example.

While ideally, the data and the total background count in Figure 9.10b should be overlapping, in this representation it is easier to see that if the fit can only has sufficient

flexibility⁴ to pull the estimate such that its uncertainty upper band overlaps with the data, then that is all it will do.




Figure 9.10: Pre-fit (9.10a) and post-fit (9.10b) plots for the example fit. Each sample that contributes to the number of events in that particular bin would get its own colour. In this case, the assumption was made that the bin was a background-only bin with only one sample contributing to events. The ideal result for the post-fit plot would be the black data point meeting the black dotted line. This would result in the blue point in the ratio plot lying on the data/MC = 1 line. Normally in a plot of this type, the y-axis would show weighted events, but there are no weights put on any events in this example for simplicity. The grey hashed line, therefore, represents the $\sqrt{N}$ error on the number of counts.

9.2.1 Adding additional systematics

When a fit has multiple systematics that act in the same region of bins, it can choose to pull any or all systematics by some amount. The functional form of the likelihood gets

⁴For example, the number of systematics are too low, or the systematic corrections are not impactful enough.

a term for each systematic introduced. In the next example, one other systematic, θ_i , will be introduced, and the likelihood presented in Equation 9.8 becomes this:

$$\mathcal{L}(\mu, \theta_i, \theta_j) = \frac{(\mathcal{N}^{exp}(\mu, \theta_i, \theta_j))^{\mathcal{N}^{data}} e^{-\mathcal{N}^{exp}(\mu, \theta_i, \theta_j)}}{\mathcal{N}^{data}!} \times \frac{1}{\sqrt{2\pi}} e^{-\frac{\theta_i^2}{2}} \times \frac{1}{\sqrt{2\pi}} e^{-\frac{\theta_j^2}{2}} \quad (9.14)$$

Sticking with $\mathcal{N}^{data} = 100$, $\mathcal{N}^{MC} = 80$, and the pre-fit event 1σ shift of $\theta_j = 10$ events, this new independent systematic θ_i will have a pre-fit event 1σ shift of 2 events. Since these systematics both act in the same bin, they are necessarily correlated. However, as their effects are different, the correlation will not be perfect.

Figure 9.11 shows the maximal likelihood for variation in both θ_i and θ_j . The particular combination of θ_i and θ_j that maximise the likelihood give rise to the central points in the pull plot in Figure 9.12.

2D Likelihood scan for $N^{data} = 100$ with Gaussian Priors on Θ_j and Θ_i




Figure 9.11: 3D plot showing the change of the likelihood with variation of two systematics: θ_i and θ_j . The maximal point in this $(\theta_i - \theta_j)$ phase space roughly corresponds to a θ_i pull of 0 and a θ_j pull of 1.





Figure 9.12: Pull plot showing the pull and variance for θ_i and θ_j . The inclusion of the red point $\theta_{j,o}$ provides a comparison to the previous fit with its only systematic θ_j . The presence of the second systematic in the same bins means the fit can be optimised by not pulling the systematic θ_j to the full extent it was pulled in the previous version of the fit.


The error bars on the central values of the systematics in Figure 9.12 are given by the 1σ variation of the NLL, where the systematic not varied is set to the value that gave rise to the globally maximised likelihood. This is shown for each systematic in the top subplots 9.13a and 9.13b of Figure 9.13.

The bottom subplots of 9.13a indicate the change in the number of events in the bin that the 1σ shifts of the systematic correspond to. The impact is obtained by dividing these numbers by the number of observed events in the bin, which in this case is an even 100. This division allows multiple systematics to be compared in a fair comparison, and their impacts to be ranked. This information is typically displayed on a ranking plot like the one shown in Figure 9.14.


In a fit with many systematics, a ranking plot is designed to display, in order of the post-fit impact, the top 10-15 systematics. The ranking plot will also show the pull of




(a)



(b)



(c)



(d)

Figure 9.13: 1σ shifts from the minimum NLL for the two systematics θ_i (top left) and θ_j (top right). The plots directly below these show the subsequent impact on the number of events in the bin for that systematic.




Figure 9.14: Ranking plot for two sets of systematics. The θ_i and θ_j set in black, blue and yellow are the impacts and pulls of the current fit. $\theta_{j,o}$ in red, peach and violet, represents the systematic pull, pre- and post-fit impact on μ for θ_j when it was the only systematic present in the bin.

the systematic, and the pre-fit impact on μ .

Coming back to the issue of correlation, understanding the correlations between different pulls is important as it will necessitate possible interactions between physical processes that were or were not expected. Let us take an example from the $VH, H \rightarrow b\bar{b}$ case in particular. If the fit has a systematic relating to the energy of a given jet, and one relating to the number of jets with a b -quark as the initial particle associated (b -tagging), then if the pulls were correlated, you could posit that the b -tagging algorithm used has an efficiency depending on the energy of the jet measured, for example.

Correlations between pairs of systematics are determined by first calculating the covariance matrix between variables. The covariance of two datasets is a measure of how much the expectation of the combination of the two measurements ($E(\mathbf{XY})$) is affected by the expectation of each dataset ($E(\mathbf{X}), E(\mathbf{Y})$) [150, p99]. It is represented by Equation 9.15.

$$Cov(\mathbf{X}, \mathbf{Y}) = E(\mathbf{XY}) - E(\mathbf{X}) \times E(\mathbf{Y}) \quad (9.15)$$

Hence if two datasets are independent then their covariance is 0. The covariance matrix is generated by the result of the double differential of the logarithm of the likelihood for each possible combination of both systematics [154, p34-35] as shown in Equation 9.16.

$$\hat{V}_{ij}(\theta) = \left[\frac{-\partial^2 \ln \mathcal{L}(\mu, \theta_i, \theta_j)}{\partial \theta^2} \Big|_{\theta=\hat{\theta}} \right]^{-1} = \left[\begin{array}{cc} \frac{\partial^2 \ln \mathcal{L}(\mu; \theta)}{\partial^2 \theta_i} & \frac{\partial^2 \ln \mathcal{L}(\mu; \theta)}{\partial \theta_j \partial \theta_i} \\ \frac{\partial^2 \ln \mathcal{L}(\mu; \theta)}{\partial \theta_i \partial \theta_j} & \frac{\partial^2 \ln \mathcal{L}(\mu; \theta)}{\partial^2 \theta_j} \end{array} \right]_{\theta=\hat{\theta}}^{-1} = \hat{\mathcal{H}}^{-1} \quad (9.16)$$

A disadvantage of covariances is that they cannot be easily compared to one another. Correlations are a way of standardising covariances such that they can be fairly com-

pared and are easier to understand as the numbers are limited to being between -1 and 1 inclusive. A variable is perfectly correlated with itself leading to a correlation value of 1. If a variable shifting 1σ upwards causes another to be shifted 1σ downwards, then those two variables are perfectly anti-correlated and are assigned a correlation of -1. The relation from the covariance matrix $\hat{V}_{ij}$ to the correlation matrix S_{ij} [155, p1-23] for two systematics is as follows.

$$\hat{S}_{ij} = \frac{\hat{V}_{ij}}{\sigma_i \sigma_j} = \begin{bmatrix} \frac{V_{ii}}{\sigma_i^2} & \frac{V_{ij}}{\sigma_i \sigma_j} \\ \frac{V_{ji}}{\sigma_j \sigma_i} & \frac{V_{jj}}{\sigma_j^2} \end{bmatrix} \quad (9.17)$$

where σ_i and σ_j are the standard deviation of the post-fit distributions of θ_i and θ_j respectively. For fits with more than two systematics, this correlation is calculated for each pair of systematics, and a correlation matrix is produced with numbers from -1 to 1 showing how correlated each pair of systematics is. Correlations can therefore be expressed in a 2D N-by-N matrix where all the off-diagonal values are non-trivial correlations. The correlation matrix for this example is shown in Figure 9.15. It shows that θ_i and θ_j are slightly anti-correlated which is expected for two systematics acting in the same direction in the same bin.




Figure 9.15: Correlation matrix for θ_i and θ_j . Since a variable is always perfectly correlated with itself, the leading diagonal will always be unity. θ_i and θ_j have a slight anti-correlation.

Once again, the last thing that needs to be seen are the pre-fit and post-fit plots for this example fit. The pre-fit, however, has not changed since it was previously shown. The post-fit plot for this latest example can be seen here in Figure 9.16.




Figure 9.16: Comparison of the post-fit plots between the single and double systematic fits. The pre-fit (Figure 9.10a) for both of these fits is the same. These plots have been zoomed in around the data to show the differences between these fits more easily (the sample level still goes to zero). In the second post-fit plot (9.16b), the data point is closer to the Monte Carlo prediction.

Figure 9.16 shows that when the two post-fit plots for the $\mathcal{N}^{data} = 100$ examples are compared, the one with more freedom of pull flexibility has a more optimised fit, as the data point is slightly more closer to the Monte Carlo prediction.

9.2.2 Adding additional bins

Since there has only been one bin in the example so far, the only types of systematics that have been explored (and the only ones that exist) are ones that affect the overall

number of events. These are called ‘normalisation’ effects. Adding additional bins allows systematics to affect multiple bins. If a systematic decreases the number of events in some bins but increases them by the same amount in other bins while not affecting the overall number of events, it is referred to as a ‘shape’ effect. A systematic can have both a normalisation and a shape effect.

Adding additional bins means the plots shown above would be made for every bin. In the analysis that consists of multiple regions with many bins in each region, this is a large output and not one easy to parse. The number of plots can be reduced to a manageable amount by combining the effect of a systematic in all the bins where that systematic is present. Normally this means either a sum via weighted averages or quadrature addition.

So in addition to the plots shown above there are two more types of plot that are produced as a result of running the fit: one where only the normalisation effects of the pull contributions are shown, known as NormVal plots; and one where only the shape effects of the pulled systematic are shown and these are referred to as ‘shape’ plots.

As mentioned in Chapter 8, for the nominal resolved analyses ($VH, H \rightarrow b\bar{b}$ and $VZ, Z \rightarrow b\bar{b}$) the variables that go into each Region of Interest (ROI) are fed into a BDT and the output from -1 to 1 is binned. It is these multivariate bins that are entered into the fit. In the Boosted analysis, however, only the key discriminating variable is entered into the fit. A shape plot will show the effects of a particular systematic in all of the bins that make up a particular ROI. An example of a shape plot is shown in Figure 9.17.

The features in the plot of Figure 9.17 will be gradually explained, but the main aim of the plot is to communicate the $+1\sigma$ fluctuation (in solid red) and the -1σ fluctuation




Figure 9.17: Illustration of the possible effect a systematic can have when it is present in multiple bins. The shape effect of the systematic is split into each sample and this plot is the effect of the first eigenvalue for the Jet energy resolution systematic on the diboson sample in the Boosted analysis. The dashed horizontal black line is the nominal MC simulated events prediction. The dashed red and blue lines are the unsmoothed $+1\sigma$ and -1σ variations of the systematic in each bin respectively. The solid red and blue line correspond to the smoothed (see Section 9.3.3) $+1\sigma$ and -1σ fluctuations of the uncertainty with respect to the nominal prediction. The left y-axis is the percentage shift in the events the systematic lines represent. The blue data points in this bin with the right y-axis being the number of events in the bin. The beige hashed lines are the errors on the number of MC simulated events events in each bin. The plot also contains information about the normalisation effect for this uncertainty in all these bins.

(in solid blue) of some systematic in every bin of a specific region (in this case the 2-jet $150 \text{ GeV} \leq p_T^V < 250 \text{ GeV}$ Signal Region of the Boosted analysis). These plots are useful for identifying the effect of a systematic on a particular sample in the fit.

The purpose of a NormVal plot (shown here as Figure 9.18) is to show the normalisation effect of a systematic in all the regions of the fit at a glance. The normalisation on each of the samples in these regions are taken separately.

Each entry in the NormVal table that has a value will have an associated shape plot.




Figure 9.18: Image showing the normalisation effect of a systematic SysJET_CR_JET_JER_EffectiveNP_1. Each sample present in a region is separated on the y-axis, and the (0L) Standalone fit regions are shown on the x-axis. Each number in the boxes is a percentage. Spaces that are empty within rows that have values in them are empty because those effects were pruned (see Section 9.3.4). If the entire row is empty this means that the systematic has no effect on that region whatsoever, or in the case of data; the effect is purposefully hidden. This systematic appears to have the largest effect in the 2-Jet $150 \text{ GeV} \leq p_T^V < 250 \text{ GeV}$ Low ΔR_{bb} control region. The size of these effects are more easily judged by looking at the number of events that this systematic can effect, which is accessed in a shape plot.

The NormVal table in Figure 9.18 encompasses the shape plot represented in Figure 9.17. The diboson normalisation effect shown in Figure 9.17 corresponds to the third value in the diboson row in Figure 9.18. This plot is useful for showing which samples are most affected by the effect of a given systematic.

To reduce the number of output plots, nominally, each systematic will be represented as one point in a pull plot and one entry in the correlation plot. This means that systematics that act over many bins are amalgamations of several individual pull effects.

Often, if an amalgamated pull is seen to be high, this is due to the effect of one particular high pull in a particular region. To check that this is the case, a systematic can be ‘decorrelated’ and its effect can be broken into several independent pulls acting in mutually exclusive bins and or regions. This allows more extreme pulls to be easily identified.

9.3 Systematic Manipulations

The systematic variations can be wide-ranging and numerous, but not all are impactful. To save time and resources, the systematics can sometimes be combined, simplified or even ignored. The last thing, therefore, that needs to be discussed before delving into the fit itself, is what happens to the systematics before and during the fit to reduce their number.

9.3.1 Eigenvector Decompositions

Some systematics that impact the fit do so via the third term in Equation 9.6. This is to say that they vary the weights associated with the event generation. Such examples include the systematics associated with the assigning of quark flavours to jets (flavour tagging - like the SysFT_EFF_Eigen_C_0 in the shape plot of Figure 9.17) and those that deal with the energy resolution or energy scale of jets in the analysis (JER or JES).

Using the flavour tagging as an example: when calibrating the b -tagging efficiency $t\bar{t}$ events are used. For each [jet, pt, eta] bin the MVA2c10 score is fitted using a likelihood constructed based on the probability that a jet is actually a b -jet given its assigned flavour. This fit is repeated for hundreds of systematics variations. In order to have the most accurate results all the systematic variations should be included, but the addition of hundreds of systematics is cumbersome.

The eigenvector decomposition [156, p49] is a way of reducing the number of systematics without reducing the information contained within them.

First, all of the hundreds of systematics, are grouped into various categories and ranked. In the case of the JET systematics for example, they are grouped into physical categories; like modelling effects and statistical uncertainties. In each of these categories the systematics are ranked according to impact with 1 being the most impactful and N being the least.

Secondly, partial covariance matrices are constructed, corresponding to each systematic, where the sums of these partial covariance matrices are taken to obtain the total covariance matrix.

Thirdly, these systematics are ‘decomposed’. This means that the grouped systematics are turned into linear combinations of uncertainties that are correlated with one other. The partial covariance matrices are combined in the same way to form one correlation matrix for the new combined systematics.

This final covariance matrix for these systematics has to be diagonalised to obtain the correlations (i.e , it needs to have an inverse). Being a symmetric, positive-definite matrix this can be considered as an eigenvalue problem. The eigenvectors that “solve” this problem can be seen as “directions” in which to carry out independent variations. The decomposition, therefore, can be thought of combining the systematics in a way such that the resulting covariance matrix is diagonalisable.

The end result is that these hundreds of systematics are grouped together to form a smaller set of eigenvector decompositions which behave like other systematics. These will be marked by number suffixes on the end of systematics according to their impact

(0 being the one with the most impact, and increasing number for decreasing impact) like `FT_Light_0` and `FT_Light_1`.

It is important to note that it is no longer clear what physical response each eigenvector decomposed systematic resembles. The input systematic uncertainties are meaningful in the sense that one corresponds maybe to the difference between “MC Generator 1” and “MC Generator 2” but after the decomposition they are linear combinations of each other. To get the final systematics that appear on the plots one may take the N most important eigenvalues separately and then add all the other eigenvalues in quadrature to form the $N+1$ th systematic.

The decompositions that are done over common $t\bar{t}$ samples are not unique to the analysis. They are also subject to changes in the latest calibration tools, so it is useful to have the decomposition scheme be uniform across the whole experiment. Hence the decorrelation schemes are defined by CP groups and are computed externally to the analyses. The weight variations that correspond to eigenvector decompositions are given to the analysers and these systematics enter the fit via the “Corrs&Systs” file.

9.3.2 Symmetrising

Systematics in general are not by definition symmetric, though it is expected that most systematics have upward and downward fluctuations that behave similarly. However some systematics are only defined as positive weight variations. These will be by definition symmetric, and to save memory, only one effect is encoded. In order to proceed, the downward variation is defined as the negative effect of the positive contribution.

9.3.3 Smoothing

After symmetrising, the next effect is smoothing. Smoothing is an averaging procedure to remove statistical fluctuations in nominal samples that have a small number of

events.

Imagine in all cases the statistics on the nominal MC simulations samples are infinite. If the systematic is a weight variation, each event is varied in the nominal, such that the systematics acts on the same infinite statistics as the nominal.

When the NLL is obtained by dividing the systematic contributions, two effects with infinite statistics are being divided against each other. It is expected that any uncertainty trend is perfectly smooth and without any need for modification. If the systematic is an alternative sample with only a few events, a fluctuation in the background could be seen solely because the systematic was determined in a low statistical environment.

Smoothing is conducted because it is expected that the trend of systematics across a range of bins is both correlated between the bins and roughly monotonic. Hence rapid fluctuations between neighbouring bins are not expected.

Systematic shapes that are rapidly changing across a range (like those represented by the dashed red and blue lines in Figure 9.17) are smoothed in three steps.

- 1) The ratios ($\frac{_{1up} - nominal}{nominal}$) for each of the problematic bins are merged until the resulting correction function is roughly monotonic (there is at most one outlier left). Here $_{1up}$ is the 1σ upward fluctuation of the systematic and *nominal* is the nominal value of the systematic.
- 2) A second round of merging is conducted to reduce the statistical uncertainty in the resultant bins to under 5%.
- 3) The merged ratios are applied to their respective original bins to preserve the bin template. This can be seen in Figure 9.17) as the last 9 bins (and independently the second, third and fourth bins) have been assigned the same variation.

If a systematic is smoothed, the pre-smoothed line is shown on the shape plot as a dotted line. For b -tagging, and many others, the variation is a weight on the nominal sample, therefore it has the same statistics as the nominal MC simulations which should be large, and therefore the smoothing can be skipped.

9.3.4 Pruning

In the middle of the fitting process, systematics whose impact on μ fall below a certain threshold are cut. The level of pruning is different between bins deemed sensitive to the signal and those that are not.

9.4 Post-fit Modelling Investigations

The aim of the fit is to correct the data-to-Monte-Carlo-event ratio (data/MC) by shifting the nominal distribution of MC simulated events. This is done by determining which of the various systematics in the fit to pull, and by how much. After the fit model is generated, the pulls are used to recalculate the data/MC by applying the corrections to the Monte Carlo templates.

If a systematic is pulled to a greater extent than expected, then this could indicate some form of modelling problem or *mismodelling* that needs to be corrected. The correlations between systematics make these corrections more tricky, as the fit may compensate for one pull by pulling systematics in another bin.

In reality, most of the systematics that are introduced to a model should have small effects. Hence any systematic pulled more than about 0.5σ , despite the penalty, requires some sort of investigation. Sometimes this means finding the source of the pull in a systematic and fixing it, and other times it means identifying the source of a pull in a

systematic and finding a physics justification for its presence.

Before the final unblinded fit is run on data in the signal region, which outputs the definitive analyses answers for μ , the fit is run multiple times without data in sensitive bins (not necessarily only in the SR) to solve these problems. Since in both the Resolved analysis and the Boosted analysis, each channel acts as a mostly separate measurement, the fit can also be run on each channel separately.

The majority of initial investigations into outlying or unexpected pulls come from decorrelating them into sub-regions. Decorrelated pulls are effectively independent of one another so they can vary freely. In the majority of cases if the outlying behaviour is not seen in most of these decorrelated pulls, then the overall pull when combined is best represented by the extrema of the pulls in the various fits.

The fit model is then examined again with the decorrelated pulls to see which, if any, other pulls are affected. If there are changes in the pulls, then there could be unforeseen correlations that could help to explain the pulls. If no other pulls are seen to change this points to a mismodelling problem.

The next part of this section will go through the set of graphs and decisions taken that mark some of my work tracking the origin of a particular pull to improve the fit model for the VH , $H \rightarrow b\bar{b}$ 0L channel of the Resolved analysis.

9.5 The 0-Lepton Resolved VHbb Fit

In the VH , $H \rightarrow b\bar{b}$ analyses, since each lepton channel is optimised separately, it makes sense that each channel can be fitted independently. These are referred to as the 0L, 1L or 2L ‘standalone’ fits (depending on the lepton channel). Each of these standalone fits

can also be run in conjunction with each other. The full fit is the combination of all three standalone fits, and is hence referred to as the Combined 012L Fit.

As detailed in the previous chapter, the main Resolved analysis standalone 0-Lepton fit is a multi-variate analysis. The variables measured in each of the events are put into a BDT, the output of which is a single variable. For the main analysis this is called the BDT_VH and this enters the fit.

The 0-Lepton channel has 12 regions fitted in parallel. With 15 bins in the 2-Jet and 3-Jet $150 \text{ GeV} \leq p_T^V < 250 \text{ GeV}$ signal regions; 8 bins in their respective $p_T^V \geq 250 \text{ GeV}$ bins; and 1 bin in every control region; there is a total of 54 bins. All channels and regions are run simultaneously and each region knows about each other regions. Each additional region essentially acts like another set of independent bins. Cross-channel systematics are determined from background processes and these are also run at the same time⁵. The main output plot from such a fit is given by Figure 9.19.




Figure 9.19: Example post-fit plots from the 0L standalone fit. The $150 \text{ GeV} \leq p_T^V < 250 \text{ GeV}$ region is on the left and the $p_T^V \geq 250 \text{ GeV}$ region is on the right. These plots were taken when the fit was blinded, so the most sensitive bins in the fit have no data points in them.

⁵There are a lot of W+jets events in the 1-lepton channel, so if there are normalisation effects from this process, these are best measured there and ported into the 0-lepton channel.

9.5.1 The 0L Flavour-tagging Fit Study

Starting from an observation of an outlying pull, a fit study will consist of several tests to determine the cause (or causes) of a pull (or set of pulls). In the following fit study, the starting point is an excess in one systematic in the Combined 012L fit.

Typically all unexpected pulls in the fits are considered, but the following flavour tagging one was the most substantial one studied and serves as an example of the general procedures involved in investigating pulls.

Figure 9.20 shows the flavour tagging pulls for the initial set of pulls produced by the Combined 012L fit. It is here that the excess in the systematic FT_Light_0 was seen and attributed to the 0L standalone fit. Figure 9.21 is a flowchart showing a quick overview of the actions and observations taken at every stage of the study.

Since the fit is complicated there are many fit studies that are undertaken at the same time, the outcome of which could affect any one study. These tangential studies will be ignored here.

The flavour tagging systematics FT_B_XX, FT_C_XX and FT_Light_XX are uncertainties on the b -tagging efficiency for b -jets, c -jets and light jets respectively. They relate to the fact that in the $t\bar{t}$ MC simulations on which these systematics are trained, for a given MV2c10 cut, there is a given efficiency for b -jets, c -jets and light jets. This efficiency is the probability a given jet is tagged as a b -jet with the b -tagger given that it is actually a b -jet, c -jet or a light jet.

These given efficiencies could differ from data for any number of reasons. The efficiencies are measured in data, and the MC simulated events are ‘calibrated’ to match




Figure 9.20: Pull plot for the systematics related to b-tagging in the Combined 012L fit. For each systematic both the combined and standalone channel fits were run. The 0L is in red, the 1L in blue, 2L in pink and the 012L in black. A pull centred in the green will have a pull of $-1 \leq x \leq 1$, and those centred in the yellow will have a pull value of between $1 < x < 2$ or $-1 < x < -2$. This set of pulls shows interesting excesses in FT_Light_0 and FT_B.21 systematic. Only the investigation into the former will be discussed.




Figure 9.21: Fit study flow chart. Fit tests are represented by blue arrows. The initial observation is in darker green, with later observations pertaining to the editing of the Light_0 systematic in a lighter green. Observations marked in orange and purple represent the study's evolution into pruning tests, and fixing new pulls caused by the new changes, respectively

it. This calibration is done independently for b -jets, c -jets and light jets as they have different efficiencies. Hence the FT_B_XX, FT_C_XX and FT_Light_XX eigenvalues are a reduced set of NP's from these calibrations.

The next port of call is to perform a 'full decorrelation' of the Light_0 pull in the fit. Since the 0L standalone Resolved analysis has a signal and two control regions for $150 \text{ GeV} \leq p_T^V < 250 \text{ GeV}$ and $p_T^V \geq 250 \text{ GeV}$ for events with both 2-jet and 3-jet final states, there are twelve possible combinations of region present in the fit. The full decorrelation of a systematic means splitting the effect of the systematic into each of these regions separately to form twelve independent systematics.

Figure 9.22 shows that when this was done the overall pull is dominated by the effect of the systematic in the 3-Jet $150 \text{ GeV} \leq p_T^V < 250 \text{ GeV}$ signal region. To check that this effect was consistent, the next check was to de-correlate the systematic into each of the p_T^V , number of jets and signal-control region separately. The result of this selective decorrelation on the Light_0 systematic can be seen in Figure 9.23.

Figure 9.23 shows the individual correlations are slightly at odds with that of the full decorrelation. While the pull trends still indicate that the pull is dominated by the signal region and the three-jet event region, the p_T^V category which has the largest impact on the pull is the $p_T^V \geq 250 \text{ GeV}$ region and not the $150 \text{ GeV} \leq p_T^V < 250 \text{ GeV}$ region. This inconsistency most likely implies that the effect causing the pull in the systematic happens to manifest in this region, rather than being a mismodelling effect unique to that region.

The clearest result as an occurrence of the individual decorrelation tests, is the conclusion that the pull comes from the signal region, and slightly pulls the control regions down to partially remedy this. However it is unclear whether the effect is indicative of




Figure 9.22: B -tagging pulls comparing the fits of the nominal fit (in black) with the fit where the Light_0 pull is fully decorrelated (in red). The name of decorrelated systematics are too long to fit on the plot, so the region part has been cut off. The naming convention is such that the systematics are organised into 4 sets of three regions. In each three the first region is CR_{Low} , the second is CR_{High} , and the final one is the SR. $J(X)$ refers to the number of jets in the category ($J3 = \text{a 3-Jet event}$). BMin150 corresponds to the $150 \text{ GeV} \leq p_T^V < 250 \text{ GeV}$ region and BMin250 corresponds to the $p_T^V \geq 250 \text{ GeV}$ region. Here the overall pull is coming from one region: The 3-Jet $150 \text{ GeV} \leq p_T^V < 250 \text{ GeV}$ SR.




Figure 9.23: Light_0 flavour pulls in three different fits. In all fits the nominal fit is in black and the respective trial fit is in red. The top fit shows the Light_0 systematic decorrelated in only signal-control region, the middle fit shows the Light_0 systematic decorrelated by jet category, and the bottom fit shows the Light_0 systematic decorrelated into the p_T^V category.

the presence of something awry in the signal region or a lack of something that should be there.

By merging the signal and control regions and repeating this test, this can be identified. If the offending pull stays the same when the signal and control regions are merged, then it shows that there is something in the signal region that should not be there. If the pull is reduced, then the presence of the control region events could be seen to ‘provide’ the merged signal region with an effect that is not present in the separated signal region. the pulls from this test are present in Figure 9.24.




Figure 9.24: Selection of flavour tagging pulls from fits where the control and signal regions are separated (black), and where the signal and control regions are merged (red). In the latter fit the outlying Light_0 pull is reduced to a far more reasonable level.

The reduction of the Light_0 pull when the signal and control regions are merged in Figure 9.24 shows that the effect is probably due a lack of something in the signal re-

gion.

In order to start narrowing the list of fit output plots to peruse, it is prudent to ascertain the nature of the effect; that is to say whether the pull is a normalisation effect, a shape effect, or a combination of the two. The simplest way to do this is to perform a single bin fit. In each of the twelve regions, the MVA classifier has multiple bins between -1 and 1. Merging all these bins into a single bin removes any shape effects of the systematic. If a pull does not differ from the nominal value in the single-bin fit, then this implies that the systematic is wholly a normalisation effect. If the pull is completely reduced to zero, then this implies the pull is wholly a shape effect. If the pull is reduced to any non-zero value, then it is likely a combination of normalisation and shape effects; the dominant effect being determined by whether the resultant pull is closer to 0 or the original pull value.




Figure 9.25: *Light_0 (and Light_2) flavour pull(s) for two different fits. The fit with the normal number of mva bins in each region is in black and the fit with only one mva bin is in red.*

Figure 9.25 shows that this effect is mostly a normalisation effect, but has a significant shape component. This means that both shape and normalisation plots will have to be looked at to see if any conclusions can be drawn.

Returning to the nominal fit, the shape plots will first be looked at. Figure 9.26 shows the three largest shape contributions. Since the MVA is trained on what the signal

should look like, the last few bins will contain the most signal, and therefore consist of events that go into various signal regions. Figure 9.26 shows that of these shapes, the sample with the largest up-down variations in these final bins is the $t\bar{t}$ one (labelled $t\bar{t}bar$). The diboson has a middling effect across the entire MVA range, and the VHSTXS (signal) has a large fluctuation in the first two MVA bins where there is little signal.




Figure 9.26: Image showing the shape and normalisation effects of the systematic `SysFT_EFF_Eigen_Light_0` or (FT_Light_0 for short) in the signal (VHSTXS), $t\bar{t}$ ($t\bar{t}bar$) and diboson samples binned by their MVA scores. This systematic has the largest impact in $t\bar{t}$ events that look signal-like.

If the $t\bar{t}$ sample is the main contributor to the pull effect, it must also have a significant normalisation effect. When returning to look at the NormVal plots (Figure 9.27), the normalisation effect of $t\bar{t}$ in the affected region is small compared to those in the samples Zl, Zcl, Wl and Wcl. These background samples correspond to the alternative situations where the boson that accompanies the Higgs candidate is reconstructed as either two light jets (Wl,Zl) or a c -jet and a light jet (Zcl,Wcl). The normalisation effect seems to be roughly 30% per light jet.




Figure 9.27: Image showing the normalisation effect of the systematic SysFT_EFF_Eigen_Light_0 (FT_Light_0 for short). This systematic appears to have the largest effect in the 2-Jet $150 \text{ GeV} \leq p_T^V < 250 \text{ GeV}$ Signal Region and the 3-Jet $p_T^V \geq 250 \text{ GeV}$ Signal Region. Since the analysis was blinded at this point, there are no normalisations recorded for data in any region.

To confirm that this effect is collectively from Zl, Zcl, Wl and Wcl (which shall now be referred to as V+light), the next test was to isolate the samples where light jets are a prominent feature from those where they are not. At the first instance, this test included the samples Zbl and Wbl, but these are included in the samples Zhf and Whf (heavy flavour), which contain any and all samples that contain a b -quark (or have two c -quarks). These samples also therefore include the Z-boson (W-boson) decaying to $b\bar{b}$, $c\bar{c}$ and bc as well as bl .

The one-bin fit will be run with the Light_0 systematic decorrelating V+jet (V+light + V+hf) backgrounds from the rest of the backgrounds. The resultant pulls can be seen in Figure 9.28, and the resultant normalisations are in Figure 9.29.



Figure 9.28: Light_0 (and Light_2) flavour pull(s) for two different fits. The single *mva* bin fit is in black and the single *mva* bin fit with the V+jets background samples decorrelated in the Light_0 systematic in red. When de-correlated in this manner the black and red systematic points for FT_Light_0 are not directly comparable, as the red point is essentially FT_Light_0_non-Vjets.




Figure 9.29: Comparison between the NormVal plots of the single bin fit (left - Figure 9.27) and the V+jets decorrelated single bin fit (top right and bottom right). The two sets of NormVal plots on the right are independent. The bottom right plot has the normalisations from all the V+jets samples in Light_0 and the right plot has the rest of the normalisations in Light_0. When compared to the single bin fit, it can be seen that the large majority of the normalisations do not change. The exception to this are the values circled black. These are the 3-jet V+hf normalisations which have all varied by less than one percent.

Figure 9.28 shows that the effect of the normalisation pull is shared between the V+jet

backgrounds and the non-V+jet backgrounds. This is attributed to the fact that (light) jets are a feature in other backgrounds outside of V+light ones (i.e top backgrounds have light jets).

To understand what kind of effects changes in the light jets can have on the fit, the light jet normalisation uncertainties (here referring to the $\vec{k}$ from Equation 9.6) can be changed. These are variables that the fit can manipulate that have wide-ranging effects in samples in the entire fit.

Figure 9.30 shows a selection of pull changes that occur when the light jet pre-fit normalisation uncertainties in the fit were doubled. Figure 9.30 also shows that the fit pulls down the normalisations (WclNorm, ZclNorm) to partially correct for the doubling of the initial light normalisations. In addition to this the fit also prefers to pull the Light_0 systematic down to correct for this effect.




Figure 9.30: Relevant pulls that have changed in the fit where the normalisation factors have doubled (red) compared to the nominal fit (black).

The next test is a compilation of results to see if a combination of effects from other pull studies can have any marked effect when taken together. In this case it is the number of jets and the samples present in the regions. The samples were split into those that

contain light jets and those that do not. This next set of de-correlations is undertaken with the doubled normalisations. These can be seen in Figure 9.31.




Figure 9.31: Pull plot where the nominal fit (black) is compared to the fit decorrelated in both n -jet and V +light categories and the normalisation systematics are doubled (red).

Figure 9.31 shows that in this decorrelation scheme, there is an overall reduction in the Light.0 pull. All new systematics are pulled in the same direction and the pull with the largest contribution is the non- V +light 3-jet pull at 0.5 sigma. The trend of pull being dominated by the non- V +lights has gone. The normalisation systematic pulls are pulled down further with respect to the fit when only the normalisation systematics are doubled.

All of these effects seem to confirm the hypothesis that there is some characteristic about events with light jets that is contributing some sort of fit effect in the signal region. To test this more, it was decided to attempt to try and exaggerate this effect; the idea being to add a systematic into the fit that acts like the Light.0 systematic but has a larger impact. This means that the fit should pull these new systematics in favour of the Light.0 ones.

Figure 9.25 shows that the maximum normalisation impact of Light_0 is about 60% so the initial normalisation of this new systematic should be larger than this for this new systematic to be pulled preferentially as the two systematics have an overlapping remit. This normalisation was set to 70%. The pulls can be seen in Figure 9.32.




Figure 9.32: Pulls from the Nominal 0L fit (black) compared to a fit where an additional systematic is added (red). This additional systematic is called DummyLight0 and is designed to look to the fit like a more amenable version of the Light_0 pull.

Figure 9.32 shows that the addition of the new systematic only sees a small reduction in the Light_0 pull, even though the added systematic is pulled significantly. Also, some of the other flavour pulls are reduced as well.

Figure 9.25 also shows that the Light_0 components outside the signal region are pruned. The effective pull could be an artefact of reduced statistics due to the fact that events that fall into these categories are being pulled.

9.5.2 0L Pruning Study

As mentioned earlier, in a fit as large as this, systematics can be pruned if their post-fit impact is small. In addition to this, the analysis has another criterion, where uncertain-

ties for samples below a particular cut-off size are pruned irrespective of the size of those uncertainties.

There are two pruning thresholds for the fit: one for bins deemed ‘sensitive’, and one for those deemed ‘not-sensitive’. If a bin is sensitive, then it contains a considerable contribution from the signal. The figure of merit for judging sensitivity is the ratio of signal over background (or S/B). If the S/B is greater than 0.02, then the bin is a sensitive one, and the threshold for the impact of a systematic for it not to be pruned is 2%. The threshold for pruning in non-sensitive bins is 0.5%.

A possible explanation for pruning in Light_0 having a large effect on the pulls is that the systematic corresponds to light components in backgrounds dominated by light jets. These backgrounds tend to be small in the 0L analysis. However, since Light_0 is a large uncertainty dedicated to light components, it should be kept even for small components, as no other systematics can compensate for changes or effects in light jets.

Given the numbers of Figure 9.25, the removing of pruning was expected to have one of two effects. The first is that the normalisation effect is spread throughout all regions where the effect is not pruned, thus vastly reducing the pull’s dependence on normalisation effects. The second is that every region allowed will follow the same trends as the SR values that were not pruned. If the latter effect occurs, then this makes the case that the pruning of this systematic is removing an important effect from the fit. The normalisations of the no pruning fit are shown in Figure 9.33.

Figure 9.33 shows that the systematics segments that were formerly pruned exhibit the same behaviour across all regions. To make the case of the removal of pruning from the fit, one must examine the pulls from the fit and a large selection of these can be found in figures 9.34 and 9.35.




Figure 9.33: Normalisation factors from all the samples separately for each region in the 0L fit. Similarly for Figure 9.27, all the samples with prominent light jet contributions have a rough normalisation of 30% per light jet.

Figure 9.34 shows that with the removal of pruning, the Light_0 systematic is reduced. To see whether a fit with no pruning should be adopted, the fit needs to be stable with this change. This means examining the other pulls in the fit as well. Some of the more interesting ones can be seen in Figure 9.35.

Figure 9.35 shows that while most of the pulls are largely consistent in the ‘no pruning’ regime, there are a few pulls of concern, the first being that two of the diboson p_T^V pulls (VVPTVME and VVPTVPSUE) have become ‘overconstrained’. These will be investigated in the next section.




Figure 9.34: Selection of flavour tagging pulls from the Nominal 0L fit (black) compared to a fit where no systematic are pruned (red). The removal of pruning reduces the pull of the *Light_0* systematic.

In the selection of pulls on the right hand side of Figure 9.35, it can be seen that some of the jet pulls have become underconstrained. These were successively investigated and was found to be related to strange patterns in input fluctuations. However given the sheer number of systematics now present in the fit, it is fewer than could be expected and none of these pulls was expected to affect the signal to any significant level. Generally this behaviour is indicative of small components while the systematics on large samples is fine. In the case of the 0-lepton standalone this seems to be the case, but then in 1- or 2-lepton fits there is a component large enough in the same region to find the correct “no-constraint”. These pulls will be touched upon again later.

In order to adopt a fit with no pruning at all, a few things need to be considered. The first issue is one of speed. The introduction of small effects makes the workspace with the fit large, and the time it takes to run the fit could be much longer. Speed tests were conducted in the 0L analysis with several different pruning regimes, and it was found that running the fit without pruning in this way was only 10% slower.




Figure 9.35: Selection of non-flavour tagging pulls from the Nominal OL fit (black) compared to a fit where no systematic are pruned (red). The lack of pruning results in some systematics with small effects now visible. The interesting pulls are those that change or become more or less constrained when pruning is removed.

Another thing to check is that the effect of the removal of pruning is stable. This is to say that the pull effect in terms of both normalisation and shape, gets smaller in higher-order Light_X systematics. To this end, a series of shape plots for all the Light_X systematics were looked at. A selection indicative of the effect trend can be seen in Figure 9.36.

Figure 9.36 shows that higher orders of Light_X values, generally, both the normalisation effect values and the magnitude of the shape effect decrease. Light_1 has a far lower normalisation effect than Light_0 with roughly a 3% effect per light jet as opposed to the rough 30% per light jet present in Light_0. This means that the decision to remove pruning is stable in the fit.

The last and most important aspect of any change is to check the effect these changes can have on the significance of the fit. Table 9.1 shows the changes in significance. Here it can be seen that implementing this change has a negative impact on the fit. This serves as an indication that the removal of pruning in its entirety is probably a bad thing for the analysis as a whole, and that a less extreme method pruning reduction should be considered. This decision was not taken at this moment as the fit was in a state of flux.

Table 9.1: Table showing the significances of the baseline standalone 0-, 1- and 2-lepton fits and the significance when the pruning is fully removed. In all analysis channels, the change results in a small loss in sensitivity.

Channel	0L	1L	2L
Baseline Fit	4.11	3.78	4.43
No Pruning Fit	4.06	3.73	4.30
% Change	-1.2%	-1.3%	-3.0%

While removing categories of pruning in the fit fixes some of the problems, however a couple of new ones arise. One of them is the overconstraining of a couple of pulls.




Figure 9.36: Comparison of shape and normalisation effects for the systematic eigenvalues Light 1 (top row), Light 2 (middle row) and Light 5 (bottom row). The four plots chosen for each eigenvalue are all from the 150 GeV $p_T^V < 250$ GeV signal region and for each row follow the ordering from left to right: Zl-2jet, Zl-2jet, Zl-3jet, Zl-3jet.

9.5.3 Diboson p_T^V Fit Study

Figure 9.35 shows that the removal of pruning in the fit, constrains some diboson transverse momentum systematics; specifically the diboson p_T^V matrix element uncertainty labelled VVPTVME, and the diboson p_T^V parton shower underlying event uncertainty labelled VVPTVPSUE.

To investigate these pulls no additional fits need to be run at first. Figure 9.37 shows the comparison between the normalisation values for both VVPTV systematics in each region of the fit with and without pruning, the normalisations that are present in both fits are identical, and the 2J systematics that appear have a normalisation of about 30%.




Figure 9.37: NormVal plots for the two systematics VVPTVME (left) and VVPTVPSUE (right) without pruning (top) and with pruning (bottom) for the 12 regions in the 0L standalone fit.

To get the full picture of the pull effect, the shapes of the uncertainties that appear when pruning is turned off should be looked at. These can be seen in Figures 9.38 and 9.39.




Figure 9.38: Figure showing the shapes for all the uncertainties for VVPTVME that appear in the fit when the pruning is removed. Only the systematics that appear in the SR have more than one bin.




Figure 9.39: Figure showing the shapes for all the uncertainties for VVPTVPSUE that appear in the fit when the pruning is removed. Only the systematics that appear in the SR have more than one bin.

Figures 9.38 and 9.39 show that there is a large discrepancy between the upward and downward fluctuations of the systematic. The thing to check after seeing this are the upward fluctuations of the systematics as stored in the diboson samples (WW , WZ and ZZ). In the 0-lepton the ZZ sample is the most important sample in the diboson contribution. The upward systematic fluctuations are stored as $\frac{1up - nominal}{nominal}$. Therefore it is expected that this distribution be centred roughly around 0.




Figure 9.40: Plots of $\frac{1up - nominal}{nominal}$ for the WZ (left) and the ZZ diboson sample (right). $1up$ is the 1σ upward fluctuation of the systematic. Within errors this is expected to be distributed around 0, which is seen for the WZ sample dominant in the 1-lepton analysis. The ZZ sample is the dominant contribution to the diboson sample in the 0- and 2-lepton analyses.

Figure 9.40b shows there is a ‘bug’ in the way the systematic is stored since the ratio $\frac{1up - nominal}{nominal}$ is centred around 1. The systematics effect in terms of number of events is compared for the ZZ sample in the current and the previous version of the analysis (also known as the ICHEP analysis) to see whether this is a new bug or not. This is shown in Figure 9.41.

According to the Figure 9.41, the sample bug is not new. The reason for the bug was identified as an erroneous scale in the input production, the result of which means that the samples need to be re-run with the bug fixed. Until the next set of samples were produced, this problem was temporarily resolved by hard-coding the correct normali-




Figure 9.41: Plots of the number of events that correspond to the 1σ upward fluctuation of the systematic ZZPTVPSUE in the previous (a) and current analysis (b). In (a), the nominal distribution is in red, and the upward fluctuation is in blue. In (b), the nominal distribution is in blue, and the upward fluctuation is in red. In both analyses, there seems to be double the number of events expected in the upwards fluctuation.

sations for ZZPTVPSUE, and removing the ZZPTVPSUE systematic from the fit.

9.6 Future Changes and Summary

The final fit model has changed since these studies were undertaken, but the fit effects described in this chapter were still relevant: in particular the decision was taken to modify the pruning of uncertainties related to small backgrounds with light jet components.

In the end it was decided to remove the pruning of systematics in samples with low statistics for Light_0 and the normalisations on light-enriched backgrounds but keep the pruning in the rest of the fit. The removal of pruning entirely has far more wide-reaching consequences. A fully un-pruned workspace, while maybe not taking long to run is still a huge object to process.

Results from analyses like this often have long usage lifetimes because the workspace is designed with the future in mind. One of the uses of the analyses results is the combina-

tion of the result with other analysis to make larger-scale measurements. This requires the combination of several workspaces to produce results. So while decisions can be made to improve the accuracy or significance of one particular analysis' result, these have to be weighed up with the additional cost in terms of the personpower, time, and understanding of other external analysers that want to use the result in the future.

When it comes to further understanding the pulls, the job of the analysers is not finished. This is because the next step is the combination of the Resolved analysis and the Boosted analysis to form the 'Legacy analysis'. One of the objectives for the Legacy analysis will be to improve aspects of modelling pulls in the fit.

Now that we have a grasp of the fit in its entirety, the last main chapter of the thesis will evaluate a potential improvement to the Boosted analysis.

E_T^{miss} triggers in the 2-Lepton channels

10.1 Motivation

Signal regions in the 2-lepton channels of the Resolved analysis and the Boosted analysis have the lowest signal yields [4, p14] of all the lepton channels in the analyses. With regards to the 0-lepton channels, the Z-boson preferentially couples to neutrinos compared with charged leptons [9, p11-12], and of the charged leptons, the analyses are only focussed on the detection of electrons and muons as taus can decay hadronically so can look like jets. This means that a-priori, there should be three times more events¹ in the 0-lepton signal region than in the 2-lepton signal region.

The yield of the 1-lepton channels are larger because the production of a W-boson from

¹Decay Branching fractions: $Z \rightarrow l\bar{l} \approx 10\%$, $Z \rightarrow \nu\bar{\nu} \approx 20\%$ but the analysis does not include taus, which means that the visible decay fraction of $Z \rightarrow ll$ is $\frac{10 \times 2}{3}\%$. The relative difference in the number of signal events between the 0-lepton and 2-lepton analyses is given by the analysis-visible decay ratio: $\frac{Z \rightarrow \nu\bar{\nu}}{Z \rightarrow ll} = \frac{20}{\frac{10 \times 2}{3}} = \frac{20 \times 3}{20} = 3$.

colliding protons does not require an interaction with a sea anti-quark, whereas the production of the Z-boson in a similar manner does. This vastly reduces the probability of $VH, H \rightarrow b\bar{b}$ where the vector boson is the Z-boson with respect to the production of the W-boson even though the Higgs boson has a larger coupling to Z-bosons. This, coupled with the fact that the semi-leptonic decay of the W-boson into $e^-\nu_e$ or $\mu^-\nu_\mu$ is $\approx 21.5\%$ - the largest vector boson signal-like decay, means that there are more 1-lepton events than 0-lepton and 2-lepton combined.

A lower signal yield means that the 2-lepton analysis has to work a lot harder to find significant deviations in the background. One of the ways of improving the significance that the 2-lepton channel brings to the entire analysis is to see if changing the trigger decisions in the event selection provides the 2-lepton channel with more raw events.

Since the start of the $VH, H \rightarrow b\bar{b}$ analyses, the 2-lepton channel has used single-lepton triggers in the electron sub-channel and muon sub-channel to define its interesting regions. The other lepton from the Z-boson is then found using event-level kinematic variables and the mass of the Z-boson as a constraint. Mostly this is for convenience. If there is a real Z-boson, and there are highly energetic leptons (electrons and muons) in the event, then there must be two of them, since the Z-boson has no charge. If an event is triggered on these energetic leptons it is almost certain that at least one of them will be recorded.

When a single lepton trigger is used on an event with two electrons in it, the requirement on the p_T^V of both of these electrons to be at least 30 GeV has a flat efficiency close to 100% throughout the detector [157]. The same is not the case for muons because the barrel region of the detector is in contact with the cavern floor (see the summary of the muon spectrometer in Chapter 4). Hence, this low muon solid-angle acceptance means there is not a high efficiency of muon object requirements across the entire detector

(about 80% on average [158]: $\sim 60\%$ in the barrel and $\sim 100\%$ in the endcaps).

10.2 Possible options

10.2.1 Di-muon trigger addition

In the current analysis, the threshold for the single lepton trigger is 26 GeV . One of the advantages of looking at a trigger that requires two muons is that the threshold for the leading muon is lowered to only 14 GeV . This means that some events that the single-lepton trigger missed could be gained. If the trigger is added alongside the single-lepton trigger, most of the time the single-lepton trigger will outperform the di-muon trigger. This means to obtain any appreciable gains the trigger should only be added in the region where the di-muon trigger will out-perform the triggers currently being used. This means to only operate the di-muon trigger in the region between $14 \text{ GeV} \leq p_T^V \leq 26 \text{ GeV}$. The problem with the inclusion of the di-muons trigger is twofold.

Firstly, the efficiency is low because in addition to the poor acceptances for muons it only takes one of the muons to be lost or misidentified for the event to fail the trigger. In addition to this, low energy muons are less likely to be isolated and hence have a higher likelihood of being misidentified. Secondly, the analysis is concerned with muons that come from the Z-boson. Operating exclusively so far away from the pole mass of the Z-boson means that the events that pass this trigger may fail other kinematic checks and hence not feature many new events in the signal region, but could increase the number of background events in the analysis more. Between these two reasons, this approach was discounted.

10.2.2 E_T^{miss} trigger alternative

In addition to the low muon solid-angle acceptance, muons produced at ATLAS are minimum ionising particles (MIPs) [159, p1]. This means that they only deposit small amounts of energy into the electromagnetic and hadronic calorimeters when they pass through.

At the L1 trigger stage, the E_T^{miss} triggers are not triggering on real E_T^{miss} , but ‘calorimeter-only’ E_T^{miss} . Calorimeter-only E_T^{miss} is calculated during the run (‘online’) from data from the ECal and the HCal only, so does not take into account the full energy of the muon which is implied from MS tracks.

Therefore the majority of the energy of the muon is included in the calculation of calorimeter-only E_T^{miss} . Since muons are a source of E_T^{miss} , one method to increase the efficiency for 2-lepton di-muon events is to try and use a E_T^{miss} trigger to recover them.

An approach like this is possible due to the fact that when the full reconstruction is done ‘offline’, energy inferred from track objects in the MS are taken into account and the muons are reconstructed and the calculation of E_T^{miss} is corrected. In order to get to this stage, however, a trigger selection is needed to pass events containing muons, and it does not matter which trigger was used to obtain these events online, just that the final state objects are of interest to the analysis.

The effect of including a E_T^{miss} trigger can be studied because the 2-lepton analysis signal MC simulation samples that are generated have no filters on them. This means that when the events were pre-processed, the information about events that passed E_T^{miss} triggers (and not muon or electron) triggers were allowed to be stored with the collection of 2-lepton events.

Figure 10.1 below illustrates the possible increases in event yield in the 2-lepton analysis by showing what muon kinematic regions could benefit from several different trigger regimes.




Figure 10.1: Graphic showing possible event yields differences with the introduction of certain triggers into the event selection. The x -axis denotes the energy of the leading lepton, and the y -axis shows the energy of the sub-leading lepton. The orange lines correspond to the lowest energy threshold of the current single lepton trigger. The addition of a dimuon trigger lowers the energy threshold of the leading lepton (to the red lines) and allows access to the blue region not yet considered by the analysis. The inclusion of a E_T^{miss} trigger would increase the likelihood an event passes a trigger at whatever point the E_T^{miss} trigger is active. Since the E_T^{miss} trigger is close to fully efficient at $p_T^V \leq 150$ GeV, the E_T^{miss} trigger will only be active in the analysis from this point. Here $p_T^V \sim P_{tL1} + P_{tL2}$ and hence the kinematically-allowed region above the purple line $P_{tL1} + P_{tL2} = 150$ GeV could see an increase in events (shown in dark green).

By asking if an event would be accepted if it passes either one of the single lepton or E_T^{miss} triggers (using a logical 'OR' of both the triggers), it is easy to see that there would be an increase in event yield as there are more trigger decisions made over the same events. This can only result in more events passing at least one of the triggers. This gain has to be offset with the cost of adding another trigger to the analysis. The logical OR of both triggers here introduces some problems due to possible double-counting of events. Additional factors will need to be calculated to combat this. The scale factors

for different regimes is a problem when a region is produced with a logical OR because the uncertainties cannot be described by either trigger scale factor. The correct trigger scale factor would be a non-obvious combination of the two². Therefore it helps to use the triggers exclusively as possible and hence, a useful starting point is to see if there is a possible gain when the muon trigger is replaced entirely by the E_T^{miss} trigger.

10.2.3 Choice of E_T^{miss} triggers

One question remains about what E_T^{miss} triggers should be used. Depending on the period that the data was taken in, different E_T^{miss} triggers are used on the data and the respective MC simulated events.

To reduce the complexity of changes to the analyses and to harmonise with the 0-lepton and 1-lepton channels in the Resolved and Boosted Analyses, the selection of E_T^{miss} triggers to be tested were those already used by the other channels. These can be seen in Table 10.1.

Table 10.1: Table showing the E_T^{miss} triggers used during the Run-II (2015-2018) data collection period. The triggers are seeded using L1_XE50 (L1_XE55) LAr and Tile calo triggers (calibrated at EM scale) with a threshold of 50 (55) GeV. The triggers with ‘mht’ and ‘pufit’ use these algorithms to calculate E_T^{miss} for the second E_T^{miss} filter. The mht algorithm makes use jets constructed at L1 in its E_T^{miss} definition, and pufit, uses topo-clusters and minimises the effect of pile-up.

Trigger name	Period	Threshold [GeV]
HLT_xe70	2015	70
HLT_xe90_mht_L1XE50	2016 (A-D3)	90
HLT_xe110_mht_L1XE50	2016 ($\geq$ D4)	110
HLT_xe110_pufit_L1XE55	2017	110
HLT_xe110_pufit_xe70_L1XE50	2018	110

²This is because the use of both calorimeter-based E_T^{miss} and muon information in the same event can risk the double-counting of the energy of the muons.

10.3 Implementation

Since the E_T^{miss} triggers are not defined in the 2-lepton channel, the first step was to add them into the framework. This entailed creating a function containing the required E_T^{miss} triggers that the 2-lepton channel could call during the event selection stage to extract whether events passed or failed these E_T^{miss} triggers. This would work alongside single-lepton trigger function.

After this, the analysis was run under different trigger regimes to see if the signal region would see an increase in events. However, since the single lepton triggers are highly-efficient for the identification of electrons, only 2-lepton events with muons need to be affected.

The MC simulated events that are affected by these changes are CxAOD's. This means these events have information on what sets of trigger requirements these events have passed at the derivation stage. In addition to this, the events have already been categorised so it is possible to filter MC simulated events by what objects they contain (i.e whether an event has muons in it). If a MC simulated event has two muons, a decision is made on whether to make that event part of the analysis depending on what triggers it passed. This study looked at the effects of replacing the triggers that 2-lepton events are required to pass from the single lepton triggers, to the E_T^{miss} triggers in certain p_T^V regimes.

Often to reduce the size of the output DxAOD, certain trigger requirements that are surplus to the analysis are not stored. For example, since the 2-lepton analysis does not use E_T^{miss} trigger requirements, the information about whether these MC simulated events passed the E_T^{miss} triggers will not be recorded and present in the CxAOD's. This is called *skimming*. If these events have been skimmed to save space, then the effect on

the yield seen when the trigger event requirement is switched from muons to E_T^{miss} will not be correct.

It is important therefore to check that these preliminary trigger requirements have been made on the 2-lepton DxAOD's with both muon and E_T^{miss} triggers. For the preliminary tests only the 2-lepton signal samples were un-skimmed. It is only possible, therefore, to see whether these changes affect the raw signal yield. To get the total effect on the significance of the analysis, a new set of derivations need to be run without the skimming on the background. Figure 10.2 shows the principle of how this can be integrated into the framework. Two options were tested.




Figure 10.2: Graphic showing differences between the new and the older trigger regime. ALU stands for Add Lowest Unprescaled triggers. Two options are presented here for the analysis for after the point where the E_T^{miss} triggers start to become 100% efficient. One with both the muon and the E_T^{miss} triggers active on the same events and one where the E_T^{miss} triggers alone are implemented.

The study was initially conducted in the Resolved analysis. Since it is expected that the impact of these changes are the same for all MC simulated events, in order to get plots with a faster turn-around, only the samples in the period MC16a were tested. If this study shows improvements in the signal yield at the cost of not much increase in background, and does not negatively impact on other areas of the analysis, then the study is then repeated using the entire MC simulated events set, called MC16ade.

10.4 Preliminary trials in the Resolved analysis

Since the E_T^{miss} is determined by what is not detected, its definition and calculation can vary. If the trigger is simulated in MC, trigger requirements have to be created to reproduce offline what is done online. If the relevant derivations for the analyses have the fully reconstructed content, then so long as the required particles/conditions are reconstructed offline, then it does not matter how events were triggered online.

The data-stream physics-main is the one used to create derivations for the main analysis. It has all events and all of the triggers, so its derivations can be used for this study. This means that this study is easily tested and integrated into the current software framework.

In order to fully test the E_T^{miss} triggers in the 2-lepton analysis, an initial E_T^{miss} proxy needs to be created. This proxy is required because there will be small levels of missing transverse energy in a 2-lepton VH , $H \rightarrow b\bar{b}$ event, and it is useful to see how much of an effect this has on the analysis. In addition to this there are other considerations such as the recalibration of jets, that can change the objects that make up E_T^{miss} offline and online.

The initial proxy is required to test the level of real E_T^{miss} in events. It can also be edited such that the offline definition of E_T^{miss} more closely resembles the online definition of E_T^{miss} ³.

In the 2-lepton channel, the p_T of the dimuon system that comes from the Z-boson (MET_{dimuon}) will serve as the initial proxy⁴. The actual reconstructed E_T^{miss} available to

³At the L1-trigger level the reconstructed E_T^{miss} in a 2-lepton event should consist of the energy of the muons in addition to the real E_T^{miss} from momentum imbalances. It was not known at this point whether the E_T^{miss} object in a 2-lepton event corrected for this fact.

⁴The subsequent figures will refer to this as ‘MET’ (with the quotation marks) as well

the analysis will be referred to as MET . Since E_T^{miss} is calculated by the vector addition of different contributions, all forms of E_T^{miss} need to be taken into the calculation as this is the information that the E_T^{miss} triggers will have when making event-level decisions.




Figure 10.3: Preliminary plots showing the effect of the E_T^{miss} study in the Resolved 2-lepton analysis. Here an event will be selected if it passes one of the muon or E_T^{miss} trigger sets. A di-muon object with missing E_T labelled as ('MET') is a stand-in for MET_{dimuon} . Note that the ratio of new to old trigger regime doesn't change for p_T^V values below 150 GeV as events at these values of p_T^V are not affected by these changes.

Figures 10.3 and 10.4 show the results of preliminary tests in the Resolved analysis when trying to implement such changes to the selection. Figure 10.3 shows the effect on the signal yield when the E_T^{miss} trigger is used in conjunction with the muon trigger to select events and Figure 10.4 shows the effect on the signal yield when the E_T^{miss} trigger replaces the muon trigger for events with $p_T^Z > 150$ GeV.

These images show two sets of information that is used to create a proxy for the E_T^{miss} triggers to trigger on in the 2-lepton analysis: MET and MET_{dimuon} , and two variables with shapes highly sensitive to these choices: p_T^V and E_T^{miss} . Here it can be seen that the reconstructed E_T^{miss} in the event is smaller in energy compared to the p_T^V of the event, and the increases in yield come from events with low E_T^{miss} and high p_T^V . This is consis-




Figure 10.4: Preliminary plots showing the effect of the E_T^{miss} trigger study in the Resolved 2-lepton analysis. Here only the E_T^{miss} trigger set is being used to select events. A di-muon object with missing E_T labelled as ('MET') is a stand-in for MET_{dimuon} . Note that the ratio of new to old trigger regime doesn't change for p_T^V values below 150 GeV as events at these values of p_T^V are not affected by these changes.

tent with the understanding of the trigger system as calibrated E_T^{miss} is of low energy and includes a muon contribution, but the trigger would only have access to the E_T^{miss} from the calorimeter which would come from muons from the Z-boson, and hence have a high p_T^V .

As expected, Figure 10.3 shows greater positive benefit for the analysis than Figure 10.4. In order to see how equivalent triggering the information from the MET_{dimuon} to triggering on actual E_T^{miss} was, both implementations were plotted on the same axes, as shown in Figure 10.5.

Here the difference between the purple and the green lines at $p_T^V = 150$ GeV confirms the presence of actual E_T^{miss} in the event. The solution, then, to obtain the best offline proxy of the E_T^{miss} object is to change the initial proxy, MET_{dimuon} by vectorally adding the event-level reconstructed E_T^{miss} to it, to combine both the contributions.




Figure 10.5: Plot showing the difference between the event selection yields when using MET_{dimuon} ('MET') as the trigger proxy, and MET. Here the event selection uses both muon and E_T^{miss} information to filter events.




Figure 10.6: Preliminary results for the impact on the decided E_T^{miss} trigger regime on the Resolved analysis. $MET + 'MET'$ refers to the final E_T^{miss} proxy used by this study which is a vector addition of MET_{dimuon} and MET.

10.4.1 Results for E_T^{miss} study in the Resolved analysis

Figure 10.6 shows that when the E_T^{miss} trigger is introduced into the event selection of the Resolved 2-lepton analysis, there is only an increase in the signal yield if the infor-

mation from both triggers are used simultaneously. Replacing the muon trigger event selection requirements with a series of E_T^{miss} ones reduces the number of signal events. To reduce the complexity of the analysis and to avoid the use of joint lepton- E_T^{miss} scale factors, in all regions where it is possible to use the E_T^{miss} triggers, they should replace the single muon triggers and not be used alongside them. The 2-lepton analysis therefore will now have the same trigger strategy as the 1-lepton analysis. Figure 10.7 shows the finalised trigger changes to be tested.




Figure 10.7: A simple graphic showing differences between the new and the older trigger regime. ALU stands for ‘add lowest un-prescaled’ triggers, which means the triggers with the lowest energy threshold not adjusted for any other effects.

However this loss of signal yield is not the final look at this study for the Resolved analysis. This test was conducted across the entire p_T^V range of the analysis. The signal regions in the version of the analysis after the ICHEP conference of 2018 is split into three distinct p_T^V categories: Medium p_T^V : $75 \text{ GeV} \leq p_T^V < 150 \text{ GeV}$, High p_T^V : $150 \text{ GeV} \leq p_T^V < 250 \text{ GeV}$ and Extreme p_T^V : $p_T^V \geq 250 \text{ GeV}$. Since each of these categories can have different trigger requirements, the next test should be to repeat the test in these separate p_T^V regions to see if there can be an increase in any one of these regions.




Figure 10.8: Final results for the impact on the decided E_{T}^{miss} trigger regime on signal and background samples present in the Resolved analysis. In this figure the study has been split into the different regions of the Resolved analysis to isolate the improvement.

Figure 10.8 shows that for the Resolved analysis only changing the triggers in the $p_T^V \geq 250 \text{ GeV}$ region will contribute to an increase in signal yield. The background yield increases as well but since the figure of merit is $\frac{s}{\sqrt{b}}$, these numbers indicate that there is likely to be an increase in the significance of the effect of the Extreme p_T^V region to the Resolved 2-lepton analysis as a whole. This effect, however will be small as both the High and Medium p_T^V regions have at least four times the number of events. Figure 10.8 also shows that the high p_T^V region suffers a loss in signal yield when the trigger changes. While the background yield decreasing by a similar amount could indicate an increase in $\frac{s}{\sqrt{b}}$, the rest of the analysis could suffer losses in significance as a result.

Before making a decision about whether this should be pursued further, one more thing must be considered. The Resolved analysis has a Top control region ($e\mu$ -CR) specifically designed to mirror the signal region and to be pure in $t\bar{t}$ events in the 2-lepton channel, so is used to inform the other analysis regions. Table 10.2 shows the current object selection for the 2-lepton analysis.

Table 10.2: Table showing the event selection criteria for the 2-lepton Resolved analysis at the time a decision was made about the study. Here the Control Region refers to the $e\mu$ -CR region in the 2-lepton analysis designed to inform the other analyses about the $t\bar{t}$ background template.

Object Selection in 2L	
Trigger	Single Lepton
Leptons	2 loose leptons with $p_T > 7 \text{ GeV}$ OR ≥ 1 lepton with $p_T > 27 \text{ GeV}$
E_T^{Miss}	-
m_{ll}	$81 \text{ GeV} < m_{ll} < 101 \text{ GeV}$
Jets	Exactly 2, OR ≥ 3
P_T^V	$75 \text{ GeV} \leq p_T^V < 150 \text{ GeV}$ (Medium), OR $150 \text{ GeV} \leq p_T^V < 250 \text{ GeV}$ (High), OR $p_T^V \geq 250 \text{ GeV}$ (Extreme)
Signal Region	Same flavour leptons AND Opposite sign charges
Control Region	Different flavour leptons AND Opposite sign charges

Table 10.2 shows that the $e\mu$ -CR has all the same triggers with only one difference, a re-


quirement for same sign lepton final states rather than different signed ones. Hence in order to keep this, any trigger changes that affect the SR should also be done in the $e\mu$ -CR. This control region, however, has actual E_T^{miss} and switching to a different trigger is not something that can be taken for granted. It risks disrupting the balance between the SR and the $e\mu$ -CR in the 2-lepton channel which has a wider impact on the entire Resolved analysis as this template is exported to other regions.

Due to the relatively small increase in the extreme p_T^V region and the clash with the selection in the $e\mu$ -CR in the $p_T^V \geq 250$ GeV region, it was decided that it should not be used in the Resolved analysis. However since the Boosted analysis is solely contained within this region, the study was continued in the Boosted analysis. Due to this additional boosting of the Z-boson system, and by extension the Higgs, in the Boosted analysis compared to the Resolved analysis, a E_T^{miss} study should have a greater impact there. Also the Boosted 2-lepton analysis does not have this control region.


10.5 Preliminary trials in the Boosted analysis

Figure 10.9a and Figure 10.9b show plots on the same axes as the plots shown in Figure 10.6 but run over events in the Boosted analysis. Figure 10.9a shows that when only the E_T^{miss} trigger acting at $p_T^V \leq 150$ GeV, there is an increase in event yield of 3.7% and Figure 10.9b shows that when the logical OR of both triggers is used to add events to the event selection, their increase is slightly larger at 4.1%. These numbers come from differences in the entire range of the boosted plots including the overflow bins not present in these plots.

Similar to the Resolved analysis, the trigger method has been decided and it would be better if this information was split into the various analysis regions. This can be seen in Figure 10.10.



(a)



(b)

Figure 10.9: Maximal percentage increase in the event signal yield from changing trigger strategy in the Boosted analysis. In (a) only the E_T^{miss} trigger makes any decisions if there are muons in the event with $p_T^V \leq 150$ GeV, and in (b) both the E_T^{miss} trigger and the muon trigger both make decisions if there are muons in the event with $p_T^V \leq 150$ GeV.




Figure 10.10: The final results for the impact on the decided E_T^{miss} trigger regime on signal and background samples present in the Boosted analysis on MC16a signal events. In this figure the study has been split into the different regions of the Boosted analysis to isolate the improvement. The region $0 \text{ GeV} \leq p_T^V < 250 \text{ GeV}$ is not in present within the Boosted analysis serves as a cross-check with the Resolved analysis.

10.5.1 E_T^{miss} Trigger Scale Factors


The use of these E_T^{miss} triggers comes with associated p_T^V -dependent trigger scale factors which mimic the efficiency of these triggers. The scale factors are already present in the 1-lepton analysis [160] and would be applied in the 2-lepton analysis but their effect is unknown. This was not currently taken into account in the study and not currently present in the Boosted 2-lepton analysis. Thus these factors are essentially assumed to be unity across the whole p_T^V range of the boosted analysis. Since the improvements seen are of the order of 4%, if the scale factors deviate from this assumption by an amount equal to or more than this, then the presented gains in switching triggers may be cancelled out with the implementation of these scale factors. The efficiency curves for these E_T^{miss} triggers are present in Figure 10.11.

Figure 10.11 shows that the scale factors for the five E_T^{miss} triggers to be used in the 2-lepton analysis plateau to a value of 1 around $p_T^V = 200 \text{ GeV}$ which is below the p_T^V threshold of all regions in the Boosted analysis. Hence the assumption that these values are unity for the entire Boosted analysis holds, and the study need not change the methodology.


Since the 2-lepton analysis uses the same triggers as the 1-lepton, these channels no longer need to be triggered separately. To ensure that these events would always be triggered in the same way, I decided to merge the trigger codes for the 1- and the 2-lepton channels and make a common trigger tool between the two.

10.5.2 Missing E_T^{miss} Trigger


A routine check in the derivations found that for the MC simulated events that correspond to the data taken in 2015 (which make up some fraction of the MC16a samples which this study was initially based on), the information about the E_T^{miss} trigger for that




(a) HLT-xe70




(b) HLT-xe90-mht-L1XE50



(c) HLT-xe110-mht-L1XE50



(d) HLT-xe110-pufit-L1XE55



(e) HLT-xe110-pufit-xe70-L1XE50

Figure 10.11: Trigger scale factor values versus p_T^V for the five E_T^{miss} triggers used in the Boosted analysis. Above 250 GeV all the scale factors for these triggers are less than a quarter of a percent from unity. Images made using [161].

data period, xe70, was missing. However the E_T^{miss} trigger that corresponds to the data taken in 2016 in MC16a, xe90_mht.L1XE50, is present.

This means that the MC simulated events that previously entered the analysis via the use of single lepton triggers in the event selection are not recovered when the event selection is changed to use the E_T^{miss} trigger information.

To check the impact that this would have on the analysis, the events in MC16 formerly triggered on xe70 are run with the next lowest unprescaled trigger that is present in the analysis: xe80_mht.L1XE50. This will be referred to simply as xe80 in plots.

Figure 10.12 shows that when the missing xe70 trigger is replaced with one present in the analysis, the resulting changes are small. Repeating this test with the xe90_mht.L1XE50 instead of the xe80_mht.L1XE50 one yields a similar sub-percentage change in the yields. Therefore there must be some xe70 data present in the CxAOD's. Tracing back the analysis steps and it was found that this was the case.

At the 2-lepton channel derivation stage, after the particular MC simulations period has been identified and the correct xAOD's have been selected, A branch for every trigger that contains the strings *e*, *mu*, *2e*, *2mu*, *xe*, *j* and *b* (single electron, single muon, di-electron, di-muon, E_T^{miss} , jet, b-jet) is created in the DxAOD. This is the first step in the derivation process and where the xe70 trigger container originates. Later on in the process, trigger skimming begins. In the initial skimming, only triggers with the strings *e*, *mu*, *2e*, *2mu* and *xe* are kept, and then later only a set of specific triggers are recorded. It is in this list that xe70 is not present. This is only the case for the 2-lepton analysis. For the 1-lepton and 0-lepton derivations, xe70 is present in the trigger list. This error only occurs for MC16a MC simulated events for the 2015 data period as the later periods use different triggers which are in this list. Once this trigger




Figure 10.12: Differences between the impacts of the two E_T^{miss} trigger regimes on signal and background samples in the top row of plots and data samples in the bottom set of plots present in the Boosted analysis. The black line represents the E_T^{miss} trigger strategy outlined in Section 10.4. The green line replaces the x_{e70} with the $x_{e80_mht_L1XE50}$ for MC16a. In this figure the study has been split into the different regions of the Boosted analysis to isolate the improvement. The region $0 \text{ GeV} \leq p_{TV} < 250 \text{ GeV}$ serves as a cross-check with the Resolved analysis.

list has been formed, each event is run through the list of triggers and gains a binary integer corresponding to the ones in the list that pass.

Later at the CxAODMaker stage, when the DxAOD's are run to produce CxAOD's, the events are run through the trigger requirements again. However each event has gained information about whether it passes or fails all the other triggers in the analysis. For example, if a MC16a event enters into the analysis because it passed the `xe80_mht_L1XE50` trigger, that event also contains the information about whether it passed the `xe70` trigger. This is extrapolated from the binary results of the trigger decisions and the analysis scale factors.

The results in Figure 10.12 show that all the events that passed the `xe80_mht_L1XE50` or `xe90_mht_L1XE50` triggers also successively passed the `xe70` trigger⁵. This could be the case because the trigger turn-on curves for the `xe70` trigger and the `xe80` trigger are similar. This means that the Reader-level requirement for the presence of events that pass the `xe70` trigger will give an answer, but it will be an underestimate of the true answer as MC16a events that made it into the analysis which exclusively passed the `xe70` trigger are non-recoverable.

The only way to rectify this error is to produce new 2-lepton derivations with this requirement added. There are possible workarounds that include using the single lepton trigger for only the part of the MC16a affected by this, but this means the temporary addition of an ugly snippet of code which will only affect a small amount of the 139 fb^{-1} making up Run-II ($\frac{\text{data}_{15}}{\text{Run-II}} = \frac{3.2 \text{ fb}^{-1}}{139 \text{ fb}^{-1}} \leq 4\%$), which then has to be removed when new derivations are produced. Since the study showed positive results for the Boosted analysis in MC16a; the code doesn't break when the `xe70` trigger information is asked for; and the running of the derivations can only increase the effect switching to E_T^{miss} trig-

⁵This is not a given a-priori as the `xe70` trigger uses a different method to calculate E_T^{miss} .

gers has on the analysis, it was decided to continue with the code-base as is, because new 2-lepton derivations were planned to be generated in the near future.


10.6 Final Results in the Boosted Analysis

When using the latest full nominal MC simulations sample set, MC16ade, and replacing the single lepton triggers with E_T^{miss} triggers regime shown in Table 10.1, the event yield in the 2-lepton channel increases. These are shown in Figure 10.13 for signal, background and data samples in the two analysis regions: $250 \text{ GeV} \leq p_T^V < 400 \text{ GeV}$ and $p_T^V \geq 400 \text{ GeV}$. Figure 10.13 shows that there is a 4-5% increase in event yield for all regions except for the one for data in the $p_T^V \geq 400 \text{ GeV}$ region which is 1.5%. Since the size of the sample is small (see Figure 10.14) this percentage increase could be a statistical effect as it looks to only gain one event in one bin.


10.6.1 Final Results from Running the Boosted WorkSpace Maker

Nominally when a study like this is conducted, in order to get the true impact on the analysis, in addition to running on the nominal MC simulations samples, one would need to run the fit with the systematics variations as well, but given the way this E_T^{miss} study is being implemented at the event selection stage, it is unlikely that any current systematics would be affected or any new systematics required. The fit with and without the changes to the event selection can be run without any systematic information in it at all with little cost to the accuracy or significance of the result. Hence to save computational time, and reduce the likelihood of runtime errors, only fluctuations from MC simulations statistics are included in the fit when determining the increase in significance.


When inputs that contain the information from the E_T^{miss} triggers are run in the fit and compared to the old standard, the total signal and background yields increases. This,




(a) Signal yield increase: 4.28%, $250 \text{ GeV} \leq p_T^V < 400 \text{ GeV}$




(b) Signal yield increase: 4.73%, $250 \text{ GeV} \leq p_T^V < 400 \text{ GeV}$




(c) Background yield increase: 4.23%, $250 \text{ GeV} \leq p_T^V < 400 \text{ GeV}$



(d) Background yield increase: 5.15%, $p_T^V \geq 400 \text{ GeV}$



(e) Data yield increase: 4.40%, $250 \text{ GeV} \leq p_T^V < 400 \text{ GeV}$



(f) Data yield increase: 1.53%, $p_T^V \geq 400 \text{ GeV}$

Figure 10.13: p_T^Z distributions in the $250 \text{ GeV} \leq p_T^V < 400 \text{ GeV}$ (left) and $p_T^V \geq 400 \text{ GeV}$ (right) regions in the 2-lepton channel for signal (top), all backgrounds (middle), and the data (bottom). The black line represents the signal yield under the use of the muon single lepton triggers. The red line shows the signal yield under the use of E_T^{miss} triggers.

alongside the improvement in significance can be seen in Table 10.3.

Table 10.3: Yields for signal and background inclusive in p_T^V in the signal regions and fit significances.

	Total (single- e & single- μ triggers)	Total (single- e & E_T^{miss} triggers)	Difference (%)
VH125 (Signal) yields	3.5652	3.7346	+4.8
Background yields	50.8405	53.7071	+5.6
MCStat Median significance	1.688	1.748	+3.4



The median statistical significance figure of merit here is an asymptotic approximation common to many analyses and has been defined in [162, p26] as

$$Z = \sqrt{2(s+b) \ln \left(1 + \frac{s}{b}\right)} - s, \quad (10.1)$$



where the total MC signal (s) and MC background (b) yields are combined to a single bin.

The impact of this trigger change maintains the ratio between the data and the MC simulated events prediction (also known as the data-MC agreement or the data-MC ratio) similar to that before the change in trigger strategy, as can be seen in Figure 10.14.

The switch from using single lepton triggers to E_T^{miss} triggers in the 2-lepton analysis results in a 3.4% improvement on the 2-lepton significance and a 4.8% acceptance on signal events. Since this increase is enough to improve the significance of the analysis as a whole (See Appendix D), the changes were adopted for the default ATLAS analysis.



(a) m_J , $250 \text{ GeV} \leq p_T^V < 400 \text{ GeV}$ single- μ trigger (b) m_J , $250 \text{ GeV} \leq p_T^V < 400 \text{ GeV}$, E_T^{miss} trigger



(c) m_J , $p_T^V \geq 400 \text{ GeV}$, single- μ trigger

(d) m_J , $p_T^V \geq 400 \text{ GeV}$, E_T^{miss} trigger

Figure 10.14: Plot showing the pre-fit m_J distributions and the data-MC ratio for the 2-lepton signal region, before (left) and after (right) the new trigger regime is implemented and for the $250 \text{ GeV} \leq p_T^V < 400 \text{ GeV}$ region (top) and the $p_T^V \geq 400 \text{ GeV}$ region (bottom). The plots on the left side were generated with full systematics, and the plots on the right were generated with only normalisation and acceptance systematics produced at the fit level. So while the systematic bands cannot be compared, the data-MC ratios are consistent.

Results and Conclusion

Perspicuously, I hope, I have informed you of my work on Higgs boson studies at the ATLAS experiment. I shall briefly summarise this journey, show and explain the results from the Resolved and Boosted analyses, and then add some concluding thoughts.

I started off in Chapter 2 with a description of the Standard Model and showed how the Higgs boson fits in with this picture to give the theoretical motivations behind further study. In Chapter 3, I set the historical context for the CERN’s accelerator complex, the LHC, and the ATLAS detector that was used in this thesis to take data, and then followed this up in Chapter 4 outlining the specifications of the ATLAS detector. Chapter 5 then described the process of generating a Standard Model response using MC simulation methods such that comparisons to the data yield measurements.

After this in Chapter 6, I took you through my work improving understanding of the behaviour of charged particles in the detector by comparing ATLAS data to MC simulations. In the following chapter I defined the experimental signatures for the Standard Model particles and higher-level physics objects such as jets. Having defined the

components, in Chapter 8 I then explained the $VH, H \rightarrow b\bar{b}$ signature that was being targeted and introduced the Resolved and Boosted analysis I worked on. These analyses are designed to combine these objects and their kinematic features into various figures-of-merit and use a statistical fit to extract the required Higgs measurements.

The statistical aspects of this combination of objects was explained in Chapter 9 where I also showed the work I did in the Resolved analysis to optimise said fit and then in Chapter 10, I showed the work I did in the Boosted analysis to design, implement and test an improvement to increase the significance of the analysis.

And with that, it is time to show the results from these analyses for a 125 GeV Higgs boson.

11.1 The Boosted analysis results

Since the last chapter was about the Boosted analysis, we shall start with its results. The Boosted analysis extracts μ_{VH}^{bb} and μ_{VZ}^{bb} simultaneously by fits of data-to-MC simulations in the m_J variable. The post-fit distributions of m_J in each of the lepton channels are shown in Figure 11.1.

When all the signal regions are taken into consideration, the observed (expected) excess of signal events over the expected SM background was 2.1σ (2.7σ). At this confidence level, the fitted value for the signal strength μ_{VH}^{bb} is $0.72^{+0.29}_{-0.28}(\text{stat.})^{+0.26}_{-0.22}(\text{syst.})$ and for μ_{VZ}^{bb} is $0.91 \pm 0.15(\text{stat.})^{+0.25}_{-0.17}(\text{syst.})$.

To show that the separation between ZH and WH processes is good enough for this extraction, the m_J distributions for data and MC simulations in all channels are combined after all the backgrounds barring the diboson ones are removed. This is shown




Figure 11.1: Mass distributions of the large- R jet in the (a,b) 0-lepton, (c,d) 1-lepton, and the (e,f) 2-lepton 2- b -tagged jet signal regions. The left hand side plots are from the $250 \text{ GeV} \leq p_T^V < 400 \text{ GeV}$ region and those on the right are from the $p_T^V \geq 400 \text{ GeV}$ region. In this plot the Low and High purity sub-regions in the 0- and 1-lepton regions have been merged. Image from [4, p14].

in Figure 11.2 where, within the statistical variations, the data matches the predictions of the (weighted) number of events for signal and diboson samples, and that there is a good separation between them.





Figure 11.2: m_J distributions in all of the lepton channels and p_T^V regions combined for data and MC simulations simulations after the removal of all non-diboson backgrounds in the Boosted analysis. The combined regions are weighted by their ratios of the Higgs-signal-to-background ratio. The hatched band shows the total uncertainty minus uncertainties on VH or VZ . Image from [4, p17].

After this, the simplified template cross sections for ZH and WH production were measured separately in the Boosted analysis STXS framework for each p_T^V region. In the STXS measurement region of $[|y_H| < 2.5, H \rightarrow b\bar{b}]$, the values of the STXS ZH and WH are shown in Figure 11.3a. This information is then shown pictorially in Figure 11.3b.

Measurement region ($ y_H < 2.5, H \rightarrow b\bar{b}$)	SM prediction [fb]	Result [fb]	Stat. unc. [fb]	Syst. unc. [fb]
$W \rightarrow \ell\nu; p_T^W \in [250, 400[\text{ GeV}$	5.83 ± 0.26	$3.3^{+4.8}_{-4.6}$	$^{+3.6}_{-3.4}$	$^{+3.2}_{-3.0}$
$W \rightarrow \ell\nu; p_T^W \in [400, \infty[\text{ GeV}$	1.25 ± 0.06	$2.1^{+1.2}_{-1.1}$	$^{+1.0}_{-0.9}$	$^{+0.6}_{-0.5}$
$Z \rightarrow \ell\ell, \nu\nu; p_T^Z \in [250, 400[\text{ GeV}$	4.12 ± 0.45	$1.4^{+3.1}_{-2.9}$	$^{+2.4}_{-2.3}$	$^{+1.9}_{-1.7}$
$Z \rightarrow \ell\ell, \nu\nu; p_T^Z \in [400, \infty[\text{ GeV}$	0.72 ± 0.05	$0.2^{+0.7}_{-0.6}$	$^{+0.6}_{-0.5}$	$^{+0.3}_{-0.3}$

(a)



(b)

Figure 11.3: (a) Measured and predicted values of the STXS's for WH and ZH in the Boosted analysis's STXS framework in each p_T^V region. This is transferred to image form (b) where the measured points are black and green points/bands and the predicted is in red. The bottom ratio plot shows the ratio with respect to the standard model (black/red vs red/red). Images from [3, p18].

11.2 The Resolved analysis results

The signal strength for the Resolved analysis is obtained from a global likelihood fit over multivariate discriminants (BDT_{VH}) from 14 signal regions and the event yields from the 28 ΔR_{bb} control regions. A selection of these for the 2-jet subcategory in the 0-, 1- and 2-lepton channels is shown in Figure 11.4 as these have the most sensitivity to the signal. From this, the value of μ_{VH}^{bb} extracted is $1.02^{+0.12}_{-0.11}(\text{stat.})^{+0.14}_{-0.13}(\text{syst.})$. This corresponds to a rejection of the background-only hypothesis to the level of 6.7σ and agrees with the expectation from the Standard Model within the uncertainties of the measurement.

This value is to be compared to the di-jet mass analysis cross check which performs a non-BDT global fit to the m_{bb} variable. When all the channels are combined the value of μ_{VH}^{bb} as measured by the cross-check is $1.17 \pm 0.16(\text{stat.})^{+0.19}_{-0.16}(\text{syst.})$. This observed (expected) excess rejects the background-only hypothesis to 5.5σ (4.9σ).

The diboson cross-check is then bootstrapped to the main analysis and the results are statistically comparable to 1.1 standard deviations. Similar to the Boosted analysis, when all the backgrounds except from the diboson ones are removed for this variable, the resulting dijet plot (Figure 11.5) shows good separation between diboson and signal, with good match to data within statistical fluctuations.

After the value of μ_{VH}^{bb} is extracted for the 1-POI fit, the fit is run with signal strengths for WH and ZH floating independently to form the 2-POI fit (the extraction of μ_{WH}^{bb} and μ_{ZH}^{bb}). μ_{WH}^{bb} was found to be $0.95 \pm 0.18(\text{stat.})^{+0.19}_{-0.18}(\text{syst.})$ which is observed (expected) to reject the background-only hypothesis at the 4.0σ (4.1σ) level. μ_{ZH}^{bb} was found to be $1.08 \pm 0.17(\text{stat.})^{+0.18}_{-0.15}(\text{syst.})$ which corresponds to an observed (expected) significance of 5.3σ (5.1σ).




Figure 11.4: Output of the multivariate BDT_{VH} distributions in the (a,b) 0-lepton, (c,d) 1-lepton, and the (e,f) 2-lepton 2-jet, 2-b-tagged jet signal region of the Resolved analysis. The left hand side plots are from the $150 \text{ GeV} \leq p_T^V < 250 \text{ GeV}$ region and those on the right are from the $p_T^V \geq 250 \text{ GeV}$ region. The predicted background before the fit is marked as a black dashed line. Image from [3, p26].

These figures are to be compared and contrasted to their diboson cross-check counterparts: μ_{WZ}^{bb} and μ_{ZZ}^{bb} , which can be seen, along with the main analysis 2-POI results in Figure 11.6.




Figure 11.5: Final m_{bb} distributions in all of the lepton channels and p_T^V regions combined for data and MC simulations after the removal of all non-diboson backgrounds in the Resolved analysis. The combined regions are weighted by their ratios of the Higgs-signal-to-background ratio. The hatched band shows the total uncertainty minus uncertainties on VH or VZ . Image from [3, p29].

Finally when the values of the μ_{ZH}^{bb} and μ_{WH}^{bb} are further segmented to their finest granularity, the 5-POI fit is formed. In the STXS measurement region of $[|y_H| < 2.5, H \rightarrow b\bar{b}]$, the values of the signal strength equivalents in the 5 STXS regions are shown in Figure 11.7a. This information is then shown pictorially in Figure 11.7b.




Figure 11.6: (a) Values of μ_{ZH}^{bb} and μ_{WH}^{bb} for the Resolved analysis 2-POI fit, and their combined 1-POI value. (b) Values of μ_{ZZ}^{bb} and μ_{WZ}^{bb} for the Resolved analysis diboson cross-check fit, and their combined 1-POI value. Images from [3, p27, p30].

11.3 Future plans for the analysis


The Resolved analysis managed to successfully increase the significance of its result, getting an observation in the $VH, H \rightarrow b\bar{b}$ channel. Following on from this success, some SMEFT operators were looked into to see if they could affect the result. This is surface-scratching work at the moment, as BSM sensitivities are low and there is no particular push from the theoretical side as to which SMEFT operators are relevant.

All the Boosted analysis results seem to be consistent with the Standard Model and small visible excesses were seen, but the analysis has yet to see evidence-level proof of $VH, H \rightarrow b\bar{b}$. The analysis, however, has passed its proof-of-concept and steps are being taken to make the analysis more robust which includes introducing a BDT multivariate analysis in line with the Resolved analysis.

The main aim in the coming years is the combination of the Resolved analysis, Boosted analysis and a sister analysis where the Higgs signal is identified by its decay into c -jets: $VH, H \rightarrow c\bar{c}$. This will form what will be known as the Run-2 Legacy Analy-

STXS region Process	$p_T^{V,t}$ interval	SM prediction		Result		Stat. unc.		Syst. unc. [fb]		
		[fb]		[fb]		[fb]		Th. sig.	Th. bkg.	Exp.
$W(\ell\nu)H$	150–250 GeV	24.0	± 1.1	19.0	± 12.1	± 7.7	± 0.9	± 5.5	± 6.0	
$W(\ell\nu)H$	> 250 GeV	7.1	± 0.3	7.2	± 2.2	± 1.9	± 0.4	± 0.8	± 0.7	
$Z(\ell\ell/\nu\nu)H$	75–150 GeV	50.6	± 4.1	42.5	± 35.9	± 25.3	± 5.6	± 17.2	± 19.7	
$Z(\ell\ell/\nu\nu)H$	150–250 GeV	18.8	± 2.4	20.5	± 6.2	± 5.0	± 2.3	± 2.4	± 2.3	
$Z(\ell\ell/\nu\nu)H$	>250 GeV	4.9	± 0.5	5.4	± 1.7	± 1.5	± 0.5	± 0.5	± 0.3	

(a)



(b)

Figure 11.7: (a) Measured and predicted values of the STXS's for VH and ZH in the Resolved analysis's STXS framework in each p_T^V region. This is transferred to image form (b) where the measured points are black and green points/bands and the predicted is in red. The bottom ratio plot shows the ratio with respect to the standard model (black/red vs red/red). Images from [3, p18].

sis. This comes with its own challenges, the most obvious being that there is a large overlap between the Resolved analysis's $p_T^V \geq 250$ GeV and the Boosted analysis's $250 \text{ GeV} \leq p_T^V < 400$ GeV region (the Resolved analysis does not seem to have many

events that enter the Boosted analysis's $p_T^V \geq 400 \text{ GeV}$ region) and it is not obvious how to ensure that events are not double counted without completely throwing away one set of events.

The aim of merging all these analysis into one is to try and come up with as complete coverage of the Higgs boson as possible such that the cross-sections can be reliably produced as a function of energy or variable. This is a *differential* measurement and is what the STXS framework essentially tries to replicate but in a simplified manner.

In addition to the combination, the analyses themselves could be made more sensitive. Ideas have been floated of things to be done to improve the analysis such as: the addition of an additional $75 \text{ GeV} \leq p_T^V < 150 \text{ GeV}$ SR in the 1-lepton channel as there are large data statistics to be exploited in this region; the possible reduction of variables that enter into the MVA as to not bias differential measurements; and actively using tau leptons to try and disentangle $Z + HF$ and $W + HF$ to try and create a pure CR in the 0-lepton channel. It is also thought that the merged analysis will enjoy more MC simulated events to reduce the effect of statistical error which is the leading systematic in the Boosted analysis.

11.4 Conclusion

The VH , $H \rightarrow b\bar{b}$ analyses had to put even more processing power into larger data volumes than previous iterations of the analyses to reach the required levels of sensitivity in their results. Part of the reason this massive undertaking was successful is thanks to: improved MC simulation techniques, more optimised data processing frameworks, and clever conversions of raw ATLAS data¹. Thanks to this, looking for a $H \rightarrow b\bar{b}$ signal at a hadron collider now and in the future is challenging but manageable.

¹and of course a dedicated and passionate team of physicists.

This thesis has been a small snapshot of the state of cutting-edge Higgs physics at the ATLAS experiment. While the decay to b -quarks has a large branching fraction, it is becoming increasingly clear that the window for BSM interference in this decay mode is narrowing. The hope for future physicists desiring to hit new-physics gold in the Higgs sector is that the HL-LHC takes enough data to make accurate measurements of the $H \rightarrow cc$ and $H \rightarrow \mu\mu$ branching fractions, and that these branching fractions exhibit statistically significant anomalies compared to the Standard Model predictions.

Appendices



From Newtonian to Lagrangian Mechanics

Newtonian Mechanics provides one main equation of note to describe a system of particles; Newton's Second Law (N2L), given by Equation A.1. Equation A.1 shows that an applied force results in the rate of change with respect to time of an object's velocity¹ (its acceleration) that depends on its mass according to:

$$F = ma = m \frac{dv}{dt} = m \frac{d^2x}{dt^2} \quad (F = ma = m\dot{v} = m\ddot{x}) \quad (\text{A.1})$$

where each dot above the letter indicates a derivative with respect to time of that quantity. A conservative force, like all the forces of the Standard Model, is mathematically defined as deriving from a potential $\mathcal{V}$. A potential here is a 'map' of how strong the interaction is at different points in space. An object inside this potential sensitive to that force will experience a force if it moves between points where the potential is different.

¹which in turn is the rate of change with respect to time of an objects position

This is represented by Equation A.2.

$$p = m\ddot{x} = -m \frac{\partial \mathcal{V}}{\partial x} \quad (\text{A.2})$$

If we multiply both sides of Equation A.2 by $\dot{x}$,

$$\begin{aligned} m\ddot{x}\dot{x} &= -\frac{\partial \mathcal{V}}{\partial x} \times \dot{x} \\ m \frac{d}{dt}(\dot{x})\dot{x} &= -\frac{\partial \mathcal{V}}{\partial x} \times \frac{dx}{dt} \end{aligned}$$

To obtain everything in terms of a time derivative we can ‘notice’ that the left hand side looks like a chain rule result of something of a higher power in $\dot{x}$ and can hence write

$$\begin{aligned} \frac{d}{dt} \left(\frac{1}{2} m \dot{x}^2 \right) &= -\frac{d\mathcal{V}}{dt} \\ \implies \frac{d}{dt} \left(\frac{1}{2} m \dot{x}^2 + \mathcal{V} \right) &= 0 \end{aligned} \quad (\text{A.3})$$

Equation A.3 shows that the rate of change of potential and kinetic energy in a closed system with respect to time is zero, which is another way of saying that energy is conserved. Usually to solve Newtons equations, one needs to be able to deal with special cases that define the field, like specific constraints on particles at the edges of the defined potential. These are called boundary conditions (A), and take the form $A_\nu(\vec{r}_i(q_\alpha, t)) = 0$ where $\vec{r}$ is a generic position vector, q_α are co-ordinates that fulfil said boundary conditions, and t is time. The indices ν and i represent the number of boundary conditions and degrees of freedom respectively.

If N2L is reformulated into q_α that fulfil all of the boundary conditions to be imposed, we can reformulate the generic position vector $\vec{r}$ as $\vec{r}(q_\alpha, t)$ and define the rate of change

with respect to time ($\dot{\vec{r}}$) as

$$\begin{aligned}\frac{d\vec{r}}{dt} &= \frac{\partial \vec{r}}{\partial t} + \sum_{\alpha} \frac{\partial \vec{r}}{\partial q_{\alpha}} \frac{dq_{\alpha}}{dt} \\ \implies \frac{\partial \vec{r}}{\partial \dot{q}_{\alpha}} &= \frac{\partial \vec{r}}{\partial q_{\alpha}}\end{aligned}\tag{A.4}$$

So now we want something similar to energy conservation for this new parametrisation. If we take Equation A.2, replace x with $\vec{r}$ and re-arrange we get

$$m\ddot{\vec{r}} + \frac{\partial \mathcal{V}}{\partial \vec{r}} = 0.\tag{A.5}$$

Multiplying Equation A.5 by $\frac{\partial \vec{r}}{\partial \dot{q}_{\alpha}}$ obtains

$$\textcolor{red}{m\ddot{\vec{r}} \frac{\partial \vec{r}}{\partial \dot{q}_{\alpha}}} + \frac{\partial \mathcal{V}}{\partial \vec{r}} \frac{\partial \vec{r}}{\partial \dot{q}_{\alpha}} = 0.\tag{A.6}$$

We can aim to replace the red part of Equation A.6 with the time and co-ordinate derivatives acting on the whole part as opposed to individual position vectors.

$$\begin{aligned}\frac{d}{dt} \left(\frac{\partial}{\partial \dot{q}_{\alpha}} \left(\frac{1}{2} m \dot{\vec{r}}^2 \right) \right) &= \frac{d}{dt} \left(m \dot{\vec{r}} \frac{\partial \vec{r}}{\partial \dot{q}_{\alpha}} \right) \\ &= \frac{d}{dt} \left(m \dot{\vec{r}} \frac{\partial \vec{r}}{\partial q_{\alpha}} \right) \\ &= m \ddot{\vec{r}} \frac{\partial \vec{r}}{\partial q_{\alpha}} + m \dot{\vec{r}} \frac{d}{dt} \left(\frac{\partial \vec{r}}{\partial q_{\alpha}} \right) \\ &= m \ddot{\vec{r}} \frac{\partial \vec{r}}{\partial q_{\alpha}} + m \dot{\vec{r}} \frac{d}{\partial q_{\alpha}} \left(\frac{\partial \vec{r}}{dt} \right) \\ &= m \ddot{\vec{r}} \frac{\partial \vec{r}}{\partial q_{\alpha}} + m \dot{\vec{r}} \frac{d}{\partial q_{\alpha}} \dot{\vec{r}} \\ &= m \ddot{\vec{r}} \frac{\partial \vec{r}}{\partial q_{\alpha}} + \frac{d}{\partial q_{\alpha}} \left(\frac{1}{2} m \dot{\vec{r}}^2 \right)\end{aligned}\tag{A.7}$$

Since the red part of Equation A.6 is the same as the first term on the right hand side of Equation A.7 we can substitute it for

$$m \ddot{\vec{r}} \frac{\partial \vec{r}}{\partial \dot{q}_{\alpha}} = \frac{d}{dt} \left(\frac{\partial}{\partial \dot{q}_{\alpha}} \left(\frac{1}{2} m \dot{\vec{r}}^2 \right) \right) - \frac{\partial}{\partial q_{\alpha}} \left(\frac{1}{2} m \dot{\vec{r}}^2 \right).\tag{A.8}$$

The blue part of Equation A.6 is easier to handle. Using the relation in Equation A.4, we can show that

$$\frac{\partial \mathcal{V}}{\partial \vec{r}} \frac{\partial \dot{\vec{r}}}{\partial \dot{q}_\alpha} = \frac{\partial \mathcal{V}}{\partial q_\alpha}. \quad (\text{A.9})$$

Using the relations in equations A.8 and A.9, Equation A.6 can be written as

$$\frac{d}{dt} \left(\frac{\partial}{\partial \dot{q}_\alpha} \left(\frac{1}{2} m \dot{\vec{r}}^2 \right) \right) - \frac{\partial}{\partial q_\alpha} \left(\frac{1}{2} m \dot{\vec{r}}^2 \right) + \frac{\partial \mathcal{V}}{\partial q_\alpha} = 0. \quad (\text{A.10})$$

Since the potential $\mathcal{V}$ is a function of q_α and not $\dot{q}_\alpha$, this implies that $\frac{\partial}{\partial \dot{q}_\alpha} \mathcal{V}(q) = 0$ and Equation A.10 can be written as Equation A.11,

$$\frac{d}{dt} \left(\frac{\partial}{\partial \dot{q}_\alpha} \left(\frac{1}{2} m \dot{\vec{r}}^2 - \mathcal{V} \right) \right) - \frac{d}{dt} \left(\frac{1}{2} m \dot{\vec{r}}^2 - \mathcal{V} \right) = 0. \quad (\text{A.11})$$

From Equation A.11 we can see that the derivatives operate on the same object. This object has a kinetic energy term ($\mathcal{T} = \frac{1}{2} m \dot{\vec{r}}^2$) and a potential energy term ($\mathcal{V}$). We define this object $L = \mathcal{T} - \mathcal{V}$ as the *Lagrangian*. Hence in what is known as *Lagrangian mechanics*, the Lagrangian contains the information about the system dynamics by being a function of carefully chosen generalized coordinates, their respective time derivatives, and time.

$$\frac{d}{dt} \frac{\partial L}{\partial \dot{q}_\alpha} - \frac{\partial L}{\partial q_\alpha} = 0. \quad (\text{A.12})$$

Equation A.12 is known as the Euler-Lagrange equation. We can define a new generalised momentum, p to go along with our generalised co-ordinate q_α .

$$p = \frac{\partial L}{\partial \dot{q}_\alpha} \quad (\text{A.13})$$

and alongside Equation A.13 we can a new function $\mathcal{H}$, shown in Equation A.14

$$\mathcal{H} = p \dot{q}_\alpha - L, \quad (\text{A.14})$$

which we insert a basic Lagrangian for a particle with kinetic energy $\mathcal{T} = \frac{1}{2}m\dot{q}_\alpha^2$ and a potential $\mathcal{V}$. For such a set-up the generalised momentum, p equals $m\dot{q}_\alpha$ and hence Equation A.14 simplifies to A.15 .

$$\begin{aligned}
\mathcal{H} &= p\dot{q}_\alpha - L \\
&= p\dot{q}_\alpha - \left(\frac{1}{2}m\dot{q}_\alpha^2 - \mathcal{V} \right) \\
&= p\dot{q}_\alpha - \left(\frac{1}{2}p\dot{q}_\alpha - \mathcal{V} \right) \\
&= p\dot{q}_\alpha - \frac{1}{2}p\dot{q}_\alpha + \mathcal{V} \\
&= \frac{1}{2}p\dot{q}_\alpha + \mathcal{V} \\
&= \frac{1}{2}m\dot{q}_\alpha^2 + \mathcal{V} = \mathcal{T} + \mathcal{V}
\end{aligned} \tag{A.15}$$

$\mathcal{H}$, therefore is an object that provides a description of the total energy of the system. It is called the *Hamiltonian*. To show that this is self consistent with classical mechanics, this Hamiltonian must have a time-derivative of 0. For this we will continue to use the previous Lagrangian definition.

$$\begin{aligned}
\frac{d\mathcal{H}}{dt} &= \frac{d}{dt} (p\dot{q}_\alpha - L) \\
&= \frac{d}{dt} \left(\frac{1}{2}p\dot{q}_\alpha + \mathcal{V} \right) \\
&= \frac{d}{dt} \left(\frac{1}{2}p\dot{q}_\alpha \right) + \cancel{\frac{d\mathcal{V}}{dt}}^0 \\
&= \frac{1}{2}\dot{p}\dot{q}_\alpha - \frac{1}{2}p\ddot{q}_\alpha \\
&= \frac{1}{2}\frac{d}{dt}(m\dot{q}_\alpha)\dot{q}_\alpha - \frac{1}{2}(m\dot{q}_\alpha)\ddot{q}_\alpha \\
&= \frac{1}{2}m\dot{q}_\alpha\ddot{q}_\alpha - \frac{1}{2}m\dot{q}_\alpha\ddot{q}_\alpha = 0
\end{aligned}$$

Hence we have recovered energy conservation. Lagrangian mechanics, is just a re-formulation of classical mechanics, designed to describe particles. This describes the internal mechanics of one particle.

APPENDIX B

Axioms of Quantum Mechanics and their consequences

Quantum mechanics aims to describe single particles as *quantum states*. In this endeavour, a new set of mathematical objects need to be defined. the Axioms of Quantum Mechanics, therefore, are five-fold [163].

- 1 The most complete knowledge one can have of a quantum system is to know its state vector $|\psi\rangle$ as defined within a ‘space of states’. This space is described by a complex Hilbert Space.
- 2 Every physical property $\mathcal{A}$ or *observable* like energy, position, momentum has a Hermitian Operator¹ ($\hat{A}$) and the result of that measurement on said Hilbert space is an eigenvalue of that operator.
- 3a The probability of finding the state $|\psi\rangle$ within another physical state $|\phi\rangle$ is given by $|\langle\phi|\psi\rangle|^2$

¹A Hermitian object is any matrix-style object in which the operator is the same as its transpose conjugate ($\hat{A} = \hat{A}^\dagger = (\hat{A}^T)^*$). For example the matrix $\mathcal{M} = \begin{pmatrix} 1 & i \\ -i & 2 \end{pmatrix}$ is Hermitian, as complex conjugate of the transpose $\begin{pmatrix} 1 & -i \\ i & 2 \end{pmatrix}^*$ equals $\begin{pmatrix} 1 & i \\ -i & 2 \end{pmatrix}$.

- 3b** If $\hat{A}$ acting on an eigenstate $|\xi\rangle$ yields eigenvalue a_ξ ($\hat{A}|\xi\rangle = a_\xi|\xi\rangle$), it does so with a probability given by $|\langle\xi|\psi\rangle|^2$.
- 4** If a measurement of $\hat{A}$ gives a particular eigenvalue a_i , then immediately afterwards the system will project out the eigenstate $|\phi_i\rangle$.
- 5** $|\psi\rangle$ has a time dependence, and its evolution with time is given by Schrödinger equation.

$$i\hbar \frac{\partial}{\partial t} |\psi(t)\rangle = \left(-\frac{\hbar^2}{2m} \frac{\partial^2}{\partial x^2} + V \right) |\psi(t)\rangle = \mathcal{H} |\psi(t)\rangle \quad (\text{B.1})$$

where $\mathcal{H}$ is the Hamiltonian (See Appendix A) of the system.

In this respect Quantum Mechanics is a combination of linear algebra and probability theory. The results of classical mechanics, however, still need to be respected. To obtain the results from classical equations of motion, one can average the results of these quantum observables.

When a measurement is made, a wavefunction will immediately collapse to an eigenstate of the operator that gives rise to the measurement. This altered state such as in the case of wave-particle dualism changes the expected result of successive measurements made upon said collapsed state. This gives rise to *commutation relations* between measurements such as position and momentum such as the one shown in Equation B.2

$$[\hat{x}, \hat{p}] = (\hat{x}\hat{p} - \hat{p}\hat{x}) |\psi\rangle = i\frac{\hbar}{2}. \quad (\text{B.2})$$

Classically this is not seen as all operators commute, so the order in which the operations are performed on a state does not matter. In addition to this, Equation (B.2) means that if two observables do not commute, there is a limit to the amount of information that can be gleaned about that system at any one time. This is given by the Heisenberg

uncertainty principle given in Equation (B.3).

$$\Delta x \Delta p = \frac{\hbar}{2}. \quad (\text{B.3})$$

Another consequence of this formalism is that the symmetries of the quantum systems manifest themselves as unitary (or anti-unitary) transformations of the wavefunction.



Basics of Group Theory

Group Theory is another handy mathematical tool used in physics to deal with the symmetries. Transformations that leave a system invariant create mathematical objects called groups. The structures and properties of these groups can be used to theoretically explore the repercussions of the symmetries irrespective of the dynamics of the system [164, p113].

C.1 Tenets of Group Theory

If a set of elements $(r, s, t, \dots)$ exhibit the following multiplicative qualities, then they form a group, $\mathcal{G}$ [164, p86].

Closure If r and s are in $\mathcal{G}$, then $u = rs$ has to also be in $\mathcal{G}$.

Associativity The order in which sets of elements are multiplied together does not matter. So for example $(rs)t = r(st)$. This is not the same as the less general trait of *Commutativity* where the order which the elements are multiplied does not matter ($rs = sr$). Groups that have commutativity are called *Abelian groups*.

Identity There exists an identity element I in $\mathcal{G}$, that when pre- or post-multiplied by

any other element, returns the other element ($rI = Ir = r$).

Invertability For every element in $\mathcal{G}$, there is an element which is the inverse of it such that the two multiplied together returns the identity element ($ss^{-1} = s^{-1}s = I$).

Certain mathematical operations can generate a set of elements that can form groups. These mathematical operations are called *generators* and groups can have several generators.

C.2 Types of Group Relevant to Physics

The realms of Group theory are vast, and there are lots of types of group and they all have many properties. This section will introduce only the groups relevant to a basic understanding of the use of group theory in particle physics.

C.2.1 Lie Groups

Every single group used in particle physics is a Lie group and Lie groups have Lie algebras. The Lie algebras are a set of equations that relate the generators of a group. These relations usually take the form of commutation relations (the same ones used in Quantum Mechanics).

C.2.2 Unitary Groups ($U(n)$)

The unitary group is a set of $n \times n$ matrices that have the Unitarian property: $UU^\dagger = U^\dagger U = 1$. For n larger than 1, the Unitary group is non-Abelian.

C.2.3 Special Unitary Groups ($SU(n)$)

If the unitary matrix also has a determinant of 1, then the group is a Special Unitary Group.



Impact of 2-Lepton significance improvements on the total significance



Equation D.1 is the formula used to represent the total expected median significance of an analysis with individual significances from three separate channels.

$$T_s = c \times \sqrt{x^2 + y^2 + z^2} \quad (\text{D.1})$$

Where x = 0-lepton significance

y = 1-lepton significance

z = 2-lepton significance

c = A correlation factor capturing the total correlation between the channels

T_s = Total significance

These numbers are given by Table D.1.

0-lepton significance	1.91
1-lepton significance	1.98
2-lepton significance	1.03
Correlation	0.98
Total significance	2.87

Table D.1: Evaluated significances of the channels in the Boosted analysis correct as of 20/11/19.

Equation D.1 can be re-arranged to

$$x^2 + y^2 + z^2 = \left(\frac{T_s}{c}\right)^2 \quad (\text{D.2})$$

Since all the significances are taken to two decimal places, The smallest increase in T_s of note is 0.01. If we ‘perturb’ this T_s by 0.01 we go from 2.87 to 2.88 which is an increase in T_s of 0.34%.

$$T'_s = 1.0034 T_s$$

Assuming the correlation between the channels does not change, then if there was to be a increase in T_s of this magnitude and it was all to come from the 2-lepton channel, what change of the 2-lepton significance will do this? This question is expressed in Equation D.3.

$$x^2 + y^2 + \sigma^2 z^2 = \left(\frac{T'_s}{c}\right)^2$$

$$x^2 + y^2 + \sigma^2 z^2 = \left(\frac{T_s}{c}\right)^2 \times 1.0034^2 \quad (\text{D.3})$$

Where σ is the increase in the 2-lepton channel that corresponds to a change in significance for the whole analysis. Subtracting Equation D.3 from D.2 obtains

$$\sigma^2 z^2 - z^2 = \left(\frac{T_s}{c}\right)^2 \times 1.0034^2 - \left(\frac{T_s}{c}\right)^2$$

$$z^2(\sigma^2 - 1) = \left(\frac{T_s}{c}\right)^2 (1.0034^2 - 1) \quad (\text{D.4})$$

Re-arranging D.4 for σ and replacing the variables with their respective values in Table D.1 obtains Equation D.5. This σ is the minimum significance increase required for a change in the 2-lepton analysis to be visible in the whole analysis ($\sigma_{min,2L}^2$).

$$\begin{aligned} \sigma_{min,2L}^2 &= \frac{\left(\frac{T_s}{c}\right)^2 (1.0034^2 - 1)}{z^2} + 1 \\ &= \frac{\left(\frac{2.87}{0.98}\right)^2 (1.0034^2 - 1)}{1.03^2} \\ &= \frac{0.0584}{1.0609} + 1 = 1.055 \end{aligned} \quad (\text{D.5})$$

$$\sigma_{min,2L} = 1.027 = 2.7\%$$

Hence any increase to the expected median significance in the Boosted 2-lepton analysis greater than 2.7% will be affect the significance of the entire Boosted analysis.

Bibliography

- [1] (ATLAS Collaboration) Aaboud M. et al. *Observation of $H \rightarrow b\bar{b}$ decays and VH production with the ATLAS detector*. Vol. B786. 2018, pp. 59–86. DOI: [10.1016/j.physletb.2018.09.013](https://doi.org/10.1016/j.physletb.2018.09.013). arXiv: 1808.08238 [hep-ex].
- [2] (ATLAS Collaboration) Aad G. et al. *Measurement of $H \rightarrow b\bar{b}$ as a function of the vector-boson transverse momentum in 13 TeV pp collisions with the ATLAS detector*. Vol. 141. updated May 2019. Mar. 2019. DOI: [https://doi.org/10.1007/JHEP05\(2019\)141](https://doi.org/10.1007/JHEP05(2019)141).
- [3] (ATLAS Collaboration) Aad G. et al. *Measurements of WH and ZH production in the $H \rightarrow b\bar{b}$ decay channel in pp collisions at 13 TeV with the ATLAS detector*. July 2020. arXiv: 2007.02873 [hep-ex].
- [4] (ATLAS Collaboration) Aad G. et al. *Measurement of the associated production of a Higgs boson decaying into b -quarks with a vector boson at high transverse momentum in pp collisions at $\sqrt{s} = 13$ TeV with the ATLAS detector*. Aug. 2020. arXiv: 2008.02508 [hep-ex].
- [5] R. Nave. *Simulations*. <http://hyperphysics.phy-astr.gsu.edu/hbase/Forces/couple.html>. (Accessed: 20th December 2020). 2020.
- [6] CERN. *The Z Boson*. <https://home.cern/science/physics/z-boson>. (Accessed: 11th September 2020). 2020.

- [7] (ATLAS Collaboration) Aad Georges et al. *Observation of a new particle in the search for the Standard Model Higgs boson with the ATLAS detector at the LHC*. Vol. B716. 2012, pp. 1–29. DOI: [10.1016/j.physletb.2012.08.020](https://doi.org/10.1016/j.physletb.2012.08.020). arXiv: [1207.7214](https://arxiv.org/abs/1207.7214) [hep-ex].
- [8] (CMS Collaboration) Chatrchyan Serguei et al. *Observation of a new boson at a mass of 125 GeV with the CMS experiment at the LHC*. Vol. B716. 2012, pp. 30–61. DOI: [10.1016/j.physletb.2012.08.021](https://doi.org/10.1016/j.physletb.2012.08.021). arXiv: [1207.7235](https://arxiv.org/abs/1207.7235) [hep-ex].
- [9] (Particle Data Group) Tanabashi M. et al. *Particle Physics Booklet*. Extracted from the Review of Particle Physics M. Tanabashi *et al.* (Particle Data Group), Phys. Rev. D **98**, no.3 , 030001 (2018). 2018. DOI: [10.1103/PhysRevD.98.030001](https://doi.org/10.1103/PhysRevD.98.030001).
- [10] (Particle Data Group) Patrignani C. et al. *Particle Physics Booklet*. Extracted from the Review of Particle Physics C. Patrignani *et al.* (Particle Data Group), Chin. Phys. C **40**, no.10 , 100001 (2016). 2016. DOI: [10.1088/1674-1137/40/10/100001](https://doi.org/10.1088/1674-1137/40/10/100001).
- [11] (Particle Data Group) Olive K.A. et al. *Particle Physics Booklet*. Extracted from the Review of Particle Physics K.A Olive *et al.* (Particle Data Group), Chin. Phys. C **38** , 090001 (2014). Chinese Physics C, 2014. DOI: [10.1088/1674-1137/38/9/090001](https://doi.org/10.1088/1674-1137/38/9/090001).
- [12] (Particle Data Group) Beringer J. et al. *Particle Physics Booklet*. Extracted from the Review of Particle Physics J. Beringer *et al.* (Particle Data Group), Phys. Rev. D **86**, 010001 (2012). APS Physics, 2012. DOI: [10.1103/PhysRevD.86.010001](https://doi.org/10.1103/PhysRevD.86.010001).
- [13] (Particle Data Group) Nakamura K. et al. *Particle Physics Booklet*. Extracted from the Review of Particle Physics K. Nakamura *et al.* (Particle Data Group), Journal of Physics G **37**, 075021 (2010). Institute of Physics Publishing, 2010. DOI: [10.1088/0954-3899/37/7A/075021](https://doi.org/10.1088/0954-3899/37/7A/075021).

- [14] (Particle Data Group) Amsler C. et al. *Particle Physics Booklet*. Extracted from the Review of Particle Physics C. Amsler *et al.* (Particle Data Group), Physics Letters **B 667**, 1 (2008). Elsevier, 2008. DOI: [10.1016/j.physletb.2008.07.018](https://doi.org/10.1016/j.physletb.2008.07.018).
- [15] (Particle Data Group) Yao Y-M. et al. *Particle Physics Booklet*. Extracted from the Review of Particle Physics Y-M. Yao *et al.* (Particle Data Group), Journal of Physics G **33**, 1 (2006). Institute of Physics Publishing, 2006. DOI: [10.1088/0954-3899/33/1/001](https://doi.org/10.1088/0954-3899/33/1/001).
- [16] (Particle Data Group) Hagiwara K. et al. *Particle Physics Booklet*. Extracted from the Review of Particle Physics K. Hagiwara *et al.* (Particle Data Group), Phys. Rev. D **66**, 010001 (2002). American Institute of Physics, 2002. DOI: [10.1103/PhysRevD.66.010001](https://doi.org/10.1103/PhysRevD.66.010001).
- [17] (Particle Data Group) Groom D.E. et al. *Particle Physics Booklet*. Extracted from the Review of Particle Physics D.E. Groom *et al.* (Particle Data Group), The European Physical Journal **C15**, 1 (2000). Springer, 2000. URL: <https://cds.cern.ch/record/788143>.
- [18] (Particle Data Group) Caso C. et al. *Particle Physics Booklet*. Extracted from the Review of Particle Physics C. Caso *et al.* (Particle Data Group), The European Physical Journal **C3**, 1 (1998). Springer, 1998. DOI: [10.1007/s10052-998-0104-x](https://doi.org/10.1007/s10052-998-0104-x).
- [19] (Particle Data Group) Aguilar-Benitez M. et al. *Particle Physics Booklet*. Extracted from the Review of Particle Properties M. Aguilar-Benitez *et al.* (Particle Data Group), Physical Review **D50**, 1173 (1994). American Institute of Physics, 1994. DOI: [10.1103/PhysRevD.50.1173](https://doi.org/10.1103/PhysRevD.50.1173).
- [20] Thomas Young. *Experimental Demonstration of the General Law of the Interference of Light*. Vol. 94. 1804. DOI: <https://doi.org/10.1098/rstl.1804.0001>.

- [21] Paul Adrien Maurice Dirac. *Quantised Singularities in the Electromagnetic Field*, vol. A133. 821. 1931, pp. 60–72. DOI: [10.1098/rspa.1931.0130](https://doi.org/10.1098/rspa.1931.0130).
- [22] Albert Einstein. *Einstein: Relativity*. Ed. by R.W Lawson. Routledge Classics, London, 2012. ISBN: 9780415253840.
- [23] Matthew D. Schwartz. *Quantum Field Theory and the Standard Model*. Cambridge University Press, Mar. 2014. ISBN: 978-1-107-03473-0, 978-1-107-03473-0.
- [24] M. Srednicki. *Quantum field theory*. Cambridge University Press, Jan. 2007. ISBN: 978-0-521-86449-7, 978-0-511-26720-8.
- [25] Peter W. Higgs. *Broken Symmetries and the Masses of Gauge Bosons*. Ed. by J.C. Taylor. Vol. 13. 1964, pp. 508–509. DOI: [10.1103/PhysRevLett.13.508](https://doi.org/10.1103/PhysRevLett.13.508).
- [26] F. Englert and R. Brout. *Broken Symmetry and the Mass of Gauge Vector Mesons*. Ed. by J.C. Taylor. Vol. 13. 1964, pp. 321–323. DOI: [10.1103/PhysRevLett.13.321](https://doi.org/10.1103/PhysRevLett.13.321).
- [27] G.S. Guralnik, C.R. Hagen, and T.W.B. Kibble. *Global Conservation Laws and Massless Particles*. Ed. by J.C. Taylor. Vol. 13. 1964, pp. 585–587. DOI: [10.1103/PhysRevLett.13.585](https://doi.org/10.1103/PhysRevLett.13.585).
- [28] Transum.org. *Graph Plotter*.
<https://www.transum.org/Maths/Activity/Graph/Desmos.asp>.
(Accessed: 1st May 2020). 2020.
- [29] (Particle Data Group) Tanabashi M. et al. *Particle Physics Booklet*. Extracted from the Review of Particle Physics M. Tanabashi *et al.* (Particle Data Group), Phys. Rev. D **98**, no.3 , 030001 (2018) and 2019 update. 2018. URL: http://pdg.lbl.gov/2019/tables/contents_tables.html.
- [30] Steven Weinberg. *Baryon- and Lepton-Nonconserving Processes*. Vol. 43. American Physical Society, Nov. 1979, pp. 1566–1570. DOI: [10.1103/PhysRevLett.43.1566](https://doi.org/10.1103/PhysRevLett.43.1566). URL: <https://link.aps.org/doi/10.1103/PhysRevLett.43.1566>.

- [31] Landon Lehman. *Extending the Standard Model Effective Field Theory with the Complete Set of Dimension-7 Operators*. Vol. 90. 12. 2014, p. 125023. DOI: [10.1103/PhysRevD.90.125023](https://doi.org/10.1103/PhysRevD.90.125023). arXiv: 1410.4193 [hep-ph].
- [32] CERN. *Where did it all begin?* <https://home.cern/about/who-we-are/our-history>. (Accessed: 2nd February 2020). 2019.
- [33] Kevin Kofal. AMS. <https://ams.nasa.gov/>. (Accessed: 12th February 2020). 2016.
- [34] CERN. NA62. <https://greybook.cern.ch/greybook/experiment/detail?id=NA62>. (Accessed: 12th February 2020). 2019.
- [35] CERN. ATLAS. <https://greybook.cern.ch/greybook/experiment/detail?id=ATLAS>. (Accessed: 12th February 2020). 2019.
- [36] CERN. CMS. <https://greybook.cern.ch/greybook/experiment/detail?id=CMS>. (Accessed: 12th February 2020). 2019.
- [37] "Science Learning Hub". *"Particle Physics Research"*. https://www.sciencelearn.org.nz/resources/1724-particle_physics_research. (Accessed: 2nd February 2020). 2011.
- [38] CERN. *The Proton Synchrotron*. <https://home.cern/science/accelerators/proton-synchrotron>. (Accessed: 2nd February 2020). 2020.
- [39] CERN Courier. *Super Proton Synchrotron Marks Its 25th Birthday*. <https://cerncourier.com/a/super-proton-synchrotron-marks-its-25th-birthday/>. (Accessed: 2nd February 2020). 2001.

- [40] Abdus Salam. *Weak and Electromagnetic Interactions*. Vol. 680519. 1968, pp. 367–377. DOI: [10.1142/9789812795915_0034](https://doi.org/10.1142/9789812795915_0034).
- [41] S.L. Glashow and Steven Weinberg. *Breaking chiral symmetry*. Vol. 20. 1968, pp. 224–227. DOI: [10.1103/PhysRevLett.20.224](https://doi.org/10.1103/PhysRevLett.20.224).
- [42] Nobel Media AB. *The Nobel Prize in Physics 1979*.
<https://www.nobelprize.org/prizes/physics/1979/summary/>.
(Accessed: 2rd February 2020). 2020.
- [43] Nobel Media AB. *The Nobel Prize in Physics 1984*.
<https://www.nobelprize.org/prizes/physics/1984/summary/>.
(Accessed: 12th February 2020). 2020.
- [44] CERN. *The Large Electron-Positron Collider*.
<https://home.cern/science/accelerators/large-electron-positron-collider>. (Accessed: 5th February 2020). 2020.
- [45] *A Combination of preliminary electroweak measurements and constraints on the standard model*. Feb. 2001. arXiv: [hep-ex/0103048](https://arxiv.org/abs/hep-ex/0103048).
- [46] CERN. *LHC Magnet Types*. https://lhc-machine-outreach.web.cern.ch/components/magnets/types_of_magnets.htm. (Accessed: 11th February 2020). 2014.
- [47] H. Wollnik. *Optics of Charged Particles*. 1st. Academic Press, 1987, pp. 33, 57.
- [48] Sarah Charley. *The LHC Does a Dry Run*.
<https://www.symmetrismagazine.org/article/march-2015/the-lhc-does-a-dry-run>. (Accessed: 2rd February 2020). 2015.
- [49] CERN. *LHCb*. <https://greybook.cern.ch/greybook/experiment/detail?id=LHCB>.
(Accessed: 12th February 2020). 2019.

- [50] CERN. *ALICE*. <https://greybook.cern.ch/greybook/experiment/detail?id=ALICE>. (Accessed: 12th February 2020). 2019.
- [51] CERN. *TOTEM*. <https://greybook.cern.ch/greybook/experiment/detail?id=TOTEM>. (Accessed: 12th February 2020). 2019.
- [52] CERN. *LHCf*. <https://greybook.cern.ch/greybook/experiment/detail?id=LHCf>. (Accessed: 12th February 2020). 2019.
- [53] CERN. *MoEDAL*. <https://greybook.cern.ch/greybook/experiment/detail?id=MoEDAL>. (Accessed: 12th February 2020). 2019.
- [54] J.J. Goodson. “Search for Supersymmetry in States with Large Missing Transverse Momentum and Three Leptons including a Z-Boson”. Presented 17 Apr 2012. PhD thesis. Stony Brook University, May 2012. DOI: [10.1103/PhysRevLett.108.261804](https://doi.org/10.1103/PhysRevLett.108.261804).
- [55] CERN. *CERN’s Accelerator Complex*. <https://home.cern/science/accelerators/accelerator-complex>. (Accessed: 10th February 2020). 2020.
- [56] CERN. *COMPASS*. <https://greybook.cern.ch/greybook/experiment/detail?id=COMPASS>. (Accessed: 12th February 2020). 2019.
- [57] P. Grafstroem and W. Kozanecki. *Luminosity determination at proton colliders*. Vol. 81. 2015, pp. 97–148. DOI: [10.1016/j.ppnp.2014.11.002](https://doi.org/10.1016/j.ppnp.2014.11.002).
- [58] CERN. *The HL-LHC project*. <https://hilumilhc.web.cern.ch/content/hl-lhc-project>. (Accessed September 29th 2020). 2020.

- [59] S. Dittmaier et al. *Handbook of LHC Higgs Cross Sections: 1. Inclusive Observables*. Jan. 2011. DOI: [10.5170/CERN-2011-002](https://doi.org/10.5170/CERN-2011-002). arXiv: [1101.0593](https://arxiv.org/abs/1101.0593) [hep-ph].
- [60] CERN. *Interim Summary Report on the Analysis of the 19 September 2008 Incident at the LHC*. Tech. rep. ATL-COM-DAQ-2018-047. Geneva: CERN, Oct. 2008. URL: https://edms.cern.ch/ui/file/973073/1/Report_on_080919_incident_at_LHC__2_.pdf.
- [61] Jurgen Schukraft. *First Results from the ALICE experiment at the LHC*. Ed. by Premomoy Ghosh, Tapan Nayak, and Raghava Varma. Vol. 862. 2011, pp. 78–84. DOI: [10.1016/j.nuclphysa.2011.05.023](https://doi.org/10.1016/j.nuclphysa.2011.05.023). arXiv: [1103.3474](https://arxiv.org/abs/1103.3474) [hep-ex].
- [62] (LHCb Collaboration) Aaij Roel et al. *Observation of $J/\psi p$ Resonances Consistent with Pentaquark States in $\Lambda_b^0 \rightarrow J/\psi K^- p$ Decays*. Vol. 115. 2015, p. 072001. DOI: [10.1103/PhysRevLett.115.072001](https://doi.org/10.1103/PhysRevLett.115.072001). arXiv: [1507.03414](https://arxiv.org/abs/1507.03414) [hep-ex].
- [63] Nobel Media AB. *The Nobel Prize in Physics 2013*. <https://www.nobelprize.org/prizes/physics/2013/summary/>. (Accessed: 14th February 2020). 2020.
- [64] CERN. *The ATLAS Collaboration*. <https://atlas.cern/discover/collaboration>. (Accessed: 3rd February 2020). 2020.
- [65] "The ATLAS Collaboration". *"Luminosity Public Results"*. <https://twiki.cern.ch/twiki/bin/view/AtlasPublic/LuminosityPublicResults>. (Accessed: 18th March 2020). 2013.
- [66] The ATLAS Collaboration. *"Luminosity Public Results Run2"*. <https://twiki.cern.ch/twiki/bin/view/AtlasPublic/LuminosityPublicResultsRun2>. (Accessed: 14th February 2020). 2019.
- [67] Joao Pequenaio. "Computer generated image of the ATLAS Muons subsystem". Mar. 2008. URL: <https://cds.cern.ch/record/1095929>.

- [68] (ATLAS Collaboration) Aad Georges et al. *Resolution of the ATLAS muon spectrometer monitored drift tubes in LHC Run 2*. Vol. 14. 09. 2019, P09011. DOI: [10.1088/1748-0221/14/09/P09011](https://doi.org/10.1088/1748-0221/14/09/P09011). arXiv: [1906.12226](https://arxiv.org/abs/1906.12226) [hep-ex].
- [69] (ATLAS Collaboration) Aad G. et al. *The ATLAS Experiment at the CERN Large Hadron Collider*. Vol. 3. 2008, S08003. DOI: [10.1088/1748-0221/3/08/S08003](https://doi.org/10.1088/1748-0221/3/08/S08003).
- [70] (ATLAS Collaboration) Aad Georges et al. *Performance of the ATLAS muon trigger in pp collisions at $\sqrt{s} = 8$ TeV*. Vol. C75. 2015, p. 120. DOI: [10.1140/epjc/s10052-015-3325-9](https://doi.org/10.1140/epjc/s10052-015-3325-9). arXiv: [1408.3179](https://arxiv.org/abs/1408.3179) [hep-ex].
- [71] Joao Pequeno. "Computer Generated image of the ATLAS calorimeter". Mar. 2008. URL: <https://cds.cern.ch/record/1095927>.
- [72] W. H. Bragg and R. Kleeman. XXXIX. *On the α particles of radium, and their loss of range in passing through various atoms and molecules*. Sept. 1905. DOI: [10.1080/14786440509463378](https://doi.org/10.1080/14786440509463378). URL: <https://doi.org/10.1080/14786440509463378>.
- [73] Alessandro Bettini. *Introduction to elementary particle physics*. 2008. ISBN: 9780521880213. URL: <http://www.cambridge.org/catalogue/catalogue.asp?isbn=9780521880213>.
- [74] The ATLAS Collaboration. *Public Liquid-Argon Calorimeter Plots on Detector Status*. 2019. URL: <https://twiki.cern.ch/twiki/bin/view/AtlasPublic/LArCaloPublicResultsDetStatus>.
- [75] J P Badiou et al. *ATLAS barrel toroid: Technical Design Report*. Technical Design Report ATLAS. Electronic version not available. Geneva: CERN, 1997. URL: <http://cds.cern.ch/record/331065>.
- [76] *ATLAS magnet system: Technical Design Report, 1*. Technical Design Report ATLAS. Geneva: CERN, 1997. URL: <http://cds.cern.ch/record/338080>.

- [77] *ATLAS end-cap toroids: Technical Design Report*. Technical Design Report ATLAS. Electronic version not available. Geneva: CERN, 1997. URL: <http://cds.cern.ch/record/331066>.
- [78] Joao Pequenaio. “Computer generated image of the ATLAS inner detector”. Mar. 2008. URL: <https://cds.cern.ch/record/1095926>.
- [79] DuPont de Nemours Inc. *Kapton® Polyimide films*. <https://www.dupont.com/electronic-materials/kapton-polyimide-film.html>. (Accessed: 29th March 2020). 2020.
- [80] A Vogel. *ATLAS Transition Radiation Tracker (TRT): Straw Tube Gaseous Detectors at High Rates*. Tech. rep. ATL-INDET-PROC-2013-005. Geneva: CERN, Apr. 2013. URL: <https://cds.cern.ch/record/1537991>.
- [81] C. L. L. Yaun. *Some New Developments On Transition Radiation Detectors for High-Energy Particles*. Oct. 1974. DOI: [10.1080/14786440509463378](https://doi.org/10.1080/14786440509463378). URL: <https://lss.fnal.gov/archive/1975/conf/fermilab-conf-75-111-e.pdf>.
- [82] Helmuth Spieler. *Semiconductor Detector Systems*. Vol. v.12. Semiconductor Science and Technology. Oxford: Oxford University Press, 2005. ISBN: 978-0-19-852784-8.
- [83] KEK Images. *Index of images*. http://rd.kek.jp/images/SOI_.jpg. (Accessed: 22rd February 2020). 2012.
- [84] Frank Hartmann. *Evolution of Silicon Sensor Technology in Particle Physics*. Vol. 231. 2009, pp. 1–204.
- [85] (ATLAS Collaboration) Aad Georges et al. *Operation and performance of the ATLAS semiconductor tracker*. Vol. 9. 2014, P08009. DOI: [10.1088/1748-0221/9/08/P08009](https://doi.org/10.1088/1748-0221/9/08/P08009). arXiv: [1404.7473 \[hep-ex\]](https://arxiv.org/abs/1404.7473).

- [86] *ATLAS inner detector: Technical Design Report, 1*. Technical Design Report ATLAS. Geneva: CERN, 1997. URL: <https://cds.cern.ch/record/331063>.
- [87] (ATLAS Collaboration) Abbott B. et al. *Production and Integration of the ATLAS Insertable B-Layer*. Vol. 13. 05. 2018, T05008. DOI: [10.1088/1748-0221/13/05/T05008](https://doi.org/10.1088/1748-0221/13/05/T05008). arXiv: [1803.00844](https://arxiv.org/abs/1803.00844) [physics.ins-det].
- [88] Joao Pequeno and Paul Schaffner. “How ATLAS detects particles: diagram of particle paths in the detector”. Jan. 2013. URL: <https://cds.cern.ch/record/1505342>.
- [89] *Identification of hadronic tau lepton decays using neural networks in the ATLAS experiment*. Tech. rep. ATL-PHYS-PUB-2019-033. Geneva: CERN, Aug. 2019. URL: <https://cds.cern.ch/record/2688062>.
- [90] (ATLAS Collaboration) Abolins M. et al. *The ATLAS Data Acquisition and High Level Trigger system*. Vol. 11. 06. 2016, P06008. DOI: [10.1088/1748-0221/11/06/P06008](https://doi.org/10.1088/1748-0221/11/06/P06008).
- [91] The ATLAS Collaboration. *ATLAS Phase II Letter of Intent*. 2013. URL: <https://inspirehep.net/files/2ecbd8cca6cab9ab09aa49ed51249420>.
- [92] The ATLAS Collaboration. *Technical Design Report for the ATLAS Inner Tracker Strip Detector*. Tech. rep. CERN-LHCC-2017-005. ATLAS-TDR-025. Geneva: CERN, Apr. 2017. URL: <http://cds.cern.ch/record/2257755>.
- [93] Lev Davidovich Landau. *On the energy loss of fast particles by ionization*. Vol. 8. 4. 1944, pp. 201–205. URL: <https://cds.cern.ch/record/216256>.
- [94] W. Dabrowski et al. *Design and performance of the ABCD chip for the binary readout of silicon strip detectors in the ATLAS semiconductor tracker*. Vol. 47. 2000, pp. 1843–1850. DOI: [10.1109/23.914457](https://doi.org/10.1109/23.914457).

- [95] L. Poley et al. *The ABC130 barrel module prototyping programme for the ATLAS strip tracker*. Vol. 15. 09. 2020, P09004. DOI: [10.1088/1748-0221/15/09/P09004](https://doi.org/10.1088/1748-0221/15/09/P09004). arXiv: [2009.03197](https://arxiv.org/abs/2009.03197) [physics.ins-det].
- [96] W. Dabrowski et al. *Design and performance of the ABCN-25 readout chip for ATLAS inner detector upgrade*. 2009, pp. 373–380. DOI: [10.1109/NSSMIC.2009.5401687](https://doi.org/10.1109/NSSMIC.2009.5401687).
- [97] AnySilicon. *eFuse IP Core*. <https://anysilicon.com/semipedia/efuse-ip-core/>. (Accessed: 20th August 2020). July 2019.
- [98] C.P. Robert. *Monte Carlo Methods in Statistics*. Sept. 2009. URL: <https://arxiv.org/pdf/0909.0389.pdf>.
- [99] Andy Buckley et al. *General-purpose event generators for LHC physics*. Vol. 504. 2011, pp. 145–233. DOI: [10.1016/j.physrep.2011.03.005](https://doi.org/10.1016/j.physrep.2011.03.005). arXiv: [1101.2599](https://arxiv.org/abs/1101.2599) [hep-ph].
- [100] Christian P. Robert and George Casella. *Monte Carlo Statistical Methods*. Ed. by R.W Lawson. 2nd. Springer, New York, 2004. ISBN: 9781475741452. URL: https://archive.org/details/springer_10.1007-978-1-4757-4145-2.
- [101] Christophe Andrieu et al. *An introduction to MCMC for machine learning*. English. Vol. 50. 1-2. Springer Verlag, Jan. 2003, pp. 5–43. DOI: [10.1023/A:1020281327116](https://doi.org/10.1023/A:1020281327116).
- [102] SLAC National Accelerator Laboratory. *Simulations*. <https://theory.slac.stanford.edu/our-research/simulations>. (Accessed: 19th September 2020). 2020.

- [103] Yuri L. Dokshitzer. *Calculation of the Structure Functions for Deep Inelastic Scattering and $e^+ e^-$ Annihilation by Perturbation Theory in Quantum Chromodynamics*. Vol. 46. 1977, pp. 641–653. URL: http://www.jetp.ac.ru/cgi-bin/dn/e_046_04_0641.pdf.
- [104] V.N. Gribov and L.N. Lipatov. *Deep inelastic $e p$ scattering in perturbation theory*. Vol. 15. 1972, pp. 438–450.
- [105] Guido Altarelli and G. Parisi. *Asymptotic Freedom in Parton Language*. Vol. 126. 1977, pp. 298–318. DOI: [10.1016/0550-3213\(77\)90384-4](https://doi.org/10.1016/0550-3213(77)90384-4).
- [106] Torbjorn Sjostrand, Stephen Mrenna, and Peter Z. Skands. *A Brief Introduction to PYTHIA 8.1*. Vol. 178. 2008, pp. 852–867. DOI: [10.1016/j.cpc.2008.01.036](https://doi.org/10.1016/j.cpc.2008.01.036). arXiv: [0710.3820 \[hep-ph\]](https://arxiv.org/abs/0710.3820).
- [107] Bo Andersson et al. *Parton Fragmentation and String Dynamics*. Vol. 97. 1983, pp. 31–145. DOI: [10.1016/0370-1573\(83\)90080-7](https://doi.org/10.1016/0370-1573(83)90080-7).
- [108] *ATLAS Pythia 8 tunes to 7 TeV datas*. Tech. rep. ATL-PHYS-PUB-2014-021. Geneva: CERN, Nov. 2014. URL: <https://cds.cern.ch/record/1966419>.
- [109] (ATLAS Collaboration) Aad G. et al. *Measurement of the Z/γ^* boson transverse momentum distribution in pp collisions at $\sqrt{s} = 7$ TeV with the ATLAS detector*. Vol. 09. 2014, p. 145. DOI: [10.1007/JHEP09\(2014\)145](https://doi.org/10.1007/JHEP09(2014)145). arXiv: [1406.3660 \[hep-ex\]](https://arxiv.org/abs/1406.3660).
- [110] John Allison et al. *Geant4 developments and applications*. Vol. 53. 2006, p. 270. DOI: [10.1109/TNS.2006.869826](https://doi.org/10.1109/TNS.2006.869826).
- [111] G. Duckeck et al., eds. *ATLAS computing: Technical design report*. June 2005. URL: <http://cds.cern.ch/record/837738>.
- [112] R. Fruhwirth. *Application of Kalman filtering to track and vertex fitting*. Vol. A262. 1987, pp. 444–450. DOI: [10.1016/0168-9002\(87\)90887-4](https://doi.org/10.1016/0168-9002(87)90887-4).

- [113] Thijs G. Cornelissen et al. *The global χ^2 track fitter in ATLAS*. Vol. 119. 2008, p. 032013. DOI: [10.1088/1742-6596/119/3/032013](https://doi.org/10.1088/1742-6596/119/3/032013).
- [114] *Modelling of Track Reconstruction Inside Jets with the 2016 ATLAS $\sqrt{s} = 13$ TeV pp dataset*. Tech. rep. ATL-PHYS-PUB-2017-016. Geneva: CERN, July 2017. URL: <https://cds.cern.ch/record/2275639>.
- [115] The ATLAS Collaboration. *Early Inner Detector Tracking Performance in the 2015 data at $\sqrt{s} = 13$ TeV*. ATL-PHYS-PUB-2015-051. Geneva, Dec. 2015. URL: <http://cds.cern.ch/record/2110140>.
- [116] (ATLAS Collaboration) Aaboud Morad et al. *Reconstruction of primary vertices at the ATLAS experiment in Run 1 proton–proton collisions at the LHC*. Vol. 77. 5. 2017, p. 332. DOI: [10.1140/epjc/s10052-017-4887-5](https://doi.org/10.1140/epjc/s10052-017-4887-5). arXiv: [1611.10235](https://arxiv.org/abs/1611.10235) [physics.ins-det].
- [117] W Lampl et al. *Calorimeter Clustering Algorithms: Description and Performance*. ATL-LARG-PUB-2008-002. ATL-COM-LARG-2008-003. Geneva, Apr. 2008. URL: <http://cds.cern.ch/record/1099735>.
- [118] The ATLAS Collaboration. *Electron and photon reconstruction and performance in ATLAS using a dynamical, topological cell clustering-based approach*. ATL-PHYS-PUB-2017-022. Geneva, Dec. 2017. URL: <https://cds.cern.ch/record/2298955>.
- [119] Byron.P Roe et al. *Boosted Decision Trees as an Alternative to Artificial Neural Networks for Particle Identification*. 2004. DOI: [10.1016/j.nima.2004.12.018](https://doi.org/10.1016/j.nima.2004.12.018). arXiv: [1302.0260](https://arxiv.org/abs/1302.0260) [hep-ph].
- [120] (ATLAS Collaboration) Aaboud Morad et al. *Electron efficiency measurements with the ATLAS detector using 2012 LHC proton–proton collision data*. Vol. 77. 3. 2017, p. 195. DOI: [10.1140/epjc/s10052-017-4756-2](https://doi.org/10.1140/epjc/s10052-017-4756-2). arXiv: [1612.01456](https://arxiv.org/abs/1612.01456) [hep-ex].

- [121] (ATLAS Collaboration) Aad Georges et al. *Muon reconstruction performance of the ATLAS detector in proton–proton collision data at $\sqrt{s}=13$ TeV*. Vol. 76. 5. 2016, p. 292. DOI: [10.1140/epjc/s10052-016-4120-y](https://doi.org/10.1140/epjc/s10052-016-4120-y). arXiv: [1603.05598](https://arxiv.org/abs/1603.05598) [hep-ex].
- [122] The ATLAS Collaboration. *Measurement of the tau lepton reconstruction and identification performance in the ATLAS experiment using pp collisions at $\sqrt{s}=13$ TeV*. ATLAS-CONF-2017-029. Geneva, May 2017. URL: <https://cds.cern.ch/record/2261772>.
- [123] (ATLAS Collaboration) Aaboud Morad et al. *Performance of missing transverse momentum reconstruction with the ATLAS detector using proton-proton collisions at $\sqrt{s}=13$ TeV*. Vol. 78. 11. 2018, p. 903. DOI: [10.1140/epjc/s10052-018-6288-9](https://doi.org/10.1140/epjc/s10052-018-6288-9). arXiv: [1802.08168](https://arxiv.org/abs/1802.08168) [hep-ex].
- [124] Kenneth G. Wilson. *Confinement of Quarks*. Ed. by J.C. Taylor. Feb. 1974, pp. 45–59. DOI: [10.1103/PhysRevD.10.2445](https://doi.org/10.1103/PhysRevD.10.2445).
- [125] John E. Huth et al. *Toward a standardization of jet definitions*. Dec. 1990, pp. 0134–136. URL: https://www.academia.edu/15824454/Toward_a_standardization_of_jet_definitions.
- [126] Gavin P. Salam. *Towards Jetography*. Vol. C67. 2010, pp. 637–686. DOI: [10.1140/epjc/s10052-010-1314-6](https://doi.org/10.1140/epjc/s10052-010-1314-6). arXiv: [0906.1833](https://arxiv.org/abs/0906.1833) [hep-ph].
- [127] Matteo Cacciari, Gavin P. Salam, and Gregory Soyez. *The Anti-k(t) jet clustering algorithm*. Vol. 04. 2008, p. 063. DOI: [10.1088/1126-6708/2008/04/063](https://doi.org/10.1088/1126-6708/2008/04/063). arXiv: [0802.1189](https://arxiv.org/abs/0802.1189) [hep-ph].
- [128] *Variable Radius, Exclusive- k_T , and Center-of-Mass Subjet Reconstruction for Higgs($\rightarrow b\bar{b}$) Tagging in ATLAS*. ATL-PHYS-PUB-2017-010. Geneva, June 2017. URL: <https://cds.cern.ch/record/2268678>.

- [129] Jessie Shelton. “Jet Substructure”. In: *Theoretical Advanced Study Institute in Elementary Particle Physics: Searching for New Physics at Small and Large Scales*. 2013, pp. 303–340. DOI: [10.1142/9789814525220_0007](https://doi.org/10.1142/9789814525220_0007). arXiv: [1302.0260](https://arxiv.org/abs/1302.0260) [hep-ph].
- [130] Mrinal Dasgupta et al. *Towards an understanding of jet substructure*. Vol. 09. 2013, p. 029. DOI: [10.1007/JHEP09\(2013\)029](https://doi.org/10.1007/JHEP09(2013)029). arXiv: [1307.0007](https://arxiv.org/abs/1307.0007) [hep-ph].
- [131] Stephen D. Ellis, Christopher K. Vermilion, and Jonathan R. Walsh. *Techniques for improved heavy particle searches with jet substructure*. Vol. 80. 2009, p. 051501. DOI: [10.1103/PhysRevD.80.051501](https://doi.org/10.1103/PhysRevD.80.051501). arXiv: [0903.5081](https://arxiv.org/abs/0903.5081) [hep-ph].
- [132] David Krohn, Jesse Thaler, and Lian-Tao Wang. *Jet Trimming*. Vol. 02. 2010, p. 084. DOI: [10.1007/JHEP02\(2010\)084](https://doi.org/10.1007/JHEP02(2010)084). arXiv: [0912.1342](https://arxiv.org/abs/0912.1342) [hep-ph].
- [133] *Jet Calibration and Systematic Uncertainties for Jets Reconstructed in the ATLAS Detector at $\sqrt{s} = 13$ TeV*. Tech. rep. ATL-PHYS-PUB-2015-015. Geneva: CERN, July 2015. URL: <https://cds.cern.ch/record/2037613>.
- [134] (ATLAS Collaboration) Aad Georges et al. *ATLAS b -jet identification performance and efficiency measurement with $t\bar{t}$ events in pp collisions at $\sqrt{s} = 13$ TeV*. Vol. 79. 11. 2019, p. 970. DOI: [10.1140/epjc/s10052-019-7450-8](https://doi.org/10.1140/epjc/s10052-019-7450-8). arXiv: [1907.05120](https://arxiv.org/abs/1907.05120) [hep-ex].
- [135] (ATLAS Collaboration) Aaboud M et al. *Measurements of b -jet tagging efficiency with the ATLAS detector using $t\bar{t}$ events in pp collisions at $\sqrt{s} = 13$ TeV*. Vol. 89. 2018. DOI: [https://doi.org/10.1007/JHEP08\(2018\)089](https://doi.org/10.1007/JHEP08(2018)089).
- [136] (ATLAS Collaboration) Aaboud M. et al. *Evidence for the $H \rightarrow b\bar{b}$ decay with the ATLAS detector*. Vol. 12. 2017, p. 024. DOI: [10.1007/JHEP12\(2017\)024](https://doi.org/10.1007/JHEP12(2017)024). arXiv: [1708.03299](https://arxiv.org/abs/1708.03299) [hep-ex].
- [137] A.D. Martin et al. *Parton distributions for the LHC*. Vol. 63. 2009, pp. 189–285. DOI: [10.1140/epjc/s10052-009-1072-5](https://doi.org/10.1140/epjc/s10052-009-1072-5). arXiv: [0901.0002](https://arxiv.org/abs/0901.0002) [hep-ph].

- [138] Simone Alioli et al. *A general framework for implementing NLO calculations in shower Monte Carlo programs: the POWHEG BOX*. Vol. 06. 2010, p. 043. DOI: [10.1007/JHEP06\(2010\)043](https://doi.org/10.1007/JHEP06(2010)043). arXiv: [1002.2581](https://arxiv.org/abs/1002.2581) [hep-ph].
- [139] T. Gleisberg et al. *Event generation with SHERPA 1.1*. Vol. 02. 2009, p. 007. DOI: [10.1088/1126-6708/2009/02/007](https://doi.org/10.1088/1126-6708/2009/02/007). arXiv: [0811.4622](https://arxiv.org/abs/0811.4622) [hep-ph].
- [140] (ATLAS Collaboration) Aad Georges et al. *Search for the $b\bar{b}$ decay of the Standard Model Higgs boson in associated $(W/Z)H$ production with the ATLAS detector*. Vol. 01. 2015, p. 069. DOI: [10.1007/JHEP01\(2015\)069](https://doi.org/10.1007/JHEP01(2015)069). arXiv: [1409.6212](https://arxiv.org/abs/1409.6212) [hep-ex].
- [141] Gerhard Bohm and Guenter Zech. *Introduction to statistics and data analysis for physicists*. Hamburg: DESY, 2014. ISBN: 978-3-935702-88-1. DOI: [10.3204/DESY-BOOK/statistics](https://doi.org/10.3204/DESY-BOOK/statistics).
- [142] D. de Florian et al. *Handbook of LHC Higgs Cross Sections: 4. Deciphering the Nature of the Higgs Sector*. Vol. 2/2017. Oct. 2016. DOI: [10.23731/CYRM-2017-002](https://doi.org/10.23731/CYRM-2017-002). arXiv: [1610.07922](https://arxiv.org/abs/1610.07922) [hep-ph].
- [143] Nicolas Berger et al. *Simplified Template Cross Sections - Stage 1.1*. June 2019. arXiv: [1906.02754](https://arxiv.org/abs/1906.02754) [hep-ph].
- [144] (ATLAS Collaboration) Aaboud M. et al. *Measurement of the Inelastic Proton-Proton Cross Section at $\sqrt{s} = 13$ TeV with the ATLAS Detector at the LHC*. Vol. 117. 18. 2016, p. 182002. DOI: [10.1103/PhysRevLett.117.182002](https://doi.org/10.1103/PhysRevLett.117.182002). arXiv: [1606.02625](https://arxiv.org/abs/1606.02625) [hep-ex].
- [145] Stefano Frixione, Paolo Nason, and Carlo Oleari. *Matching NLO QCD computations with Parton Shower simulations: the POWHEG method*. Vol. 11. 2007, p. 070. DOI: [10.1088/1126-6708/2007/11/070](https://doi.org/10.1088/1126-6708/2007/11/070). arXiv: [0709.2092](https://arxiv.org/abs/0709.2092) [hep-ph].

- [146] (ATLAS Collaboration) Aad G. et al. *Jet energy scale and resolution measured in proton-proton collisions at $\sqrt{s} = 13$ TeV with the ATLAS detector*. July 2020. arXiv: [2007.02645 \[hep-ex\]](#).
- [147] S. Dittmaier et al. *Handbook of LHC Higgs Cross Sections: 2. Differential Distributions*. Jan. 2012. DOI: [10.5170/CERN-2012-002](#). arXiv: [1201.3084 \[hep-ph\]](#).
- [148] J R Andersen et al. *Handbook of LHC Higgs Cross Sections: 3. Higgs Properties*. Ed. by S Heinemeyer et al. July 2013. DOI: [10.5170/CERN-2013-004](#). arXiv: [1307.1347 \[hep-ph\]](#).
- [149] Eilam Gross and Ofer Vitells. *Trial factors for the look elsewhere effect in high energy physics*. Vol. 70. 2010, pp. 525–530. DOI: [10.1140/epjc/s10052-010-1470-8](#). arXiv: [1005.1891 \[physics.data-an\]](#).
- [150] G. D’Agostini. *Bayesian reasoning in data analysis: A critical introduction*. Nov. 2003. URL: <https://epdf.pub/bayesian-reasoning-in-data-analysis-a-critical-introduction.html>.
- [151] Borek Puza. *Bayesian Methods for Statistical Analysis*. ANU Press, 2015. ISBN: 9781921934254. URL: <http://www.jstor.org/stable/j.ctt1bgzbn2>.
- [152] Glen Cowan et al. *Asymptotic formulae for likelihood-based tests of new physics*. Vol. 71. [Erratum: Eur.Phys.J.C 73, 2501 (2013)]. 2011, p. 1554. DOI: [10.1140/epjc/s10052-011-1554-0](#). arXiv: [1007.1727 \[physics.data-an\]](#).
- [153] S.S. Wilks. *The Large-Sample Distribution of the Likelihood Ratio for Testing Composite Hypotheses*. Vol. 9. 1. 1938, pp. 60–62. DOI: [10.1214/aoms/1177732360](#).

- [154] Olaf Behnke et al., eds. *Data analysis in high energy physics: A practical guide to statistical methods*. Weinheim, Germany: Wiley-VCH, 2013. ISBN: 978-3-527-65343-0.
- [155] Luca Lista. *Statistical Methods for Data Analysis in Particle Physics*. Vol. 941. Springer, 2017. DOI: [10.1007/978-3-319-62840-0](https://doi.org/10.1007/978-3-319-62840-0).
- [156] (ATLAS Collaboration) Aad Georges et al. *Jet energy measurement and its systematic uncertainty in proton-proton collisions at $\sqrt{s} = 7$ TeV with the ATLAS detector*. Vol. 75. 2015, p. 17. DOI: [10.1140/epjc/s10052-014-3190-y](https://doi.org/10.1140/epjc/s10052-014-3190-y). arXiv: [1406.0076](https://arxiv.org/abs/1406.0076) [hep-ex].
- [157] (ATLAS Collaboration) Aad G. et al. *Performance of electron and photon triggers in ATLAS during LHC Run 2*. 2019. arXiv: [1909.00761](https://arxiv.org/abs/1909.00761) [hep-ex].
- [158] Marcus Morgenstern and Mark Owen. *Muon trigger efficiency in early 2018 data taking*. ATL-COM-DAQ-2018-047. Geneva, May 2017. URL: <https://cds.cern.ch/record/2318500>.
- [159] *Muon Performance in Minimum Bias pp Collision Data at $\sqrt{s} = 7$ TeV with ATLAS*. Tech. rep. ATLAS-CONF-2010-036. Geneva: CERN, July 2010. URL: <https://cds.cern.ch/record/1277675>.
- [160] (ATLAS Collaboration) Aad G. et al. *Object Definitions and Selections for Standard Model $VH \rightarrow ll/l\nu/\nu\nu + b\bar{b}$* . 2018. URL: <https://cds.cern.ch/record/2317182>.
- [161] J. D. Hunter. *Matplotlib: A 2D graphics environment*. Vol. 9. 3. IEEE COMPUTER SOC, 2007, pp. 90–95. DOI: [10.1109/MCSE.2007.55](https://doi.org/10.1109/MCSE.2007.55).
- [162] Glen Cowan et al. *Asymptotic formulae for likelihood-based tests of new physics*. Vol. C71. [Erratum: Eur. Phys. J.C73,2501(2013)]. 2011, p. 1554. DOI: [10.1140/epjc/s10052-011-1554-0](https://doi.org/10.1140/epjc/s10052-011-1554-0). arXiv: [1007.1727](https://arxiv.org/abs/1007.1727) [physics.data-an].

- [163] Paola Cappellaro. *Lecture Notes on Quantum Theory of Radiation Interactions*. (Accessed 11th April 2020). MIT OpenCourseWare, Oct. 2012. URL: https://ocw.mit.edu/courses/nuclear-engineering/22-51-quantum-theory-of-radiation-interactions-fall-2012/lecture-notes/MIT22_51F12_Ch3.pdf.
- [164] Ta-Pei Cheng, Ling-Fong Li, and David Gross. *Gauge Theory of Elementary Particle Physics*. Vol. 38. Jan. 1985. DOI: [10.1063/1.2814821](https://doi.org/10.1063/1.2814821).