Using Upgrades to the ATLAS Hardware Trigger to Search for Rare Couplings of the Higgs Boson through Displaced Jets and Photon Decay Channels

by

Anthony Carroll

A dissertation accepted and approved in partial fulfillment of the

requirements for the degree of

Doctor of Philosophy

in Physics

Dissertation Committee:

Laura Jeanty, Chair

David Strom, Advisor

Spencer Chang, Core Member

David Levin, Institutional Representative

University of Oregon

Summer 2026

© 2026 Anthony Carroll
This work, including text and images of this document but not including supplemental files (for example, not including software code and data), is licensed under a Creative Commons
Attribution 4.0 International License.


DISSERTATION ABSTRACT

Anthony Carroll

Doctor of Philosophy in Physics

Title: Using Upgrades to the ATLAS Hardware Trigger to Search for Rare Couplings of the Higgs Boson through Displaced Jets and Photon Decay Channels

The LHC collides protons to create an event rate of 40 MHz. This generates more data than we are able to save for later analysis. The ATLAS experiment implements a two-tiered trigger system, divided into hardware and software triggers, to reduce the rate of accepted events. Improvements to the trigger will increase the sensitivity to rare physics processes, such as rare couplings and decays of the Higgs boson. The Higgs is a portal to new physics, requiring an increase in selection of Higgs events to study. This thesis presents work done both on the Phase-I upgrades currently operating in Run 3 of the LHC as well as on the Phase-II upgrades in development for the coming HL-LHC Run 4 scheduled to begin in 2030. The Phase-I work consists of implementing a validation framework for the Global Feature Extractor (gFEX), which uses information from the entire ATLAS detector to compute global values such as large radius jets, missing transverse energy, and total scalar energy. The improvements to sensitivity for new physics searches are assessed in the study of trigger efficiencies of a gFEX seeded prompt track fraction (PTF) trigger applied to a search for displaced jets in the ATLAS inner detector. The benchmark model for this search is a Higgs boson pair producing long-lived neutral scalars that decay back into SM quarks. Studies showing improvement to expected limits in the ggF region of this analysis are shown. The Phase-II work consists of studying the efficiency of the Global Trigger upgrade's diphoton algorithm for both prompt and long-lived processes, and studying the performance of a machine learning (ML) based version of the diphoton algorithm.

This dissertation includes unpublished and co-authored material.

ACKNOWLEDGEMENTS

I am incredibly grateful that I got to spend the last several years of my life pursuing this degree. Not only did it challenge me and push me to further my understanding of underlying physics concepts, but it gave me the opportunity to meet some of the most incredible, passionate, smart, and kind people I have ever had the pleasure of working with, and truly feel like I am a part of an amazing international community. I grew not only professionally, but personally throughout my years as a graduate student at the University of Oregon, and that is in no small part due to the people surrounding me.

First and foremost I would like to thank my advisor, David Strom. His enthusiasm for work on gFEX inspired me and helped shape my research interests as I began working with him at the beginning of my time on ATLAS. He is incredibly kind, even though he did try to hit me with his car one time. I would also like to thank the rest of the UO ATLAS group: the faculty Laura Jeanty, Stephanie Majewski, Eric Torrence, and Jim Brau, and all the post docs and students, both past and current, for countless helpful discussions, questions, and venting sessions throughout the duration of my time at UO. Especially thank you to my cohort, Tim, Jennifer, Kehang, and Nathan, for always being willing to help, debug, or just talk. A special thank you to Jen for introducing me to the wonder that is Takoyaki in a cake pop maker. Revolutionary. To the UO post docs, both past and present, Jochen, Dani, Ondra, Despina, Federica, and Giulia, for being so welcoming and making me feel at home at CERN. I'm looking forward to seeing you all again in person soon!

I'd like to thank the gFEX team. David, David, and Sabine did an excellent job in leading the group through some very rocky times during the commissioning of the system at the start of Run 3. To my fellow gFEX students, Greg, Emily, Kristin, and Rabia, thank you for not only teaching me everything I know about gFEX, but also being phenomenal friends and people. Thank you to our engineers, Kirk, who taught me what little bit of VHDL I needed to learn before HLS became widespread, and to Dan, for being the main reason gFEX actually works (and being a good friend too). Thank you to the post docs I worked with on gFEX during my time, Jared, Simone, Despina, and Cecilia, I learned so much from all of you and I'm glad we got to work together for so many years.

I'd like to thank the Displaced Jets Analysis team. Specifically our fearless leaders, Dominique, Elena, and Melissa. This summer is the EB, I just know it. I'd like to thank the other students on the analysis, Emily, Jesse, Connor, Gerwin, and Sidney. I spent so much time working on gFEX that I came into this analysis completely unexperienced, and you all made it easy to learn and grow through both your expertise and your kindness. I know we are going to produce some amazing results when all is said and done. I'd also like to thank Jackson, who has been my default LLP expert to bother since before this analysis even existed.

I'd like to thank the Omega group at BNL. The entire group was so welcoming that I immediately forgot how much I don't like Long Island. I probably owe Liza Brost a full page of acknowledgments, so we will just start typing and see how long it gets. Liza was the best mentor I could have possibly asked for. She helped me revamp my entire professional trajectory, and I can confidently say that I wouldn't have been prepared to find a post doc position and defend this thesis without her guidance over the last year. I'm excited to follow in her footsteps along the Oregon to NIU to BNL pipeline when she eventually hires me once Michael retires and crowns her group leader. I'm looking forward to continuing working with her on diphoton trigger shenanigans with team awesome, even though she insists that I am a millennial (I am not), And I will never learn my lesson about opportunities :p I'd also like to thank Chris. Lunch time catfishing with Chris and Liza in the kitchen were some of my favorite times during the winter and spring at BNL before the flock of shrieking interns returned. Chris's kindness and humor were a large part of what made my time at BNL go by so quickly. I'd like to thank Mounia and Ganesh for being great office mates, I hope to see you both around CERN sometime in the near future! I'd like to thank the other post docs I had the opportunity to work with at BNL: Abe, Ang, Sahibjeet, Punit, and Rea, just for being great people who inspire me to do good physics. I'd like to thank Haider for taking the best worst picture of me ever at Lepton Photon. I use it for all of my official profile pictures now, and I will probably never change it. I'd like to thank Michael Begel for running an incredible group at BNL and being genuinely kind and approachable, a rare quality in someone so accomplished.

I'd like to thank all the friends I have made in my time at Oregon, CERN, and BNL. Jennifer, Nathan, Kehang, Tim, Kristin, Billy, Emily, Mason, Greg, Meg, Casey (and DaVinci), Nik, Ezra, Yvonne, Sylvia, Michelle, Jacob, Emily, Jesse, Phil, Dani, Jochen, Ondra, Despina, Iza, Mark, Andi, Calla, Juliet, Abdullah, Ella, Daniel, and Noe, you all kept me more sane than I had any right to be.

I'd like to thank my family: My parents Ann Marie and James, my sister Elizabeth, and my wife Carol. Their love and support is really the only reason I was able to finish this work, and for that I am incredible grateful.

Finally, I'd like to thank our dog, Booker. He deserves and honorary PhD for the emotional support he provided. Even though he can't read, he makes up for it in terms of snuggles.

This material is based upon work supported by the U.S. Department of Energy Grant DE-SC0017996.

This material is based upon work supported by the U.S. Department of Energy, Office of Science, Office of Workforce Development for Teachers and Scientists, Office of Science Graduate Student Research (SCGSR) program. The SCGSR program is administered by the Oak Ridge Institute for Science and Education for the DOE under contract number DE-SC0014664.

This thesis is dedicated to Carol, who put up with my complaining (both in person and remotely) throughout the whole ordeal.

TABLE OF CONTENTS

Chapter Page

I. INTRODUCTION 24

II. THEORY BACKGROUNDS AND MOTIVATION 27

2.1. Theoretical Formalism: Particles, Fields, and Quantum Field Theory 27

2.1.1. Natural Units 27

2.1.2. Quantum Field Theory Formalism 28

2.1.3. Scattering 29

2.1.4. Gauge Symmetries and Symmetries 31

2.2. The Standard Model of Particle Physics 32

2.2.1. The Electroweak Sector 34

2.2.2. Electroweak Symmetry Breaking and the Higgs Mechanism . . 36

2.2.3. Quantum Chromodynamics 41

2.3. Beyond Standard Model Physics Motivations 44

2.3.1. Motivations for BSM Physics 44

2.3.2. The Higgs as a Portal to New Physics 44

2.4. Long-Lived Particles 45

2.4.1. What is a Long-Lived Particle? 45

2.4.2. The Higgs Boson Coupling to LLPs 46

III. ACCELERATOR AND DETECTOR 49

3.1. Circular Hadron Colliders 49

3.2. The Large Hadron Collider 49

3.2.1. CERN Accelerator Complex 49

3.2.2. LHC Magnets, Acceleration, and Proton Bunches 51

Chapter Page

3.3. The ATLAS Detector . . . 53
3.3.1. Coordinate System . . . 53
3.3.2. The Inner Detector . . . 54
3.3.2.1. The Silicon Pixel Detector . . . 55
3.3.2.2. The Semiconductor Tracker . . . 56
3.3.2.3. The Transition Radiation Tracker . . . 56
3.3.3. Calorimeters and Calorimetry . . . 56
3.3.3.1. The Electromagnetic Calorimeter . . . 57
3.3.3.2. The Hadronic Calorimeter . . . 59
3.3.4. The Muon Spectrometer . . . 60
3.3.5. The Magnet System . . . 62
3.4. The Trigger and Data Acquisition System . . . 63
3.4.1. Level-1 Trigger System . . . 64
3.4.1.1. L1Calo . . . 65
3.4.1.2. L1Muon . . . 67
3.4.1.3. L1Topo . . . 67
3.4.2. High Level Trigger System . . . 67
IV. GFEX ALGORITHM VALIDATION . . . 69
4.1. Motivation for gFEX . . . 69
4.2. gFEX Hardware and Infrastructure . . . 70
4.2.1. ATCA Architecture . . . 70
4.3. gFEX Algorithms . . . 72
4.3.1. gFEX Tower Builder . . . 73
4.3.2. Jet Finder . . . 74
4.3.3. Missing Transverse Energy: Jets without Jets . . . 76

Chapter Page

4.3.4. Special Timing Triggers . . . 78
4.4. gFEX Validation Framework . . . 78
4.4.1. Obtaining gFEX Data for Comparisons . . . 79
4.4.2. The Bitwise C Simulation . . . 80
4.4.3. gFEX Compare . . . 80
4.4.3.1. Parser . . . 80
4.4.3.2. The Comparison Framework, Reports, and Plots . . . 81
4.4.4. The Validation Work in Action: Debugging gFEX Inter-FPGA Communication . . . 82
4.5. gFEX Performance and Usage in Run 3 . . . 84
V. ANALYSIS: DISPLACED JETS IN THE ATLAS INNER DETECTOR . . . 86
5.1. Analysis Motivation and Introduction . . . 86
5.2. Datasets and MC Samples . . . 90
5.2.1. Datasets . . . 90
5.2.1.1. Event cleaning . . . 90
5.2.2. Monte Carlo samples . . . 90
5.2.2.1. Signal samples . . . 90
5.3. Object Reconstruction and Definitions . . . 92
5.3.1. Tracks and Tracking . . . 92
5.3.2. Jet Reconstruction . . . 93
5.3.3. Leptons and Photons . . . 94
5.4. Analysis Triggers . . . 94
5.4.1. Emerging Jets Trigger . . . 95
5.4.2. Displaced Jet Triggers . . . 95

Chapter Page

5.4.3. VBF Trigger . . . 96

5.5. Analysis Strategy . . . 97

5.5.1. The VBF Channel . . . 97

5.5.2. The ggF Channel . . . 98

5.5.3. Boosted vs Resolved Regimes . . . 99

5.6. Determination of the EJ 2024 Scale Factor . . . 100

5.7. Trigger Studies to Asses the Utility of the EJ trigger . . . 104

5.7.1. Signal Efficiency Trigger Studies . . . 105

5.7.1.1. DJ Symmetric and Asymmetric Trigger comparisons 105

5.7.1.2. VBF Preselection Studies . . . 106

5.7.1.3. EJ and DJ Trigger Complimentarity . . . 106

5.7.2. Defining Trigger Strategy for ggF Channel - Trigger Buckets . . . 108

5.8. Background Estimation Attempts and Preliminary Expected Limits . . . 112

5.8.1. ggF Background Estimation and EJ Expected Limits Contribution . . . 113

5.9. Current Status of Analysis and Next Steps . . . 118

VI. UPGRADES TO ATLAS FOR THE HIGH-LUMINOSITY LHC AND THE PHOTON POINTING ALGORITHM . . . 120

6.1. The High-Luminosity LHC . . . 120

6.1.1. Detector and TDAQ changes for HL-LHC . . . 120

6.2. ATLAS Global Trigger Upgrade . . . 122

6.3. Diphoton Trigger Algorithm for Global Trigger . . . 124

6.3.1. Motivation for Creating the Diphoton Trigger . . . 125

6.3.2. Feasibility of Using Diphoton Triggers on Exotic Samples . . 126

6.4. The Baseline Diphoton Algorithm for Global Trigger . . . 127

Chapter Page

6.4.1. The PPZ Algorithm 128

6.4.2. High Level Synthesis Implementation 130

6.5. A Neural Net Approach to Photon Pointing 130

6.5.1. NN Architecture 131

6.5.2. Input Variables and Event Cleaning 132

6.5.3. Signal Background Classifier 134

6.5.4. PPZ Regression Model 136

6.6. Future Work on NN PPZ 140

VII. CONCLUSION 141

APPENDICES

A. AN EXAMPLE GFEX ROB FRAGMENT 143

B. HIGGS TO DISPLACED JET RUN 3 ANALYSIS SAMPLES LIST 147

C. DISPLACED JETS ANALYSIS TRIGGER CHAINS 150

D. INCLUSIVE TRIGGER EFFICIENCY STUDIES 151

E. BACKGROUND CONTAMINATION AND EVENT YIELDS OF THE TRIGGER BUCKET STRATEGY 155

F. SAMPLES FOR DIPHOTON TRIGGER STUDIES 163

G. ASSESSING E_{RATIO} CUTS USING THE SIGNAL BACKGROUND CLASSIFIER MODEL 164

REFERENCES CITED 168

LIST OF FIGURES

Figure Page

1. Standard Model particles graphic from [1]. The three generations of fermions are shown on the left, increasing in generation from left to right. The Spin-1 gauge bosons are shown on the right, and the Higgs boson is shown in the center. . . . 33

2. Potential of the field ϕ [2]. Before spontaneous symmetry breaking, the potential sits on its unstable equilibrium at 0. Small fluctuations make the potential drop into a new local minimum. This breaks the symmetry locally while preserving the overall global symmetry. Since there is no preferred direction to the symmetry breaking, it is called spontaneous. . 38

3. All of the Higgs sector (left) and electroweak sector (right) tree level couplings of the SM. The V bosons are generic massive vector bosons that can refer to W or Z. 41

4. All of the tree level couplings in QCD. The gluon can couple to pairs of quarks as well as itself in a trilinear and a quartic coupling. 43

5. Toy diagram of a Higgs decay to two of the new scalar particles, which in turn decay into SM fermions [3]. 47

6. A schematic of the CERN accelerator complex, showing the full proton injection chain for the LHC [4]. 50

7. Plot of the luminosity-weighted distributions of the pileup for Run 3 data-taking years as well as the total across Run 3. The average value increases in later data-taking years as the Run 3 system became more mature. 53

8. Diagram of a radial slice of the ATLAS Inner Detector system. The distance from the relevant subsystems and layers from the beamline are listed [5]. 55

9. Cutaway diagram of the ATLAS calorimeter system with major subcomponents labeled [5]. 57

Figure Page

10. The cell break down by layer of a central EMB trigger tower. The accordion like structure of the LAr calorimeter can be seen on the right side of the graphic [6]. 58

11. A Legacy trigger tower 0.1×0.1 in $\eta - \phi$ (a) and a schematic of a Super Cell (b) from [6]. 59

12. A module of the Tile Calorimeter, showing the PMTs connected to the scintillator material in it's alternating pattern with the steel [7]. 60

13. Cutaway diagram of the Muon Spectrometer with labeled subcomponents [5]. 62

14. Schematic of the magnet system in ATLAS, with the cylindrical solenoid magnet in the center, surrounding the ID, and the superconducting coils of the toroids arranged around the barrel and arranged in a disk on each endcap [8]. 63

15. TDAQ data flow diagram for Run 3 [5]. 64

16. The L1 Trigger logic diagram for Run 3 [5]. 65

17. From left to right, a production eFEX, jFEX, and gFEX board [9]. 66

18. The ATCA shelf setup used for the gFEX board [10]. 71

19. gFEX FPGA link connections [11]. 72

20. A map of $\eta - \phi$ space covered by the gFEX system. The central regions are made of mostly 0.2×0.2 gTowers corresponding to pFPGAs A and B. The forward regions are covered by pFPGA C, and it should be noted the the gTowers become less granular in the forward region. Overlayed on the map are several hypothetical large-R jet candidates for the trigger algorithm. While a nominal gJet is seeded by a 3×3 gBlock, those centered on the border of a jet engine (such as those on the right side of this image) are only seeded by a 2×3 gBlock. 74

21. An example gJet calculation for a candidate centered at the gTower marked with an X. Below is the clock cycle diagram for the process to aid in visualizing how the gTowers are summed within the firmware. 76

Figure Page

22. Format for Jet TOBs sent to L1Topo from pFPGAs. Each half FPGA sends its TOBs on one Fiber. The ρ density value is calculated per FPGA and so is sent only once. The LSBs are not listed in the table, but are nominally 200 MeV for these TOB words. Taken from gFEX Data Formats Documentation. . 76

23. The TOB formats that the Zynq FPGA sends to L1Topo (left) and the TOB formats that each pFPGS sends to the Zynq (right). The LSBs are not listed in the table, but are nominally 800 MeV for these TOB words. Taken from gFEX Data Formats Documentation. 78

24. Diagram showing the logical flow of the validation framework. The arrows with green text indicate steps that are completed by the validation framework guided by a shell script. . . . 79

25. An example of unpacking a hexadecimal gFEX TOB into its components in binary. The validation framework does this conversion to run a full bitwise comparison. 82

26. Example output of the comparison framework, showing the scatter plot (left) and distribution comparison plot (right) of the left half of pFGPA A's Leading Large-R Jet E_T. This is data dumped from Run 476589 of ATLAS taken in May of 2024. 82

27. The Large-R Jet E_T comparisons for a run where the FPGA-intercommunication was not working. This plot was specially modified to only show the E_T mod 8. 83

28. The back calculated line numbers of data shared between FPGAs B and C. The distribution should be uniform, so the gap was a clear indication of a misalignment in the data sharing. 84

29. Left: Turn on curves for gFEX Large-R jet triggers compared to Legacy Small-R J100 and the L1Topo Simple Cone algorithm. Note that the lowest unprescaled gFEX trigger, gLJ140, boasts a sharper turn on and faster plateau than its legacy equivalent. Right: Turn on curves for gFEX, jFEX, and legacy MET algorithms, with both gFEX and jFEX showing sharper turn ons than the legacy equivalent [9]. 85

30. Branching ratios of a single scalar s to the SM as a function of its mass. Image taken from [12]. 88

Figure Page

31. 95% confidence level upper limits on the Higgs boson branching to a pair of long-lived neutral scalars S as a function of $c\tau$, for a Hidden Sector model where (a) $m_S = 5$ GeV, (b) $m_S \in [15, 20]$ GeV, (c) $m_S \in [30, 40]$ GeV, and (d) $m_S \in [50, 60]$ GeV. Plots taken from [13]. 89

32. The efficiency plots used to determine the EJ trigger SF for 2024, with the efficiency parameterized by the Leading RCTopo Jet p_T (left) and the Leading RCTopo Jet PTF (right). The EJ trigger decision for this plot takes the logical OR of all three possible L1 seeds to the EJ trigger. 102

33. The efficiency plots used to determine the EJ trigger SF for 2024, with the efficiency parameterized by the Leading RCTopo Jet p_T (left) and the Leading RCTopo Jet PTF (right). The EJ trigger decision for this plot uses only the gFEX large-R jet for its L1 seed. 103

34. The efficiency plots used to determine the EJ trigger SF for 2024, with the efficiency parameterized by the Leading RCTopo Jet p_T (left) and the Leading RCTopo Jet PTF (right). The EJ trigger decision for this plot uses only the jFEX small-R jet for its L1 seed. 103

35. The efficiency plots used to determine the EJ trigger SF for 2024, with the efficiency parameterized by the Leading RCTopo Jet p_T (left) and the Leading RCTopo Jet PTF (right). The EJ trigger decision for this plot uses only the L1Topo Simple Cone algorithm for its L1 seed. 104

36. Distributions of the LLP decay length L_{xy} for different selections of the DJ symmetric and asymmetric trigger. The left plot shows the exclusive distributions, where the trigger not selected is required to fail. The right plot shows the inclusive distributions, which look largely identical. This provides justification for moving forward with the asymmetric DJ trigger as the only DJ trigger to consider for the signal region. 106

Figure Page

37. Summary plots of the inclusive trigger efficiency on ggF production mode signal samples for the EJ and DJ triggers, with both preselections applied and left out to gain a complete picture. The EJ clearly shows an increase in performance for the low lifetime samples, suggesting including it in the analysis can be greatly beneficial to gaining additional sensitivity in that region of phase space. 107

38. The significance plots for the DJ trigger bucket in the boosted region for signal samples organized by lifetime, with 1 mm samples on the top left, 10 mm samples on the top right, and 100 mm samples on the bottom. The large-R DJ tagger score extends to a maximum value of 0.99999, which was found to maximize the analysis sensitivity while maintaining high signal acceptance. 109

39. The significance plots for the EJ trigger bucket in the boosted region for signal samples organized by lifetime, with 1 mm samples on the top left, 10 mm samples on the top right, and 100 mm samples on the bottom. The large-R DJ tagger score extends to a maximum value of 0.99999, which was found to maximize the analysis sensitivity while maintaining high signal acceptance. 110

40. The significance plots for the DJ trigger bucket in the resolved region for signal samples organized by lifetime, with 1 mm samples on the top left, 10 mm samples on the top right, and 100 mm samples on the bottom. The small-R DJ tagger score extends to a maximum value of 0.999, which was found to maximize the analysis sensitivity while maintaining high signal acceptance. 111

41. The significance plots for the EJ trigger bucket in the resolved region for signal samples organized by lifetime, with 1 mm samples on the top left, 10 mm samples on the top right, and 100 mm samples on the bottom. The small-R DJ tagger score extends to a maximum value of 0.999, which was found to maximize the analysis sensitivity while maintaining high signal acceptance. 112

42. The generated mistag rate for the specified mistag map. This was generated for the or of DJ and EJ buckets in order to assess if adding EJ made the estimation of DJ close better. 114

Figure Page

43. Closure tests of the mistag map for the DJ bucket (top left), EJ bucket (top right), and the combined DJ or EJ selection (bottom). As can be seen, the closure in the signal region is poor, indicating that this mistag-map approach was not sufficiently accurate for the final background estimation. 115

44. The signal contamination for the mistag map in the DJ Bucket (left) and the or of the DJ and EJ regions (right). This shows very little signal contamination, suggesting that it can be used in the limit framework. 116

45. Preliminary expected limits derived from the mistag map data driven background estimation strategy, using only the DJ Bucket (left) as well as the DJ or EJ buckets together (right). Since this background estimation strategy has not been able to close, a 50% uncertainty has been included on the background estimate. 118

46. The ratio of the two expected limits from the previous Figure 45. The numerator of the plot is the expected limits in the ggF region using both the DJ and EJ triggers to define the mistag map, and the denominator is the expected limits in the ggF region using only the DJ trigger to define the mistag map. While not providing additional sensitivity to the 10 and 100 mm lifetimes, the inclusion of the EJ trigger improves the limits set on the 1 mm lifetime by a full order of magnitude in most cases, up to 2 orders of magnitude for one mass sample (10 GeV). 119

47. Schematic sketch of the ITk layers in z-r space with different values of η marked with lines [14]. 121

48. A schematic showing the full dataflow of the Phase-II TDAQ system, starting with detector input to the trigger system and ending with permanent storage. Of particular interest is the Global Trigger, which will take in inputs from the rest of L0 and the calorimeters to run final hardware selection algorithms before passing trigger objects to the CTP[15]. . . . 122

49. A schematic showing the dataflow path of Global Trigger, where inputs are fed into MUXes which in turn sort the information by BCID to the GEPs where algorithms are run. The resulting decisions are sent to CTP for final decision[16]. . . . 123

Figure Page

50. A schematic showing the time-multiplexing of the MUXes into the GEPs. Events are multiplexed such that one GEP is handling the information for one BCID. The CTP interface will demultiplex them such that they show up in order to the CTP[16]. 124

51. A diagram of the Maximum Possible Value (MPV) latency of the Global system, showing when different inputs will be available to the GEP[17]. 124

52. Signal efficiency turn on curves for the asymmetric diphoton trigger on exotic Higgs samples as a proxy for ALPs decaying to diphoton. 127

53. A sketch of the PPZ algorithm main steps. Image credit Liza Brost. 129

54. Resource estimates for a version of the PPZ algorithm implemented in HLS using Vitis software. 130

55. The NN score output distributions for different E_T bins (left) and the mean NN output as a function of E_T (right) with E_T included as a training variable. A vertical line at 10 GeV shows a theoretical desired cut on e/γ objects in Global. 135

56. ROC curves binned by E_T (left) and the same mean NN score as a function of E_T (right) for the training without E_T included as an input variable. There is still a clear preference for high E_T photons. 136

57. The PPZ value targets (left) and predictions (right) show a good first attempt at regression overall. 137

58. The overall distribution of target and prediction values for the sample. In general, the regression does fairly well, as it is clustered around the slope 1 line of agreement for the large bulk of events. 138

59. The PPZ value targets and predictions 2d map for photons in the barrel region, split into two different η bins to look at highly central structure. These on average perform better than the full distribution. 139

Figure Page

60. The PPZ value targets and predictions 2d map for photons in the endcap region, split into two different η bins to look at more forward structure. These on average perform worse than the full distribution. 139

E.1. DJ Bucket Boosted Region Signal Contamination, ordered in increasing signal sample lifetime. 155

E.2. EJ Bucket Boosted Region Signal Contamination, ordered in increasing signal sample lifetime. 156

E.3. DJ Bucket Resolved Region Signal Contamination, ordered in increasing signal sample lifetime. 157

E.4. EJ Bucket Resolved Region Signal Contamination, ordered in increasing signal sample lifetime. 158

E.5. DJ Bucket Boosted Region Event Yields, ordered in increasing signal sample lifetime. 159

E.6. EJ Bucket Boosted Region Event Yields, ordered in increasing signal sample lifetime. 160

E.7. DJ Bucket Resolved Region Event Yield, ordered in increasing signal sample lifetime. 161

E.8. EJ Bucket Resolved Region Event Yield, ordered in increasing signal sample lifetime. 162

G.1. NN output of the S/B classifier parameterized by E_{ratio} for E_T bins 1 and 2 corresponding to [0,10) GeV and [10,15) GeV. 164

G.2. NN output of the S/B classifier parameterized by E_{ratio} for E_T bins 3 and 4 corresponding to [15,20) GeV and [20,25) GeV. 164

G.3. NN output of the S/B classifier parameterized by E_{ratio} for E_T bins 5 and 6 corresponding to [25,30) GeV and [30,35) GeV. 165

G.4. NN output of the S/B classifier parameterized by E_{ratio} for E_T bins 7 and 8 corresponding to [35,40) GeV and [40,45) GeV. 165

G.5. NN output of the S/B classifier parameterized by E_{ratio} for E_T bins 9 and 10 corresponding to [45,50) GeV and [50,60) GeV. 166

G.6. NN output of the S/B classifier parameterized by E_{ratio} for E_T bins 11 and 12 corresponding to [60,80) GeV and [80,100) GeV. 166

Figure Page

G.7. NN output of the S/B classifier parameterized by E_{ratio} for E_T bin 13 corresponding to [100, ∞) GeV. 167

LIST OF TABLES

Table Page

1. The particles of the SM and their charges under the gauge symmetries. Only the first generation of fermions is shown for simplicity, and have further been split into their left hand (doublet) and right hand (singlet) representations. 34

2. R&D signal samples. The first two lines show the filtered samples that were generated, while the last two lines show the unfiltered samples, which have double the events of the filtered ones to allow for sufficient statistics after the relevant triggers are applied. 91

3. Signal grid extension. 91

4. Cuts made on physics objects used to define control regions. 94

5. Online requirements for the VBF trigger. 97

6. Selections for VBF Trigger offline requirements. 97

7. Selections for DJ Trigger offline requirements. 98

8. Selections for EJ Trigger offline requirements. 99

9. Signal Regions for the analysis as they currently stand. 100

10. CR_{μ} selections. 101

11. RCTopo Jet selections for CR_{μ}. 101

12. Selections for each variable's efficiency plots. 102

13. SFs for EJ Trigger 2024 split by L1 seed. 104

14. SFs for the EJ trigger for the relevant data taking years for the analysis. Years 2022 and 2023 are taken from the on plateau values considered in the Emerging Jets Run 3 Analysis [18]. 104

15. Light Higgs Resonance decays to diphoton samples. These are used as a proxy for an ALP decaying to diphotons in trigger studies as these samples were readily available. 127

Table Page

16. Simple quality cuts made on eFEX TOBs before being used as inputs to NN training. . . . 133

17. Quality cuts made on E_{ratio} before being used as input to NN training. . . . 134

B.1. Datasets used for each data-taking year (Physics Main and VBF Delayed streams). . . . 147

B.2. Datasets, corresponding Good Run Lists, and integrated luminosities. . . . 148

B.3. List of LHE files used for the different Higgs boson production mechanisms. . . . 149

D.1. Inclusive trigger efficiencies for the Emerging Jet and Asymmetric Displaced Jet triggers with and without preselection requirements for the full mass and lifetime signal sample grid. This covers mc campaigns mc23a and mc23d for the ggF production modes. Rates shown are unique, so passing one trigger means failing the other. EJ shows significant gains in the low lifetime samples over DJ, and comparable rates for longer lifetimes. . . . 152

D.2. Inclusive trigger efficiencies for the symmetric and asymmetric DJ triggers for the full mass/lifetime sample grid. This covers mc campaigns mc23a and mc23d for the ggF production modes Exclusive denotes that the other trigger must fail so that we obtain a rate of unique events captured by each trigger. The asymmetric DJ trigger shows better acceptance overall, so it is taken as the main DJ trigger for other studies. . . . 153

D.3. Inclusive trigger rates for the VBF trigger and preselection options for the full mass/lifetime grid. Calculated using mc23a and mc23d for the VBF production mode. This study played a a role in determining the final preselection cuts to be applied: 100 GeV for leading forward jet and 90 GeV for subleading. . . . 154

E.1. Background Monte Carlo samples used in signal studies. . . . 155

F.1. Light Higgs Resonance decays to diphoton samples. . . . 163

CHAPTER I
INTRODUCTION

The Standard Model (SM) of particle physics is our current best understanding of the universe and its underlying forces. It is incredibly successful, with the most famous prediction being fulfilled with the discovery of the Higgs boson by the ATLAS and CMS collaborations at the Large Hadron Collider (LHC) in 2012. Despite its robustness in many measurements, there are still open questions that the SM does not currently answer, such as the existence of dark matter, the matter-antimatter asymmetry of the universe, and the naturalness of the Higgs sector. These all motivate us to look for physics beyond the Standard Model (BSM).

The LHC was made to probe the highest energies (and therefore the smallest distances) possible. Among the experiments at the LHC is the ATLAS detector, which captures nearly the whole solid angle around an LHC collision point to measure all the decay products of a particle physics interaction. Quarks and gluons produced in these collisions cannot exist as free particles due to color confinement (see section 2.2.3), and instead undergo a process called hadronization, producing collimated sprays of particles known as jets. Identifying and correctly reconstructing jets is an important task of the ATLAS detector and plays a critical role in both the trigger system and many physics analyses. The LHC collides proton bunches which are spaced at 25 ns intervals, creating a bunch crossing rate of 40 MHz. While not all of these bunches are filled, the resulting size of output data is far too much for ATLAS to record and save for further analysis. To ensure that ATLAS picks interesting physics events out of this data stream, it has implemented a trigger and data acquisition (TDAQ) system to aid in making decisions about which events to keep. The trigger system, which reduces that initial rate of 40 MHz down to a manageable size for readout and storage, is two-tiered, with the first tier being a hardware based trigger that must make decisions on the order of a few μs, and the second tier being a software based trigger that further reduces the rate by running more stringent algorithms on the order of one second. In this way, determining which algorithms run in the trigger ultimately decide what physics we get to search for.

The LHC organizes its data taking periods into multi-year long "Runs" punctuated by Long Shutdown (LS) periods where maintenance and upgrades are performed. Run 1 began in 2010, and ended in 2013. LS1 followed immediately after Run 1 and lasted until the beginning of Run 2 in 2015. Run 2 lasted until 2018. During LS2 between Runs 2 and 3, upgrades including the new L1 trigger system were installed. These upgrades are referred to as "Phase-I" upgrades in preparation for data collection with higher instantaneous luminosity. Run 3 spanned from 2022 until its recent end in 2026, bringing us to the beginning of LS3 at the time of this dissertation. LS3 will consist of major upgrades to the ATLAS detector and TDAQ system, referred to as "Phase-II" upgrades. LS3 is scheduled

to end in 2030, marking the beginning of the High-Luminosity LHC which will collect an unprecedented amount of data beginning with Run 4.

The Run 3 upgrade of the ATLAS Level-1 calorimeter trigger introduced a new generation of Field Programmable Gate Array (FPGA) based feature extractors (FEXes) that are capable of running programmable algorithms at hardware like speeds. These systems allow for finer selections of physics objects at the hardware trigger level, allowing for decreased thresholds without causing large trigger rate increases. Run 3 of the LHC has just ended in Summer 2026, and as such we are entering the era of the High-Luminosity LHC (HL-LHC), where we will collect even more data due to a higher number of proton interactions occurring in each bunch crossing. The ATLAS TDAQ system will again undergo upgrades. Specifically, ATLAS will install the new Global Trigger, which will be able to exploit full calorimeter granularity at the hardware level for the first time. Its advanced hardware design will also allow for more complex algorithms to run within strict hardware level latency requirements. Alongside advances in detector and trigger technologies, machine learning (ML) has become an increasingly important tool in experimental particle physics analyses. Common ML tools are used for object identification, reconstruction, event classification, and can even be implemented in trigger decisions.

This dissertation brings together these developments through contributions ranging from Run 3 detector operations, Run 3 trigger development and validation, Run 3 searches for new physics, and ML applications for the Run 4 trigger system. The first part of this thesis describes the theoretical motivation for studying physics beyond the standard model, focusing in particular on couplings of the Higgs boson to long-lived particles (LLPs). It then introduces the LHC, the ATLAS detector, and the Run 3 trigger architecture before presenting work on the commissioning and validation of the Global Feature Extractor (gFEX), which comprises the development of a standalone python framework used to validate firmware algorithms and ensure the gFEX system is ready for data taking. The dissertation then presents studies supporting an ongoing ATLAS search for Higgs bosons decaying to a pair of long-lived neutral scalars which decay hadronically, creating a displaced jet signature in the ATLAS inner detector. Specifically, it examines the use of newly available triggers in Run 3, including dedicated displaced jet, emerging jet, and vector boson fusion triggers to improve sensitivity to previously untouched phase space of the Higgs branching ratio to LLPs. The chapter describing this search contains unpublished and co-authored material that will be submitted for publication by the ATLAS collaboration upon its completion. Finally, it presents a project in trigger development for the HL-LHC, presenting studies of machine learning approaches to a proposed diphoton trigger algorithm to be run in the Global Trigger environment and exploring how neural net regression techniques may improve the reconstruction of photon pointing information within stringent hardware trigger latency and resource constraints.

These studies presented illustrate the close relationship between detector instrumentation, trigger algorithm development, and physics analysis. Advances in real-time event selection directly expand the discovery potential of the ATLAS experiment, enabling increasingly sensitive searches for rare physics processes, both in the current LHC programme and in the imminent HL-LHC era.

CHAPTER II
THEORY BACKGROUNDS AND MOTIVATION

This chapter describes the theoretical motivations for the hardware work and physics analyses shown in the following chapters. Section 2.1 describes the foundations of Quantum Field Theory to provide context. Section 2.2 describes the Standard Model of particles physics, going into detail about the particles and interactions that have been predicted by this model and discovered by experiments. Section 2.3 describes motivations for physics searches beyond the Standard Model, highlighting mysteries that the Standard Model is yet to be able to answer. Finally, section 2.4 describes a class of particles called Long-Lived Particles, where they arise in theory, and their detector signatures.

2.1 Theoretical Formalism: Particles, Fields, and Quantum Field Theory

The large majority of this section is a summary of the mathematical basis for the study of particle physics. It is not presented in an extremely rigorous manner, but rather a summary of key points, with enough math to (hopefully) make it coherent. Concepts summarized are taken from a variety of canonical QFT and Particle Physics text books, including but not limited to, *An Introduction to Quantum Field Theory* by Peskin and Schroeder, *Modern Particle Physics* by Thomson, and the (as of yet) unpublished *Fundamentals of QFT* by Cohen and Luty.

2.1.1 Natural Units.

Throughout this dissertation, natural units will be assumed, where the reduced planck constant (lovingly referred to as "h-bar") and the speed of light c, will be set equal to one, $\hbar = c = 1$. In SI units, $\hbar$ has units of Joule-seconds (energy times time), and the speed of light of course has units of meters per second (distance divided by time). This gives us a nice benefit of being able to define a relationship between the units of energy, momentum, and mass. First, we will write the units of mass, energy, and momentum in SI units, and then convert to an $\hbar$, c, and GeV (Gigaelectron Volt) basis:

$$\begin{aligned} [Energy] &= kg * \frac{m^2}{s^2} = GeV \\ [Momentum] &= kg * \frac{m}{s} = \frac{GeV}{c} \\ [Mass] &= kg = \frac{GeV}{c^2} \end{aligned} \tag{2.1}$$

which means that when setting c to 1, we can discuss the mass, momentum, and energy of a particle using only one unit, GeV (or another scale of eV, such as MeV or TeV).

The convention of natural units further implies that we obtain the following relationships between energy, time, and distance:

$$\begin{aligned} [Energy \times Time] = 1 &\implies [Energy] = [Time]^{-1} \\ [\frac{Distance}{Time}] = 1 &\implies [Distance] = [Time] \\ &\implies [Energy] = [Distance]^{-1} \end{aligned} \tag{2.2}$$

As shown above, in natural units, energy and distance have an inverse relationship. This means that in order to probe smaller and smaller distances, we must push to higher and higher energies.

2.1.2 Quantum Field Theory Formalism.

To describe phenomena in particle physics, which involves physical phenomena that occur on very small distance scales at very high energies (and therefore very high velocities), we need a way to combine quantum mechanics with special relativity. This is arguably most successfully accomplished by beginning with relativistic field theory, such as ones that we use to describe gravity and electromagnetism, and then quantizing it by promoting classical fields to quantum operators. For a general field theory, we can consider the action S of the fields ϕ and their derivatives $\partial_\mu \phi$, which is a functional defined as the full 4-dimensional integral of the Lagrangian density.

$$S = \int d^4x \mathcal{L}(\phi, \partial_\mu \phi) \tag{2.3}$$

Extremizing this action,

$$\frac{\partial S}{\partial \phi} = 0 \tag{2.4}$$

results in the Euler-Lagrange equations of motion for the fields ϕ

$$\frac{\partial}{\partial x^\mu}(\frac{\partial \mathcal{L}}{\partial(\partial_\mu \phi)}) = \frac{\partial \mathcal{L}}{\partial \phi} \tag{2.5}$$

From here, we can follow a classical mechanics approach to calculate the conjugate momenta,

$$\pi \equiv \frac{\partial \mathcal{L}}{\partial \dot{\phi}} \tag{2.6}$$

which allows us to find the Hamiltonian density,

$$\mathcal{H} = \pi\dot{\phi} - \mathcal{L} \tag{2.7}$$

From here, we can take the usual quantum mechanics approach and promote π and ϕ to operators, and use them to define commutation relations. We can then use the commutation relations to construct creation and annihilation operators. These creation and annihilation operators can be acted on the ground state in a time ordered way to obtain any valid final state.

We can calculate the two point Green's Function to propagate to a generic point x_1 from another generic point x_2:

$$\Delta_F = <0|T[\hat{\phi}(x_1)\hat{\phi}(x_2)]|0> \tag{2.8}$$

where T denotes the time-ordered product of the ϕ fields. We then apply Wick's theorem in order to equate the time-ordered product to the normal ordered sum of all possible contractions of the fields ϕ. This will be extremely useful to us as we begin to consider the formalism for the scattering of two particles, a highly relevant occurrence in particle physics experiments.

2.1.3 Scattering.

The physical observable of relevance in a scattering experiment is the S-Matrix, which we will now construct. We begin by considering some initial state $|i>$ that consists of n particles that are well separated so that we can consider interactions between them to be negligible. We allow them to collide, measure the final state particles from that collision, and assume that at sufficiently long time the final state particles also do not interact. Then we can describe the probability of the state $|i>$ evolving into a final state $|f>$ via some quantum mechanical operator $\hat{S}$

$$Prob(i \to f) = |\langle f|\hat{S}|i\rangle|^2 \tag{2.9}$$

where $\hat{S}$ is defined as

$$\hat{S} = \lim_{\substack{t_i \to -\infty \\ t_f \to +\infty}} e^{i\hat{H}_0 t_f} e^{-i(\hat{H}_0 + \hat{H}_{int})(t_f - t_i)} e^{-i\hat{H}_0 t_i} \tag{2.10}$$

Here, we have taken advantage of the assumption that the interactions are small enough to be considered a perturbation on the base Hamiltonian. These exponentials are the time evolution operator in the interaction picture, $\hat{U}_{IP}(t_f, t_i)$.

In order to make this more useful, we can consider the time evolution condition of the Schrödinger equation on this operator

$$\frac{\partial}{\partial t}\hat{U}_{IP}(t,t_i) = -i\hat{H}_{IP}(t)\hat{U}_{IP}(t,t_i) \tag{2.11}$$

which, with normalization conditions, has the solution of a time ordered exponential

$$\hat{U}_{IP}(t,t_i) = T\exp\{-i\int_{t_i}^{t} dt'\hat{U}_{IP}(t')\} \tag{2.12}$$

This form finally allows us to write the S-matrix as

$$\hat{S} = T\exp\{-i\int d^4x\mathcal{H}_{int}(\hat{\phi}_{IP}(x),\hat{\pi}_{IP}(x))\} \tag{2.13}$$

where we have translated the Interaction Picture Hamiltonian back into the the Interaction Picture Hamiltonian density. Finally, if we assume that $\mathcal{L}_{int}$ does not involve any time derivatives (common for theories we are primarily interested in), we can write our final version of the S-matrix:

$$\hat{S} = T\exp\{i\int d^4x\mathcal{L}_{int}(\hat{\phi}_{IP}(x))\} \tag{2.14}$$

The important takeaway is that we only need to define the free and perturbative parts of the Lagrangian for a given theory to calculate interactions of the particles. This will be evident in how we describe the Standard Model in the following section.

Now that we have a form for the S-matrix that we can consider in a perturbative way, we can finally compute S-matrix elements for phenomena we care about in particle physics experiments, such as scattering amplitudes, decay rates, and cross sections. Evaluating the S-matrix operator for the interaction, we find that we pick up a geometric factor from the integrals, a delta function enforcing the conservation of 4-momentum, and a Lorentz invariant matrix element we will call $\mathcal{M}$

$$\langle f|\hat{S}|i\rangle = (2\pi)^4\delta(P_i - P_f)i\mathcal{M}_{fi} \tag{2.15}$$

With this form in hand, we can now begin to think about observables we as experimentalists can measure in a laboratory. To do this, we introduce the concept of a *cross-section*, σ. At a conceptual level, the cross-section is the number of interactions per unit time per target particle divided by the incident flux. Considering the case of a collision between two particles (with momenta k_1 and k_2) that also results in two outgoing particles (with momenta k_1' and k_2'), we can describe the differential cross-section in the center of mass frame as

$$d\sigma_{\text{cm}} = \frac{1}{4E_{\text{cm}}|\mathbf{k_1}_{,\text{cm}} + \mathbf{k_2}_{,\text{cm}}|}|\mathcal{M}_{fi}|^2 d\Phi_f \tag{2.16}$$

where $d\Phi_f$ is the Lorentz invariant phase space measure for integrating over all final states,

$$d\Phi_f = (\prod_f (dp_f))(2\pi)^4\delta^4(P_f - P_i) \tag{2.17}$$

Similarly, keeping in mind that the differential cross section of a process can be related to the differential cross-section by the flux of incident particles, we can find a form of the differential decay rate of a particle in its own rest frame as

$$d\Gamma_{\text{rest}} = \frac{1}{2M_i}|\mathcal{M}_{fi}|^2 d\Phi_f \tag{2.18}$$

Integrating over this phase space gives us the total decay rate of a particle in its rest frame,

$$\Gamma_{\text{rest}} = \frac{1}{2M_i}\int |\mathcal{M}_{fi}|^2 d\Phi_f \tag{2.19}$$

We can use this value to calculate the probability of a particle decaying. Cross-sections and decay rates are values that we can measure at particle colliders, giving us a way to test the predictions of the underlying theory, the Standard Model of particle physics.

2.1.4 Gauge Symmetries and Symmetries.

The final piece of formalism we need to start discussing the Standard Model is that of Lie groups and gauge symmetries. Gauge symmetry refers to the invariance of the Lagrangian under local transformations belonging to a Lie group (a group whose elements are labeled by continuous parameters). Elements of this group can be expressed as finite dimensional matrices, allowing us to interpret

them as quantum mechanical operators acting on the space. We can consider transformations close to the identity and then expand in powers of the parameters θ,

$$g(\theta) = \mathbb{1} - i\sum_{a=1}^{n} \theta_a T_a + \mathcal{O}(\theta^2) \tag{2.20}$$

where T_a are the generators of the group.

The linear combinations of the generator matrices form the elements of the Lie algebra associate with the group. Group closure asserts that the commutation of two generators must also be in the Lie algebra, giving us the following commutation relations

$$[T_a, T_b] = i\sum_{c=1}^{n} C_{abc} T_c \tag{2.21}$$

where C_{abc} are the structure constants. In the event that this structure constant is 0 (i.e., the commutator vanishes), the group is Abelian. The gauge group for Abelian theories is the Unitary Group, $U(1)$. This is relevant for us in that it describes the force of electromagnetism in the Standard Model.

Nonzero structure constants therefore imply that the group is non-Abelian. Of particular interest to the standard model of particle physics is a subset of the Lie algebras that have matrix determinant 1. These are referred to as Special Unitary groups, denoted $SU(N)$. The generators of $SU(N)$ are the set of $N \times N$ unitary matrices with determinant 1. The fundamental representation of $SU(N)$ is an N-dimensional complex vector. $SU(2)$ and $SU(3)$ appear in the Standard Model in terms of the electroweak and strong sectors, respectively. Now that we have built up all the formalism that we need, we can put it into action to describe our best theory of how matter behaves on a fundamental level: The Standard Model of particle physics.

2.2 The Standard Model of Particle Physics

Our best understanding of matter and its interactions at the most fundamental level is described by the Standard Model (SM) of particle physics. The SM is a Quantum Field Theory, specifically, and Effective Field Theory (EFT), that describes 3 of the 4 fundamental forces in the Universe, namely the electromagnetic, the weak nuclear, and the strong nuclear forces. The fourth fundamental force, gravity, is well described by General Relativity, but is not incorporated into the Standard Model in a complete way. For the purposes of our collider experiments, the force of gravity on and between particles is negligible. Gravity would become non-negligible at energies ordering the Planck Scale (10^{19}

GeV). The following section contains a description of the fundamental particles of the SM.

The Standard Model is really a model of symmetries. It consists of fermionic fields of matter that are invariant under a $SU(3)_{\mathrm{C}} \times SU(2)_{\mathrm{L}} \times U(1)_{\mathrm{Y}}$ gauge symmetry corresponding to the bosonic fields that mediate interactions. There are 12 fermions (Spin-$\frac{1}{2}$) split into 3 generations, four types of vector gauge bosons (Spin-1), and one scalar boson (Spin-0). A visual representation of the SM particles can be found in Figure 1. The fermions are classified as either leptons or quarks. The $SU(2)_{\mathrm{L}} \times U(1)_{\mathrm{Y}}$ describes the electroweak symmetry before spontaneous symmetry breaking from the Higgs field, which will be described in the following section. After spontaneous symmetry breaking, all of the fermions interact with the massive gauge bosons (the Z and W bosons), and the scalar boson (the Higgs) if they gain their mass through interactions with the Higgs. Charged fermions (all fermions except the neutrinos) interact with the massless gauge boson of the resulting $U(1)_{\mathrm{EM}}$ symmetry (the photon) which governs electromagnetic interactions. Additionally, the quarks have a color charge (C) that implies interactions with the gauge bosons of the $SU(3)$ group (the gluons). The field content of the SM underlying gauge groups can be found in Table 1.

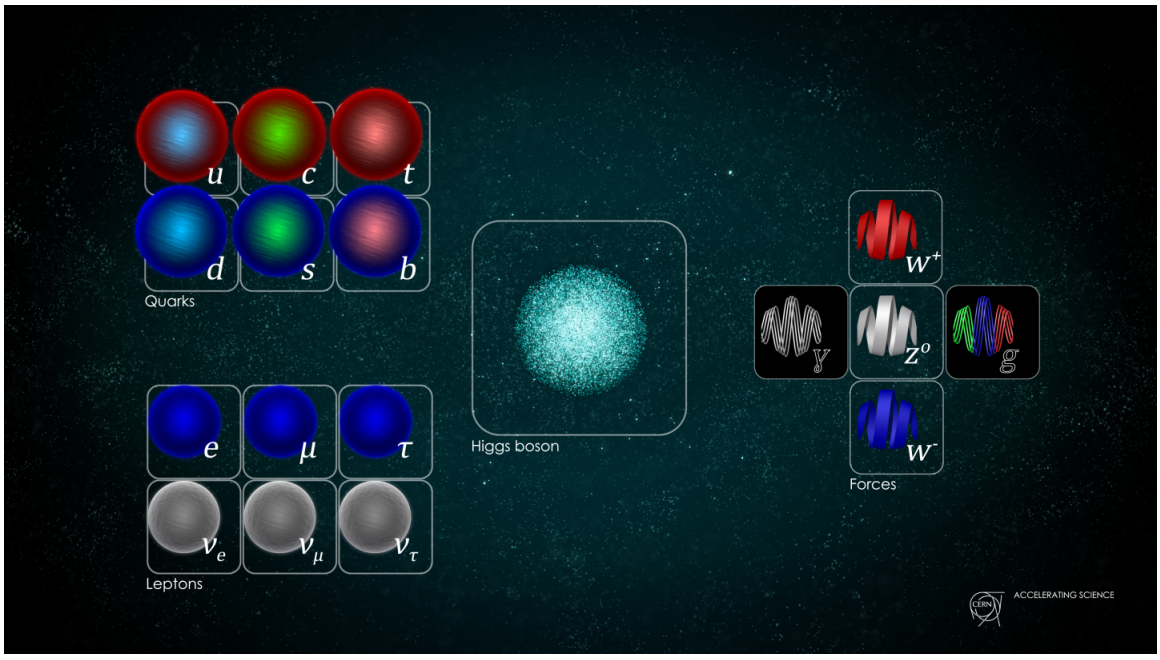


Figure 1. Standard Model particles graphic from [1]. The three generations of fermions are shown on the left, increasing in generation from left to right. The Spin-1 gauge bosons are shown on the right, and the Higgs boson is shown in the center.

Field	Spin	SU(3)	SU(2)	U(1)
B_μ	1	1	1	0
W_μ^a	1	1	3	0
G_μ^a	1	8	1	0
$Q_L = \begin{pmatrix} u \\ d \end{pmatrix}$	$\frac{1}{2}$	3	2	$+\frac{1}{6}$
u_R	$\frac{1}{2}$	$\bar{3}$	1	$\frac{2}{3}$
d_R	$\frac{1}{2}$	$\bar{3}$	1	$-\frac{1}{3}$
$L_L = \begin{pmatrix} \nu \\ e \end{pmatrix}$	$\frac{1}{2}$	1	2	$-\frac{1}{2}$
e_L	$\frac{1}{2}$	1	1	-1
ϕ	0	1	2	1

Table 1. The particles of the SM and their charges under the gauge symmetries. Only the first generation of fermions is shown for simplicity, and have further been split into their left hand (doublet) and right hand (singlet) representations.

The SM Lagrangian can be written out in full detail spanning several lines (or pages), but can be more easily simplified into major sections:

$$\mathcal{L}_{\text{SM}} = \mathcal{L}_{\text{EW}} + \mathcal{L}_{\text{Higgs}} + \mathcal{L}_{\text{QCD}} \tag{2.22}$$

We will discuss these sections in order below.

2.2.1 The Electroweak Sector.

We will begin this section by considering a simple U(1) symmetry so that we can understand the structure of fermions in a basic example. A fermion ψ which transforms under a $U(1)$ symmetry with the fundamental representation

$$\psi(x) \to e^{i\alpha(x)}\psi(x) \tag{2.23}$$

To make sure this field is invariant under the transformation, we introduce a gauge field $A_\mu(x)$ that transforms as:

$$A_\mu(x) \to A_\mu(x) - \frac{1}{g}\partial_\mu\alpha(x) \tag{2.24}$$

We can then define the covariant derivative on ψ as

$$D_\mu \psi \equiv (\partial_\mu + igA_\mu(x))\psi \tag{2.25}$$

and the field strength tensor as

$$F_{\mu\nu} = \partial_\mu A_\nu - \partial_\nu A_\mu \tag{2.26}$$

This is recognizable as what we expect from classical electromagnetism when viewed through classical field theory. When considering groups that are non-Abelian, we must introduce the generators of the group to the covariant derivative, giving us the more general covariant derivative and field strength tensor

$$D_\mu \psi \equiv (\partial_\mu + igt^a A^a_\mu(x))\psi \tag{2.27}$$

$$F^a_{\mu\nu} = \partial_\mu A^a_\nu - \partial_\nu A^a_\mu - g f^{abc} A^b_\mu A^c_\nu \tag{2.28}$$

where t^a and f^{abc} are a generator and the (now non-zero) structure constant, respectively.

With this in mind, we now pivot back to the full electroweak sector. The electroweak sector is comprised of the $SU(2)_{\mathrm{L}} \times U(1)_{\mathrm{Y}}$ symmetry of the SM. At high energies (such as that of the early universe), this symmetry is unbroken. This results in 4 massless gauge fields, called W^1, W^2, W^3 bosons corresponding to $SU(2)_{\mathrm{L}}$ and the B boson corresponding to the $U(1)_{\mathrm{Y}}$. We can write the field strength tensors of these bosons as

$$W_{\mu\nu} = \partial_\mu W_\nu - \partial_\nu W_\mu - g W_\mu \times W_\nu \tag{2.29}$$

$$B_{\mu\nu} = \partial_\mu B_\nu - \partial_\nu B_\mu \tag{2.30}$$

where g is the $SU(2)_L$ coupling constant. With these tensors defined, we now can write down the gauge portion of the electroweak Lagrangian:

$$\mathcal{L}_{\text{EW}} \supset -\frac{1}{4}B_{\mu\nu}B^{\mu\nu} - \frac{1}{4}W_{\mu\nu}W^{\mu\nu} \tag{2.31}$$

Now that we have described the bosonic portion of the EW sector, we need to still describe the fermionic part. To do this, we begin with the Dirac equation, which describes the quantum mechanical motion of relativistic fermions:

$$(i\gamma^\mu\partial_\mu - m)\psi = 0 \tag{2.32}$$

where γ^μ are the Dirac gamma matrices. Using the conjugate fermion field $\bar{\psi}$, we can write the kinematic term of the Lagrangian as

$$\mathcal{L} = \bar{\psi}(i\gamma^\mu D_\mu - m)\psi \tag{2.33}$$

where we have also defined a covariant derivative in a similar way to our $U(1)$ example. Specifically, to remain invariant under $SU(2)_{\text{L}} \times U(1)_{\text{Y}}$, we get the resulting covariant derivative for the fermions:

$$D_\mu = \partial_\mu - ig'\frac{Y}{2}B_\mu - i\frac{g}{2}(\sigma \cdot W_\mu) \tag{2.34}$$

where g' is the $U(1)_{\text{Y}}$ coupling constant, Y is the $U(1)_{\text{Y}}$ hypercharge, and σ are the two-by-two Pauli matrices.

Now we can write the full EW Lagrangian as

$$\mathcal{L}_{\text{EW}} = \bar{\psi}(i\gamma^\mu D_\mu)\psi - \frac{1}{4}B_{\mu\nu}B^{\mu\nu} - \frac{1}{4}W_{\mu\nu}W^{\mu\nu} \tag{2.35}$$

This fully describes the EW sector portion of the SM. Notably, masses for the fermions are missing. We know experimentally that fermions have mass, so there must be a mechanism which allows for mass terms. This is known as the Higgs mechanism.

2.2.2 Electroweak Symmetry Breaking and the Higgs Mechanism.

The previous section described the electroweak sector as one unified symmetry, which holds at high energies. At low energies (such as that of the universe today), a scalar field (the Higgs) spontaneously breaks this symmetry into an untouched $U(1)_{\text{EM}}$ and the weak sector, whose bosons have now gained mass.

Mathematically, this occurs through the introduction of a complex doublet field called ϕ with hypercharge $Y = 1$. We can represent this field in the following way such that the two component fields are hypercharge $Y = \frac{1}{2}$:

$$\phi(x) = \begin{pmatrix} \phi^+(x) \\ \phi^0(x) \end{pmatrix} \tag{2.36}$$

This allows us to write the following Lagrangian

$$\mathcal{L}_{\text{Higgs,bosons}} = (D^\mu \phi)^\dagger (D_\mu \phi) - V(\phi) \tag{2.37}$$

which has the following potential term:

$$V(\phi) = \frac{1}{2}\mu^2 \phi^\dagger \phi + \frac{1}{4}\lambda(\phi^\dagger \phi)^2 \tag{2.38}$$

A visualization of this potential can be seen in Figure 2.

The covariant derivative is taken from the form we derived for the unbroken $SU(2)_{\text{L}} \times U(1)_{\text{Y}}$:

$$D_\mu = \partial_\mu - i\frac{g'}{2}B_\mu - i\frac{g}{2}(\sigma^a \cdot W_\mu^a) \tag{2.39}$$

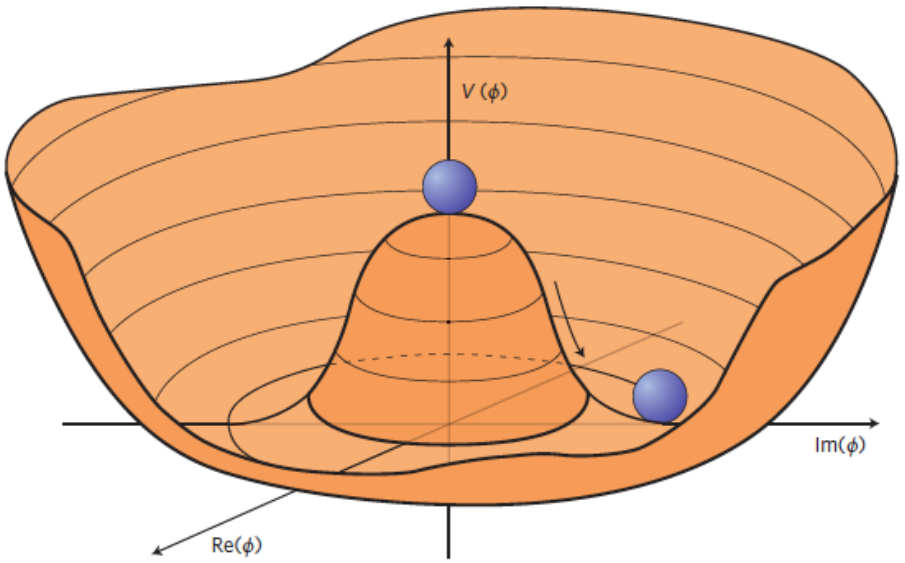


Figure 2. Potential of the field ϕ [2]. Before spontaneous symmetry breaking, the potential sits on its unstable equilibrium at 0. Small fluctuations make the potential drop into a new local minimum. This breaks the symmetry locally while preserving the overall global symmetry. Since there is no preferred direction to the symmetry breaking, it is called spontaneous.

The shape of the Higgs potential is clearly determined by μ and λ. If $-\mu^2 < 0$ and $\lambda > 0$, then the minimum potential is unstable. We can make a global transformation of this field that brings the vacuum expectation value (vev) into only the second component of ϕ, resulting in the form

$$\langle 0|\phi|0\rangle = \frac{1}{\sqrt{2}}\begin{pmatrix}0\\v\end{pmatrix} \tag{2.40}$$

where the vev v is:

$$v = \sqrt{\frac{-\mu^2}{\lambda}} \tag{2.41}$$

So we can express ϕ as small deviations about this vev as so:

$$\phi = \frac{1}{\sqrt{2}}\begin{pmatrix}0\\v+H\end{pmatrix} \tag{2.42}$$

where H is a real scalar mass eigenstate. From this, the potential now takes the form:

$$V(H) = \frac{1}{2}\mu^2 H + \lambda_3 \nu H^3 + \frac{1}{4}\lambda H^4 \tag{2.43}$$

giving us the kinetic term as well as the trilinear and quartic Higgs self-couplings. Plugging this into the kinetic term in the Lagrangian and turning the mathematical crank, we will see the SM gauge bosons appear. Namely, the W, Z, and A (photon) bosons are the following linear combinations of the W_μ and B_μ bosons of the unbroken symmetry:

$$W_\mu^\pm = W_\mu^1 \mp i W_\mu^2 \tag{2.44}$$

$$Z_\mu = \frac{1}{\sqrt{g^2 + g'^2}}(g W_\mu^3 - g' B_\mu) = \cos\theta_W W_\mu^3 - \sin\theta_W B_\mu \tag{2.45}$$

$$A_\mu = \frac{1}{\sqrt{g^2 + g'^2}}(g' W_\mu^3 - g B_\mu) = \cos\theta_W B_\mu + \sin\theta_W W_\mu^3 \tag{2.46}$$

where we have defined the weak mixing angle θ_W via

$$\cos\theta_W = \frac{g}{\sqrt{g^2 + g'^2}}, \sin\theta_W = \frac{g'}{\sqrt{g^2 + g'^2}} \tag{2.47}$$

From this we also find the masses acquired by the $W^\pm$ and Z bosons:

$$m_W = \frac{gv}{2}, m_Z = \frac{v}{2}\sqrt{g'^2 + g^2} \tag{2.48}$$

and note that the photon lacks a coupling to the physical Higgs field, leaving it to remain massless. We can also note for completeness that the original 12 degrees of freedom from the unbroken symmetry and complex Higgs field (two polarizations for each massless $SU(2)_\mathrm{L} \times U(1)_\mathrm{Y}$ boson and 4 degrees of freedom for the complex Higgs doublet) remain in the new formulation, where three of the Higgs field's degrees of freedom are "eaten" by the $W^\pm$ and Z bosons as they gain mass (and

therefore a longitudinal polarization). With these interactions, we also get the Higgs self interactions, which occur through trilinear and quartic couplings, which results in a mass term for the Higgs of

$$m_f = \sqrt{2\lambda}v \tag{2.49}$$

This Higgs field is now free to interact with the fermion fields, generating the anticipated mass terms via Yukawa interactions, seen in the Lagrangian in the form:

$$\mathcal{L}_{\text{Higgs,fermions}} = -\lambda_f(\bar{f}_L\phi f_R + \bar{f}_R\phi^\dagger f_L) \tag{2.50}$$

where we represent the fermions in their left hand doublet and right hand singlet representation when interacting with the Higgs scalar doublet, each scaled by the appropriate Yukawa coupling strength λ_f. Plugging in our Higgs doublet we can see that the mass of a fermion is given by

$$m_f = \frac{\lambda_f v}{\sqrt{2}} \tag{2.51}$$

With the Higgs fully incorporated into the SM framework, we can also now revisit the electroweak sector, noting that we can rewrite the interacting part of the Lagrangian in terms of the $W^\pm$ and Z bosons, giving interactions between the SM fermions and these bosons, as well as trilinear and quartic couplings between the massive bosons themselves. Of particular interest is the coupling of $W^\pm$ to fermions, as it is the only interaction in the SM that allows the coupling of up-like to down-like quarks and the coupling of charged leptons to neutrinos. For the quark couplings specifically, the weak eigenstates are not necessarily the mass eigenstates, meaning that the interactions come with a possibility of changing quark generation. The Cabibbo-Kobayashi-Maskawa (CKM) matrix describes the probability of a particular quark pair interaction:

$$\begin{pmatrix} V_{ud} & V_{us} & V_{ub} \\ V_{cd} & V_{cs} & V_{cb} \\ V_{td} & V_{ts} & V_{tb} \end{pmatrix} \tag{2.52}$$

This matrix can be described by three rotation angles and a complex phase. These can be measured experimentally [19], and the numeric approximation of the matrix is

$$\begin{pmatrix} V_{ud} & V_{us} & V_{ub} \\ V_{cd} & V_{cs} & V_{cb} \\ V_{td} & V_{ts} & V_{tb} \end{pmatrix} \approx \begin{pmatrix} 0.974 & 0.225 & 0.004 \\ 0.225 & 0.973 & 0.041 \\ 0.009 & 0.040 & 0.999 \end{pmatrix} \tag{2.53}$$

where we notice that the on diagonal (i.e., the quarks do not change generation) are highly preferred, but changing generations favored enough to be fairly common in our context.

Our picture of the electroweak and Higgs sectors are now complete, as we have described the kinematics and interaction terms of broken electroweak symmetry by introducing the Higgs. The full collection of first order vertices allowed by all these interactions can be found in Figure 3. Now we can move on to the final piece of the standard model, the interactions governed by the remaining $SU(3)$ symmetry.

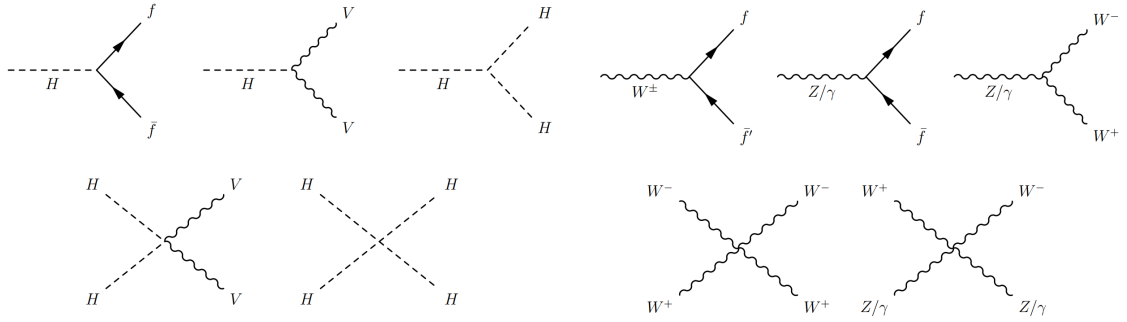


Figure 3. All of the Higgs sector (left) and electroweak sector (right) tree level couplings of the SM. The V bosons are generic massive vector bosons that can refer to W or Z.

2.2.3 Quantum Chromodynamics.

Quarks are additionally charged under the $SU(3)$ symmetry, which we colloquially refer to as color, coining the term "Quantum Chromodynamics." It bears a similarity to QED in that quarks interact through a massless force carrier, in this case, the gluon. As we mentioned in the formalism section, $SU(3)$ has 8 generators (which correspond to gauge bosons). These are the 8 gluons of QCD. The associated charge of the symmetry, called color, can take one three possible values, usually called red (r), green (g), and blue (b), as well as respective anti-color charges for anti-quarks ($\bar{r}, \bar{g}, \bar{b}$). The gluons themselves carry both a color and anti-color charge, resulting in the following possible combinations: $r\bar{g}, r\bar{b}, g\bar{r}, g\bar{b}, b\bar{r}, b\bar{g}, r\bar{r}, g\bar{g}, b\bar{b}$. Linear combinations of the first six options form 6 of the 8 gluons. The last three options are color neutral and are functionally indistinguishable, resulting in one color singlet state

$$\frac{1}{\sqrt{3}}(r\bar{r} + g\bar{g} + b\bar{b}) \tag{2.54}$$

which does not correspond to a physical gluon. Linear combinations of the color neutral options result in the remaining 2 SM gluons, listing all 8 together:

$$\begin{aligned}
&\frac{1}{\sqrt{2}}(r\bar{g} + g\bar{r}), \frac{-i}{\sqrt{2}}(r\bar{g} - g\bar{r}) \\
&\frac{1}{\sqrt{2}}(r\bar{b} + b\bar{r}), \frac{-i}{\sqrt{2}}(r\bar{b} - b\bar{r}) \\
&\frac{1}{\sqrt{2}}(g\bar{b} + b\bar{g}), \frac{-i}{\sqrt{2}}(g\bar{b} - b\bar{g}) \\
&\frac{1}{\sqrt{2}}r\bar{r} - g\bar{g}), \frac{1}{\sqrt{6}}(r\bar{r} + g\bar{g} - 2b\bar{b})
\end{aligned} \tag{2.55}$$

Similarly to our previous exercise in the electroweak sector, we can write the field strength tensor for QCD

$$G^a_{\mu\nu} = \partial_\mu G^a_\nu - \partial_\nu G^a_\mu - g_s f^{abc} G^b_\mu G^c_\nu \tag{2.56}$$

where the index a is running over all the gluons, g_s is the QCD coupling strength, and f^{abc} is the structure constant, and then construct the Lagrangian for both the gauge field and the fermions

$$\mathcal{L}_{\text{QCD}} = -\frac{1}{4} G^a_{\mu\nu} G^{a\mu\nu} + \bar{\psi} i\gamma^\mu D_\mu \psi \tag{2.57}$$

with the covariant derivative

$$D_\mu = \partial_\mu + i g_s \frac{\lambda_a}{2} G^a_\mu \tag{2.58}$$

where λ_a are the Gell-Mann matrices, which are related to the generators of the $SU(3)$ group. Expanding out this Lagrangian describes all the interactions of QCD, including the gluon self couplings. The diagrams depicting these interactions can be found in Figure 4.

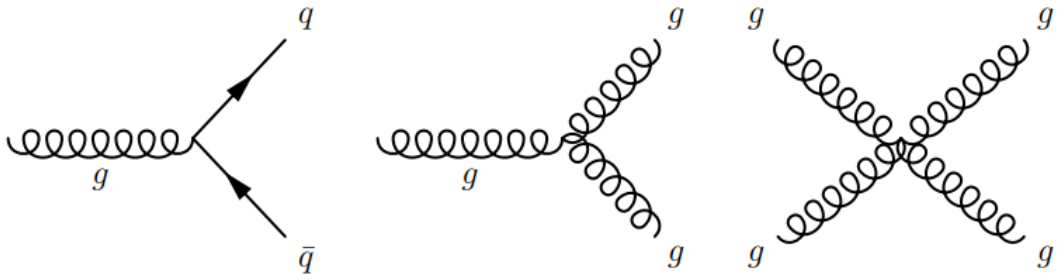


Figure 4. All of the tree level couplings in QCD. The gluon can couple to pairs of quarks as well as itself in a trilinear and a quartic coupling.

Now, we consider that the coupling "constant" g_s, or a more commonly used factor, $\alpha_s = \frac{g_s^2}{4\pi}$, is dependent on the energy scale. Differing from QED, the coupling actually increases at low energies. A direct consequence of the structure of this constant is that of color confinement, where only color neutral particles, called hadrons, are observed in nature. There are two ways quarks bind themselves into hadrons in the SM. A collection of three (anti-)quarks that all have different (anti-)color charges form a three quark state called an (anti-)baryon. A quark anti-quark pair with complementary color charges can also form a two quark state called a meson. As the distance between two color charged particles increases, the QCD force between them increases in an attempt to keep them together. Once this energy passes the threshold for quark-antiquark pair production, it becomes energetically favorable for new quark and gluon pairs to be created. These particles couple to each other and the already existing quarks to form hadrons, in a process known as hadronization. This creates a cascading effect, forming a jet of particles that leave a signature in collider experiments such as ATLAS. One final consequence of color confinement is that of the structure of a proton (the lightest SM baryon). It is commonly understood that a proton is "made" of three quarks, two up quarks and one down quark, referred to as the valence quarks of the proton. However, this is a simplification. The proton is really a cloud of quarks, anti-quarks, and gluons, constantly interacting and changing (also called sea quarks). The consequence of this is that when colliding protons, the probability of a given hard scatter process is described by a set of parton distribution functions (PDFs) that describe the probability of interacting with a specific particle within the proton. This depends on energy scale, and for example, at energy scales like those of the LHC, the gluons carry the largest portion of the proton's momentum. With this we have described the basics of the Standard Model of particle physics. It feels extraordinarily complete, but there are several things we observe in the Universe that the SM cannot explain.

2.3 Beyond Standard Model Physics Motivations

While incredibly successful in describing a large number of observations, there are still a few phenomena that the SM cannot yet explain. Just to name a few: we know that neutrinos have mass, but how they gain this mass is still unknown. We do not have a theory of quantum gravity which fits into our understanding of the SM. We know that dark matter and dark energy exist, but do not know their nature or how they connect to the SM. We cannot explain why the observed Higgs mass, a free parameter in the SM, is so low when quantum corrections to the Higgs mass should cause it to increase. We have an asymmetry between matter and antimatter, called baryon asymmetry, which is also unexplained by the SM. It is clear that our work to describe the Universe is not yet done. This section will highlight DM specifically and justification for searching for physics that exists beyond the Standard Model (BSM) through the Higgs boson.

2.3.1 Motivations for BSM Physics.

Perhaps the most compelling piece of evidence for BSM physics is the existence of dark matter (DM). A wide variety of astrophysical and cosmological evidence for DM has been observed, with the most famous being the observation of galaxy rotation curves, finding that the speed at which galaxies rotate is faster than what one could expect from the amount of visible matter in the galaxy [20]. Additionally, anisotropies in the cosmic microwave background show us the relative abundance of DM to normal baryonic matter (about 5 times more DM) [21]. The SM matter accounts for only about 5% of the energy content of the universe. Dark matter takes up about 27%. The remaining 68% is attributed to dark energy, which drives the accelerating expansion of the universe. This motivates us to search for new particles that can account for these missing pieces.

Another motivation for BSM physics of particular relevance to this thesis is the hierarchy problem. The hierarchy problem describes the apparent lack of "naturalness" in the measured Higgs mass compared to the scale of new physics. The Higgs mass should be sensitive to the largest energy scale that the Higgs couples with, which in theory includes a hypothetical coupling to quantum gravity at the Planck scale. This is clearly much different than the Higgs mass we measure at LHC experiments ([22, 23]), which is on the order of 100 GeV. Quantum corrections to the Higgs mass from particles at this energy scale would push the Higgs mass to much larger orders of magnitude, requiring a fine-tuning of the tree level mass of the Higgs in order to achieve the observed value. One way to avoid this fine tuning is to introduce the concept of naturalness [24], which states that a small parameter (in our case, the Higgs mass compared to the Planck scale), is allowed to be small if it stems from a deeper symmetry. This motivates many neutral naturalness models [25], some of which will be discussed in Section 5.1. The mystery of the Higgs mass scale provides a strong theoretical motivation for us to search for BSM couplings to the Higgs boson.

2.3.2 The Higgs as a Portal to New Physics.

The Higgs boson is a compelling candidate for BSM physics couplings for a few reasons. For example, new particles can be massive, and therefore they can gain their mass through interactions with the Higgs boson. Since all SM particles except perhaps neutrinos gain their mass through the Higgs mechanism, it is reasonable to search for such couplings involving the Higgs to BSM particles in general. Additionally, there is a purely experimental reason why we should search for new couplings to the Higgs boson. The dominant decays of the Higgs boson to SM particles have been measured, giving us a reasonable understanding the relative likelihoods of different Higgs decays [19]. However, if the Higgs were to couple to BSM particles, these particles would contribute to the Higgs width, proportionately decreasing the expected number of observed decays for all Higgs to SM processes, motivating precision measurements of these couplings. Additionally, there are decay modes of the Higgs that cannot be detected, either because they are invisible to the detector, like the decay of Higgs bosons to an all neutrino final state ($\sim 0.1\%$ of Higgs decays in the SM), or they are hard to find due to high background, such as the Higgs decaying to gluons or light quarks. When treated as a free parameter, the upper bound on the Higgs branching ratio to such undetected decays is $\lesssim 12\%$ [26]. This leaves ample room to search for new physics in Higgs boson couplings.

2.4 Long-Lived Particles

The search for new physics can be highly targeted, such as looking for specific DM candidates of a specific theory. However, it can also be fairly general, using the experimental tools available to us to search for any signatures we are not expecting or are not consistent with the SM. This section will describe a class of particles referred to as "Long-Lived", and why we should consider looking for them broadly in the search for BSM physics.

2.4.1 What is a Long-Lived Particle?.

In its own rest frame, the probability of a particle decaying at a time t is given by a simple exponential function

$$P(t) = e^{-t/\tau} \tag{2.59}$$

where τ is the mean proper lifetime of the particle. Standard Model particles have proper lifetimes spanning many orders of magnitude. In the context of detecting particles in a collider, they can be categorized into three broad categories: detector prompt, where the particles decay on a timescale short enough to be indistinguishable from the primary vertex, detector long-lived, where the particles decay within the volume of the detector but beyond the immediate interaction point, and detector stable, where the particle travels through the entirety of the detector without decay. The second and third scenarios result in the class of particles referred to as Long-Lived Particles (LLPs). There are several reasons why a particle may be long-lived. A particle may decay via a heavy particle. The typical

SM example of this is the muon, which decays primarily to an electron and two neutrinos via the much heavier W-boson. A particle may also be long-lived if it has a limited kinematic phase space to decay into. For example, the Kaons (both short and long), decay into pions. The difference of the mass of the Kaons and their respective decay final states (two pions for K short and three pions for K long) is 220 MeV and 80 MeV respectively. This limited amount of phase space causes the Kaons to have a longer lifetime. Finally, a particle can have a long lifetime due to small couplings with their decay products. The famous SM examples of this scenario are the B mesons. A b quark within a B meson decays via a W boson and a c quark. However, this coupling requires a change in quark generation, which corresponds to a weak coupling in the off-diagonal part of the the CKM quark mixing matrix shown in Section 2.2.2. It is evident that known SM particles span the range of lifetimes considered to be long-lived, so it is also natural to assume that BSM particles could have lifetimes in this regime as well.

2.4.2 The Higgs Boson Coupling to LLPs.

The couplings of the Higgs boson to LLPs is well motivated by both top-down (stemming from a fully developed fundamental theory) as well as bottom-up (driven by experimental observations to consider available phase space). They are called for in top-down models such as Neutral Naturalness models [25] which seek to explain the little hierarchy problem (why have we not found any new physics at the LHC scale if the Higg's mass is 125 GeV?) by introducing colorless partners to the top quark. This results in the Higgs becoming a portal between the SM and the new sector, ranging over a variety of lifetimes. Additionally, for a bottom-up approach, one of the simplest possible models to extend the SM is to add a singlet scalar S. This scalar can mix with the standard model Higgs, resulting in the following Lagrangian interaction term [3]

$$\mathcal{L} \supset sin\theta \frac{m_f}{v_h} s f \bar{f} \tag{2.60}$$

where $sin\theta$ is the mixing angle between the scalar and the Higgs. If the scalar is less than half the mass of the Higgs [27], it can be pair produced from Higgs decays, resulting in processes like the figure below.

In a similar vein, we can also imagine that the portal is a pseudo scalar particle, a, which could arise for example as a pseudo Goldstone Boson in theories that have a spontaneously broken global symmetry. This particle is often referred to as an Axion-like particle (ALP), due to the fact that they are similar to the pseudo scalar particle the axion, which was proposed by Weinberg as a consequence of the Peccei-Quinn U(1) symmetry to solve the strong CP problem [28]. Unlike the axion, ALP masses and couplings are not tightly theoretically constrained, so a wide phase space is available for searches over several energy scales. They are popular BSM objects to search for because depending on the region of phase space, they can be dark matter mediators, or even dark matter candidates themselves.

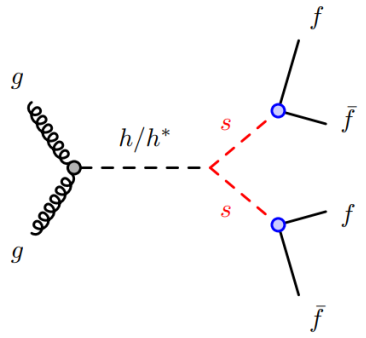


Figure 5. Toy diagram of a Higgs decay to two of the new scalar particles, which in turn decay into SM fermions [3].

These particles would naturally interact with SM fermions via the Lagrangian term [3]

$$\mathcal{L} \supset \frac{\partial_\mu a}{f} \sum_i \frac{c_i}{2} \bar{\psi}_i \gamma_\mu \gamma_5 \psi \tag{2.61}$$

Going to higher mass dimension for more exotic interactions, we can find that the Higgs will decay to two a's, giving a similar experimental signature to the scalar case. A popular benchmark for experimental LLP searches that will be discussed later in detail is where the LLP scalar/pseudoscalar decays hadronically to the heaviest available SM quark pair (often $b\bar{b}$), creating jet signatures which are displaced from the primary interaction vertex. The status for an ATLAS search for such signatures will be discussed in Chapter 5.

ALPs do not only have to decay hadronically. In fact, a very common decay channel for ALPs, particularly lighter ALPs, is to two photons. We can see how the ALP would couple to photons through the following effective Lagrangian term after electroweak symmetry breaking [29]

$$\mathcal{L}_{eff} \supset e^2 C_{\gamma\gamma} \frac{a}{\Lambda} F_{\mu\nu} \tilde{F}_{\mu\nu} \tag{2.62}$$

where $C_{\gamma\gamma}$ is a Wilson coefficient defined as $C_{\gamma\gamma} = C_{WW} + C_{BB}$, the Wilson coefficients for the ALP coupling to the unbroken SU(2) symmetry and the unbroken SU(1) symmetry, respectively. One can choose a benchmark model where $C_{WW} = 0$, which means that the ALP mainly couples to electroweak gauge bosons, and thus decays primarily to photons (see section 6.3.1).

We have shown throughout this section that studying exotic decays of the Higgs boson is a promising avenue toward new physics discovery. In the following section, we will take a look at the ATLAS experiment at the LHC.

CHAPTER III
ACCELERATOR AND DETECTOR

This chapter describes the experimental setup of the ATLAS detector at the Large Hadron Collider in Geneva, Switzerland. Section 3.1 provides some context for circular hadron colliders in general. Section 3.2 describes the LHC accelerator complex, including all the stages required to accelerate particles to their collision energies. Section 3.3 describes the ATLAS detector and its subsystems. The configurations described are for Run 3 of the Large Hadron Collider and ATLAS. A discussion of some upgrades to be made to these for the future Run 4 will be described in Chapter 6.

3.1 Circular Hadron Colliders

In general, particle colliders transform energy into matter. The goal of this is to create very exotic matter that was present in the earliest moments of the universe. While it is more complex to design a circular accelerator as opposed to something like a linear fixed target, where one beam of particles is accelerated in a straight line towards a stationary object, there is a large payoff in terms of collision energy. In a fixed target experiment, the energy of the collision is proportional to the square root of the beam energy:

$$E \propto \sqrt{E_{beam}} \tag{3.1}$$

In a collider, where two beams of particles are accelerated independently and then collided together, the energy is simply the sum of the two energy beams:

$$E = E_{beam1} + E_{beam2} \tag{3.2}$$

So by colliding two beams together, we are able to push our collision energies higher, which allows processes involving higher rest mass particles.

3.2 The Large Hadron Collider

The Large Hadron Collider (LHC) is a 27 km circumference circular particle accelerator located at the European Organization for Nuclear Research (CERN) [30]. It is about 100 m underground and spans across the border between France and Switzerland. Utilizing superconducting magnets, the LHC accelerates and collides two beams of hadron bunches, primarily protons but also heavier ions, at a maximum energy of 7 TeV per beam for a total maximum center of momentum energy of 14 TeV. There are 4 main experiments which collect data at interaction points along the LHC ring: ATLAS [8], CMS [31], ALICE [32], and LHCb [33]. ATLAS and CMS are general purpose detectors covering a 4π solid angle around the interaction points, ALICE studies heavy-ion collisions, and LHCb studies flavor physics, particularly focusing on b-quarks. In LHC Run 3, spanning from Summer 2022 until Summer 2026, the LHC collided proton bunches at an energy of 13.6 TeV in order to collect a large Run 3 data set on the order of 300 fb^{-1} (the units of inverse femtobarns will be discussed later).

3.2.1 CERN Accelerator Complex.

Before protons reach the LHC, they undergo several steps at smaller accelerators within the CERN site. Negative hydrogen ions (H^-) are obtained from a source in Linear Accelerator 4 (LINAC4) [34]. LINAC4 accelerates the H^- ions in four steps to an energy of 160 MeV, at which point they are injected into the Proton Synchrotron Booster (PSB). At the time of injection into the PSB, the electrons are stripped away from the ions, leaving only protons. The PSB consists of 4 superimposed synchrotron rings that accelerate the protons to an energy of 2 GeV before injecting into the Proton Synchrotron (PS). The Proton Synchrotron is the first proper particle accelerator in the chain. It is a circular accelerator measuring 628 m in circumference and consists of 277 room temperature electromagnets, with 100 of those being dipoles to bend the beam around the ring. PS accelerates the protons to an energy of 26 GeV before sending them to the Super Proton Synchrotron (SPS). The SPS is (as the name suggests) a larger version of the PS. It is around 7 km in circumference and consists of 1317 room temperature electromagnets, 744 of which are dipoles for bending the beam around the ring. Protons are accelerated in the SPS to an energy of 450 GeV, at which point the beams are fed into the LHC for the final ramping. A schematic of the full complex can be found in Figure 6.

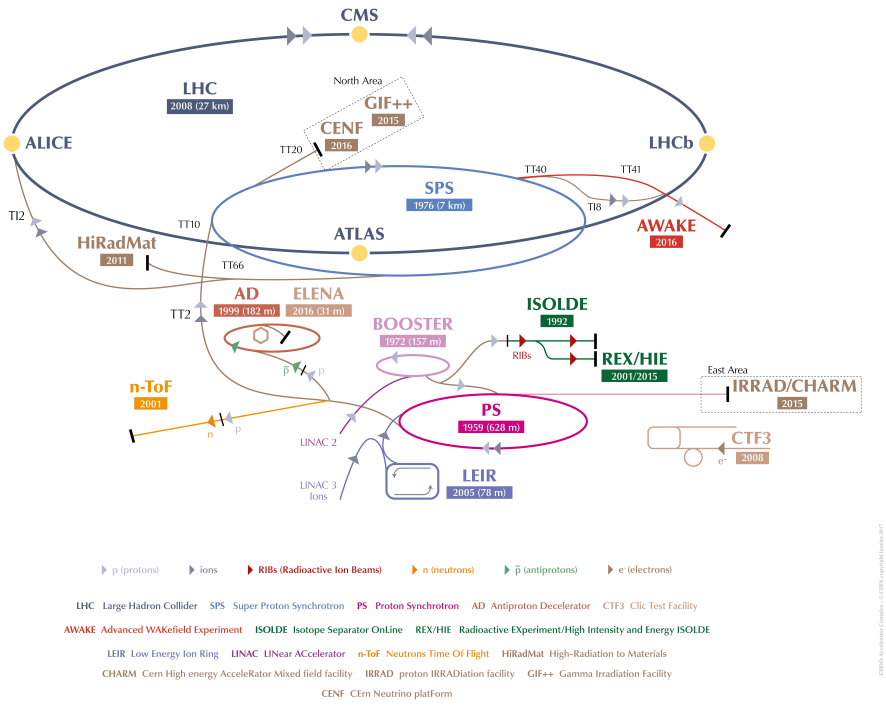


Figure 6. A schematic of the CERN accelerator complex, showing the full proton injection chain for the LHC [4].

3.2.2 LHC Magnets, Acceleration, and Proton Bunches.

The 27 km long ring consists of about 10,000 superconducting magnets, including 1232 dipoles to guide the beam along the ring, and 392 quadrupoles to keep the beam focused. Since the magnetics are superconducting, the LHC uses liquid helium to cool the magnets to their operating temperature of 1.9 K. The particle beams are accelerated using Radio Frequency (RF) cavities that apply a 5 MV/m electric field alternating at 400.8 MHz with a revolution frequency of 11.245 kHz [35]. The alternating field creates "buckets" where protons will cluster together longitudinally into bunches, each containing 10^{11} protons. There are 35640 possible buckets in the machine. The bunches are spaced at 25 ns intervals (meaning only 1 in every 10 possible bunches could be filled), resulting in a total of 3564 possible locations of bunches [36]. In practice, not all of these slots are filled due to necessary gaps needed for injection from the SPS and a 3 μs abort gap in the bunch train needed for the LHC kicker magnet to ramp up its field for a beam dump. This results in a maximum possible filling scheme of about 2800 bunches. Furthermore, because of these gaps, bunches are not filled uniformly, but rather injected in batches called bunch trains to leave room for these injection gaps and other machine stability considerations. Taking all these constraints into account, a typical physics run in Run 3 consisted of 2448 filled bunches for ATLAS.
The collision rate of these proton-proton interactions is the instantaneous luminosity, $\mathcal{L}_{\mathrm{i}}$, and can be expressed in terms of beam parameters for a Gaussian beam distribution [30]

$$\mathcal{L}_i = \frac{N_b^2 n_b f_{\mathrm{rev}} \gamma_r}{4\pi \mathcal{E}_n \beta *} F \tag{3.3}$$

where N_b is the number of particles per bunch, n_b is the number of bunches per beam, f_{rev} is the frequency of revolution, γ_r is the relativistic factor, $\mathcal{E}_n$ is the normalized transverse beam emittance, $\beta *$ is the beta function at the collision point, and F is the geometric luminosity reduction factor due to the interaction point crossing angle. This instantaneous luminosity is directly related to the number of events generated per second in LHC collisions

$$N_{\mathrm{event}} = \mathcal{L}_i \sigma_{\mathrm{event}} \tag{3.4}$$

where σ_{event} is the cross section for a given event. For Run 3, the instantaneous luminosity is around 2×10^{34} cm^{-2}s^{-1}. This instantaneous luminosity is also related to the overall amount of data taken by the experiments. Specifically, the integrated luminosity,

$$\mathcal{L} = \int \mathcal{L}_i dt \tag{3.5}$$

which one can note has units of cm^{-2} after integrating over time. We often refer to the amount of data taken over an entire dataset in terms of the total integrated luminosity. Rather than express this value in terms of cm^{-2}, we choose a different unit of area, called a barn, where one inverse barn is equal to 10^{24}cm^{-2}. For the scale of LHC data-taking runs, we often use inverse femtobarns ($\text{fb}^{-1} = 10^{-15}\text{b}^{-1}$). For example, the total delivered integrated luminosity for Run 3 proton-proton collisions from the LHC is about 350fb^{-1} [37]. The units of inverse femtobarns to describe the amount of data taken is convenient because we can then easily calculate the number of events for a process with a known cross section (measured in femtobarns) by multiplying the cross section by the integrated luminosity.
As the luminosity increases, the probability that multiple proton-proton interactions occur in one bunch crossing also increases. The multiple interactions from the current bunch crossing as well as observed interactions from adjacent bunch crossings are referred to as "pileup." Pileup can be characterized by the mean number of interactions for a given cross section

$$\langle \mu \rangle = \frac{\langle \mathcal{L}_i \rangle \sigma_{\text{inel}}}{f_{\text{rev}}} \tag{3.6}$$

where σ_{inel} is the cross section of inelastic proton-proton collisions. Figure 7 [37] shows the pileup profile for most of the data-taking years for the LHC Run 3 as measured by the ATLAS detector. In general, the mean and most probable value of the pileup has increased for subsequent data-taking years, allowing for the collection of more data, and for important practice in higher pileup environments in order to prepare for the High-Luminosity LHC (more details in Chapter 6).

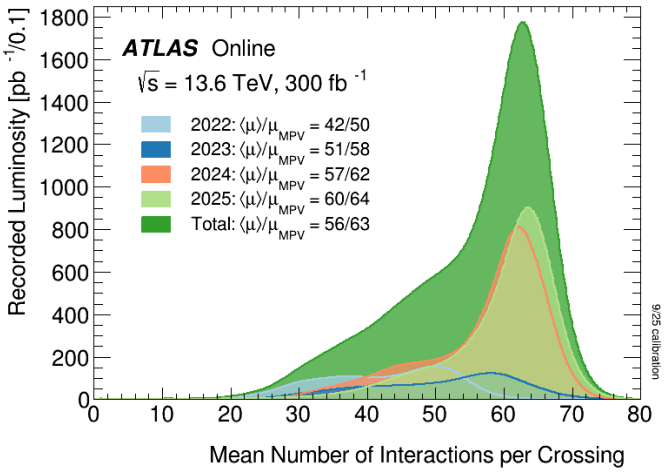


Figure 7. Plot of the luminosity-weighted distributions of the pileup for Run 3 data-taking years as well as the total across Run 3. The average value increases in later data-taking years as the Run 3 system became more mature.

3.3 The ATLAS Detector

The ATLAS Detector is a cylindrical all purpose detector located at interaction point 1 of the LHC. It covers nearly the entire solid angle around the interaction point, meaning that it can capture essentially all the information about a particle collision that occurs within its center. ATLAS is a massive machine that measures 46 m in length, 25 m in diameter, and weighs approximately 7000 tonnes. The detector is built in layers: the inner most layers are used for tracking charged particles, the electromagnetic and hadronic calorimeters capture energy from most of the decaying particles, and the muon spectrometer captures tracks from muons. Various aspects of the detector components and data acquisition system were upgraded in the time between Run 2 and Run 3, known as Long Shutdown 1 (LS1). The configuration of ATLAS for Run 3 will be discussed here in detail.

3.3.1 Coordinate System.

ATLAS is a cylinder, so it makes sense to use a cylindrical coordinate system (r, ϕ, z) where the z-axis is aligned along the beam line, and the origin of the coordinate system is centered at the interaction point (IP). However, rather than describe physics objects in these coordinates, it makes more sense to consider a special quantity, the *pseudorapidity*, η, which effectively measures the angle of the particle from the beamline defined as

$$\eta = -ln(tan(\frac{\theta}{2})) \tag{3.7}$$

where θ represents the polar angle. The pseudorapidity itself is the massless limit of true rapidity, which is defined as

$$y = \frac{1}{2} ln(\frac{E + p_{\text{Longitudinal}}}{E - p_{\text{Longitudinal}}}) \tag{3.8}$$

where E is the energy of the particle, and $p_{\text{Longitudinal}}$ is the momentum of the particle along the beam axis. While rapidity requires knowing the precise energy and longitudinal momentum of a given particle, pseudorapidity is purely geometric, making it a useful quantity for describing any particle's position with respect to the beam line. While this is exact for massless particles, the high energy of collisions compared to the rest mass energy of SM particles makes pseudorapidity a good approximation of rapidity for all particles. Differences in particle pseudorapidity are Lorentz invariant, meaning this quantity does not change due to particles traveling extremely fast along the beam axis. The pseudorapity and the azimuthal angle ϕ can be used to define the distance of two objects in $\eta - \phi$ space, referred to as ΔR

$$\Delta R = \sqrt{\Delta\eta^2 + \Delta\phi^2} \tag{3.9}$$

which can be used to refer to the angular separation between two objects as well as define the radius of jet objects.

3.3.2 The Inner Detector.

The Inner Detector (ID) is used to track the paths of all charged particles in the ATLAS Detector. It is 7.024 m long, surrounds the beam line radially outward to 1.150 m, and provides tracking including pattern recognition, momentum resolution, and primary and secondary vertex measurements for particles within $|\eta| < 2.5$. The ID is comprised of three subdetectors: the high granularity silicon pixel detector [38], the micro-strip based Semiconductor Tracker (SCT) [39], and the gaseous straw tube Transition Radiation Tracker (TRT) [40]. The entire ID is immersed in a 2T magnetic field from the solenoid magnet, causing the charged particles to bend according to their charge and mass. A diagram of the ID with all its subcomponents can be found in Figure 8.

Tracking is performed by measuring the charge, timing information, and position (given by which sensor was hit) of a charged particle as it passes through layers of the ID. When a particle passes through a layer of the ID, it ionizes the detector material and creates electron-hole pairs. The electrons drift due to the electric field from a reverse bias voltage difference across sensors as well as the existing magnetic field. These charges are collected and registered as "hits" in the detector. Clusters of hits from multiple layers are used to reconstruct a "track" for the particle. The charged particle will bend in accordance to how it moves in the magnetic field, so the radius of curvature of the track can be used to find the particle's momentum, and the direction of it's track tells the charge. We can additionally use trackers to determine whether a particle came from the primary vertex or a secondary vertex, as well as some limited particle identification (PID) capabilities through the detection of transition radiation X-ray photons [5].

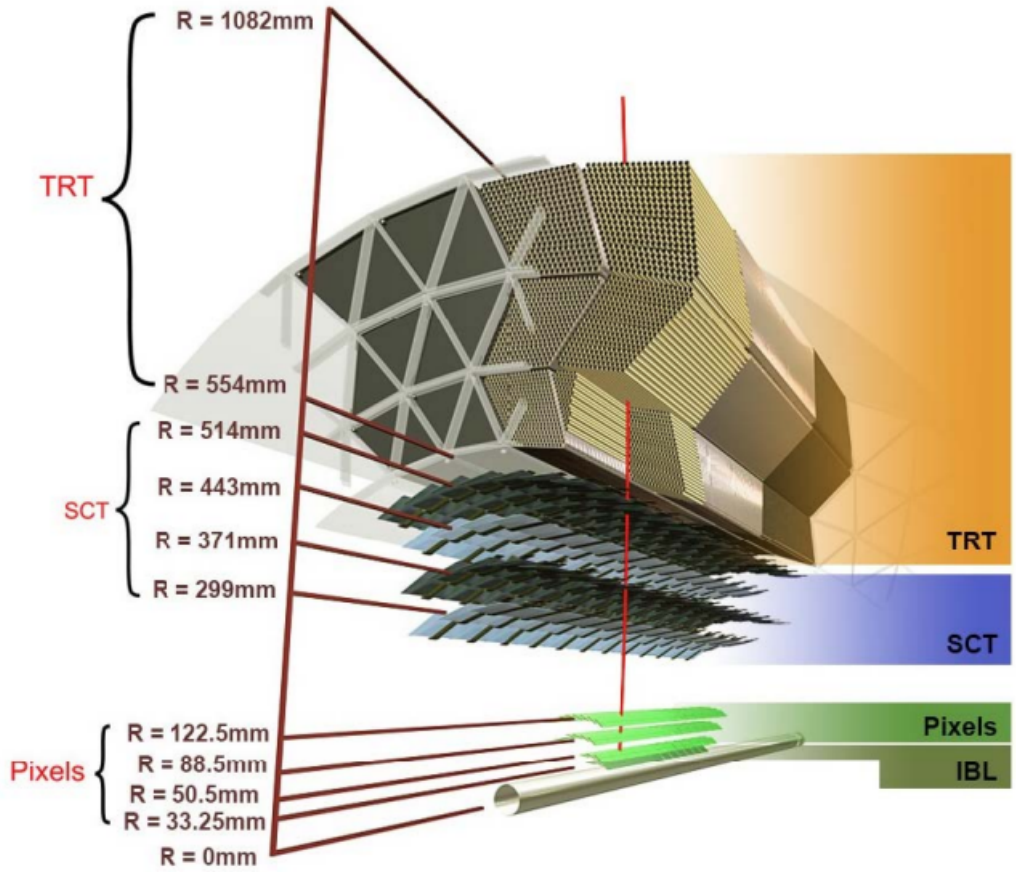


Figure 8. Diagram of a radial slice of the ATLAS Inner Detector system. The distance from the relevant subsystems and layers from the beamline are listed [5].

3.3.2.1 The Silicon Pixel Detector.

The innermost part of the ID is the Silicon Pixel Detector. The Pixel Detector itself has four overlapping layers in the barrel region. From the beampipe radially outwards, the layers are the Insertable B-Layer (IBL) at a distance of 33.25 mm from the beampipe, the Pixel B Layer (Layer 0) at a distance of 50 mm, Pixel Layer 1 at a distance of 88.5 mm, and Pixel Layer 2 at a distance of 122.5 mm. Additionally, three discs on each side of the detector provide end cap coverage.

The IBL was installed during LS1 to improve the Pixel Detector's performance in Runs 2 and 3. It accomplishes this in several ways. The decreased distance of the IBL to the beam line with respect to the Pixel B Layer increases the resolution of track impact parameters. This improves offline vertex reconstruction and flavor-tagging, which is highly important for both precision SM measurements and BSM searches, particularly for searches involving LLPs that decay farther away from the beam line. The addition of the highly granular IBL also helps mitigate performance degradation in the other Pixel Layers caused by radiation

damage and irreparable module failures. The IBL consists of pixels 50 μm $\times$ 250 μm organized in 80 $\times$ 336 arrays, with some pixels longer than the nominal to fill gaps between front-ends. In order to make room for the IBL, the original beam pipe was also replaced, going from an inner radius of 29 mm down to 23.5 mm. The three remaining Pixel Layers have identical pixel sensors that are slightly coarser, being comprised of 50 μm $\times$ 400 μm pixels. In total, there are 92 million pixel readout channels, with 12 million coming from the IBL, and the remaining 80 million from Layers 0-2. The granularity of the overlapping sensors gives about a 10 μm precision in the $r\phi$ plane and 115 μm in the z direction.

3.3.2.2 The Semiconductor Tracker.

The next layer of the ID is the Semiconductor Tracker (SCT), which spans 299 mm $< r <$ 514 mm radially outward from the beamline. The SCT is a silicon microstrip tracker comprised of four layers covering the barrel region and nine end cap disc layers to cover the pseudorapidity up to $|\eta| < 2.5$, with each of the layers providing a measurement for a charged track. The SCT consists of approximately 6.3 million readout channels, and is designed to aid the Pixel Detector in measuring particle momentum and vertex identification. Its lower granularity results in slightly worse precision than the Pixel Detector, with 17 μm precision in the $r\phi$ plane and 580 μm in the z direction.

3.3.2.3 The Transition Radiation Tracker.

The outermost section of the ID is the Transition Radiation Tracker (TRT), which spans 563 mm $< r <$ 1066 mm radially outward from the beam line. The TRT is composed of drift tubes filled with gas. There are about 300,000 tubes interleaved with transition radiation material that provide on average 30 points with 130 μm precision in the $r\phi$ plane for charged particle tracks for $|\eta| < 2$ and $p_T > 0.5$ GeV. Each tube is 4 mm in diameter and filled with mostly xenon gas as well as a central grounded wire made of tungsten plated in gold. As a charged particle passes through the gas, the particle ionizes the gas, and the electrons are collected by the grounding wire, providing the hits. In addition to aiding in the p_T resolution, the TRT provides electron identification capabilities through its detection of transition radiation X-ray photons. As the charged particles pass through different media of the TRT, radiation in the form of photons is produced at the boundaries of different media. Since the energies of these resulting photons are dependent on the mass of the particle, the TRT can help distinguish electrons from heavier charged hadrons.

3.3.3 Calorimeters and Calorimetry.

The next major layer of the ATLAS detector outside the ID are the calorimeters. The calorimeters absorb and measure the energy of particles after they pass through the tracking system. ATLAS takes advantage of two main sampling calorimeter technologies: Liquid Argon (LAr), which is used for all electromagnetic calorimetry as well as hadronic calorimetry in the end caps, and scintillating tiles (Tile), which is used for hadronic calorimetry in the barrel region

[41]. Sampling refers to the structure of the calorimeters, meaning they alternate between an active medium and an absorber material in layers. Figure 9 shows the major components of the calorimeter system in ATLAS. The calorimeter system covers the entirety of the detector up to $|\eta| < 4.9$. The calorimeter system is designed to capture the energy of all particles that emerge from the collisions. It is used to measure the energy and position of electrons, photons, and hadrons stemming from other particle decays. All species of particles are accounted for by the calorimeters except muons and neutrinos. Muons are visible in the calorimeter, but are minimally ionizing particles that deposit energy negligible compared to its energy and the energy of particles which the calorimeters are designed to measure. Neutrinos fly through the calorimeters (and the rest of the detector) with no interaction due to their very small interaction probabilities with other particles. Muons have a dedicated system that will be discussed in following sections, and neutrinos interact so sparingly that their experimental signature in the detector is missing transverse energy. The following sections will discuss the electromagnetic calorimeter (ECAL) and the hadronic calorimeter (HCAL).

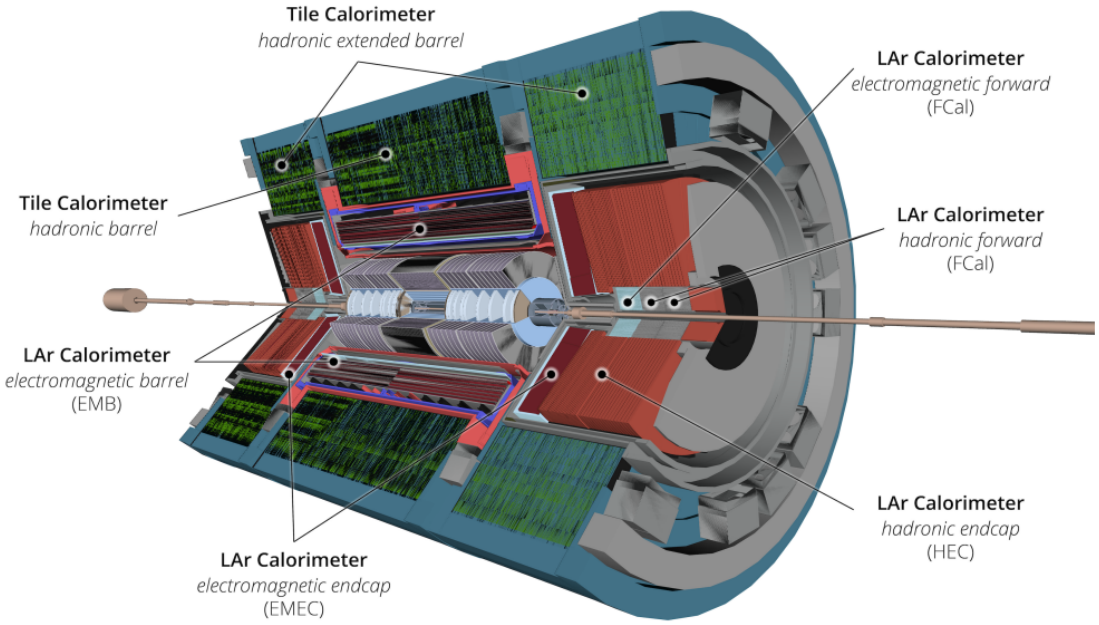


Figure 9. Cutaway diagram of the ATLAS calorimeter system with major subcomponents labeled [5].

3.3.3.1 The Electromagnetic Calorimeter.

ATLAS utilizes Liquid Argon (LAr) for electromagnetic calorimetry in the barrel (the Electromagnetic Barrel Calorimeter, or EMB) and end caps (the Electromagnetic Endcap Calorimeter, or EMEC). The liquid argon serves as

the active medium while steel acts as the dense absorber layer. These layers are arranged in an accordion like pattern, with LAr filling the gaps. The LAr calorimeter measures the energy of electrons and photons primarily, as well as some τ leptons and hadrons as they interact with the calorimeter material. The ECAL is comprised of a presampler layer and three subsequent layers, covering an η range of $|\eta| < 1.52$ in the barrel region and a range of $1.5 < |\eta| < 3.2$ in the end cap region. It is designed to contain the entirety of an electromagnetic shower, so its corresponding absorption power is 24-27 radiation lengths (X_0).

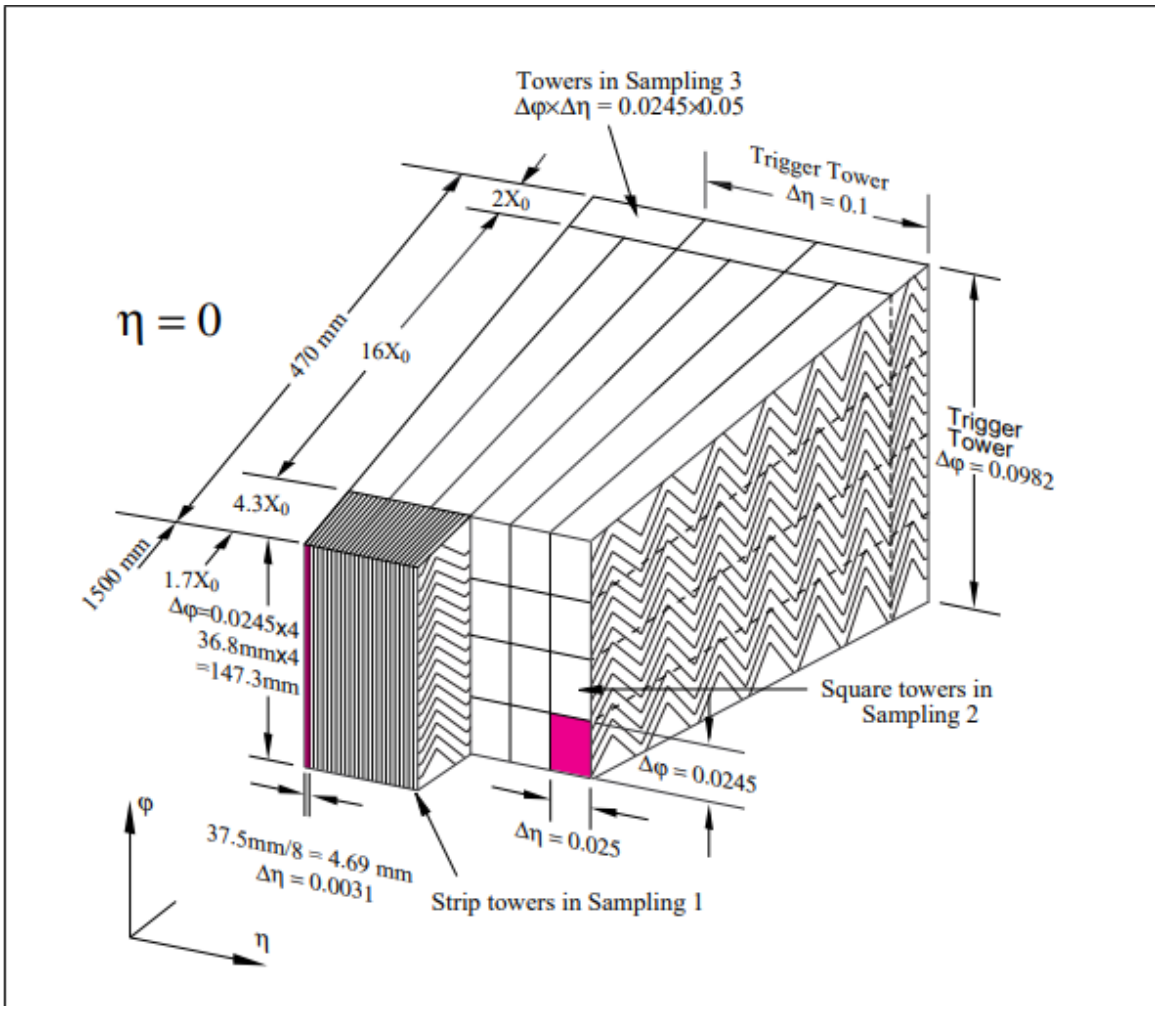


Figure 10. The cell break down by layer of a central EMB trigger tower. The accordion like structure of the LAr calorimeter can be seen on the right side of the graphic [6].

The ECAL calorimeter components are comprised of elementary calorimeter cells that vary in size according to layer and η position. In Run 3, to provide information to the Level-1 Trigger System, the elementary cells are combined into

lower granularity Super Cells as input to the digital trigger readout path. This is an improvement compared to previous Runs of the LHC, where the highest granularity information LAr provided was 0.1×0.1 in $\eta - \phi$ space trigger towers. Figure 11 shows a schematic comparing the legacy trigger towers to the Run 3 Super Cells. This new functionality to provide Super Cells to the trigger system is due to the upgraded LAr digital trigger system, which ran in parallel with the legacy system at the start of Run 3 [5].

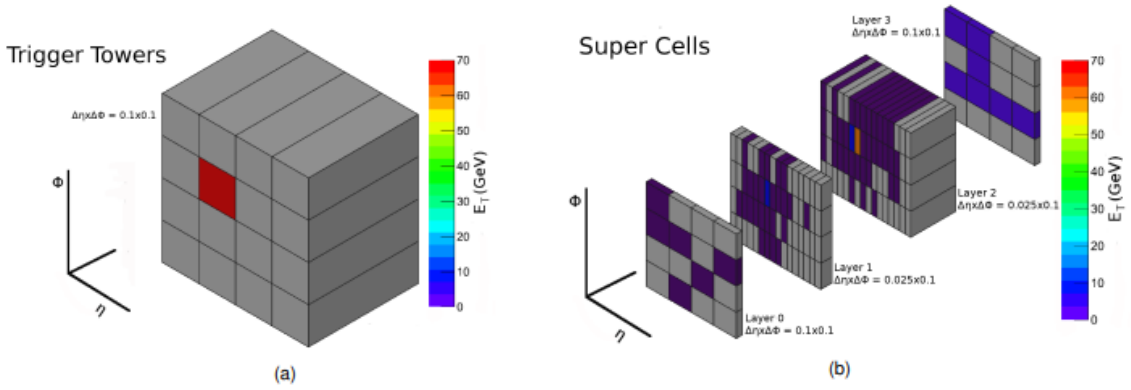


Figure 11. A Legacy trigger tower 0.1×0.1 in $\eta - \phi$ (a) and a schematic of a Super Cell (b) from [6].

3.3.3.2 The Hadronic Calorimeter.

The goal of the HCAL in ATLAS is to capture the energy of hadrons like pions and kaons that punch through the electromagnetic calorimeters. The particles that punch through the ECAL interact hadronically with the dense medium of a hadronic calorimeter, resulting in hadronic showers of particles. For hadronic calorimeters, we are typically concerned with the interaction length λ, the analog to the ECAL's radiation length X_0. λ is the characteristic distance where the hadron has a $1 - \frac{1}{e}$ chance of interacting with the material. The HCAL is designed to capture all the hadronic showers of the particles passing through, with a combined absorption power of $\geq 10\lambda$ in the entire $|\eta| < 4.9$ coverage of the detector.

The HCAL comprises three different calorimeters to cover different regions of pseuorapidity. In the barrel region, the ATLAS detector utilizes scintillating plastic tiles in a steel matrix (Tile) in the central region of $|\eta| < 1.7$. Tile is comprised of three layers in both its main barrel region ($|\eta| < 1$) as well as the extended barrel region $0.8 < |\eta| < 1.7$. The signal from Tile is carried by photons and collected by photomultiplier tubes (PMTs). The photons come from the excitations of molecules in the scintillating plastic when they return to their ground state. In the end cap region spanning $1.5 < |\eta| < 3.2.$, the Hadronic Endcap (HEC) sits just after the EMEC, and uses liquid argon technology similar to the ECAL. Finally the Forward Calorimeter (FCAL) covers the extreme forward region

of $3.1 < |\eta| < 4.9$, and uses LAr as the active medium alternating with copper or tungsten. The first copper layer is designed to aid in EM measurements, while the two tungsten layers are for hadronic measurements.

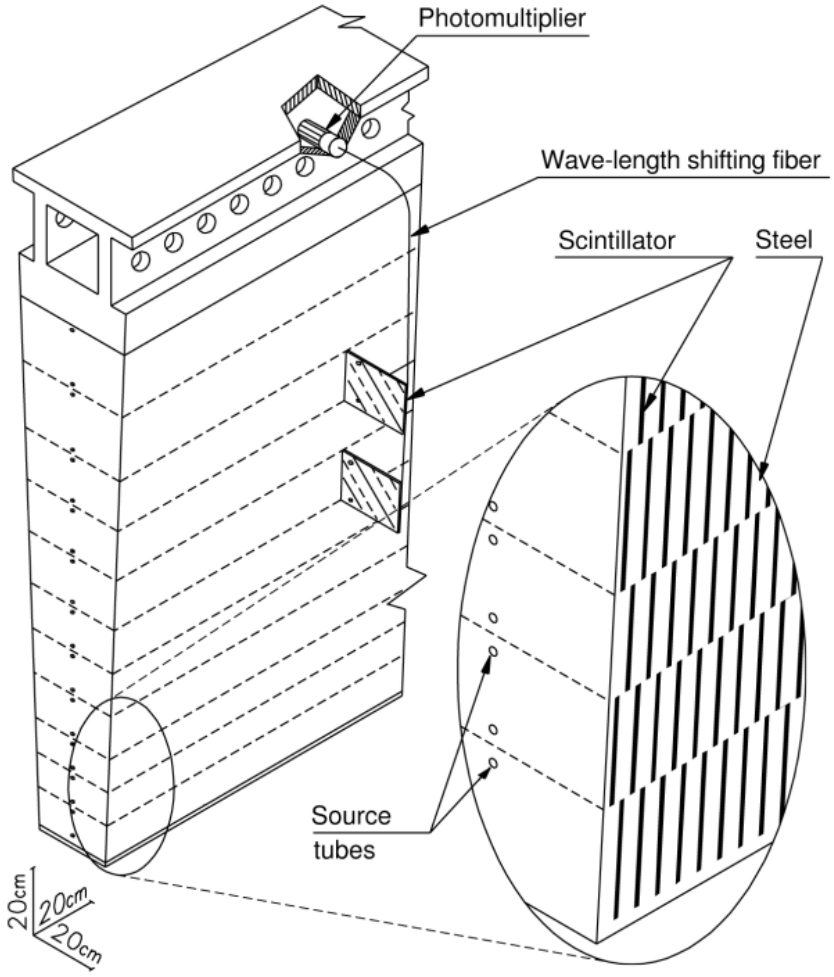


Figure 12. A module of the Tile Calorimeter, showing the PMTs connected to the scintillator material in it's alternating pattern with the steel [7].

3.3.4 The Muon Spectrometer.

The Muon Spectrometer (MS) is the outermost layer of the detector, and is of course dedicated to the identification and momenta measurements of muons. Since muons from collisions are traveling near the speed of light, they pass through the rest of the ATLAS detector with minimal energy loss, making them one of the few particles we can detect directly in ATLAS as opposed to measuring their decay products. The MS consists of a barrel containing three concentric stations (called inner, middle, and outer barrels) as well as two end caps each containing

three disks (called inner, middle, and outer wheels). The MS provides coverage up to $|\eta| < 2.7$. Figure 13 shows a cutaway diagram of the muon system with labeled subcomponents.

The muon system relies on deflection of muon tracks in large superconducting air-core toroid magnets, which contain both trigger chambers and high-precision tracking chambers. Within $|\eta| < 1.4$, the bending comes from the large barrel toroid. For $1.6 < |\eta| < 2.7$, the bending comes from the smaller endcap toroid magnets at the ends of the barrel toroid. In the space between these two regions, a combination of the barrel and endcap toroids contribute to the bending. All three barrel stations as well as the middle and outer wheels are made of multilayered Monitored Drift Tube (MDT) chambers to precisely measure the bending of muon tracks. The MDTs are 30 mm diameter aluminum tubes filled with Ar and CO_2 gases (93% and 7% respectively), that provide excellent spatial resolution of about 10 μm. The trigger chambers in the middle and outer barrel stations are covered by Resistive Plate Chambers (RPCs) and in the middle wheel endcap by Thin Gap Chambers (TGCs). RPCs are made of two parallel plates with a thin gap between them filled with gas (primarily $C_2H_2F_4$), which creates an avalanche of electrons, resulting in a fast response and high temporal resolution. TGCs are multiwire proportional chambers with a 2.8 mm gas gap, containing CO_2 and n-pentane. The short gap provides a small drift distance which results in a fast response and good time resolution.

New for Run 3 is an upgrade to the inner most wheel, called the small wheel. For Run 3, this small wheel was replaced with the New Small Wheel (NSW) which uses two different chamber technologies: small-strip TGCs (sTGCs) and micro-mesh gaseous structure (Micromegas) for both precision tracking and triggering capabilities [42]. The NSW upgrade was one of the primary hardware upgrades for Run 3, primarily motivated by the need for improved p_T resolution in the muon trigger and increased background rejection. This allows for a lower muon p_T threshold at a manageable L1 trigger rate.

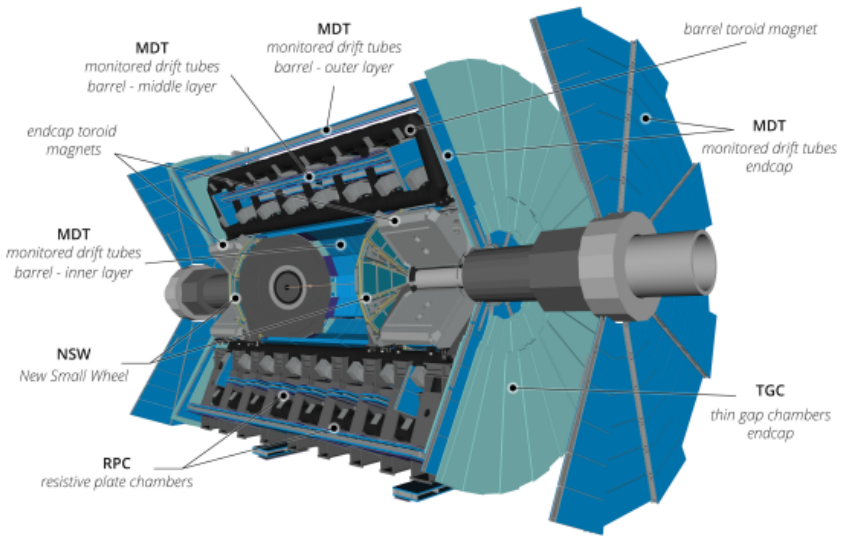


Figure 13. Cutaway diagram of the Muon Spectrometer with labeled subcomponents [5].

3.3.5 The Magnet System.

Aspects of the ATLAS magnet system have been mentioned in passing multiple times throughout the previous sections, but it will be described here as well for completeness. The magnet system comprises a solenoid magnet surrounding the ID, as well as large air-core toroidal magnets surrounding the calorimeters and parts of the MS [8]. All magnets use superconducting coils cooled using liquid helium to a temperature of 4.5 K. The solenoid provides a 2 T magnetic field that immerses the ID, providing the magnetic field used for charged particle tracking. The barrel and endcap toroids surround the calorimeters. They are made of 8 superconducting coils and are arranged around the outside of the calorimeters in the case of the barrel and in a disk shape at each endcap, providing a magnetic field ranging 2-6 T across the majority of the detector. Figure 15 shows a schematic of the magnet system.

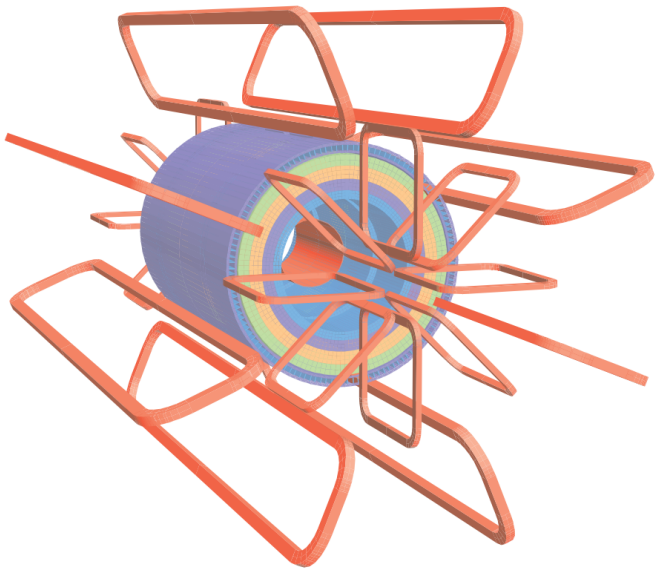


Figure 14. Schematic of the magnet system in ATLAS, with the cylindrical solenoid magnet in the center, surrounding the ID, and the superconducting coils of the toroids arranged around the barrel and arranged in a disk on each endcap [8].

3.4 The Trigger and Data Acquisition System

We have established that the rate of digitizations for proton-proton collisions at the LHC is 40 MHz (40 million bunch crossings per second). The size of one event in the RAW data format is about 2 MB. In order to save all of those events, ATLAS would have to write out 80 TB of data per second. This is unfeasible, both to write out and to store, so we can only save a portion of all events. Since the majority of the events will simply be QCD background (quarks and gluons scattering to other quarks and gluons), we need to be smart about which events we choose to save. ATLAS implements a two-tiered trigger system that runs on each bunch crossing and is designed to filter for interesting physics events by using selection algorithms for different physics objects. This section will describe the two tiers of the trigger system: the Level-1 (L1) hardware trigger and the High Level Trigger (HLT) software trigger.

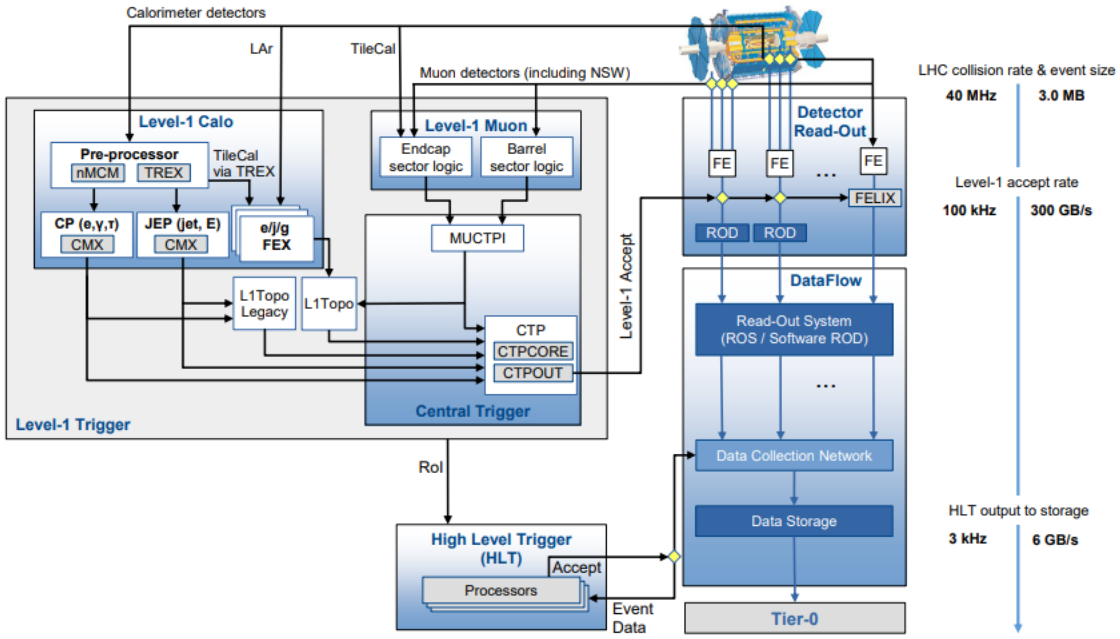


Figure 15. TDAQ data flow diagram for Run 3 [5].

3.4.1 Level-1 Trigger System.

The Level-1 Hardware trigger is responsible for reducing the initial 40 MHz event rate down to a 100 kHz "Level-1 Accepts" (L1As) that will be passed further down the trigger chain to the High Level Trigger. To achieve this rate, the L1 system has a total 2.5 μs latency to make a trigger decision, which includes the both the time needed to run algorithms that filter events according to specific criteria (often $\lesssim$200 ns) as well as the time to communicate back to the system whether to keep or delete the information of an event. In Run 3, much of the hardware trigger was upgraded to replace the legacy system's Application-Specific Integrated Circuits with Field Programmable Gate Arrays (FPGAs) to perform more complex algorithms. FPGAs allow for customizable algorithms that can be loaded via firmware. The L1 Hardware Trigger system is split into two main subsystems: the Calorimeter Trigger system (L1Calo) and the Muon Trigger system (L1Muon), which both feed into the L1 Topological Processor (L1Topo). L1Topo performs final algorithms before sending information forward to the L1 Central Trigger Processor (L1CTP). A schematic of the entire L1 system can be found in Figure 16.

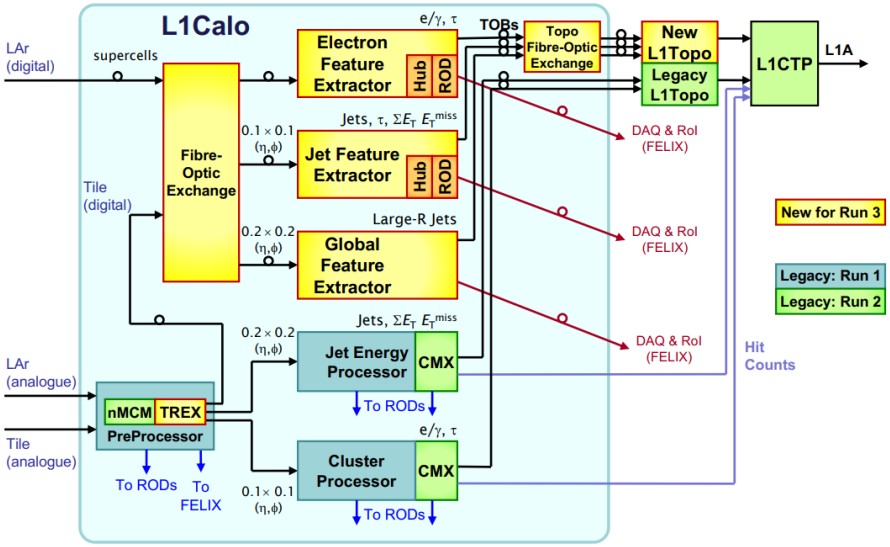


Figure 16. The L1 Trigger logic diagram for Run 3 [5].

3.4.1.1 L1Calo.

In Run 3, L1Calo consists of a few pieces of the Run 2 trigger system (referred to as the "legacy" system) as well as several new feature extractors (FEXes) and infrastructure. The FEXes provide Trigger Objects (TOBs) to L1Topo for further processing. The legacy L1 Trigger path was active in the beginning of Run 3 until 2024 when the Phase-I system was fully commissioned and validated. It consists of the Pre-Processor Module (PPM), the Jet Energy Processor (JEP), and the Cluster Processor (CP). The PPM receives analog signal from the calorimeters and digitizes it, then forming it into 0.1×0.1 trigger towers before sending to the JEP and CP modules for identification of physics objects. The CP identifies EM and hadronic candidates, while the JEP is responsible for creating potential jet clusters.

The Phase-I upgrades primarily consist of the new FEXes designed to run firmware algorithms on FPGAs to select specific physics objects. The electron feature extractor (eFEX) identifies electron and photon candidates as well as τ leptons. The jet feature extractor (jFEX) identifies small radius ($R = 0.4$), missing transverse energy, as well as handling more forward electrons and taus. eFEX and jFEX are direct equivalents to the legacy trigger system, and can be thought of as the successors of the legacy CP and JEP, respectively. There is also a new global feature extractor (gFEX), which is designed to identify large radius ($R = 0.9$) jets as well as provide differing missing transverse energy algorithms. gFEX architecture, algorithms, and commissioning will be discussed in greater detail in Chapter 4. For both large-R and small-R jets, the stated radii are not exact due to the geometry of the constituent towers, but are correct within rounding error. Figure 17 shows a production board of each FEX.

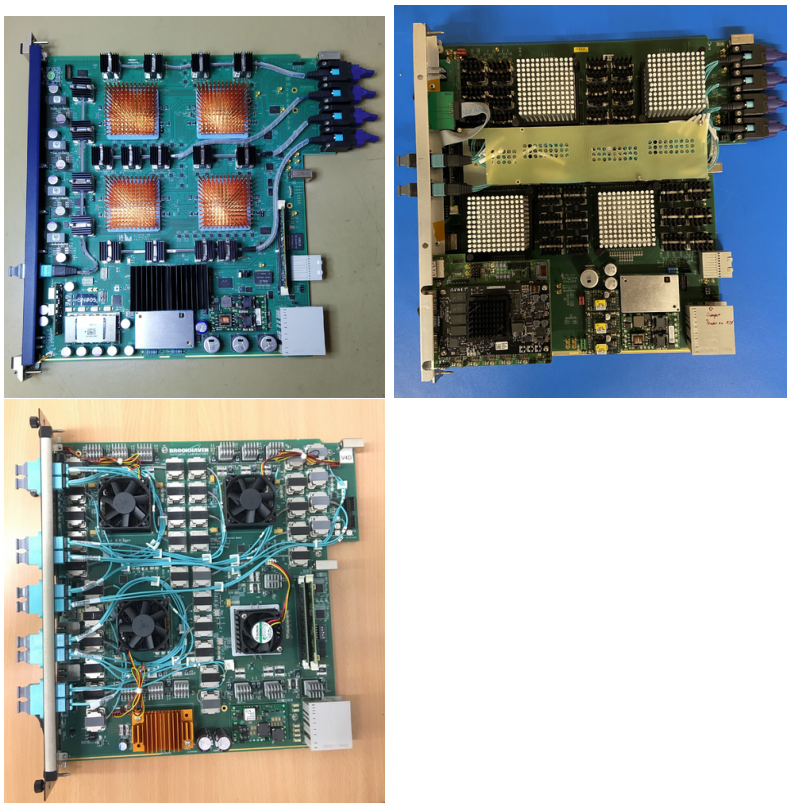


Figure 17. From left to right, a production eFEX, jFEX, and gFEX board [9].

The eFEX system consists of 24 ATCA boards each hosting 4 Virtex 7 processor FPGAs. It is designed to use the full available granularity of the Run 3 calorimeter system, ranging from 0.1×0.1 to 0.025×0.1 in $\eta - \phi$ depending on the layer of the calorimeter. eFEX defines a 0.3×0.3 window to search for isolated energy clusters indicative of a photon, electron, or τ. Additionally for the τ, a boosted decision tree is implemented to further aid in isolation criteria.

The jFEX system consists of 6 ATCA boards each hosting 4 Virtex Ultrascale+ processor FPGAs. It takes 0.1×0.1 input towers and uses a 0.9×0.9 sliding window algorithm to identify local maxima as small-R jet and tau candidates. It additionally provides E_T^{miss}, $\sum E_T$, and forward electron triggers.

The gFEx system consists of 1 ATCA board hosting 3 Virtex Untrascale+ processor FPGAs. Unique to the FEXes, it has access to the entire range of the calorimeters on one board through the use of coarse 0.2×0.2 input towers. It uses this coarse input to identify large radius jets, small radius jets, E_T^{miss}, and $\sum E_T$.

In addition to the FEX modules, the Phase-I upgrade includes other new modules to facilitate the transfer of optical data to L1Calo and through the entire readout chain of the system. These include the Tile Rear Extension Module (TREX), the Fiber Optic Exchange box (FOX), the Front End Link Exchange (FELIX), and the software read out device (SWROD). The TREX digitizes the

Tile calorimeter data, and this data along with the data from LAr, which is digitized on Detector before coming to L1Calo, are routed through the FOX from the calorimeters to the correct FEX module. The FELIX receives the optical output from the FEX systems, and the SWROD receives data from the FELIX and forwards this data to both the HLT and ultimately to the CERN Tier 0 data storage facility.

3.4.1.2 L1Muon.

L1Muon receives input from the MS in order to identify muon candidates. It does this by measuring the degree to which a candidate's path deviates in $r - \phi$ from the pattern of hits expected from a muon with infinite momentum [5]. These deviations are inversely proportional to the p_T of the candidate muon. "Muon roads" are defined using expected deviations for different muon p_T thresholds. In Run 3, there are 15 programmable thresholds available in the endcap, an increase from Run 2's available 6.

3.4.1.3 L1Topo.

L1Topo is technically part of L1Calo, but deserves its own section here, as it takes the outputs of the L1Calo FEXes as well as the L1Muon system as its input. It consists of 3 ATCA boards each with two Virtex UltraScale+ FPGAs. One board is responsible for identifying L1Calo trigger multiplicities, and the other two perform topological selections. The main categories of L1Topo algorithms are: *sort/select/no-sort*, where TOBs from all the inputs are formatted into a single global data format, and can be sorted by E_T or p_T or selected for passing a certain threshold. They are then used as inputs for the following decision and multiplicity algorithms: *Decision*, where the algorithm determines if a given trigger condition has been satisfied, and *multiplicity*, where the algorithm counts the number of objects passing a given threshold. Both send their decision or multiplicity bits onward to the CTP. The latencies for these algorithms are extremely tight, being limited to 2, 1, and 3 bunch crossings (or 50, 25, and 75 ns).

3.4.2 High Level Trigger System.

After passing the L1 trigger, events are sent to the High Level Trigger (HLT) for further refinement. The HLT uses software based algorithms implemented on a computing farm to reconstruct physics objects in Regions of Interest (RoIs) identified by the L1 system. Events are distributed to HLT on nodes that run object reconstruction using Athena, the offline reconstruction software [43]. The online reconstruction happens in sequence, where dedicated fast trigger algorithms reject a majority of events and are then followed by the more CPU-intensive algorithms that are similar or identical to those that run offline. In Run 3, Athena was upgraded to include multi-threading capability, which allows the HLT to operate at an event rate of 3 kHz, increased from the Run 1 and 2 software trigger output of 1 kHz [44]. Events in the HLT are processed on the order of 700 ms, where this limit is driven by the number of available CPU cores in the farm rather than the latency of the HLT hardware (which is on the order of several

minutes). If an event passes all the HLT algorithm requirements, it is saved to disk and stored to be used in offline analyses.

CHAPTER IV
GFEX ALGORITHM VALIDATION

This chapter describes in detail the Global Feature Extractor subsystem of the ATLAS hardware trigger system [11]. It begins with motivation for the gFEX project. It covers the gFEX hardware and infrastructure, as well as a description of the firmware algorithms that run on the board. Finally, it covers the validation of these algorithms using a standalone python framework to make bitwise comparisons between the firmware algorithm outputs and the independent C-based bitwise simulation of the algorithms. It concludes with a brief discussion of gFEX Run 3 performance and looking toward uses for gFEX triggers.

4.1 Motivation for gFEX

Quarks and gluons produced in the proton-proton collisions at the LHC cannot propagate as free particles due to color confinement (see section 2.2.3). Instead, they undergo parton showering and hadronization, producing collimated sprays, or jets, of stable hadrons. In ATLAS, these hadrons deposit energy in the calorimeters. In offline analyses, the energy deposits of these hadrons are clustered together algorithmically to identify jets. The anti-k_T algorithm [45] is the standard jet reconstruction algorithm used by ATLAS because it is infrared and collinear safe, which means that the algorithm can correctly identify jets in an event regardless of the presence of soft particles or splitting of objects along the same line. The algorithm is characterized by a jet radius parameter R, which determines the characteristic angular size of the jet in the $\eta - \phi$ plane. In ATLAS analyses, jets with $R = 0.4$ are referred to as small radius (small-R) jets. These jets are typically designed to contain all the energy from a single quark or gluon. However, when a particle with a large Lorentz boost (high p_T) is produced, its decay products will be collimated in the same direction. In this regime, the decay products merge into a jet with a larger radius. In ATLAS anaylses, a radius of $R = 1.0$ is typically chosen to capture the full energy of such decays. These objects are referred to as large radius (large-R) jets. Offline jet reconstruction in the context of an ATLAS analysis is covered in more detail in 5.3.2.

The Global Feature Extractor (gFEX) of the L1 Calo trigger system is designed to search for energy patterns corresponding to large radius jets ($R \sim 0.9$) as well as global values of interest such as pile up, missing transverse energy, and the sum of transverse energy. Interesting physics events can occur in a boosted regime, where the parent particle had a high enough momentum that its decay products come out collimated. Boosted decays of Higgs, W and Z bosons, top quarks, and even BSM exotic particles can result in a large radius (large-R) jet signature. Providing a reliable L1 trigger for these objects can therefore help us lower p_T thresholds while maintaining a reasonable trigger rate and gain sensitivity to these boosted processes.

While offline jet reconstruction utilizes the anti-k_T algorithm on calibrated calorimeter clusters, gFEX only has on the $\mathcal{O}(100)$ ns to perform its algorithms.

Consequently, gFEX must use a more rudimentary strategy that can be implemented on FPGAs as firmware using coarse granularity information available at L1. This is less refined than the offline choice, but still provides sufficient information to identify boosted hadronic signatures for quick trigger decisions.

4.2 gFEX Hardware and Infrastructure

This section describes the gFEX hardware and surrounding infrastructure necessary to make the system function within ATLAS. The gFEX module runs on one board containing three processor Field Programmable Gate Arrays (FPGAs) and one Zynq FPGA. This section describes these pieces of hardware as well as some of the supporting infrastructure for the gFEX board as a whole.

4.2.1 ATCA Architecture.

gFEX operates on one Advanced Telecommunications Computing Architecture (ATCA) board. This is in contrast to the other L1Calo feature extractors, which have 6 (jFEX) and 24 (eFEX) respectively. There are a total of 3 production gFEX boards, labeled v4B, v4C, and v4D, with the intention of one board being installed and running in the ATLAS TDAQ environment, and the other two being available for testing developments and as possible hot spares in case of malfunction of the operating board. There is also an additional v4A board, which is a partial production board lacking the processor FPGAs of the other three.

The ATCA boards themselves fit into a double slot ATCA shelf. Along with the board, the ATCA shelf consists of a shelf manager, a backplane, cooling fans, and power entry modules (PEMs). The shelf manager has full monitoring and control capabilities of the shelf, allowing for changes of individual modules, fans, and PEMs. The backplane provides access for direct connections, signal routing, and can provide power to individual modules. Each module also has an Intelligent Platform Management Controller (IPMC) [10] which communicates with the shelf manager and helps monitor the status of the board. A diagram of this layout can be seen in Figure 18.

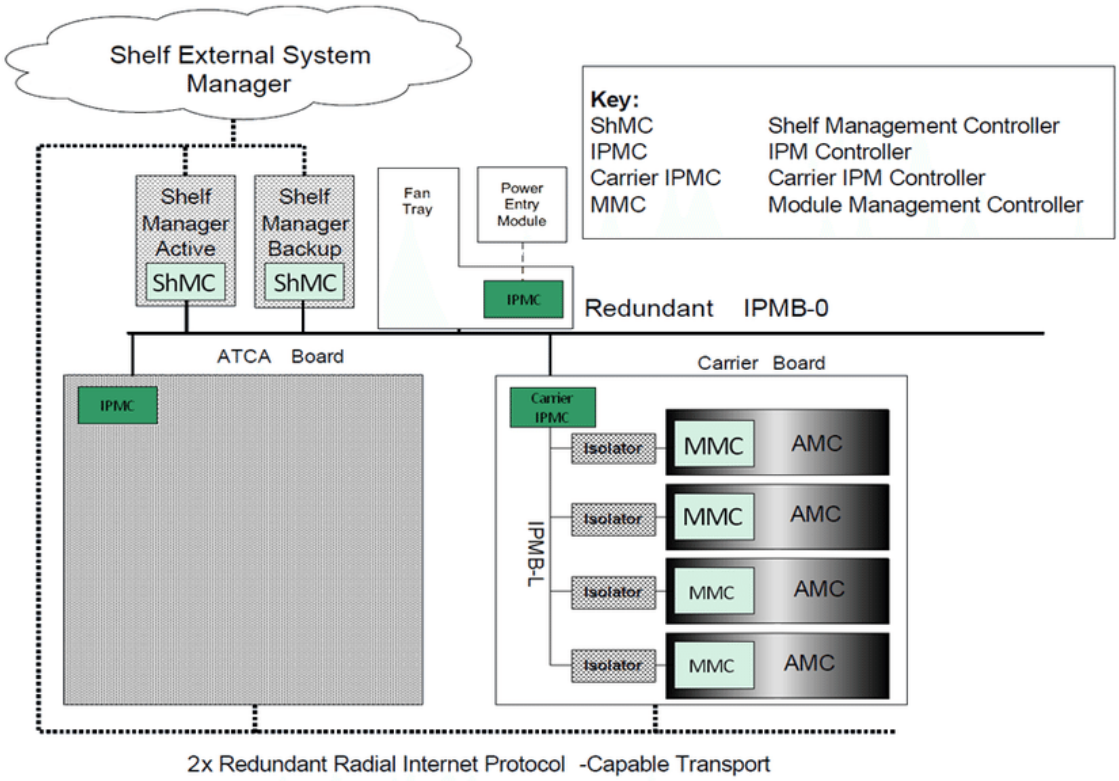


Figure 18. The ATCA shelf setup used for the gFEX board [10].

The gFEX is much coarser than other feature extractors in the phase 1 system (see tower granularity for the FEXes described in section 3.4.1.1); one ATCA board consisting of 4 FPGAs covers the entirety of the detector, allowing for algorithms to calculate global values on an event by event basis. This granularity was chosen in part because the core of most jets fit nicely within a single 0.2×0.2 gTower [46]. Three of the four FPGAs (processor FPGAs labeled A, B, and C) are AMD Virtex^{TM} Ultrascale+^{TM} FPGAs that cover the full eta/phi space of the detector between them, while the final FPGA (labeled Z) is a Zynq+ Multi-Processor System on Chip (MPSoC) that handles monitoring, control, and some global calculations [11]. The gFEX board runs on a custom operating system developed using the Yocto Project [47]. These FPGAs handle the bulk of the data processing and algorithms for gFEX, but there are also many smaller pieces that facilitate data flow as well as pieces that regulate and monitor voltage, current, and temperature of the components.

To receive data from the calorimeters and send data out to L1Topo and FELIX, gFEX has external optical fibers organized into various multi-fiber push on (MPO) connectors. The MPOs used by gFEX are MTP ® connectors, a trademarked MPO from USCONEC [48]. Specifically, eight MTP 48 connectors, one MTP 24 connectors, and one spare MTP 12 connector are used. Each connector has a fiber bundle which routes 12 fibers to individual

Avago^{TM} miniPODs^{TM}, which are high performance fiber optics modules used for parallelization in multi-lane data communications. There are a total of 35 miniPODs on the board, 26 for receiving external signals and 11 for transmitting signals off the board. The miniPODs are connected to Multi Gigabit Transceivers (MGTs) on the FPGAs, which allow for the firmware to decode and send signals to and from the miniPODs. Figure 19 shows these link connections between FPGAs and on/off the board. The MGTs are organized into groups of four referred to as Quads, and serve as the sole path for dataflow through the gFEX system.

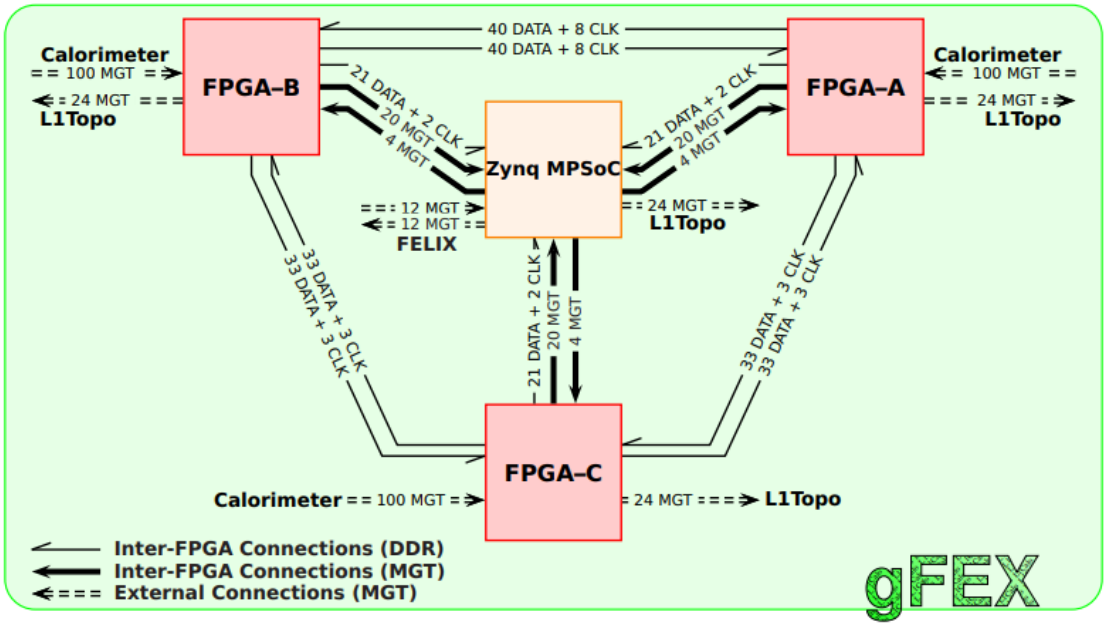


Figure 19. gFEX FPGA link connections [11].

4.3 gFEX Algorithms

The primary gFEX algorithms are described in the following section. The most fundamental of the algorithms is the tower builder, which takes calorimeter inputs and organizes them into nominally 0.2×0.2 "gTowers" in $\eta - \phi$ space (in the more forward region, these towers become larger than $0.2\times$ 0.2). These towers form the basis of the other algorithms. The jet finder algorithm uses these gTowers to calculate all possible "gBlocks," which are 3×3 gTowers, essentially creating small radius jets with a radius of $R = 0.3$. For each gBlock passing a configurable seed threshold, a gJet is calculated, consisting of 69 gTowers, the highest p_T of which becomes the large radius jet candidate for that region. This algorithm also implements a pileup estimation and subtraction from the energy of the jet. gFEX also implements a novel E_T^{miss} algorithm based on Jets without Jets (JwoJ) [49], and additionally has the capability to implement simpler E_T^{miss} algorithms such as

a noise cut approach (NCMET) or ρ plus root mean square (ρ+ RMS). Later in Run 3, special timing triggers were designed and implemented to trigger one or two bunch crossings in the future to aid in Heavy Ion collisions. Algorithms in gFEX are pipelined, such that a given step of an algorithm for one event is happening in parallel with other steps for different events.

4.3.1 gFEX Tower Builder.

The gFEX Tower Builder is the basis for all following gFEX algorithms. Towers are constructed via a complex mapping of the input fibers from the calorimeters to create flat 0.2×0.2 gTowers in $\eta - \phi$ space. This algorithm runs on the three pFPGAs, each of which builds the towers for a subset of the full η range. The gFEX receives Super Cell sums from the LAr Digital Trigger system, which internally are denoted as gCaloTowers, as well as Trigger Towers from TREX. This data is sent to each pFPGA via 100 optical fibers each sending 7 32-bit words. The gCaloTower energies are transmitted as 17-bit signed integer values with a least significant bit (LSB) corresponding to 25 MeV. When gFEX receives this data, it converts it to a 16-bit integer with an LSB of 50 MeV. The gCaloTowers nominally have a 0.1×0.1 spatial resolution in $\eta - \phi$ space, although this becomes coarser in the more forward regions. The TREX Trigger Towers come as a 16-bit unsigned integer with an LSB of 1 GeV and nominally have a 0.2×0.2 spatial resolution in $\eta - \phi$ space. When forming gTowers, there are also subtleties in overlap or extended regions, which result in some gTowers being more complicated to sum. In general, each gTower E_T value consists of: one EM E_T, one hadronic E_T, *possibly* one HEC E_T, which overlaps with Tile, *possibly* one EMEC E_T, which extends $|\eta|$ in $|\eta| \in [2.4, 2.5]$, and *possibly* one HEC E_T, which extends $|\eta|$ in $|\eta| \in [2.4, 2.5]$. This makes the most complicated gTowers have four possible contributions to its energy sum. These are combined to form a grid in each pFPGA of 12×32 gTowers, resulting in 384 gTowers per FPGA and 1152 gTowers total for all of gFEX. In pFPGAs A and B, which cover the central regions of the detector, most of these gTowers are 0.2×0.2 gTowers in $\eta - \phi$ space. In pFPGA C, which covers the forward region, these towers become less granular at larger values of $|\eta|$. A map of the $\eta - \phi$ space covered by the gFEX pFPGAs as well as the divisions of their respective gTowers can be seen in Figure 20.

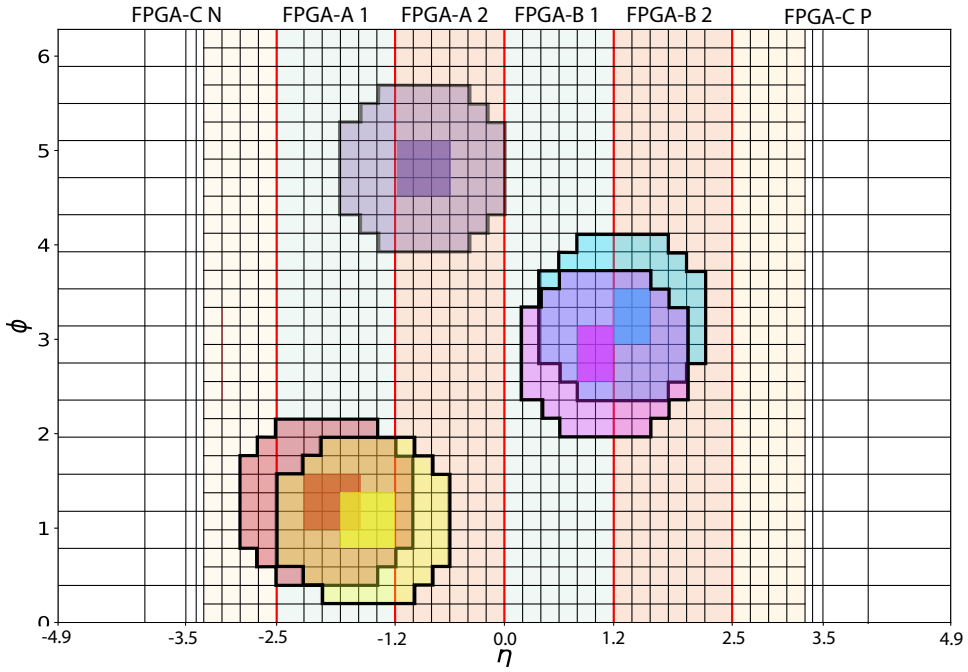


Figure 20. A map of $\eta - \phi$ space covered by the gFEX system. The central regions are made of mostly 0.2×0.2 gTowers corresponding to pFPGAs A and B. The forward regions are covered by pFPGA C, and it should be noted the the gTowers become less granular in the forward region. Overlayed on the map are several hypothetical large-R jet candidates for the trigger algorithm. While a nominal gJet is seeded by a 3×3 gBlock, those centered on the border of a jet engine (such as those on the right side of this image) are only seeded by a 2×3 gBlock.

4.3.2 Jet Finder.

Perhaps the most important algorithm gFEX performs is the Jet Finder, which seeks to identify large-R jet candidates by identifying local energy maxima in the $\eta - \phi$ plane. Each pFPGA is divided into 64 "jet engines" consisting of 6 gTowers. The gTowers all share the same ϕ bin and span six contiguous η bins, effectively splitting each pFPGA in half with two columns of 32 jet engines. They are organized in this way such that a new gTower is added to the relevant sum every clock cycle, where there are 7 clock cycles total (the seventh tick providing a reset). Using the input gTowers, the algorithm first constructs 3×3 gBlocks, which have a radius of $R = 0.3$ assuming a 0.2×0.2 standard gTower. The algorithm calculates all possible gBlocks. The two highest p_T gBlocks per half-pFPGA are output as small-R jet candidates. Once the gBlocks are calculated, gJets are constructed around the center of gBlocks which pass a configurable seed threshold. gJets are a collection of 69 gTowers that form a rough circle of $R = 0.9$ in $\eta - \phi$ space. Each jet engine handles jets whose centers are within its 6 gTower purview, and the jet finder is then responsible for sending the relevant partial sums of these

jet engines to its neighboring jet engines to complete a gJet sum. A diagram with an example gJet calculation and an accompanying clock cycle diagram can be found in Figure 21. It is highly likely that a gJet can span two separate pFPGA regions, and so information on the borders of these regions are shared between the pFPGAs using so called inter-FPGA communication over serial lines. The highest p_T gJet for each half pFPGA is selected as the large-R jet candidate. In this way, pFPGAs A and B provide a total of 4 small-R jet TOBs and 2 large-R jet TOBs to both L1Topo as well the Zynq FPGA to be read out using FELIX. Outputs of pFPGA C are not sent onward, but the FPGA still plays an important role in the inter-FPGA communication for jets in pFPGAs A and B. Additionally, the Jet Finder algorithm has the capability to apply a correction to the energy of a jet to account for pileup. The pileup density per gTower ρ, is estimated as:

$$\rho = \frac{\sum_{\text{towers below threshold}} E_T}{N_{\text{towers below threshold}}} \tag{4.1}$$

and so the energy of the gJet becomes:

$$E_{T,\text{corrected}} = E_T - \rho * 69 \tag{4.2}$$

where 69 corresponds to the number of gTowers in a gJet. In addition to the Jet TOB information, this ρ is sent in one of the otherwise unused 32-bit TOB words as a 16-bit signed integer value. The format in which TOBs are sent to L1Topo for the jet information can be found in Figure 22.

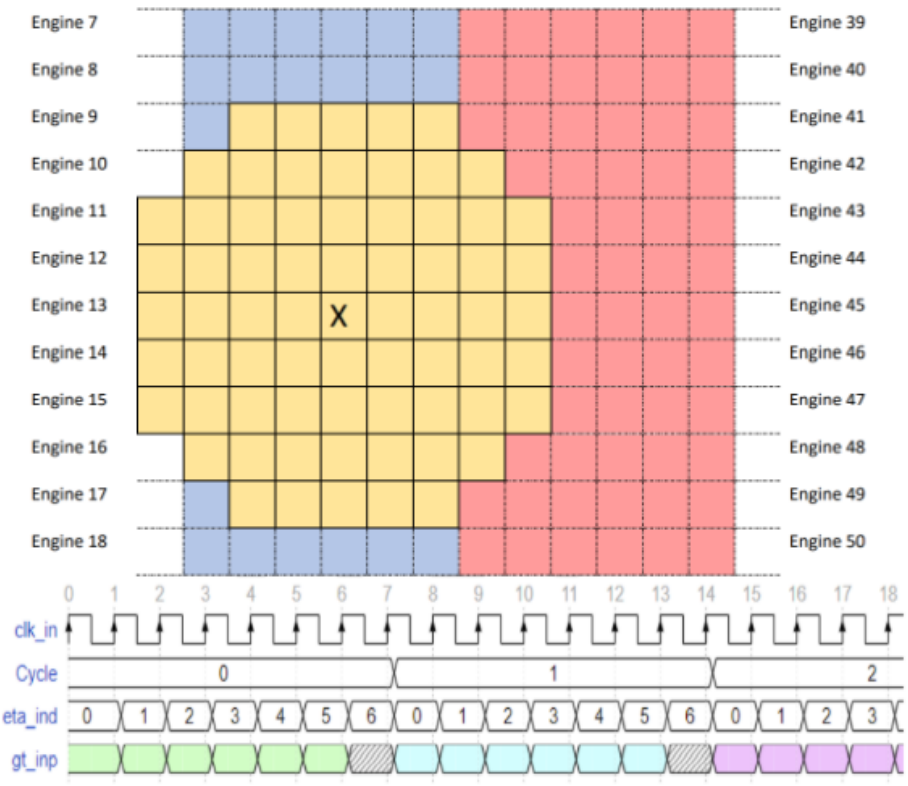


Figure 21. An example gJet calculation for a candidate centered at the gTower marked with an X. Below is the clock cycle diagram for the process to aid in visualizing how the gTowers are summed within the firmware.

Word	Fiber # 0	LSB	Fiber # 1	LSB
0	unused		rho in FPGA	xxx MeV
1	leading gBlock in eta-bins 0-5	xxx MeV	leading gBlock in eta-bins 6-11	xxx MeV
2	second leading gBlock in eta-bins 0-5	xxx MeV	second leading gBlock in eta-bins 6-11	xxx MeV
3	leading large-R jet in eta-bins 0-5	xxx MeV	leading large-R jet in eta-bins 6-11	xxx MeV
4	unused		unused	
5	unused		unused	
6	Trailer		Trailer	

Figure 22. Format for Jet TOBs sent to L1Topo from pFPGAs. Each half FPGA sends its TOBs on one Fiber. The ρ density value is calculated per FPGA and so is sent only once. The LSBs are not listed in the table, but are nominally 200 MeV for these TOB words. Taken from gFEX Data Formats Documentation.

4.3.3 Missing Transverse Energy: Jets without Jets.

gFEX utilizes a novel algorithm to calculate missing transverse energy, based on Jets without Jets (JwoJ) [49]. The main idea is to replace jet clustering with event-shape observables that can mimic the behavior of jet observables

without actually having to explicitly define jets. This is computationally simple and easy to parallelize, making it a good choice for trigger applications. The JwoJ algorithm works by splitting the energy of gTowers that fall above or below a certain threshold into hard (MHT) and soft (MST) terms. This is done at the gBlock level, such that

1. For a given gTower, consider the E_T of the tower as well as the surrounding gBlock E_T
2. Check if the gBlock E_T is above a certain threshold. If it is, then add the gTower E_T to MHT
3. If the gBlock E_T is below the threshold, then add the gTower E_T to MST
4. MET x and y components are constructed from MHT and MST using sine and cosine look-up tables (LUTs) that store these values for each discrete possbility of ϕ
5. Compute the weighted sum of the hard and soft terms: $E_{T,x/y} = a * \text{MHT}_{x/y} + b * \text{MST}_{x/y}$, where a and b are configurable weights that can change for each pFPGA (but do not differ between x and y components)
6. pFPGAs send these components to the Zynq FPGA
7. In the Zynq, the components are summed to compute the full missing transverse energy: $(E_T^{\text{miss}})^2 = E_{T,x}^2 + E_{T,y}^2$
8. The Zynq sends the final E_T^{miss} to L1Topo and to FELIX for readout

Splitting into the hard and soft components in this way allows us to balance contributions from the hard-scatter and pileup for a given event. Near the end of Run 3, updates to the JWOJ algorithm were made to include a similar pileup subtraction scheme as the Jet Finder.

Additional E_T^{miss} algorithms are also considered for gFEX, namely the Noise-Cut MET (NCMET) and the ρ+RMS MET. As of the end of Run 3, the NCMET algorithm has been implemented in gFEX firmware, providing an alternative to the JwoJ approach. ρ+RMS has not yet been implemented due to latency constraints on the needed calculations. All E_T^{miss} algorithms can be used to find the scalar sum of energy, providing $\sum E_T$ triggers as well. The TOB formats sent to both the Zynq FPGA and to L1Topo can be found in Figure 23

Word	Fiber # 0	LSB	Fiber # 1	LSB
0	A+B+C JwJ MHTx | MHTy	xxx MeV	A+B+C ncMETx | ncMETy	xxx MeV
1	A+B+C JwJ MSTx | MSTy	xxx MeV	A+B+C rmsMETx | rmsMETy	xxx MeV
2	A+B+C JwJ METx | METy	xxx MeV	A+B+C ncMET | ncSumET (set to 0)	xxx MeV
3	A+B+C JwJ MET^2 | JwJ SumET	xxx MeV	A+B+C rmsMET | rmsSumET(set to 0)	xxx MeV
4	unused | gRISTRETTO		unused	
5	unused | gESPRESSO		unused	
6	Trailer		Trailer	

Word	Link # 0	Link # 1
0	JwJ pFPGA MHTx | MHTy	ncMETx | ncMETy
1	JwJ pFPGA MSTx | MSTy	rmsMETx | rmsMETy
2	JwJ pFPGA METx | METy	ncSumET | rmsSumET
3	unused | JwJ pFPGA SumET	unused
4	unused | JwJ pFPGA SumHT	unused
5	unused | JwJ pFPGA SumST	unused
6	Trailer	

Figure 23. The TOB formats that the Zynq FPGA sends to L1Topo (left) and the TOB formats that each pFPGS sends to the Zynq (right). The LSBs are not listed in the table, but are nominally 800 MeV for these TOB words. Taken from gFEX Data Formats Documentation.

4.3.4 Special Timing Triggers.

During heavy ion collisions in Run 3, an issue involving trigger shadowing began to affect our data taking efficiency. In heavy ion collisions, there are two extreme types of events, known as Ultra Central Collisions (UCCs) and Ultra Peripheral Collisions (UPCs). UCC events occur when two heavy ions collide essentially head on, creating very large energy events with many decay products. Conversely, UPCs occur with large distances between nuclei, mostly interactions of photons in the respective clouds around the heavy ions, leaving lower energy detector signatures [50]. The extremely busy UCC collisions would cause the low energy UPC triggers to fire one or two bunch crossings early. This trigger would then activate simple dead time of 4 BCs, meaning that the UCC triggers could not fire on the true UCC event. To combat this, the gFEX implemented two new triggers to make use of the mistimed signals: gEspresso, which uses signals from two BCs in the future, and gRistretto, which uses signals from one BC in the future. These triggers were then used as a veto in order to block the low energy UPC triggers from firing and initiating simple dead time. This was possible due to gFEX's multislice TOB readout capabilities.

4.4 gFEX Validation Framework

In order for the Phase-I L1Calo system to officially take over the seeding of HLT trigger chains from the legacy L1Calo system, each subsystem needed to prove that their algorithms provided reliable trigger decisions with performance that matched or improved on legacy performance. As the only feature extractor with no legacy equivalent, gFEX had the unique challenge of proving our algorithms matched expectations for performance without the ability to directly compare to legacy L1 triggers. Validating the baseline gFEX algorithms at the start of Run 3 (namely the Jet Finder and the JwoJ MET algorithms) was my qualification task to become an ATLAS author. To show that the gFEX algorithms were performing as expected, I designed a python based framework that compares the hexadecimal coded output of the FW algorithms to output generated by the standalone bitwise C Simulation using the same input data. This is accomplished by dumping the readout data bytestream of gFEX, which contains both the input data (either real inputs from the calorimeters, or set input data patterns from the testing

setup in the STF) as well as the TOB words produced by the algorithms on the pFPGAs before they are sent to the Zynq FPGA and L1Topo. The input data and output TOBs are parsed apart, and the input data is formatted and sent as input into the bitwise C Simulation. The C Simulation produces its own version of the TOBs according to its version of the algorithm. The two sets of output TOBs are then unpacked from hexadecimal and compared line by line at a bitwise level by the validation framework. The framework generates detailed text file reports categorizing every mismatch so they can be followed up in great detail. Additionally, the framework generates scatter plots and distribution plots that show overall agreement of the events in a given data comparison. This workflow is summarized in Figure 24.

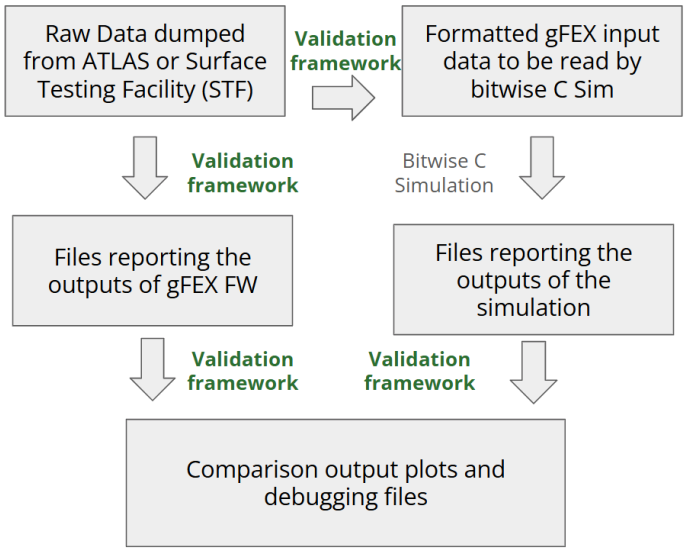


Figure 24. Diagram showing the logical flow of the validation framework. The arrows with green text indicate steps that are completed by the validation framework guided by a shell script.

4.4.1 Obtaining gFEX Data for Comparisons.

Data for gFEX comparison was obtained using a dedicated dumping script that can dump a given number of events from the gFEX readout system while a gFEX board was running, both in live runs or tests occurring at P1 or in tests with preset patterns from the STF [51]. This script strips the gFEX Readout Buffer (ROB) from raw data files and saves them in a text file that can be processed by the comparison framework. These ROB fragments contain header and footer information, the input fiber data for all three pFGPAs, and the output TOBs as calculated by the firmware for the current BC as well as BC+1 and BC+2. This results in 2462 32-bit words per event that need to be properly parsed. An example event written out in full detail can be seen in Appendix A.

4.4.2 The Bitwise C Simulation.

The Bitwise C Simulation, often referred to as the CSim, is a complete simulation of the gFEX algorithms designed to replicate every step of the firmware running on a gFEX board [52]. In the context of the comparison script, the gJet.C file is run as a ROOT macro [53], with a few different input data format options, including already calibrated gTowers, fiber data at the point of input to the tower builder, or fiber data with shared quad mapping. The last option corresponds to the fiber data received from the dumping script that is parsed by the comparison framework. This simulation produces the output TOBs in an easy to read format, with the ability to decode the TOBs to make them human readable and create many diagnostic plots to evaluate performance. The TOB files produced became the standard format used as input for the comparison framework. These TOB outputs are created as a separate text file for each FPGA and both Jet and Global TOBs, meaning in total 6 output TOB files are created: FPGA A Jet TOBs, FPGA B Jet TOBs, FPGA C Jet TOBs, FPGA A Global TOBs, FPGA B Global TOBs, and FPGA C Global TOBs. These files will be used in the comparison to the gFEX ROB Fragment Readout TOBs.

4.4.3 gFEX Compare.

The gFEX compare framework is a python based framework designed to read in text files that contain byte stream readout from the gFEX FW. The compare framework can be divided into two major parts: the parser, which takes in byte stream data from gFEX and formats it properly, and the comparison, which compares output from both gFEX FW and the C Simulation to make direct bitwise comparisons between the two and make detailed logs and summary plots of the comparisons. To run the framework, both it and the bitwise CSim are used as git submodules and called by a lightweight shell script which only requires two simple setup and run commands that point to the desired data file [54].

4.4.3.1 Parser.

The first step of the comparison is to take the byte stream data, which looks like a large string of hexadecimal coded words, and properly parse out the input information, the header information, and the output information. The framework accomplishes this by identifying line breaks which signify different events. Then arrays are formed for all the data, including header information that will not be used in comparisons. These arrays are structured such that each new entry in every array is a new event, and order is preserved from the original data dump to ensure information remains aligned. The run number is parsed from this information and saved to use for input to the CSim, which requires a run number as input to the ROOT macro. The rest of the event information is parsed into relevant arrays and then saved to output text files. The input fiber data is formatted into the shared quads format expected by the CSim and written to a text file identified by the run number. The TOBs are written into their respective files such that a file is created for Jet and Global TOBs for each FPGA, resulting in 6 total files

that have identical format to those that will be produced by the CSim. In some cases, the data of a given fragment could be corrupted or incomplete. To address this without causing the events to become misaligned, any event that is found to have incomplete information is not included in any final write out. Instead, the corrupted or missing fragment information is written to a separated "missing data" file that can be used for debugging. Once the parser and the CSim have been run, there should be 12 total TOB files (6 from the readout and 6 from the CSim) that are saved in respective directories expected by the remaining comparison framework.

4.4.3.2 The Comparison Framework, Reports, and Plots.

With the TOB files generated, the true comparisons now begin. The comparison framework reads in the respective gFEX readout and CSim output TOB files, parses out the TOBs, and translates them into binary strings. An example of this translation and the corresponding physical values of a jet TOB can be found in Figure 25. The strings are then directly compared, and any differences are identified; this is the bitwise level checking of the algorithm output. The comparison framework generates a large amount of output for any number of debugging purposes. The most low-level outputs are the text file debug reports, which take any identified mismatches, and log the mismatch, which word it was found in, and any other identifying information that might be necessary for follow up debugging. These files are incredibly helpful when the reason behind mismatches were non physics related, such as a misaligned word causing a bit shift. In addition to the text based reports, both scatter plots and distribution plots were made to look at one variable, such as pT of the leading large-R jet. The scatter plots display the readout TOB value of the desired item on the x-axis, and the CSim TOB value of the same on the y-axis. In this way, a line with a slope of 1 shows exact agreement between the firmware calculated value and the CSim calculated value, and any differences could be identified and followed up on. The distribution plots were added to show additional information of the physics content of the items being compared. The readout and CSim distributions were displayed on top of each other, so any differences could be inspected visually while also examining the true distribution of the TOB values, which was particularly helpful for determining if the algorithms were producing meaningful values. An example of a scatter plot and distribution plot can be found in Figure 26. The Analog to Digital Converter (ADC) counts, which is the direct translation of the hexadecimal words to base-10 are plotted. At this point, the TOBs have an LSB of 200 MeV, so one can multiply the value by 0.2 to obtain the E_T value in GeV. With this framework, the gFEX algorithms were fully validated. Additionally, this framework was used frequently to validate new firmware builds, even when the changes did not directly touch algorithms, as a way of verifying that the builds did not break anything unintentionally.

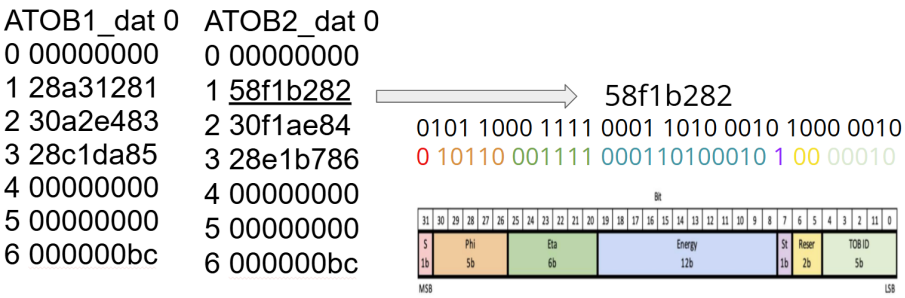


Figure 25. An example of unpacking a hexadecimal gFEX TOB into its components in binary. The validation framework does this conversion to run a full bitwise comparison.

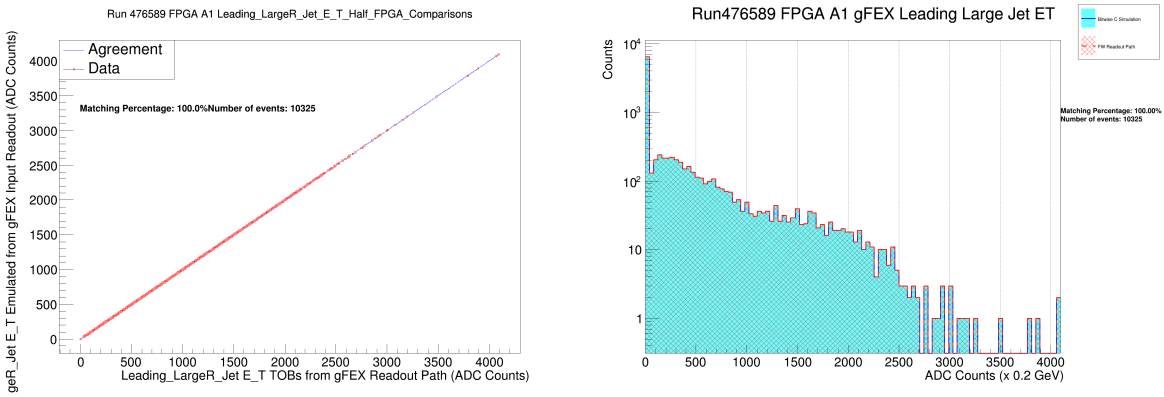


Figure 26. Example output of the comparison framework, showing the scatter plot (left) and distribution comparison plot (right) of the left half of pFGPA A's Leading Large-R Jet E_T. This is data dumped from Run 476589 of ATLAS taken in May of 2024.

4.4.4 The Validation Work in Action: Debugging gFEX Inter-FPGA Communication.

To highlight the importance of this comparison framework to validating and debugging the gFEX algorithms, an example involving the gFEX Inter-FPGA Communication is explained here. In 2023, during Phase-I commissioning, there were consistent mismatches between the readout TOBs and the CSim TOBs for some large-R jet E_T, η, and ϕ values, but exact matching for the gBlocks/small-R jets. The comparison framework revealed that the mismatches in the large-R E_T manifested in vertical bands, meaning that for a given readout E_T value, there were several possible CSim values. These bands can be seen in the plot in Figure 27. We identified that these bands occurred in multiples of 8; this corresponds to the LSB of the data shared across FPGAs during inter-communication. This prompted a search into the data sharing to determine if there could be misalignment. Using

the η and ϕ bins of mismatched events, we back calculated the clock cycle and data line number that the mismatched events would have been transferred on. There was a clear gap in some of the data lines, which can be seen in Figure 28. This was a clear indication of a misalignment of the shared data. This allowed firmware engineers to investigate the timing thoroughly and implement a permanent fix in the firmware.

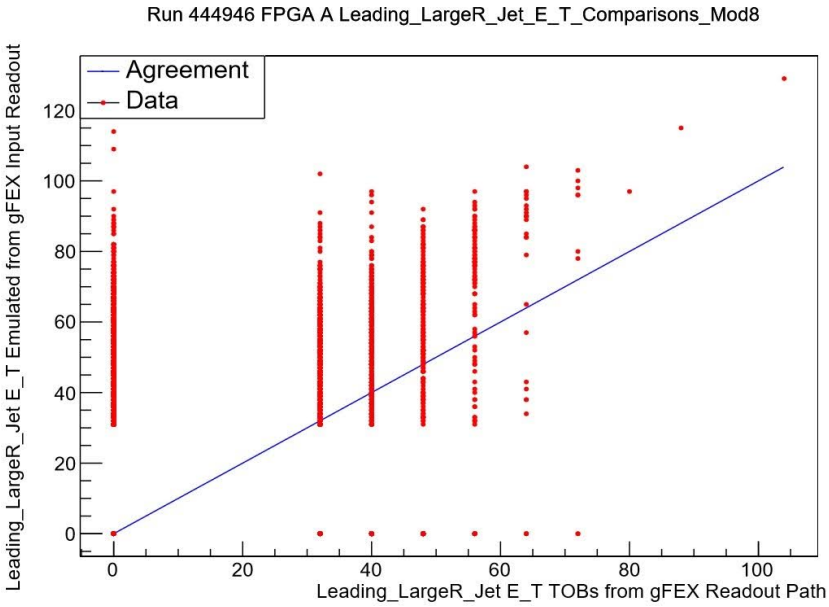


Figure 27. The Large-R Jet E_T comparisons for a run where the FPGA-intercommunication was not working. This plot was specially modified to only show the E_T mod 8.



Figure 28. The back calculated line numbers of data shared between FPGAs B and C. The distribution should be uniform, so the gap was a clear indication of a misalignment in the data sharing.

4.5 gFEX Performance and Usage in Run 3

In early 2024, the Phase-I L1Calo system officially took over the seeding of HLT algorithms from the Legacy system. The gFEX system is performing quite well on its main deliverables. Figure 29 shows both the turn ons for several gFEX Large-R Jet triggers compared to an equivalent legacy small-R jet trigger and the L1Topo Simple Cone Large-R jet algorithm, as well as the turn on curves for the gFEX JwoJ MET trigger along with the jFEX MET algorithm compared to the legacy system before it was decommissioned. Now that we are confident that gFEX works as intended and performs well in the Run 3 environment, we need to decide how to use it. In the next chapter, we will look at the Run 3 Higgs to Displaced Jets Analysis, where we will show that the Emerging Jet Prompt Track Frac Trigger, which is now seeded by 3 Phase-I L1 seeds as of 2024, can aid in the search for new BSM physics.

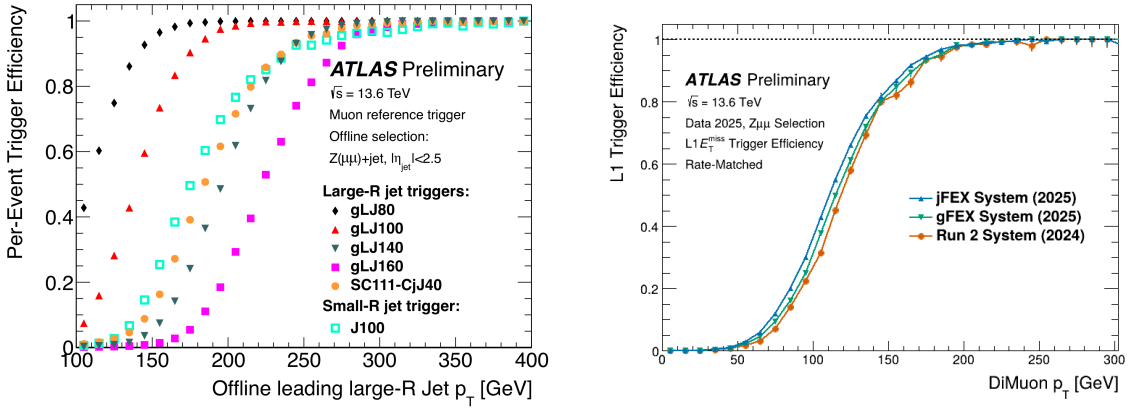


Figure 29. Left: Turn on curves for gFEX Large-R jet triggers compared to Legacy Small-R J100 and the L1Topo Simple Cone algorithm. Note that the lowest unprescaled gFEX trigger, gLJ140, boasts a sharper turn on and faster plateau than its legacy equivalent. Right: Turn on curves for gFEX, jFEX, and legacy MET algorithms, with both gFEX and jFEX showing sharper turn ons than the legacy equivalent [9].

CHAPTER V
ANALYSIS: DISPLACED JETS IN THE ATLAS INNER DETECTOR

At the time of writing this thesis, the Higgs to Displaced Jets Run 3 analysis is still active. This Chapter contains unpublished studies based on work from the Higgs to Displaced Jet Analysis team[1] as part of the ATLAS Collaboration. The analysis motivations will be presented, followed by a description of the Datasets and Monte Carlo (MC) samples used. A description of the analysis strategy, including region and object definitions, trigger strategy, and details on a novel tagger, will be discussed. The bulk of my contributions to this analysis thus far have been various trigger studies in the particularly challenging gluon-gluon fusion (ggF) channel of the analysis, so these will be presented in depth. Preliminary limits of the analysis will be shown. Finally, a brief discussion of the remaining steps will be discussed.

5.1 Analysis Motivation and Introduction

The Standard Model (SM) of particle physics is remarkably successful in predicting the observations of collider experiments. The most notable of which is likely the discovery of the 125 GeV Higgs boson in 2012 by the ATLAS [22] and CMS [23] collaborations. However, many questions remain unanswered by the SM, including the nature of dark matter, the Higgs mass hierarchy problem, and observed matter-antimatter asymmetry. These questions motivate the existence of physics beyond the Standard Model (BSM). The Higgs discovery opened the door to a new chapter of particle physics research that aims to understand both the properties of the Higgs as well as the BSM physics which can answer these open questions. Many BSM theories have been proposed, with several suggesting that the Higgs could be a portal to new physics. In these theories, the exotic decay of the Higgs is a crucial means of discovery.

Traditional BSM searches assume promptly decaying particles, produced at the pp interaction point. However, another compelling possibility is the existence of long-lived particles (LLPs), particles with lifetimes long enough to travel a macroscopic distance in the detector before decaying. LLPs present unique experimental signatures such as displaced vertices and displaced jets. These signatures are strikingly different from signatures produced by prompt decays, necessitating unconventional analysis techniques and specialized triggers to observe them.

LLPs arise naturally in many top-down scenarios via models with small mass splittings or weak couplings. Several theories with these qualities solve the Higgs mass hierarchy problem using the concept of Neutral Naturalness, including Folded SUSY [55], Quirky Little Higgs[56], and some Twin Higgs models [57, 58].

[1] Analysis team members: Jackson Burzynski, Anthony Carroll, Matthias Danninger, Emily Duden, Gabriel Facini, Jesse Farr, Achmad Gerwin, Efstathios Karentzos, Benjamin Kerridge, Andreas Korn, Elena Pompa Pacchi, Despoina Sampsonidou, Ariel Schwartzman, Gabriella Sciolla, David Strom, John Stupak, Dominque Trischuk, Connor Waits, and Melissa Yexley.

In such theories, the top partner states that provide cancellations to the top quadratic correction to the Higgs mass in typical Supersymmetry (SUSY) models [59, 60] does not carry a SM color charge. The production cross section for these uncolored quark partners is significantly reduced, resulting in much weaker existing experimental constraints while remaining within the reach of LHC discovery. In Neutral Naturalness models, a BSM $SU(3)_B$ gauge group matches the usual SM $SU(3)_C$ color gauge group. In the absence of any light BSM charged matter, the lowest energy bound states are glueballs (neutral bound states made of self interacting mediators from $SU(3)_B$, analogous to SM glueballs formed of gluons). These lie on the bottom of the BSM mass spectrum and therefore must decay back into SM fermions. Glueball masses in the range of $\sim 10 - 60$ GeV are theoretically preferred, and could be produced in Higgs decays [61].

Bottom-up scenarios, driven by experimental limits, are also well motivated. Existing constraints on the SM Higgs allow for significant couplings of the Higgs to BSM physics [12]. Notably, the SM Higgs field allows the renormalizable portal interaction $\epsilon|S|^2 H^\dagger H$, coupling the SM to a hidden sector not charged under SM gauge groups. This simple extension to the SM provides a generic example of scenarios where the Higgs boson decays to LLPs and arises naturally in the low-energy limit of more complex models such as 2HDM+a [57] and Folded SUSY [62].

In this analysis, we consider the model agnostic addition of a scalar field S to the SM, where S couples to the Higgs. Given two scalar fields with nonzero vacuum excpectation values $S = v_s + s$ and $H = (v_h + h)/\sqrt{2}$, the mixing angle between the fields is:

$$\sin\theta = \frac{\epsilon v_h v_s}{m_h^2 - m_s^2}, \tag{5.1}$$

allowing the scalar s to decay into SM particles like the Higgs with a decay width decreased by a factor of $\sin^2\theta$. The branching ratios (BRs) for s decays to the SM are shown in Figure 30; since s couples to the SM in the same way as the Higgs, the dominant decay is to a pair of quarks and dependent on the scalar mass, with $4b$ dominating at m_s greater than 10 GeV. Depending on the free parameters of the model, the lifetime of s can be large enough to decay at a large displacement, but small enough to decay within the fiducial volume of the ATLAS detector. This gives rise to the experimental signature of two displaced hadronic jets originating in the ATLAS ID.

Several ATLAS and CMS searches have investigated this benchmark model, setting stringent limits on the branching ratio for $H \rightarrow aa$ at proper lifetimes above 10^{-1} m. In this regime, dedicated triggers designed to select events with LLPs decaying in the calorimeters or in the Muon Spectrometer (MS) drive sensitivity [63, 64, 65, 66, 67, 68]. Probes into signatures with shorter lifetime decays (less than 10^{-2} m) are limited in ATLAS by a lack of triggers to identify LLP signatures originating within the tracking detector. As a result, the previous ATLAS search for Higgs decays to long-lived scalars decaying to jets with these

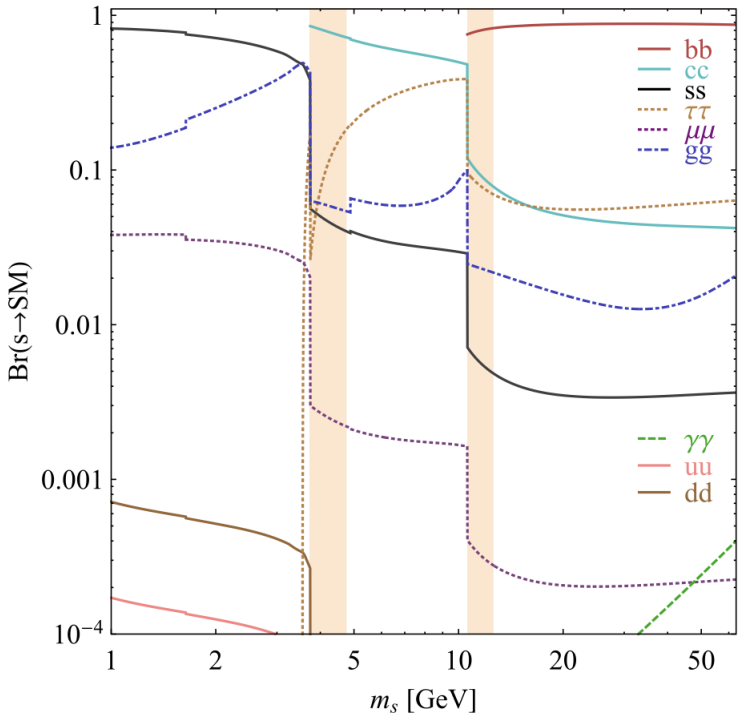


Figure 30. Branching ratios of a single scalar s to the SM as a function of its mass. Image taken from [12].

lifetimes relied on triggers targeting the VH and vector boson fusion (VBF) Higgs production modes, reducing the available statistics [69]. Additionally, these triggers have limited sensitivity to lower LLP masses, restricting the phase space of investigations. In CMS, dedicated displaced jet triggers have been available since LHC Run 2, with improved triggers introduced for Run 3. These triggers facilitate an impressive increased sensitivity across masses up to 55 GeV [70, 71], but limits at low masses and lifetimes are weakened by constraints on large-radius tracking and reconstruction challenges. A summary of previous ATLAS searches across a range of scalar masses is shown in Figure 31.

Several factors contribute to potential increases in sensitivity at proper lifetimes 1 mm $\leq c\tau \leq$ 100 mm. Recent improvements to Large Radius Tracking (LRT) in ATLAS have enabled the introduction of new displaced jet triggers in LHC Run 3 as well as improved offline reconstruction. With these triggers, it is now possible to target Higgs bosons produced with the ggF production mode, which accounts for the majority of the Higgs boson production at the LHC. The existing Emerging Jets trigger can add additional sensitivity to this production mode. In addition to improvements in LRT and the expansion of available production modes, data volume is also significantly increased by the VBF trigger, which in data-taking years 2022-2024 is available for 3× the luminosity of LHC

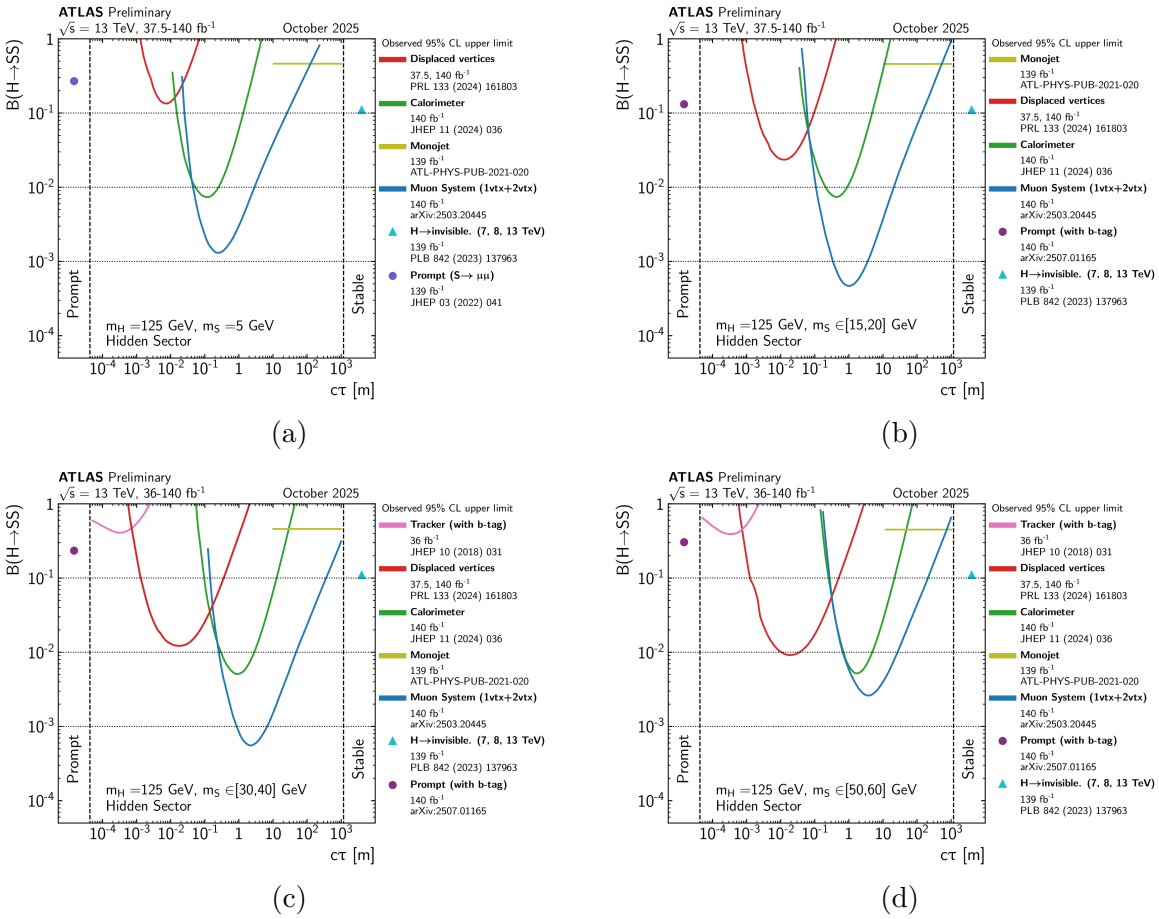


Figure 31. 95% confidence level upper limits on the Higgs boson branching to a pair of long-lived neutral scalars S as a function of $c\tau$, for a Hidden Sector model where (a) $m_S = 5$ GeV, (b) $m_S \in [15, 20]$ GeV, (c) $m_S \in [30, 40]$ GeV, and (d) $m_S \in [50, 60]$ GeV. Plots taken from [13].

Run 2. Finally, the introduction of a Graph Neural Network (GNN) originally used for flavor tagging tasks uses jet and track-level variables to distinguish SM jets from jets produced by BSM LLP decays, providing higher discrimination power than the Boosted Decision Tree (BDT) employed in the Run 2 analysis.

5.2 Datasets and MC Samples

This section describes the datasets and Monte Carlo (MC) generated simulation samples used in the analysis.

5.2.1 Datasets.

This analysis uses partial Run 3 pp collisions at $\sqrt{s} = 13.6\,\text{TeV}$ data collected between 2022–2024 by ATLAS. The analysis exploits two different data streams, depending on the triggers being used. The displaced jet trigger and the emerging jet trigger, used in the gluon-gluon fusion channel, are in the trigger menu of the Physics Main data stream, while the VBF trigger is part of the trigger menu of the VBF Delayed stream. The datasets used for both streams in the data taking years probed in this analysis are listed in Table B.1 in Appendix B.

5.2.1.1 Event cleaning.

Only events recorded during stable beam conditions and with all ATLAS sub-systems fully operational are considered. The list of such events is stored in Good Run Lists (GRLs). Standard event cleaning is performed on events in the GRLs, checking for incomplete events (which could for example occur in the first few Lumi-Blocks (LBs) after a TTC restart), corrupted ones (e.g. Tile corrupted events) or events with LAr noise bursts. The Good Run Lists used in this analysis for the different data taking years are listed, along with their associated luminosities, in Table B in Appendix B. Additionally, only events with a good Primary Vertex (PV) are retained, where the PV must have at least two associated tracks and its $\sum p_{\text{T},i} > 500\,\text{MeV}$, where i loops over the associated tracks to the PV.

5.2.2 Monte Carlo samples.

Signal Monte Carlo samples are used to evaluate the trigger performance and estimate the expected signal acceptance across the parameter space of the benchmark model. Simulated samples were produced centrally by the ATLAS collaboration for several Higgs production mechanisms, LLP masses, and LLP proper lifetimes. In addition, background samples for processes including QCD dijets, $t\bar{t}$, W boson to single charged lepton, Z boson to two lepton, and photon + jet were also produced to define control regions for various studies as well as assess background contamination in preliminary signal regions.

5.2.2.1 Signal samples.

Higgs boson production is simulated using centrally produced Les Houches Event (LHE) files corresponding to the relevant Standard Model production mechanisms, as listed in Table B.3 in Appendix B. The LHE files are then interfaced with `Pythia v8.312` [72] (`AthGeneration 23.6.36`) to simulate the parton showering as well as the Higgs boson decay into two long-lived scalars. A

common control file is included in a per-process JO, through the name of which the details of the relevant process are extracted (Higgs production mode and LLP mass, mean proper lifetime and decay mode). The common control file includes generation filters relevant to different processes.

Two MC requests were performed: an R&D request with only two mass points (16 GeV and 55 GeV) and the final production where the full mass grid was requested. The first MC request was used to inform the second one in the choice of the generator filters to be used, as explained in the following. The list of samples can be found in Table 2 and Table 3.

m_{LLP} [GeV]	$c\tau$ [mm]	Final state	Campaign	Production mode	Filter
16	1,10,100	4b	a,d	ggF/VBF	Y
55	1,10,100	4b	a,d	ggF/VBF	Y
16	10	4b	a,d	ggF/VBF	N
55	10	4b	a,d	ggF/VBF	N

Table 2. R&D signal samples. The first two lines show the filtered samples that were generated, while the last two lines show the unfiltered samples, which have double the events of the filtered ones to allow for sufficient statistics after the relevant triggers are applied.

For the first MC production, $m_{\text{LLP=16 GeV}}$ and $m_{\text{LLP=55 GeV}}$ were considered, assuming $BR(\text{LLP} \rightarrow b\bar{b}) = 100\%$. The ggF and VBF production modes are considered and only a and d campaigns corresponding to data taking years 2022 and 2023 were generated, with the same statistics to reflect the luminosity of the two campaigns. Lifetimes of 1, 10 and 100 mm were considered for the filtered samples while 10 mm was used for the unfiltered samples.

m_{LLP} [GeV]	$c\tau$ [mm]	Final state / BR	Campaign	Production	Filter
3	1,10,100	4s	a,d,e	ggF/VBF	Y
5	1,10,100	4c	a,d,e	ggF/VBF	Y
10	1,10,100	4b	a,d,e	ggF/VBF	Y
15	1,10,100	4u/4c/4b/4g/4τ	a,d,e	ggF/VBF	Y
20	1,10,100	4b	a,d,e	ggF/VBF	Y
30	1,10,100	4b	a,d,e	ggF/VBF	Y
40	1,10,100	4b	a,d,e	ggF/VBF	Y
55	10	4b	e	ttH	Y
55	10	4b	e	WmH/WpH/ ggZllH/qqZllH	Y
55	1,10,100	4b	e	ggF/VBF	Y

Table 3. Signal grid extension.

In the second MC production, additional mass points were added, going down to 3 GeV. Different final states were produced depending on the LLP mass. Multiple final states were generated for the 15 GeV to assess the analysis sensitivity for different final states at the same m_{LLP}. The ggF and VBF production modes are considered for most of the samples, while ttH, W^+H, W^-H, $gg \to Z(ll)H$ and $q\bar{q} \to Z(ll)H$ are considered for $m_{\text{LLP}} = 55\,\text{GeV}$ and $c\tau = 10\,\text{mm}$ (the most sensitive signal points in the previous analysis) for signal contamination studies in Control Regions. Here a, d and e campaigns were generated corresponding with data taking years 2022, 2023, and 2024, with relative statistics to reflect the luminosity of the three campaigns. Lifetimes of 1, 10 and 100 mm were considered for all the samples, besides, as said, the ones used for signal contamination studies.

5.3 Object Reconstruction and Definitions

This section describes the reconstruction and definition of physics objects used in this analysis. It will describe the tracks used, both standard tracks as well as large radius tracks (LRT), small-R jets and large-R jets that are relevant for the analysis signal. It will also describe the objects used for control regions, namely muons, electrons, and photons.

5.3.1 Tracks and Tracking.

This section describes the tracks used in this analysis. Tracks are nominally the most important physics objects for this search because we use them to define a jet as displaced or not. In general, tracks are used to construct essentially every other physics object we use, so they are truly fundamental. There are two types of tracks used in this analysis: standard ATLAS tracks (STD) as well as special large radius tracks (LRT).

In this analysis, both standard tracks and LRT tracks are used. Tracks in ATLAS are represented with a dedicated parametrization called the perigee representation, defined as $\mathbf{q} = (d_0, z_0, \phi, \theta, q/p)$. The variables d_0 and z_0 describe the local position of the track with respect to the beam spot, where d_0 is the transverse impact parameter and z_0 is the longitudinal impact parameter. The variables ϕ and θ describe the track's direction of propagation, where ϕ is the transverse angle and θ is longitudinal angle. The final variable q/p is the signed charge of total momentum, and can describe the track's curvature. The uncertainties of these measurements and their correlations are represented by a covariance matrix:

$$\begin{pmatrix} \sigma^2(d_0) & cov(d_0, z_0) & cov(d_0, \phi) & cov(d_0, \theta) & cov(d_0, q/p) \\ . & \sigma^2(z_0) & cov(z_0, \phi) & cov(z_0, \theta) & cov(z_0, q/p) \\ . & . & \sigma^2(\phi) & cov(\phi, \theta) & cov(\phi, q/p) \\ . & . & . & \sigma^2(\theta) & cov(\theta, q/p) \\ . & . & . & . & \sigma^2(q/p) \end{pmatrix}$$

In general, the processes for finding tracks in ATLAS follow a set order. The first step is pattern recognition: identifying seeds and tracks. Seeds are formed from

groups of 3 detector hits, and are passed on to the track finding stage if they pass quality cuts including a minimum 500 MeV p_T threshold, a maximum d_0,$z_0 \lesssim 5$ mm cut. These seeds are made into search roads along the rough expected trajectory. Then in the track finding step, tracks are extended from seeds along this search road using a Combinatorial Kalman Filter (CKF). Track parameters are continually updated as the filter propagates. After the final layer of the tracker is reached the initial fitting process is done. Quality checks are made on the tracks including at least 8 silicon hits, $d_0 < 5$ mm, no more than one Pixel "hole" or two combined Pixel+SCT "holes" where a track passes through a layer with no hit. Then the tracks undergo some more cleaning including ambiguity resolution, where track momentum, number of holds and hits, and shared modules are used to compute a quality score, where a low score track is likely to be a combinatorial fake. Finally, a global χ^2 is performed to obtain the final high precision track parameter estimate.

Track reconstruction in ATLAS happens in a few "passes" of ordered Athena algorithms [43]. The first pass is an "inside-out" pass where seeds from the Pixel Detector and SCT are extended outward to form the majority of tracks. The second pass is an "outside-in" pass which uses the remaining unused hits extending inward from the TRT seeds in the vicinity of calorimeter regions of interest (ROIs). Tertiary passes are for special use cases, such as LRT, which uses relaxed requirements on $|d_0|$ and $|z_0|$ to provide an opportunity for reconstructing displaced tracks. Specifically, $|d_0|$ is extended from a < 5 mm cut to < 300 mm, and z_0 is extended from a < 200 mm to a < 500 mm. This comes with some more stringent cuts, such as a minimum p_T requirement of 1 GeV as opposed to 500 MeV, fewer allowed holes, and seeding only done in the SCT, to reduce the computational load from the increased combinatorics.

For the dedicated Displaced Jet triggers in this analysis, LRT tracks are defined as prompt if they meet $|d_0| < 3$ mm and $|sin(\theta) * dz_0| \leq 3$, where $dz_0 = z_0 + v_z - z_{PV}$ is the difference of the track from the primary vertex on the z-axis. The requirements on tracks for the Emerging Jets trigger will be described in a later section, but it utilizes standard tracks to define a fraction of prompt to total track momentum within a large-R jet.

5.3.2 Jet Reconstruction.

When quarks or gluons are produced in a collision, they immediately hadronize due to color confinement, creating sprays of hadrons that we reconstruct as "jets." In this way, jets are the physical observable corresponding to QCD processes. Hadrons within the jets interact with nuclei in the detector material, as described in section 3.3.3.2.

Jet reconstruction in ATLAS uses a topological clustering algorithm to build topo-clusters, which are formed based on the ratio of a calorimeter cell to the surrounding noise. The clusters are calibrated at the EM energy scale, which corresponds to the energy deposited by electromagentic showers, but underestimates hadrons. The clusters are fed into an anti-k$_T$ algorithm with

$R = 0.4$ to form small-R Jets [45]. The jets are then calibrated and the energy is corrected to the Jet Energy Scale (JES) to account for the calorimeter response to hadrons. ATLAS additionally uses a particle flow (PFlow) algorithm which provides an improvement on the accuracy of charged hadron measurements by using inner tracker momenta measurements instead of calorimeter energy deposits [73]. The PFlow jets are now standard to use in ATLAS analyses, however, this analysis will still use the traditional EMTopo jets for the chosen small-R Jet collection because PFlow jet collections containing LRT Tracks are not currently available in our derived analysis objects, and indeed the PFlow jets may perform worse with LRT due to decreased efficiency. For the large-R jets, uncalibrated small-R EMTopo jets are reclustered according to an anti-k$_T$ algorithm now with a larger radius cone of $R = 1.0$. These Reclustered Topo Jets (RCTopo jets) are calibrated by calibrating their constituent small-R jets and summing the result. This collection was chosen to synergize well with the Emerging Jet trigger. For the purposes of this analysis, small-R jets in the central region have a p_T minimum cut of 20 GeV and an η cut of $|\eta| < 2.5$. Large-R jets are chosen to have a minimum p_T cut of 200 GeV and have a center that falls within $|\eta| < 1.8$. The p_T cut is chosen to ensure that the particle decaying is boosted enough to ensure that all its decay products fall within one large-R jet, and the $|\eta|$ cut is chosen to ensure that the outer edge of the jet still remains within the central region of the detector.

5.3.3 Leptons and Photons.

Muons, electrons, and photons are used in this analysis to form control regions for various studies. Muons are reconstructed from tracks in the MS and ID, and are very clean due to their unique signature. Electrons and photons are reconstructed via EMTopo clusters as described in the jet selection, but of course do not need the JES calibration to correct for their energy. Table 4 describes the cuts made on these objects in the analysis

Cut	Muon	Electron	Photon
p_T	> 20 GeV	> 20 GeV	> 20 GeV
$\|\eta\|$	< 2.5	$< 2.5, \notin [1.37, 1.52]$	$< 2.5, \notin [1.37, 1.52]$
Identification	Medium ID	Loose ID	Loose ID
Isolation	PFlowLoose_VarRad isolation	FixedCut Loose isolation	PFlowLoose_VarRad isolation

Table 4. Cuts made on physics objects used to define control regions.

5.4 Analysis Triggers

This section describes the triggers used to define the signal regions in the analysis. The table of the full trigger chains and their purpose can be found in Appendix C. The following subsections describe these trigger algorithms and selections in depth.

5.4.1 Emerging Jets Trigger.

The Emerging Jet Prompt Track Fraction (EJ PTF, or just EJ) Trigger was designed to identify events with a large radius jet that has a low fraction of prompt track pT compared to the total pT of the tracks (prompt track fraction, or PTF). It is the dedicated trigger for the Emerging Jets Analysis [18], but it also warrants study in this analysis because it should do an excellent job of identifying displaced jets, which inherently have a low PTF value. The HLT trigger chain for this trigger is HLT_j200_0eta180_emergingPTF0p08dR1p2_a10sd_cssk_pf_jes_ftf_preselj200. In 2022 and 2023, it was seeded by a legacy small-R jet trigger at L1: L1J100. In 2024 and onward, it is seeded by three new Phase-I L1 seeds: a gFEX large-R jet with restricted η range: L1gLJ140p0ETA25, a jFEX small-R jet: L1jJ160, and the L1Topo Simple Cone (SC) algorithm which draws a large-R jet centered on a small-R jet seed from jFEX: L1SC111-CjJ40. This trigger was introduced at the beginning of the Run 3 to target an emerging jet signature, therefore allowing a lower p_T threshold compared to the standard single jet triggers [44]. The EJ trigger selects events containing at least one reconstructed large-R Particle Flow (PFlow) jet with constituent subtracted soft killer (CSSK) and SoftDrop procedures applied at the trigger level satisfying $p_T > 200$ GeV, $|\eta| < 1.8$, and PTF < 0.08. The PTF is defined as

$$PTF = \frac{\sum_{\text{trk}\subseteq\Delta R<1.2} p_T^{\text{trk}}(|d_0| < 2.5\sigma_{d_0,\text{bkg}}(p_T^{\text{trk}}))}{p_T^{\text{jet}}} \tag{5.2}$$

where tracks in the sum satisfy

1. $p_T > 1$ GeV
2. $\Delta z = |z_{\text{HS PV}} - z_0| < 10$ mm
3. $d_0 < 2.5\sigma_{d_0,\text{bkg}}(p_T^{\text{trk}})$

Here, Δz is distance between the z position of the Hard Scatter Primary Vertex and the track z_0 value, and $\sigma_{d_0,\text{bkg}}(p_T^{\text{trk}})$ is the significance of the track's d_0 value. The threshold of $2.5\sigma_{d_0,\text{bkg}}(p_T^{\text{trk}})$ was determined using background SM jet MC samples at the time of trigger creation.

5.4.2 Displaced Jet Triggers.

The inclusion of LRT at the HLT level in Run 3 allows for the development of triggers that target jets with a high multiplicity of displaced tracks (classified by a large impact parameter). Some such new triggers are the dedicated displaced jet (DJ) triggers. The two DJ trigger chains studied in this analysis are HLT_j180_2dispjet50_3d2p ("symmetric" trigger) and HLT_j180_dispjet50_3d2p_dispjet50_1p ("asymmetric" trigger). In 2022 and 2023, these were seeded by the legacy L1J100 trigger at the hardware level. In 2024, when the Phase-I L1 triggers began seeding HLT algorithms, this changed to the jFEX L1jJ160 trigger, which serves as the equivalent to the legacy. The two trigger

chains require one jet (prompt or displaced) with $p_{\mathrm{T}} > 180$ GeV and two displaced jets with $p_{\mathrm{T}} > 50$ GeV, where a jet is defined as displaced depending on the multiplicity of prompt and displaced tracks associated with it.

A prompt track has to satisfy:

$$|d_0| < 3 \text{ mm and } |\sin(\theta) \times \Delta z_0| \leq 3, \tag{5.3}$$

where d_0 is the transverse impact parameter of the track and

$$\Delta z_0 = z_0 + v_z - z_{\mathrm{PV}}, \tag{5.4}$$

where z_0 is the z position of the track perigee, v_z is the reference point for track reconstruction (namely, the beamspot position) and z_{PV} is the z position of the primary vertex. The tracks associated to a jet that are considered to be prompt are the ones reconstructed with the standard tracking algorithm and satisfy the above requirements. It should be noted that this is a different definition of prompt than that of the EJ trigger. A displaced track has to satisfy

$$|d_0| > 3 \text{ mm and } d_{0\,\text{significance}} > 10, \tag{5.5}$$

where $d_{0\,\text{significance}}$ is the significance of d_0, taking into account the uncertainty of track parameters and the uncertainty of beamspot parameters. The tracks associated to a jet that are considered to be displaced are the ones reconstructed with the LRT algorithm and satisfy the above requirements.

The multiplicity of prompt and displaced tracks is then used to define a displaced jet (DJ) based on the specific trigger. The HLT_j180_2dispjet50_3d2p symmetric trigger defines a DJ as having a maximum of two prompt tracks (2p) and a minimum of three displaced tracks (3d), as specified in the 3d2p displaced jet requirement. Whereas the HLT_j180_dispjet50_3d2p_dispjet50_1p asymmetric trigger requires two DJs which have to satisfy two different requirements, one is having a maximum of two prompt tracks and a minimum of three displaced tracks (3d2p), and the other is having a maximum of one prompt track with no requirement on displaced track multiplicity (1p).

5.4.3 VBF Trigger.

The Vector Boson Fusion (VBF) trigger is a Run 3 trigger designed to identify events with a VBF topology, specifically the VBF production of the Higgs that is characterized by two energetic forward jets. This topology includes the kinematics of a forward jet pair, namely: the jet pair invariant mass M_{jj}, the leading forward jet p_T, the subleading forward jet p_T, as well as η and ϕ cuts on the forward jet pair. The HLT trigger chain for this trigger is HLT_j70_j50a_j0_DJMASS1000j50dphi200x400deta with the legacy L1 seed of L1MJJ_500_NFF for data taking years 2022 and 2023 and the Phase-I L1 seed of L1jMJJ_500_NFF for data taking year 2024 and onward. Table 5 shows the cuts made by the HLT trigger algorithm. Since M_{JJ}, involves combinations of two jets, it is calculated in L1Topo. The NFF (No forward forward) requirement helps reduce combinatorics by requiring that the leading jet of the pair not be in the FCal. Still with these requirements, the trigger rate is too high for this algorithm

to run the prompt reconstruction in the physics main stream, so this trigger makes use of the delayed physics stream.

Leading Jet	$p_T > 70$ GeV, $\eta < 3.2$
Subleading Jet	$p_T > 50$ Gev, $\eta < 4.9$
Forward Jet Pair Invariant Mass	$M_{jj} > 1000$ GeV
$\Delta\eta_{jj}$ and $\Delta\phi_{jj}$	$\Delta\eta < 4$, $\Delta\phi < 2$

Table 5. Online requirements for the VBF trigger.

5.5 Analysis Strategy

Signal Regions (SRs) for the analysis are split into two channels: one targeting ggF Higgs production, and the other targeting VBF Higgs production. Each channel is further split into two signal regions: resolved and boosted. The Higgs production mode regions are defined using offline trigger preselections as well as a Recursive Neural Net (RNN) tagger designed to identify VBF topologies. Within each proudction region, two versions of the GN2 based Displaced Jet Tagger, one trained for small-R jets and one trained for large-R jets, are used to define the resolved and boosted topology regions, respectively. The analysis also utilizes several control regions for use in physics studies. This section will describe the selections made to define the regions for the analysis.

5.5.1 The VBF Channel.

The VBF channel targets VBF production mode of the Higgs. The region is defined by two main discriminators: the VBF trigger with associated offline selection cuts, and a score cut on an RNN VBF tagger. The first discriminator in determining the VBF region is the passing of the VBF trigger, which also requires offline selections to ensure good event quality of the events which passed the trigger online. Studies using a range of cuts on the forward dijet mass M_{jj}, leading forward jet p_T, and subleading forward jet p_T were conducted to find suitable cuts on the VBF trigger variables to achieve a good trigger efficiency. The cuts chosen are summarized in the following Table 6.

VBF Variable	Cut
M_{jj}	> 1200 GeV
Leading Forward Jet p_T cut	> 100 GeV
Subleading Forward Jet p_T	> 90 GeV

Table 6. Selections for VBF Trigger offline requirements.

The second discriminator is a simple RNN designed to give high scores to events with VBF topologies signified by two forward jets. This additional handle on the region is helpful in ensuring that the VBF and ggF regions remain orthogonal.

This is needed because the VBF trigger runs in the Delayed Physics data stream, as opposed to the ggF region triggers, which run in physics main. In principle, these can have overlapping events, but cannot be made orthogonal by requiring opposing triggers to fail due to the different trigger menus. The tagger score of > 0.7 was chosen to define the VBF signal region, as it provided good separation of the VBF signal from background.

5.5.2 The ggF Channel.

The ggF channel targets the ggF production mode of the Higgs. This region is defined completely by the trigger and corresponding offline selections. In the case of only including the DJ trigger bucket, this is decided entirely by the asymmetric DJ trigger and its offline selections. These offline selections are summarized in Table 7. It should be noted that while the online trigger requires jets to have at least 3 displaced tracks and at most 2 prompt tracks (3d2p), differences between online and offline track reconstruction (namely, the lack of ambiguity solving for LRT online) resulted in the loosening of this requirement offline to 2d2p. The 1p requirement (at most one prompt track) remains consistent between online and offline. The symmetric DJ trigger is omitted from the signal region definition (see section 5.7.1.1).

Asymmetric DJ Variable	Cut
Number of jets with 2d2p	≥ 2
p_T of jets with 2d2p	> 70 GeV
Number of jets with 1p	> 100 GeV
p_T of jets with 1p	> 70 GeV
Any additional jet p_T	> 200 GeV

Table 7. Selections for DJ Trigger offline requirements.

In the case of including the EJ trigger in the SR via a trigger bucketing strategy (see 5.7.2), the following offline selections in Table 8 are applied to the EJ bucket. This ensures that at least one large-R jet in the event has both the p_T and PTF requirements of the trigger, even if it is not the highest p_T jet in the event. These were taken from the on plateau requirements of the Run 3 Emerging Jets Analysis [18] in order to maintain consistency in the use of the EJ trigger. The EJ analysis additionally included an off plateau cut of 250 GeV for the large-R jet p_T requirement, but this results in a variable scale factor that would be difficult to harmonize with the other aspects of the Displaced Jets Analysis. It was verified during the scale factor determination that these cuts still provided good efficiency performance (see 5.6).

EJ Variable	Cut
Number of RCTopo jets	≥ 1
RCTopo jet p_T	> 300 GeV
RCTopo jet PTF	< 0.04

Table 8. Selections for EJ Trigger offline requirements.

5.5.3 Boosted vs Resolved Regimes.

Within each Higgs production region, we split further into two different topologies: the resolved case, where the LLPs are separated enough that their decays are well separated from each other, resulting in two small-R displaced jets, and the boosted case: where a high momentum Higgs results in collimated LLPs, and their decays can be identified within one large-R jet. These displaced jets are identified using a per-jet Graph Neural Network (GNN) based on the ATLAS GN2 FTAG algorithm [74]. The tagger's primary task is to label jets as displaced or as QCD background after being trained on MC simulation samples. It takes two jet-level variables (p_T and η) as well as 15 track-level variables as input for training. To ensure that jets were not evaluated in the same network that they were used to train, jets are split by even and odd event number into two separate taggers, and only evaluated on the tagger they did not train. For each tagger, 80% of jets are used for training and the other 20% is used for validation. There are two version of this tagger, one trained for small-R jets, and the other trained for large-R jets. The trigger studies in the following sections cite the DJ tagger score thresholds used to originally define regions, where a small-R jet is identified as having a score of > 0.999, and a large-R jet is identified as having a score of > 0.99999. These scores are very high, and can sometimes run into problems with machine precision (high scores being rounded to 1), so the analysis is currently moving toward a logit score difference to define the displaced jets. The logit scores are the raw, un-normalized scores of the last layer in the GNN. They serve as the inputs to get the output scores that range from 0 to 1. The differences between the LLP jet logit and the QCD background jet logit are unbounded numbers, so we do not run into any constraints on a range like we do with the constrained sigmoid score. Currently a displaced large-R jet is defined as having a logit difference of 16, and a displaced small-R jet is defined as having a logit difference of 6.9. The boosted region within each Higgs production region requires $\geq$ 1 large-R DJ, and the resolved region requires $\geq$ 2 small-R DJs. To maintain orthogonality, the ggF boosted region requires < 2 small-R DJs, and the VBF resolved requires < 1 large-R DJ.

Table 9 shows the full breakdown of the SRs in the analysis.

Selection	ggF Boosted	ggF Resolved	VBF Boosted	VBF Resolved
Trigger	Asymm DJ (EJ)	Asymm DJ (EJ)	VBF	VBF
Central Jets	Asymm DJ (EJ) presels	Asymm DJ (EJ) presels	- -	- -
Forward Jets	- -	- -	VBF presels	VBF presels
Displaced Jets	$\geq$ 1 large-R DJ && $<$ 2 small-R DJs	$\geq$ 2 small-R DJs	$\geq$ 1 large-R DJ	$\geq$ 2 small-R DJs && $<$ 1 large-R DJ
VBF RNN Score	-	-	0.7	0.7

Table 9. Signal Regions for the analysis as they currently stand.

5.6 Determination of the EJ 2024 Scale Factor

To account for differences between MC simulation and actual data taking conditions, we determine a flat multiplicative scale factor (SF) by taking the ratio of the Emerging Jet trigger efficiency in data and MC in a defined control region. Since this analysis is using data from 2022-2024, the SF determined by the partial Run 3 Emerging Jets Analysis [18] for 2022 and 2023 can be copied. However, no one has yet used the EJ trigger for an analysis involving 2024 data, so the SF must be calculated here. Preselection cuts on the two relevant variables: transverse momentum of the leading RCTopo large R Jet, and the prompt track fraction of the leading RCTopo large-R jet, are made so that the scale factor is determined on the plateau of the trigger efficiency curves. These studies are conducted in a single muon control region (CR_μ). CR_μ is defined using the lowest unprescaled single muon trigger and having exactly one muon reconstructed satisfying medium identification and tight isolation requirements. The tight isolation requirement is chosen so that the offline requirement is tighter than the onlime trigger requirement. Additionally, the isolation will ensure that the muon is not produced within a jet, reducing correlation between the muon trigger with analysis jet triggers. The cuts and selections made can be found in Table 10. The CRμ is enriched of $W(\mu\nu)$ + jets and $t\bar{t}$ single-lepton events, and hence those relevant MC samples are considered for this study.

Single μ trigger	HLT_mu24_ivarmedium_L1MU14FCH
μ identification	medium
μ isolation	PflowTight_VarRad
Number of μ	1
μ selection	$p_{\mathrm{T}} > 27$ GeV, $\lvert\eta\rvert < 2.5$

Table 10. CR_{μ} selections.

In addition to the CR_{μ} requirements, additional requirements are placed on the selection criteria for large-R jets to ensure good object selections. We require one large-R RCTopo Jet in the region, with a p_T of at least 200 GeV, with the center of the jet falling within $|\eta| < 1.8$. This η requirement ensures that the edge of the RCTopo Jet (which has an effective radius of $R \sim 1.2$) only goes to a maximum of $|\eta| = 2.4$, which is the limit for LRT tracks. Finally, the RCTopo Jet must pass photon overlap removal to make sure any isolated photons within the jet radius do not contribute to the jet energy. These additional requirements are summarized in Table 11.

Number of RCTopo Jets	≥ 1
Leading RCTopo Jet p_T	≥ 200 GeV
Leading RCTopo Jet $\lvert\eta\rvert$	≤ 1.8
Passes γ OR	true

Table 11. RCTopo Jet selections for CR_{μ}.

The efficiency curves used to determine the scale factor are parameterized in terms of variables, the leading RCTopo Jet p_T, and the leading RCTopo Jet PTF. These variables correspond to the large-R jet p_T and large-R jet PTF requirements for the EJ trigger. The efficiency for each variable is expressed as

$$\epsilon_{var_{1/2}} = \frac{\text{Events passing CR}\mu\text{ AND passing var}_{2/1}\text{ cut AND passing EJ trigger}}{\text{Events passing CR}\mu\text{ AND passing var}_{2/1}\text{ cut}} \quad (5.6)$$

The selections for both variables are summarized in Table 12. These cuts were made to match the on-plateau requirements of the Emerging Jets Run 3 Analysis in their EJ Trigger scale factor determinations for 2022 and 2023. The efficiencies are calculated for both MC (using mc23e samples) and Data24. The ratio of the data and MC efficiencies is fit with a flat line to test the hypothesis of a flat SF. The efficiency plots for the EJ trigger can be found in Figure 32. For the EJ trigger decision, an OR of all three possible L1 seeds for the trigger are taken. Figures 33, 34, and 35 show the resulting scale factor efficiency plots for the gFEX, jFEX, and L1Topo SC seeds, respectively. No major discrepancies between the L1 seeds are observed, so the OR of all three is suitable to use for the overall scale factor. Furthermore, the scale factors determined from the two different parameterizations

are compatible within error, so we take their average and sum their errors in quadrature to achieve the final scale factor for the EJ trigger in 2024. Table 13 shows the scale factors for the leading jet p_T and PTF as well as their average, and Table 14 lists the EJ scale factors for 2022, 2023, and 2024 to be used in this analysis.

	Leading RCTopo Jet p_{T}	Leading RCTopo Jet PTF
Baseline selection	CRμ_{DJ}	
Leading RCTopo Jet p_T cut	-	> 300 GeV
Leading RCTopo Jet PTF	< 0.04 GeV	-

Table 12. Selections for each variable's efficiency plots.

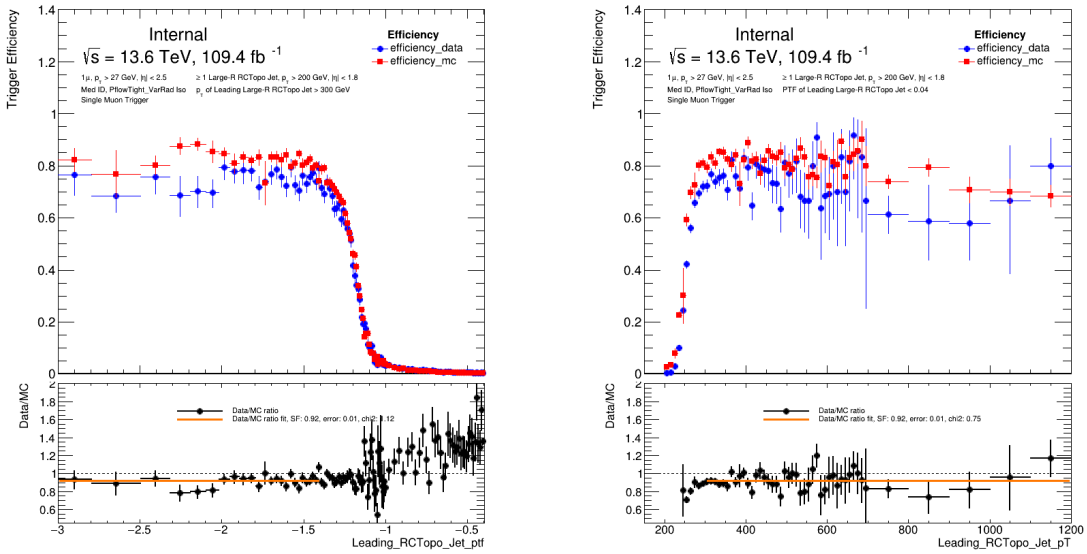


Figure 32. The efficiency plots used to determine the EJ trigger SF for 2024, with the efficiency parameterized by the Leading RCTopo Jet p_T (left) and the Leading RCTopo Jet PTF (right). The EJ trigger decision for this plot takes the logical OR of all three possible L1 seeds to the EJ trigger.

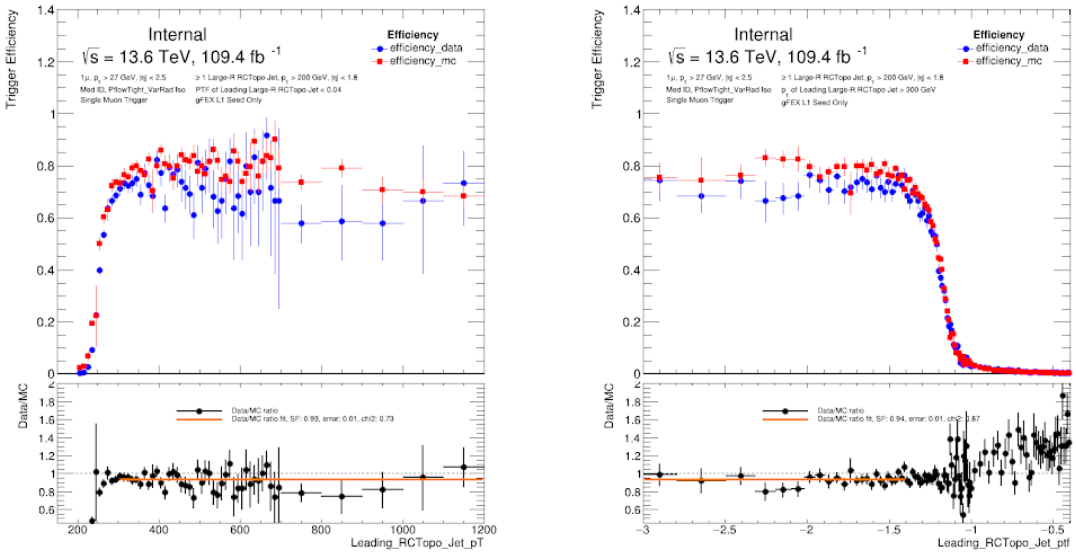


Figure 33. The efficiency plots used to determine the EJ trigger SF for 2024, with the efficiency parameterized by the Leading RCTopo Jet p_T (left) and the Leading RCTopo Jet PTF (right). The EJ trigger decision for this plot uses only the gFEX large-R jet for its L1 seed.

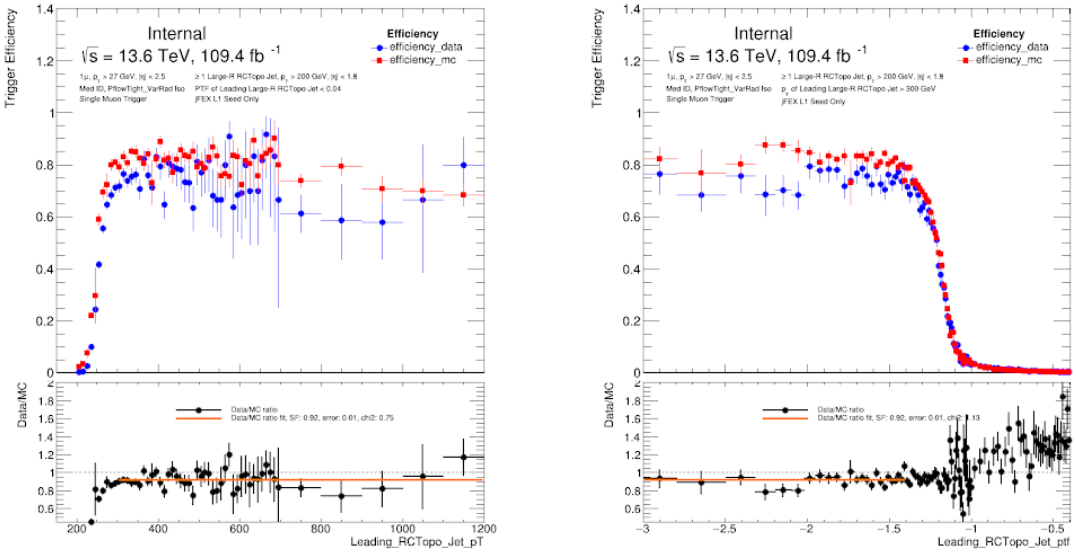


Figure 34. The efficiency plots used to determine the EJ trigger SF for 2024, with the efficiency parameterized by the Leading RCTopo Jet p_T (left) and the Leading RCTopo Jet PTF (right). The EJ trigger decision for this plot uses only the jFEX small-R jet for its L1 seed.

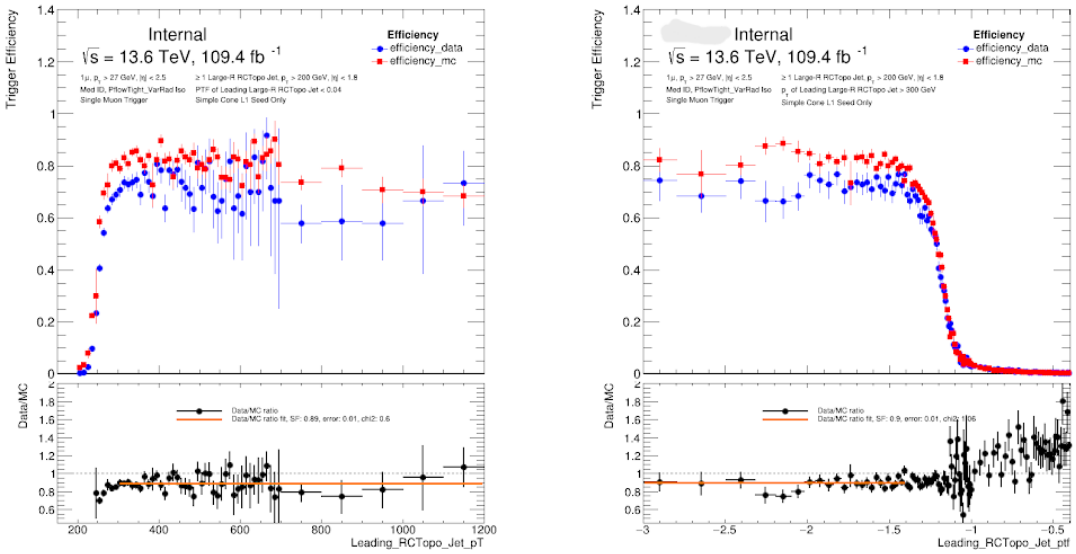


Figure 35. The efficiency plots used to determine the EJ trigger SF for 2024, with the efficiency parameterized by the Leading RCTopo Jet p_T (left) and the Leading RCTopo Jet PTF (right). The EJ trigger decision for this plot uses only the L1Topo Simple Cone algorithm for its L1 seed.

EJ Trigger 2024 SF	gFEX L1 Seed	jFEX L1 Seed	SC L1 Seed	OR of all L1 Seeds
Leading RCTopo Jet p_T	0.93 ± 0.01	0.92 ± 0.01	0.89 ± 0.01	0.92 ± 0.01
Leading RCTopo Jet PTF	0.94 ± 0.01	0.92 ± 0.01	0.90 ± 0.01	0.92 ± 0.01
Average SF	0.94 ± 0.01	0.92 ± 0.01	0.90 ± 0.01	0.92 ± 0.02

Table 13. SFs for EJ Trigger 2024 split by L1 seed.

EJ Trigger Year	2022	2023	2024
SF	0.96±0.01	0.96±0.01	0.92±0.02

Table 14. SFs for the EJ trigger for the relevant data taking years for the analysis. Years 2022 and 2023 are taken from the on plateau values considered in the Emerging Jets Run 3 Analysis [18].

5.7 Trigger Studies to Asses the Utility of the EJ trigger

In order to help define our trigger strategy, the efficiencies of our triggers in signal Monte Carlo (MC) are studied and compared. The emerging jet and

displaced jet triggers are studied using signal samples where the Higgs is produced by gluon-gluon fusion, and the VBF trigger is studied using signal samples where the Higgs is produced by vector boson fusion. A list of all samples with their DSIDs can be found in Appendix B. The workflow for these studies is as follows: analysis ntuples are centrally produced from DAOD_LLP1 derivations using our analysis framework [75], which is based off of EasyJet[76]. Information relevant to the trigger efficiency studies are extracted into histograms via a custom Displaced Jet Frame (DJFrame) in the FastFrames framework from TopCPToolkit [77]. The resulting histograms are then processed by dedicated python scripts to create trigger efficiency tables and plots, signal sensitivity plots, and scale factor plots.

5.7.1 Signal Efficiency Trigger Studies.

Inclusive trigger efficiency studies on signal MC samples were performed to determine how to best use the triggers in this analysis. Selections on the signal MC samples of different trigger selections both with and without preselection cuts were made to determine the effect of overall signal efficiency from these cuts. Additionally, comparisons between triggers were made by examining their relative exclusive trigger efficiency on signal samples, This allowed us to determine if there was significant overlap between any of the triggers. These studies were completely carried out using mc23a and mc23d samples, corresponding to data taking years 2022 and 2023. This is because these studies were conducted near the beginning of the analysis timeline, and so 2024 production samples were not yet available.

5.7.1.1 DJ Symmetric and Asymmetric Trigger comparisons.

One of the first trigger efficiency studies I conducted in this analysis was examining the overlap between the two DJ triggers, dubbed the symmetric and the asymmetric. To do this, I made simple trigger selections with no additional cuts on the full mass and lifetime grid of ggF production mode signal samples to find the overall trigger efficiency for the following categories: passing the DJ symmetric trigger, passing the DJ symmetric trigger and failing the DJ asymmetric trigger, passing the DJ asymmetric trigger, and passing the DJ asymmetric trigger and failing the DJ symmetric trigger. In this way, we can see how many of the events passing one trigger can be "caught" by the other. The results can be found in Table D.2 in Appendix D. From the results, it is clear that the DJ symmetric trigger offers very little unique contribution to the overall trigger efficiency. This provided motivation to leave the symmetric trigger out of signal region definitions and only rely on the asymmetric trigger. This has the benefit of simplifying things such as the DJ trigger scale factor as well as background estimation in the future. To ensure that the very little rate provided by the symmetric trigger did not account for a specific region of LLP kinematic phase space, distributions of the calculated lifetime of the signal sample LLPs were made for the different trigger selections. These distributions can be found in Figure 36. While the exclusive plot has too low statistics to determine any feature, the inclusive distributions look very

similar, suggesting that we can leave out the symmetric DJ trigger without fear of ignoring important corners of phase space.

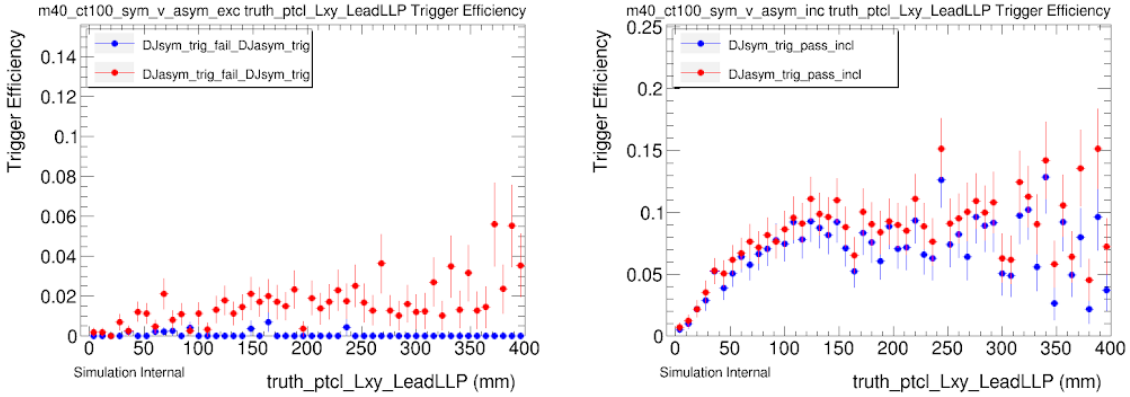


Figure 36. Distributions of the LLP decay length L_{xy} for different selections of the DJ symmetric and asymmetric trigger. The left plot shows the exclusive distributions, where the trigger not selected is required to fail. The right plot shows the inclusive distributions, which look largely identical. This provides justification for moving forward with the asymmetric DJ trigger as the only DJ trigger to consider for the signal region.

5.7.1.2 VBF Preselection Studies.

Another study I conducted was to determine the effect of different p_T cuts of the forward jets in the VBF trigger preselection to help determine what to use. Similar to the previous study, the different trigger selections were made looking at the full mass and lifetime grid of the signal samples, but this time using the VBF production mode samples. Two different preselection cuts were tested: one with the leading forward jet at a p_T of 110 GeV and the subleading at 90 GeV, and another with the leading forward jet at 100 GeV and the subleading at 90. Table D.3 in Appendix D shows the results. Ultimately, the differences are small, but there was enough marginal benefit to the looser p_T cut to proceed with 100 GeV for the VBF preselection.

5.7.1.3 EJ and DJ Trigger Complimentarity.

The most important goal of these inclusive trigger efficiency studies was to justify the use of the emerging jet trigger in this analysis. In order to be worth including, the EJ trigger has to contribute uniquely to the trigger efficiency in the ggF region, since the VBF trigger is by far the most sensitive to the VBF signal samples. The first step in assessing the EJ trigger's performance is to simply look at the unique events captured by the EJ trigger as compared to the asymmetric DJ trigger. The trigger selections were run over the full mass and lifetime grid for the ggF production mode signal samples, with the following selections: EJ trigger pass with DJ asymmetric failing, EJ trigger and preselection cuts pass with DJ

asymmetric failing, DJ asymmetric trigger pass with EJ failing, and DJ asymmetric trigger and preselection cuts pass with EJ failing. This gives us a complete picture of how useful the EJ trigger can be. Table D.1 in Appendix D shows the full detail of the outputs. In general, it can be seen that the EJ trigger does a particularly good job in the low lifetime samples where the DJ trigger struggles to capture any of the signal. This is incredibly promising that the EJ trigger can add much needed sensitivity to the part of phase space this analysis is trying to probe. Performance on longer lifetime samples are also comparable to the DJ performance. Figure 37 shows summary plots of this study's findings. It is clear through this study that EJ can be beneficial in gaining more signal, but we also need to verify that it does not also let in a disproportionate amount of background events into the signal region, so we must continue with more detailed studies.

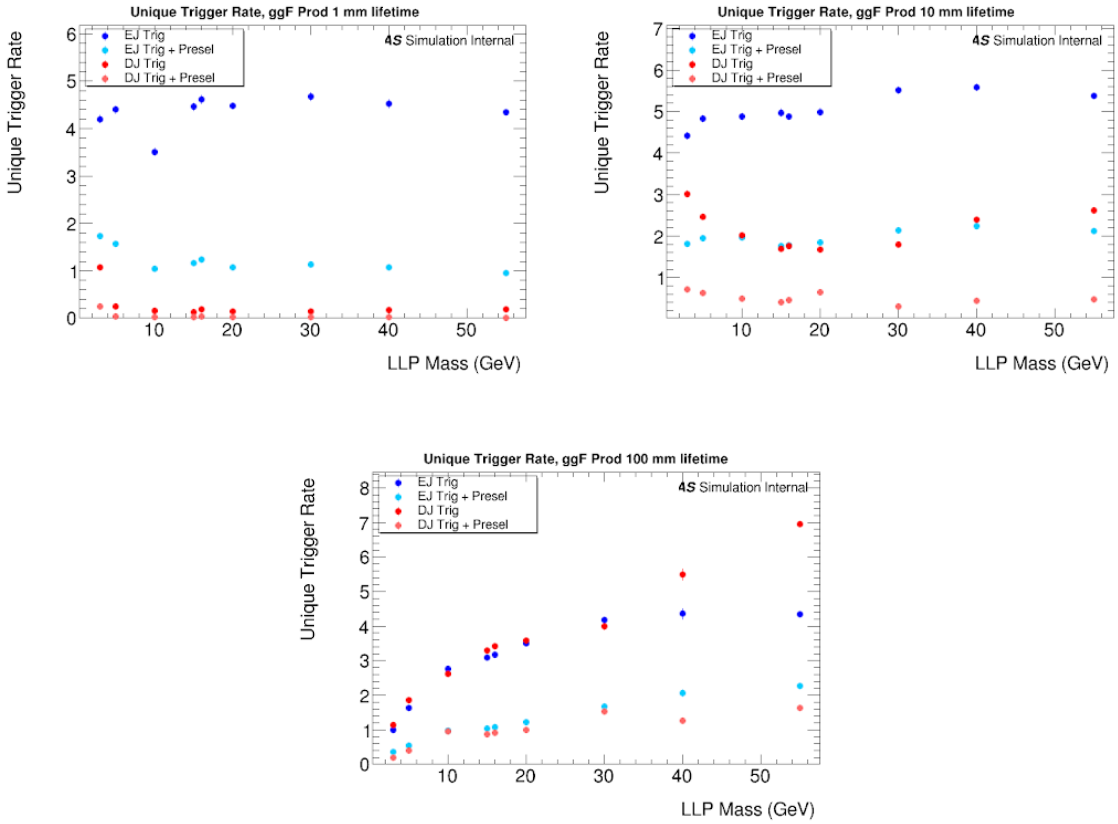


Figure 37. Summary plots of the inclusive trigger efficiency on ggF production mode signal samples for the EJ and DJ triggers, with both preselections applied and left out to gain a complete picture. The EJ clearly shows an increase in performance for the low lifetime samples, suggesting including it in the analysis can be greatly beneficial to gaining additional sensitivity in that region of phase space.

5.7.2 Defining Trigger Strategy for ggF Channel - Trigger Buckets.

In the last section, we showed that the emerging jet trigger can capture signal that the displaced jet trigger misses, but we still need to verify that including this trigger does not introduce a disproportionately large amount of background events into the signal region or overcomplicate the background estimation. The DJ and EJ triggers target very different topologies, so simply combining them into one SR by taking the or of their trigger decisions can cause complications in the preselection requirements, scale factor determinations, or even introduce double counting of events if not careful. Therefore, to apply both triggers to the same region, we define different *trigger buckets*, where the DJ trigger and its preselections define one bucket, and the EJ trigger and its preselections define the other. To ensure orthogonality of the buckets, they are given an ordered priority, where in general, "Bucket 1" is checked first, and if the preselections for Bucket 1 are not met, then the preselections for "Bucket 2" are checked. In the studies conducted, both possibilities of ordering DJ or EJ were first attempted. We observed negligible difference between the two strategies, so for ease of other analysis workflows, the DJ trigger was given priority, as it is considered to be the default trigger for the ggF region. This strategy has the additional benefit of allowing us to use the individual trigger efficiency scale factors.

To asses the feasibility of this trigger bucket strategy, we created signal significance plots by taking the number of events passing a given trigger bucket for a ggF production signal MC sample and dividing it by the square root of the number of events passing that trigger bucket for background MC samples, where we have scaled the samples appropriately by luminosity and generator weight factors. The signal samples additionally have an extra factor to account for more realistic coupling of the Higgs to LLPs, as the branching ratio is assumed to be 100% for the purposes of signal generation. This is chosen to be 12% (the branching ratio of Higgs to undetected) or the best limit that exists for a given sample, whichever is more stringent. For these studies mc23a and mc23d samples are used in order to cross check initial studies done for the ggF region before the inclusion of the trigger bucketing strategy. The MC samples used are $t\bar{t}$, QCD dijet, and $\gamma+$ jets samples in order to get a good approximation of the multijet background that could contaminate the signal region. The results are shown in cutflow diagrams for the different buckets, with the final several cuts in any cutflow being dedicated to increasing DJ tagger score thresholds, requiring two small-R jets above a given score for the resolved scenario, and one small-R jet above a given score for the boosted scenario. This was done to aid in finding the appropriate threshold for the tagger scores, which are used for region definitions, which maximize our signal sensitivity and reject the most background. Figures 38, 39, 40, and 41 show these cutflows for the DJ Bucket in the boosted scenario, the EJ Bucket in the boosted scenario, the DJ Bucket in the resolved scenario, and the EJ bucket in the resolved

scenario, respectively. By comparing the DJ and EJ buckets, we can see that EJ continues to be extremely beneficial to the overall significance, especially in the shorter lifetime LLP samples. The EJ bucket performs noticeably worse than the DJ only in the 100 mm lifetime samples in the boosted regime, but still boasts significance values between 1 and 10 for the the different signal samples. With these studies, we have definitively shown that the analysis strategy benefits from including the EJ trigger by adding sensitivity to regions missed by the DJ trigger without introducing excess background events, which will allow us to place very stringent limits on the branching ratio of Higgs to LLPs. The next step of this study is to determine how to estimate the background in the ggF signal regions, which has proven difficult both with and without the inclusion of the EJ trigger. Appendix E contains the associated Signal Contamination (S/B) and Event Yield plots that were generated using the same cutflows and selections.

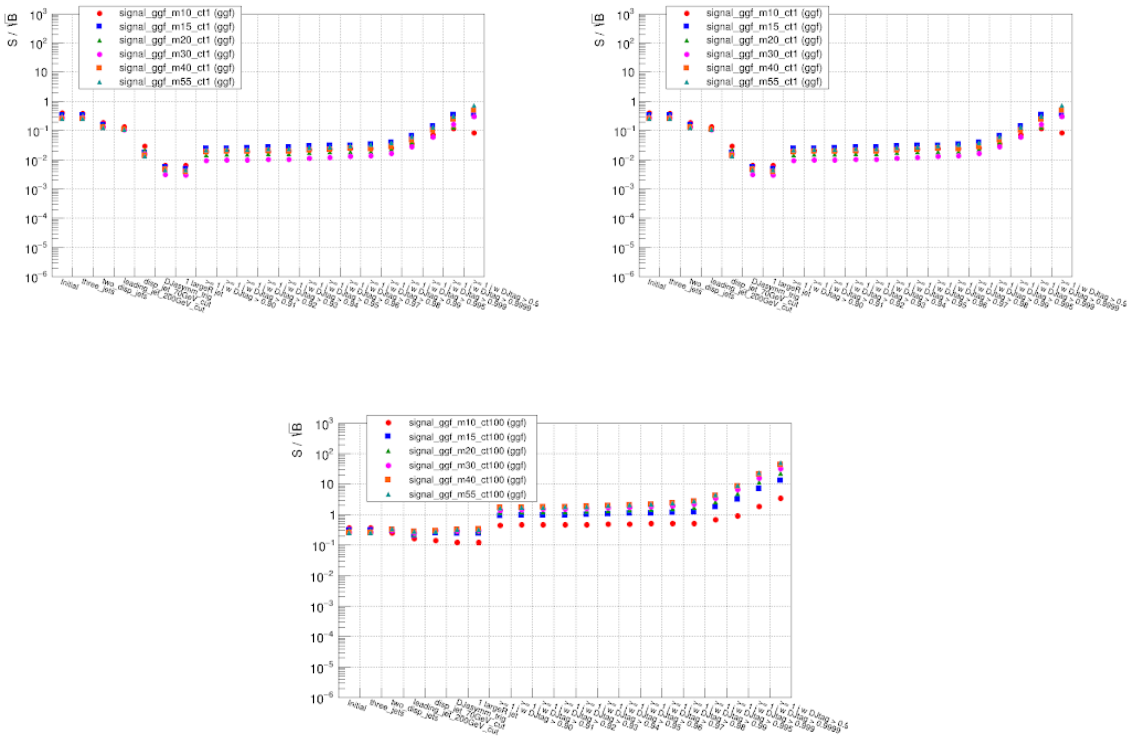


Figure 38. The significance plots for the DJ trigger bucket in the boosted region for signal samples organized by lifetime, with 1 mm samples on the top left, 10 mm samples on the top right, and 100 mm samples on the bottom. The large-R DJ tagger score extends to a maximum value of 0.99999, which was found to maximize the analysis sensitivity while maintaining high signal acceptance.

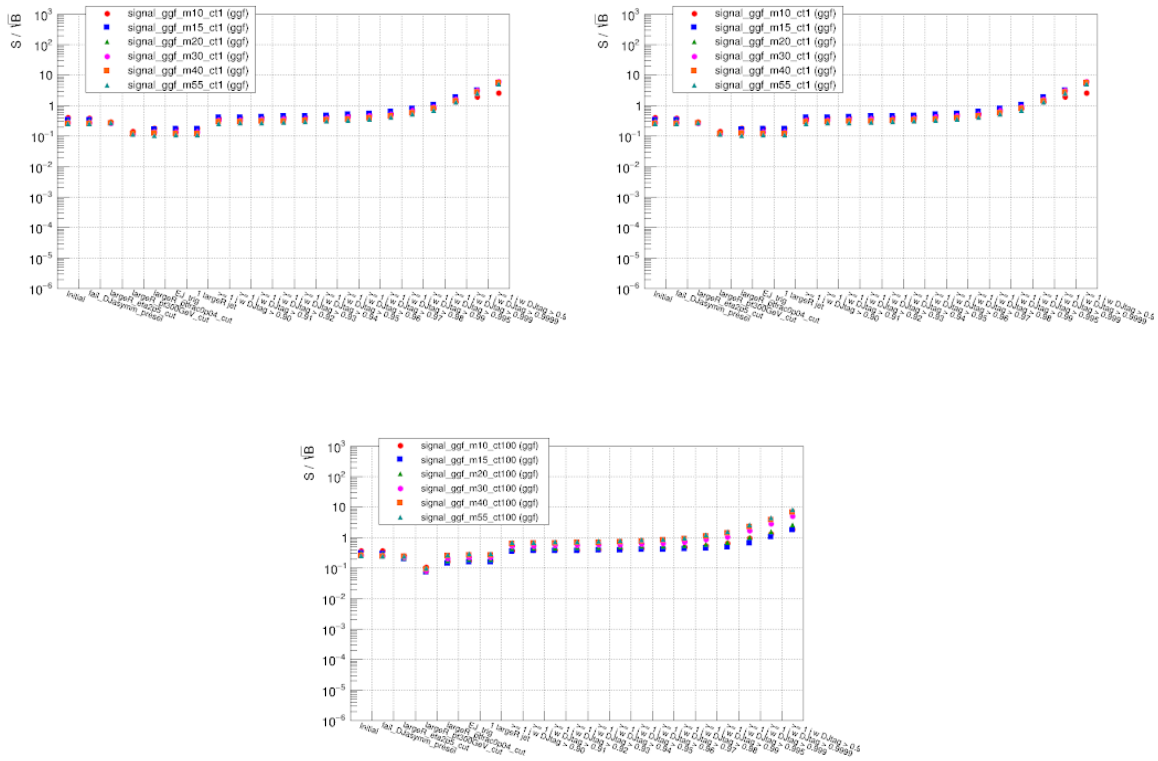


Figure 39. The significance plots for the EJ trigger bucket in the boosted region for signal samples organized by lifetime, with 1 mm samples on the top left, 10 mm samples on the top right, and 100 mm samples on the bottom. The large-R DJ tagger score extends to a maximum value of 0.99999, which was found to maximize the analysis sensitivity while maintaining high signal acceptance.

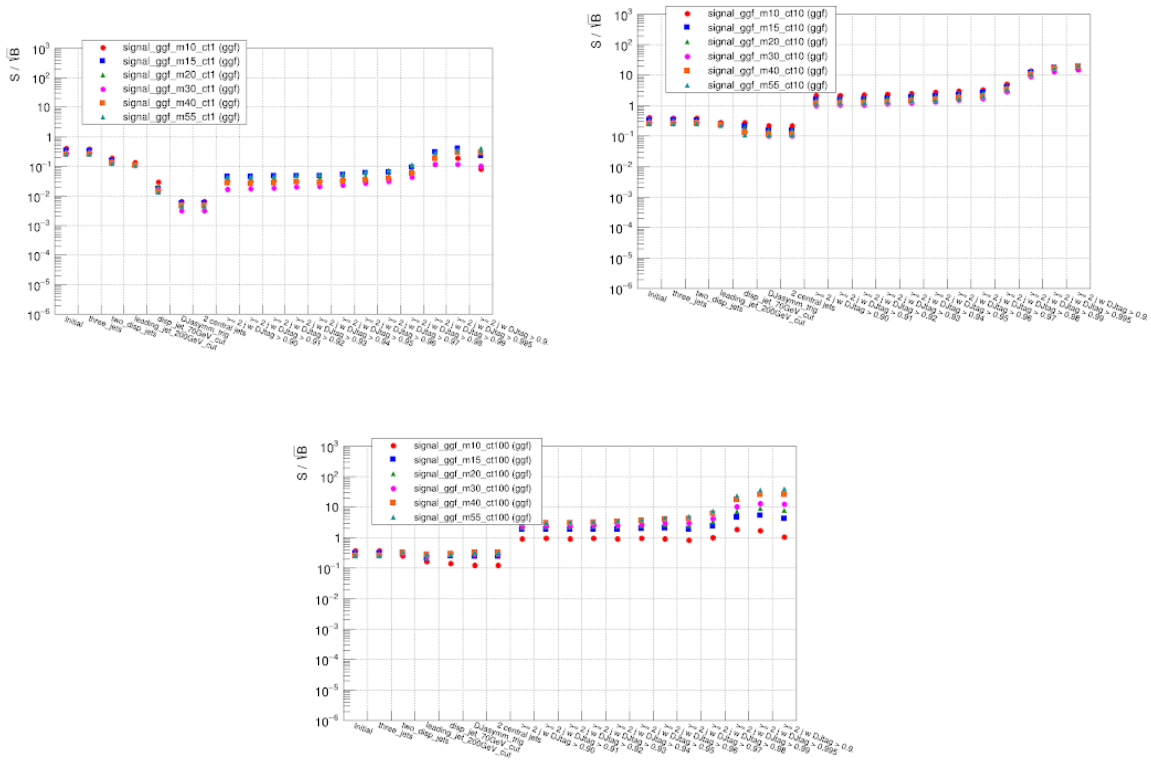


Figure 40. The significance plots for the DJ trigger bucket in the resolved region for signal samples organized by lifetime, with 1 mm samples on the top left, 10 mm samples on the top right, and 100 mm samples on the bottom. The small-R DJ tagger score extends to a maximum value of 0.999, which was found to maximize the analysis sensitivity while maintaining high signal acceptance.

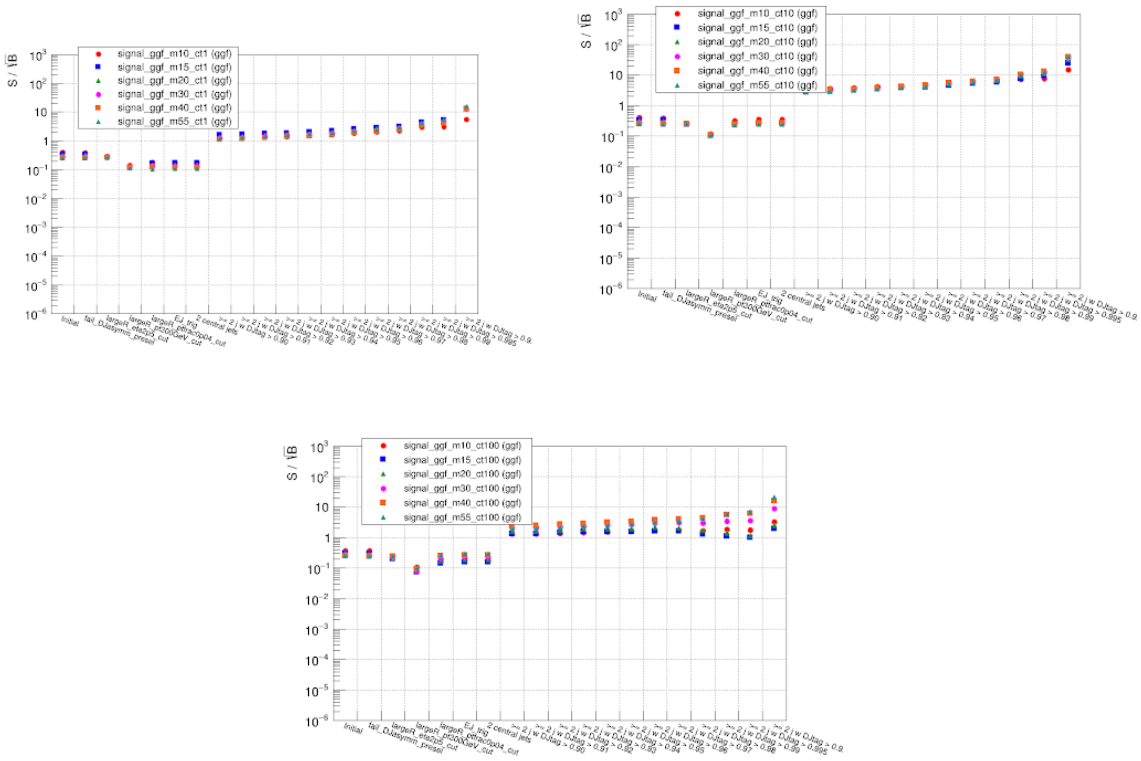


Figure 41. The significance plots for the EJ trigger bucket in the resolved region for signal samples organized by lifetime, with 1 mm samples on the top left, 10 mm samples on the top right, and 100 mm samples on the bottom. The small-R DJ tagger score extends to a maximum value of 0.999, which was found to maximize the analysis sensitivity while maintaining high signal acceptance.

5.8 Background Estimation Attempts and Preliminary Expected Limits

This section describes the data driven background estimation strategies for the two channels of the analysis. The ggF channel is working to define a background estimation strategy that gives good closure. Both a mistag map strategy and an ABCD plane strategy have been attempted for the resolved ggF region, but for reasons to be discussed in this section, a suitable method has not yet been found for scenarios both including and not including the EJ trigger bucket strategy. However, an estimation of the expected limits with and without the EJ trigger is shown, which once and for all confirms that the EJ trigger should be included in the ggF analysis region. The VBF channel implements a combined mistag map and ABCD plane strategy for the resolved case, and just an ABCD plane strategy for the boosted. It is much further along than the ggF, but since the majority of my work has been on defining the trigger strategy for the ggF region, it will not be mentioned here.

5.8.1 ggF Background Estimation and EJ Expected Limits Contribution.

In the ggF resolved region, which requires two small-R displaced jets (defined by a GNN tagger score of $>$ 0.999), a mistag map (MM) method is employed to attempt to estimate the number of background jets mistakenly tagged as displaced by the analysis taggers. To do this, a region orthogonal to the ggF resolved signal region (consisting of at most one displaced jet) is used to create a map parameterized by three jet-level observables. The goal of this is to capture correlations between the DJ GNN tagger score and these observables. In the boosted region, this method does not work because the region only requires one large-R jet, so an ABCD method is attempted instead. In both cases, good candidates for the mistag map parameterizations and ABCD axes have not yet been identified. Attempts to do so have resulted in poor closure, resulting in extremely large uncertainties that negatively impact our expected limits. Studies of this were done in only the DJ Bucket, but efforts to include the EJ trigger bucket as well have not improved the lack of closure. A best attempt at achieving a MM in the resolved region consists of the following parameterization with the corresponding bin edges:

- Displaced tracks count: -0.5, 1.5, 3.5, 7.5, 9.5,11.5,13.5,15.5
- Pseudorapidity: -2.5,-1.5,-0.5,0.5,1.5,2.5
- GN2 b-tag score: -5,1,6,12

The mistag rate is computed directly in data for each bin of the mistag map by dividing the number of tagged jets by all jets in this region. Figure 42 shows the 3d map of the mistag rate for this parameterization taking a logical or of the EJ and DJ buckets. Similar maps are made for the EJ and DJ Buckets alone. Strong correlation between mistag rate and the axis variables motivates their inclusion in the map parameterization. After calculating the mistag map, the probability of tagging exactly zero jets in an event is given by:

$$P\left(0 \text{ tag}|\text{event}\right) = \prod_{i}^{n_{jet}} \left(1 - P\left(\text{tag}|j_i\right)\right) \tag{5.7}$$

where j_i is jet i of the event and $P(\text{tag}|j_i)$ is the mistag rate in the mistag map bin corresponding to j_i. The probability of tagging exactly one jet is

$$P\left(1 \text{ tag}|\text{event}\right) = \sum_{k}^{n_{jet}} P\left(\text{tag}|j_k\right) \times \prod_{i \neq k} \left(1 - P\left(\text{tag}|j_i\right)\right). \tag{5.8}$$

The probability of a given event having $\geq$ 2 tagged jets and thus being in the signal region is then

$$P\,(2\text{ tag}|\text{event}) = 1 - P\,(0\text{ tag}|\text{event}) - P\,(1\text{ tag}|\text{event})\,. \tag{5.9}$$

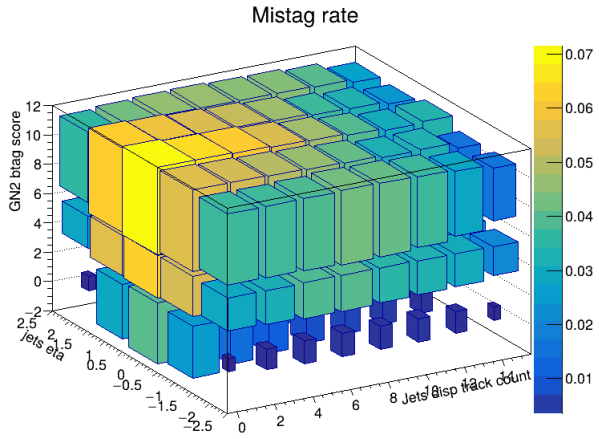


Figure 42. The generated mistag rate for the specified mistag map. This was generated for the or of DJ and EJ buckets in order to assess if adding EJ made the estimation of DJ close better.

To predict the number of events that have 0, 1, or ≥ 2 tagged jets, the above probabilities are calculated and applied as per-event weights, where the sum of all event weights comprise the final background prediction. This prediction is then compared to the observed DJ in the $t\bar{t}$ control region to assess closure. Figure 43 shows the closure plots for the DJ Bucket, EJ Bucket, and DJ Bucket or EJ Bucket, unfortunately showing no good closure.

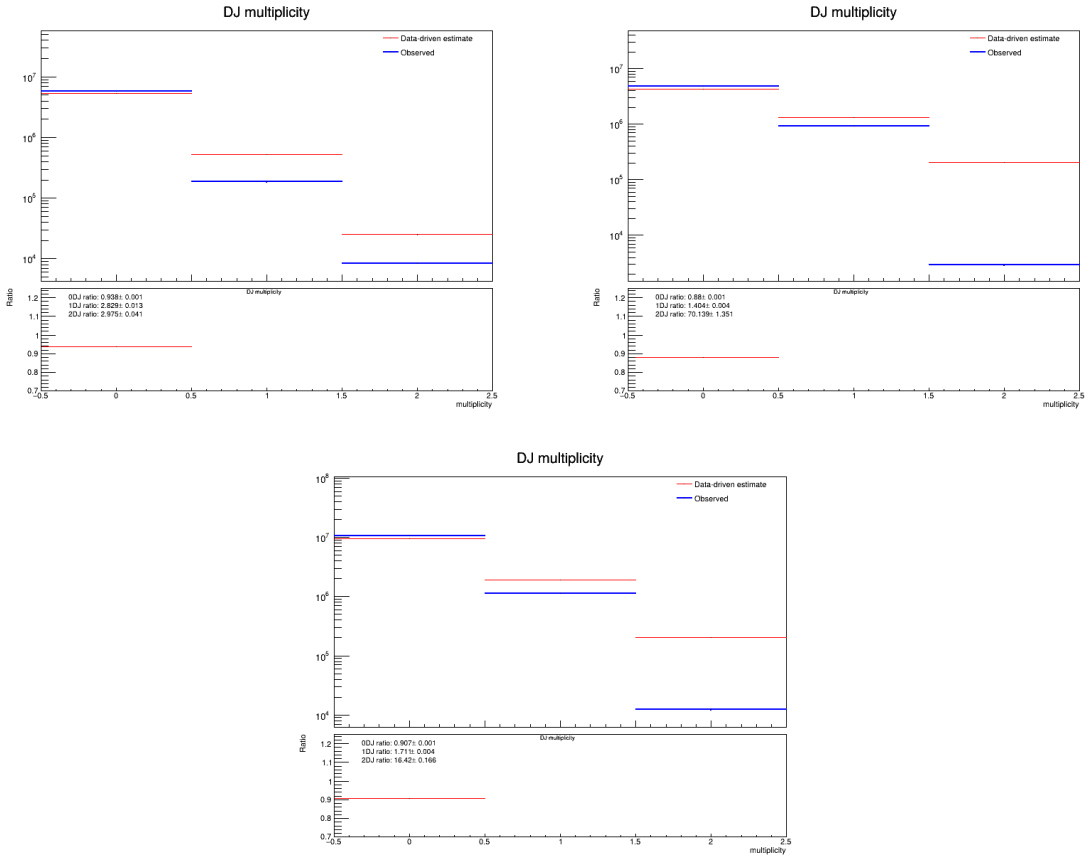


Figure 43. Closure tests of the mistag map for the DJ bucket (top left), EJ bucket (top right), and the combined DJ or EJ selection (bottom). As can be seen, the closure in the signal region is poor, indicating that this mistag-map approach was not sufficiently accurate for the final background estimation.

Efforts are ongoing to understand why it is difficult to find a winning combination of variables to parameterize a mistag map. Since the DJ and EJ triggers select jets that are highly displaced within their selections, it is possible that there are correlations we do not yet understand that violate the independence assumption used in the mistag map method. The motivation behind using these methods was to directly adapt the strategy used in the VBF channel (which will be described briefly below), but it may require a dedicated alternative method, which is under current study.

For the purposes of this thesis, I wanted to obtain a comparison of expected limits in the ggF region both with and without the EJ trigger. This should be the final justification needed to fully implement the EJ bucket into the SR, even as alternative background estimation methods are considered. To accomplish this, I defined a different mistag map using the following variables and their binning:

– Small-R Jet p_T: 0, 50, 100, 200, 300, 400 GeV

– Pseudorapidity: -2.5,-1.5,-0.5,0.5,1.5,2.5

– PTF: 0.0, 0.1, 0.2, 0.4, 0.6, 0.8, 1.0

This was chosen in an attempt to see if adding the PTF variable could help achieve closure. This map was first derived in dijet and $t\bar{t}$ MC and then injected with signal to check the parameterization for contamination. The results of the contamination check for the mistag map in both the DJ Bucket and the or of the DJ and EJ Buckets can be found in Figure 44.

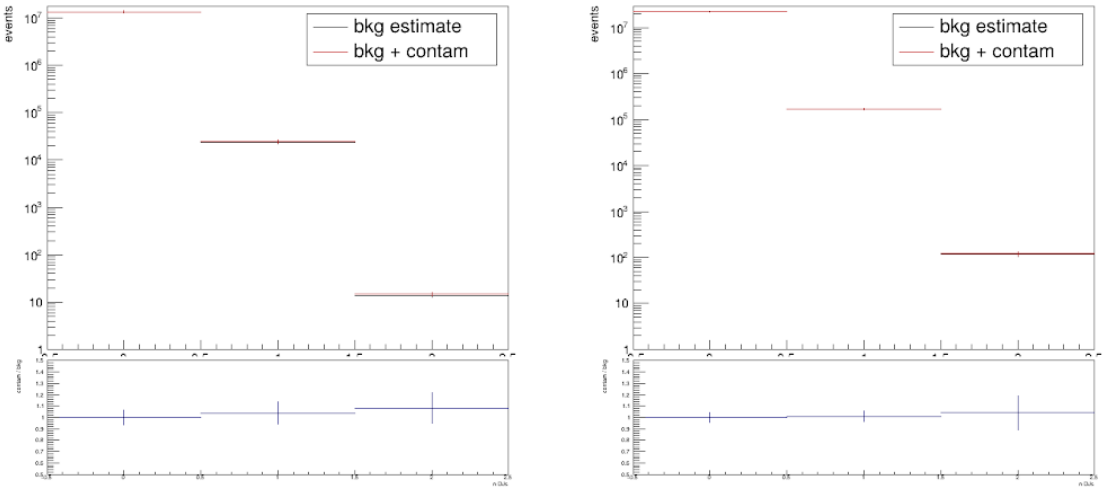


Figure 44. The signal contamination for the mistag map in the DJ Bucket (left) and the or of the DJ and EJ regions (right). This shows very little signal contamination, suggesting that it can be used in the limit framework.

This map is then rederived in data and applied to the relevant ggF SR candidates. To achieve limits, the resulting mistag rate predictions are fed into our limit setting framework, which utilizes FastFrames built in synergy with TRExFitter [78]. It works by performing simultaneous maximum likelihood fits on all provided signal and control regions to determine the signal strength, which is interpreted as the branching ratio of Higgs to a pair of scalar LLPs. This strength is treated as a free parameter, and nuisance parameters allow variations of the expectations for provided signal and background yields according to systematic uncertainties. Upper limits are derived at 95% Confidence Level (CL). For most cases, the parameter of interest (POI), μ (where the expected limit is the 95% CL of the determined μ) was given a range of -0.5 to 1 for fitting. In some particular points, this did not provide a good result, and the upper range of μ was expanded to allow for a larger range so that a fit could be obtained. This could result in some mass/lifetime points yielding an unphysical limit (a branching ratio $>$ 1). The

samples that were given a larger maximum POI range to explore in the fit are the following:

1. The 3 GeV 100 mm sample was given a POI max of 25
2. The 5 GeV 1 mm sample was given a POI max of 10
3. The 5 GeV 100 mm sample was given a POI max of 10
4. The 10 GeV 1 mm sample was given a POI max of 25
5. The 30 GeV 1 mm sample was given a POI max of 10

Figure 45 shows the expected limits for this mistag map using just the DJ trigger bucket and using the DJ or EJ trigger bucket. Since the background strategy is not yet confirmed, a conservative 50% error on the data driven estimate from the mistag map is included in the fit. Figure 46 shows a ratio of the two fits: the numerator being the expected limits set in the DJ or EJ region, and the denominator being the expected limits set in the DJ region only. This allows us to see the relative effect of including the EJ trigger even without having the background strategy finalized. For the samples that had a non-standard POI maximum, this maximum was kept consistent across both the numerator and denominator so as to not introduce any external influence on the ratio. As can be seen, the EJ trigger helps tremendously in the small lifetime case, improving the limits by about an order of magnitude on average. While it does not seem to help the 10 mm or 100 mm lifetime sample limits due to a worse signal to background ratio from the EJ trigger for these samples, the results are still within an order of magnitude, showing at least similar performance for the highly unoptimized strategy. In principle, adding more information, such as additional contribution from the EJ trigger, should never weaken limits. It should be noted that the weakening of the limits for the 10 mm and 100 mm lifetimes are likely due to a worse closure for the mistag map method when using the EJ trigger. The EJ trigger has a worse signal to background ratio in this regime than the DJ, and so the relative increase in background events likely leaks into the map, making the uncertainties worse. This in turn makes the expected limits worse as the limits are highly sensitive to the uncertainty. Once a good background estimation strategy with smaller closure error is found, these limits will be improved to obtain at worst the same expected limits as the DJ bucket alone, and likely show some marginal improvement. The order of magnitude increase in sensitivity in the small lifetime region even with these challenges shows great promise in using the EJ trigger.

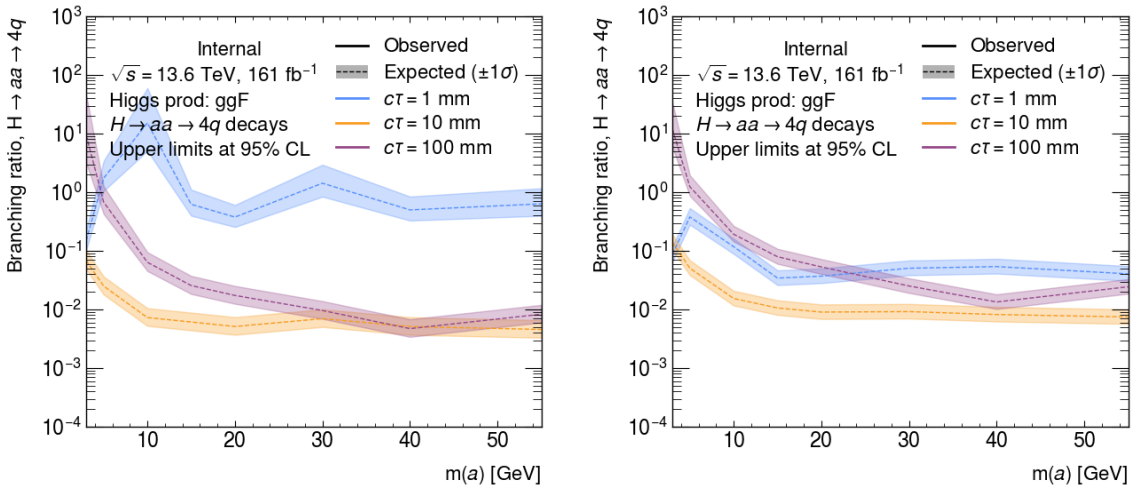


Figure 45. Preliminary expected limits derived from the mistag map data driven background estimation strategy, using only the DJ Bucket (left) as well as the DJ or EJ buckets together (right). Since this background estimation strategy has not been able to close, a 50% uncertainty has been included on the background estimate.

5.9 Current Status of Analysis and Next Steps

As of the writing of this thesis, the analysis is still underway. Due to the difficulties of defining a background estimation for the ggF region as described in the previous section, it is falling behind the progress of the VBF channel. The focus of the analysis is shifting to finish tasks for the VBF region, as the background estimation is essentially finalized. Then the focus will shift back to the ggF channel, as combining the two is still expected to give an extremely strong constraint on the Higgs branching ratio to LLPs compared to recent searches [71]. There is still some work to be done, but this analysis will provide some of our best limits to date of the coupling of the Higgs boson to LLPs using the new trigger capabilities of the ATLAS Run 3 system.

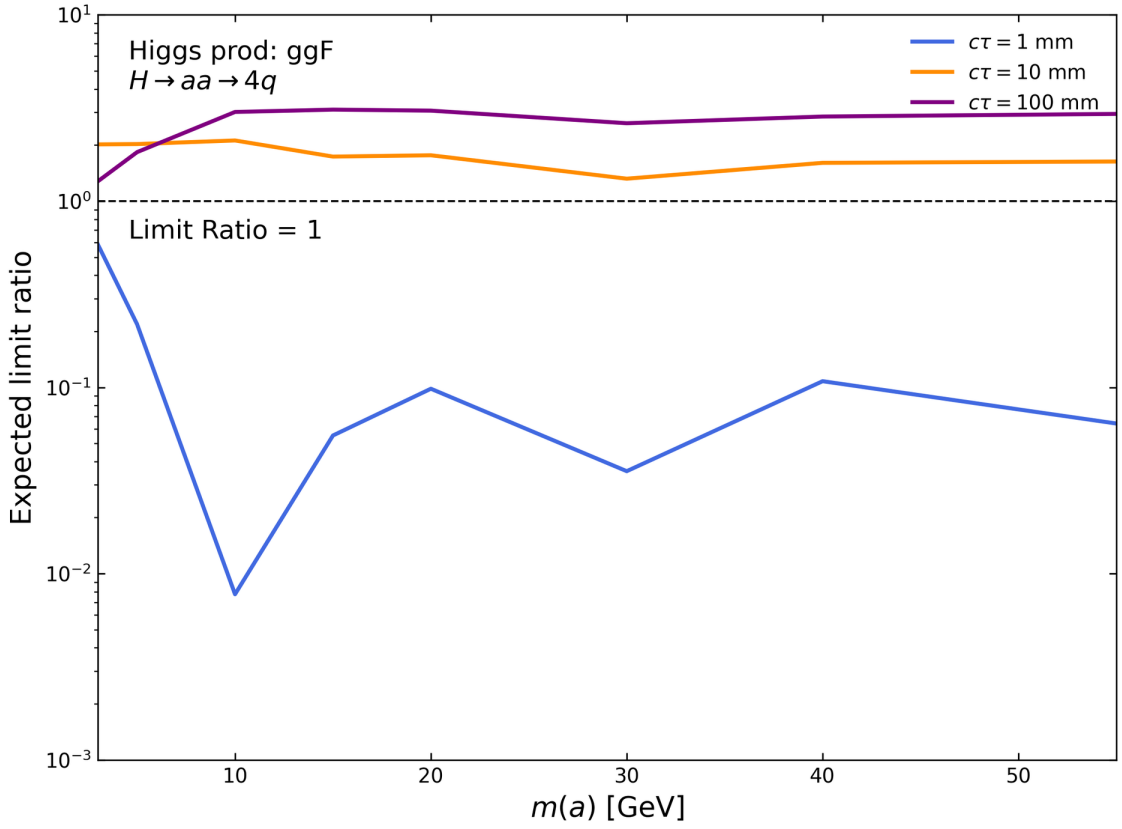


Figure 46. The ratio of the two expected limits from the previous Figure 45. The numerator of the plot is the expected limits in the ggF region using both the DJ and EJ triggers to define the mistag map, and the denominator is the expected limits in the ggF region using only the DJ trigger to define the mistag map. While not providing additional sensitivity to the 10 and 100 mm lifetimes, the inclusion of the EJ trigger improves the limits set on the 1 mm lifetime by a full order of magnitude in most cases, up to 2 orders of magnitude for one mass sample (10 GeV).

CHAPTER VI
UPGRADES TO ATLAS FOR THE HIGH-LUMINOSITY LHC AND THE PHOTON POINTING ALGORITHM

This chapter describes the coming upgrades to the ATLAS in the wake of Run 4, the beginning of the High-Luminosity LHC. It describes the major detector and TDAQ upgrades, namely the new Global Trigger, which will act as the Run 4 version of the Run 3 L1 Topological Trigger. Then, motivation for a new trigger algorithm for Global: the photon pointing based diphoton trigger, will be introduced, including a discussion of SM and BSM interest in diphoton signatures. The algorithm will then be described, including its current implementation, as well as attempts to create a machine learning based version of the algorithm. It ends with a discussion of future work to be done on this trigger algorithm.

6.1 The High-Luminosity LHC

Run 4 of the LHC is scheduled to start in 2030. This marks the beginning of the High-Luminosity LHC (HL-LHC). The LHC will be upgraded to increase the instantaneous luminosity, resulting in a higher number of inelastic proton-proton collisions. The full HL-LHC dataset will be about 3ab^{-1}, or 10 times the integrated luminosity of the full Run 1-3 dataset. The HL-LHC will reach instantaneous luminosities ranging from $5 - 7.5 \times 10^{34}\text{cm}^{-2}s^{-1}$, resulting in a peak value of 200 inelastic proton-proton collisions per bunch crossing, also referred to as pileup [8]. This is about a factor of three larger than the current average pileup for Run 3, which will make triggering and event reconstruction significantly more challenging.

6.1.1 Detector and TDAQ changes for HL-LHC.

To not only maintain, but improve physics performance for the HL-LHC, ATLAS will undergo several upgrades to its detector components and the TDAQ system to withstand these much higher intensities. These are referred to as the Phase-II upgrades of the ATLAS Detector. The largest upgrade to the ATLAS detector itself is the complete replacement of the ID with a new fully silicon tracker called Inner Tracker (ITk) [79]. The ITk will provide higher granularity tracking and will be able to provide efficient pattern recognition and track reconstruction in the high pileup environment and resulting radiation dose. The ITk will consist of two silicon tracking subsystems: an inner Pixel Detector and an outer Strips Detector. High granularity pixel sensors are used in the innermost barrel layers and endcaps closest to the interaction point, where the particle occupancy is the greatest, while the silicon strip sensors cover the outer layers and endcaps to provide coverage over the rest of the tracker's volume. A diagram of the ITk layout with the different pixel and strip layers can be found in Figure 47. ITk assembly and construction is currently under way. There will also be an entirely new subsystem: the High-Granularity Timing Detector (HGTD) [80], designed to augment the performance of ITk in the forward region by providing high granularity timing information for charged particles, which will aid in pileup mitigation. The bulk of the calorimeters will remain the same, with only the

readout electronics changing to accommodate the more granular information it will pass to the trigger systems.

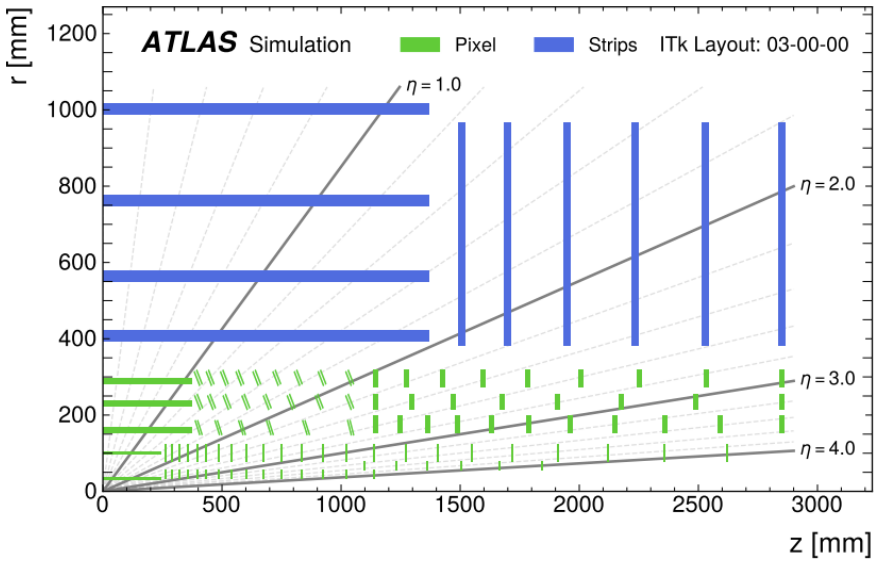


Figure 47. Schematic sketch of the ITk layers in z-r space with different values of η marked with lines [14].

The ATLAS TDAQ system for the Phase-II is getting a very large overhaul to capitalize on the increased number of events available [16]. The rate of accepted events at each tier of the trigger system will increase, with the hardware trigger increasing from the 100 kHz rate of L1A output in Run 3 to a new 1 MHz output rate, and the software trigger will take in this new 1 MHz rate as input, subsequently increasing from an output rate of about 3 kHz in Run 3 to a new 10 kHz rate to be saved to tape. This will require upgrades to both levels of the trigger system. The software trigger will be implemented as the Event Filter (EF), which consists of a CPU and GPU farm to perform highly complex trigger algorithms, including tracking information. The hardware trigger will gain a few changes and new pieces of hardware: a new Forward Feature Extractor (fFEX) designed to specifically target physics objects in the forward region, upgrades to the Muon trigger system readout and electronics, and the Global Trigger, which will replace the existing L1Topo and receive input from all the FEXes, the Muon system, and (new for Run 4) the calorimeters themselves. The Global Trigger will be discussed in greater detail in the following sections. The full schematic of the Phase-II TDAQ system can be found in Figure 48.

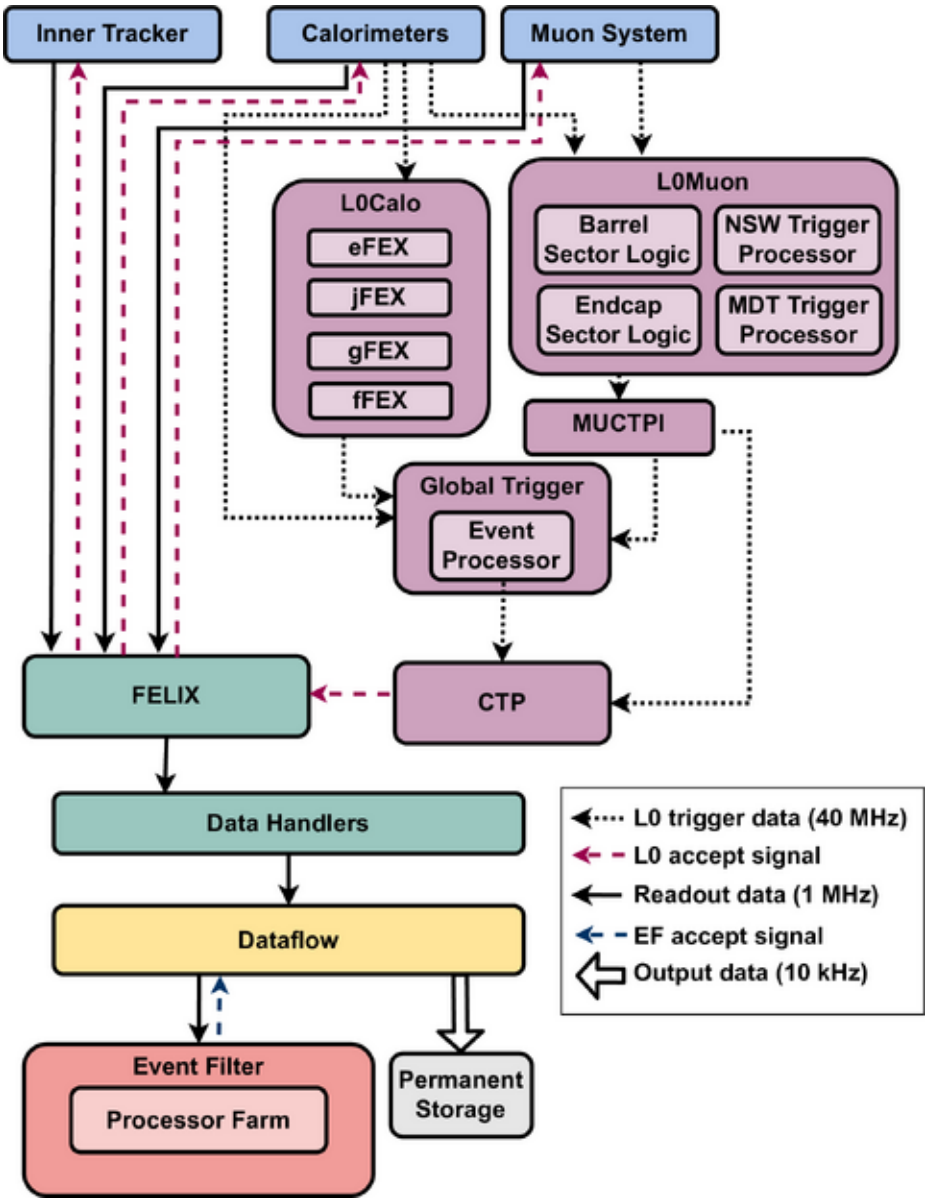


Figure 48. A schematic showing the full dataflow of the Phase-II TDAQ system, starting with detector input to the trigger system and ending with permanent storage. Of particular interest is the Global Trigger, which will take in inputs from the rest of L0 and the calorimeters to run final hardware selection algorithms before passing trigger objects to the CTP[15].

6.2 ATLAS Global Trigger Upgrade

As part of the upgrade plan for the ATLAS detector, the new Global Trigger will be implemented at the hardware trigger level (L0). The Global Trigger will replace the current L1Topo, taking in information from the FEXes and Muon system to make trigger decisions. In addition to information from

earlier parts of the hardware trigger, Global Trigger will also receive inputs from the ATLAS calorimeters themselves, providing calorimeter-cell level granularity information to be used in trigger algorithms. The goal of Global Trigger is to bring offline-like algorithm capabilities to the L0 trigger system to better refine trigger objects before final selection in the CTP. To achieve this, input data including full granularity cell data from the LAr and Tile Calorimeters, L0Calo FEX trigger objects, and L0Muon information will be sent to Global through time-multiplexers (MUXes). The MUXes will distribute events organized by their bunch crossing ID into 48 Global Event Processors (GEPs) for parallel processing such that all the information for one event is contained on one GEP. These 48 processors mean that for a 25 ns bunch spacing, each event has about 1 μs to process each step before moving on to the next BC. The multiplexing is fully deterministic, meaning that the full workflow for an event is handled on one processor. This allows for the processor to use fully pipelined, low-latency logic for running algorithms. Each GEP node is housed on a Xilinx® Versal Premium® VP1802 FPGA to allow for flexible firmware algorithms. Both MUXes and GEPs are built on a common piece of hardware, the Global Common Module (GCM) such that only the firmware differs between them. Diagrams of the time-multiplexing structure of Global can be seen in Figures 49 and 50. The full latency budget of the L0 system for Run 4 is about 10 μs, with GEP having about 4-5 μs total to process an event and get it to the CTP. A diagram of the L0 latency budget can be found in Figure 51. Specifically, the green arrow points to the latency due to the cell level information coming from LAr. This constraint limits the amount of information that LAr can send. This will affect the diphoton algorithm discussed later in restricting the access the algorithm has to the η values of individual cells. This is why in the algorithm many different combinations of cell and tower η values were tested.

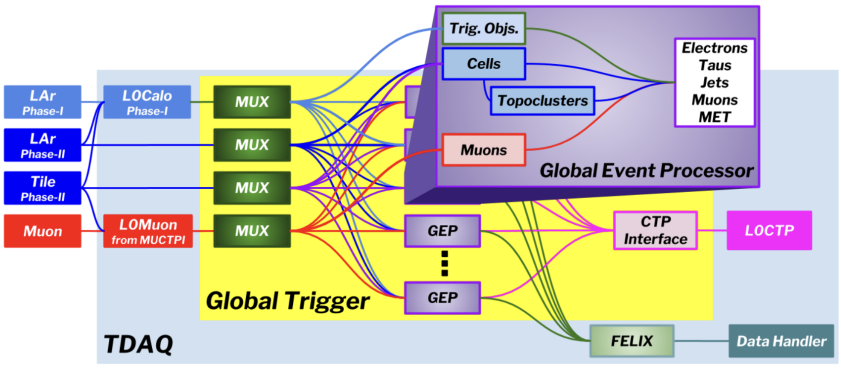


Figure 49. A schematic showing the dataflow path of Global Trigger, where inputs are fed into MUXes which in turn sort the information by BCID to the GEPs where algorithms are run. The resulting decisions are sent to CTP for final decision[16].

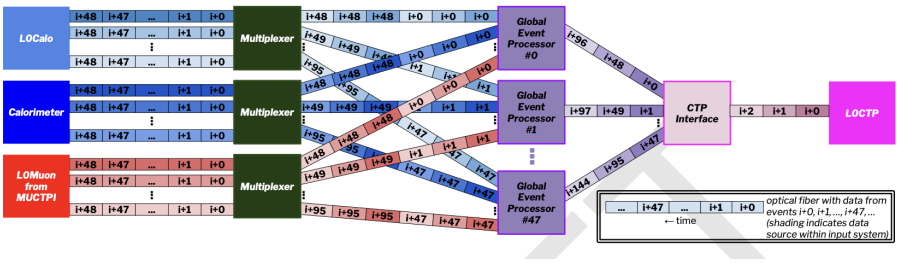


Figure 50. A schematic showing the time-multiplexing of the MUXes into the GEPs. Events are multiplexed such that one GEP is handling the information for one BCID. The CTP interface will demultiplex them such that they show up in order to the CTP[16].

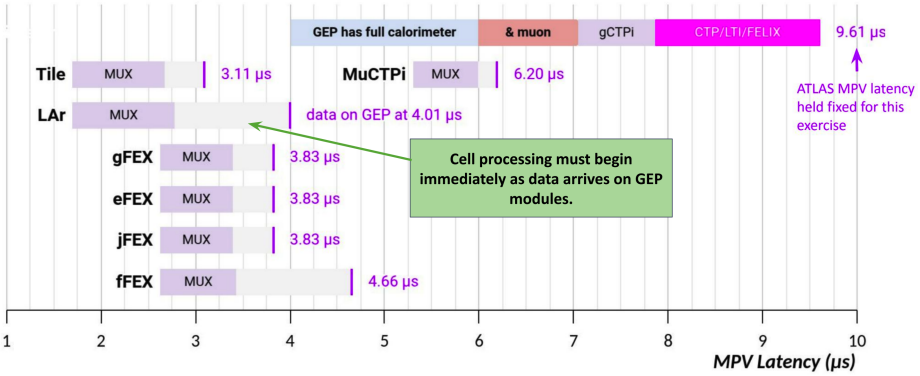


Figure 51. A diagram of the Maximum Possible Value (MPV) latency of the Global system, showing when different inputs will be available to the GEP[17].

6.3 Diphoton Trigger Algorithm for Global Trigger

A diphoton trigger, which bases its trigger decision on the kinematics of two photon objects, is of particular interest for the Global Trigger upgrade. Diphoton triggers are extremely relevant for the detection of diHiggs events: events where two Higgs bosons are produced via the rare trilinear vertex of the Higgs self-coupling, an important value we need to study in the SM that also has BSM applications. The decay channel of $HH \rightarrow bb\gamma\gamma$ where one of the Higgs decays to a pair of photons through a heavy quark loop, is highly relevant to diHiggs discovery due to its high signal to background ratio. Lowering the p_T threshold required to trigger on diphoton events can therefore help the chance of discovering diHiggs at the HL-LHC. Additionally, in a BSM context, softer diphoton pairs can also be a result of a lighter BSM particle decaying to diphoton, such as an ALP. In previous Runs of the LHC, diphoton candidates could be identified offline using photon pointing: where highly granular cell information is used to determine the point of origin along the beam line of a photon. If two photons come from a sufficiently close enough

window, it can be assumed that they came from the same decay. The Global diphoton trigger is being developed by the Global Trigger diphoton algorithm team at Brookhaven National Laboratory (BNL). I participated in a DOE Office of Science Graduate Student Research (SCGSR) grant to work with the team there on this algorithm, including investigating machine learning implementations of this trigger algorithm and assessing the feasibility of this trigger on LLP samples. The following section will describe in more detail the motivation for diphoton searches.

6.3.1 Motivation for Creating the Diphoton Trigger.

Following the discovery of the 125 GeV SM Higgs boson by the ATLAS [22] and CMS [23] Collaborations, much of the LHC physics programme has been dedicated to studying its properties. One of the central predictions of the SM is the form the Higgs potential, which is relevant not just to its mass but also its self-couplings. After Electroweak Symmetry Breaking (EWSB), the SM Higgs potential takes the form $\lambda_3 = \lambda = \lambda_{SM} = \frac{m_H^2}{2v^2}$ where the mass of the Higgs m_H is 125 GeV and the Higgs vev $\sim$ 246 GeV. While many expected SM Higgs couplings have been measured in experiments, direct experimental access to the Higgs self-coupling has remained limited due to statistics [81]. The trilinear Higgs coupling, λ_3, enters directly in Higgs boson pair production (see 3) and so a search for diHiggs production will provide the first opportunity to directly probe the Higgs self-coupling at collider experiments. Deviations from the SM prediction of this coupling would be indications of BSM physics.

Studying the Higgs self-coupling is not only an important test of the SM, but can also shed light on several BSM topics such as the vacuum stability of the universe [82], inflation [83], and baryogenesis [84]. These extensions to the SM predict modifications to the Higgs potential, leading to anomalous Higgs self-couplings. These anomalous couplings can affect the diHiggs production cross section as well as the kinematic properties of the diHiggs pair, making measurements of HH a sensitive probe for BSM physics.

At the LHC, diHiggs pairs are produced predominantly through gluon gluon fusion. At leading order, this process proceeds through two Feynman diagrams: the triangle diagram containing the trilinear Higgs self-coupling and the box diagram involving a virtual top-quark loop with no trilinear Higgs vertex. These amplitudes interfere destructively, making the total production cross section highly sensitive to the value of the Higgs self-coupling. The sensitivity of the diHiggs process to the trilinear Higgs self-coupling is greatest at relatively small invariant masses of the Higgs boson pair, where the destructive interference between the triangle and box amplitudes is strongest. In this kinematic regime, the decay products of the Higgs bosons, including photons from $H \to \gamma\gamma$, typically possess lower transverse momenta. Consequently, reducing the diphoton trigger thresholds increases the acceptance for the region of phase space that is most sensitive to measurements of the Higgs self-coupling [85]. As discussed in Chapter 2, the Higgs boson can couple to ALPs. These ALPs can then in turn commonly decay to photons, prompting

searches for Higgs to photon final states [86]. These searches can therefore also benefit from an improved diphoton trigger.

6.3.2 Feasibility of Using Diphoton Triggers on Exotic Samples.

In order to establish a use case for a diphoton trigger on exotic Higgs samples in general, studies were conducted using available exotic Higgs samples and triggers in Run 3. In Run 3 2024, a new diphoton trigger (referred from here onward as the asymmetric diphoton trigger) was introduced to the trigger menu with the following HLT trigger chain: HLT_g45_medium_g20_medium_L1eEM40L_2eEM18L. This increased the threshold on the higher p_T photon in exchange for lowering the threshold on the lower p_T photon with respect to the original diphoton trigger: HLT_g35_medium_g25_medium_L12eEM24L . This was done with the intention of improving sensitivity to the Run 3 $HH \rightarrow bb\gamma\gamma$ and $H \rightarrow \gamma\gamma$ analyses. Current Higgs to displaced photon via an axion-like particle (ALP) searches in ATLAS Exotics typically look for an $H \rightarrow Za$ signal, where the Z boson decays leptonically. These leptons typically provide the objects used for triggering in the analysis. This works well for associated production of the ALP, but cannot be used to search for direct pair production from the Higgs: $H \rightarrow aa \rightarrow 4\gamma$. A possibility for this process could be to trigger on the pair of photons from one ALP directly using a diphoton trigger. In the case of light ALPs, the photons can become "merged," meaning they are reconstructed as one photon in analysis. If the Higgs were to be high enough momentum such that the light ALPs come out fairly boosted, the two "merged" photons from each ALP could also potentially be identified by a diphoton trigger. With these topologies in mind, it is worth exploring the use of existing diphoton triggers on exotic Higgs samples in order to see if a Run 4 diphoton trigger could aid in future versions of displaced photon analyses.

At the time of this study, there were no available MC samples for Higgs decaying to any kind of diphoton signature involving ALPs. As a proxy, samples of light resonance of Higgs decaying to diphoton m directly ($H \rightarrow \gamma\gamma$) were used to assess the performance of the Run 3 asymmetric diphoton trigger. Table 15 below lists the considered samples, which are all reconstructed in MC campaign mc23e, corresponding to the 2024 data taking conditions.

DSID	Higgs Mass	Decay
603519	60 GeV	$\gamma\gamma$
603520	70 GeV	$\gamma\gamma$
603521	80 GeV	$\gamma\gamma$
603522	90 GeV	$\gamma\gamma$
603516	100 GeV	$\gamma\gamma$
603517	110 GeV	$\gamma\gamma$
603518	120 GeV	$\gamma\gamma$

Table 15. Light Higgs Resonance decays to diphoton samples. These are used as a proxy for an ALP decaying to diphotons in trigger studies as these samples were readily available.

To assess the performance of the asymmetric diphoton trigger on these samples, trigger efficiency turn on curves were made parameterized by the leading and subleading p_T of the photon pair. The denominator selections were made to ensure good events were chosen, including basic event cleaning, and the presence of a diphoton pair with tight identification and isolation requirements on both photons. The numerators are the same selections with the additional requirement of passing the diphoton trigger. Figure 52 shows the turn on curves for the asymmetric diphoton trigger as a function of leading and subleading photon p_T for a subset of the listed samples. In particular, the sharp turn on of the subleading photon p_T suggests that we could benefit from lowering the threshold of the subleading photon even further for processes involving light scalars or pseudoscalars decaying to two photons.

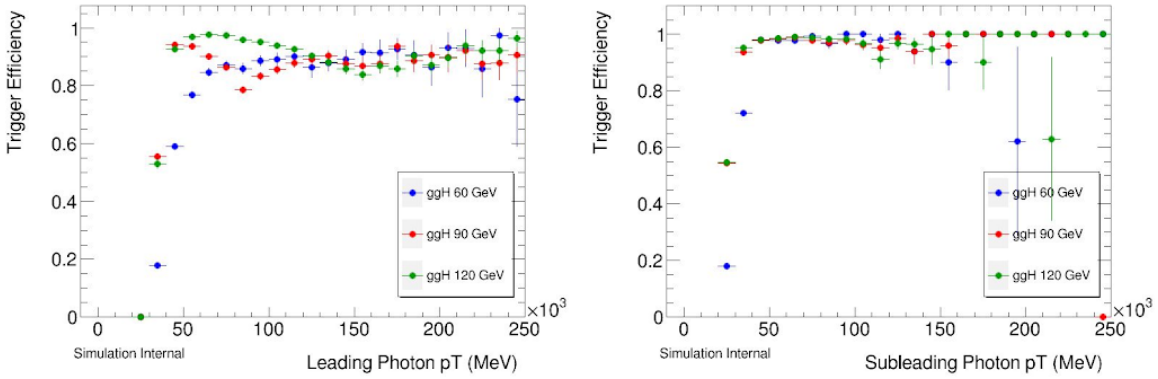


Figure 52. Signal efficiency turn on curves for the asymmetric diphoton trigger on exotic Higgs samples as a proxy for ALPs decaying to diphoton.

6.4 The Baseline Diphoton Algorithm for Global Trigger

This section describes the current implementation of the Photon Pointing in Z (PPZ) algorithm. The general outline of the algorithm will be discussed,

followed by a discussion of inputs as they pertain to the information accessible to the algorithm in the Global Trigger environment. This algorithm has been run through High Level Synthesis (HLS) to obtain resources estimations for running it on FPGAs in Global Trigger.

6.4.1 The PPZ Algorithm.

The general idea of the Global diphoton algorithm is to take advantage of the newly available calorimeter cell-level information to be more selective of diphoton candidate pairs at the hardware trigger level. Rather than relying on calorimeter energy deposits alone, the access to cell information can allow us to do a version of *photon pointing*, where a photon's position along the beam axis (z) can be determined by drawing a straight line back through two calorimeter cell layer hits. Before the HL-LHC, this would only be possible offline. The formula to calculate PPZ is a function of only the η value of the corresponding hits in layers 1 and 2 of the calorimeter. It is calculated as follows:

$$PPZ(\eta_1, \eta_2) = Z_1 Z_2 \frac{\sinh(\eta_1) - \sinh(\eta_2)}{\sinh(\eta_1) Z_2 - \sinh(\eta_2) Z_1} \tag{6.1}$$

Where $Z_i = R_i(\eta_i)\sinh(\eta_i)$, and $R_i(\eta_i) = c_0 + c_1\eta_i + c_2\eta_i^2$ is the taylor expanded response of the calorimeter as a function of η. The constants c_i are measured in test beam data, and differ between endcap and barrel region. In the endcaps, the parameterization is done only in terms of Z_i, so to avoid division (which will be relevant when discussing HLS conversion), everything is done is terms of Z_i from the onset.

The algorithm follows this general workflow:

1. Different values of η are calculated for layer 1 and 2 during the Global tower-building phase
2. eFEX inputs to Global will be tower matched, and those tower's pre-calculated η values are used in the PPZ calculation
3. ΔPPZ is calculated for all pairs of e/γ objects to use in the trigger decision
4. The pair with the lowest ΔPPZ is the event level trigger candidate

Figure 53 shows a nice visual representation of the algorithm

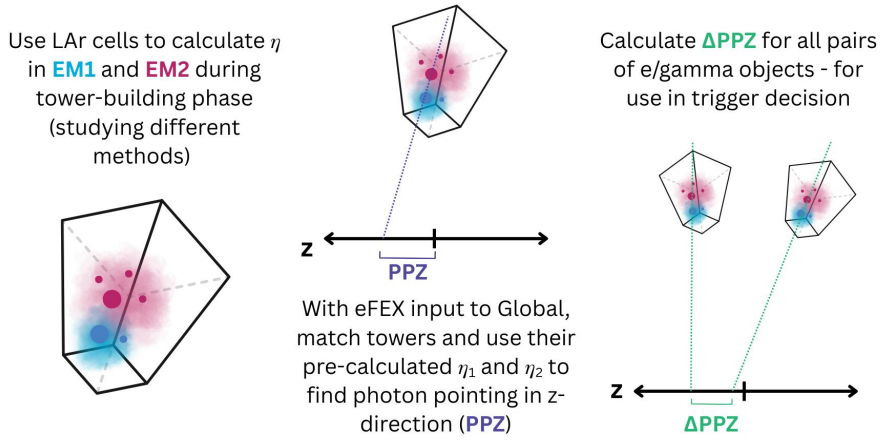


Figure 53. A sketch of the PPZ algorithm main steps. Image credit Liza Brost.

Due to the computational and latency related challenges of providing a large number of calorimeter cells for any given event, we need to determine a scheme for receiving only a subset of relevant η values in each layer to use in the algorithm. This needs to happen before the actual algorithm, at the tower builder step, as this is the only time in the Global Trigger framework that all the cells will be available for calculations. Several studies were conducted by the BNL team using a Run 4 sample of $HH \rightarrow bb\gamma\gamma$ with an average pileup of 200 to mimic peak HL-LHC conditions to determine a method for extracting η_1 and η_2 values during tower building. For the purposes of these studies, eFEX TOBs are treated as e/γ objects with loose requirements used in Run 3 (these will be discussed in the machine learning portion). The two main methods considered are the so called "Max Cell" and "Energy Weighted Barycenter" algorithms. In the Max Cell algorithm towers that fall within a $\Delta R < 0.1$ of an eFEX TOB are considered. The η for each layer can either be taken directly from the max cell in the tower collection (max cell), or can find the tower with the highest E_T and then use the η in both layer 1 and 2 corresponding to the cell with the max(E_T) within that tower (max tower max cell). In the weighted barycenter algorithms, a few approaches were considered including taking the average of the top 5 cells across all towers that meet our eFEX matching cuts (Barycenter Average Over Top 5 Cells), calculating the barycenter for each tower and then taking the sum later E_T weighted average of those barycenters (Barycenter Average of Tower Barycenters Top 5 Cells), and a ringer-like geometric approach, where the max cell in L2 and its direct neighbors form the inputs to the barycenter calcuation (Geometric Ring L2), based on the existing ringer algorithm [87]. These methods were all evaluated by constructing Receiver Operator Characteristic (ROC) curves to evaluate the background rejection at different signal efficiency points of the different methods. Based on these studies, it was determined that using the max cell strategy in layer 1 and the average of individual tower barycenters in layer 2 gave the best

performance while still being attainable from a resource perspective. The most recent studies show that the algorithm can run with input at 11 bit integer and 4 bit decimal precision, producing reliable minΔPPZ calculations.

6.4.2 High Level Synthesis Implementation.

In order to run this algorithm in the Global Trigger environment, it needs to be translated from $C++$ simulation code to VHSIC (Very High Speed Integrated Circuit) Hardware Descriptive Language (VHDL) that is suitable to run as firmware (FW) on FPGAs. Third party tools such as the Xilinx Vitis HLS tool from AMD [88], provides clean interfaces and resource estimates for simulating the running of algorithms on proprietary FPGAs that are used to construct Global Trigger hardware.

One of the biggest challenges in converting the PPZ algorithm using HLS is the restriction on floating point precision in η. Since the calorimeter layers are highly granular, there needs to be a significant precision on the decimal of η values to properly distinguish nearby cells. However, full floating point precision is not feasible to implement within latency and resource requirements. Studies were conducted to find suitable precision for values needed in the algorithm, eventually finding precision values that gave a good trade off of performance and resource and timing usage. Deciding on final precision values will be a discussion that has to happen across all the Global algorithm teams as the system continues to mature in the wake of Run 4. Figure 54 shows the usage of resources for a version of the PPZ algorithm in the Vitis HLS software. The calculation of PPZ carries the bulk of the latency required for the algorithm. This motivates searching for machine learning alternatives to the algorithm that can achieve similar or better performance in a less resource intensive way.

Performance & Resource Estimates

Modules Loops Hide empty columns

MODULES & LOOPS	LATENCY(NS)	ITERATION LATENCY	INTERVAL	TRIP COUNT	PIPELINED	STRUCTURE	BRAM	DSP	FF	LUT	URAM
PPZCalcStream (3)	396.000		96		no	function	0	25	7,003	14,508	1
PPZCalcStream_Pipeline_PPZCalc_Loop (1)	217.000		16		loop auto-rew	function	0	25	5,836	4,931	1
SortPPZArr (1)	70.839		16		loop auto-rew	function	0	0	58	6,797	0
PPZCalcStream_Pipeline_VITIS_LOOP_47_1 (1)	75.006		15		loop auto-rew	function	0	0	497	1,825	0

Figure 54. Resource estimates for a version of the PPZ algorithm implemented in HLS using Vitis software.

6.5 A Neural Net Approach to Photon Pointing

Another way to tackle the task of obtaining a photon point trigger is to apply Machine Learning (ML). In the following section, we will look at two ML methods used to try and provide a good value that can determine a diphoton trigger decision at the hardware trigger level. The first method is a simple Signal Background Classifier, which was trained on $HH \rightarrow bb\gamma\gamma$ as well as $Z \rightarrow ee$ MC signal samples to give a floating point per photon score ranging from 0 to 1,

with 0 being "background like" and 1 being "signal like." The second method is a Linear Regression, which attempts to use inputs to determine the PPZ value per photon. As this is the most computationally intensive step to implement in FW, it was chosen to be the first part of the diphoton algorithm we attempted to use ML to replace, since the following steps (calculating ΔPPZ and then determining min(ΔPPZ)) can be done fairly simply with some combinatorics and sorting algorithms.

6.5.1 NN Architecture.

Both methods explored in this study use a Feed Forward Neural Network (FFNN), which is a ML model in which the information only flows in one direction, from input to output. An FFNN was chosen for its simplicity and ease of training, with the intent that it can be made fairly lightweight for eventual HLS. The architecture consists of an input layer, three dense hidden layers, and a single output layer that is either a Sigmoid in the case of the Signal Background Classifier or a Linear Activation in the case of PPZ Regression. Rectified Linear Units (ReLU) activations are used in each hidden layer because they are computationally efficient, generally provide good performance for deep networks, and can help mitigate the "vanishing gradient problem" where different functions like sigmoid or tanh can have gradients that approach zero and can cause early layers in a deep network to stop learning. A dropout of 0.1 is also applied after each activation. This deactivates 10% of neurons randomly at each iteration to help prevent overfitting. The full architecture is as follows:

Input Layer (training features, customizable)

Input $\rightarrow$ Hidden Layer 1 (1000 nodes)

Hidden Layer 1 $\rightarrow$ ReLU

ReLU $\rightarrow$ Dropout(0.1)

Dropout(0.1) $\rightarrow$ Hidden Layer 2 (600 nodes)

Hidden Layer 2 $\rightarrow$ ReLU

ReLU $\rightarrow$ Dropout(0.1)

Dropout(0.1) $\rightarrow$ Hidden Layer 3 (50 nodes)

Hidden Layer 3 $\rightarrow$ ReLU

ReLU $\rightarrow$ Dropout(0.1)

Dropout(0.1) $\rightarrow$ Output Layer (1 node) (sigmoid for S/B, linear for Regr.)

The hidden layers progressively learn nonlinear combinations of the input variables. The decreasing layer widths compress the learned representation in stages before producing the final prediction. The specific number of nodes per hidden layer was inherited from previous iterations of this study, and have not been optimized. For the purposes of the model, there is little impact to keeping this many neurons, but a dedicated study to assess how lightweight these layers can be made without sacrificing performance in order to run through HLS tools will be necessary in the future.

This model uses Adaptive Moment Estimation (Adam) as an optimizer to adjust learning rates during training[89]. Adam combines the advantages of two different optimization techniques, Momentum [90], which accelerates gradient descent and speeds up convergences by using a moving averages of past gradients, and Root Mean Square Propagation (RMSprop) [91], which uses an exponentially weighted moving average of squared gradients to prevent the learning rate from decreasing too quickly during training. Due to its combination of methods, Adam provides a balanced and efficiency learning rate optimization. The learning rate of 10^{-3} was chosen as a standard, which is common for most Adam optimizers. A standardized batch size of 128 was chosen to trade off benefits of a smaller batch size such as better generalization with benefits of a larger batch size such as being able to run many events simultaneously by taking advantage of GPU acceleration. A max epoch of 50 was chosen to be safe, although most models did not need that many epochs to successfully train. The inputs for both models were normalized by a standard scaler, which was saved for reproducibility. In the case of the regression model, the same scaler was used to unscale the outputs, allowing the result of the regression to be interpreted as a PPZ value in mm. Once training ended, the model with the best learning rate would be saved to an output file so it could be used later.

On average, these models are fairly heavy, having a total of over 600,000 parameters, with most of that weight coming from the first to second hidden layer step. While this is far too dense of an architecture to implement on an FPGA using HLS, the focus of this study for now is on obtaining good performance of the ML method, and focusing on trade offs between performance and resource usage later. The model, training scripts, and plotting functions were all implemented in a python framework [92] using Keras [93], which was chosen for its synergy with tools we will eventually use to convert this model to VHDL using HLS4ML [94, 95].

6.5.2 Input Variables and Event Cleaning.

Both methods take in eFEX TOB variables as input, and apply Run 3 loose quality cuts to clean events for training. Additionally, other event selections were made to ensure consistency with the baseline PPZ algorithm studies. In addition to the eFEX ROIs, the regression model uses the layer information specified in the baseline diphoton algorithm as training input to aid in its goal of regressing to a per photon PPZ value. The eFEX TOBs contain the following information:

kinematic information E_T, η, ϕ, and shower shape informationR_η, R_{had}, $w_{s,\text{tot}}$, and E_{ratio}. The shower shape information is defined as follows:

R_η is the lateral shower containment, defined as

$$R_\eta = 1 - \frac{E_T^{\text{clust}}}{E_T^{\text{env}}} \tag{6.2}$$

where E_T^{clust} is the transverse energy summed over a 3×2 SC region in the second layer centered on the TOB seed and its highest-E_T neighbor in ϕ, and E_T^{env} is the transverse energy summed over the same layer and center covering 7×3 SC region.

R_{had} is the ratio of E_T in the HCAL to the E_T of the EM cluster, defined as

$$R_{\text{had}} = \frac{E_T^{\text{had}}}{E_T^{\text{EM}} + E_T^{\text{had}}} \tag{6.3}$$

where E_T^{EM} is the transverse energy summed over the four ECAL layers, and E_T^{had} is the transverse energy is a 0.3×0.3 window in $\eta - \phi$ in the HCAL.

$w_{s,\text{tot}}$ is the total lateral shower width, defined as

$$w_{s,\text{tot}} = \sqrt{\frac{\sum E_i (i - i_{\text{max}})^2}{\sum E_i}} \tag{6.4}$$

where i indexes through all the strips in a 5×3 SC window in $\eta - \phi$.

E_{ratio} is the energy balance between strip maxima, defined as

$$E_{\text{ratio}} = \frac{E_{1,\text{max}} - E_{2,\text{max}}}{E_{1,\text{max}} + E_{2,\text{max}}} \tag{6.5}$$

where $E_{i,\text{max}}$ is the max transverse energy in the ith strip.
Most of the cuts are simple, Table 16 summarizes them

Variable	Cut
E_T	> 5 GeV
$\lvert\eta\rvert$	< 1.37 (in barrel) < 2.5 (in endcap)
R_{had}	< 0.08
R_η	> 0.1
$w_{s,\text{tot}}$	< 4.0

Table 16. Simple quality cuts made on eFEX TOBs before being used as inputs to NN training.

E_{ratio} cuts are more complex, as they change depending on E_T. Within the barrel region, the cuts are shown in 17.

E_T Range	E_{ratio} Cut
< 10 GeV	> 0.59
$[10, 15)$ GeV	> 0.59
$[15, 20)$ GeV	> 0.65
$[20, 25)$ GeV	> 0.71
$[25, 30)$ GeV	> 0.73
$[30, 35)$ GeV	> 0.75
$[35, 40)$ GeV	> 0.77
$[40, 45)$ GeV	> 0.75
$[45, 50)$ GeV	> 0.79
$[50, 60)$ GeV	> 0.77
$[60, 80)$ GeV	> 0.81
$[80, 100)$ GeV	> 0.83
≥ 100 GeV	> 0.89

Table 17. Quality cuts made on E_{ratio} before being used as input to NN training.

Additionally, when performing the regression method, a quality cut was made on the PPZ value used as targets of $|PPZ| < 3000$ mm since anything greater than 3 m will not ever be relevant for photon pointing. Finally, there is also a "good event" selection cut, where only photons from events which have a best matched leading and subleading photon to EFEX TOBs are used in the training. After all these cuts, the number of events used for training and validation (with an 80-20 split) is about 50,000.

6.5.3 Signal Background Classifier.

For the Signal Background Classifier, the loss function is a Binary Cross Entropy (BCE) function, which for a given data point is defined as

$$L = -[y * \log(\hat{y}) + (1 - y) * \log(1 - \hat{y})] \tag{6.6}$$

where y is the true label (0 or 1) and $\hat{y}$ is model's predicted probability that the point belongs in the 1 label.

The signal background classifier was chosen as an initial model to have something that was easy to verify while building up the plotting infrastructure and get a baseline understanding of the performance of the network on our signal samples. This was trained on both $HH \to bb\gamma\gamma$ samples simulated at $\langle\mu\rangle = 200$ to mimic the maximum HL-LHC pileup conditions. Initially, all of the eFEX TOB information was passed as input. Figure 55 shows the resulting distributions of NN scores binned by E_T and the mean NN output score as a function of E_T. Since E_T

was included in the training, it is not surprising that it had a high dependence on E_T.

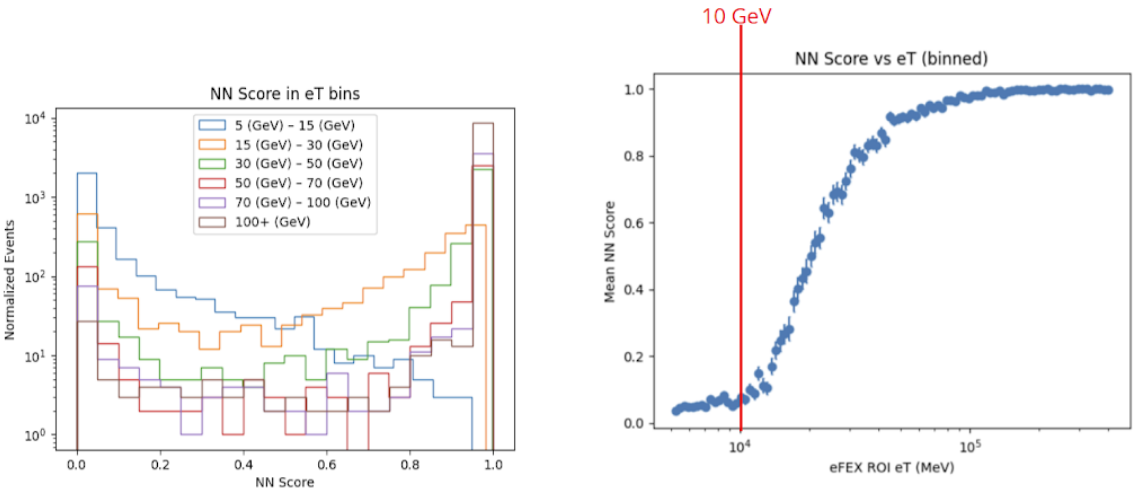


Figure 55. The NN score output distributions for different E_T bins (left) and the mean NN output as a function of E_T (right) with E_T included as a training variable. A vertical line at 10 GeV shows a theoretical desired cut on e/γ objects in Global.

After this, E_T was excluded from training, as it was essentially just teaching the NN that high E_T corresponds to signal. This is not helpful when attempting to lower overall photon p_T thresholds at the hardware trigger level. The training was repeated after removing E_T as a training variable. At the same time $Z \rightarrow ee$ samples were included to give the model a more varied input. Since e/γ is indistinguishable at the hardware trigger level, it should perform well on this sample. The resulting plots can be seen in Figure 56.

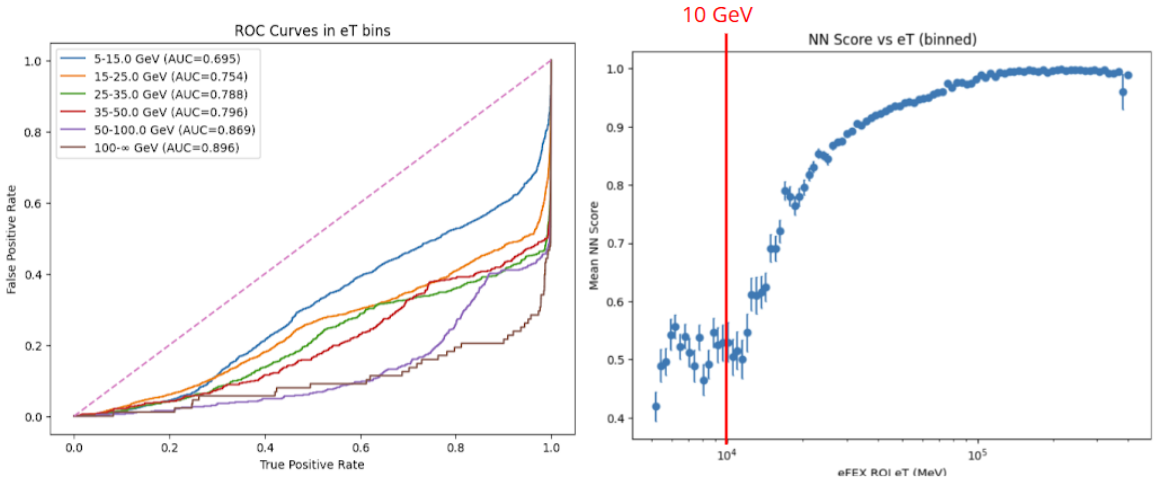


Figure 56. ROC curves binned by E_T (left) and the same mean NN score as a function of E_T (right) for the training without E_T included as an input variable. There is still a clear preference for high E_T photons.

It is clear that just providing the eFEX TOBs to identify signal is not really teaching the model anything useful, but this exercise allowed for me to build up the plotting framework and other infrastructure needed to train and evaluate the model while being fairly easy to interpret the results. This makes it easier to move the model toward a regression and quickly evaluate its performance. Additionally, doing this check allowed us to assess the quality of the complex E_{ratio} cuts. Some plots showing the mean NN score as a function of E_{ratio} for different E_T bins can be found in Appendix G

6.5.4 PPZ Regression Model.

The most intensive piece of the photon pointing algorithm, calculating the PPZ for each photon in an event, is a prime candidate for utilizing a neural net regression model. Then the remaining pieces, calculating the ΔPPZ and identifying the minimum, can be done afterwards. For the regression model, the loss function is the Mean Squared Error (MSE) function, which for a data set of n points is

$$\text{MSE} = \frac{1}{n}\sum_{i=1}^{n}(y_i - \hat{y}_i)^2 \tag{6.7}$$

where $(y_i - \hat{y}_i)$ is the residual between the truth value and the predicted value for the ith data point.

After making minimal changes to the Signal Background Classifier in order to convert it to a regression model, it was trained on the $HH \rightarrow bb\gamma\gamma$ sample for both barrel and endcap photons. Figure 57 shows the target and prediction distributions for the regression output. The predictions have a narrower width

than the target, suggesting they are learning more about the bulk low PPZ value photons than the outliers with larger PPZ values. Figure 58 shows the 2d distribution of the targets and predictions. Overall the agreement is decent for a first attempt at regression, but this model can clearly be improved upon.

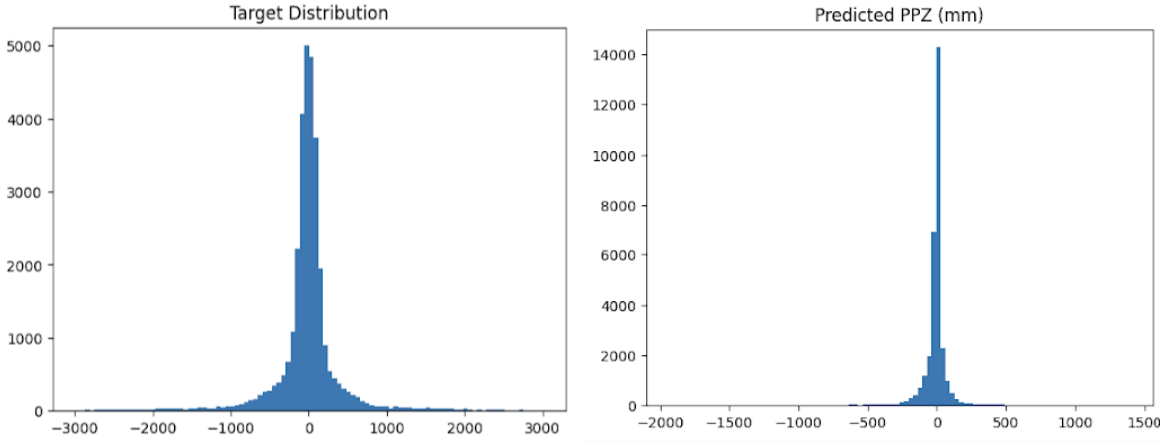


Figure 57. The PPZ value targets (left) and predictions (right) show a good first attempt at regression overall.

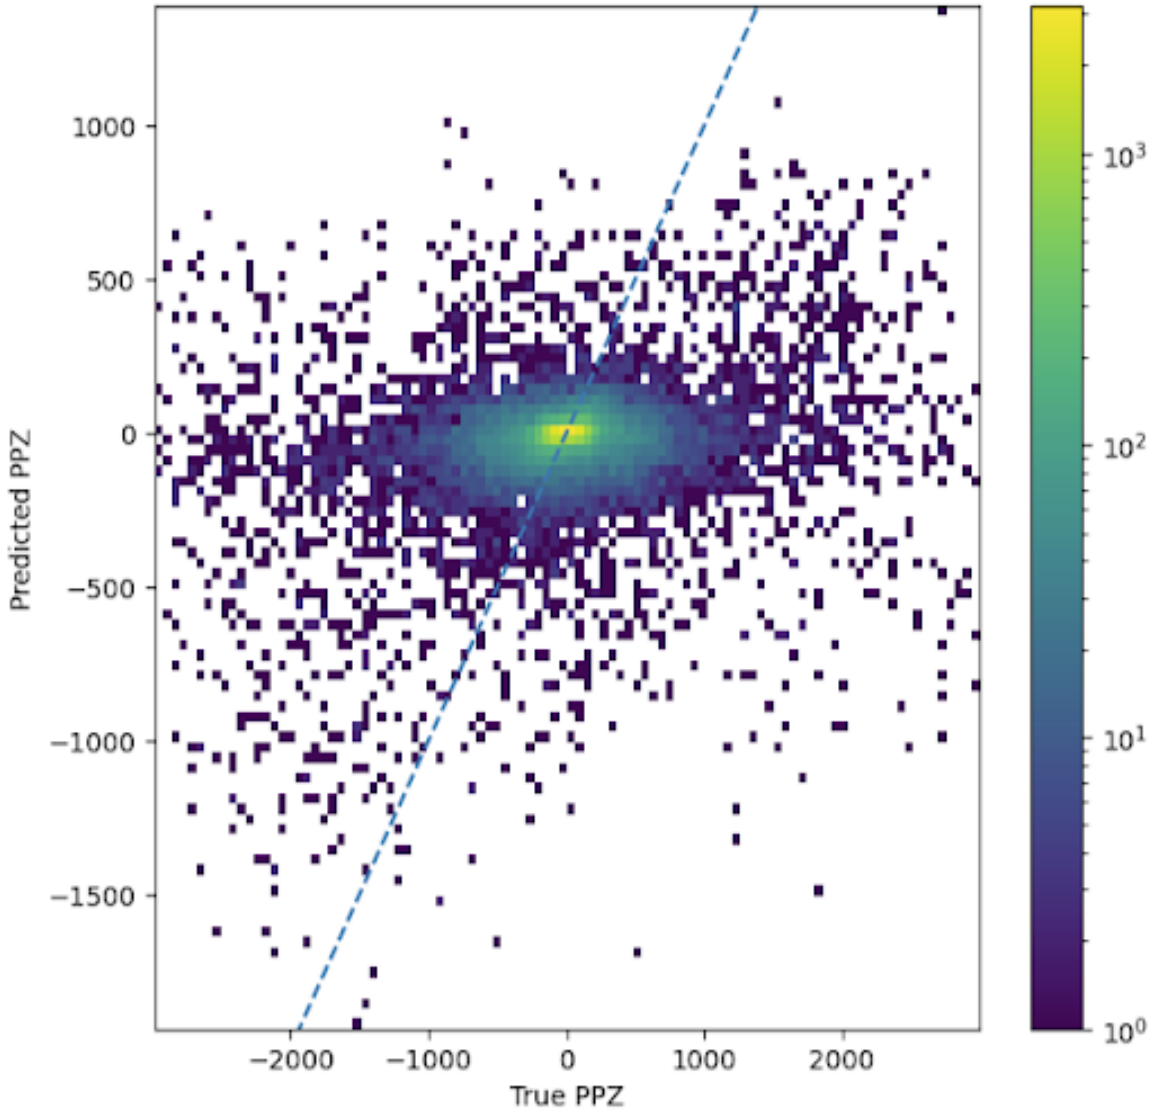


Figure 58. The overall distribution of target and prediction values for the sample. In general, the regression does fairly well, as it is clustered around the slope 1 line of agreement for the large bulk of events.

Since the barrel and end cap photons were trained together, we can also look at the prediction and target distributions for different η slices. Figure 59 shows the distribution for photons in the barrel region $|\eta| < 1.37$ and Figure 60 shows the same for the end cap region $1.37 < |\eta| < 2.5$. We can see that the endcaps perform significantly worse than the barrel region, strongly motivating us to create two different models for training those regions separately.

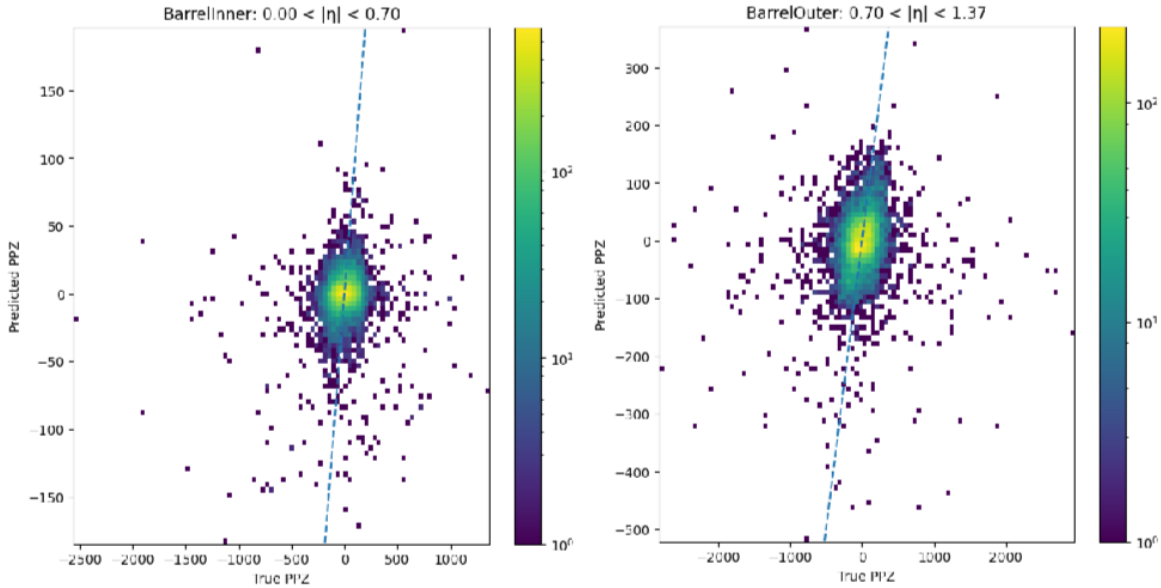


Figure 59. The PPZ value targets and predictions 2d map for photons in the barrel region, split into two different η bins to look at highly central structure. These on average perform better than the full distribution.

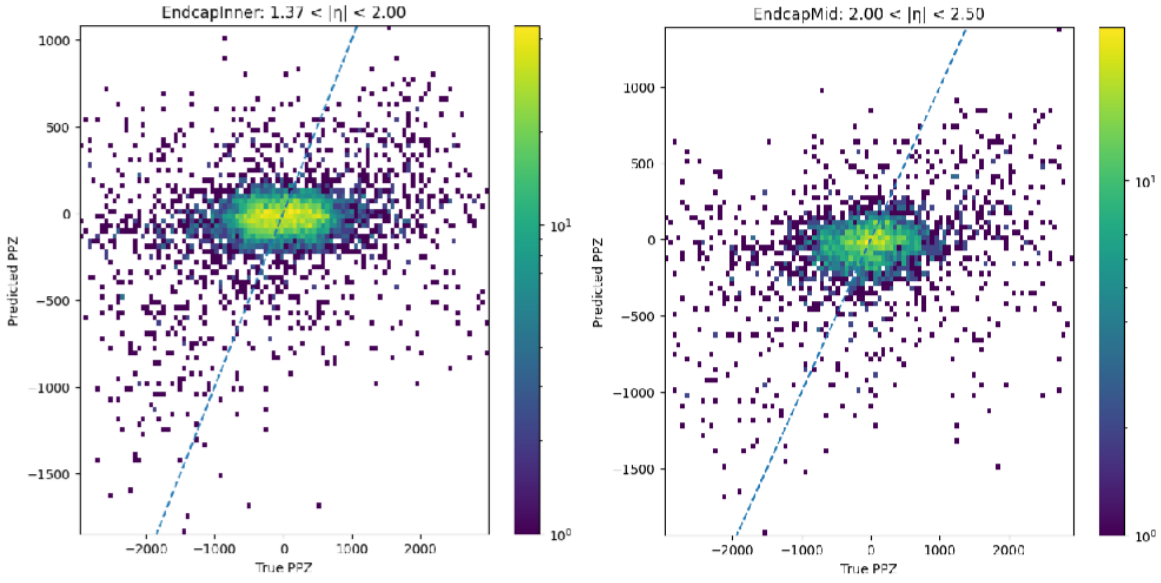


Figure 60. The PPZ value targets and predictions 2d map for photons in the endcap region, split into two different η bins to look at more forward structure. These on average perform worse than the full distribution.

6.6 Future Work on NN PPZ

By the end of this project, I have helped create a FFNN regression model that can calculate a per photon PPZ value and build up a framework of plotting infrastructure to help analyze performance. There is still much work to do on this, however. The model can be improved further by splitting the barrel and endcap, and allowing for more training (for the regression, all 50 epochs were used, as opposed to the signal background classifier, which typically capped around 20). Additional optimizations to the model can also be considered. Once a model has shown good enough performance, we also need to prepare it for deployment on an FPGA by using HLS4ML or a similar HLS tool. Finally, a new production of Run 4 samples are currently underway. Among these are some exotic Higgs couplings to ALPs that decay to photons. Running the model on these samples will help us determine its effectiveness on finding diphotons coming from LLPs. The work on this algorithm from the BNL group continues, and as a soon to be post doctoral researcher at Northern Illinois University, I will also continue to work with the BNL group on this model.

CHAPTER VII
CONCLUSION

This dissertation presented work done to improve and use the ATLAS hardware trigger system to search for physics beyond the Standard Model (BSM). It began with describing the theoretical building blocks of particle physics experiments, starting with fundamentals of the underlying theory, describing our best understanding of the universe (the SM), and providing some motivation for looking for BSM physics. In particular, we discussed the Higgs boson, and how both full top down theories as well as simple bottom up extensions of the SM call for exotic couplings of the Higgs boson to BSM particles to explain a number of outstanding questions that the SM does not answer, including the little hierarchy problem, which is concerned with how the Higgs mass is so small compared to the apparent scale of new physics.

We then discussed the LHC, the worlds largest machine, and how experiments like the ATLAS detector are built to capture the decay products of collisions that occur within the LHC. We described the system in great detail, and put a heavy focus on a particular upgrade to the Run 3 TDAQ system: the global feature extractor (gFEX). We described gFEX's goal of providing hardware trigger sensitivity to high momentum physics objects that create boosted topologies resulting in large radius jet energy patterns. The gFEX hardware and infrastructure was described in detail. The task of validating the gFEX algorithms was then described, involving the creation of a standalone python framework which has immense robustness in identifying problems with gFEX firmware builds. This resulted in the full validation of the gFEX system, allowing it to join the other Phase-I L1Calo hardware trigger systems in taking over the seeding of HLT algorithms starting in 2024.

In order to benefit from this inclusion, the progress on an analysis searching for the coupling of a Higgs boson to a long-lived neutral scalar particle (LLP) was included. The analysis utilizes new custom triggers that were not available before Run 3 including dedicated displaced jet triggers which take advantage of Large Radius Tracking being available at the HLT level, the emerging jet trigger, which uses new L1 seeds from gFEX, jFEX, and L1Topo to provide good seeding for a large-R jet which has a low fraction of prompt track p_T relative to the whole jet, and the Vector Boson Fusion (VBF) trigger, which takes advantage of new L1 capabilities to trigger on forward jets at a reasonable rate. Studies were done involving the emerging jet trigger in an effort to justify its use as a complimentary trigger to the displaced jet triggers targeting the ggF Higgs production mode. A trigger bucket strategy was defined and presented. Although the analysis is not yet complete, and the ggF region does not have a defined background estimation strategy, a mistag map method was derived in order to show the ratio of expected limits in the ggF region both with and without the inclusion of the emerging jet trigger, showing a 2 order of magnitude increase in limit setting power for the

shortest lifetime sample LLPs, while only moderately weakening the limits on the longer lifetime LLPs. The longer lifetime LLPs tend to have stronger limits already anyway, so pushing for the short lifetime limit could prove to be extremely helpful. While the VBF channel was not discussed at length, it is much more mature than the ggF, and so the analysis could be split into two separate papers. This is under discussion with the analysis team.

Finally, now that Run 3 has ended, we must also turn out attention to the coming Run 4, which marks the beginning of the High-Luminosity LHC, marked by significant increase in pileup that will result in far more data. To handle this increasingly intense environment, ATLAS is upgrading its detector and TDAQ system accordingly. Specifically, the Global Trigger will replace the current L1Topo and take inputs from the existing L1 system as well as direct input from the calorimeters. Time multiplexing will allow it to run complex algorithms. One such algorithm is the diphoton algorithm, which seeks to find pairs of nearby photons by exploiting calorimeter cell level information to do a form of photon pointing in the hardware. This algorithm exists and can be implemented on FPGAs, but here studies showing a machine learning approach to the diphoton algorithm are presented. A signal background classifier model shows a large amount of E_T dependence, but provided important foundations for moving towards a regression model which finds the per photon PPZ value. There is still work to do to optimize this model and make it suitable for HLS conversion, but this work is ongoing with the diphoton trigger algorithm team at BNL.

APPENDIX A
AN EXAMPLE GFEX ROB FRAGMENT

This table shows an example gFEX ROB fragment for one event. This contains some header and footer information, as well as the input fibers for all three pFPGAs, and the calculated TOBs for the current BC and the two next BCs.

```
aa1234aa 0000099e 00000017 05000000 007c0000 00000000 00000000 69133da6 00000000 00000000
00000000 0000000f 0007c94e 00000000 01000054 000005c0 00000005 00000000 00000987 00000000
00000000 00000000 00000000 dd1234dd 00000858 00000009 05000000 00933000 00000002 00000000
00000000 00000000 ee1234ee 00000009 03011004 00933000 0007c94e 01000054 000005c0 00000005
00000000 a14002bc 00000000 00000000 00000000 00000000 00000000 00000000 00400000 00000000
00000000 00000000 00000000 00000000 00000000 00400000 00000000 00000000 00000000 00000000
00000000 00000000 00400000 00000000 00000000 00000000 00000000 00000000 00000000 00400000
05005000 000c0000 01100600 0b006005 500b0050 00600500 39c058bc 0500b00c 00060000 00000600
0500b005 000000b0 00600000 db4058bc 00000005 001100b0 00000a00 0600500b 500500c0 00000600
324058bc 06005000 00060000 00b00000 0a006006 600a00a0 00000000 c34058bc 0600000c 00050050
00000000 0f006000 00120050 00601000 cdc058bc 00006006 50050060 00000500 0a006005 00000060
00000b00 a2c058bc 0a00f00f c00b0050 00b00600 0b00600a 00060060 01000500 6cc058bc 05006005
b0000000 00500000 00005006 600a0060 00500500 c1c058bc 00000000 00000000 00000000 00000000
00000000 00000000 00400000 00000000 00000000 00000000 00000000 00000000 00000000 00400000
00000000 00000000 00000000 00000000 00000000 00000000 00400000 00000000 00000000 00000000
00000000 00000000 00000000 00400000 00000000 00000000 00000000 00000000 00000000 00000000
00400000 00000000 00000000 00000000 00000000 00000000 00000000 00400000 00000000 00000000
00000000 00000000 00000000 00000000 00400000 00000000 00000000 00000000 00000000 00000000
00000000 00400000 00000000 00000000 00000000 00000000 00000000 00000000 00400000 00000000
00000000 00000000 00000000 00000000 00400000 00000000 00000000 00000000 00000000 00000000
00000000 00000000 00400000 00000000 00000000 00000000 00000000 00000000 00000000 00400000
00000000 00000000 00000000 00000000 00000000 00000000 00400000 00000000 00000000 00000000
00000000 00000000 00000000 00400000 00000000 00000000 00000000 00000000 00000000 00000000
00400000 00000000 00000000 00000000 00000000 00000000 00000000 00400000 06000000 00000050
00a00b00 0600b006 a00600c0 00000500 cb4058bc 00006005 60050050 00500600 0000c000 600000b0
00600000 b74058bc 00000005 600600a0 00000000 0b006000 000c00b0 00000500 494058bc 06006000
500600a0 00500500 05000005 00000060 00600500 bec058bc 0600a005 60060060 00a00500 05000006
b00c00b0 00000a00 834058bc 0a005005 600000a0 00500500 00000000 b00f00b0 00500500 524058bc
1000500a 500a0060 00b00500 06006000 60150000 00500500 38c058bc 0a00c005 a00b0000 00000000
0500600a f00000c0 00a00600 334058bc 00000000 00000000 00000000 00000000 00000000 00000000
00400000 00000000 00000000 00000000 00000000 00000000 00000000 00400000 00000000 00000000
00000000 00000000 00000000 00000000 00400000 00000000 00000000 00000000 00000000 00000000
00000000 00400000 00000000 00000000 00000000 00000000 00000000 00000000 00400000 00000000
00000000 00000000 00000000 00000000 00000000 00400000 00000000 00000000 00000000 00000000
00000000 00000000 00400000 00000000 00000000 00000000 00000000 00000000 00000000 00400000
00000000 00000000 00000000 00000000 00000000 00000000 00400000 00000000 00000000 00000000
00000000 00000000 00000000 00400000 00000000 00000000 00000000 00000000 00000000 00000000
00400000 00000000 00000000 00000000 00000000 00000000 00000000 00400000 00000000 00000000
00000000 00000000 00000000 00000000 00400000 00000000 00000000 00000000 00000000 00000000
00000000 00000000 00400000 00000000 00000000 00000000 00000000 00000000 00000000 00400000
00000000 00000000 00000000 00000000 00000000 00000000 00400000 00000000 00000000 00000000
00000000 00000000 00000000 00400000 00000000 00000000 00000000 00000000 00000000 00000000
00400000 00000000 00000000 00000000 00000000 00000000 00000000 00400000 00000000 00000000
00000000 00000000 00000000 00000000 00400000 00000000 00000000 00000000 00000000 00000000
00000000 00400000 00000000 00000000 00000000 00000000 00000000 00000000 00400000 00000000
00000000 00000000 00000000 00000000 00000000 00400000 00000000 00000000 00000000 00000000
00000000 00000000 00400000 00000000 00000000 00000000 00000000 00000000 00000000 00400000
00000000 00000000 00000000 00000000 00000000 00000000 00400000 00000000 00000000 00000000
00400000 00000000 00000000 00000000 00000000 00000000 00000000 00400000 00000000 00000000
00000000 00000000 00000000 00400000 00000000 00000000 00000000 00000000 00000000 00000000
00000000 00400000 00000000 00000000 00000000 00000000 00000000 00000000 00400000 00000000
```

```
00000000 00000000 00000000 00000000 00000000 00400000 00000000 00000000 00000000 00000000
00000000 00000000 00400000 00000000 00000000 00000000 00000000 00000000 00000000 00400000
00000000 00000000 00000000 00000000 00000000 00000000 00400000 00000000 00000000 00000000
00000000 00000000 00000000 00400000 00000000 00000000 00000000 00000000 00000000 00000000
00400000 00000000 00000000 00000000 00000000 00000000 00000000 00400000 00000000 00000000
00000000 00000000 00000000 00000000 00400000 00000000 00000000 00000000 00000000 00000000
00000000 00400000 00000000 00000000 00000000 00000000 00000000 00000000 00400000 00000000
00000000 00000000 00000000 00000000 00000000 00400000 00000000 00000000 00000000 00000000
00000000 00000000 00400000 00000000 00000000 00000000 00000000 00000000 00000000 00400000
00000000 00000000 00000000 00000000 00000000 00000000 00400000 00000000 00000000 00000000
00000000 00000000 00000000 00400000 00000000 00000000 00000000 00000000 00000000 00000000
00400000 00000000 00000000 00000000 00000000 00000000 00000000 00400000 00000000 00000000
00000000 00000000 00000000 00000000 00400000 00000000 00000000 00000000 00000000 00000000
00000000 004000bc b14002bc 00000000 00000000 00000000 00000000 00000000 00000000 00400000
00000000 00000000 00000000 00000000 00000000 00000000 00400000 00000000 00000000 00000000
00000000 00000000 00000000 00400000 00000000 00000000 00000000 00000000 00000000 00000000
00400000 00000000 00000000 00000000 00000000 00000000 00000000 00400000 00000000 00000000
00000000 00000000 00000000 00400000 00000000 00000000 00000000 00000000 00000000 00000000
00000000 00400000 00000000 00000000 00000000 00000000 00000000 00000000 00400000 00000000
00000000 00000000 00000000 00000000 00000000 00400000 00000000 00000000 00000000 00000000
00000000 00000000 00400000 00000000 00000000 00000000 00000000 00000000 00000000 00400000
00000000 00000000 00000000 00000000 00000000 00000000 00400000 00000000 00000000 00000000
00000000 00000000 00400000 00000000 00000000 00000000 00000000 00000000 00000000 00000000
00400000 00000000 00000000 00000000 00000000 00000000 00000000 00400000 00000000 00000000
00000000 00000000 00000000 00000000 00400000 0a000005 600b00c0 00a00b00 0c005010 60000060
00500600 67c058bc 06005005 00050000 00500b00 0600a006 c0060060 00500000 114058bc 0b005006
600a0000 00501100 10000006 500c0050 00600500 99c058bc 0b00b005 000a0050 00b00600 05011006
00000050 00000a00 a94058bc 10000005 600b0050 00a00600 0a000006 500600a0 00000000 0fc058bc
0000000b 000b00b0 00b00600 0a000006 60000050 00b00500 4cc058bc 05005000 10000050 00a00c01
0600000a b00b00c0 00a01500 814058bc 0500000c 50060050 00500500 05000011 a0120060 00500000
44c058bc 00000000 00000000 00000000 00000000 00000000 00000000 00400000 00000000 00000000
00000000 00000000 00000000 00000000 00400000 00000000 00000000 00000000 00000000 00000000
00000000 00400000 00000000 00000000 00000000 00000000 00000000 00000000 00400000 00000000
00000000 00000000 00000000 00000000 00000000 00400000 00000000 00000000 00000000 00000000
00000000 00000000 00400000 00000000 00000000 00000000 00000000 00000000 00000000 00400000
00000000 00000000 00000000 00000000 00000000 00000000 00400000 00000000 00000000 00000000
00000000 00000000 00000000 00400000 00000000 00000000 00000000 00000000 00000000 00000000
00400000 00000000 00000000 00000000 00000000 00000000 00000000 00400000 00000000 00000000
00000000 00000000 00000000 00000000 00400000 00000000 00000000 00000000 00000000 00000000
00000000 00400000 00000000 00000000 00000000 00000000 00000000 00000000 00400000 0b005005
600500a0 00600000 0000600b 500b00b0 00600600 254058bc 0c00b000 c0000050 00b00600 0000000a
00000060 01000000 1b4058bc 00000000 00000000 00000000 00000000 00000000 00000000 00400000
00000000 00000000 00000000 00000000 00000000 00000000 00400000 05000005 60060000 00600b00
06006005 000000a0 00601100 9a4058bc 0000a00b 000b0060 00500000 1200500b 600a0050 00000b00
06c058bc 0500500b 00050050 00600a00 0500a000 000000b0 00500500 1ec058bc 0000000a 50050060
00500000 0600c006 a0050050 00b00600 fa4058bc 0b000000 000600a0 00a00b00 0600b006 50060060
00000500 3c4058bc 0a00c000 000600b0 00a00c00 1100b000 60000050 00600600 65c058bc 00000000
00000000 00000000 00000000 00000000 00000000 00400000 00000000 00000000 00000000 00000000
00000000 00000000 00400000 00000000 00000000 00000000 00000000 00000000 00000000 00400000
00000000 00000000 00000000 00000000 00000000 00000000 00400000 00000000 00000000 00000000
00000000 00000000 00000000 00400000 00000000 00000000 00000000 00000000 00000000 00000000
00400000 00000000 00000000 00000000 00000000 00000000 00000000 00400000 00000000 00000000
00000000 00000000 00000000 00000000 00400000 00000000 00000000 00000000 00000000 00000000
00000000 00400000 00000000 00000000 00000000 00000000 00000000 00000000 00400000 00000000
00000000 00000000 00000000 00000000 00000000 00400000 00000000 00000000 00000000 00000000
00000000 00000000 00400000 00000000 00000000 00000000 00000000 00000000 00000000 00400000
00000000 00000000 00000000 00000000 00000000 00000000 00400000 00000000 00000000 00000000
00000000 00000000 00000000 00400000 00000000 00000000 00000000 00000000 00000000 00000000
00400000 00000000 00000000 00000000 00000000 00000000 00000000 00400000 00000000 00000000
00000000 00000000 00000000 00000000 00000000 00400000 00000000 00000000 00000000 00000000
00000000 00400000 00000000 00000000 00000000 00000000 00000000 00000000 00400000 00000000
00000000 00000000 00000000 00000000 00000000 00400000 00000000 00000000 00000000 00000000
00000000 00000000 00400000 00000000 00000000 00000000 00000000 00000000 00000000 00400000
00000000 00000000 00000000 00000000 00000000 00000000 00400000 00000000 00000000 00000000
00400000 00000000 00000000 00000000 00000000 00000000 00000000 00400000 00000000 00000000
00000000 00000000 00000000 00000000 00400000 00000000 00000000 00000000 00000000 00000000
00000000 00400000 00000000 00000000 00000000 00000000 00000000 00000000 00400000 00000000
00000000 00000000 00000000 00000000 00000000 00400000 00000000 00000000 00000000 00000000
00000000 00000000 00400000 00000000 00000000 00000000 00000000 00000000 00000000 00400000
00000000 00000000 00000000 00400000 00000000 00000000 00000000 00000000 00000000 00000000
```

00400000 00000000 00000000 00000000 00000000 00000000 00000000 00400000 00000000 00000000
00000000 00000000 00000000 00000000 00400000 00000000 00000000 00000000 00000000 00000000
00000000 00400000 00000000 00000000 00000000 00000000 00000000 00000000 00400000 00000000
00000000 00000000 00000000 00000000 00000000 00400000 00000000 00000000 00000000 00000000
00000000 00000000 00400000 00000000 00000000 00000000 00000000 00000000 00000000 00400000
00000000 00000000 00000000 00000000 00000000 00000000 00400000 00000000 00000000 00000000
00000000 00000000 00400000 00000000 00000000 00000000 00000000 00000000 00000000 00000000
00400000 00000000 00000000 00000000 00000000 00000000 00000000 00400000 00000000 00000000
00000000 00000000 00000000 00000000 00400000 00000000 00000000 00000000 00000000 00000000
00000000 00400000 00000000 00000000 00000000 00000000 00000000 00000000 00400000 00000000
00000000 00000000 00000000 00000000 00000000 00400000 00000000 00000000 00000000 00000000
00000000 00000000 004000bc c14002bc 00000000 00000000 00000000 00000000 00000000 00000000
00400000 00000000 00000000 00000000 00000000 00000000 00000000 00400000 00000000 00000000
00000000 00000000 00000000 00000000 00400000 00000000 00000000 00000000 00000000 00000000
00000000 00400000 00000000 00000000 00000000 00000000 00000000 00000000 00400000 00000000
00000000 00000000 00000000 00000000 00000000 00400000 00000000 00000000 00000000 00000000
00000000 00000000 00400000 00000000 00000000 00000000 00000000 00000000 00000000 00400000
00000000 00000000 00000000 00000000 00000000 00000000 00400000 00000000 00000000 00000000
00400000 00000000 00000000 00000000 00000000 00000000 00000000 00400000 00000000 00000000
00000000 00000000 00000000 00000000 00400000 00000000 00000000 00000000 00000000 00000000
00000000 00400000 00000000 00000000 00000000 00000000 00000000 00000000 00400000 00000000
00000000 00000000 00000000 00000000 00000000 00400000 00000000 00000000 00000000 00000000
00000000 00000000 00400000 00000000 00000000 00000000 00000000 00000000 00000000 00400000
00000000 00000000 00000000 00000000 00000000 00000000 00400000 00000000 00000000 00000000
00000000 00400000 00000000 00000000 00000000 00000000 00000000 00000000 00400000 00000000
00400000 00000000 00000000 00000000 00000000 00000000 00000000 00400000 00000000 00000000
00000000 00000000 00000000 00400000 00000000 00000000 00000000 00000000 00000000 00000000
00000000 00000000 00400000 00000000 00000000 00000000 00000000 00000000 00000000 00400000
00000000 00000000 00000000 00000000 00000000 00400000 00000000 00000000 00000000 00000000
00000000 00000000 00000000 00000000 00400000 00000000 00000000 00000000 00000000 00000000
00000000 00400000 00000000 00000000 00000000 00000000 00000000 00000000 00400000 00000000
00000000 00000000 00000000 00000000 00000000 00400000 00000000 00000000 00000000 00000000
00000000 00000000 00400000 00000000 00000000 00000000 00000000 00000000 00000000 00400000
00000000 00000000 00000000 00000000 00000000 00000000 00400000 00000000 00000000 00000000
00000000 00000000 00000000 00400000 00000000 00000000 00000000 00000000 00000000 00000000
00400000 00000000 00000000 00000000 00000000 00000000 00000000 00400000 00000000 00000000
00000000 00000000 00000000 00000000 00400000 00000000 00000000 00000000 00000000 00000000
00000000 00400000 00000000 00000000 00000000 00000000 00000000 00000000 00400000 00000000
00000000 00000000 00000000 00000000 00000000 00400000 00000000 00000000 00000000 00000000
00000000 00000000 00400000 00000000 00000000 00000000 00000000 00000000 00000000 00400000
00000000 00000000 00000000 00000000 00000000 00000000 00400000 00000000 00000000 00000000
00000000 00000000 00000000 00400000 00000000 00000000 00000000 00000000 00000000 00000000
00400000 00000000 00000000 00000000 00000000 00000000 00000000 00400000 00000000 00000000
00000000 00000000 00000000 00000000 00400000 00000000 00000000 00000000 00000000 00000000
00000000 00400000 00000000 00000000 00000000 00000000 00000000 00000000 00400000 00000000
00000000 00000000 00000000 00000000 00000000 00400000 00000000 00000000 00000000 00000000
00000000 00000000 00400000 00000000 00000000 00000000 00000000 00000000 00000000 00400000
00000000 00000000 00000000 00000000 00000000 00000000 00400000 00000000 00000000 00000000
00400000 00000000 00000000 00000000 00000000 00000000 00000000 00400000 00000000 00000000
00000000 00000000 00000000 00000000 00400000 00000000 00000000 00000000 00000000 00000000
00000000 00400000 00000000 00000000 00000000 00000000 00000000 00000000 00400000 00000000
00000000 00000000 00000000 00000000 00000000 00400000 00000000 00000000 00000000 00000000
00000000 00000000 00400000 00000000 00000000 00000000 00000000 00000000 00000000 00400000
00000000 00000000 00000000 00000000 00000000 00000000 00400000 00000000 00000000 00000000
00000000 00000000 00000000 00400000 00000000 00000000 00000000 00000000 00000000 00000000
00400000 00000000 00000000 00000000 00000000 00000000 00000000 00400000 00000000 00000000
00000000 00000000 00000000 00000000 00400000 00000000 00000000 00000000 00000000 00000000
00000000 00400000 00000000 00000000 00000000 00000000 00000000 00000000 00400000 00000000
00000000 00000000 00000000 00000000 00000000 00400000 00000000 00000000 00000000 00000000
00000000 00000000 00400000 00000000 00000000 00000000 00000000 00000000 00000000 00400000
00000000 00000000 00000000 00000000 00000000 00000000 00400000 00000000 00000000 00000000
00000000 00000000 00000000 00400000 00000000 00000000 00000000 00000000 00000000 00000000
00400000 00000000 00000000 00000000 00000000 00000000 00000000 00400000 00000000 00000000
00000000 00000000 00000000 00000000 00400000 00000000 00000000 00000000 00000000 00000000
00000000 00400000 00000000 00000000 00000000 00000000 00000000 00000000 00400000 00000000
00000000 00000000 00000000 00000000 00000000 00400000 00000000 00000000 00000000 00000000
00000000 00000000 00400000 00000000 00000000 00000000 00000000 00000000 00000000 00400000
00000000 00000000 00000000 00000000 00000000 00000000 00400000 00000000 00000000 00000000
00000000 00000000 00000000 00400000 00000000 00000000 00000000 00000000 00000000 00000000
00400000 00000000 00000000 00000000 00000000 00000000 00000000 00400000 00000000 00000000
00000000 00000000 00000000 00000000 00400000 00000000 00000000 00000000 00000000 00000000
00000000 00400000 00000000 00000000 00000000 00000000 00000000 00000000 00400000 00000000
00000000 00000000 00000000 00000000 00000000 00400000 00000000 00000000 00000000 00000000

```
00000000 00000000 00400000 00000000 00000000 00000000 00000000 00000000 00000000 00400000
00000000 00000000 00000000 00000000 00000000 00000000 00400000 00000000 00000000 00000000
00000000 00000000 00000000 00400000 00000000 00000000 00000000 00000000 00000000 00000000
00400000 00000000 00000000 00000000 00000000 00000000 00000000 00400000 00000000 00000000
00000000 00000000 00000000 00000000 00400000 00000000 00000000 00000000 00000000 00000000
00000000 00400000 00000000 00000000 00000000 00000000 00000000 00000000 00400000 00000000
00000000 00000000 00000000 00000000 00000000 00400000 00000000 00000000 00000000 00000000
00000000 00000000 00400000 00000000 00000000 00000000 00000000 00000000 00000000 00400000
00000000 00000000 00000000 00000000 00000000 00000000 00400000 00000000 00000000 00000000
00000000 00000000 00000000 004000bc a1400003 00000000 d0000002 d0000003 b1400003 00000000
d0000002 d0000003 c1400003 00000000 d0000002 d0000003 0000000c 00000837 00000001 dd1234dd
0000012f 00000009 05000000 00943000 00000002 00000000 00000000 00000000 ee1234ee 00000009
03011004 00943000 0007c94e 01000054 000005c0 00000005 00000000 a140002a 00000000 04c18e81
0cc14b83 34d41185 00000000 00004000 000030bc 0002c000 2cf19582 58f17084 51061986 00000000
00004000 000030bc 00000000 4cc15c81 78c12c83 4cd59585 00000000 00004100 000031bc 0002a400
51022482 1cf20684 4507c186 00000000 00004100 000031bc 00000000 1cc17e81 3cc14a83 68d48c85
00000000 00004200 000032bc 00028a00 7501f082 7cf1a784 69073686 00000000 00004200 000032bc
b140002a 00000000 71619481 71818e83 5d76e185 00000000 00004000 000030bc 0002c500 59a19a82
61a16f84 65a51186 00000000 00004000 000030bc 00000000 1d819c81 2d618083 71765a85 00000000
00004100 000031bc 0002a100 79a15e82 05a14584 05a4f586 00000000 00004100 000031bc 00000000
4d81e481 01519283 2576fa85 00000000 00004200 000032bc 00027d00 31a19082 21a16684 29a52886
00000000 00004200 000032bc c140002a 00000000 00600001 00600003 00200005 00000000 00004000
000030bc 00000000 02400002 02400004 02000006 00000000 00004000 000030bc 00000000 00600001
00600003 00200005 00000000 00004100 000031bc 00000000 02400002 02400004 02000006 00000000
00004100 000031bc 00000000 00600001 00600003 00200005 00000000 00004200 000032bc 00000000
02400002 02400004 02000006 00000000 00004200 000032bc 1140002a fef2002f 00000000 fef7002e
00001b03 00001b94 00000000 00173000 fef2002f 00000000 00000000 00000000 00000000 00000000
00173000 fefdffb0 000dffdf ff08ffa4 00001e42 00001ee4 00000022 00173100 ff0aff8f 00000000
00000000 00000000 00000000 00000000 00173100 ffb3ff95 00000000 ffb4ff97 00001d40 00001ddd
00000000 00173200 ffb3ff95 00000000 00000000 00000000 00000000 00000000 00173200 2140002a
0047fe5a 0000ffec 0045fe5a 00001d5b 00001df8 00000013 00173000 0047fe46 00000000 00000000
00000000 00000000 00000000 00173000 00ff0073 0019000a 01040075 00001bd8 00001c6d 0000001b
00173100 0119007d 00000000 00000000 00000000 00000000 00000000 00173100 ffc10019 ffd7000b
ffb2001d 00001d6e 00001e0b 00000051 00173200 ff990024 00000000 00000000 00000000 00000000
00000000 00173200 3140002a 00000000 00000000 00000000 00000000 00000000 00000000 00173000
00000000 00000000 00000000 00000000 00000000 00000000 00173000 00000000 00000000 00000000
00000000 00000000 00000000 00173100 00000000 00000000 00000000 00000000 00000000 00000000
00173100 00000000 00000000 00000000 00000000 00000000 00000000 00173200 00000000 00000000
00000000 00000000 00000000 00000000 00000000 00173200 a1400003 00000000 d0000002 d0000003 b1400003
00000000 d0000002 d0000003 c1400003 00000000 d0000002 d0000003 11400003 00000000 d0000002
d0000003 21400003 00000000 d0000002 d0000003 31400003 00000000 d0000002 d0000003 00000018
00000102 00000001
```

APPENDIX B
HIGGS TO DISPLACED JET RUN 3 ANALYSIS SAMPLES LIST

This Appendix describes the Data and Monte Carlo samples used in the Higgs to Displaced Jet Run 3 analysis.

Year	**Datasets**
2024	data24_13p6TeV.periodAllYear.physics_Main. PhysCont.DAOD_LLP1.grp24_v01_p7079 data24_13p6TeV.periodAllYear.physics_VBFDelayed. PhysCont.DAOD_LLP1.grp24_v01_p7079
2023	data23_13p6TeV.periodAllYear.physics_Main. PhysCont.DAOD_LLP1.grp23_v01_p7079 data23_13p6TeV.periodAllYear.physics_VBFDelayed. PhysCont.DAOD_LLP1.grp23_v01_p7079
2022	data22_13p6TeV.periodAllYear.physics_Main. PhysCont.DAOD_LLP1.grp22_v01_p7079 data22_13p6TeV.periodAllYear.physics_VBFDelayed. PhysCont.DAOD_LLP1.grp22_v01_p7079

Table B.1. Datasets used for each data-taking year (Physics Main and VBF Delayed streams).

Dataset	**Good Run List**	**Luminosity**
data22	data22_13p6TeV.periodAllYear_DetStatus-v134-pro2809_MERGED_PHYS_StandardGRL_All_Good_25ns.xml	26328.8 pb^{-1}
data23	data23_13p6TeV.periodAllYear_DetStatus-v133-pro3111_MERGED_PHYS_StandardGRL_All_Good_25ns.xml	25204.3 pb^{-1}
data24	physics_25ns_data24_periodsEtoO_reproc.xml	107.9 fb^{-1}

Table B.2. Datasets, corresponding Good Run Lists, and integrated luminosities.

Higgs production	**LHE file**
ggF	mc23_13p6TeV.601485.Ph_PDF4LHC21_ggH125_MiNLO_LHE.evgen.TXT.e8537
VBF	mc23_13p6TeV.601401.Ph_PDF4LHC21_VBFH125_LHE.evgen.TXT.e8537
ttH	mc23_13p6TeV.601452.Ph_PDF4LHC21_dilep_ttH125_LHE.evgen.TXT.e8537
W^+H	mc23_13p6TeV.602408.Ph_PDF4LHC21_WpH125_VpTbias_LHE.evgen.TXT.e8544
W^-H	mc23_13p6TeV.602407.Ph_PDF4LHC21_WmH125_VpTbias_LHE.evgen.TXT.e8544
$qqZ(\to \ell\ell)H$	mc23_13p6TeV.602409.Ph_PDF4LHC21_qqZllH125_VpTbias_LHE.evgen.TXT.e8544
$ggZ(\to \ell\ell)H$	mc23_13p6TeV.602411.Ph_PDF4LHC21_ggZllH125_LHE.evgen.TXT.e8544

Table B.3. List of LHE files used for the different Higgs boson production mechanisms.

APPENDIX C
DISPLACED JETS ANALYSIS TRIGGER CHAINS

This appendix contains the full HLT trigger chains for the triggers used to define the Signal Regions in the Higgs to Displaced Jets partial Run 3 analysis.

Year	HLT Trigger Chain	Trigger Name/Description	Purpose
2022	HLT_j200_0eta180_emergingPTF0p08dR1p2_a10sd_cssk_pf_jes_ftf_preselj200_L1J100	EJ Trigger	ggF SR
2023	HLT_j200_0eta180_emergingPTF0p08dR1p2_a10sd_cssk_pf_jes_ftf_preselj200_L1J100	EJ Trigger	ggF SR
2024	HLT_j200_0eta180_emergingPTF0p08dR1p2_a10sd_cssk_pf_jes_ftf_preselj200_L1gLJ140p0ETA25	EJ Trigger gFEX seed	ggF SR
2024	HLT_j200_0eta180_emergingPTF0p08dR1p2_a10sd_cssk_pf_jes_ftf_preselj200_L1SC111-CjJ40	EJ Trigger Simple Cone seed	ggF SR
2024	HLT_j200_0eta180_emergingPTF0p08dR1p2_a10sd_cssk_pf_jes_ftf_preselj200_L1jJ160	EJ Trigger jFEX seed	ggF SR
2022	HLT_j180_dispjet50_3d2p_dispjet50_1p_L1J100	Asym DJ Trigger	ggF SR
2023	HLT_j180_dispjet50_3d2p_dispjet50_1p_L1J100	Asym DJ Trigger	ggF SR
2024	HLT_j180_dispjet50_3d2p_dispjet50_1p_L1jJ160	Asym DJ Trigger	ggF SR
2022	HLT_j180_2dispjet50_3d2p_L1J100	Sym DJ Trigger	ggF SR
2023	HLT_j180_2dispjet50_3d2p_L1J100	Sym DJ Trigger	ggF SR
2024	HLT_j180_2dispjet50_3d2p_L1jJ160	Sym DJ Trigger	ggF SR

APPENDIX D
INCLUSIVE TRIGGER EFFICIENCY STUDIES

Mass/Lifetime	EJ Excl	EJ + presel Excl	DJ Excl	DJ + presel Excl
m = 3GeV $c\tau$ = 1mm	4.20 ± 0.08 %	1.97 ± 0.06 %	1.08 ± 0.04 %	0.31 ± 0.02 %
m = 5GeV $c\tau$ = 1mm	4.40 ± 0.08 %	1.59 ± 0.05 %	0.24 ± 0.02 %	0.04 ± 0.01 %
m = 10GeV $c\tau$ = 1mm	3.51 ± 0.08 %	1.06 ± 0.04 %	0.15 ± 0.02 %	0.02 ± 0.01 %
m = 15GeV $c\tau$ = 1mm	4.46 ± 0.08 %	1.16 ± 0.04 %	0.13 ± 0.01 %	0.03 ± 0.01 %
m = 16GeV $c\tau$ = 1mm	4.61 ± 0.09 %	1.24 ± 0.05 %	0.18 ± 0.02 %	0.05 ± 0.01 %
m = 20GeV $c\tau$ = 1mm	4.48 ± 0.08 %	1.07 ± 0.04 %	0.14 ± 0.02 %	0.02 ± 0.01 %
m = 30GeV $c\tau$ = 1mm	4.68 ± 0.09 %	1.15 ± 0.04 %	0.14 ± 0.02 %	0.02 ± 0.01 %
m = 40GeV $c\tau$ = 1mm	4.53 ± 0.09 %	1.09 ± 0.04 %	0.17 ± 0.02 %	0.02 ± 0.01 %
m = 55GeV $c\tau$ = 1mm	4.35 ± 0.06 %	0.97 ± 0.03 %	0.18 ± 0.01 %	0.01 ± 0.00 %
m = 3GeV $c\tau$ = 10mm	4.42 ± 0.09 %	2.43 ± 0.06 %	3.02 ± 0.07 %	0.98 ± 0.04 %
m = 5GeV $c\tau$ = 10mm	4.83 ± 0.09 %	2.43 ± 0.06 %	2.46 ± 0.06 %	0.89 ± 0.04 %
m = 10GeV $c\tau$ = 10mm	4.88 ± 0.09 %	2.31 ± 0.06 %	2.01 ± 0.06 %	0.75 ± 0.04 %
m = 15GeV $c\tau$ = 10mm	4.97 ± 0.09 %	2.01 ± 0.06 %	1.69 ± 0.05 %	0.61 ± 0.03 %
m = 16GeV $c\tau$ = 10mm	4.88 ± 0.07 %	2.05 ± 0.05 %	1.76 ± 0.04 %	0.68 ± 0.03 %
m = 20GeV $c\tau$ = 10mm	4.99 ± 0.09 %	2.07 ± 0.06 %	1.68 ± 0.05 %	0.64 ± 0.03 %
m = 30GeV $c\tau$ = 10mm	5.51 ± 0.09 %	2.44 ± 0.06 %	1.79 ± 0.05 %	0.46 ± 0.03 %
m = 40GeV $c\tau$ = 10mm	5.58 ± 0.09 %	2.53 ± 0.06 %	2.39 ± 0.06 %	0.62 ± 0.03 %
m = 55GeV $c\tau$ = 10mm	5.37 ± 0.07 %	2.37 ± 0.05 %	2.62 ± 0.05 %	0.65 ± 0.03 %
m = 3GeV $c\tau$ = 100mm	1.00 ± 0.04 %	0.41 ± 0.03 %	1.13 ± 0.05 %	0.24 ± 0.02 %
m = 5GeV $c\tau$ = 100mm	1.63 ± 0.06 %	0.73 ± 0.04 %	1.85 ± 0.06 %	0.52 ± 0.03 %
m = 10GeV $c\tau$ = 100mm	2.76 ± 0.07 %	1.34 ± 0.05 %	2.62 ± 0.07 %	0.96 ± 0.04 %
m = 15GeV $c\tau$ = 100mm	3.10 ± 0.08 %	1.46 ± 0.06 %	3.29 ± 0.08 %	1.23 ± 0.05 %
m = 16GeV $c\tau$ = 100mm	3.18 ± 0.06 %	1.50 ± 0.04 %	3.41 ± 0.06 %	1.33 ± 0.04 %
m = 20GeV $c\tau$ = 100mm	3.50 ± 0.08 %	1.68 ± 0.05 %	3.59 ± 0.08 %	1.48 ± 0.05 %
m = 30GeV $c\tau$ = 100mm	4.17 ± 0.09 %	2.34 ± 0.07 %	3.99 ± 0.09 %	1.53 ± 0.06 %
m = 40GeV $c\tau$ = 100mm	4.36 ± 0.15 %	2.77 ± 0.12 %	5.49 ± 0.16 %	1.91 ± 0.10 %
m = 55GeV $c\tau$ = 100mm	4.35 ± 0.06 %	2.95 ± 0.05 %	6.96 ± 0.08 %	2.15 ± 0.05 %

Table D.1. Inclusive trigger efficiencies for the Emerging Jet and Asymmetric Displaced Jet triggers with and without preselection requirements for the full mass and lifetime signal sample grid. This covers mc campaigns mc23a and mc23d for the ggF production modes. Rates shown are unique, so passing one trigger means failing the other. EJ shows significant gains in the low lifetime samples over DJ, and comparable rates for longer lifetimes.

Mass/Lifetime	DJ Symmetric	DJ Symmetric Excl	DJ Asymmetric	DJ Asymmetric Excl
m = 3GeV $c\tau$ = 1mm	1.47 ± 0.05 %	0.37 ± 0.03 %	1.50 ± 0.05 %	0.40 ± 0.03 %
m = 5GeV $c\tau$ = 1mm	0.30 ± 0.02 %	0.04 ± 0.01 %	0.31 ± 0.02 %	0.05 ± 0.01 %
m = 10GeV $c\tau$ = 1mm	0.18 ± 0.02 %	0.03 ± 0.01 %	0.18 ± 0.02 %	0.03 ± 0.01 %
m = 15GeV $c\tau$ = 1mm	0.15 ± 0.02 %	0.01 ± 0.00 %	0.16 ± 0.02 %	0.01 ± 0.00 %
m = 16GeV $c\tau$ = 1mm	0.21 ± 0.02 %	0.02 ± 0.01 %	0.21 ± 0.02 %	0.02 ± 0.01 %
m = 20GeV $c\tau$ = 1mm	0.17 ± 0.02 %	0.01 ± 0.00 %	0.18 ± 0.02 %	0.02 ± 0.01 %
m = 30GeV $c\tau$ = 1mm	0.16 ± 0.02 %	0.02 ± 0.01 %	0.16 ± 0.02 %	0.02 ± 0.01 %
m = 40GeV $c\tau$ = 1mm	0.19 ± 0.02 %	0.02 ± 0.01 %	0.20 ± 0.02 %	0.03 ± 0.01 %
m = 55GeV $c\tau$ = 1mm	0.21 ± 0.01 %	0.02 ± 0.00 %	0.21 ± 0.01 %	0.02 ± 0.00 %
m = 3GeV $c\tau$ = 10mm	2.80 ± 0.07 %	0.10 ± 0.01 %	4.34 ± 0.09 %	1.64 ± 0.05 %
m = 5GeV $c\tau$ = 10mm	2.50 ± 0.06 %	0.11 ± 0.01 %	3.62 ± 0.08 %	1.23 ± 0.05 %
m = 10GeV $c\tau$ = 10mm	2.61 ± 0.07 %	0.16 ± 0.02 %	3.11 ± 0.07 %	0.65 ± 0.03 %
m = 15GeV $c\tau$ = 10mm	2.46 ± 0.06 %	0.18 ± 0.02 %	2.57 ± 0.06 %	0.28 ± 0.02 %
m = 16GeV $c\tau$ = 10mm	2.60 ± 0.05 %	0.17 ± 0.01 %	2.69 ± 0.05 %	0.26 ± 0.02 %
m = 20GeV $c\tau$ = 10mm	2.48 ± 0.06 %	0.15 ± 0.02 %	2.55 ± 0.06 %	0.22 ± 0.02 %
m = 30GeV $c\tau$ = 10mm	2.58 ± 0.07 %	0.14 ± 0.02 %	2.61 ± 0.07 %	0.18 ± 0.02 %
m = 40GeV $c\tau$ = 10mm	3.27 ± 0.07 %	0.08 ± 0.02 %	3.29 ± 0.07 %	1.89 ± 0.10 %
m = 55GeV $c\tau$ = 10mm	3.41 ± 0.06 %	0.16 ± 0.01 %	3.39 ± 0.06 %	0.13 ± 0.01 %
m = 3GeV $c\tau$ = 100mm	0.89 ± 0.04 %	0.02 ± 0.01 %	1.36 ± 0.05 %	0.49 ± 0.03 %
m = 5GeV $c\tau$ = 100mm	1.48 ± 0.05 %	0.04 ± 0.01 %	2.35 ± 0.07 %	0.92 ± 0.04 %
m = 10GeV $c\tau$ = 100mm	2.29 ± 0.07 %	0.05 ± 0.01 %	3.67 ± 0.08 %	1.43 ± 0.05 %
m = 15GeV $c\tau$ = 100mm	2.86 ± 0.08 %	0.03 ± 0.01 %	4.61 ± 0.10 %	1.78 ± 0.06 %
m = 16GeV $c\tau$ = 100mm	3.18 ± 0.06 %	0.05 ± 0.01 %	4.87 ± 0.07 %	1.74 ± 0.04 %
m = 20GeV $c\tau$ = 100mm	3.60 ± 0.08 %	0.08 ± 0.01 %	5.31 ± 0.09 %	1.78 ± 0.06 %
m = 30GeV $c\tau$ = 100mm	4.59 ± 0.10 %	0.07 ± 0.01 %	6.41 ± 0.11 %	1.90 ± 0.06 %
m = 40GeV $c\tau$ = 100mm	6.33 ± 0.17 %	0.08 ± 0.02 %	8.14 ± 0.20 %	1.89 ± 0.10 %
m = 55GeV $c\tau$ = 100mm	8.51 ± 0.09 %	0.12 ± 0.01 %	9.27 ± 0.09 %	0.88 ± 0.03 %

Table D.2. Inclusive trigger efficiencies for the symmetric and asymmetric DJ triggers for the full mass/lifetime sample grid. This covers mc campaigns mc23a and mc23d for the ggF production modes Exclusive denotes that the other trigger must fail so that we obtain a rate of unique events captured by each trigger. The asymmetric DJ trigger shows better acceptance overall, so it is taken as the main DJ trigger for other studies.

Mass/Lifetime	VBF Trigger	Presel 100/90	VBF Trig + Presel 100/90	Presel 110/90	VBF Trig + Presel 110/90
m = 3GeV $c\tau$ = 1mm	20.80 ± 0.17 %	4.73 ± 0.09 %	4.44 ± 0.08 %	4.40 ± 0.08 %	4.15 ± 0.08 %
m = 5GeV $c\tau$ = 1mm	19.25 ± 0.16 %	4.46 ± 0.08 %	4.14 ±0.08 %	4.19 ± 0.08 %	3.92 ± 0.08 %
m = 10GeV $c\tau$ = 1mm	15.48 ± 0.15 %	3.55 ± 0.08 %	3.27 ±0.07 %	3.31 ± 0.07 %	3.07 ±0.07 %
m = 15GeV $c\tau$ = 1mm	18.45 ±0.16 %	4.40 ±0.08 %	4.12 ±0.08 %	4.14 ±0.08 %	3.89 ±0.08 %
m = 16GeV $c\tau$ = 1mm	18.31 ±0.12 %	4.51 ±0.07 %	4.18 ± 0.06 %	4.23 ± 0.06 %	3.94 ±0.06 %
m = 20GeV $c\tau$ = 1mm	18.44 ±0.16 %	4.44 ±0.08 %	4.12 ±0.08 %	4.18 ±0.08 %	3.90 ±0.08 %
m = 30GeV $c\tau$ = 1mm	18.29 ±0.16 %	4.78 ±0.09 %	4.43 ±0.08 %	4.48 ±0.08 %	4.16 ±0.08 %
m = 40GeV $c\tau$ = 1mm	17.82 ±0.16 %	4.68 ±0.09 %	4.36 ±0.08 %	4.35 ±0.08 %	4.08 ±0.08 %
m = 55GeV $c\tau$ = 1mm	17.60 ±0.13 %	4.69 ± 0.07 %	4.33 ±0.07 %	4.41 ±0.07 %	4.09 ±0.07 %
m = 3GeV $c\tau$ = 10mm	20.40 ±0.17 %	4.57 ±0.09 %	4.24 ±0.08 %	4.26 ±0.08 %	3.97 ±0.08 %
m = 5GeV $c\tau$ = 10mm	19.44 ±0.16 %	4.37 ±0.08 %	4.09 ±0.08 %	4.08 ±0.08 %	3.84 ±0.08 %
m = 10GeV $c\tau$ = 10mm	15.69 0.15 %	3.56 ±0.08 %	3.30 ±0.07 %	3.33 ±0.07 %	3.11 ±0.07 %
m = 15GeV $c\tau$ = 10mm	18.55 ±0.16 %	4.43 ±0.08 %	4.13 ±0.08 %	4.18 ±0.08 %	3.91 ±0.08 %
m = 16GeV $c\tau$ = 10mm	18.39 ±0.12 %	4.46 ±0.07 %	4.14 ±0.06 %	4.20 ±0.06 %	3.91 ±0.06 %
m = 20GeV $c\tau$ = 10mm	18.20 ±0.22 %	4.75 ±0.12 %	4.36 ±0.12 %	4.46 ±0.12 %	4.11 ±0.12 %
m = 30GeV $c\tau$ = 10mm	18.51 ±0.16 %	4.72 ±0.09 %	4.41 ±0.08 %	4.39 ±0.08 %	4.12 ±0.08 %
m m = 40GeV $c\tau$ = 10mm	17.87 ±0.16 %	4.73 ±0.09 %	4.39 ±0.08 %	4.41 ±0.08 %	4.10 ±0.08 %
m = 55GeV $c\tau$ = 10mm	17.41 ±0.14 %	4.70 ±0.08 %	4.31 ±0.08 %	4.44 ±0.08 %	4.10 ±0.08 %
m = 3GeV $c\tau$ = 100mm	16.74 ±0.23 %	4.30 ±0.13 %	3.91 ±0.12 %	4.00 ±0.12 %	3.68 ±0.12 %
m = 5GeV $c\tau$ = 100mm	16.87 ±0.16 %	4.09 ±0.09 %	3.75 ±0.08 %	3.80 ±0.08 %	3.51 ±0.08 %
m = 10GeV $c\tau$ = 100mm	14.46 ±0.15 %	3.18 ±0.08 %	2.94 ±0.07 %	2.97 ±0.07 %	2.77 ±0.07 %
m = 15GeV $c\tau$ = 100mm	18.63 ±0.23 %	4.29 ±0.12 %	4.07 ±0.12 %	3.97 ±0.12 %	3.79 ±0.12 %
m = 16GeV $c\tau$ = 100mm	17.96 ±0.13 %	4.19 ±0.07 %	3.88 ±0.06 %	3.95 ±0.06 %	3.67 ±0.06 %
m = 20GeV $c\tau$ = 100mm	18.58 ±0.20 %	4.45 ±0.11 %	4.15 ±0.10 %	4.17 ±0.10 %	3.91 ±0.10 %
m = 30GeV $c\tau$ = 100mm	18.56 ±0.16 %	4.56 ±0.09 %	4.27 ±0.08 %	4.26 ±0.08 %	4.00 ±0.08 %
m = 40GeV $c\tau$ = 100mm	18.24 ± 0.16 %	4.68 ±0.09 %	4.32 ±0.08 %	4.38 ±0.08 %	4.05 ±0.08 %
m = 55GeV $c\tau$ = 100mm	17.78 ±0.15 %	4.72 ±0.08 %	4.37 ±0.08 %	4.42 ±0.08 %	4.11 ±0.08 %

Table D.3. Inclusive trigger rates for the VBF trigger and preselection options for the full mass/lifetime grid. Calculated using mc23a and mc23d for the VBF production mode. This study played a a role in determining the final preselection cuts to be applied: 100 GeV for leading forward jet and 90 GeV for subleading.

APPENDIX E
BACKGROUND CONTAMINATION AND EVENT YIELDS OF THE TRIGGER BUCKET STRATEGY

Table E.1. Background Monte Carlo samples used in signal studies.

Process	**DSIDs**
$t\bar{t}$	601229, 601230, 601237
QCD dijet	801167, 801168, 801169, 801170, 801171, 801172, 801173, 801174
$\gamma +$ jets	801649, 801650, 801651, 801652, 801653, 801654, 801655, 801656, 801657, 801658, 801659, 801660, 801663, 801664, 801665, 801666, 801667, 801668, 801669, 801670, 801671, 801672, 801673, 801674, 801675, 801676

The following plots show the signal contamination and event yields for the DJ and EJ trigger buckets as described in 5.7.2. These plots are supplementary to the significance plots shown in the text, and included here for completeness

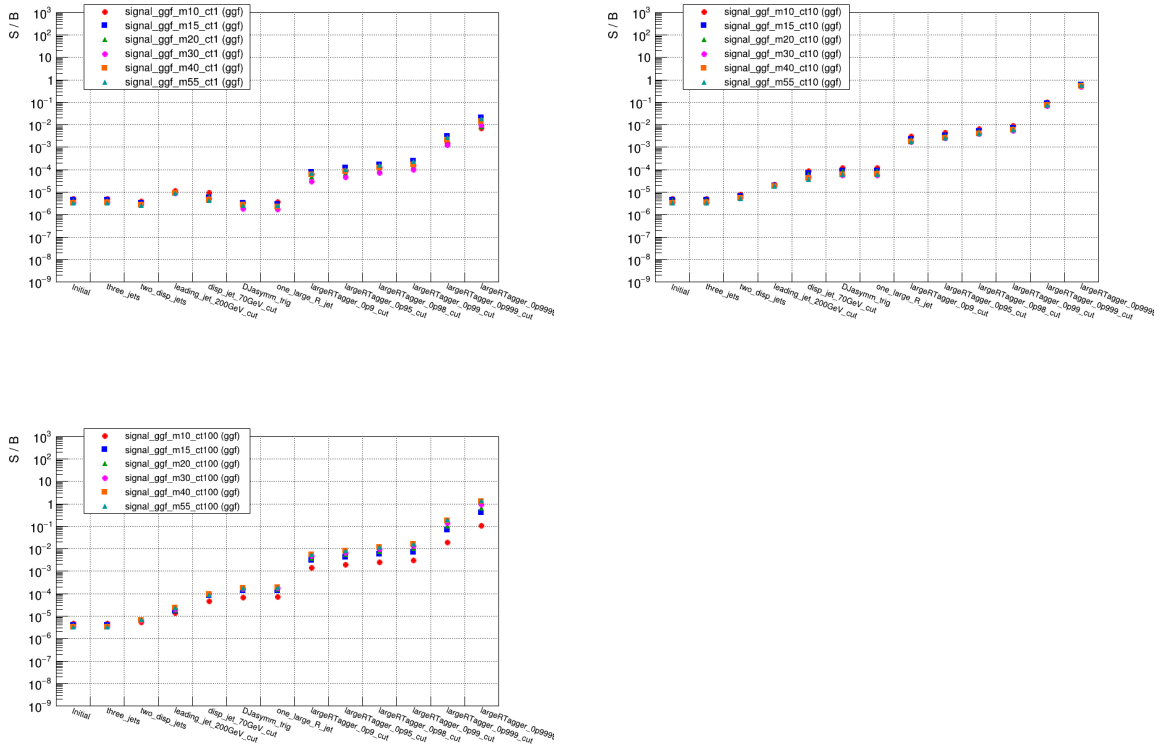


Figure E.1. DJ Bucket Boosted Region Signal Contamination, ordered in increasing signal sample lifetime.

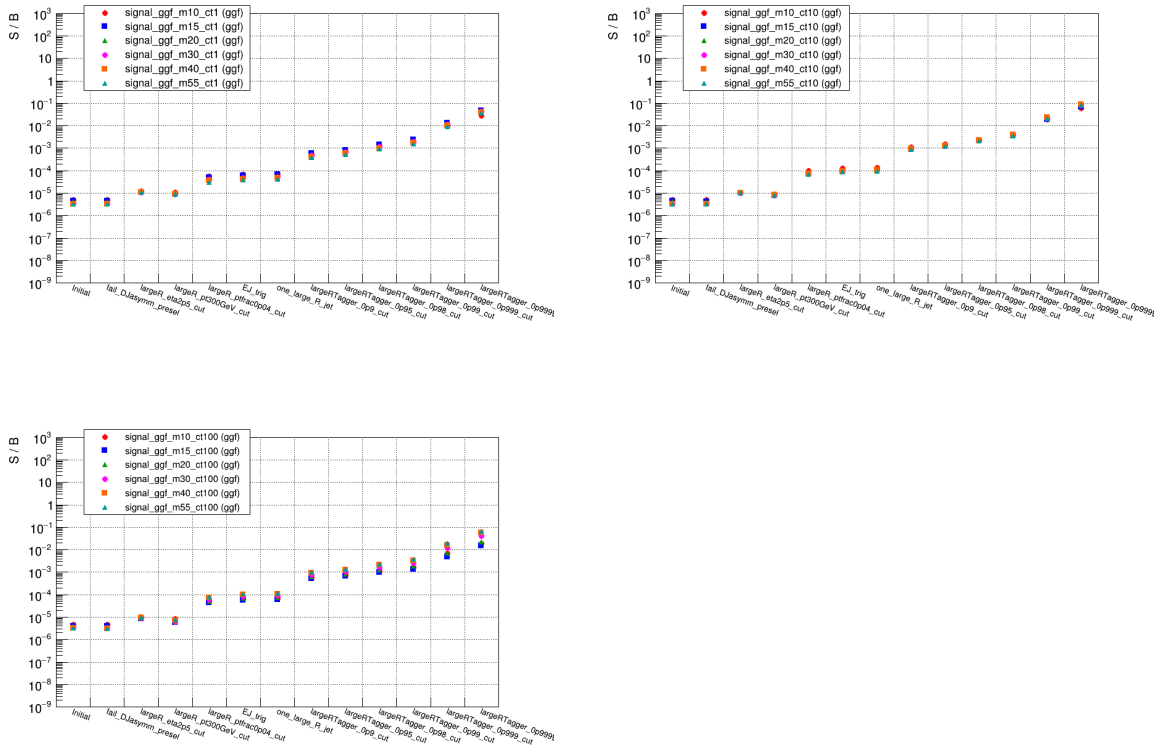


Figure E.2. EJ Bucket Boosted Region Signal Contamination, ordered in increasing signal sample lifetime.

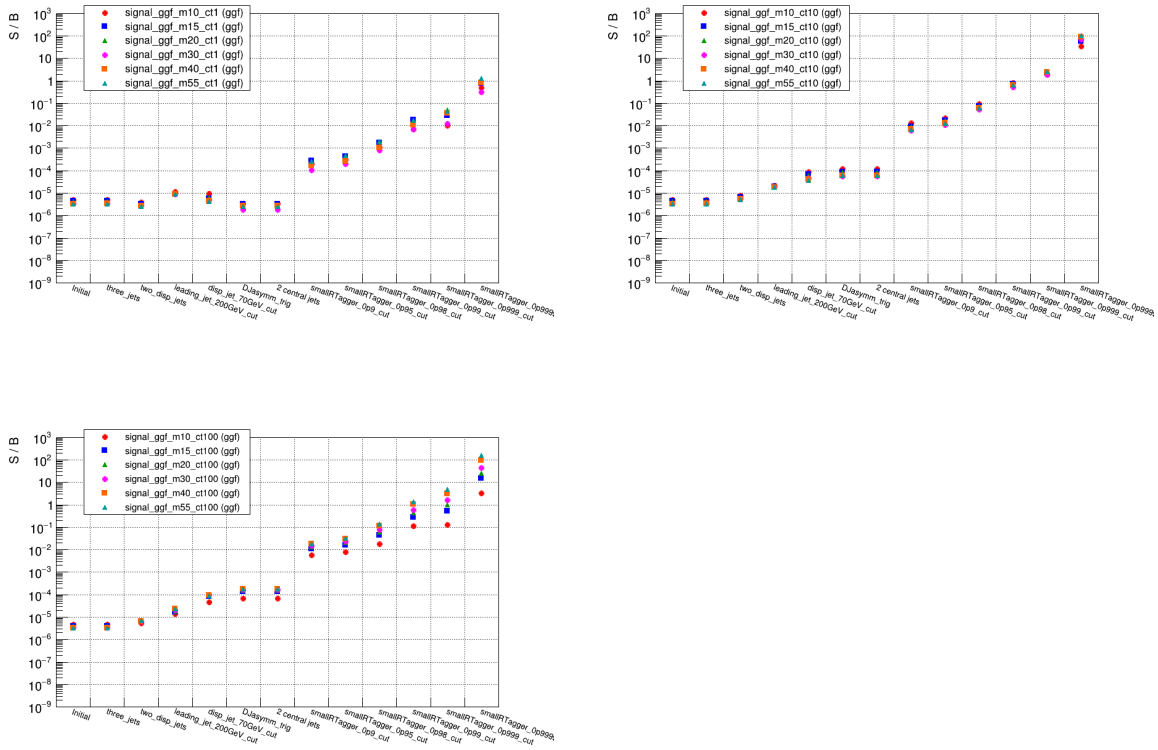


Figure E.3. DJ Bucket Resolved Region Signal Contamination, ordered in increasing signal sample lifetime.

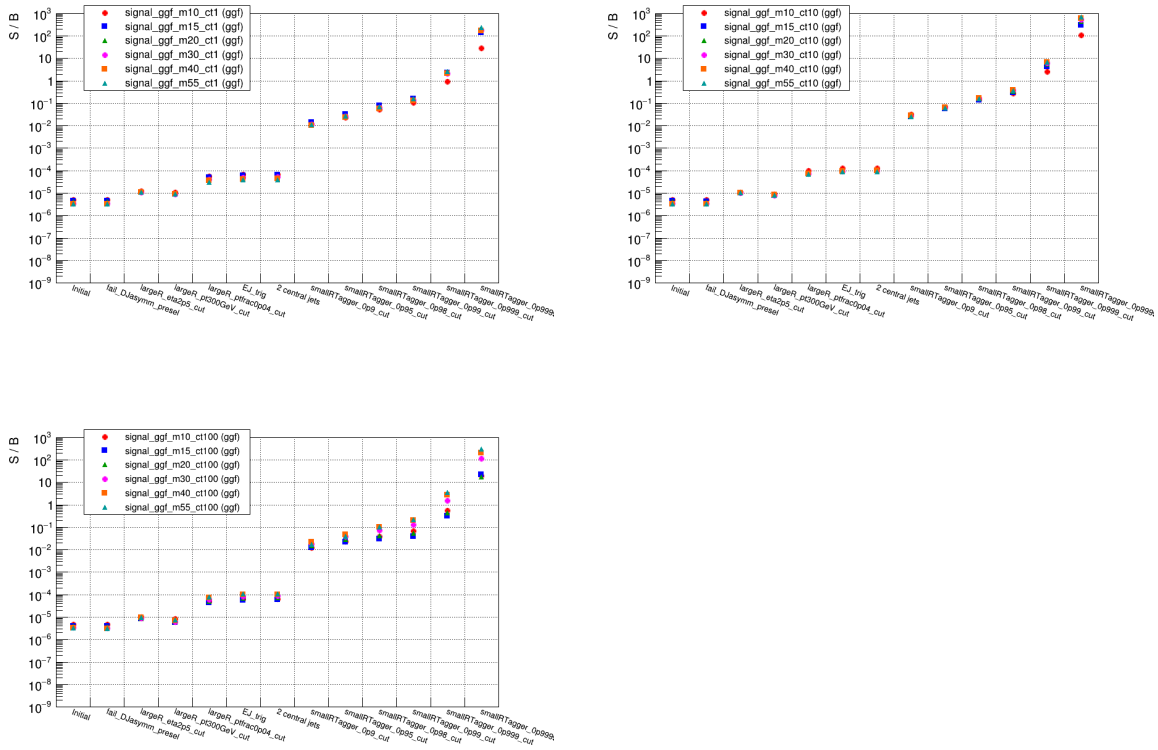


Figure E.4. EJ Bucket Resolved Region Signal Contamination, ordered in increasing signal sample lifetime.

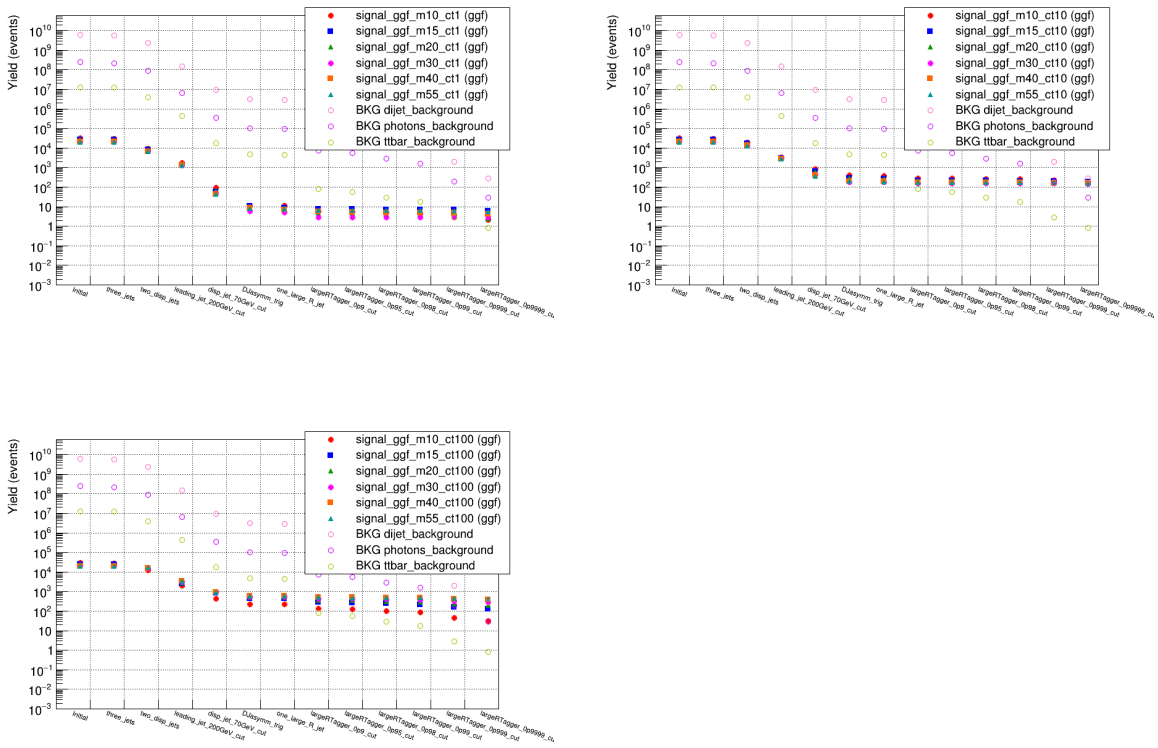


Figure E.5. DJ Bucket Boosted Region Event Yields, ordered in increasing signal sample lifetime.

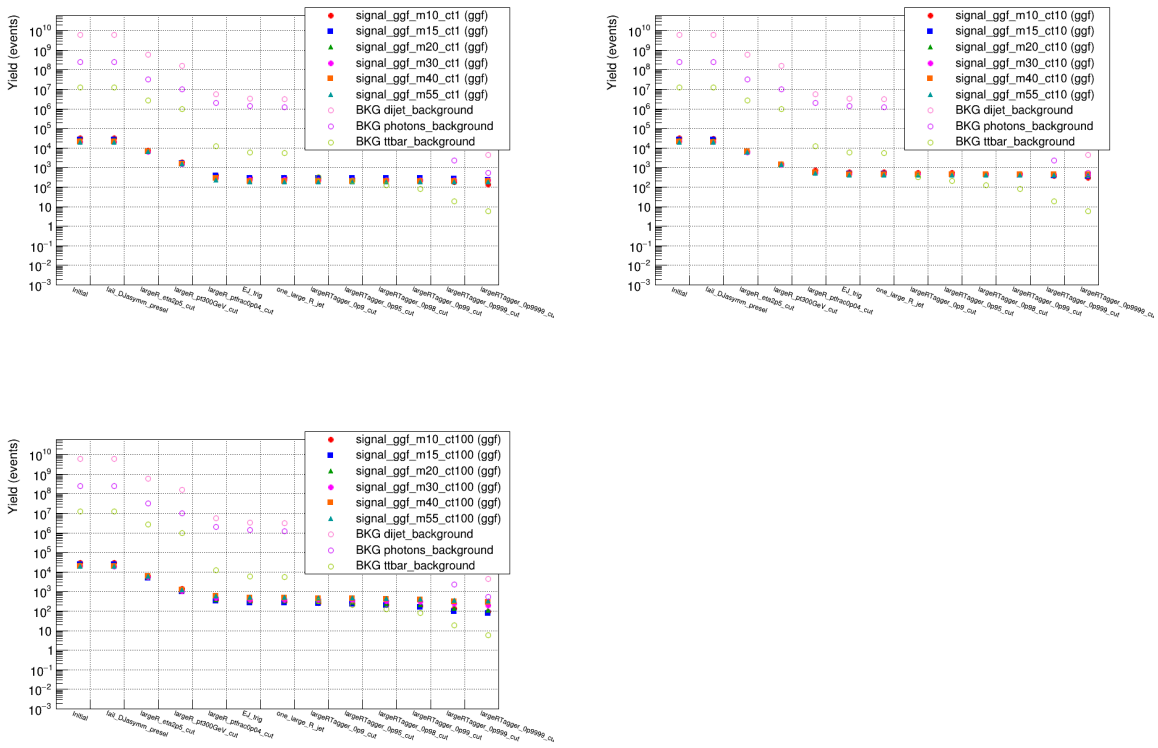


Figure E.6. EJ Bucket Boosted Region Event Yields, ordered in increasing signal sample lifetime.

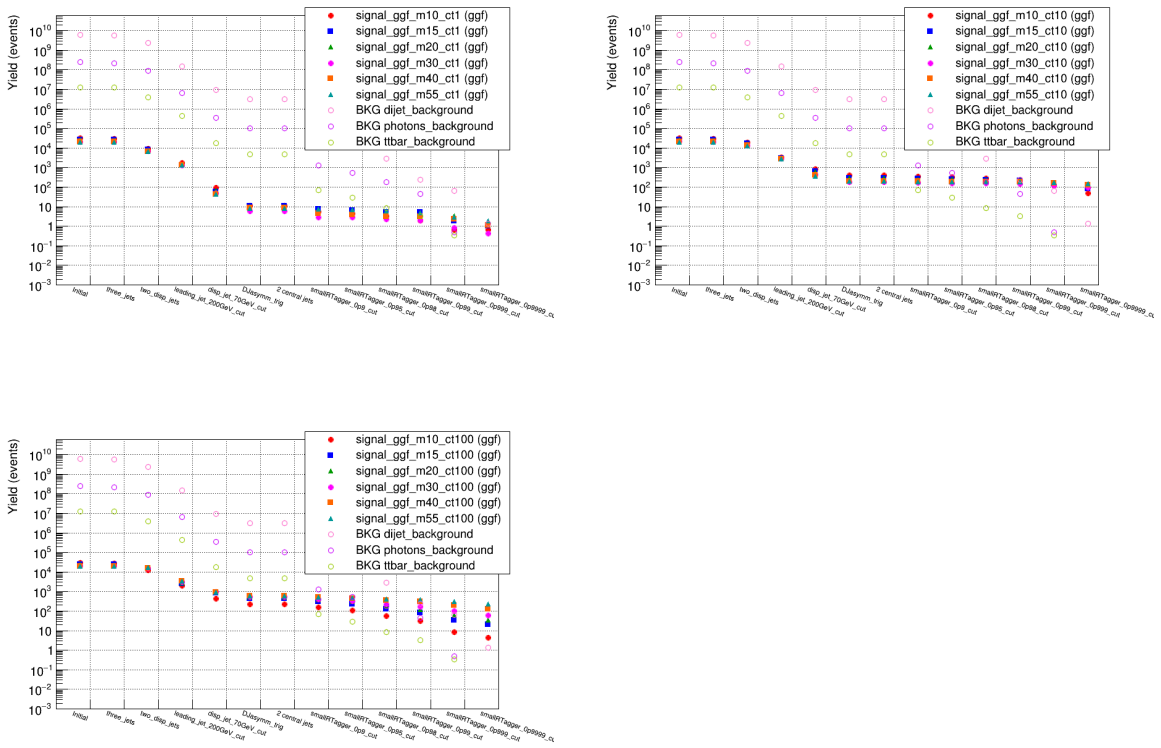


Figure E.7. DJ Bucket Resolved Region Event Yield, ordered in increasing signal sample lifetime.

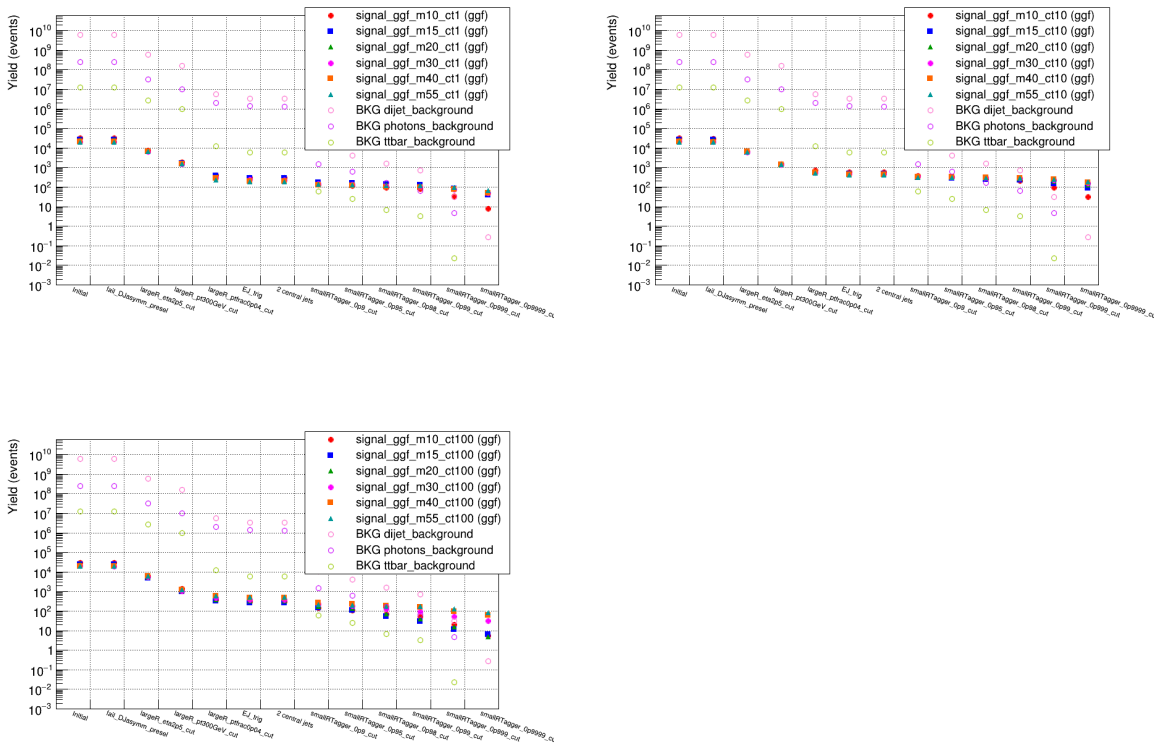


Figure E.8. EJ Bucket Resolved Region Event Yield, ordered in increasing signal sample lifetime.

APPENDIX F
SAMPLES FOR DIPHOTON TRIGGER STUDIES

This Appendix contains the samples used in trigger studies for the Run 3 asymmetric diphoton trigger. The samples are from light Higgs resonances decaying to two photons, and are used as a proxy for an ALP decaying to two photons to assess the feasibility of using a diphoton trigger for such processes.

DSID	Process
603519	mc.PhPy8EG_PDF4LHC21_ggH60_MiNLO_gamgam.deriv. DAOD_PHYS.e8557_e8586_s4369_s4370_r16083_r15970_p6522
603520	mc.PhPy8EG_PDF4LHC21_ggH70_MiNLO_gamgam.deriv. DAOD_PHYS.e8557_e8586_s4369_s4370_r16083_r15970_p6522
603521	mc.PhPy8EG_PDF4LHC21_ggH80_MiNLO_gamgam.deriv. DAOD_PHYS.e8557_e8586_s4369_s4370_r16083_r15970_p6522
603522	mc.PhPy8EG_PDF4LHC21_ggH90_MiNLO_gamgam.deriv. DAOD_PHYS.e8557_e8586_s4369_s4370_r16083_r15970_p6522
603516	mc.PhPy8EG_PDF4LHC21_ggH100_MiNLO_gamgam.deriv. DAOD_PHYS.e8557_e8586_s4369_s4370_r16083_r15970_p6522
603517	mc.PhPy8EG_PDF4LHC21_ggH110_MiNLO_gamgam.deriv. DAOD_PHYS.e8557_e8586_s4369_s4370_r16083_r15970_p6522
603518	mc.PhPy8EG_PDF4LHC21_ggH120_MiNLO_gamgam.deriv. DAOD_PHYS.e8557_e8586_s4369_s4370_r16083_r15970_p6522

Table F.1. Light Higgs Resonance decays to diphoton samples.

APPENDIX G
ASSESSING E_{RATIO} CUTS USING THE SIGNAL BACKGROUND CLASSIFIER MODEL

This Appendix contains additional output plots of the signal background classifier FFNN that show the mean NN score as a function of E_{ratio}. Plots are split into E_T bins which correspond to the dividing bins of the variable E_{ratio} cuts found in Table 17. The purpose of this check is to see if these cuts are too aggressive for our diphoton algorithms and if we could benefit from looser selections on E_{ratio}. For the purposes of this study, the E_{ratio} cuts were not included in the training,

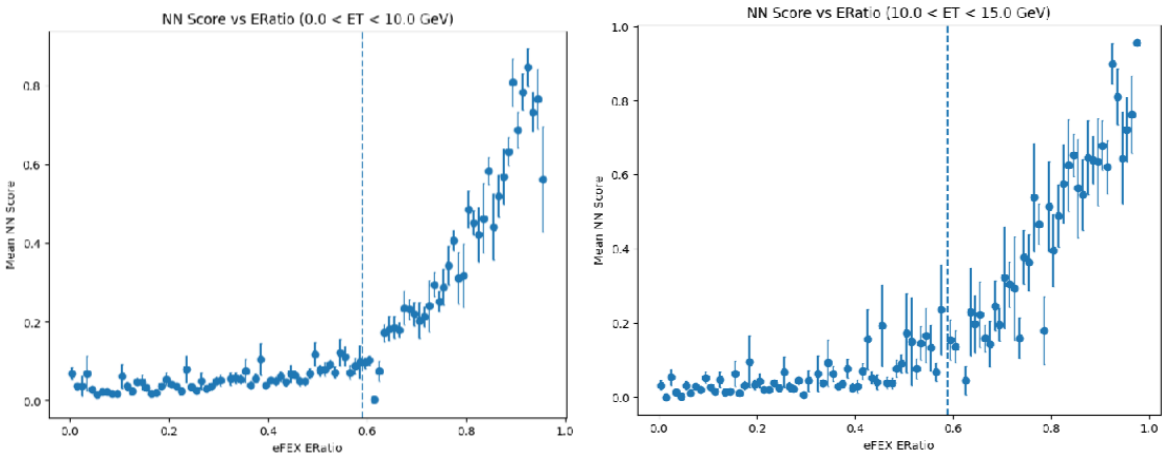


Figure G.1. NN output of the S/B classifier parameterized by E_{ratio} for E_T bins 1 and 2 corresponding to [0,10) GeV and [10,15) GeV.

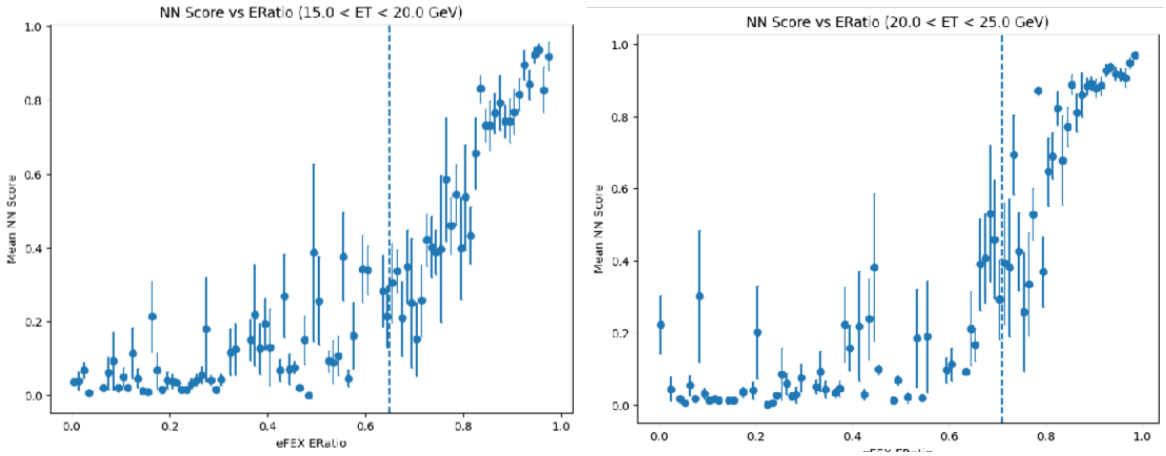


Figure G.2. NN output of the S/B classifier parameterized by E_{ratio} for E_T bins 3 and 4 corresponding to [15,20) GeV and [20,25) GeV.

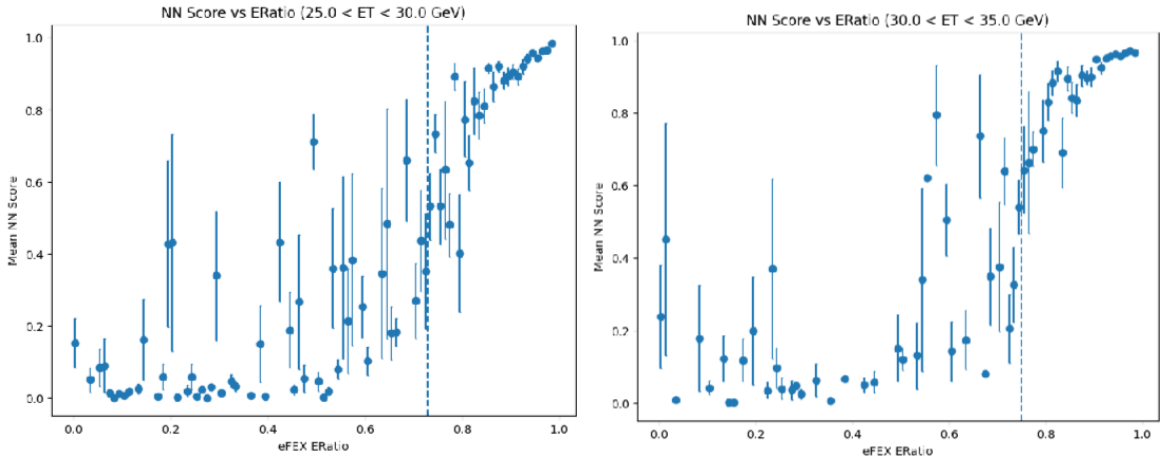


Figure G.3. NN output of the S/B classifier parameterized by E_{ratio} for E_T bins 5 and 6 corresponding to [25,30) GeV and [30,35) GeV.

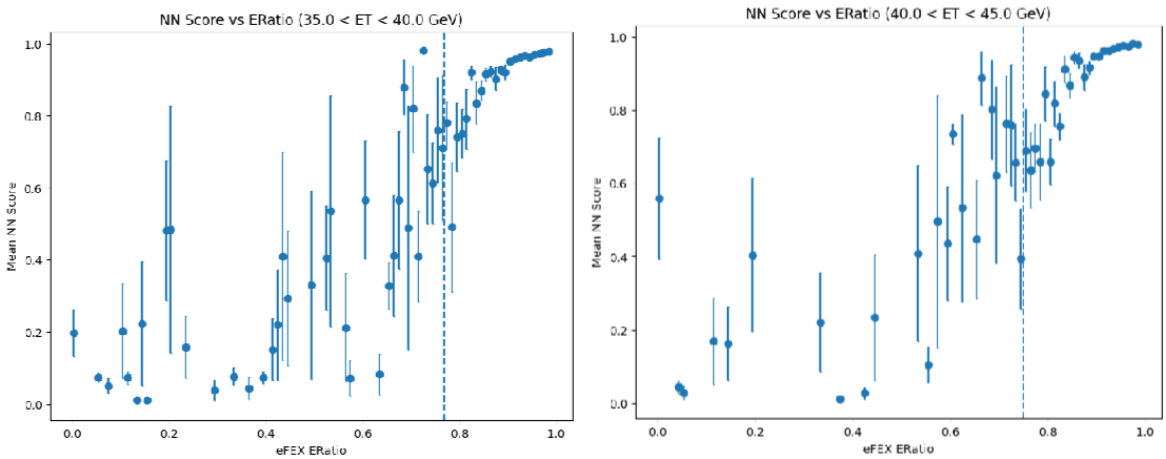


Figure G.4. NN output of the S/B classifier parameterized by E_{ratio} for E_T bins 7 and 8 corresponding to [35,40) GeV and [40,45) GeV.

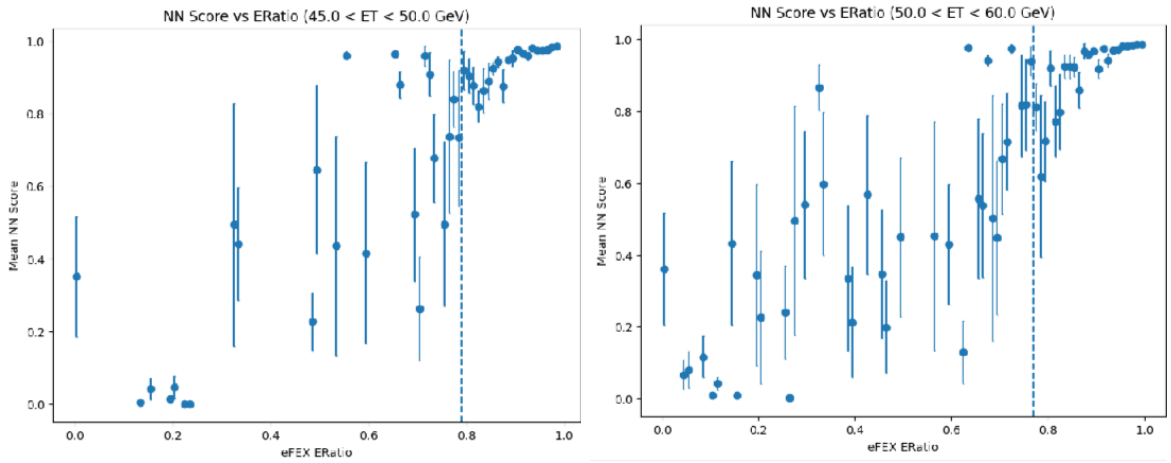


Figure G.5. NN output of the S/B classifier parameterized by E_{ratio} for E_T bins 9 and 10 corresponding to [45,50) GeV and [50,60) GeV.

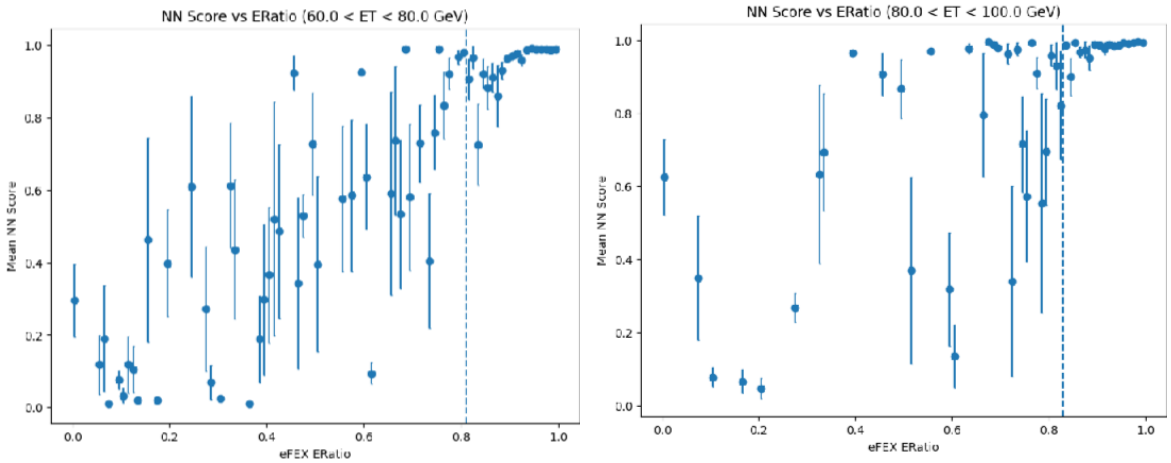


Figure G.6. NN output of the S/B classifier parameterized by E_{ratio} for E_T bins 11 and 12 corresponding to [60,80) GeV and [80,100) GeV.

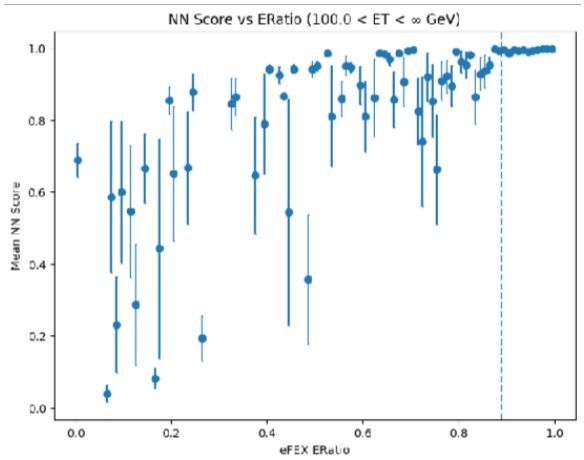


Figure G.7. NN output of the S/B classifier parameterized by E_{ratio} for E_T bin 13 corresponding to [100, ∞) GeV.

Straight lines are drawn on each plot corresponding to that E_T bin's E_{ratio} cut to asses the quality of NN score on either side of the cut. For essentially every E_T bin, the cuts do not prohibitively cut events that perform well. As E_T increases, higher score events begin to appear on the left side of the cut line with increasing frequency. This is not necessarily a problem, however, since we know that this model is highly dependent on E_T, so there is an increasing likelihood of higher scores in those bins anyway.

REFERENCES CITED

[1] D. Dominguez, *Standard Model*, . https://cds.cern.ch/record/2002395. General Photo.

[2] J. Ellis, M. K. Gaillard, and D. V. Nanopoulos, *A Historical Profile of the Higgs Boson*, arXiv:1504.07217. https://cds.cern.ch/record/2012465. 22 pages, 5 figures, update of arXiv:1201.6045, to be published in the volume.

[3] L. Lee, C. Ohm, A. Soffer, and T.-T. Yu, *Collider searches for long-lived particles beyond the Standard Model*, Progress in Particle and Nuclear Physics **106** (may, 2019) 210–255. https://doi.org/10.1016%2Fj.ppnp.2019.02.006.

[4] E. Mobs, *The CERN accelerator complex*, . https://cds.cern.ch/record/2197559. General Photo.

[5] ATLAS Collaboration, *The ATLAS experiment at the CERN Large Hadron Collider: a description of the detector configuration for Run 3*, Journal of Instrumentation **19** (May, 2024) P05063. http://dx.doi.org/10.1088/1748-0221/19/05/P05063.

[6] M. C. Aleksa, W. P. Cleland, Y. T. Enari, M. V. Fincke-Keeler, L. C. Hervas, F. B. Lanni, S. O. Majewski, C. V. Marino, and I. L. Wingerter-Seez, *ATLAS Liquid Argon Calorimeter Phase-I Upgrade*, tech. rep., 2013. https://cds.cern.ch/record/1602230. Final version presented to December 2013 LHCC.

[7] ATLAS Collaboration, *Technical Design Report for the Phase-II Upgrade of the ATLAS Tile Calorimeter*, tech. rep., CERN, Geneva, 2017. https://cds.cern.ch/record/2285583.

[8] ATLAS Collaboration, *The ATLAS Experiment at the CERN Large Hadron Collider*, Journal of Instrumentation **3** (2008) no. 08, S08003. https://dx.doi.org/10.1088/1748-0221/3/08/S08003.

[9] *ATLAS L1CaloTrigger Public Results web page*, https://twiki.cern.ch/twiki/bin/view/AtlasPublic/L1CaloTriggerPublicResults.

[10] J. Mendez, V. Bobillier, S. Haas, M. Joos, and F. Vasey, *Evaluation of a commercial AdvancedTCA board management controller solution (IPMC)*, Journal of Instrumentation **11** (02, 2016) C02044–C02044.

[11] M. Begel, S. Majewski, B. W. Allen, B. T. Carlson, R. Camacho Toro, H. Chen, K. Chen, C. Doglioni, H. Evans, N. M. Harrison, T. M. Hong, W. Hopkins, C. Johnson, C. W. Kalderon, R. Kopeliansky, V. Kouskoura, S. E. Kroeker, B. M. Kunkler, S. Lammers, F. Lanni, R. A. Linck, D. Miller, J. Roloff, S. Silverstein, N. Sinev, I. M. Snyder, G. Stark, D. Strom, H. Takai, F. Tang, S. Tang, R. Van Kooten, and W. Wu, *Global Feature Extractor of the Level-1 Calorimeter Trigger*, tech. rep., CERN, Geneva, 2016. https://cds.cern.ch/record/2233958.

[12] D. Curtin, R. Essig, S. Gori, P. Jaiswal, A. Katz, T. Liu, Z. Liu, D. McKeen, J. Shelton, M. Strassler, and et al., *Exotic decays of the 125 GeV Higgs boson*, Physical Review D **90** (10, 2014) . http://dx.doi.org/10.1103/PhysRevD.90.075004.

[13] ATLAS Collaboration, *Summary plots for Higgs boson mediated hidden sector involving long-lived particles*, tech. rep., CERN, Geneva, 2025. https://cds.cern.ch/record/2946396. All figures including auxiliary figures are available at https://atlas.web.cern.ch/Atlas/GROUPS/PHYSICS/PUBNOTES/ATL-PHYS-PUB-2025-039.

[14] ATLAS Collaboration, *Expected tracking performance of the ATLAS Inner Tracker at the High-Luminosity LHC*, Journal of Instrumentation **20** (feb, 2025) P02018. https://doi.org/10.1088/1748-0221/20/02/P02018.

[15] ATLAS Collaboration, *The ATLAS Upgrade for the HL-LHC*, .

[16] ATLAS Collaboration, *Technical Design Report for the Phase-II Upgrade of the ATLAS TDAQ System*. Geneva, 2017. https://cds.cern.ch/record/2285584.

[17] M. Begel and D. Sankey, *System Specification for the Global Trigger*, tech. rep., CERN, Geneva, 2021. https://cds.cern.ch/record/2791119.

[18] ATLAS Collaboration, Aad, Georges et al., *Search for emerging jets in pp collisions at $\sqrt{s} = 13.6$ TeV with the ATLAS experiment*, Rept. Prog. Phys. **88** (2025) no. 9, 097801, arXiv:2505.02429 [hep-ex].

[19] Particle Data Group Collaboration, F. Takahashi et al., *Review of Particle Physics*, Int. J. Mod. Phys. A **41** (2026) 2630011.

[20] A. Borriello and P. Salucci, *The dark matter distribution in disc galaxies*, Monthly Notices of the Royal Astronomical Society **323** (may, 2001) 285–292. https://doi.org/10.1046%2Fj.1365-8711.2001.04077.x.

[21] W. Hu and S. Dodelson, *Cosmic Microwave Background Anisotropies*, Annual Review of Astronomy and Astrophysics **40** (2002) no. Volume 40, 2002, 171–216. https://www.annualreviews.org/content/journals/10.1146/annurev.astro.40.060401.093926.

[22] ATLAS Collaboration, *Observation of a new particle in the search for the Standard Model Higgs boson with the ATLAS detector at the LHC*, Physics Letters B **716** (2012) no. 1, 1–29. http://dx.doi.org/10.1016/j.physletb.2012.08.020.

[23] CMS Collaboration, *Observation of a new boson at a mass of 125 GeV with the CMS experiment at the LHC*, Physics Letters B **716** (2012) no. 1, 30–61. http://dx.doi.org/10.1016/j.physletb.2012.08.021.

[24] G. 't Hooft, *Naturalness, chiral symmetry, and spontaneous chiral symmetry breaking*, NATO Sci. Ser. B **59** (1980) 135–157.

[25] B. Batell, M. Low, E. T. Neil, and C. B. Verhaaren, *Review of neutral naturalness*, Physics Reports **1165** (Apr., 2026) 1–49. http://dx.doi.org/10.1016/j.physrep.2025.12.001.

[26] ATLAS Collaboration, G. Aad et al., *A detailed map of Higgs boson interactions by the ATLAS experiment ten years after the discovery*, Nature **607** (2022) no. 7917, 52–59, arXiv:2207.00092 [hep-ex]. [Erratum: Nature 612, E24 (2022)].

[27] G. Branco, P. Ferreira, L. Lavoura, M. Rebelo, M. Sher, and J. P. Silva, *Theory and phenomenology of two-Higgs-doublet models*, Physics Reports **516** (jul, 2012) 1–102. https://doi.org/10.1016%2Fj.physrep.2012.02.002.

[28] S. Weinberg, *A New Light Boson?*, Phys. Rev. Lett. **40** (Jan, 1978) 223–226. https://link.aps.org/doi/10.1103/PhysRevLett.40.223.

[29] M. Bauer, M. Heiles, M. Neubert, and A. Thamm, *Axion-like particles at future colliders*, The European Physical Journal C **79** (Jan., 2019) . http://dx.doi.org/10.1140/epjc/s10052-019-6587-9.

[30] L. Evans and P. Bryant, *LHC Machine*, Journal of Instrumentation **3** (aug, 2008) S08001. https://dx.doi.org/10.1088/1748-0221/3/08/S08001.

[31] CMS Collaboration, Chatrchyan, S et al., *The CMS experiment at the CERN LHC*, JINST **3** (2008) S08004. https://cds.cern.ch/record/1129810. Also published by CERN Geneva in 2010.

[32] ALICE Collaboration, *The ALICE experiment at the CERN LHC*, Journal of Instrumentation **3** (aug, 2008) S08002. https://doi.org/10.1088/1748-0221/3/08/S08002.

[33] LHCb Collaboration, Alves, A Augusto et al., *The LHCb Detector at the LHC*, JINST **3** (2008) S08005. https://cds.cern.ch/record/1129809. Also published by CERN Geneva in 2010.

[34] Arnaudon, L et al., *Linac4 Technical Design Report*, tech. rep., CERN, Geneva, 2006. https://cds.cern.ch/record/1004186. revised version submitted on 2006-12-14 09:00:40.

[35] J. Gareyte, *LHC main parameters*, Part. Accel. **50** (1995) 61–68. https://cds.cern.ch/record/304825.

[36] M. Ferro-Luzzi, W. Herr, and T. Pieloni, *LHC bunch filling schemes for commissioning and initial luminosity optimization*, tech. rep., CERN, Geneva, 2008. https://cds.cern.ch/record/1114612.

[37] *ATLAS Luminosity Public Results Run 3 web page*, https://twiki.cern.ch/twiki/bin/view/AtlasPublic/LuminosityPublicResultsRun3.

[38] G. A. et al., *ATLAS pixel detector electronics and sensors*, Journal of Instrumentation **3** (jul, 2008) P07007. https://doi.org/10.1088/1748-0221/3/07/P07007.

[39] ATLAS Collaboration, *Operation and performance of the ATLAS semiconductor tracker*, Journal of Instrumentation **9** (aug, 2014) P08009. https://doi.org/10.1088/1748-0221/9/08/P08009.

[40] ATLAS TRT Collaboration, *The ATLAS Transition Radiation Tracker (TRT) proportional drift tube: design and performance*, Journal of Instrumentation **3** (feb, 2008) P02013. https://doi.org/10.1088/1748-0221/3/02/P02013.

[41] ATLAS Collaboration, *ATLAS tile calorimeter*. Technical design report. ATLAS. CERN, Geneva, 1996. https://cds.cern.ch/record/331062.

[42] T. Kawamoto, S. Vlachos, L. Pontecorvo, J. Dubbert, G. Mikenberg, P. Iengo, C. Dallapiccola, C. Amelung, L. Levinson, R. Richter, and D. Lellouch, *New Small Wheel Technical Design Report*, tech. rep., 2013. https://cds.cern.ch/record/1552862. ATLAS New Small Wheel Technical Design Report.

[43] ATLAS Collaboration, *The ATLAS Collaboration Software and Firmware*, Tech. Rep. ATL-SOFT-PUB-2021-001, 2021.

[44] ATLAS Collaboration, Aad, Georges et al., *The ATLAS trigger system for LHC Run 3 and trigger performance in 2022*, JINST **19** (2024) no. 06, P06029, arXiv:2401.06630 [hep-ex].

[45] M. Cacciari, G. P. Salam, and G. Soyez, *The anti-k_t jet clustering algorithm*, JHEP **04** (2008) 063, arXiv:0802.1189 [hep-ph].

[46] *ATLAS L0Calorimeter Trigger Public Results web page*, https://twiki.cern.ch/twiki/bin/view/AtlasPublic/L0CalorimeterTriggerPublicReults.

[47] *Yocto Project*, 2011. https://www.yoctoproject.org/.

[48] *USCONEC*, https://www.usconec.com.

[49] D. Bertolini, T. Chan, and J. Thaler, *Jet observables without jet algorithms*, Journal of High Energy Physics **2014** (Apr., 2014) . http://dx.doi.org/10.1007/JHEP04(2014)013.

[50] A. J. Baltz et al., *The Physics of Ultraperipheral Collisions at the LHC*, Phys. Rept. **458** (2008) 1–171, arXiv:0706.3356 [nucl-ex].

[51] *gFEX Raw Data Dumping Script*, https://gitlab.cern.ch/atlas-l1calo/gfex/gfex_ops/-/blob/master/online_analysis/dump_rawFile.py?ref_type=heads.

[52] *gFEX Bitwise C Simulation*, https://gitlab.cern.ch/atlas-l1calo/gfex/fw_gfex_bitw_model.

[53] R. Brun and F. Rademakers, *ROOT: An object oriented data analysis framework*, Nucl. Instrum. Meth. A **389** (1997) 81–86.

[54] *gFEX Comparison framework*, https://gitlab.cern.ch/atlas-l1calo/gfex/gfex-run-csim-comparisons.

[55] T. Cohen, N. Craig, H. K. Lou, and D. Pinner, *Folded supersymmetry with a twist*, Journal of High Energy Physics **2016** (03, 2016) . http://dx.doi.org/10.1007/JHEP03(2016)196.

[56] H. Cai, H.-C. Cheng, and J. Terning, *A quirky little Higgs model*, Journal of High Energy Physics **2009** (05, 2009) 045–045. http://dx.doi.org/10.1088/1126-6708/2009/05/045.

[57] U. Haisch and L. Schnell, *Long-lived particle phenomenology in the 2HDM+a model*, arXiv:2302.02735 [hep-ph].

[58] Z. Chacko, H.-S. Goh, and R. Harnik, *Natural Electroweak Breaking from a Mirror Symmetry*, Physical Review Letters **96** (2006) no. 23, . http://dx.doi.org/10.1103/PhysRevLett.96.231802.

[59] S. P. Martin, *A Supersymmetry Primer*, Adv. Ser. Direct. High Energy Phys. **18** (1998) 1, arXiv:hep-ph/9709356.

[60] P. Fayet, *Supersymmetry and weak, electromagnetic and strong interactions*, Phys. Lett. B **64** (1976) 159.

[61] D. Curtin and C. B. Verhaaren, *Discovering uncolored naturalness in exotic Higgs decays*, Journal of High Energy Physics **2015** (12, 2015) 1–36. http://dx.doi.org/10.1007/JHEP12(2015)072.

[62] G. Burdman, Z. Chacko, H.-S. Goh, and R. Harnik, *Folded supersymmetry and the LEP paradox*, Journal of High Energy Physics **2007** (02, 2007) 009–009. http://dx.doi.org/10.1088/1126-6708/2007/02/009.

[63] ATLAS Collaboration, Aaboud, Morad et al., *Search for long-lived particles produced in pp collisions at $\sqrt{s} = 13$ TeV that decay into displaced hadronic jets in the ATLAS muon spectrometer*, Phys. Rev. D **99** (2019) no. 5, 052005, arXiv:1811.07370 [hep-ex].

[64] ATLAS Collaboration, Aaboud, Morad et al., *Search for long-lived neutral particles in pp collisions at $\sqrt{s}$ = 13 TeV that decay into displaced hadronic jets in the ATLAS calorimeter*, Eur. Phys. J. C **79** (2019) no. 6, 481, arXiv:1902.03094 [hep-ex].

[65] ATLAS Collaboration, Aad, Georges et al., *Search for long-lived neutral particles produced in pp collisions at $\sqrt{s} = 13$ TeV decaying into displaced hadronic jets in the ATLAS inner detector and muon spectrometer*, Phys. Rev. D **101** (2020) no. 5, 052013, arXiv:1911.12575 [hep-ex].

[66] Aad, Georges et al., *Search for neutral long-lived particles in pp collisions at $\sqrt{s}$ = 13 TeV that decay into displaced hadronic jets in the ATLAS calorimeter*, Journal of High Energy Physics **2022** (2022) no. 6, . http://dx.doi.org/10.1007/JHEP06(2022)005.

[67] CMS Collaboration, Sirunyan, Albert M, et al., *Search for long-lived particles decaying into displaced jets in proton-proton collisions at $\sqrt{s} = 13$ TeV*, Phys. Rev. D **99** (2019) no. 3, 032011, arXiv:1811.07991 [hep-ex].

[68] CMS Collaboration, Sirunyan, Albert M et al., *Search for long-lived particles using displaced jets in proton-proton collisions at $\sqrt{s} = 13$ TeV*, Phys. Rev. D **104** (2021) no. 1, 012015, arXiv:2012.01581 [hep-ex].

[69] ATLAS Collaboration, Aad, Georges et al., *Search for Light Long-Lived Particles in pp Collisions at s=13 TeV Using Displaced Vertices in the ATLAS Inner Detector*, Phys. Rev. Lett. **133** (2024) no. 16, 161803, arXiv:2403.15332 [hep-ex]. [Erratum: Phys.Rev.Lett. 135, 259901 (2025)].

[70] CMS Collaboration, Tumasyan, Armen et al., *Search for long-lived particles produced in association with a Z boson in proton-proton collisions at $\sqrt{s}$ = 13 TeV*, JHEP **03** (2022) 160, arXiv:2110.13218 [hep-ex].

[71] CMS Collaboration, A. Hayrapetyan et al., *Search for light long-lived particles decaying to displaced jets in proton–proton collisions at $\sqrt{s} = 13.6$ TeV*, Rept. Prog. Phys. **88** (2025) no. 3, 037801, arXiv:2409.10806 [hep-ex].

[72] C. Bierlich et al., *A comprehensive guide to the physics and usage of PYTHIA 8.3*, SciPost Phys. Codeb. **2022** (2022) 8, arXiv:2203.11601 [hep-ph].

[73] ATLAS Collaboration, M. Aaboud et al., *Jet reconstruction and performance using particle flow with the ATLAS Detector*, Eur. Phys. J. C **77** (2017) no. 7, 466, arXiv:1703.10485 [hep-ex].

[74] ATLAS Collaboration, Aad, Georges et al., *Transforming jet flavour tagging at ATLAS*, Nature Commun. **17** (2026) no. 1, 541, arXiv:2505.19689 [hep-ex].

[75] *Higgs to Displaced Jet Analysis Repository*, https://gitlab.cern.ch/atlas-phys/exot/ueh/DisplacedJetsRun3/displaced-jets-analysis.

[76] ATLAS Collaboration, *easyjet analysis framework*, 2025. https://doi.org/10.5281/zenodo.15927813. Repository: https://gitlab.cern.ch/easyjet/easyjet.

[77] ATLAS Collaboration, *FastFrames*, 2025. https://doi.org/10.5281/zenodo.14773274. Repository: https://gitlab.cern.ch/atlas-amglab/fastframes.

[78] ATLAS Collaboration, *TRExFitter*, 2025. https://doi.org/10.5281/zenodo.14845712. Repository: https://gitlab.cern.ch/TRExStats/TRExFitter.

[79] ATLAS Collaboration, *Technical Design Report for the ATLAS Inner Tracker Strip Detector*, Tech. Rep. CERN-LHCC-2017-005, ATLAS-TDR-025, 4, 2017.

[80] ATLAS Collaboration, *A High-Granularity Timing Detector for the ATLAS Phase-II Upgrade: Technical Design Report*, Tech. Rep. CERN-LHCC-2020-007, ATLAS-TDR-031.

[81] ATLAS Collaboration, G. Aad et al., *Study of Higgs boson pair production in the $HH \to b\bar{b}\gamma\gamma$ final state with 308 fb^{-1} of data collected at $\sqrt{s} = 13$ TeV and 13.6 TeV by the ATLAS experiment*, Phys. Lett. B **876** (2026) 140280, arXiv:2507.03495 [hep-ex].

[82] J. Elias-Miró, J. R. Espinosa, G. F. Giudice, G. Isidori, A. Riotto, and A. Strumia, *Higgs mass implications on the stability of the electroweak vacuum*, Physics Letters B **709** (2012) no. 3, 222–228. https://www.sciencedirect.com/science/article/pii/S037026931200144X.

[83] F. Bezrukov and M. Shaposhnikov, *The Standard Model Higgs boson as the inflaton*, Physics Letters B **659** (2008) no. 3, 703–706. https://www.sciencedirect.com/science/article/pii/S0370269307014128.

[84] D. E. Morrissey and M. J. Ramsey-Musolf, *Electroweak baryogenesis*, New Journal of Physics **14** (2012) no. 12, 125003. https://dx.doi.org/10.1088/1367-2630/14/12/125003.

[85] E. Brost and L. Cadamuro, *Searching for Pairs of Higgs Bosons in the LHC Run 2 Dataset*, Symmetry **14** (2022) no. 7, 1467.

[86] ATLAS Collaboration, G. Aad et al., *Search for short- and long-lived axion-like particles in $H \to aa \to 4\gamma$ decays with the ATLAS experiment at the LHC*, Eur. Phys. J. C **84** (2024) no. 7, 742, arXiv:2312.03306 [hep-ex].

[87] ATLAS Collaboration, G. Aad et al., *Performance of electron and photon triggers in ATLAS during LHC Run 2*, Eur. Phys. J. C **80** (2020) no. 1, 47, arXiv:1909.00761 [hep-ex].

[88] Advanced Micro Devices, Inc. (AMD), *AMD Vitis HLS*, https://www.amd.com/en/products/software/adaptive-socs-and-fpgas/vitis/vitis-hls.html.

[89] D. P. Kingma and J. Ba, *Adam: A Method for Stochastic Optimization*, 2017. arXiv:1412.6980 [cs.LG], https://arxiv.org/abs/1412.6980.

[90] E. S. Gardner Jr., *Exponential smoothing: The state of the art*, Journal of Forecasting **4** (1985) no. 1, 1–28, https://onlinelibrary.wiley.com/doi/pdf/10.1002/for.3980040103. https://onlinelibrary.wiley.com/doi/abs/10.1002/for.3980040103.

[91] S. Ruder, *An overview of gradient descent optimization algorithms*, CoRR **abs/1609.04747** (2016) , 1609.04747. http://arxiv.org/abs/1609.04747.

[92] *Global Diphoton Trigger Repo*, 2026. https://gitlab.cern.ch/atlas-tdaq-l0global-offlinesw/algorithms/photonpointing.

[93] M. Watson, F. Chollet, D. Sreepathihalli, S. Saadat, R. Sampath, G. Rasskin, , S. Zhu, V. Singh, L. Wood, Z. Tan, I. Stenbit, C. Qian, J. Bischof, et al., *KerasHub*, https://github.com/keras-team/keras-hub, 2024.

[94] FastML Team, *hls4ml*, Mar., 2026. https://fastmachinelearning.org/hls4ml.

[95] J. Duarte, S. Han, P. Harris, S. Jindariani, E. Kreinar, B. Kreis, J. Ngadiuba, M. Pierini, R. Rivera, N. Tran, and Z. Wu, *Fast inference of deep neural networks on FPGAs with hls4ml*, JINST **13** (2018) no. 07, P07027.