Search for Displaced Muon Pairs from the Decays of Low-Mass Long-Lived Dark Photons with the ATLAS Experiment

by

Kehang Bai

A dissertation accepted and approved in partial fulfillment of the

requirements for the degree of

Doctor of Philosophy

in Physics

Dissertation Committee:

James Brau, Chair

Laura Jeanty, Advisor

Timothy Cohen, Core Member

Cristopher Niell, Institutional Representative

University of Oregon

Spring 2026

© 2026 Kehang Bai
All rights reserved.

DISSERTATION ABSTRACT

Kehang Bai

Doctor of Philosophy in Physics

Title: Search for Displaced Muon Pairs from the Decays of Low-Mass Long-Lived Dark Photons with the ATLAS Experiment

The Standard Model of particle physics, while extraordinarily successful in describing the fundamental constituents of matter and their interactions, leaves several deep questions unanswered, including the nature of dark matter, the origin of neutrino mass, and the matter–antimatter asymmetry. Potential solutions to these problems may exist within an entirely new sector of particles, known as a dark sector. Such a sector does not interact with the strong and electroweak forces, but may interact with the Standard Model through mediator particles. A well-motivated mediator is a dark photon, a massive vector boson that mixes kinetically with the Standard Model photon. This work explores new regions of dark photon parameter space with the ATLAS experiment at the Large Hadron Collider, targeting small dark photon masses and kinetic mixing strengths where dark photons are long-lived.

This dissertation presents the complete analysis strategy, background estimation, and expected sensitivity of the first ATLAS search for collimated, displaced dimuon vertices arising from low-mass, long-lived dark photons. The analysis uses proton–proton collision data at center-of-mass energies of $\sqrt{s}$ = 13 and 13.6 TeV, corresponding to integrated luminosities of 140 and 162 fb^{-1} collected during Run 2 and Run 3, respectively. The expected sensitivity extends beyond existing dark photon searches in the targeted regions of parameter space. The analysis further establishes novel reconstruction methods for displaced muons and secondary vertices, which can also benefit future long-lived particle searches at ATLAS.

This dissertation includes previously published co-authored material.

ACKNOWLEDGEMENTS

I am incredibly grateful for the opportunity of this doctoral degree which allowed me to think deeply about fundamental physics and conduct experiments to probe the nature of the universe. To me, this has been a dream come true. For enabling the pursuit of fundamental science, I give my greatest gratitude to the University of Oregon, to the Hong Kong University of Science and Technology, to the education I received in Hohhot, Inner Mongolia, and to the experimental collaborations at the Large Hadron Collider at CERN.

I will never forget the support, inspiration, and friendship of each individual I met throughout this journey, who together shaped me into the scientist I am today. It is the passion, dedication, and kindness of my advisor, my mentors, and my colleagues and friends that I have always admired. It is also their joy in physics and life that makes me feel hopeful for the future.

To Laura Jeanty: thank you for being the most amazing advisor that I could have ever asked for. Your trust, encouragement, and guidance have helped me gradually find my way to become a more independent, capable, and happier physicist. You are always open-minded about me trying something new, but always able to keep me on the right track. You are patient with me when I get stuck or lost, and you are always empathetic when I am feeling down. Many times, you understood me more than I had understood myself. You have set the best example of a great experimentalist and I have been so lucky to be your student.

To Tim Cohen: thank you for being the theory member of my committee, for teaching me QFT, and for believing in me. You always saw a theorist in me and reminded me that this is also a sign of a good experimentalist. Thank you for patiently answering all my questions, for being there for me through the ups and downs, and for fostering my vision for particle physics.

To Jim Brau: thank you for being the chair of my committee, for holding my work to a high standard and for introducing me to the Oregon ATLAS group. When we first met in Hong Kong, it was your enthusiasm for future colliders, your kindness in talking with a nervous

undergraduate, and your incredible joy in discussing physics that convinced me that a PhD in particle physics is a meaningful life decision.

To Cris Niell: thank you for being my external committee member. You are an incredible scientist, and I am honored to work with you on my dissertation.

To Simone Pagan Griso: thank you for the significant role you played in my PhD through continuous mentorship and constant inspiration. With you, a chat by the coffee machine always turns into new ideas, a lunch discussion always touches upon the past, present and future of particle physics, a climbing session is always filled with laughter and thrills, and even a "Kehang joke" can be funny sometimes. Thank you for helping me forge the idea of my dissertation topic, and succeed in following it through.

To Timon Heim: thank you for taking me on adventures in hardware and teaching me the very basics of silicon detectors. Your door is always open for me to ask questions. Your enthusiasm and talent in instrumentation are contagious. Despite your sarcastic humor, you are the true glue of the group, both at work and outside of work. Thank you for taking particularly good care of me.

To Marjorie Shapiro: thank you for joining my analysis and being an important mentor of mine. It has been a huge honor and a lot of fun to do a search with you. You are always so sharp and curious. I look up to your passion and dedication for physics. Thank you for always being available to talk and for giving me the advice that I needed.

To Javier Montejo Berlingen: thank you for being one of my first mentors in ATLAS and for teaching me everything. I was very lucky to work with you as a CERN Summer Student, and later during my early PhD. Thank you for sharing your love for physics, and for inspiring me to one day be a great physicist like you.

To the dark photon analysis team: thank you for being part of this journey with me and for all your excellent work. Maria, Giulia, Axel, Alic, Khurshid, Haoran, Prathik and Christina, thank you for making the analysis possible through your amazing work. Laura, Simone, Marjorie,

Colin, Kate, Rebeca and Ben, thank you for providing feedback and supervision. Special thanks to Colin for your indispensable help in tackling the difficult background composition. Thank you, Maria and Giulia, for coordinating this ambitious project with me, thank you, Alic, for always coming in for help when we needed it, and thank you, Axel, for taking on many critical tasks. I have learned so much from all of you, and it has been an honor to work together.

To the ATLAS Pixel operations team: thank you for helping me learn about our detector, and letting me take on the important role of taking care of it. James, Kerstin, and Martin, thank you for your patience and for teaching me everything about Pixel operations. Yue, Elia, Simone, Chris, Sergi, Marcelo, Vadim and Tobias, thank you for being amazing to work with.

I want to thank all the Oregon High Energy Physics faculty for cultivating a community of excellence and warmth. David, thank you for all the lunch chats at CERN, where I learned so much about particle physics and more importantly, had someone to talk to. Stephanie, thank you for always having my back, and for giving me invaluable advice on leadership and being a woman in physics. Eric, thank you for sharing your love for taus, and for all the fun physics and non-physics discussions at CERN. Graham, thank you for the enjoyable lectures, for all the dinners and climbing sessions, and for many conversations that shaped the way I think about physics. Tien-Tien, Spencer and Ben Farr, thank you for being amazing faculty members of IFS and for always being so friendly to me. Additional thanks to Anne McGinley and Claire Staley for your help with all the administrative work.

I would also like to thank many of the staff scientists and faculty members of the ATLAS group at LBNL. Elisabetta, thank you for introducing me to ITk and being my QT supervisor, and for much of your care and support. Maurice, thank you for providing guidance to my instrumentation work, for giving me career advice, and for hosting a wonderful dumpling dinner. Carl, thank you for the greatest stories and food. Zach, thank you for answering all my questions and all the fun chats at CERN. Fabio, thank you for inviting me to the coffee chats, so

we get to talk about interesting physics every week. Haichen, thank you for always asking great questions so I can learn. Kevin and Sandra, thank you for your friendliness and hospitality.

To the Nachman group: thank you for welcoming me to the HEP ML community and for all the creative work. Ben, thank you for introducing me to the group, brainstorming with me, and helping me throughout our project. I have learned so much from you. Radha, thank you for being a great collaborator and a friend to talk to. Vini, Aishik, Elham, Fernando, Sascha, Dennis, Mariel and Krish, thank you for all the interesting discussions and memorable get-togethers.

I want to thank many theory friends and colleagues that I crossed paths with over the past years for all the wonderful conversations and your appreciation of experiments. Simon Knapen, thank you for finding the "triangle of hope" with me for my dissertation topic, and for many intriguing discussions on collider physics. Christiane Scherb, thank you for your help with the Hidden Valley models and for being an incredible friend. Dan Carney, thank you for the donation of a blackboard. Kailash and Orion, thank you for being buddies. I am grateful for many members of the CERN theory group for useful discussion and for inviting me to journal clubs, lunches and social activities, including Miguel Escudero Abenza, Joe Davighi, Guilherme Guedes, Itay Bloch, Mario Reig Lopez, Marc Riembau, Ryan Plestid and Matthew McCullough. Special thanks to Maria Ramos and Rachel Houtz for being inspiring women physicists and for your friendship. Lastly, thank you, everyone at the ICTP Summer School.

I want to thank my friends at Oregon, especially those whom I took QFT and particle physics with and the ATLAS students I overlapped with. It has been fun to learn with you and learn from you. In particular, I thank Tom Bouley for helping me with QFT and for being a great friend. Thank you, Cristien, Kevin and Chester, for being the late night homework crew, and for doing everything else together in Eugene. Jennifer, Anthony, Nathan, Tim, Haidar and Austin, it has been fun to overlap and graduate together. Additionally, thank you Robin and Mei-Li for being my friends outside of physics and for making Eugene feel like home for me.

I am tremendously grateful for all my friends at Berkeley. My time as a visitor at LBNL has been the highlight of my PhD, and I am so glad I got to spend it with all of you. Emily, thank you for being my first and best friend when I arrived, for saying yes to every adventure, and for bringing so much positivity to my life. Miha, thank you for the countless physics discussions, for the enormous amount of help you've given me, and for your friendship. Anne, thank you for all the long walks home, for being so caring and kind, and for always being there for me. Ian, thank you for your mysterious friendship, and for always being interested in what I am working on. Aleksandra, Angira, Miao, Maria Giovanna, Maria Mironova, and Rebecca, thank you for the brunches, the night-outs, and for every empowering conversation we had. Julien, Elliot, Zhicai and Xiangyang, thank you for all the fun memories and outdoor adventures. Liam, Laura, Luc, Charlie, Hadley, and Luke, thank you for your great energies in the offices, at CERN and at climbing gyms. Louis-Guillaume, Carlo and Tomohiro, thank you for always keeping your lunch table open for me at CERN.

To my Saint-Genis crew at CERN: thank you for being there for me through a stressful time and at a challenging place. Sylvia and Michelle, thank you for being there with me from the start till the end, for always being by my side, and for growing together. Jacob, Siddharth, and Jared, thank you for keeping us together, and for always bringing laughter, energy or a bottle of wine. Ryan, thank you for the spontaneous coffees, lunches and dinners, for pushing me to be a better physicist, and for being my friend through everything.

Thank you to all the Oregon students and postdocs at CERN, for always being helpful and friendly. Dani, Jochen and Ondra, Despoina, Federica and Goetz, thank you for your help both at work and outside of work, and for making the B32 basement a nice place to be. Juliet, Calla and Anthony, thank you for being great buddies and keeping me company.

To many wonderful colleagues and friends from everywhere else in ATLAS and CERN: thank you for making my experience in ATLAS and at CERN so memorable. Cristiano, Jackson and Elena, thank you for being the most supportive subgroup conveners, and for being extremely

fun to hang out with outside of work. Giordon, Lingxin, and Hide, thank you for your help on ITk and for the fun module QC hackathons. Haider, thank you for always giving me timely advice and pushing me to do my best. Tadej, thank you for being my friend and looking out for my Athena merge requests. James, Julia and Audrey, thank you for being so warm and kind to me. Dan and Claire, thank you for the enjoyable dipole lunches. Iacopo and Guglielmo, thank you for the consistent climbing sessions. Ben Rosser, thank you for never running out of physics to talk about. Alexis, thank you for the journal clubs and for crossing the finish line with me.

To all my mentors from undergraduate: thank you for encouraging me to pursue research and for seeing potential in me. Kirill Prokofiev and Ki Lie, thank you for your supervision of my final-year capstone project with the HKUST ATLAS group. Till Eifert, thank you for your supervision at CERN. Andy Buckley, thank you for introducing me to ATLAS, and for your supervision at Glasgow. Tao Liu, thank you for taking me on as your student. Zhiyu Yang, thank you for providing me my first research project. Andrew Horner, thank you for the best C++ lectures and for keeping me inspired. Andy Cohen, thank you for your guidance in finding my path in particle physics.

To friends who knew me before graduate school: thank you for never losing touch, for visits or calls across continents, and for your precious, lasting friendship.

To my family: thank you for your love, trust, and support throughout my life. I apologize for being away from home for so long, but you always understand and encourage me to follow my dream. I want to especially thank my mom, who has been the biggest hero in my life.

To Kevin Langhoff: thank you for being the one who understands me the most, who supports me the most, and who cares for me the most, and thank you for being the one.

Dedicated to my parents and grandparents.

TABLE OF CONTENTS

Chapter Page

1. Introduction 33
2. Theory Motivations 35
2.1. Standard Model of Particle Physics 35
2.1.1. The formalism 36
2.1.2. Total parameters 39
2.2. Physics Beyond the Standard Model 40
2.2.1. Dark Matter and Dark Sector 40
2.2.2. Dark Photon 41
3. The Large Hadron Collider 47
3.1. Historical prelude 47
3.2. The Large Hadron Collider 53
3.2.1. Basic principles and layout 53
3.2.2. Operations 56
3.3. High Luminosity LHC 57
4. The ATLAS Experiment 60
4.1. Coordinate System 61
4.2. Magnet System 62
4.3. Detector system 63
4.3.1. Inner Detector 63
4.3.2. Calorimeter 67
4.3.3. Muon Spectrometer 70
4.4. Trigger system 73

Chapter Page

4.5. Physics Program . . . 74

5. Event reconstruction . . . 76

5.1. Track and vertex reconstruction . . . 76

5.1.1. Track reconstruction . . . 77

5.1.2. Vertex reconstruction . . . 82

5.1.2.1. Primary vertex reconstruction . . . 82

5.1.2.2. Secondary vertex reconstruction . . . 83

5.2. Muon reconstruction . . . 85

5.2.1. Standard Combined (CB) Muons . . . 85

5.2.2. Large Radius Tracking (LRT) Muons . . . 85

5.2.3. Zero Pixel Hit Muons . . . 86

5.3. Electron and photon reconstruction . . . 86

5.4. Jet reconstruction . . . 87

5.5. Tau leptons . . . 90

5.6. Missing transverse momentum . . . 90

6. Search for Low-Mass, Long-Lived Dark Photons to Muons . . . 92

6.1. Analysis Overview . . . 92

6.2. Data and Simulation Samples . . . 93

6.2.1. Data . . . 93

6.2.2. Simulated Samples . . . 94

6.2.2.1. Signal . . . 95

6.2.2.2. Background . . . 97

6.3. Analysis objects . . . 98

6.3.1. Muons . . . 98

6.3.2. Secondary vertex reconstruction with muons . . . 100

Chapter Page

6.3.3. Vertex pulling . . . 103
6.3.4. Triggers . . . 105
6.4. Event selections . . . 106
6.4.1. Baseline Selections . . . 106
6.4.2. Signal Region Definition . . . 107
6.4.2.1. Muon selections . . . 107
6.4.2.2. Displaced vertex selections . . . 113
6.4.2.3. 2-muon Channel Signal Region . . . 116
6.4.2.4. 4-muon Channel Signal Region . . . 118
6.4.3. Control region definitions . . . 121
6.4.4. Signal vs background comparisons . . . 122
6.5. Background Composition . . . 128
6.5.1. Heavy flavor background . . . 129
6.5.2. Material interactions . . . 131
6.5.3. Prompt background . . . 133
6.5.4. Accidental crossings . . . 136
6.5.5. Cosmic muon background . . . 136
6.5.6. Summary of background composition . . . 136
6.6. Background Estimation . . . 137
6.6.1. 2-muon channel low-mass region . . . 138
6.6.2. Signal Modeling . . . 140
6.6.2.1. Interpolation of signal parametrization . . . 141
6.6.2.2. Signal injection test . . . 143
6.6.3. Background Templates . . . 144
6.6.4. Background Modeling . . . 146

Chapter Page

6.6.4.1. Background model criteria . . . 146
6.6.4.2. Goodness of the background-only fit . . . 147
6.6.4.3. Spurious signal test . . . 150
6.6.5. 2-muon channel high-mass region . . . 153
6.6.6. 4-muon data-driven background estimation . . . 155
6.7. Systematic Uncertainties . . . 156
6.7.1. Uncertainties on background estimation methods . . . 156
6.7.1.1. Spurious signal uncertainty . . . 156
6.7.1.2. ABCD estimation uncertainty . . . 160
6.7.2. Uncertainties on the signal process . . . 161
6.7.2.1. Experimental uncertainties . . . 161
6.7.2.2. Theoretical and modeling uncertainties . . . 162
6.8. Statistical Analysis and Expected Limits . . . 162
6.8.1. 2-muon channel: low-mass region . . . 163
6.8.2. 2-muon channel: high-mass region . . . 165
6.8.3. 4-muon channel . . . 166
7. Conclusions and Outlook . . . 168

APPENDICES

A. Characterization of Vertex Pulling Effect in ATLAS Secondary Vertex Reconstruction . 170
A.1. Characterization with the SPG simulation . . . 171
A.2. Characterization with signal simulation . . . 173
A.3. Characterization with control region data . . . 175
B. Analog Readback for the ATLAS Inner Tracker Upgrade Pixel Module Electrical Quality Control . . . 177
B.1. The ATLAS Inner Tracker upgrade and Pixel Modules . . . 177

Chapter Page

B.2. Analog readback in module QC . . . 179

B.2.1. Chip sensing and monitoring . . . 180

B.3. Temperature measurement . . . 182

B.3.1. External NTC temperature sensor measurement . . . 183

B.3.2. DCS external NTC . . . 184

B.3.3. MOS sensors . . . 184

B.4. VDDA and VDDD measurement . . . 185

B.5. Ring Oscillator measurement . . . 186

C. HL-LHC ATLAS Inner Tracker Pixel Detector ToT-to-Charge Calibration and Tracking Performance . . . 188

C.1. Time-over-threshold and charge reconstruction . . . 188

C.2. ToT-versus-charge response in the ITk Pixel front-end . . . 190

C.3. Charge calibration at pixel-level vs at FE-level . . . 191

C.4. Impact on tracking performance . . . 194

C.5. Challenges for future investigations . . . 197

D. Non-resonant Anomaly Detection with Background Extrapolation . . . 198

D.1. Introduction . . . 198

D.2. Methodology . . . 201

D.2.1. Locating overdensities . . . 201

D.2.2. Constructing background templates . . . 203

D.2.3. Network architectures and training specifications . . . 205

D.3. Application with dark QCD jets . . . 207

D.3.1. Dataset specifications . . . 208

D.3.2. Results . . . 211

D.3.2.1. Closure tests . . . 211

D.3.2.2. Signal detection tests . . . 213

D.4. Conclusions 214

D.5. Comparing different signal parameters 216

D.6. Modifying the simulation 219

REFERENCES CITED 222

LIST OF FIGURES

Figure Page

1. Kinetic mixing of dark photon with the SM photon. Figure (a) shows how the kinetic mixing can be generated by integrating out heavy new physics interacting with both γ and A' at the 1-loop level. Figure (b) shows the kinetic mixing in low energy EFT, below the mass scale of the particles in the loop. 41

2. Dark photon production and decay diagrams for two models. Figure (a) shows direct production from the minimal model. Figure (b) shows Higgs-portal production from the HAHM model. 43

3. Decay branching fraction of A' from [6]. 43

4. Dark photon decay length (a) and lifetime (b) as a function of kinetic mixing coupling, ε, and dark photon mass, $m_{A'}$. Figures taken from [7]. 44

5. Current constraints on visible dark photons from different experiments [7]. 45

6. Current constraints on HAHM dark photon from ATLAS [8, 9] (a), and CMS [10] (b). 45

7. Estimated signal production rate from $A' \to \mu^+\mu^-$ events in the minimal dark photon model. Figure shows the number of events satisfying jet $p_{\mathrm{T}} >$ 40 GeV and displacement $L_{\mathrm{dis}} \in [1,300]$ mm across the $(m_{A'}, \varepsilon)$ parameter space. . . . 46

8. The discovery photograph of the positron in a cloud chamber by Carl Anderson on August 2, 1932 [11]. In this photograph, a cosmic-ray particle moves upward, passes through the lead plate at the center, and curves to the left under a magnetic field. The curvature indicates its charge is positive, and the ionization shows that the particle is not a proton. The positron was the first evidence of anti-matter. Anderson won the Nobel Prize in Physics in 1936 for this discovery. 48

9. The design of the cyclotron by Ernest Lawrence from his 1934 patent [12]. 49

10. Maximum available center-of-mass energy versus time for different accelerator experiments [13] before the Large Hadron Collider. The Superconducting Super Collider (SSC), planned for construction in Texas, United States, was canceled in 1993 before completion. 51

Figure Page

11. The discovery of the J/ψ meson in November 1974. Figure (a) shows the J particle discovered at Brookhaven National Laboratory led by Samuel Ting [14], observed as a sharp resonance in the invariant mass spectrum of e^+e^- pairs produced in collisions of protons on a beryllium target. Figure (b) shows the ψ particle discovered at SLAC led by Burton Richter [15], observed as a sharp peak in the e^+e^- annihilation cross section at $E_{\text{cm}} \approx 3.1$ GeV for final states of hadrons, e^+e^-, and $\mu^+\mu^-$, from top to bottom, respectively. This simultaneous discovery of the J/ψ is famously referred to as the "November Revolution". 52

12. The cross section of $e^+e^- \to$ hadrons as a function of center-of-mass energy from a combination of experimental measurements. The horizontal axis shows the center-of-mass energy in GeV and the vertical axis shows the cross section σ in mb. Prominent resonances corresponding to the ω, ρ, ϕ, J/ψ, $\psi(2S)$, Υ, and Z bosons are clearly visible. Figure taken from the Particle Data Group [16]. 53

13. The schematic layout of the LHC. There are eight insertion regions (IRs), each labeled with its main purpose. Figure taken from [19]. 55

14. The LHC dipole magnet, showing the cross-sectional view (a) and the magnetic flux distribution (b). Figures taken from the LHC Design Report [20] and CERN Courier [21]. 55

15. The luminosity-weighted distribution of the mean number of interactions per crossing for proton–proton collisions data recorded by ATLAS during 2011–2025 LHC runs. Figure taken from the ATLAS Luminosity Public Results page [22]. 58

16. Evolution of the phase-space ellipses at position s along the ring. The transverse beam size is $\sigma(s) = \sqrt{\beta(s)\varepsilon}$, where $\beta(s)$ (emittance) modulates the envelope of the transverse oscillations and ε is the phase-space volume of the beam. Figure adapted from [24]. 59

17. A schematic cut-away view of the ATLAS detector in Run 3. The detector has a cylindrical structure and is 25 m in height and 44 m in length. Figure taken from [26]. 60

18. The ATLAS coordinate system showing the cylindrical coordinates (r, ϕ, z) (a), and the definition of pseudorapidity η as a function of polar angle θ (b). 61

Figure Page

19. The geometry of the ATLAS magnetic system. The barrel and end-cap toroids and the central solenoid inside the tile calorimeter are shown in red (a). The magnetic field maps are shown in the transverse plane at $z = -20$ cm (b) and in the r–z plane at $y = 10$ cm (c). Figures taken from [27, 28]. 62

20. A cut-away view of the ATLAS Inner Detector layout [26]. It provides high-precision measurements of charged-particle trajectories. 63

21. Illustrations of the ATLAS Pixel detector system. A picture of the barrel pixel Layer 2 is shown in (a). The installation of the IBL around the beam pipe is shown in (b), and its layout in (c). Figures taken from [27, 30, 31]. 64

22. A schematic illustration of a pixel detector module. A schematic illustration of a charged particle passing through a single pixel cell (a). The layout of the ATLAS barrel pixel module is shown in (b). Figures taken from [27, 30, 31]. 65

23. ATLAS SCT detector modules. A picture of the end-cap disk layer of SCT is shown in (a). The layout of a barrel SCT module is shown in (b). Figures taken from [27]. 66

24. ATLAS TRT detector. A picture of the TRT barrel structure is shown in (a). A schematic view of the straw-tube layout is shown in (b). Figures taken from [27, 33]. 66

25. A cut-away view of the ATLAS Calorimeter layout [26]. 67

26. Schematic illustrations of the ATLAS electromagnetic calorimeter (ECAL) layout. The three-layer cell structure of the barrel ECAL at $\eta = 0$ is shown in (a). The accordion geometry of the LAr ECAL, with alternating lead absorbers and kapton readout electrodes separated by LAr gaps as the active medium, is shown in (b). Figures taken from [27, 34]. 69

27. Schematic illustration of a tile calorimeter module, showing the alternating steel absorbers and scintillating tiles with the optical readout. Figure taken from [34]. 69

28. A cross-sectional view of the ATLAS Muon Spectrometer, showing the arrangement of MDT and CSC precision tracking chambers and the RPC and TGC trigger chambers in the barrel and end-cap regions [26]. 70

29. Schematic views of an ATLAS MDT tube. Figure (a) shows a cross-sectional view of an MDT tube. Figure (b) shows a longitudinal view of an MDT tube. Figures taken from [27]. 71

30. Schematic views of the ATLAS MS trigger chambers. Figure (a) shows the layout of an RPC. Figure (b) shows the layout of a TGC. Figures taken from [27]. 72

31. Schematic views of the two NSW detector technologies. The small-strip Thin Gap Chambers consist of pad and strip cathode boards with sandwiching anode wires at a pitch of 1.8 mm in a 1.4 mm gas gap, shown in (a). The Micromegas detectors consist of a 5 mm conversion/drift gap between the drift electrode and the micromesh, with a 128 μm amplification gap at 505 V above the readout strip PCB, shown in (b). Figures taken from [35]. 72

32. A schematic illustration of the ATLAS trigger and data acquisition system in Run 3. Figure taken from [26]. 73

33. Summary of SM production cross section measurements by ATLAS, corrected for branching fractions and compared to theoretical predictions. The measurements span a wide range of processes and center-of-mass energies. Figure taken from [36]. 75

34. A detailed layout of the ATLAS Inner Detector [37]. 77

35. Illustration of the five track parameters ($d_0, z_0, \phi_0, \theta, q/p$) used to parameterize charged-particle trajectories at the perigee. Figure (a) shows the transverse and longitudinal impact parameters, d_0 and z_0, of the particle track (orange), where $\vec{l}$ is the vector from the reference point to the perigee and $\vec{p}$ is the track momentum vector. Figure (b) illustrates the azimuthal angle ϕ_0 and polar angle θ_0 of the momentum vector $\vec{p}$ at the perigee. Figures taken from [38]. 79

36. A simplified schematic illustration of inside-out tracking sequence's pattern recognition stage shown in the transverse plane of the Inner Detector. Figures taken from [38]. 80

37. A schematic illustration of secondary vertex reconstruction shown in the transverse plane of the Inner Detector [42]. 84

38. Schematic illustration of an electron traversing the ATLAS tracker and calorimeter. The electron forms an ID track (red), emits a bremsstrahlung photon (dashed red), and deposits energy in clusters of calorimeter cells (green), which are used to build the final supercluster. Figure taken from [48]. 88

39. ATLAS simulation of a dijet event where at least one jet enters the calorimeter and topo-clusters are formed in three different stages: (a) cells with signal significance larger than four are selected to seed topo-clusters, (b) cells with signal significance larger than 2 are added to the topo-clusters, and (c) topo-clusters form by all clustered cells. Figures taken from [45]. 89

40. A candidate $W^- \to \mu^- \bar{\nu}_\mu$ event recorded by ATLAS in proton–proton collisions at $\sqrt{s} = 13$ TeV during a low pile-up run. The orange line shows the reconstructed muon track and the red arrow indicates the missing transverse momentum, consistent with an escaping neutrino. Figure taken from [53]. 91

41. Feynman diagrams of HAHM signal in the 2-muon (a) and 4-muon (b) final states. 95

42. The full signal grid for the HAHM model. The z-axis shows the number of signal events at generator level, assuming a Higgs-to-dark-photon branching ratio $BR(H \to \gamma_D \gamma_D) = 1\%$ and transverse displacement of the dimuon vertex within 1–300 mm. 96

43. HAHM signal muon kinematic distributions at reconstruction level: muon transverse impact parameter d_0 (a), leading muon p_T (b), subleading muon p_T (c), dimuon vertex L_{xy} (d), and dimuon opening angle ΔR (e), for the 1.2 GeV dark photon with 0.1, 1.0, and 5.0 mm lifetime. 96

44. Muon reconstruction efficiency as a function of L_{xy} for different muon types. Events are selected to satisfy muon trigger requirements. 100

45. Vertexing efficiency for the four-muon final state as a function of L_{xy} for dark photon mass and lifetime (a) $m_{\gamma_D} = 300$ MeV, $c\tau = 0.1$ mm, (b) $m_{\gamma_D} = 600$ MeV, $c\tau = 1.0$ mm, (c) $m_{\gamma_D} = 1200$ MeV, $c\tau = 5.0$ mm, and (d) $m_{\gamma_D} = 4500$ MeV, $c\tau = 5.0$ mm. 102

46. The L_{xy} distribution of HAHM $2\mu 2e$ signal for a 300 MeV dark photon with (a) $c\tau = 0.01$ mm, (b) $c\tau = 0.1$ mm, (c) $c\tau = 1$ mm, and (d) $c\tau = 5$ mm. In (a), the signal has a very short lifetime, and a fraction of DVs near the IP are being reconstructed at IBL. In (b), (c), and (d), the signal has a longer lifetime, therefore more DVs are being reconstructed at every pixel layer. The vertex pulling effects in signal MC are manifested in the peaks at the pixel layers in these distributions. No other selections have been applied to these distributions except for requiring a pair of truth-matched opposite-sign (OS) muons. 104

47. Comparison of trigger efficiency of full trigger chain (all low-threshold, unprescaled single and dimuon triggers) and selected triggers of an example 4-muon channel sample (left) and 2-muon channel sample (right) 105

48. Kinematic distributions of signal and heavy flavor simulation for the 2-muon channel after the baseline selections: muon p_{T}, isolation, dimuon invariant mass, and dimuon p_{T}. 108

49. Kinematic distributions of signal and heavy flavor simulation for the 2-muon channel after the baseline selections: DV transverse displacement L_{xy}, L_{xy} significance, dimuon opening angle ΔR, dimuon pointing angle $\Delta\phi(\overrightarrow{DV} - \overrightarrow{PV}, \vec{p}_{\mu\mu})$, DV track multiplicity, and muon $d_0/(L_{\mathrm{xy}} m_{\mu\mu}/p_{\mathrm{T}}^{\mu\mu})$. 109

50. Kinematic distributions of signal and heavy flavor simulation for the 4-muon channel after the baseline selections: leading muon p_{T}, sub-leading muon p_{T}, and dimuon $m_{\mu\mu}$ for leading and sub-leading dimuons. 110

51. Kinematic distributions of signal and heavy flavor simulation for the 4-muon channel after the baseline selections: DV L_{xy}, muon track-based isolation, and muon calorimeter-based isolation for leading and sub-leading dimuons. 111

52. Signal region boundaries in the dimuon mass plane, defined by $m_{\mu\mu}^{\mathrm{Leading}}$ and $m_{\mu\mu}^{\mathrm{Subleading}}$. 119

53. Data in `fail Isolation CR`. (a) dimuon invariant mass vs L_{xy} and (b) cutflow. The cutflow shows all selections applied in the `fail Isolation CR`, which are the same selection as SR except for inverting the isolation requirement. 121

54. Kinematic distributions of 2-muon signal and heavy flavor MC under SR selections, and data 22–24 under CR Iso Loose selections: muon p_{T}, dimuon p_{T}, and opening angle. 124

55. Kinematic distributions of 2-muon signal and heavy flavor MC under SR selections, and data 22–24 under CR Iso Loose selections: dimuon invariant mass, transverse displacement, L_{xy} significance, and pointing angle. 125

56. Kinematic distributions of 4-muon signal and heavy flavor MC under SR selections, and data 22–24 under CR Iso Loose selections: muon p_{T}, dimuon p_{T}, and opening angle. 126

Figure Page

57. Kinematic distributions of 4-muon signal and heavy flavor MC under SR selections, and data 22–24 under CR Iso Loose selections: dimuon invariant mass, transverse displacement, dimuon mass difference, and pointing angle. 127

58. Distribution of the dimuon invariant mass $m_{\mu\mu}$ in different L_{xy} windows in 2022–2024 data in the Fail Loose Isolation CR. 128

59. Parent particles of the leading and subleading muons from HF background simulation with a loose 2μ channel SR selection. 129

60. Feynman diagrams for decays of b quarks, and a semi-leptonic decay of B^0 meson. 130

61. The dimuon invariant mass distributions in data and HF simulation. Data distribution is shown in the fail-isolation CR, while HF MC is shown in both fail-isolation CR and pass-isolation SR. Figure (a) shows distributions normalized to unity. Figure (b) shows the distributions where MC in CR is scaled to have the same number of events as data under the J/ψ mass window, and MC in SR is scaled by the same factor as the CR. 130

62. Kinematic distributions of material interaction backgrounds: L_{xy} (left), $m_{\mu\mu}$ (center), and leading muon p_{T} (right). The top row shows the hadronic τ sample ($W \to \tau\nu$ and $Z \to \tau\tau$) with Geant4 selections and photon conversions removed; the bottom row shows the $Z+\gamma$ sample with photon conversion selections applied. 132

63. A Feynman diagram for Drell–Yan production decaying to two muons. 133

64. The L_{xy} and dimuon invariant mass distributions in low-mass DY MC with SR selections corresponding to the full Run 2 luminosity. Figure (a) shows the L_{xy} distributions, with an existing L_{xy} requirement at 1 mm. Figure (b) shows the dimuon invariant mass distribution of the selected events. . . . 135

65. The L_{xy}, L_{xy} significance, and dimuon invariant mass distributions in low-mass DY MC with SR selections but removing L_{xy} and L_{xy} significance requirements. 135

Figure Page

66. Fits to known SM resonances in the dimuon mass spectrum used to determine the masking windows. The dataset used is the `Fail Isolation Loose` CR data from 2022–2024. Shown are the J/ψ (a), $\psi(2S)$ (b), ρ (c), and ϕ (d) resonances. Each resonance window is fitted with a Double-Sided Crystal Ball signal function plus a power-law background (solid line), with signal and background components shown separately as dashed lines. For the ϕ and ρ resonances, the fits are only performed in the L_{xy} = 1–2 mm region to obtain the best mass resolution. 139

67. Signal MC fits with DSCB functions, inclusive in L_{xy}. The number of entries corresponds to data 22–24 luminosity with $BR(H \to A'A') = 1\%$. 142

68. Double-Sided Crystal Ball (DSCB) shape parameters, mean (μ), width (σ), tail parameters (α_{lo}, α_{hi}) and power-law exponents (n_{lo}, n_{hi}), as a function of the dark photon signal mass m_{γ_D}, shown for four signal lifetimes $c\tau = 0.1, 1, 5, 20$ mm. The parameters are extracted from DSCB fits to the $m_{\mu\mu}$ distribution of simulated $H \to \gamma_D\gamma_D \to \mu\mu + X$ signal events at each mass point. Dashed lines show linear fits to the mass dependence, used to parameterize the signal shape as a function of mass. Points excluded from the linear fit by a 2σ iterative clipping procedure are shown as open markers. 143

69. Comparison of the dimuon invariant mass $m_{\mu\mu}$ (a) and transverse decay length L_{xy} (b) distributions between the SR J/ψ mass window (2.9–3.2 GeV) and scaled CR datasets, where the CR datasets are scaled to have the same number of events in the J/ψ mass window. 145

70. Background-only fits for the 0.5 GeV mass window using the `Power0` fit function, for the `Fail Isolation Loose` dataset. Shown are $L_{xy} \in$ [0,2] mm (a), [2,10] mm (b), [10,20] mm (c), [20,40] mm (d), and [40,300] mm (e). . . 148

71. Background-only fits for the 1.8 GeV mass window using the `Power1` fit function, for the `Fail Isolation Loose` dataset. Shown are $L_{xy} \in$ [0,2] mm (a), [2,5] mm (b), [5,10] mm (c), [10,20] mm (d), [20,40] mm (e), [40,60] mm (f), and [60,300] mm (g). 149

72. Background-only fits for the 4.5 GeV mass window using the `Power2` fit function, for the `Fail Isolation Loose` dataset. Shown is $L_{xy} \in$ [0,300] mm (a). . . 149

73. $S/\delta S$ distributions for the 0.5 GeV mass window using the `Power0` fit function, for the `Fail Isolation Loose` dataset scaled to expected SR yields. Shown are $L_{xy} \in$ [0,2] mm (a), [2,10] mm (b), [10,20] mm (c), [20,40] mm (d), and [40,300] mm (e). 151

Figure Page

74. $S/\delta S$ distributions for the 1.8 GeV mass window using the `Power1` fit function, for the `Fail Isolation Loose` dataset scaled to expected SR yields. Shown are $L_{xy} \in [0,2]$ mm (a), [2,5] mm (b), [5,10] mm (c), [10,20] mm (d), [20,40] mm (e), [40,60] mm (f), and [60,300] mm (g). 152

75. $S/\delta S$ distribution for the 4.5 GeV mass window using the `Power2` fit function, for the `Fail Isolation Loose` dataset scaled to expected SR yields. Shown is $L_{xy} \in [0,300]$ mm (a). 152

76. The ABCD method used for background estimation in the high-mass region. Illustration of the ABCD plane defined by muon isolation and the pointing angle $\Delta\phi(PV\text{–}DV, \vec{p}_{\mu\mu})$ (a). Dimuon invariant mass distributions in data in control regions B, C, and D (b). 154

77. Summary of ABCD background estimates for the 4μ channel. Observed event counts in the ABCD control regions for Run 2 and Run 3. 156

78. Spurious signal test results for the $m_{\gamma_D} = 1.8$ GeV signal mass window and $L_{xy} \in [1,2]$ mm region as a function of the template normalization, using the `Power0` function. (a) Absolute spurious signal size S as a function of normalization. (b) Relative spurious signal size $S/N_{\text{background}}$ as a function of normalization. 158

79. Expected 95% CL upper limits on the branching ratio BR($H \to \gamma_D\gamma_D$) as a function of the proper decay length $c\tau$, for dark photon masses m_{γ_D} = 0.4, 0.6, 1.8, and 4.5 GeV in the 2-muon channel low-mass region. The expected limits are derived from the partial Run 3 dataset ($\sqrt{s} = 13.6$ TeV, 162.4 fb^{-1}) under the background-only hypothesis. The solid line shows the median expected limit, while the green and yellow bands indicate the $\pm1\sigma$ and $\pm2\sigma$ uncertainty intervals, respectively. 164

80. Expected 95% CL upper limits on the branching ratio BR($H \to \gamma_D\gamma_D$) as a function of the dark photon mass m_{γ_D} for a fixed proper decay length $c\tau = 1$ mm. The expected limits are obtained from a sliding mass window parametric-fit across signal masses using the partial Run 3 dataset ($\sqrt{s}$ = 13.6 TeV, 162.4 fb^{-1}) under the background-only hypothesis. The solid line shows the median expected limit, while the green and yellow bands indicate the $\pm1\sigma$ and $\pm2\sigma$ uncertainty intervals, respectively. The gray shaded regions indicate excluded mass windows due to known SM resonances (ρ, ϕ, J/ψ, and $\psi(2S)$), as defined in Table 16. 165

81. Expected 95% CL upper limits on the branching ratio BR($H \to \gamma_D\gamma_D$) as a function of the proper decay length $c\tau$, shown for dark photon masses m_{γ_D} = 7 GeV and m_{γ_D} = 10 GeV in the 2-muon channel high-mass region. The expected limits are derived from the partial Run 3 dataset ($\sqrt{s}$ = 13.6 TeV, 162.4 fb^{-1}) under the background-only hypothesis. The solid line shows the median expected limit, while the green and yellow bands indicate the $\pm1\sigma$ and $\pm2\sigma$ uncertainty intervals, respectively. 166

82. Expected 95% CL upper limits on the branching ratio BR($H \to \gamma_D\gamma_D$) as a function of the dark photon mass m_{γ_D}, shown for four fixed proper decay lengths $c\tau$ = 0.01, 0.1, 1, and 5 mm in the 4-muon channel. The expected limits are derived from the partial Run 3 dataset ($\sqrt{s}$ = 13.6 TeV, 162.4 fb^{-1}) under the background-only hypothesis. The solid line shows the median expected limit, while the green and yellow bands indicate the $\pm1\sigma$ and $\pm2\sigma$ uncertainty intervals, respectively. 167

A.83. The $m_{\mu\mu}$, L_{xy}, and L_{xy} significance distributions in the prompt-muon SPG sample. 171

A.84. dimuon $\Delta\eta$ and $\Delta\phi$ distributions in the prompt-muon SPG sample, for vertices reconstructed at (a) small displacement ($L_{xy} < 10$ mm) and (b) large displacement ($L_{xy} > 20$ mm). The large-displacement region is dominated by vertex pulling, and shows that the muon pairs are more collimated. The x-axis is scaled by 10^{-3}. 172

A.85. The $m_{\mu\mu}$, dimuon $\Delta\phi$, and L_{xy} significance as a function of L_{xy} in the prompt-muon SPG sample. 173

A.86. Reconstructed L_{xy} distributions for the m_{γ_D} = 0.3 GeV signal at $c\tau$ = 0.01, 0.1, and 1 mm, shown before and after requiring L_{xy} significance > 2 and > 3. The dashed vertical lines label the pixel detector layers (IBL and PIX 1–3). 173

A.87. Reconstructed L_{xy} distributions for the m_{γ_D} = 0.6 and 1.2 GeV signals at $c\tau$ = 1 and 5 mm, shown before and after requiring L_{xy} significance > 2 and > 3. The dashed vertical lines label the pixel detector layers (IBL and PIX 1–3). 174

A.88. Reconstructed L_{xy} distributions for the m_{γ_D} = 0.3 GeV signal at $c\tau$ = 0.01, 0.1, and 1 mm, shown before and after requiring dimuon $\Delta\phi > 0.005$. The dashed vertical lines label the pixel detector layers (IBL and PIX 1–3). 174

A.89. Reconstructed L_{xy} distributions for the m_{γ_D} = 0.6 and 1.2 GeV signals at $c\tau$ = 1 and 5 mm, shown before and after requiring dimuon $\Delta\phi > 0.005$. The dashed vertical lines label the pixel detector layers (IBL and PIX 1–3). 175

A.90.Distributions of $m_{\mu\mu}$ and L_{xy} with $\Delta\phi$ required to be larger than 0.005 and 0.006 in the fail-L_{xy}-significance CR using 2022–2023 data. 176

B.91.A schematic illustration of the ATLAS ITk detector. Figure taken from [66]. 178

B.92.Detailed layout of the ITk. Figure (a) shows one quadrant of the ITk in the r–z plane, with the pixel layers in red and the strip layers in blue [66]. Figure (b) shows a comparison between the current Run 2-Run 3 ATLAS Inner Detector, and ITk in the transverse plane [67]. 178

B.93.An ITk Pixel quad module. Figure (a) shows a photograph of an assembled quad module, and Figure (b) shows a schematic layout, in which four front-end chips are bump-bonded to a single sensor and wire-bonded to the module PCB. 179

B.94.Diagram of the RD53B FE chip monitoring block. The current multiplexer (IMUX) and analog voltage multiplexer (VMUX) feed the input of the analog-to-digital converter (ADC). Figure taken from [65]. 180

B.95.Location of the main elements on the ITkPixV1.1 front-end chip, including the monitoring block, the analog and digital shunt-LDO regulators, the temperature and radiation sensors and the Ring Oscillators. Figure taken from [65]. 182

B.96.Placement of the NTC temperature sensors on the ITkPixV1.1 front-end chip. 183

B.97.Diagram of a MOS temperature sensor on the ITkPixV1.1 front-end chip. Figure taken from [65]. 184

B.98.Measured VDDA (a) and VDDD (b) as a function of the trim setting. 185

B.99.Diagram of a ring oscillator block in the ITkPixV1.1 front-end chip. Figure taken from [65]. 186

B.100Example ring oscillator measurements for the Bank A and Bank B ROSCs, showing the oscillation frequency as a function of VDDD. The figures correspond to Bank A (a), Bank B left (b), Bank B right (c), Bank B FF (d), Bank B LVT (e), and Bank B injection capacitor (f). 187

C.101 A schematic illustration of the charge reconstruction chain at the ATLAS Pixel detector. The charge deposition in the pixel front-ends is measured by time-over-threshold (ToT), which has limited resolution. In the offline software, simulations (SIM) of events go through a digitization (DIGI) step to convert analog charge into digitized ToT based on the FE configuration. Both simulated and raw data events are then reconstructed (RECO), where discrete ToT values are converted back to charge. 189

C.102 ToT response of the ITk Pixel front-end as a function of injected charge, measured on a single-chip module with a 50×50 μm^2 3D sensor tuned to a threshold of 1000 e^-. Figure (a) shows the ToT as a function of injected charge, and Figure (b) shows the injected charge per ToT count. 191

C.103 ToT-to-charge calibration look-up tables for a single front-end chip. Figure (a) shows the pixel-level calibration look-up table, where the reconstructed charge on the z-axis is given for each pixel channel on the x-axis and ToT value on the y-axis. Figure (b) shows the FE-level calibration look-up table, obtained by averaging the injected charge over all channels for each ToT value. 192

C.104 Closure test comparing the FE-level (per-chip) and pixel-level calibrations. The reconstructed charge Q as a function of the injected charge Q_0 is shown for the FE-level (a) and pixel-level (b) calibrations at the top. The corresponding relative residuals $(Q - Q_0)/Q_0$ are shown for the FE-level (c) and pixel-level (d) calibrations at the bottom. 193

C.105 RMS of the reconstructed-charge Q distribution as a function of the injected charge Q_0, for the pixel-level (per pix) and FE-level (per chip) calibrations. 193

C.106 Pixel hit position residual in x, comparing the baseline (**Original DIGI + Original RECO**) and the per-chip calibration (**Per-chip DIGI + Per-chip RECO**). . . . 195

C.107 Pixel hit position residual in x, comparing the per-chip (**Per-chip DIGI + Per-chip RECO**) and per-pixel (**Per-pixel DIGI + Per-pixel RECO**) calibrations. 196

C.108 Pixel hit position residual in x, comparing the matched per-pixel configuration (**Per-pixel DIGI + Per-pixel RECO**) and the mismatched configuration (**Per-pixel DIGI + Per-chip RECO**). 196

D.109 A schematic of the setup for non-resonant anomaly detection. The signal region (SR) is defined by a one-sided cut on a context variable $m > c$. While the figure shows a one-dimensional context variable, in practice the context can be multidimensional. A number of other features x are used for the CWoLa anomaly detection classifier. It is also possible to include m in the classifier. 203

D.110 A histogram of m_{jj}, corresponding to the reconstructed mass of the 4 TeV Z' particle, for a selection of choices within the Hidden Valley model with parameter r_{inv}. In this study, we set $r_{\text{inv}} = 1/3$, at which point the Z' signal is no longer clearly resonant. 209

D.111 Histograms of the observables used in the Hidden Valley non-resonant anomaly detection task, for background ("data" and "simulation") and signal events). 210

D.112 Feature distributions as constructed by the Reweight, Generate, and Morph methods in the absence of signal. Errorbars in the ratio plots reflect Poisson statistics. Note that the values of the context and features have all been minmax-scaled to the range (-2.5, 2.5). 212

D.113 Results for a classifier tasked with discriminating extrapolated background samples from true background. The "Ideal" corresponds to an idealized classifier, which has been trained to discriminate two statistically identical sets of background, and thus represents a realistic random classifier. Errorbars and uncertainty bands show a 68-percentile spread across 20 independent classifier runs. 213

D.114 Metrics for an anomaly detection classifier tasked with discriminating extrapolated background samples from data in the SR. Uncertainty bands represent the 68-percentile spread over 10 different signal injections, all of which have been score-averaged over 20 independent classifier runs. 215

D.115 Histograms of the observables used in the Hidden Valley non-resonant anomaly detection task, for background ("data" and "simulation") and signal events). 218

D.116 Metrics for an anomaly detection classifier tasked with discriminating extrapolated background samples from data in the SR. The background is QCD dijet events. The signal is a 2 TeV Z' producing two dark QCD jets with $2m_{q_{\text{D}}} = m_{\pi_{\text{D}}} = m_{\rho_{\text{D}}} = \Lambda_{\text{D}} = 200$ GeV. Uncertainty bands represent the 68-percentile spread over 10 different signal injections, all of which have been score-averaged over 20 independent classifier runs. 218

D.117 Metrics for an anomaly detection classifier tasked with discriminating extrapolated background samples from data in the SR. The background is QCD dijet events. The signal is a 3 TeV Z' producing two dark QCD jets with $2m_{q_D} = m_{\pi_D} = m_{\rho_D} = \Lambda_D = 300$ GeV. Uncertainty bands represent the 68-percentile spread over 10 different signal injections, all of which have been score-averaged over 20 independent classifier runs. 219

D.118 Histograms of the observables used in the Hidden Valley non-resonant anomaly detection task, for background ("data" and "simulation") and signal events). The QCD dijet simulation has been modified to appear different from QCD dijet data. 221

D.119 SIC at a rejection of 10^3 for an anomaly detection classifier tasked with discriminating extrapolated background samples from data in the SR. Uncertainty bands represent the 68-percentile spread over 10 different signal injections, all of which have been score-averaged over 20 independent classifier runs. 221

LIST OF TABLES

Table Page

1. Particle content of the Standard Model gauge group $SU(3) \times SU(2) \times U(1)$. Fermions come in three generations. 37

2. Key parameters of the LHC during its operational runs and the planned High-Luminosity upgrade (HL-LHC) [16]. 56

3. Comparison of key parameters between the nominal LHC and the planned High-Luminosity LHC (HL-LHC) for proton–proton collisions [23]. 59

4. Key differences in selection criteria used for primary (inside-out sequence only) and LRT tracking [40]. 81

5. Center-of-mass energy and integrated luminosity passing GRL quality requirements for each data-taking year used in this analysis. 94

6. Simulated background samples used in this analysis. 97

7. Summary of key muon identification requirements applied in this analysis for each reconstruction category. 99

8. Key differences between `VSI_Leptons` and `VSI_BoostedMuons`. 100

9. Summary of chosen muon triggers by year 105

10. Baseline selections for the 2μ and 4μ channels. 106

11. Common baseline selections for dimuon DVs. 107

12. Signal region selection for the 2μ analysis channel. "DV" denotes the displaced vertex. "Muons" refers to the two muons forming the DV. χ^2_{ID} and χ^2_{CB} are for the inner-detector and combined tracks, respectively. 117

13. Signal region selection for the 4μ analysis channel. "DV" denotes each displaced vertex. "Muons" refers to the two muons forming that DV. χ^2_{ID} and χ^2_{CB} are for the inner-detector and combined tracks, respectively. 120

14. Summary of control regions 122

15. Qualitative comparison of background contributions in the 2μ channel SR, estimated for the combined Run 2+3 luminosity of ~300 fb^{-1}. 137

Table Page

16. Known SM resonances. 138

17. Background templates from data control regions, and the corresponding scale factors (SFs) used to normalize the templates to the SR-like reference yield in the J/ψ peak. 144

18. L_{xy} regions used in each mass range. Adjacent default L_{xy} bins (Section 6.6.1) are merged where the event yield is expected to fall below $\mathcal{O}(10)$ events per $(m_{\mu\mu}, L_{xy})$ bin. Mass intervals overlapping the SM resonance vetoes (Table 16) are excluded. 145

19. Definitions of empirical functions for the background modeling. 146

20. List of tested empirical functional forms for the background modeling. 146

21. Selected background PDF for each mass range in the 2-muon channel low-mass parametric fit. 153

22. Observed events in the control regions and estimated expected events in the signal region A for the Run 3 dataset corresponding to 174.8 fb^{-1}. 154

23. Background estimates in the 4μ signal region from the ABCD method with Bayesian pseudo-experiment sampling. N_{bkg} is defined as the median of the pseudo-experiment ensemble, and the uncertainties correspond to the 16% and 84% intervals. 155

24. Spurious signal uncertainty derived from `Fail Isolation Loose` CR background template 159

25. Spurious signal uncertainty derived from `Fail Isolation Loose && Fail Pointing Angle` CR background template 159

26. Spurious signal uncertainty derived from `Fail Pointing Angle` CR background template 160

B.27 Current multiplexer (I_mux) assignments for the ITk Pixel FE chip [65]. 181

B.28 Voltage multiplexer (V_mux) assignments for the ITk Pixel FE chip [65]. 181

D.29 Breakdown of events generated for the Hidden Valley model. We found that large training sets in the CR were necessary for good performances of the Generate and Morph (the flow-based) methods, while the Reweight method worked well even with a smaller set. We did not explore oversampling in the SR, which may improve performance [125]. 209

Chapter 1

Introduction

To address the question of what the universe is made of, decades of experiments have discovered a set of elementary particles and their interactions described in the language of quantum field theory, known as the Standard Model. While remarkable, it is incomplete. Not only do many mysteries lie within the Standard Model itself, but there is also strong evidence for the existence of new physics that the Standard Model does not explain, such as dark matter. The search for physics beyond the Standard Model is therefore one of the most important scientific goals in fundamental particle physics.

One possible explanation for the origin of dark matter is the *dark sector* scenario: a new sector of particles that interacts with the Standard Model only through mediator particles, called portals. One of the simplest renormalizable portals is the dark photon, a massive vector boson that interacts with the visible photon through kinetic mixing. The Higgs boson, as the only fundamental scalar in the Standard Model, provides a second renormalizable portal to the dark sector through a mixing of the Higgs field with a new scalar. This Higgs portal motivates a class of models in which the Higgs boson can decay directly into pairs of dark photons, making high-energy proton–proton collisions at the Large Hadron Collider (LHC) the most suitable place to study such dark sector scenarios.

In the parameter space with a sufficiently small kinetic mixing coupling and relatively low dark-photon masses, dark photons produced at the LHC are predicted to have a macroscopic lifetime before decaying back into Standard Model particles. This long-lived feature provides a distinctive experimental signature where the decay of dark photons to fermions leads to displaced vertices in the detector volume, suppressing background from prompt Standard

Model processes. This work explores this unconventional experimental signature to probe an underexplored region of dark photon parameter space via the decay to muons.

This dissertation presents the complete analysis strategy, background estimation, and expected sensitivity of a search for long-lived dark photons decaying to muon pairs using the ATLAS detector. The analysis is performed using a proton–proton collision dataset from the full Run 2 and partial Run 3 of Large Hadron Collider data-taking periods. The analysis targets dark-photon masses from 300 MeV to 10 GeV and is sensitive to kinetic mixing parameters of order 10^{-4}–10^{-6}. At the time of completing this dissertation, the analysis has not yet been unblinded. The results presented in this dissertation are based on expected sensitivity derived from background estimates and signal simulation. The final statistical interpretation of observed collision data will follow the completion of the internal review process.

This dissertation presents the analysis with the Hidden Abelian Higgs Model (HAHM) interpretation, while the minimal dark photon model interpretation will be included in the full search. The analysis consists of two complementary channels: a 2-muon (2μ) channel targeting both the HAHM and the minimal dark photon model, and a 4-muon (4μ) channel targeting the HAHM exclusively. Together, these channels are expected to provide the most sensitive ATLAS search for long-lived dark photons in the targeted regions of parameter space.

Chapter 2 reviews the Standard Model, motivations for dark sector extensions, and the theoretical and experimental status of dark photon searches. Chapter 3 describes the Large Hadron Collider, its current operation, and the planned High-Luminosity upgrade. Chapter 4 describes the ATLAS detector and its sub-systems. Chapter 5 describes the reconstruction of physics objects in ATLAS. Chapter 6 presents the dark photon search, covering the analysis strategy, event selection, background estimation, systematic uncertainties, and expected results. Chapter 7 summarizes the results and discusses prospects for future searches at the High-Luminosity LHC. Appendices A, B, C and D describe additional work carried out during the doctorate, with Appendix D presenting published co-authored material.

Chapter 2

Theory Motivations

2.1 Standard Model of Particle Physics

The Standard Model (SM) is an effective field theory (EFT) describing the dynamics of the observed elementary particles. There are three generations of fermions, four types of vector bosons, and a scalar boson. Their interactions are constrained to be invariant under three local (or gauge) transformations. The four vector bosons, also known as gauge bosons, mediate these interactions. They include the photon, gluon, W boson, and Z boson. The fermions are further categorized into quarks and leptons. Quarks and gluons interact through the strong force, mediated by gluons. All the electrically charged particles interact through the electromagnetic force, mediated by photons. All fermions can interact through the weak force, mediated by W and Z bosons. The fermions and gauge bosons acquire mass through their coupling to the Higgs field.

The interactions of the SM particles determine how they can be observed. In fact, most elementary particles cannot be observed directly. Quarks and gluons are confined by the strong interaction, as described by quantum chromodynamics (QCD). Therefore, once produced at a laboratory, they immediately combine into hadrons. The collimated stream of hadrons that is observed by a detector is known as a QCD jet. The weak vector bosons and the Higgs boson have extremely short lifetimes, and will decay immediately into fermions when they are produced. Neutrinos on the other hand rarely interact with ordinary matter. Therefore, they are usually "observed" as missing energy.

2.1.1 The formalism

The SM is built on the formalism of Quantum Field Theory (QFT). In the relativistic quantum regime, fields are defined at every point in spacetime, giving rise to an infinite number of degrees of freedom. Particles are excitations of quantum fields, and the interactions of these fields are local. The SM is understood to be an *effective field theory* (EFT), which means that low-energy observables are insensitive to physics at much higher energy scales. The equations of motion of this EFT are then determined by identifying its matter content and symmetries, and writing down the finite number of interactions which can affect low-energy observables.

Symmetries are transformations that leave the equations of motion of a physical system invariant. Two classes of symmetry are relevant for the SM: *Poincaré symmetries* of Minkowski spacetime and *gauge symmetries*, or gauge redundancies. The former is a combination of Lorentz transformations and spacetime translations. Requiring that particles transform in representations of the Poincaré group implies particles with half-integer spin, where all observed fundamental particles carry spin 0, $\frac{1}{2}$, or 1. The gauge symmetries are local transformations that describe the interactions of the spin-1 fields.

The gauge symmetries of the SM are:

$$\mathscr{G}_{\mathrm{SM}} = SU(3) \times SU(2) \times U(1), \tag{2.1}$$

where $SU(3)$ describes the strong interaction of the gluon, G_μ^A ($A = 1,\ldots,8$), and $SU(2) \times U(1)$ describes the electroweak interactions of W_μ^a ($a = 1,2,3$) and B_μ. The electroweak gauge bosons mix after spontaneous symmetry breaking to give the physical photon, Z, and $W^\pm$ bosons.

All observed fermionic and scalar matter fields are described as representations of $\mathscr{G}_{\mathrm{SM}}$, as listed in Table 1. Each fermion field has three replicas, called *generations*.

The Lagrangian which describes the interactions of the theory can be decomposed into

$$\mathscr{L}_{\mathrm{SM}} = \mathscr{L}_{\mathrm{kin}} + \mathscr{L}_{\mathrm{Yuk}} - V(H). \tag{2.2}$$

Table 1. Particle content of the Standard Model gauge group $SU(3) \times SU(2) \times U(1)$. Fermions come in three generations.

Spin	**Field**	$SU(3)$	$SU(2)$	$U(1)$
0	H	**1**	**2**	$\frac{1}{2}$
$\frac{1}{2}$	$Q = \begin{pmatrix} u \\ d \end{pmatrix}$	**3**	**2**	$+\frac{1}{6}$
	$\bar{u}$	$\bar{\mathbf{3}}$	**1**	$-\frac{2}{3}$
	$\bar{d}$	$\bar{\mathbf{3}}$	**1**	$+\frac{1}{3}$
	$L = \begin{pmatrix} \nu \\ e \end{pmatrix}$	**1**	**2**	$-\frac{1}{2}$
	$\bar{e}$	**1**	**1**	1
1	B_μ	**1**	**1**	0
	W_μ^a	**1**	**3**	0
	G_μ^A	**8**	**1**	0

The kinetic terms are

$$\mathscr{L}_{\text{kin}} = \sum_{\psi=Q,\bar{u},\bar{d},L,\bar{e}} \left[\sum_{f=1}^{3} i\psi_f^\dagger \bar{\sigma}^\mu D_\mu \psi_f \right] + |D_\mu H|^2 - \frac{1}{4} B_{\mu\nu} B^{\mu\nu} - \frac{1}{4} W_{\mu\nu}^a W^{a\mu\nu} - \frac{1}{4} G_{\mu\nu}^A G^{A\mu\nu}, \quad (2.3)$$

where

$$D_\mu = \partial_\mu + i\, g' B_\mu Y + i\, g W_\mu^a T^a + i\, g_s G_\mu^A T^A,$$

$$B_{\mu\nu} \equiv \partial_\mu B_\nu - \partial_\nu B_\mu,$$

$$W_{\mu\nu}^a = \partial_\mu W_\nu^a - \partial_\nu W_\mu^a + g\epsilon^{abc} W_\mu^b W_\nu^c,$$

$$G_{\mu\nu}^A = \partial_\mu G_\nu^A - \partial_\nu G_\mu^A + g_s f^{ABC} G_\mu^B G_\nu^C, \text{ and}$$

ϵ^{abc} and f^{ABC} are the structure constants for $SU(2)$ and $SU(3)$, respectively.

These terms describe the kinetic energy terms and the gauge interactions, where g', g and g_s are the gauge couplings. The fermions additionally have interactions with the Higgs field via

the Yukawa terms:

$$-\mathscr{L}_{\text{Yuk}} = \sum_{i,j=1}^{3} \left(Q_i Y^u_{ij} \bar{u}_j H + Q_i Y^d_{ij} \bar{d}_j \widetilde{H} + L_i Y^e_{ij} \bar{e}_j \widetilde{H} \right) + \text{h.c.} \tag{2.4}$$

The last term is the Higgs potential, defined as

$$V(H) = -\frac{1}{2}\mu^2 |H|^2 + \frac{\lambda}{4}|H|^4, \tag{2.5}$$

whose minimum occurs at $\langle H \rangle = (0,\, v)^T$ for $v = \sqrt{\mu^2/\lambda}$ which is the vacuum expectation value.

The Higgs field, H, has four real degrees of freedom, three of which can be transformed into the longitudinal polarizations of the massive W and Z gauge bosons, after applying a transformation to the unitary gauge. This is the so-called spontaneous electroweak symmetry breaking, where the electroweak symmetry to electromagnetism as $SU(2) \times U(1) \rightarrow U(1)_{\text{EM}}$. The electromagnetic charge can be written in terms of the $U(1)$ hypercharge (Y) in Table 1 and the diagonal weak iso-spin generators (T^3) as $Q = Y + T^3$.

To obtain the gauge boson mass terms, H is expanded around its vacuum expectation value in the unitary gauge as $H = \frac{1}{\sqrt{2}}(0,\, v + h)^T$, where h is a real scalar mass eigenstate with $m_h = \mu$. After replacing the Higgs field with $(0,\, v)$, the $|D_\mu H|^2$ term gives:

$$\frac{v^2}{4}\left[g^2 \left(\left(W^1_\mu\right)^2 + \left(W^2_\mu\right)^2 \right) + \left(g' B_\mu - g W^3_\mu \right)^2 \right] \tag{2.6}$$

The mass eigenstates are identified through the field redefinitions

$$W^\pm_\mu = \frac{1}{\sqrt{2}}\left(W^1_\mu \mp i W^2_\mu \right) \tag{2.7}$$

$$A_\mu = \cos\theta_W B_\mu + \sin\theta_W W^3_\mu \tag{2.8}$$

$$Z_\mu = -\sin\theta_W B_\mu + \cos\theta_W W^3_\mu \tag{2.9}$$

where θ_W is the weak mixing angle, and

$$\sin\theta_W = \frac{g'}{\sqrt{g^2 + g'^2}}, \quad \cos\theta_W = \frac{g}{\sqrt{g^2 + g'^2}} \tag{2.10}$$

The $W^\pm$ and Z bosons acquire masses

$$m_W = \frac{g v}{2}, \quad m_Z = \frac{v}{2}\sqrt{g'^2 + g^2} \tag{2.11}$$

while the photon A_μ remains massless.

Similarly, the Yukawa interactions after electroweak symmetry breaking take the form

$$-\mathscr{L}_{\text{Yuk}} \to \sum_{i,j=1}^{3} \left[u_i (Y^u_{ij}\, v) \bar{u}_j + d_i (Y^d_{ij}\, v) \bar{d}_j + e_i (Y^e_{ij}\, v) \bar{e}_j \right] + \text{h.c.} + \mathscr{O}(h/v) \tag{2.12}$$

The mass matrices $Y_{ij}\, v$ can be diagonalized by performing a unitary transformation on both left- and right-handed quarks and charged leptons. Since independent transformations must be applied to u and d belonging to the same quark doublet Q, the charged current interaction of quarks becomes non-diagonal in the mass basis. The rotation between the up- and down-type mass eigenstates is parametrized by the Cabibbo–Kobayashi–Maskawa (CKM) matrix, V, such that the interaction in the mass basis (denoted by a prime) is of the form

$$\mathscr{L}_{\text{CC}} = u^\dagger_i V_{ij} W^-_\mu \bar{\sigma}^\mu d_j \tag{2.13}$$

The CKM matrix is a unitary matrix with 4 degrees of freedom in the form of 3 mixing angles and 1 CP violating phase.

2.1.2 Total parameters

The SM is fully specified by 18 free parameters[1], summarized as follows:

- Three Gauge couplings: g', g, and g_s.
- Two Higgs potential: μ and λ.
- Nine Fermion masses: m_u, m_c, m_t, m_d, m_s, m_b, m_e, m_μ, and m_τ.
- Four CKM parameters: mixing angles θ_{12}, θ_{13}, θ_{23} and CP-violating phase δ_{CP}.

These parameters have been extracted from a wide range of experimental measurements and found to be mutually consistent. It stands as one of the most precisely tested theories in physics.

[1]There is also a strong CP violating angle in $\theta G^A_{\mu\nu} \tilde{G}^{A\mu\nu}$ leading to a non-zero neutron electric dipole moment at low energies. Absence of an observed neutron electric dipole moment is the *strong CP problem*.

2.2 Physics Beyond the Standard Model

There are many mysteries in fundamental physics that the SM cannot explain. The renormalizable SM, defined without right-handed neutrinos, does not generate neutrino masses. A consistent theory of quantum gravity requires heavy new physics below the Planck scale, which couples to the Higgs boson and would naively generate corrections to its mass up to 16 orders of magnitude larger than the observed value. The universe has a baryon asymmetry, with more matter than antimatter observed, for which there is no dynamical explanation within the SM alone. Evidence exists for dark matter and dark energy, which are not described by the SM at all.

There is no doubt that physics Beyond the Standard Model (BSM) exists. One way to account for some of these mysteries is by introducing new particles and interactions. The rest of this section explores the possibility of having a dark sector where unknown particles exist to provide dark matter candidates and explain some of the outstanding questions in particle physics.

2.2.1 Dark Matter and Dark Sector

One of the most convincing experimental pieces of evidence for BSM physics is the cosmological observations of dark matter (DM) [1]. For example, the galactic rotation curves show that there is additional mass that causes galaxies to rotate faster than expected from only visible matter [2]. By studying the cosmic microwave background and the large-scale structure of the universe, we have been able to determine the relative abundance of DM in the universe [3].

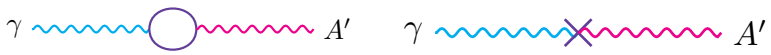
It turns out that SM matter only makes up 5% of the energy content of the universe, while dark matter makes up 27%. The remaining 68% is what is called dark energy, whose existence is inferred by the current accelerated expansion of the universe. The cosmological properties of dark matter have been studied extensively. It interacts weakly or not at all with the SM,

except that it behaves the same as SM matter gravitationally. It is also non-baryonic, cold, and sufficiently stable. Many theoretical models have been proposed to explain the abundance of the DM.

DM may interact with the SM through a mediator particle, often referred to as a portal, without being directly charged under the SM gauge group. In this framework, a *dark sector* of particles exists, and only interact with the SM via the portal. The existence of a dark sector can be motivated not only by providing dark matter candidates but also by addressing other BSM problems.

2.2.2 Dark Photon

One of the simplest portals between SM and the dark sector is through a kinetic mixing of the visible photon (γ) with a vector boson in the dark sector, called the dark photon (A' or γ_D) [4]. The visible photon is the gauge boson of the U(1) symmetry of electromagnetism, while the dark photon is the gauge boson of a broken U(1)$'$ symmetry in the dark sector. The kinetic mixing is illustrated in Figure 1. The mixing strength is denoted ε, and the dark photon mass $m_{A'}$.



(a) (b)

Figure 1. Kinetic mixing of dark photon with the SM photon. Figure (a) shows how the kinetic mixing can be generated by integrating out heavy new physics interacting with both γ and A' at the 1-loop level. Figure (b) shows the kinetic mixing in low energy EFT, below the mass scale of the particles in the loop.

Minimal dark photon model In the minimal model, only an additional $U(1)'$ gauge symmetry is introduced. The relevant kinetic terms of the Lagrangian become

$$\mathscr{L} \supset -\frac{1}{4}F'^{\mu\nu}F'_{\mu\nu} + \frac{1}{2}m^2_{A'}A'^{\mu}A'_{\mu} + \frac{\epsilon}{2\cos\theta_W}B^{\mu\nu}F'_{\mu\nu}, \tag{2.14}$$

where the third term parametrizes the kinetic mixing between the dark photon and the SM B_μ gauge boson. Performing a field redefinition on A'_μ and B_μ to diagonalize the kinetic energy terms results in a dark photon coupling to the electromagnetic current with coupling strength εe:

$$\mathscr{L} \supset -\frac{1}{4}F'_{\mu\nu}F'^{\mu\nu} + \frac{1}{2}m^2_{A'}A'^{\mu}A'_{\mu} + \varepsilon e A'_\mu J^\mu_{\mathrm{EM}} + \mathscr{O}(m^2_{A'}/m^2_Z), \tag{2.15}$$

where the corrections of order $m^2_{A'}/m^2_Z$ are neglected when $m_{A'} \ll m_Z$.

Hidden abelian Higgs model The Hidden Abelian Higgs Model (HAHM) [5] introduces a Higgs portal to the SM, where a dark scalar S charged under $U(1)'$ with gauge coupling g_d mixes with the SM Higgs via $\kappa|S|^2|H|^2$. When S breaks $U(1)'$, it obtains a vacuum expectation value v_S and the dark photon acquires mass $m_{A'} = g_d v_S$. Following electroweak symmetry breaking, the gauge eigenstates (H and S) generate an off-diagonal mass term of the form $4\kappa v_S v_H h's'$, such that transforming to the mass basis (h and s without primes) requires a rotation in field space by a mixing angle approximated by

$$\sin\theta_h \approx \frac{\kappa}{2}\frac{v v_S}{v_S^2\lambda_S - v^2\lambda}, \tag{2.16}$$

where λ and λ_S are the quartic couplings of H and S, respectively. The effective Lagrangian term for the interaction between the SM Higgs and dark photon is

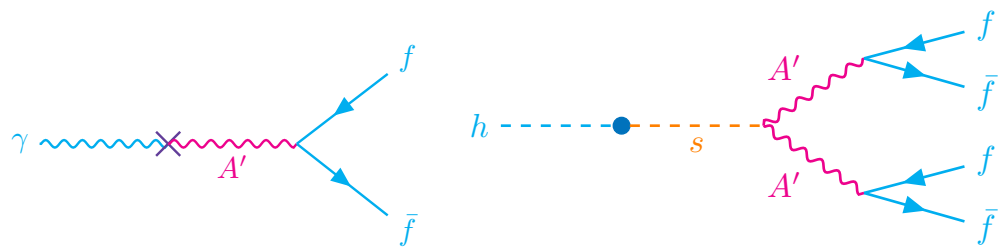
$$\mathscr{L} \supset 2\sin\theta_h \frac{m^2_{A'}}{v_s} h A'_\mu A'^{\mu}. \tag{2.17}$$

Consequently, the partial decay width for $h \to A'A'$ is given by

$$\Gamma(h \to A'A') = \frac{\kappa^2}{32\pi}\frac{v^2}{m_h} + \mathscr{O}\left(\frac{m_s^2}{m_h^2}, \frac{m^2_{A'}}{m_h^2}\right). \tag{2.18}$$

which is independent of the kinetic mixing parameter ε.

For both models, we assume that the dark photon decays only visibly to SM fermions, as shown in Fig. 2. Since the dark photon must decay via the SM photon, its decay is suppressed



(a) (b)

Figure 2. Dark photon production and decay diagrams for two models. Figure (a) shows direct production from the minimal model. Figure (b) shows Higgs-portal production from the HAHM model.

by ε^2. The decay width of A' to SM leptons ℓ is

$$\Gamma\left(A' \to \ell^+\ell^-\right) = \frac{1}{3}\alpha\varepsilon^2 m_{A'}\sqrt{1-\frac{4m_\ell^2}{m_{A'}^2}}\left(1+\frac{2m_\ell^2}{m_{A'}^2}\right) \tag{2.19}$$

for $m_{A'} > 2m_e$. The branching fractions for A' decaying into SM fermions are shown in Figure 3.

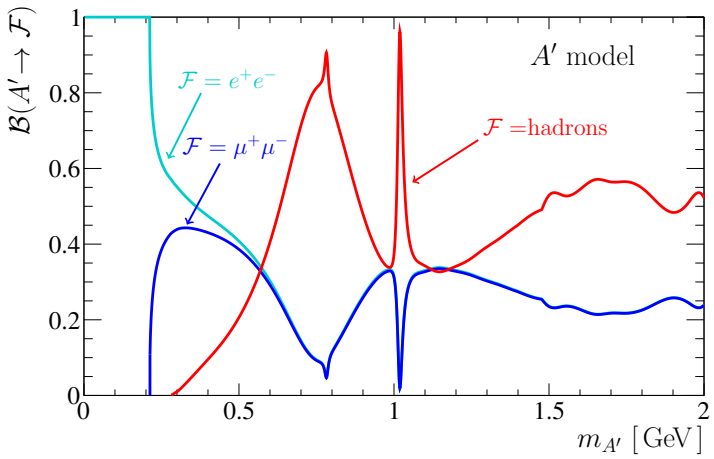


Figure 3. Decay branching fraction of A' from [6].

The lifetime of the dark photon is given by the inverse of the decay width in Equation (2.19). Therefore, in the limit of $\frac{2m_\ell}{m_{A'}} \to 1$, the lifetime of A' corresponds to

$$\tau_{A'} \propto \frac{1}{\varepsilon^2 m_{A'}}. \tag{2.20}$$

The dark photon lifetime and proper decay length as a function of ϵ and $m_{A'}$ are shown in Figure 4. Below twice the electron mass, the dark photon can only decay via the 1-loop suppressed process $A' \to \gamma\gamma\gamma$ and would have a significantly longer lifetime.

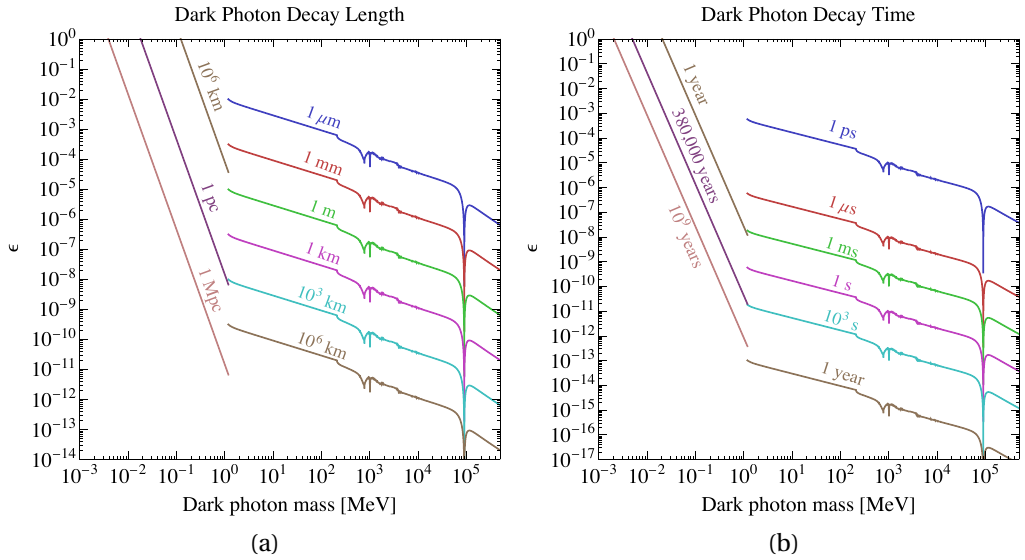


Figure 4. Dark photon decay length (a) and lifetime (b) as a function of kinetic mixing coupling, ϵ, and dark photon mass, $m_{A'}$. Figures taken from [7].

A natural hypothesis that can be made about the kinetic mixing between γ and A' is that it arises from the quantum effects of particles at much higher mass scales. In this case, a heavy particle that is charged under both the SM and dark sector symmetries will generate the mixing at the 1-loop level, as shown in Figure 1a. In grand unified theories, the kinetic mixing often appears at the 2-loop level. The mixing strength at 1- to 2-loop levels corresponds to $10^{-6} < \epsilon < 10^{-2}$, which is the most motivated range for ϵ [4].

Current experimental constraints on the minimal dark photon and HAHM parameter spaces are summarized in Figures 5 and 6, respectively. The current accelerator-based experimental constraints on the minimal model have largely come from prompt dark photon searches, where the kinetic mixing coupling is relatively large. In order to probe regions with smaller couplings, the dark photon lifetime must be taken into consideration. The rest of

the work in the dissertation will explore new experimental signatures to target such regions of parameter space.

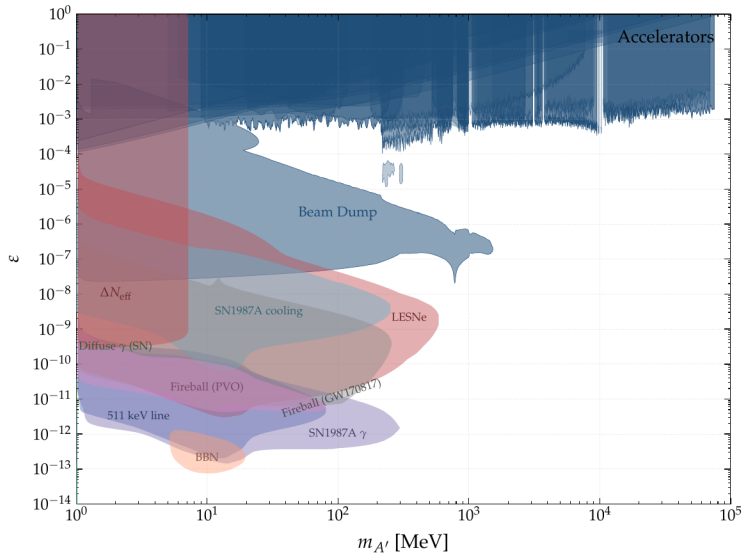
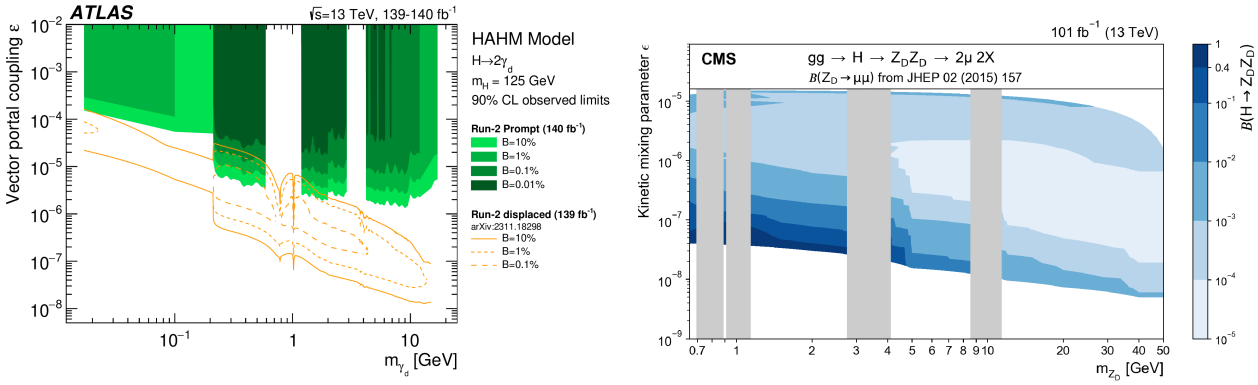


Figure 5. Current constraints on visible dark photons from different experiments [7].



(a) (b)

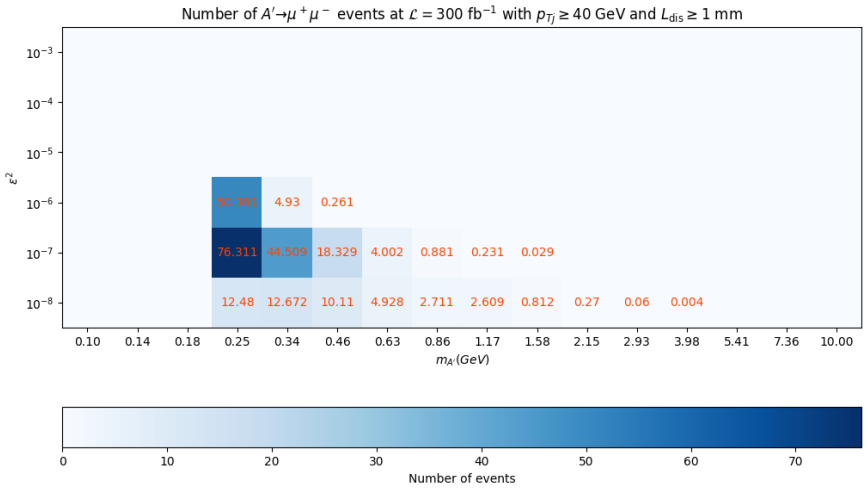
Figure 6. Current constraints on HAHM dark photon from ATLAS [8, 9] (a), and CMS [10] (b).

Sensitivity to the minimal model The expected signal rate from direct A' production via Drell–Yan with an initial-state radiation jet can be estimated by combining the A' production cross section obtained from Monte Carlo simulation, the dark photon branching fraction shown in Figure 3, and the detector acceptance. An integrated luminosity of $\mathscr{L} = 300$ fb^{-1} is assumed,

corresponding to the combined ATLAS Run 2 and partial Run 3 dataset. The jet p_{T} is used as a proxy for muon p_{T}, and the three-dimensional displacement L_{dis} is used to define the detector fiducial volume. Figure 7 shows the estimated signal rate without detector efficiency corrections based on

$$\mathscr{N}_{\text{direct prod.}}\left(A' \to \mu^+\mu^-\right) = \sigma\left(pp \to A'\right)\mathscr{L}\mathscr{B}\left(A' \to \mu^+\mu^-\right)\mathscr{A}, \tag{2.21}$$

where σ is the production cross section of A', $\mathscr{L}$ is the integrated luminosity, $\mathscr{B}$ is the decay branching fraction, and $\mathscr{A}$ is the detector acceptance. These estimates indicate sensitivity in the previously underexplored region $\varepsilon \in [10^{-4}, 10^{-3}]$ and $m_{A'} \in [200, 600]$ MeV with a displaced dimuon signature, a range that falls within the theoretically motivated window $10^{-6} < \varepsilon < 10^{-2}$ discussed above.



(a)

Figure 7. Estimated signal production rate from $A' \to \mu^+\mu^-$ events in the minimal dark photon model. Figure shows the number of events satisfying jet $p_{\mathrm{T}} > 40$ GeV and displacement $L_{\mathrm{dis}} \in [1, 300]$ mm across the $(m_{A'}, \varepsilon)$ parameter space.

Chapter 3

The Large Hadron Collider

Colliders allow us to explore the fundamental constituents of matter through particle scattering. In quantum mechanics, a relativistic particle of energy E probes distances set by the *Compton wavelength*, $\lambda = \hbar c/E$. Higher scattering energies therefore allow us to probe smaller distances, and enable the creation of heavier particles through $E = mc^2$. Scattering experiments have thus provided the primary means to discover and study fundamental particles.

The Large Hadron Collider (LHC) is the largest scattering experiment ever built, accelerating protons along a 27 km underground tunnel to nearly the speed of light. To fully appreciate the LHC, it is worth reviewing how decades of interplay between theory and experiment led to the evolution of collider technology and shaped our understanding of the Standard Model.

This chapter begins with a historical overview of how particle experiments evolved into modern collider experiments over the past century in Section 3.1. Section 3.2 then describes the design and operation of the LHC, followed by Section 3.3, which outlines the planned high-luminosity upgrade.

3.1 Historical prelude

The first hint of subatomic structure came from J.J. Thomson's study of cathode rays in 1897, identifying the electron as the first known fundamental particle. The next breakthrough in scattering experiments was by Ernest Rutherford during 1906–1913, who measured how alpha particles (He^{2+}) scattered off atomic targets in the form of thin metal foils. Rutherford's results suggested the existence of dense nuclei, which were later found to be composed of protons and neutrons. Since then, scattering experiments have evolved to accelerate particles at ever-higher energies to probe ever-smaller constituents of matter.

Before collider experiments were established, many particles were first discovered through nature's own high-energy particle source: cosmic rays. Cosmic rays are high-energy particles produced by galactic and extra-galactic sources such as supernovae. They travel through space to the Earth and interact with the atmosphere to produce showers of secondary particles detectable at the surface. Early cosmic ray experiments marked the beginning of particle physics, leading to a remarkable sequence of discoveries. Victor Hess first identified cosmic rays during a hot-air balloon flight to over 5000 meters in 1912. Carl Anderson discovered the positron using a cloud chamber in 1932 (see Figure 8). Then, the muon, charged pions, and kaons were subsequently identified. These discoveries revealed a rich subatomic world and gave hints of the existence of new particles at higher energies.

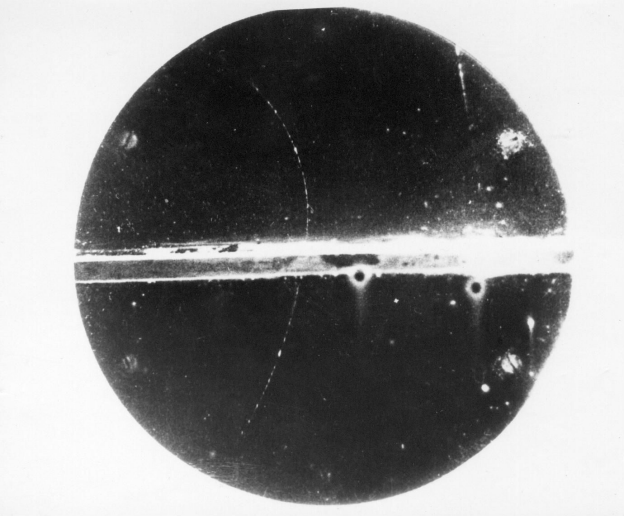


Figure 8. The discovery photograph of the positron in a cloud chamber by Carl Anderson on August 2, 1932 [11]. In this photograph, a cosmic-ray particle moves upward, passes through the lead plate at the center, and curves to the left under a magnetic field. The curvature indicates its charge is positive, and the ionization shows that the particle is not a proton. The positron was the first evidence of anti-matter. Anderson won the Nobel Prize in Physics in 1936 for this discovery.

The first major breakthrough in particle accelerators came with the invention of the cyclotron by Ernest Lawrence in 1930 (see Figure 9). A cyclotron confines particles to a circular path perpendicular to a fixed magnetic field, within two D-shaped electrodes, accelerating

them with an AC voltage each time they pass from one to the other. The frequency at which the particle completes each revolution is determined by

$$\omega = \frac{qB}{m}. \tag{3.1}$$

The AC voltage is therefore tuned to a fixed radio-frequency (RF), and the particle spirals outward at an increasing radius as it gains velocity until it exits. Based on these simple principles of electromagnetism, the cyclotron was able to accelerate protons to 20 MeV by the end of the 1930s.

As particle energies increase, relativistic effects become significant: the particle mass increases with its velocity by a factor of $\gamma = E/mc^2$, reducing the revolution frequency to

$$\omega = \frac{qB}{\gamma m}. \tag{3.2}$$

The synchrocyclotron was developed as an upgrade to the cyclotron by varying the magnetic field in order to keep the frequency constant as the energy increases. Lawrence's first synchrocyclotron was able to accelerate alpha particles to 720 MeV. Subsequently, synchrotrons were developed to accelerate particles at a fixed radius, where accelerating RF cavities were placed along the ring and operated at an increasing frequency as the particle velocity increased.

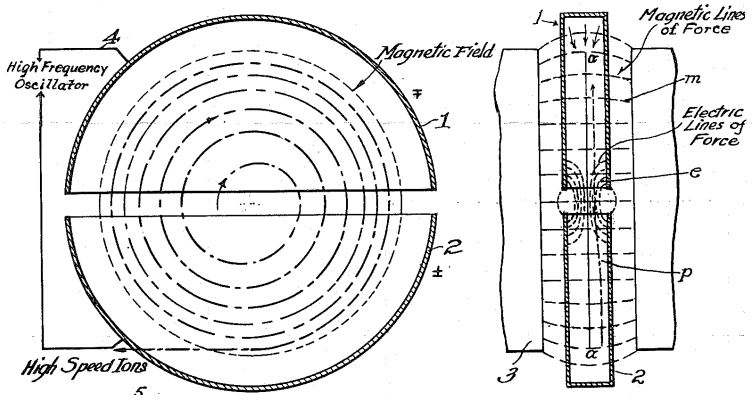


Figure 9. The design of the cyclotron by Ernest Lawrence from his 1934 patent [12].

Until the 1970s, particle accelerators had been designed for fixed-target experiments, where an accelerated beam of charged, stable particles scatters off a stationary target, producing secondary particles to be detected and studied. Fixed-target experiments are limited in center-of-mass energy (E_{cm}) compared to the so-called *collider* experiments where two accelerated beams scatter.

In a fixed-target experiment with an incident beam of energy E_{beam} and a stationary target of mass m_{target}, the center-of-mass energy scales as $E_{\text{cm}} \approx \sqrt{2E_{\text{beam}}m_{\text{target}}}$. In contrast, a collider with two beams each of energy E_{beam} achieves $E_{\text{cm}} \approx 2E_{\text{beam}}$, scaling linearly rather than as the square root of the beam energy. Collider experiments have a clear advantage in reaching higher energy.

Collider experiments, however, pose significant technical challenges. The most critical challenge is achieving sufficient luminosity. The instantaneous luminosity $\mathscr{L}$ is defined as the number of incident particles per second times the number of target particles per unit area, characterizing the rate at which collision events are accumulated. For observing rare processes, the luminosity is therefore just as important as the center-of-mass energy. Compared to fixed-target experiments, where the target naturally provides a high density of nucleons, achieving high luminosity in a collider requires careful cooling, focusing, and steering of both beams.

In the 1970s, collider experiments began operating at ever-higher energies as accelerator technology advanced. Different types of machine, including electron–positron (e^+e^-), hadron (pp, $p\bar{p}$), and electron–proton (ep), were developed and built at laboratories around the world. Figure 10 shows the maximum available center-of-mass energy as a function of time, indicating the transition from fixed-target to collider experiments.

In addition to reaching higher center-of-mass energies, collider experiments opened up new opportunities for resonance searches. In particular, e^+e^- machines can precisely tune the beam energy such that E_{cm} coincides with the rest-mass energy of a new particle, greatly enhancing the annihilation probability and giving rise to a *resonance*. A famous example is

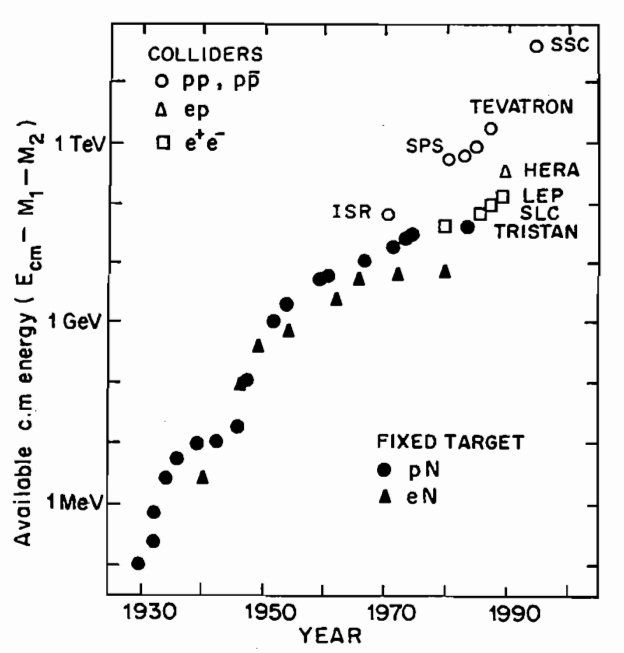


Figure 10. Maximum available center-of-mass energy versus time for different accelerator experiments [13] before the Large Hadron Collider. The Superconducting Super Collider (SSC), planned for construction in Texas, United States, was canceled in 1993 before completion.

the simultaneous discovery of the J/ψ meson at Brookhaven and SLAC in 1974, shown in Figure 11 [14, 15]. More broadly, decades of experiments measuring the e^+e^- annihilation cross section into hadrons, summarized in Figure 12, have helped shape our understanding of the Standard Model.

Before the Large Hadron Collider (LHC), the highest-energy hadron collider ever built, was constructed, the same 27-kilometer tunnel was first used for the Large Electron–Positron (LEP) collider, the largest e^+e^- accelerator ever built. LEP operated between 1989 and 2000, providing some of the most precise measurements of Standard Model parameters, in particular properties of the $W^\pm$ and Z bosons. These results suggested the possible mass of the Higgs boson, as well as motivations for physics beyond the Standard Model. The scientific progress thus far set the stage for the LHC to push the energy frontier further and address open questions about the fundamental laws of nature.

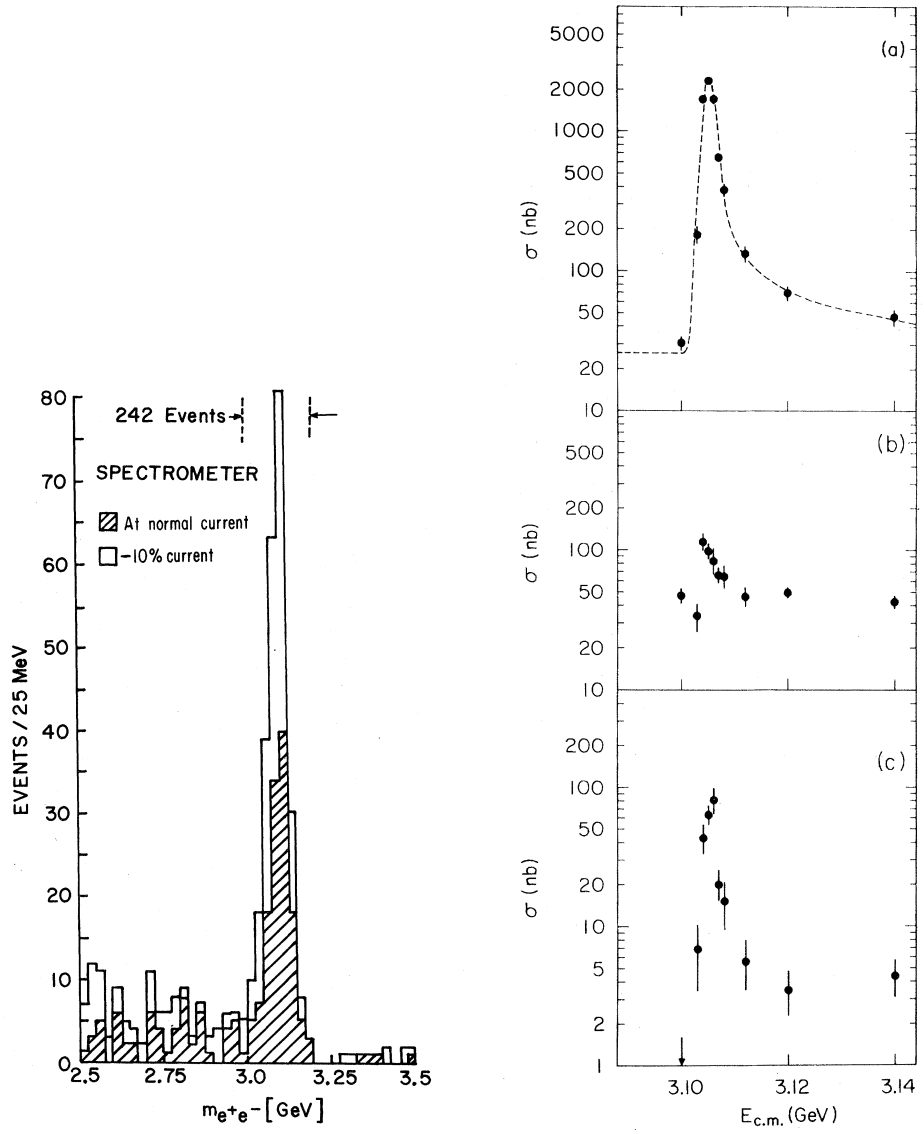


Figure 11. The discovery of the J/ψ meson in November 1974. Figure (a) shows the J particle discovered at Brookhaven National Laboratory led by Samuel Ting [14], observed as a sharp resonance in the invariant mass spectrum of e^+e^- pairs produced in collisions of protons on a beryllium target. Figure (b) shows the ψ particle discovered at SLAC led by Burton Richter [15], observed as a sharp peak in the e^+e^- annihilation cross section at $E_{\text{cm}} \approx 3.1$ GeV for final states of hadrons, e^+e^-, and $\mu^+\mu^-$, from top to bottom, respectively. This simultaneous discovery of the J/ψ is famously referred to as the "November Revolution".

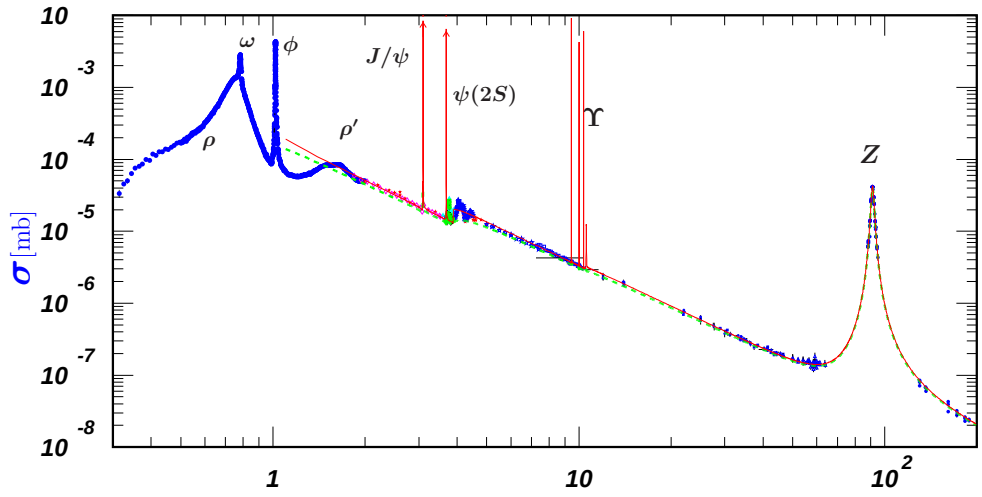


Figure 12. The cross section of $e^+e^- \rightarrow$ hadrons as a function of center-of-mass energy from a combination of experimental measurements. The horizontal axis shows the center-of-mass energy in GeV and the vertical axis shows the cross section σ in mb. Prominent resonances corresponding to the ω, ρ, ϕ, J/ψ, $\psi(2S)$, Υ, and Z bosons are clearly visible. Figure taken from the Particle Data Group [16].

3.2 The Large Hadron Collider

The Large Hadron Collider (LHC) [17] is a superconducting particle accelerator designed primarily to accelerate proton beams to up to 7 TeV in beam energy. It is part of the larger accelerator complex at CERN [18], and the highest hadron center-of-mass energy collider to date.

3.2.1 Basic principles and layout

The LHC was installed in the existing 26.7 km tunnel used for LEP, located 45–170 m underground at the border between France and Switzerland. It is used for both proton–proton and ion–ion collisions.

In a circular accelerator, particles gain energy by repeatedly passing through the high-gradient electric field of radio-frequency (RF) cavities, with the beam path guided by a magnetic field. As the beam energy increases, the magnetic field must increase to maintain the beam

trajectory at the fixed bending radius (ρ):

$$p = qB\rho, \quad E_{\text{beam}} \approx pc, \tag{3.3}$$

or

$$E_{\text{beam}}\ [\text{GeV}] = 0.3(B\rho)\ [\text{Tm}] \tag{3.4}$$

in the ultra-relativistic limit where $pc \gg mc^2$ and setting $q = e$. The bending radius at the LHC is $\rho = 2800$ m. The main bending magnets at the LHC are superconducting, made with niobium-titanium (Nb-Ti) Rutherford cables, achieving a maximum field of 8.3 T when cooled to 1.9 K using liquid helium.

The layout of the LHC is shown in Figure 13. There are two rings, each carrying a beam traveling in opposite directions. The tunnel alternates between eight *arcs* (~ 2460 m each) and *straight sections* (~ 528 m each), connected by dispersion suppressors (~ 147 m each). A straight section together with two dispersion suppressors on each side is called an *insertion region* (IR). Four of the eight IRs have an interaction point (IP) at their center, where the two beams cross, hosting the main LHC experiments: ATLAS at IR1, ALICE at IR2, CMS at IR5, and LHCb at IR8.

The configuration of magnets placed along the acceleration ring is referred to as the magnetic *lattice*. The magnets fall into two categories, *lattice magnets* used in the arcs and *insertion magnets* used in the straight sections. Lattice magnets maintain the beam trajectory and quality, with dipole magnets bending the beam and quadrupoles focusing it. Sextupole, octupole, and decapole magnets provide further corrections to the dipole magnetic fields. Insertion magnets are mainly used for squeezing the beam before collision. At each of the four IPs used for the main LHC experiments, two triplets of quadrupoles are placed on each side, known as the *inner triplets*, which squeeze the beam from 0.2 mm down to 16 μm before collision. The cross-sectional view and the magnetic flux distribution of the main LHC dipole magnets are illustrated in Figure 14.

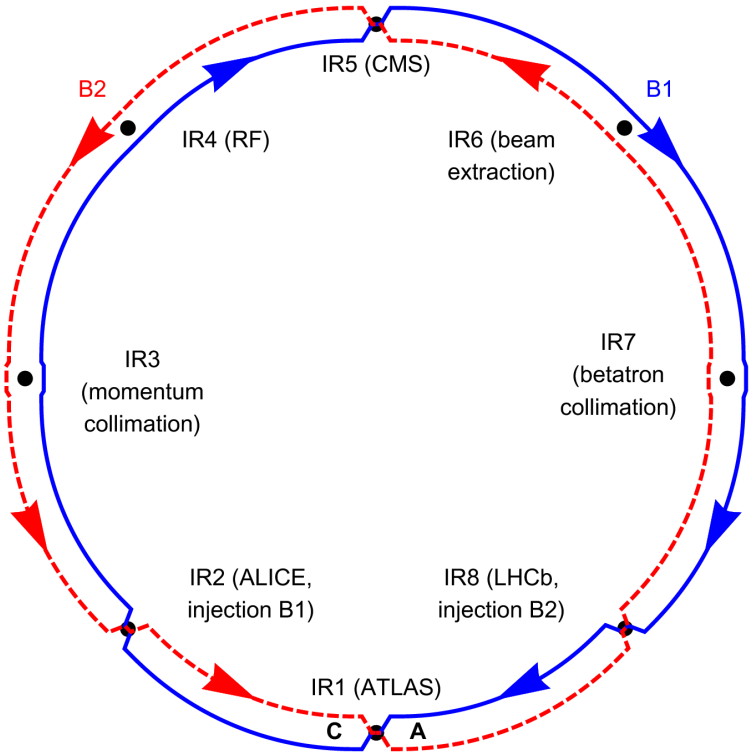
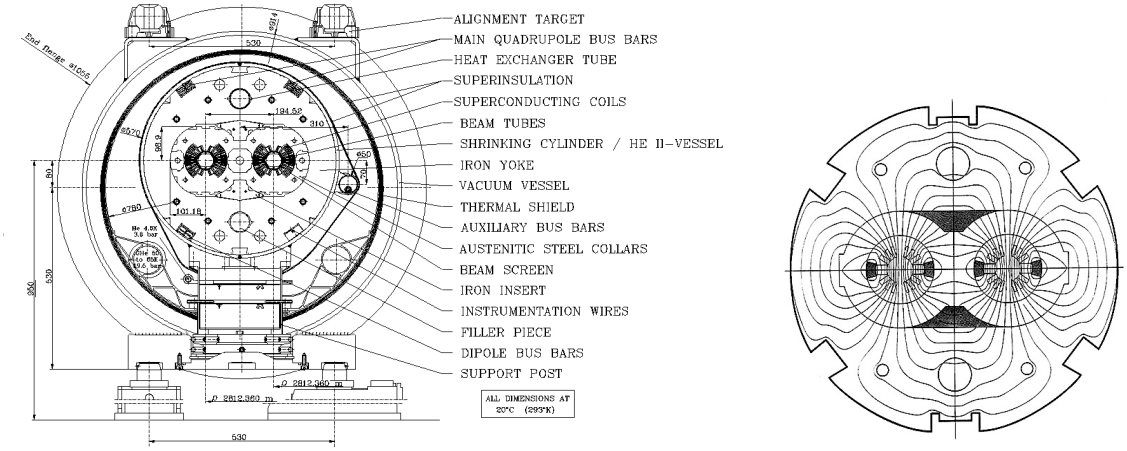


Figure 13. The schematic layout of the LHC. There are eight insertion regions (IRs), each labeled with its main purpose. Figure taken from [19].



(a) (b)

Figure 14. The LHC dipole magnet, showing the cross-sectional view (a) and the magnetic flux distribution (b). Figures taken from the LHC Design Report [20] and CERN Courier [21].

The beam injection systems are located at IR2 and IR8, each for injecting one of the two beams arriving from below the LHC reference plane. Beam cleaning is performed at IR3 and IR7, where collimators absorb particles that have scattered due to large momentum offsets. The 400 MHz superconducting RF cavities are located at Point 4, where they capture, accelerate, and store the beam. The beam abort systems are located at IR6, where both beams are deflected away into absorbers in a separate tunnel.

3.2.2 Operations

The operations of the LHC have gone through several phases as listed in Table 2. As of 2026, it is operating with a beam energy of 6.8 TeV in its last year of operation before the third long shutdown for the High-Luminosity LHC (HL-LHC) upgrade.

Table 2. Key parameters of the LHC during its operational runs and the planned High-Luminosity upgrade (HL-LHC) [16].

	Run 1	Run 2	Run 3	HL-LHC
Data-taking years	2011–12	2015–18	2022–26	2030–41
Max. beam energy (TeV)	4.0	6.5	6.8	7.0
Peak luminosity (10^{33} cm^{-2}s^{-1})	7.73	21.0	26.1	50 (leveled)
Delivered integrated luminosity (fb^{-1})	28.26	156	338	250/year
Time between collisions (ns)	49.90	24.95	24.95	24.95

The beam energy is driven by the synchrotron condition in Equation (3.4). This means that the magnetic field of the superconducting magnets sets the limitation in its energy reach. The designed maximum beam energy at the LHC is 7 TeV, which will be achieved in the Run-4 period of data-taking.

Another important parameter for a collider is its luminosity, as introduced in Section 3.1. The total number of expected physics events of a given process is determined by its production cross section (σ_{exp}) times the *integrated luminosity* over the data-taking time period ($\mathscr{L}_{\text{int}}$):

$$N_{\text{exp}} = \sigma_{\text{exp}} \cdot \mathscr{L}_{\text{int}}, \quad \mathscr{L}_{\text{int}} = \int \mathscr{L}(t)\,dt, \tag{3.5}$$

where $\mathscr{L}(t)$ is the *instantaneous luminosity*, defined as the number of collisions per unit area per second. It can be expressed as

$$\mathscr{L}(t) = f_{\text{coll}} \frac{N_1 N_2}{4\pi\sigma_x^*\sigma_y^*} \mathscr{F}, \tag{3.6}$$

where N_1 and N_2 are the number of particles in each colliding bunch, f_{coll} is the colliding frequency, σ_x^* and σ_y^* are the transverse root-mean-square beam sizes at the IP in the horizontal and vertical directions, respectively, and $\mathscr{F} \sim \mathscr{O}(1)$ is a correction factor due to the inefficient geometric overlap of the beams.

The LHC employs a bunched beam structure, where each beam contains up to thousands of particle bunches, and the bunch crossing frequency is given by $f_{\text{coll}} = n_{\text{b}} f_0$, where n_{b} is the number of bunches per beam and f_0 is the revolution frequency. In Run-3 of the proton–proton data-taking period, the LHC operated with up to $n_{\text{b}} = 2400$ bunches per beam. The typical bunch population is about 10^{11} (100 billion) protons.

The instantaneous luminosity therefore depends on the bunch population, the number of bunches per beam, the beam focusing at the IP, and the bunch crossing frequency. In practice, however, increasing these parameters is constrained by the capabilities of the particle detectors. Higher luminosity leads to a larger number of simultaneous interactions per bunch crossing, known as *pile-up*, which places stringent requirements on the detector readout and trigger systems. The luminosity-weighted distribution of average pile-up for each data-taking period at ATLAS is shown in Figure 15.

3.3 High Luminosity LHC

For the Run 4 data-taking period and beyond, both the accelerator and the main experiments at the LHC are undergoing major upgrades to operate with unprecedented instantaneous and integrated luminosity at the LHC. This new machine configuration is called

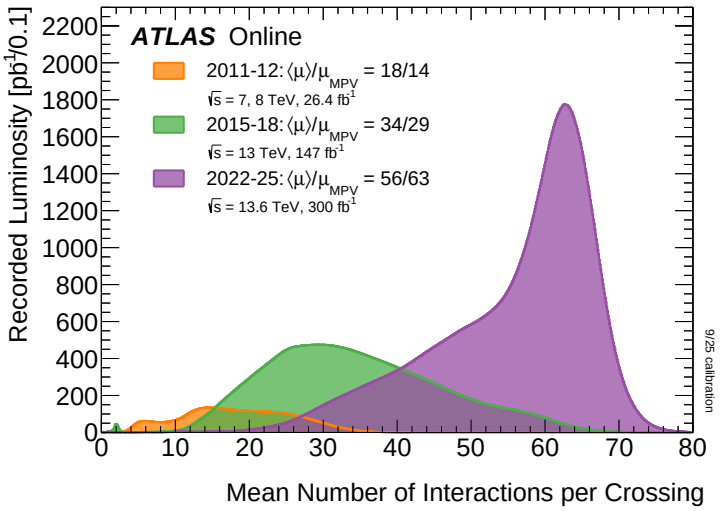


Figure 15. The luminosity-weighted distribution of the mean number of interactions per crossing for proton–proton collisions data recorded by ATLAS during 2011–2025 LHC runs. Figure taken from the ATLAS Luminosity Public Results page [22].

the *High-Luminosity LHC* (HL-LHC) [23]. The instantaneous luminosity will be increased by a factor of five and the integrated luminosity by a factor of ten.

As illustrated in Figure 16, the instantaneous luminosity is closely related to two key beam properties:

- *Emittance* (ε): the phase-space volume of the beam, measuring its spread in position and momentum. In the absence of radiation and dissipative effects, it is a conserved quantity according to Liouville's theorem.

- *Beta-function* $\beta(s)$: the function that modulates the envelope of the transverse oscillations of particles as a function of position s along the ring. It reaches the minimum value β^* at the interaction point.

The transverse beam size at the IP can then be expressed as $\sigma^* = \sqrt{\varepsilon\beta^*}$, and the instantaneous luminosity defined in Equation (3.6) can be written as

$$\mathscr{L} = \gamma \frac{n_b N^2 f_0}{4\pi\varepsilon_n \beta^*} \mathscr{F} \tag{3.7}$$

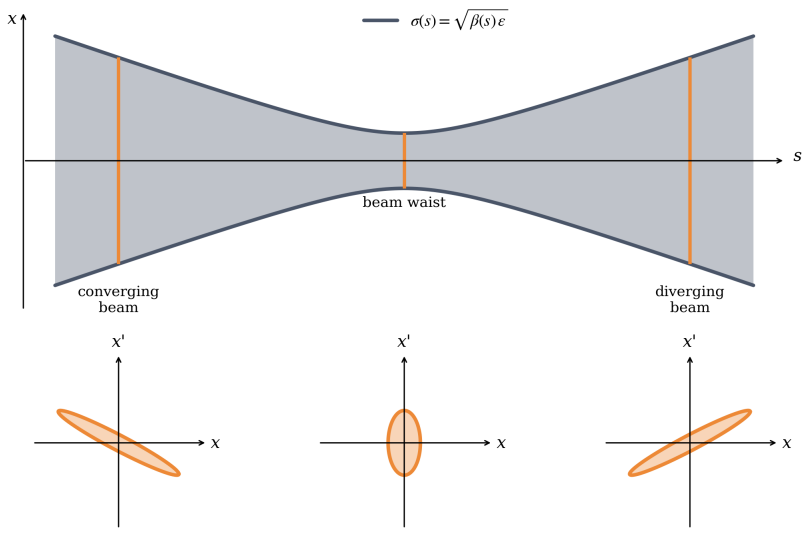


Figure 16. Evolution of the phase-space ellipses at position s along the ring. The transverse beam size is $\sigma(s) = \sqrt{\beta(s)\varepsilon}$, where $\beta(s)$ (emittance) modulates the envelope of the transverse oscillations and ε is the phase-space volume of the beam. Figure adapted from [24].

The primary strategy to increase the luminosity is to reduce the beta-function at the IP, β^*. New inner triplets made with niobium-tin (Nb_3Sn) quadrupoles and new collimators will be installed at IR1 (ATLAS) and IR5 (CMS) to achieve lower β^*. Crab cavities will also be introduced to compensate for the geometric reduction factor $\mathscr{F}$ and provide luminosity leveling. In addition, new 11 T dipoles will be installed in the dispersion suppressor regions to create space for additional collimators. A comparison of the key parameters between the nominal LHC and the HL-LHC is shown in Table 3, and the full accelerator update descriptions can be found in [23].

Table 3. Comparison of key parameters between the nominal LHC and the planned High-Luminosity LHC (HL-LHC) for proton–proton collisions [23].

Parameter	LHC (nominal)	HL-LHC (standard)
Number of bunches	2808	2760
Beam current (A)	0.58	1.1
Transverse emittance ϵ (μm)	3.75	2.50
Minimum β^* (m)	0.55	0.15
Events per crossing $\langle\mu\rangle$	27	131

Chapter 4

The ATLAS Experiment

ATLAS is a general-purpose detector located at one of the four interaction points of the LHC [25]. It observes the physics processes produced from proton–proton collisions, and measures the kinematic properties of these particles through their interactions with materials in the detector. There are four main components of the ATLAS detector: the Inner Detector (ID), the electromagnetic calorimeter (ECAL), the hadronic calorimeter (HCAL), and the Muon Spectrometer (MS). Figure 17 shows a cut-away view of the full detector layout, illustrating the arrangement of each subsystem.

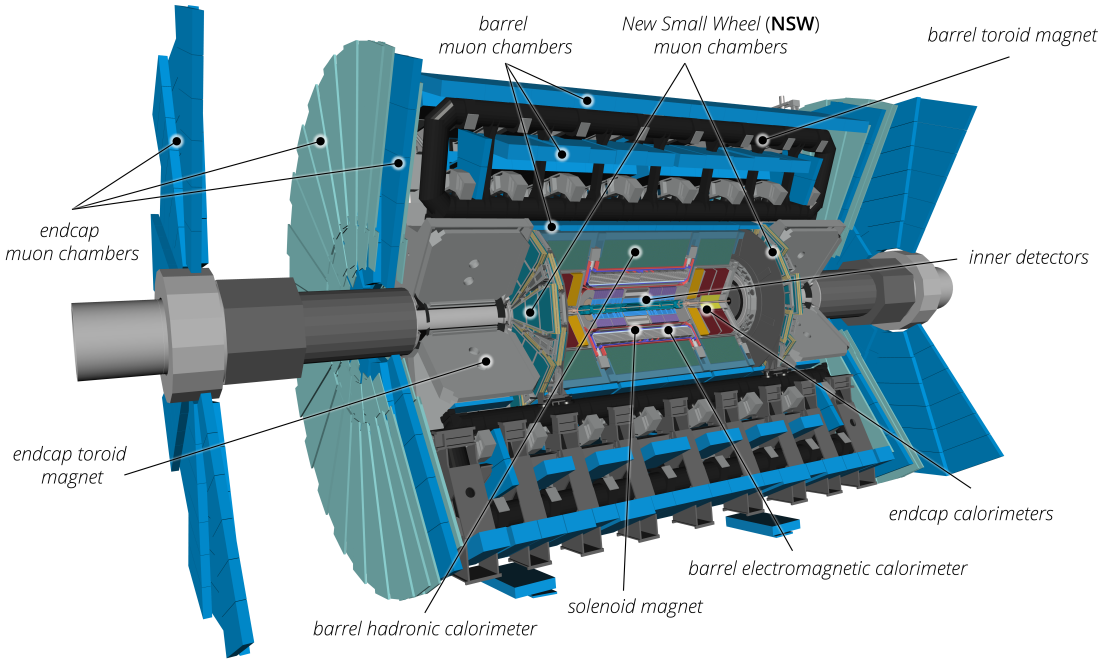


Figure 17. A schematic cut-away view of the ATLAS detector in Run 3. The detector has a cylindrical structure and is 25 m in height and 44 m in length. Figure taken from [26].

4.1 Coordinate System

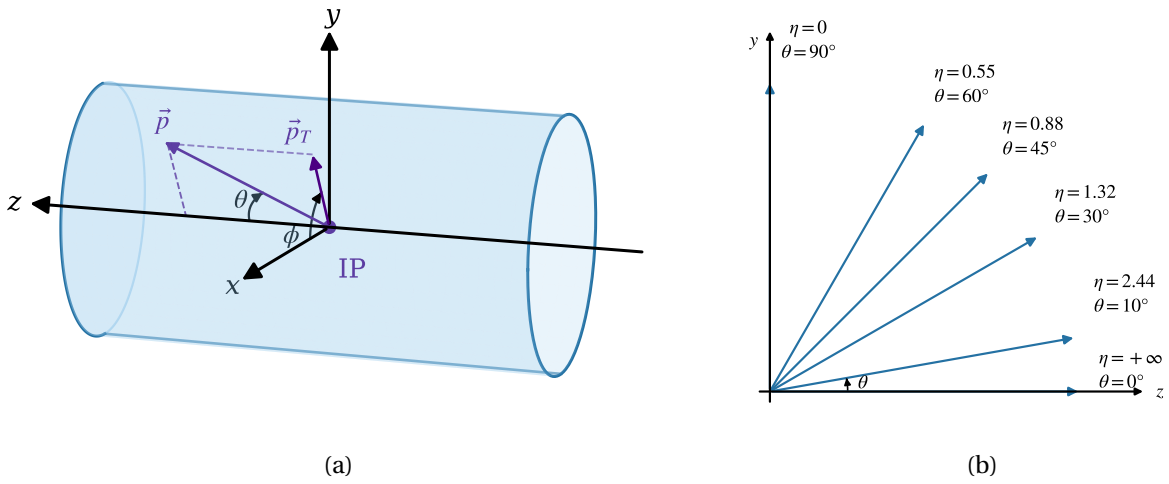
The ATLAS coordinate system, illustrated in Figure 18, is right-handed with its origin at the nominal interaction point (IP). The positive z-axis points along the beam pipe toward the A-side, the positive x-axis points from the IP toward the center of the LHC ring, and the positive y-axis points upward. Cylindrical coordinates (r, ϕ, z) are used, where r is the radial distance from the z-axis in the transverse plane and ϕ is the azimuthal angle measured from the positive x-axis. The *pseudorapidity* is defined in terms of the polar angle θ, measured from the z-axis, as

$$\eta = -\ln\tan\left(\frac{\theta}{2}\right), \tag{4.1}$$

which is equivalent to the *rapidity*,

$$y = \frac{1}{2}\ln\frac{E + p_z}{E - p_z}, \tag{4.2}$$

in the massless limit. Differences in rapidity are invariant under Lorentz boosts along the z-axis, and therefore η is preferable to θ for describing particle kinematics in the massless limit.



(a) (b)

Figure 18. The ATLAS coordinate system showing the cylindrical coordinates (r, ϕ, z) (a), and the definition of pseudorapidity η as a function of polar angle θ (b).

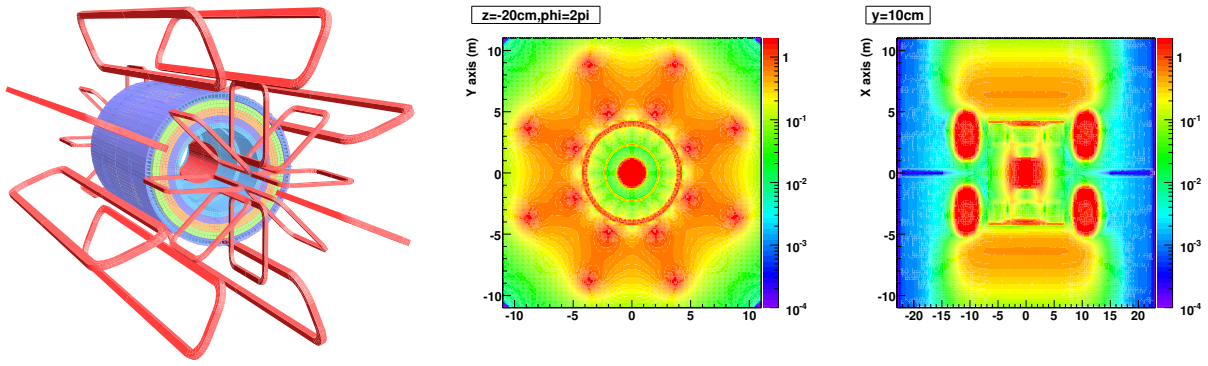
Particle momenta are typically expressed in terms of (p_T, η, ϕ), corresponding to the transverse momentum, pseudorapidity, and azimuthal angle, respectively. A commonly used measure of angular separation between two objects is

$$\Delta R = \sqrt{(\Delta\eta)^2 + (\Delta\phi)^2}, \tag{4.3}$$

which is Lorentz invariant in the massless limit.

4.2 Magnet System

ATLAS has a magnet system consisting of a central solenoid, a barrel toroid, and two end-cap toroids to provide the magnetic field required for momentum measurement of charged particles [27]. The general layout is shown in Figure 19a. There are eight barrel toroid coils surrounding the calorimeter volume with one end-cap toroid on each side, providing a toroidal field for the Muon Spectrometer of ~ 0.5 T in the barrel region and ~ 1 T in the end-cap regions. The solenoid is positioned between the Inner Detector and the calorimeter, providing a 2 T axial field for the Inner Detector. The magnetic field maps in the transverse plane and the r–z plane are shown in Figures 19b and 19c, respectively.



(a) (b) (c)

Figure 19. The geometry of the ATLAS magnetic system. The barrel and end-cap toroids and the central solenoid inside the tile calorimeter are shown in red (a). The magnetic field maps are shown in the transverse plane at $z = -20$ cm (b) and in the r–z plane at $y = 10$ cm (c). Figures taken from [27, 28].

4.3 Detector system

The ID reconstructs the trajectories and momenta of charged particles within a pseudorapidity range of $|\eta| < 2.5$. The calorimeters measure energies of both charged and neutral particles. Particles lose energy in calorimeter materials by producing showers of particles through electromagnetic or nuclear interactions. The energy of electrons and photons is measured in the ECAL, while the energy of hadrons such as pions and kaons is measured in the HCAL. Muons can pass through the calorimeters and reach the MS, which is the outermost detector dedicated to measuring muon momenta.

4.3.1 Inner Detector

Figure 20 shows a schematic cut-away view of the ID. It consists of three sub-systems. The innermost is the silicon Pixel detector, followed by the Semiconductor Tracker (SCT), and the Transition Radiation Tracker (TRT) at the outermost radius. The ID is primarily used for charged particle track reconstruction through ionization, which is detailed in Section 5.1.1.

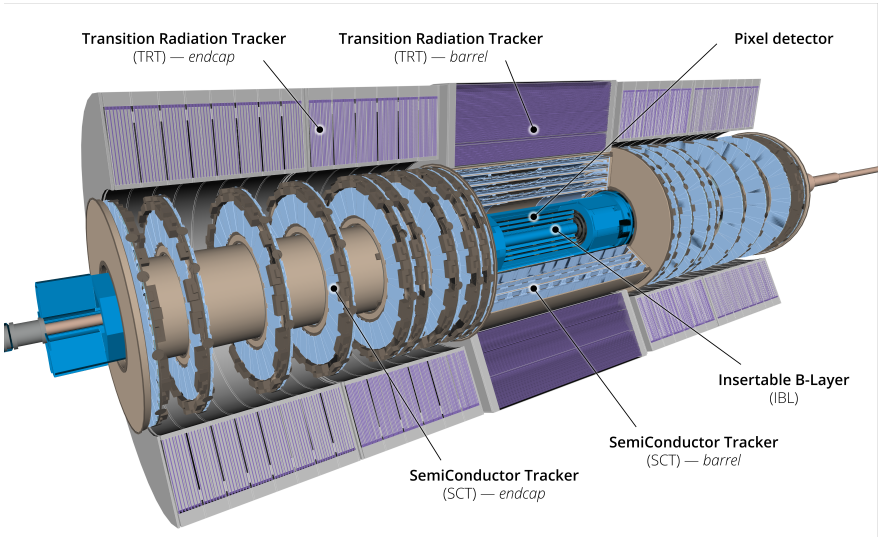
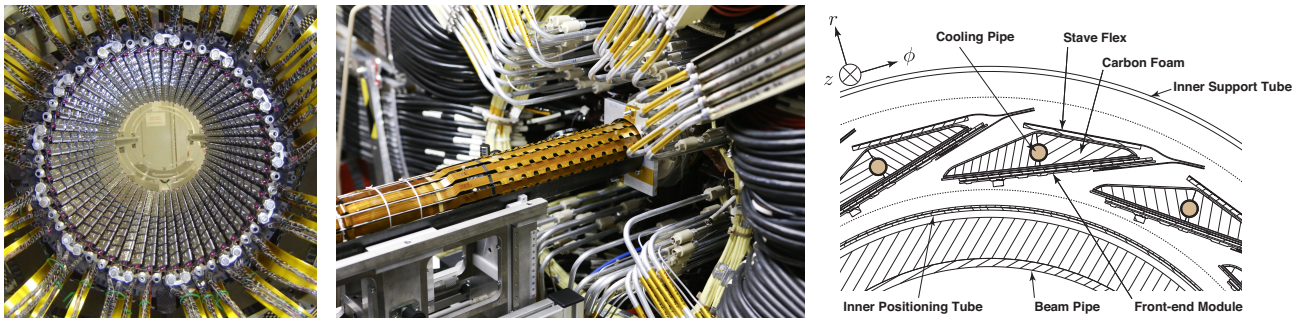


Figure 20. A cut-away view of the ATLAS Inner Detector layout [26]. It provides high-precision measurements of charged-particle trajectories.

The Pixel detector consists of four barrel layers (Insertable B-Layer, B Layer 0, Layer 1, Layer 2) and three end-cap disk layers on each side. Figure 21a shows a view of the barrel Layer 2. The Insertable B-Layer (IBL) [29] was installed after Run 1, directly around a new, narrower beam pipe at r = 33.5 mm to improve radiation hardness and tracking and vertexing performance of the Pixel detector. The beam pipe inner radius was reduced from 29 mm to 23.5 mm to accommodate the IBL. The IBL installation and layout are shown in Figures 21b and 21c.



(a) (b) (c)

Figure 21. Illustrations of the ATLAS Pixel detector system. A picture of the barrel pixel Layer 2 is shown in (a). The installation of the IBL around the beam pipe is shown in (b), and its layout in (c). Figures taken from [27, 30, 31].

The Pixel detector is a semiconductor detector with high-granularity pixel sensors that detect charged particles through ionization. As illustrated in Figure 22a, an ionizing particle passing through the sensor creates electron-hole pairs that drift in the depletion region under a reverse bias voltage. Pixel modules in Layers 0–2 use planar silicon sensors with FE-I3 readout chips, as shown in Figure 22b. The IBL uses planar sensors in the central region and 3D sensors in the forward region. Both sensors are read out by FE-I4 chips, which provide improved radiation tolerance and hit-rate capability required at the innermost layer of the detector [32].

The Pixel detector provides the highest-granularity position measurements with a total of ~ 90 million readout channels. The pixel sensors have a pitch of 50×400 μm^2 for the majority of layers, and 50×250 μm^2 for IBL. The Pixel detector extends out to r = 122.5 mm.

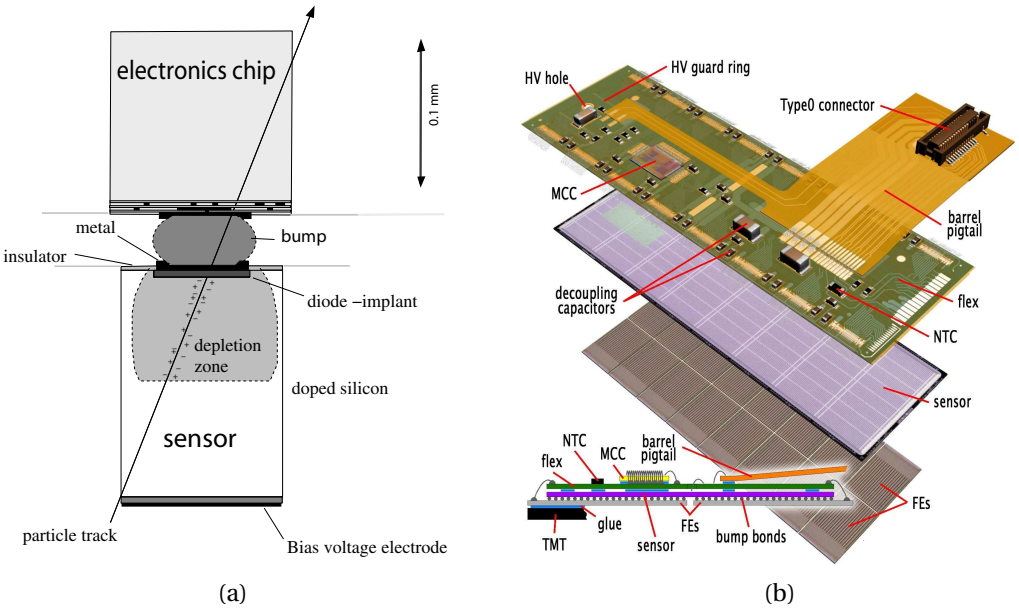
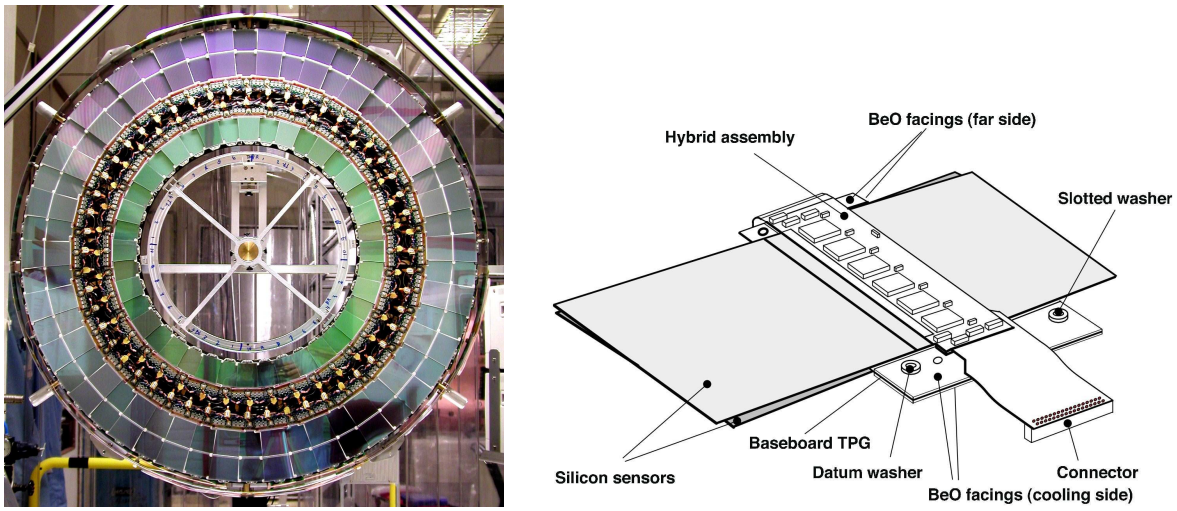


Figure 22. A schematic illustration of a pixel detector module. A schematic illustration of a charged particle passing through a single pixel cell (a). The layout of the ATLAS barrel pixel module is shown in (b). Figures taken from [27, 30, 31].

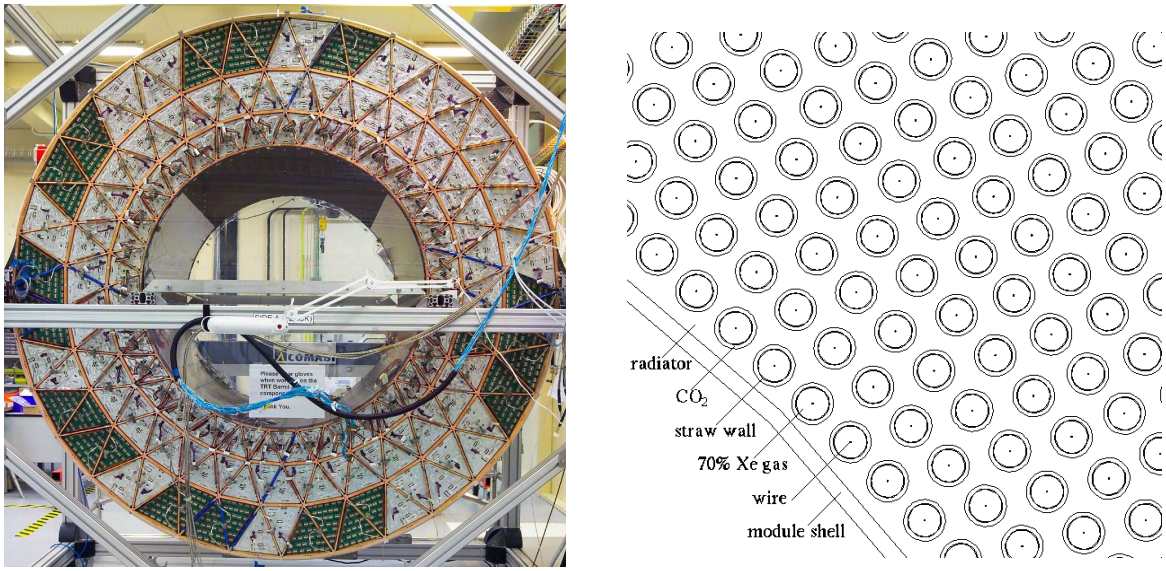
The SCT is a silicon microstrip detector comprising four barrel layers and nine end-cap disk layers on each side. Each module consists of two layers of back-to-back strip sensors with a stereo angle of 40 mrad between the two sides to provide an additional constraint on the position along the strip direction, as illustrated in Figure 23b. A typical strip sensor has a pitch of 80 μm. There are a total of 6 million readout channels. The SCT covers r = 299–514 mm.

The TRT is a gaseous drift tube detector positioned between r = 563 mm and r = 1066 mm, consisting of ~ 3 million gas-filled straw tubes. Each straw tube is 4 mm in diameter and has a 30 μm gold-plated tungsten wire at the center. The structure of the TRT detector is shown in Figure 24. Although the drift tubes were originally filled with a xenon-based gas mixture, an argon-based mixture was adopted from Run 2 onward following observed leakage in a fraction of the tubes in order to reduce cost. In addition to providing track hits, the straws produce transition radiation photons that are used to identify electrons from pions up to a momentum of ~ 100 GeV.



(a) (b)

Figure 23. ATLAS SCT detector modules. A picture of the end-cap disk layer of SCT is shown in (a). The layout of a barrel SCT module is shown in (b). Figures taken from [27].



(a) (b)

Figure 24. ATLAS TRT detector. A picture of the TRT barrel structure is shown in (a). A schematic view of the straw-tube layout is shown in (b). Figures taken from [27, 33].

4.3.2 Calorimeter

The ATLAS calorimeter system is divided into an electromagnetic calorimeter (ECAL) and a hadronic calorimeter (HCAL), as shown in Figure 25, providing energy and position measurements for photons, electrons, and hadrons. Muons and neutrinos traverse the calorimeter with minimal or no interaction, while all other particle species are typically fully absorbed. Object reconstruction from the calorimeter is discussed in Sections 5.3 and 5.4.

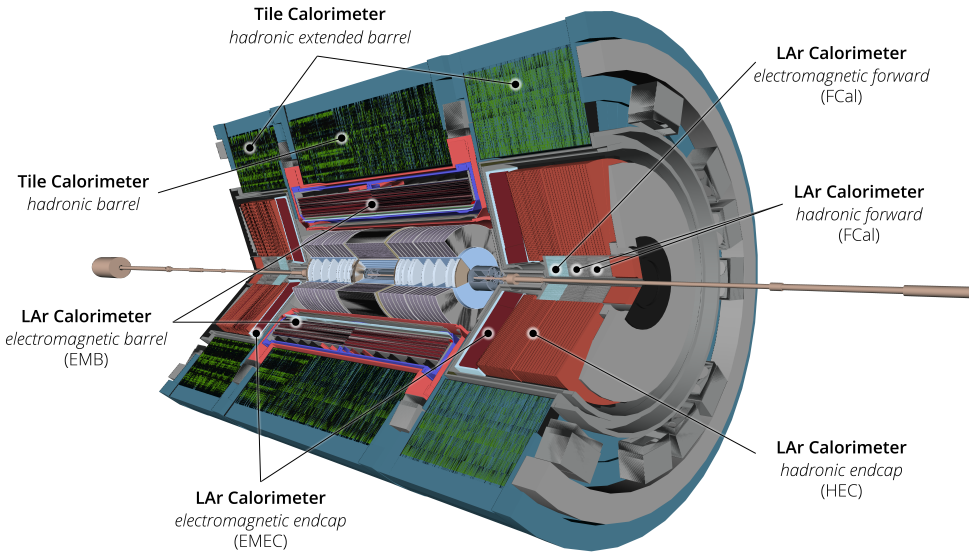


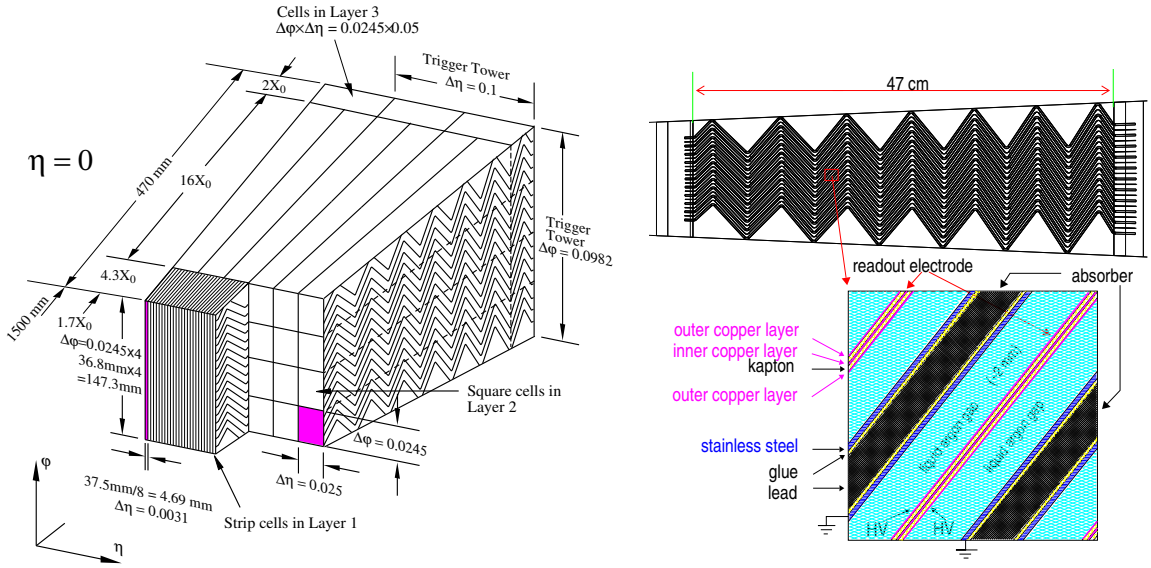
Figure 25. A cut-away view of the ATLAS Calorimeter layout [26].

Both the ECAL and HCAL are sampling calorimeters, consisting of alternating layers of dense absorber material and active medium. The absorber layers initiate electromagnetic or hadronic showers, where cascades of secondary particles are produced. The active medium detects the resulting signal to measure the deposited energy. The ECAL uses lead absorbers with liquid argon (LAr) as the active medium. The barrel HCAL uses steel absorbers with scintillating tiles, while the end-cap HCAL uses LAr-based technology.

The ECAL covers $|\eta| < 1.475$ in the barrel and $1.375 < |\eta| < 3.2$ in the end-caps. The lead absorbers and readout electrodes are arranged in an accordion geometry with active LAr gaps

in between, as shown in Figures 26a and 26b, providing full azimuthal coverage and high granularity in η–ϕ. The ECAL is segmented into three radial layers of varying cell granularity, as illustrated for the barrel in Figure 26a. The first layer contains narrow strip cells with the finest η granularity, optimized for distinguishing $\pi^0 \to \gamma\gamma$ decays from single photons or electrons. The second layer contains the bulk of the shower energy in square cells, while the third layer provides coarser sampling of the shower tail. The total depth is designed to fully contain electromagnetic showers, corresponding to > 22 radiation length (X_0) in the barrel and $> 24\,X_0$ in the end-caps.

The HCAL consists of three subsystems covering complementary η ranges. The tile calorimeter covers $|\eta| < 1.7$, with scintillation signals read out by wavelength-shifting fibers and photomultiplier tubes, as shown in Figure 27. It extends from $r = 2.28$ m to $r = 4.25$ m with a total depth of $\sim 9.7\,\lambda$ at $\eta = 0$. The LAr hadronic end-cap calorimeter (HEC) covers $1.5 < |\eta| < 3.2$ using copper absorbers with 8.5 mm LAr gaps. The forward calorimeter (FCal) covers $3.1 < |\eta| < 4.9$, using copper absorbers in the first module for electromagnetic measurements and tungsten absorbers in the remaining two modules for hadronic measurements, with LAr as the active medium throughout. The HEC and FCal together provide a total depth of $\sim 10\,\lambda$ in the end-cap region.



(a) (b)

Figure 26. Schematic illustrations of the ATLAS electromagnetic calorimeter (ECAL) layout. The three-layer cell structure of the barrel ECAL at $\eta = 0$ is shown in (a). The accordion geometry of the LAr ECAL, with alternating lead absorbers and kapton readout electrodes separated by LAr gaps as the active medium, is shown in (b). Figures taken from [27, 34].

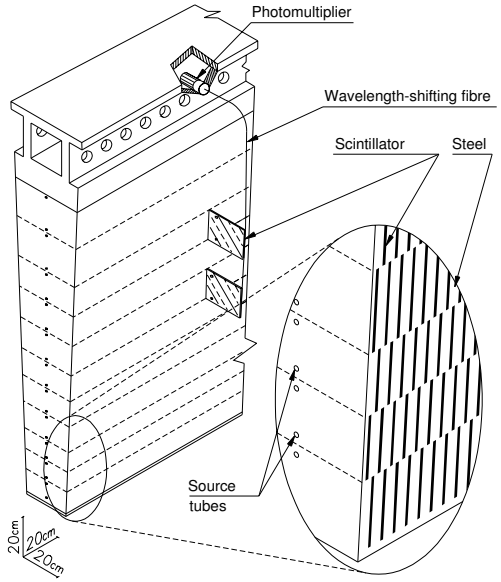


Figure 27. Schematic illustration of a tile calorimeter module, showing the alternating steel absorbers and scintillating tiles with the optical readout. Figure taken from [34].

4.3.3 Muon Spectrometer

The Muon Spectrometer (MS) is the outermost ATLAS detector system, designed to identify muons and measure their momenta. Since muons are the only charged particles that typically traverse the calorimeter without being absorbed, the MS provides both standalone muon tracking and muon identification for triggering, as described in Section 4.4. Figure 28 shows a schematic layout of the MS, where the barrel region ($|\eta| < 1.05$) consists of three cylindrical stations, and the end-cap regions ($1.05 < |\eta| < 2.7$) each consist of three disk layers. There are two types of detector technologies: precision tracking chambers and trigger chambers.

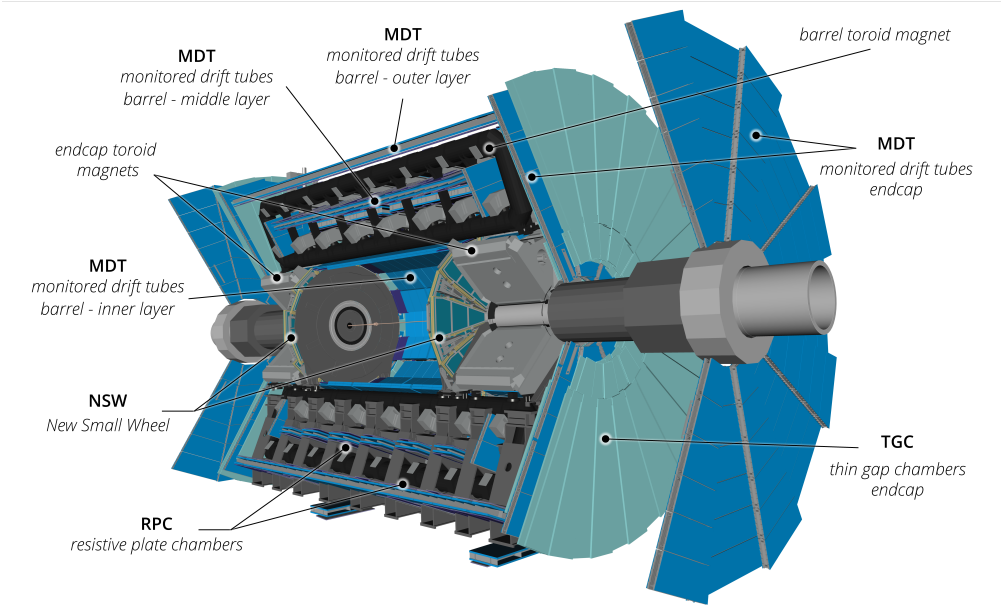


Figure 28. A cross-sectional view of the ATLAS Muon Spectrometer, showing the arrangement of MDT and CSC precision tracking chambers and the RPC and TGC trigger chambers in the barrel and end-cap regions [26].

The precision tracking chambers are Monitored Drift Tubes (MDTs) in all three layers in the barrel region and the two outermost layers in the end-cap regions, and Cathode Strip Chambers (CSCs) in the innermost end-cap layer during Run 1 and Run 2, replaced by the New Small Wheel (NSW) for Run 3. MDTs are aluminum drift tubes of ~3 cm in diameter, filled with 93% Ar and 7% CO_2 and a central tungsten-rhenium (W-Re) wire of 50 μm in diameter at a

3080 V potential, where electrons from ionization are collected, as shown in Figure 29. CSCs are multiwire proportional chambers in which high-voltage wires are strung in parallel through a gas volume bounded by conducting cathode planes. The signal is read out by segmented strips on the cathode planes. CSCs provide high position and timing resolution and higher rate tolerance compared to MDTs, suited for the high-occupancy innermost end-cap layer.

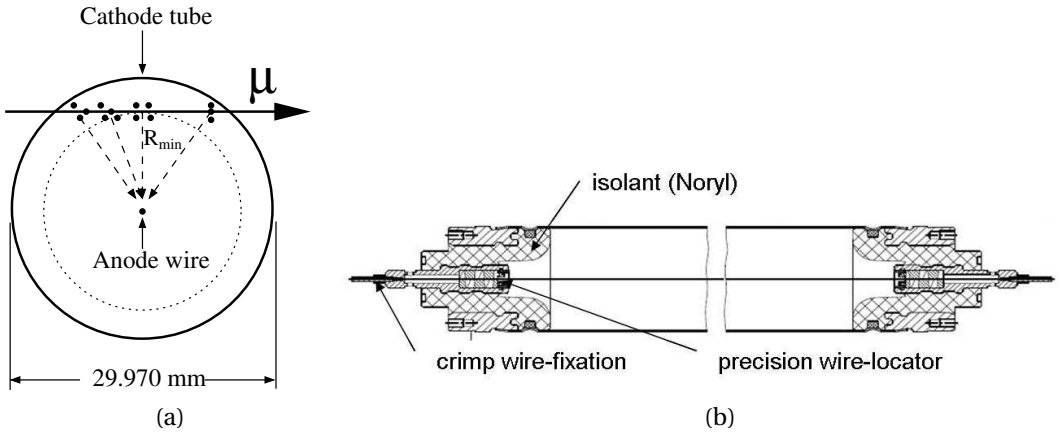
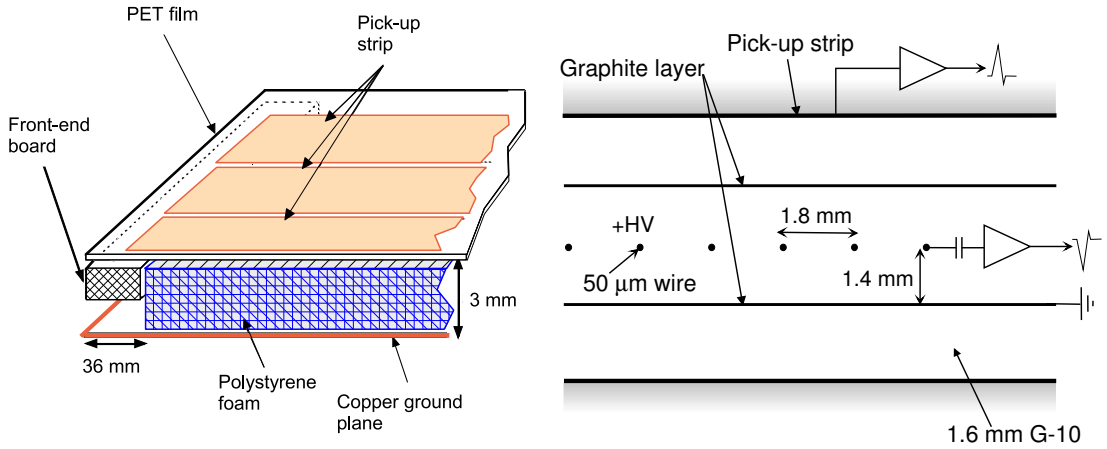


Figure 29. Schematic views of an ATLAS MDT tube. Figure (a) shows a cross-sectional view of an MDT tube. Figure (b) shows a longitudinal view of an MDT tube. Figures taken from [27].

The trigger chambers are Resistive Plate Chambers (RPCs) in the barrel and Thin Gap Chambers (TGCs) in the end-caps, as shown in Figure 30. RPCs consist of two parallel resistive electrode plates with a narrow gas gap in between, where ionizing particles form electron avalanches in the electric field between the plates; the signal is read out by two sets of metallic strips mounted on the outer surfaces of the plates. TGCs are multiwire proportional chambers with a smaller wire-to-wire pitch than wire-to-cathode distance, providing good timing resolution. Together, the RPCs and TGCs provide fast muon identification within the 25 ns LHC bunch crossing interval.

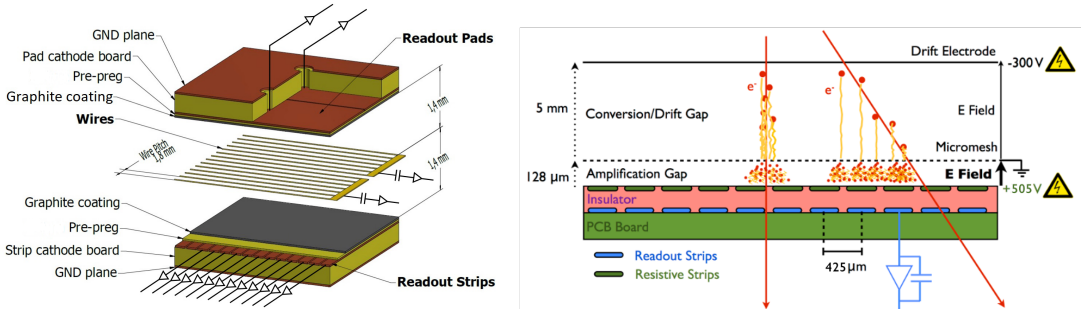
The NSW was installed after Run 2 [26, 35], replacing the innermost end-cap layer and covering $1.0 < |\eta| < 2.7$. It uses two new detector technologies: small-strip Thin Gap Chambers (sTGCs) and Micromegas (MM) detectors, shown in Figure 31, both used for triggering and precision tracking. The NSW upgrade was motivated by the need to suppress background



(a) (b)

Figure 30. Schematic views of the ATLAS MS trigger chambers. Figure (a) shows the layout of an RPC. Figure (b) shows the layout of a TGC. Figures taken from [27].

from low-momentum charged particles produced in hadronic showers in the forward shielding material. Such particles can be misreconstructed as muons, and the rate of these fake muons scales with instantaneous luminosity. The improved p_{T} resolution and timing precision of the NSW help maintain the muon trigger threshold and purity in high-pile-up conditions.



(a) (b)

Figure 31. Schematic views of the two NSW detector technologies. The small-strip Thin Gap Chambers consist of pad and strip cathode boards with sandwiching anode wires at a pitch of 1.8 mm in a 1.4 mm gas gap, shown in (a). The Micromegas detectors consist of a 5 mm conversion/drift gap between the drift electrode and the micromesh, with a 128 μm amplification gap at 505 V above the readout strip PCB, shown in (b). Figures taken from [35].

4.4 Trigger system

The trigger and data acquisition (TDAQ) system is designed to select and record collision events at a rate compatible with the LHC bunch crossing frequency of 40 MHz. ATLAS uses a two-level trigger system consisting of a hardware-based Level-1 (L1) trigger and a software-based high-level trigger (HLT). The L1 trigger processes reduced-granularity calorimeter and Muon Spectrometer inputs to reduce the event rate to ~ 100 kHz and define Regions of Interest (RoI) for HLT. Then, the HLT reconstructs these events using the full detector readout within the RoIs, reducing the rate to ~ 3 kHz for storage. The full layout of the ATLAS TDAQ system is shown in Figure 32.

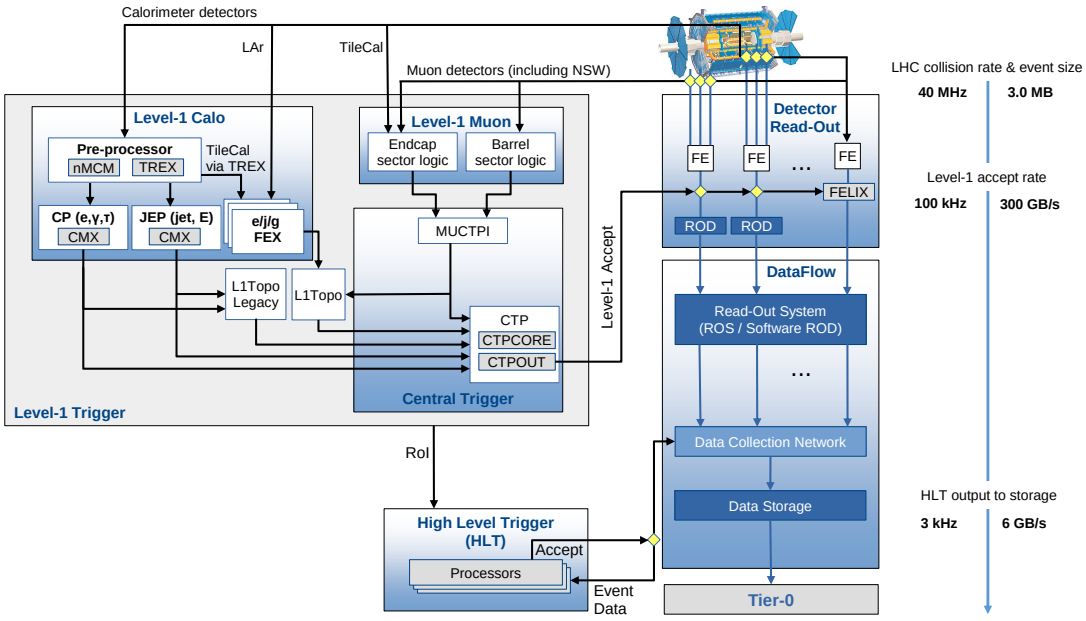


Figure 32. A schematic illustration of the ATLAS trigger and data acquisition system in Run 3. Figure taken from [26].

The L1 trigger includes three sub-systems: the L1 calorimeter trigger (L1Calo), the L1 muon trigger (L1Muon), and the central trigger processor (CTP). L1Calo and L1Muon independently reconstruct trigger-level objects, including electrons, photons, muons, jets, missing transverse energy, and total transverse energy, and forward their results to CTP. The CTP combines these inputs and makes an L1 accept decision based on predefined object thresholds. The maximum time allowed for this decision is the L1 latency, within which detector data are buffered in on-detector memories, while only events accepted by the L1 trigger are saved to the Readout Drivers (RODs). The L1 latency is 2.5 μs per event.

The HLT trigger is a software-based (or CPU-based) system that processes events accepted by L1 and further reduces the output rate by applying more stringent selection requirements. HLT performs more precise object reconstruction using full-granularity detector information in RoIs identified by L1 triggers. Based on the trigger decision, the DAQ system transfers data from detector electronics to offline processing.

4.5 Physics Program

The ATLAS collaboration pursues a broad physics program covering precision measurements of the Standard Model, characterization of the Higgs boson, and searches for BSM physics. The LHC has completed two major data-taking periods and is currently operating in Run 3, scheduled to continue until mid-2026, followed by the High-Luminosity LHC (HL-LHC) upgrade. A summary of ATLAS measurements of key SM production cross sections, compared to theoretical predictions, is shown in Figure 33.

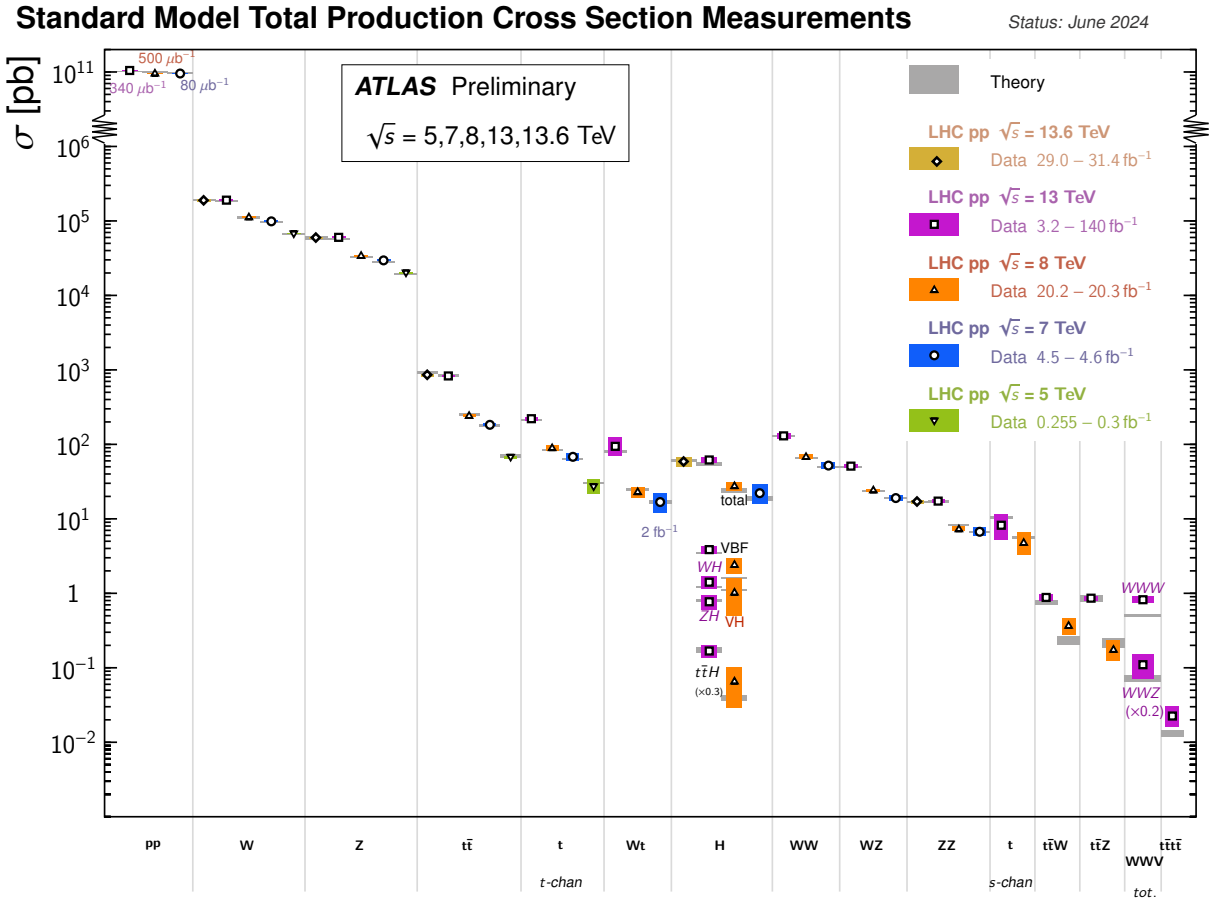


Figure 33. Summary of SM production cross section measurements by ATLAS, corrected for branching fractions and compared to theoretical predictions. The measurements span a wide range of processes and center-of-mass energies. Figure taken from [36].

Chapter 5

Event reconstruction

This chapter describes the reconstruction of collision events from raw ATLAS detector data into physics objects. Event reconstruction begins with charged-particle tracking in the Inner Detector (ID), from which primary and secondary vertices are determined. Higher-level objects, including electrons, photons, muons, jets, and missing transverse momentum, are then built from combinations of ID tracks, calorimeter energy deposits, and Muon Spectrometer segments.

5.1 Track and vertex reconstruction

A fundamental component of event reconstruction in multi-purpose high-energy physics detectors such as ATLAS is the reconstruction of charged-particle trajectories. These trajectories, or tracks, serve both as standalone physics objects and as inputs to the reconstruction of more complex objects, and are used for the identification of their production vertices.

The Inner Detector (ID), described in Section 4.3.1, is the primary tracking system of ATLAS, located closest to the proton–proton interaction point (IP) and operating in a 2 T solenoidal magnetic field. It provides charged-particle track reconstruction and vertex reconstruction within $|\eta| < 2.5$, $|z| < 3512$ mm, and $r < 1150$ mm. A detailed layout of all three subsystems of ID is shown in Figure 34.

The ID is designed to sustain the high hit occupancies and radiation conditions arising from pile-up (see Section 3.2.2). At LHC design luminosity, each bunch crossing produces on average 30–50 simultaneous inelastic interactions, each corresponding to a *primary vertex*. The ID must reconstruct tracks from all pile-up interactions and resolve all associated primary vertices.

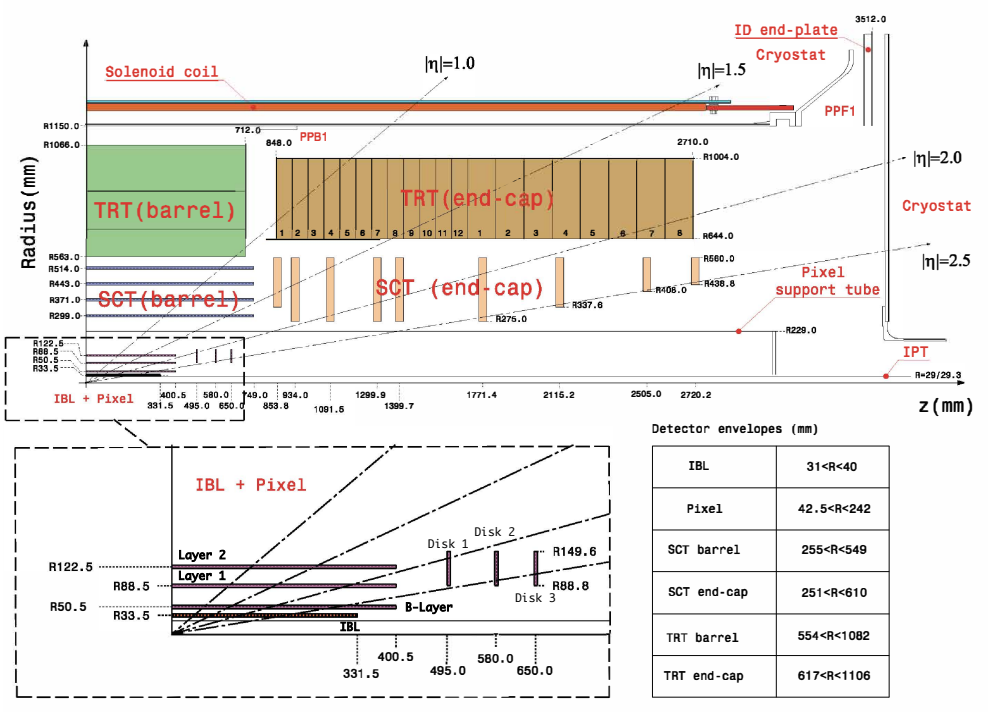


Figure 34. A detailed layout of the ATLAS Inner Detector [37].

5.1.1 Track reconstruction

The reconstruction of charged particle trajectories begins with the measurement of energy deposited in sensor material via ionization processes. In silicon-based detectors, the energy deposition follows the Landau distribution, with the mean energy loss described by the Bethe–Bloch formula [16]:

$$-\left\langle \frac{\mathrm{d}E}{\mathrm{d}x} \right\rangle = Kz^2 \frac{Z}{A} \frac{1}{\beta^2} \left(\frac{1}{2} \ln \frac{2m_\mathrm{e} c^2 \beta^2 \gamma^2 T_\mathrm{max}}{I^2} - \beta^2 + \cdots \right), \tag{5.1}$$

where dE/dx is the energy loss per unit path length, $K = 0.307$ MeV mol^{-1} cm^2 is a constant coefficient for dE/dx, z is the particle charge, Z and A are the atomic number and mass of the absorbing medium (14 and 28 for silicon, respectively), β and γ are the velocity and Lorentz factor of the traversing particle, T_max is the maximum energy transfer to an electron in a single collision, and $I = 137$ eV is the mean excitation energy of silicon.

The collected charge is then digitized in the front-end chip. The chip first converts the charge into a voltage via a preamplifier, which subsequently returns to its baseline value via a feedback circuit that controls the discharge rate. The time during which the preamplifier voltage is above the threshold is called time-over-threshold (ToT), which is proportional to the deposited charge and is used for charge calibration.

As a charged particle passes through the detector, adjacent readout channels above threshold form a cluster at each detector layer. The charge measurements at each channel are used to estimate cluster hit position through *analog clustering*, which performs a charge-weighted computation of the point at which the particle traverses the detector. The simplest algorithm computes the hit position as the weighted mean of the activated channel positions. More advanced neural-network-based methods are used when multiple charged particles share the same cluster, referred to as merged clusters.[Cite tracking in dense env.] .

The charged particle trajectories are reconstructed from the cluster hits by fitting a helical model consistent with the uniform solenoidal magnetic field within the ID volume. Each track is parameterized by five quantities: $(d_0, z_0, \phi_0, \theta_0, q/p)$ evaluated at the *perigee*, the point of closest approach of the helix to the global z direction centered on the beam spot. As illustrated in Figure 35, the five parameters are defined as follows:

- d_0 is the transverse impact parameter, defined as the distance of closest approach in the transverse plane.
- z_0 is the longitudinal impact parameter, defined as the z-coordinate of the perigee.
- ϕ_0 is the azimuthal angle of the momentum vector at the perigee.
- θ_0 is the polar angle of the momentum vector at the perigee.
- q/p is the ratio of charge (q) and momentum (p) of the track, determined from the curvature of the helix in the magnetic field.

The curvature of the helical trajectory is characterized by *sagitta s*, defined as the perpendicular distance from the midpoint of the arc to the chord connecting two hits on the trajectory. The sagitta relates to the transverse momentum p_{T} and particle charge q by:

$$s \propto \frac{L^2 qB}{8p_{\mathrm{T}}}, \tag{5.2}$$

where B is the magnitude of the uniform magnetic field and L is the length of the chord connecting the two hits.

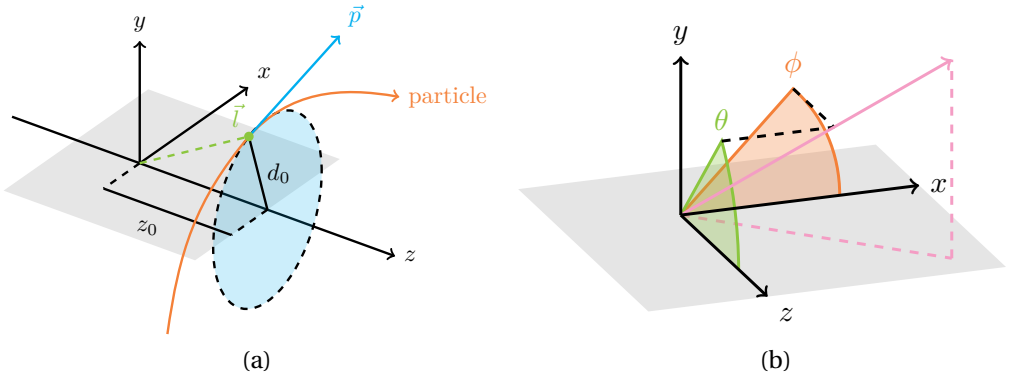


Figure 35. Illustration of the five track parameters (d_0, z_0, ϕ_0, θ, q/p) used to parameterize charged-particle trajectories at the perigee. Figure (a) shows the transverse and longitudinal impact parameters, d_0 and z_0, of the particle track (orange), where $\vec{l}$ is the vector from the reference point to the perigee and $\vec{p}$ is the track momentum vector. Figure (b) illustrates the azimuthal angle ϕ_0 and polar angle θ_0 of the momentum vector $\vec{p}$ at the perigee. Figures taken from [38].

Once space points are formed from detector measurements, tracks are reconstructed following several sequences targeting charged particles from different processes. All sequences involve a common set of procedures: pattern recognition (seed finding and track finding), ambiguity solving, and track fitting. Track seeds are formed from detector hits, or space points, and passed to the track-finding stage if they satisfy minimum quality criteria. During track finding, additional hits are added to the track candidates using the Kalman filter algorithm [39], which propagates the trajectory and performs a sequential update of the track parameters at each new hit. The ambiguity solving algorithm then selects the final set of track candidates and

resolves cases where a cluster is shared between multiple candidates. Finally, a global χ^2 fit is performed to provide the best estimate of the track parameters, accounting for energy loss and scattering in the detector material.

The primary tracking sequence in ATLAS, also referred to as the *inside-out* sequence, is designed to reconstruct tracks originating directly from the interaction point. It begins by forming track seeds from three space points in Pixel and SCT, then extends each track candidate into the TRT by collecting hits along the extrapolated trajectory and performing a global refit including the TRT hits. A schematic illustration of the pattern recognition procedure during the inside-out sequence is shown in Figure 36.

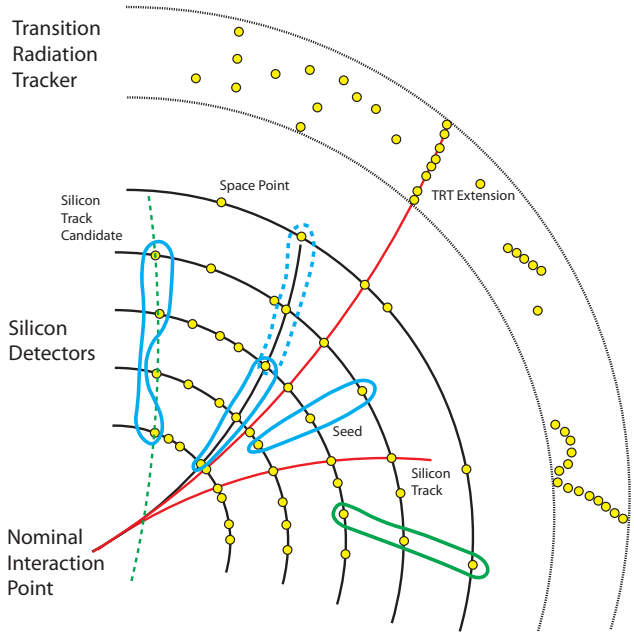


Figure 36. A simplified schematic illustration of inside-out tracking sequence's pattern recognition stage shown in the transverse plane of the Inner Detector. Figures taken from [38].

The secondary sequence, also called the *back-tracking* sequence, uses an outside-in approach. It first forms seeds from TRT hits not assigned to tracks in the primary sequence, and in regions of interest identified by energy deposits in the electromagnetic calorimeter. The back-tracking sequence is designed to reconstruct tracks from particles produced away from the interaction point, such as electrons from photon conversions in the detector material.

Finally, an additional sequence called *Large Radius Tracking* (LRT) is deployed after primary tracking to reconstruct displaced tracks not reconstructed by either the primary or secondary sequences, using only detector hits not assigned to primary tracks. LRT uses a looser set of track selection criteria compared to primary tracking to improve the reconstruction efficiency for displaced tracks. The key differences in track selection criteria between the three sequences are summarized in Table 4 [40]. One of the most noticeable differences is that LRT does not impose an explicit requirement on d_0, whereas primary tracking requires $d_0 < 5$ mm. Consequently, LRT effectively extends track reconstruction to tracks with $d_0 > 5$ mm, typically originating from displaced objects.

Table 4. Key differences in selection criteria used for primary (inside-out sequence only) and LRT tracking [40].

Selection criteria	Primary (*inside-out*)	LRT
max. d_0 [mm]	5	300
max. z_0 [mm]	200	500
min. p_{T} [GeV]	0.5	1
max. $\lvert\eta\rvert$	2.7	3.0
max. silicon holes	2	1
max. double holes	1	0
max. holes gap	2	1
road width [mm]	12	5
seeding	Pixels and SCT	SCT only
max. seeds per Pixel SP	1	–
max. seeds per SCT SP	5	1
Common selection criteria		
min. silicon hits	8	
min. unshared silicon hits	6	
max. track χ^2/DoF	9	
keep all confirmed seeds	true	

5.1.2 Vertex reconstruction

5.1.2.1 Primary vertex reconstruction

The reconstructed tracks are further used to determine the positions of the proton–proton interaction points, known as primary vertices (PVs). In each bunch crossing, the PV that corresponds to the most energetic interaction is referred to as the hard-scatter PV, while the others are referred to as pile-up PVs. The hard-scatter PV is selected to be the one with the largest $\sum p_{\mathrm{T}}^2$ of associated tracks.

The PV reconstruction procedure consists of four steps: seeding, track association, vertex fitting, and vertex selection. First, a subset of high-quality tracks is used to form seed vertices, providing an initial estimate of the PV positions. Nearby tracks are then associated with each seed based on their longitudinal and transverse compatibility with the seed position. The resulting vertex candidates are fitted, assessing the compatibility of each associated track with the vertex position and resolving overlaps between nearby vertices. Vertices satisfying the selection criteria are selected, and their associated tracks are removed from the seed pool. This procedure is repeated iteratively until no further vertices can be reconstructed from the remaining tracks.

The vertex fitting algorithm used for PV reconstruction is the *adaptive multi-vertex fitter* (AMVF) [41]. The vertex position is determined by minimizing a weighted sum of squared distances between the vertex position and the point of closest approach of each associated track. The three-dimensional Euclidean distance (d_i) is normalized by its uncertainty, $\chi_i = d_i/\sigma_{d_i}$. The weight assigned to the i-th track is given by:

$$\omega_i(\chi_i^2) = \frac{e^{-\frac{1}{2}\chi_i^2}}{\sum_j e^{-\frac{1}{2}\chi_j^2} + e^{-\frac{1}{2}\chi_0^2}}, \tag{5.3}$$

where the sum runs over χ_j, the normalized distances of the i-th track to all vertices with which it is associated (including the current vertex), and χ_0 is a cutoff parameter, typically set to three, that suppresses the contribution of incompatible tracks.

The PV reconstruction selects ID tracks that can reliably determine the positions of all PVs, while the remaining tracks in the event can be evaluated for their compatibility with a given vertex via the Track-To-Vertex Association (TTVA) procedure. This allows analyses to incorporate additional tracks using looser selection criteria specific to their physics requirements. The TTVA procedure assesses track-to-vertex compatibility using the transverse and longitudinal impact parameters, d_0 and z_0, as well as the track weight ω from the AMVF vertex fit, defined in Equation (5.3).

5.1.2.2 Secondary vertex reconstruction

Dedicated vertex reconstruction algorithms are used to reconstruct interaction vertices displaced from the proton–proton interaction point. Particles with a macroscopic lifetime may be produced at the IP but decay at a measurable distance from it, giving rise to secondary vertices (SVs). Such particles are referred to as *long-lived*. Many particles within the SM have sizable lifetimes, such as pions, kaons, and B-hadrons. Other processes, such as photon conversions and hadronic interactions with the detector material, can also produce charged particles away from the IP, forming SVs. In BSM scenarios, many hypothetical new particles are predicted to be long-lived, motivating searches for displaced tracks and vertices as signatures of new, long-lived particle decays.

This section describes the secondary vertex reconstruction algorithm used in searches for new long-lived particles, implemented in the Vertex Secondary Inclusive (VSI) package [42]. Other secondary vertex reconstruction algorithms, such as those used for identifying b-jets, are based on similar principles but are tailored to their specific physics needs.

The VSI algorithm begins by forming two-track seed vertices from tracks selected based on a series of quality criteria. Once all two-track seed vertices are formed, they are iteratively merged to form all possible n-track vertices. At this stage, a single track may be associated with multiple DVs. The ambiguity arising from tracks shared among multiple vertices is resolved by

iteratively removing tracks and merging vertices. Finally, tracks not initially selected for seeding are recovered by attaching them to compatible vertices.

VSI allows the track selection criteria and SV property requirements to be configured both for specific analysis needs and to suppress the large combinatorics due to high track multiplicities in collision events. These include requirements on track properties such as p_{T}, d_0, and associated physics objects. Unlike PV reconstruction, which uses only primary tracks, both primary and LRT tracks are typically used as input. However, tracks associated with the PV are excluded to suppress contributions from prompt particle tracks.

To further improve the purity of reconstructed SVs from physical displaced interactions, requirements are applied to suppress SVs formed by accidental crossings of prompt tracks. For example, Figure 37 illustrates an SV formed between the Pixel B-Layer and Pixel Layer-1, where the algorithm checks whether hits are present on the detector layers at smaller radii than the SV position along the track trajectories. The presence of such hits would indicate that the tracks forming the SV do not originate from the same interaction, and SVs failing this criterion are rejected.

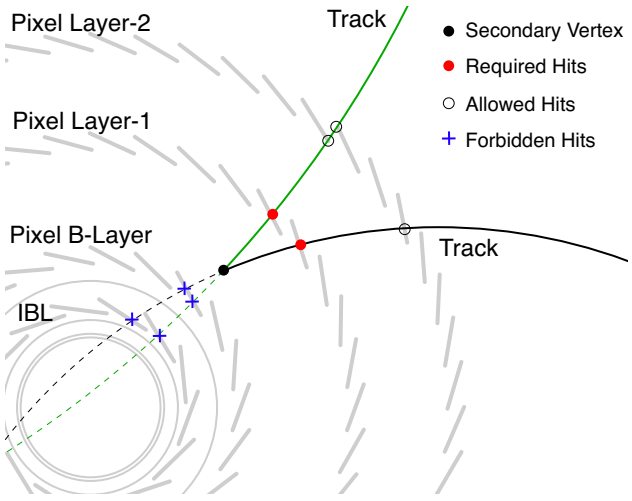


Figure 37. A schematic illustration of secondary vertex reconstruction shown in the transverse plane of the Inner Detector [42].

A more detailed description of the SV reconstruction configuration used for the search for displaced dimuons is discussed in Section 6.3.2.

5.2 Muon reconstruction

Muons are identified and reconstructed from track segments in the Muon Spectrometer (MS), which are combined with reconstructed tracks from the Inner Detector (ID) where available. As discussed in Section 5.1.1, ID tracks are reconstructed following several sequences depending on the origin of the particle. Since muons are the only charged particles that traverse the full detector without depositing all of their energy in the calorimeters, they leave clean signatures in both the ID and MS, and their combination provides precise momentum measurements.

Three types of muon reconstruction, based on the ID track type used, are discussed below. For other muon reconstruction types, please refer to [43].

5.2.1 Standard Combined (CB) Muons

Standard CB muon reconstruction combines ID tracks with MS–only extrapolated tracks to form combined muon objects. Additional selection requirements are applied to the ID tracks, including a transverse momentum of $p_{\mathrm{T}} > 2$ GeV and a transverse impact parameter of $|d_0| <$ 50 mm. Standard CB muons are further required to have at least one pixel hit, three SCT hits, and a total of four silicon hits. Consequently, standard CB muons are not reconstructed beyond the outermost pixel layer, located at a radius of approximately 125 mm. Therefore, it is crucial to include other muon types dedicated for displaced muons.

5.2.2 Large Radius Tracking (LRT) Muons

LRT muons are used to improve signal efficiency by relaxing the Inner Detector (ID) track requirements applied to standard CB muons. The LRT muon reconstruction follows the same

procedure of matching an ID track to a Muon Spectrometer (MS) track, while only ID tracks reconstructed with the LRT algorithm are considered.

LRT muons select ID tracks with $p_{\mathrm{T}} > 2$ GeV and $|d_0| < 10^4$ mm, effectively removing the maximum d_0 requirement. LRT muons also apply modified hit requirements, with no explicit requirement on the number of pixel hits, while requiring at least four SCT hits and a total of four silicon hits. This relaxed selection allows more displaced muons originating beyond the outermost pixel layer to be reconstructed.

5.2.3 Zero Pixel Hit Muons

As discussed previously, the standard *Combined* Muons reconstruction requires the ID track of the muon to have at least one pixel hit. Though this requirement does not apply to LRT muons, most signal muons are reconstructed with primary ID tracks ($d_0 < 5$ mm) due to the boosted nature of the signal. The pixel hit requirement on the ID track limits the ability to reconstruct muons beyond the last pixel layer, resulting in a significant drop in muon reconstruction efficiency for $L_{xy} > 125$ mm.

To overcome this limitation, a custom muon reconstruction has been implemented, called *Zero Pixel Hit (ZPH) Muons* [44]. ZPH muons are reconstructed by matching MS `StandAlone` muons with an ID track satisfying $|\Delta R(\text{ID track}, \text{MS track})| < 0.1$ and consistent charge assignment, $q_{\text{reco, ID track}} = q_{\text{reco, MS track}}$. Most of the ZPH muon properties are taken from the original `StandAlone` muons, including muon p_{T}. However, the muon η and ϕ are taken from the ID track.

5.3 Electron and photon reconstruction

Electrons and photons are reconstructed from energy deposits in the electromagnetic calorimeter. At high energies (above tens of MeV), electrons lose energy primarily through bremsstrahlung, while photons do so via e^+e^- pair production. The electromagnetic

calorimeter provides absorbing layers that induce electromagnetic showers, in which cascades of bremsstrahlung and pair-production processes generate secondary electrons and photons of progressively lower energy.

The electromagnetic showers produced by electrons and photons in the EM calorimeter form clusters of cells with energy deposits. ATLAS uses a topological clustering algorithm [45] to select groups of calorimeter cells with connected energy deposits, called topo-clusters. These topo-clusters are matched to reconstructed ID tracks from electrons and to secondary vertices from photon conversions. The ID tracks matched to the topo-clusters are re-fitted using a Gaussian Sum Filter [46] to account for energy loss due to bremsstrahlung. The matched clusters are then calibrated and used as inputs to form superclusters, used for building the final electron and photon objects [47].

Electron and photon superclusters are built simultaneously, and the corresponding physics objects must then be identified. A supercluster is classified as an electron if it is matched to a good-quality ID track and no good-quality photon conversion vertex is found. It is classified as a photon if it is not matched to any good-quality ID track. When ambiguity arises, both electron and photon objects are built, allowing each analysis to further classify them based on its specific requirements. A schematic illustration of this reconstruction procedure is shown in Figure 38.

5.4 Jet reconstruction

When quarks or gluons are produced in a collision, they immediately undergo a parton shower and hadronize, forming collimated sprays of hadrons that are reconstructed as jets in the detector. The hadrons within these jets interact with the nuclei in the detector material via inelastic strong interactions, producing secondary particles that develop into hadronic showers measured in the calorimeters.

Hadronic showers can begin in the EM calorimeter, but are typically fully contained within the hadronic calorimeter, which is designed to provide sufficient nuclear interaction lengths for

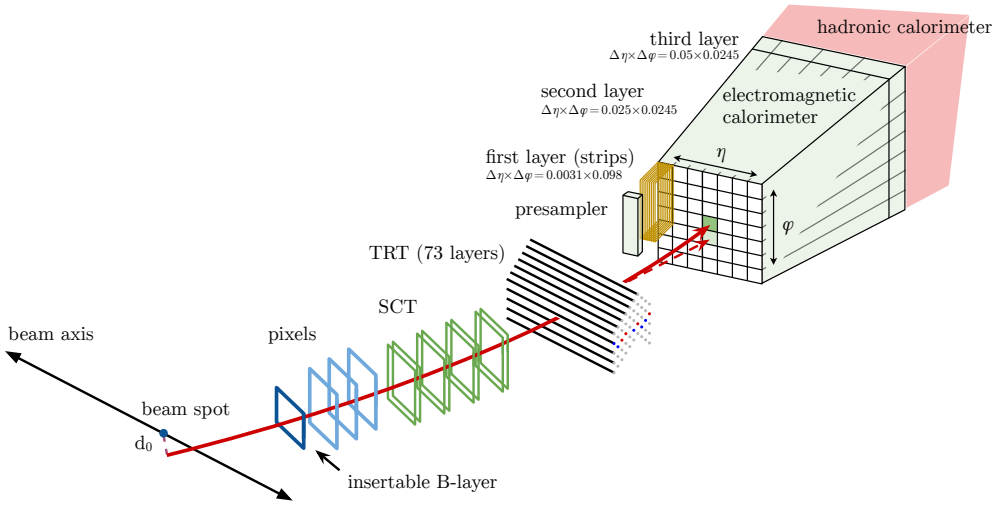
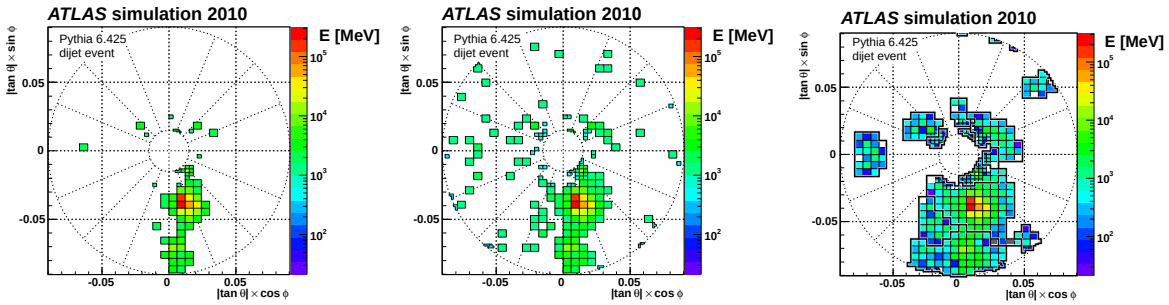


Figure 38. Schematic illustration of an electron traversing the ATLAS tracker and calorimeter. The electron forms an ID track (red), emits a bremsstrahlung photon (dashed red), and deposits energy in clusters of calorimeter cells (green), which are used to build the final supercluster. Figure taken from [48].

full shower development. Some hadrons, such as π^0 mesons, decay into photon pairs ($\pi^0 \to \gamma\gamma$) and initiate electromagnetic sub-showers in the EM calorimeter. As a result, hadronic showers deposit energy across both the EM and hadronic calorimeters, and cells from both are used as inputs to jet reconstruction. Occasionally, shower particles are not fully absorbed and leave tracks in the Muon Spectrometer, referred to as *punch-through*.

Jet reconstruction in ATLAS uses the same topological clustering algorithm described in Section 5.3 to build topo-clusters using both EM and hadronic calorimeter cells. Topo-clusters are formed based on the signal significance of each calorimeter cell, defined as the ratio of the cell energy to its expected noise, suppressing background contributions such as electronic noise and fluctuations due to pile-up. The clustering algorithm seeds clusters from cells with high significance and iteratively adds neighboring cells with lower significance to capture the full shower shape. Figure 39 shows examples of the stages in which topo-clusters are formed in ATLAS simulation [45].



(a) (b) (c)

Figure 39. ATLAS simulation of a dijet event where at least one jet enters the calorimeter and topo-clusters are formed in three different stages: (a) cells with signal significance larger than four are selected to seed topo-clusters, (b) cells with signal significance larger than 2 are added to the topo-clusters, and (c) topo-clusters form by all clustered cells. Figures taken from [45].

The energy of topo-clusters is initially measured at the electromagnetic scale, which corresponds to energy deposits from electromagnetic interactions and is larger than the response to hadrons of the same energy. Local cluster calibrations are therefore applied to correct for the hadronic energy response, energy losses in passive material, and other detector effects. The η and ϕ coordinates of each topo-cluster are also corrected based on the hard-scatter primary vertex.

To further improve the jet energy resolution, a *particle flow* algorithm [49] is applied, combining measurements from the ID and calorimeters into particle flow objects that more precisely represent the constituent particles of a jet. For charged particles, their calorimeter energy deposits are replaced by ID track measurements, which have better momentum resolution at low transverse momenta. These particle flow objects are then clustered using the anti-k_t algorithm [50] at a given radius to form the final jet objects, which is infrared and collinear safe.

5.5 Tau leptons

Tau leptons have a proper decay length of 87.03 μm and decay promptly via the weak interaction into either leptonic or hadronic final states. In leptonic decays ($\sim$ 35% of decays), the tau decays into an electron or muon and two neutrinos, $\tau^{\pm} \to \ell^{\pm} + \nu_{\ell} + \nu_{\tau}$. In hadronic decays ($\sim$ 65% of decays), the tau predominantly decays into one or three charged pions or kaons, along with up to two neutral pions and a tau neutrino, $\tau^{\pm} \to \text{hadrons} + \nu_{\tau}$. ATLAS uses dedicated algorithms to reconstruct hadronically decaying taus by explicitly reconstructing the charged and neutral hadrons in the final state, distinguishing them from electrons and jets [51]. Since leptonically decaying taus are indistinguishable from prompt electrons or muons, their neutrino contributions to the missing transverse energy are typically accounted for at the analysis level.

5.6 Missing transverse momentum

Missing transverse energy typically arises from neutrinos escaping the detector undetected, but can also indicate the presence of new, detector-stable particles predicted by BSM scenarios. A candidate $W^- \to \mu^- \bar{\nu}_{\mu}$ event, where the neutrino manifests as $p_{\mathrm{T}}^{\mathrm{miss}}$, is shown in Figure 40.

Missing transverse momentum ($p_{\mathrm{T}}^{\mathrm{miss}}$) is defined as the negative vector sum of the transverse momenta of all reconstructed and calibrated *hard* physics objects (electrons, photons, muons, tau-leptons, and jets), plus a *soft* term from ID tracks originating from the primary vertex but not associated with any hard object [52]. This is given by

$$\boldsymbol{p}_{\mathrm{T}}^{\mathrm{miss}} = -\underbrace{\sum_{e,\gamma,\mu,\tau,\mathrm{jets}} \boldsymbol{p}_{\mathrm{T}}^{\mathrm{obj}}}_{\text{hard term}} - \underbrace{\sum_{\text{unused tracks}} \boldsymbol{p}_{\mathrm{T}}^{\mathrm{track}}}_{\text{soft term}}. \tag{5.4}$$

An overlap removal procedure is applied to resolve ambiguities where tracks, topo-clusters, and particle flow objects are shared between different reconstructed physics objects, ensuring that $p_{\mathrm{T}}^{\mathrm{miss}}$ is constructed from mutually exclusive detector inputs.

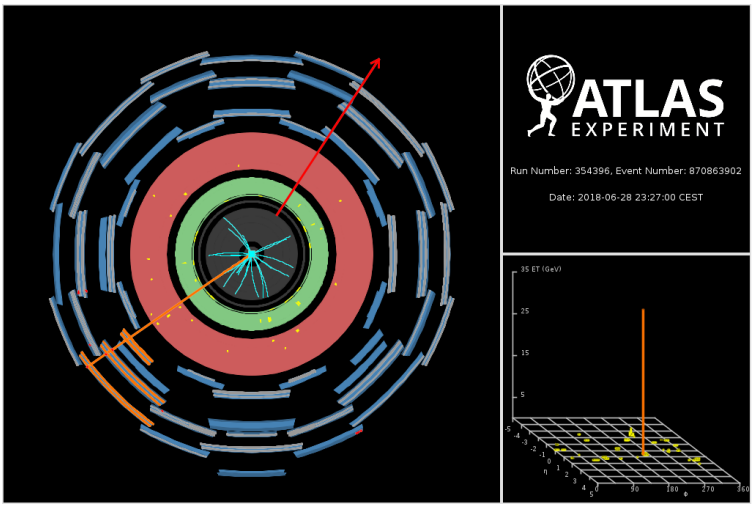


Figure 40. A candidate $W^- \to \mu^- \bar{\nu}_\mu$ event recorded by ATLAS in proton–proton collisions at $\sqrt{s} = 13$ TeV during a low pile-up run. The orange line shows the reconstructed muon track and the red arrow indicates the missing transverse momentum, consistent with an escaping neutrino. Figure taken from [53].

Chapter 6

Search for Low-Mass, Long-Lived Dark Photons to Muons

6.1 Analysis Overview

The searches for dark photons at the LHC conventionally fall into two categories:

– look for a GeV–TeV scale new vector boson in the central region with ATLAS and CMS;

– look for a MeV–GeV scale new vector boson in the forward region with LHCb and FASER.

The difference in mass scale between the detectors is due to the kinematic properties of dark photon production. However, by exploring additional models [54], and using special triggers [55], both ATLAS and CMS were able to add sensitivity to unconventional parts of dark photon parameter space.

To utilize the full potential of general-purpose detectors, it is worth exploring the dark photon lifetime (see Section 2.2.2), given by

$$\tau_{A'} \propto \frac{1}{\varepsilon^2 m_{A'}}, \tag{6.1}$$

in the limit of $\frac{m_\ell}{m_{A'}} \to 1$. For a detectable lifetime, the dark photon must have a sufficiently small coupling and a relatively low mass. These requirements will affect the signal production rate and characteristics, depending on a given model.

This chapter presents a new analysis to search for long-lived dark photons in the mass range 300 MeV–10 GeV, decaying into a pair of muons. This analysis extends the experimental sensitivity to the Hidden Abelian Higgs Model (HAHM) beyond existing results. The analysis also targets the minimal dark photon model with the same signature, while the sensitivity is currently constrained by the level of background. Signal events are triggered using single and

dimuon triggers. Events containing two isolated muons consistent with originating from a common displaced vertex are selected. To ensure the track reconstruction performance, the transverse displacement (L_{xy}) of the dimuon vertex is required to be between 0.5–300 mm.

Since the HAHM model predicts the production of two dark photons, each of which can decay into a pair of muons, the analysis is divided into two orthogonal channels:

- **2-muon channel**: Events with at least one selected displaced dimuon vertex.
- **4-muon channel**: Events with exactly two selected displaced dimuon vertices compatible with the SM-like Higgs boson mass.

This analysis strategy, optimized for a collimated displaced dimuon signature, provides the foundation for a forthcoming interpretation of the minimal model in the 2-muon channel, which will be the first such result at ATLAS.

6.2 Data and Simulation Samples

This analysis uses proton–proton collision data collected by the ATLAS detector during Run 2 (2015–2018) and Run 3 (2022–2024), together with Monte Carlo (MC) simulations for signal modeling and background studies.

6.2.1 Data

This analysis is based on the full ATLAS Run 2 and partial Run 3 collision datasets. Not every recorded luminosity block is suitable for physics analysis: detector conditions vary over time, and individual sub-systems may be temporarily non-operational. Good-Run-Lists (GRLs) are used to select only luminosity blocks in which all relevant ATLAS sub-detectors were fully operational and the data quality was certified as good. The integrated luminosity passing GRL requirements for each data-taking year is listed in Table 5. The total GRL-selected luminosity

is 140.1 fb^{-1} with a 0.83% relative uncertainty for the Run 2 dataset, and 159.8 fb^{-1} with a 1.9% relative uncertainty for the combined 2022–2024 dataset.

Table 5. Center-of-mass energy and integrated luminosity passing GRL quality requirements for each data-taking year used in this analysis.

Year	$\sqrt{s}$ [TeV]	Luminosity [fb^{-1}]
2015	13	3.24
2016	13	33.40
2017	13	44.63
2018	13	58.79
2022	13.6	26.1
2023	13.6	25.8
2024	13.6	107.9

The data samples are processed into a reduced format known as Derived Analysis Object Data (DAOD). For analyses targeting long-lived particle signatures, a dedicated DAOD format, `DAPD_LLP1`, is used, which retains additional track and vertex information together with custom reconstructions of unconventional object topologies not preserved in the standard format.

6.2.2 Simulated Samples

Monte Carlo simulations are used in two complementary roles in this analysis. Signal MC samples are generated across a wide grid of dark photon masses and lifetimes to characterize the expected detector response and optimize the event selection. Background MC samples model the primary processes that can produce displaced dimuon vertices, and are used to understand background composition and develop data-driven estimation strategies. The background yields in the signal regions are ultimately derived from data rather than taken directly from simulation. All MC samples are processed into the `DAOD_LLP1` format.

6.2.2.1 Signal

The signal samples are produced using the Hidden Abelian Higgs Model (HAHM, see Section 2.2.2) [5] implementation with `MadGraph5 aMC@NLO 3.5.3` to simulate the signal process $H \to \gamma_D\gamma_D \to 4f$, where $f = \mu, e, q$. Only $\gamma_D \to \mu^+\mu^-$ decays are considered in this analysis, as the resulting oppositely charged muon pairs provide clean, unambiguous displaced-vertex signatures in the ATLAS detector. Figure 41 shows the 2μ and 4μ signal diagrams.

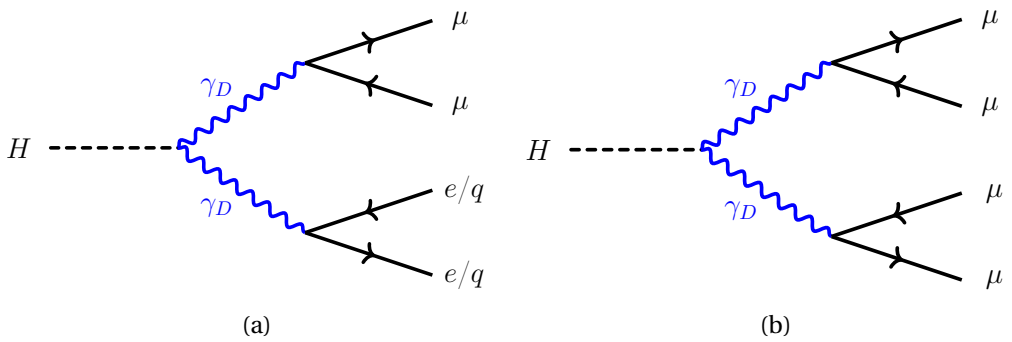


Figure 41. Feynman diagrams of HAHM signal in the 2-muon (a) and 4-muon (b) final states.

The signal MC samples are processed into the `DAOD_LLP1` format with Release `25.0.38`. The signal samples are split into three types of final states: 4μ, $2\mu+2e$, and $2\mu+2q$, though only $\gamma_D \to \mu^+\mu^-$ decays are reconstructed in this search. The full signal grid has 125 points characterized by dark photon mass, lifetime, and decay final state, totaling 15M simulated events. The simulated signal points are shown in Figure 42.

Since the dark photon mass is small relative to the Higgs boson mass, the muons are highly boosted, resulting in sufficiently large transverse momenta to satisfy trigger requirements and relatively small transverse impact parameters d_0 accessible to standard track reconstruction. Figure 43 shows the leading and sub-leading muon p_{T}, d_0, L_{xy}, and dimuon opening angle ΔR for the $m_{\gamma_D} = 1.2$ GeV signal with lifetimes of 0.1, 1.0, and 5.0 mm at reconstruction level.

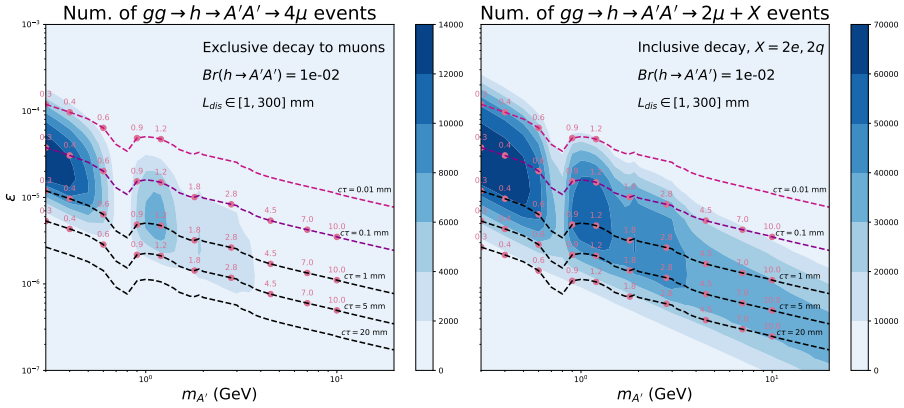


Figure 42. The full signal grid for the HAHM model. The z-axis shows the number of signal events at generator level, assuming a Higgs-to-dark-photon branching ratio $BR(H \to \gamma_D\gamma_D) = 1\%$ and transverse displacement of the dimuon vertex within 1–300 mm.

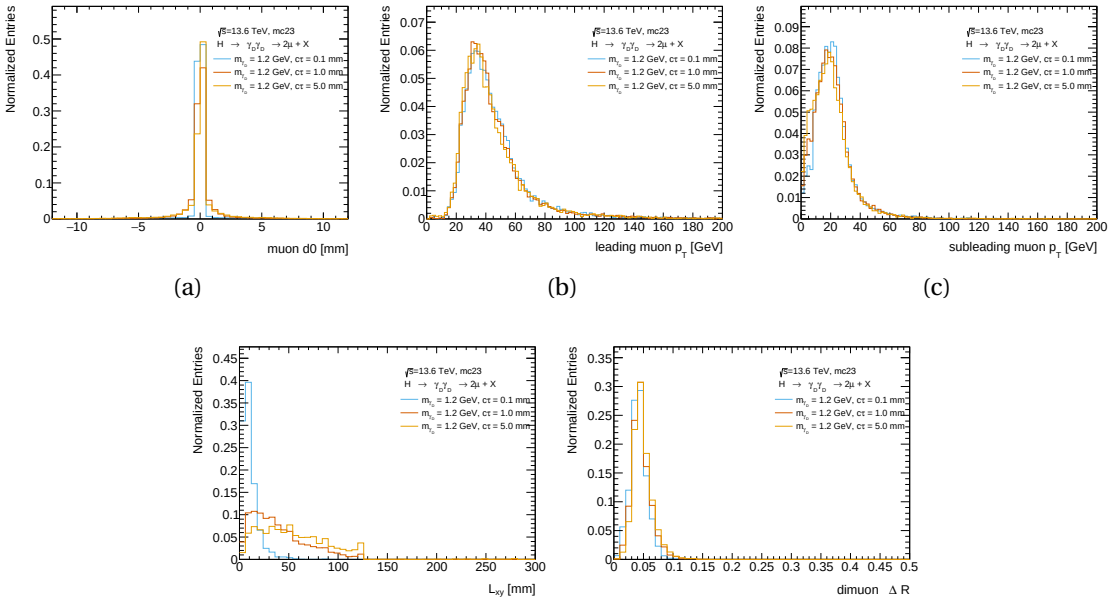


Figure 43. HAHM signal muon kinematic distributions at reconstruction level: muon transverse impact parameter d_0 (a), leading muon p_T (b), subleading muon p_T (c), dimuon vertex L_{xy} (d), and dimuon opening angle ΔR (e), for the 1.2 GeV dark photon with 0.1, 1.0, and 5.0 mm lifetime.

6.2.2.2 Background

The primary sources of background are processes that can produce real or fake displaced dimuon vertices mimicking the signal topology: secondary dimuon vertices from heavy flavor (HF) decays, prompt Drell–Yan production, uncorrelated muon crossings, and material interactions. MC simulations are used to study the characteristics of these background sources, while the expected background rates are estimated using data-driven techniques detailed in Section 6.6. The background samples are summarized in Table 6.

Table 6. Simulated background samples used in this analysis.

Process	Generator	Background source	Key kinematic range
Heavy flavor ($b\bar{b}$, $c\bar{c}$)	Pythia8+EvtGen	HF semileptonic decays	$m_{\mu\mu} \in [0.1, 15]$ GeV
$Z/\gamma^* \to \mu\mu$	Sherpa 2.2.14	Prompt Drell–Yan	$m_{\mu\mu} \in [3.6, 10]$ GeV
Dijet	Pythia8	Hadronic material interactions	jet $p_{\mathrm{T}} \in [60, 800]$ GeV
$W(\tau\nu)$+jets, $Z(\tau\tau)$+jets	Sherpa 2.2.14	Hadronic τ decay	–
$\gamma + Z(\mu\mu)$	Sherpa 2.2.14	Photon conversions	$p_{T,\gamma} > 7$ GeV

Heavy flavor decays of b- and c-hadrons constitute the dominant background source. Semileptonic decays of B and D mesons produce genuine displaced vertices with associated muons, closely mimicking the signal topology. These are modeled using a dedicated Pythia8 sample of hard QCD $b\bar{b}$ and $c\bar{c}$ production, with generator-level muon filters targeting the signal-relevant dimuon mass range.

Prompt dimuon production from low-mass Drell–Yan ($Z/\gamma^* \to \mu^+\mu^-$) is modeled with Sherpa 2.2.14, and is used to study background contributions from prompt vertices at small displacement.

Material interactions with the detector produce displaced vertices through two mechanisms. Hadronic interactions of secondary particles, such as pions and kaons produced in jets or τ decays, can create displaced vertices in detector material, studied using dijet and hadronic τ-decay ($W \to \tau\nu$ and $Z \to \tau\tau$) samples. Photon conversions in detector material layers can

also produce displaced muon pairs, studied using a $Z(\mu\mu)+\gamma$ sample, where dimuon pairs originating from photons converting within the detector material are selected.

6.3 Analysis objects

6.3.1 Muons

The analysis uses `Combined` muons reconstructed by matching track segments in the Muon Spectrometer (MS) with tracks in the Inner Detector (ID). The muon ID tracks are used in the reconstruction of secondary vertices for the displaced dimuon vertex signature. Three complementary muon reconstruction strategies are considered: *standard Combined* (stdCB), *Large Radius Tracking* (LRT), and *Zero Pixel Hit* (ZPH) muons, each targeting a distinct regime of muon displacement, as described in Section 5.2.

An overlap removal procedure is applied to remove duplicate candidates arising when both a standard and an LRT track are matched to the same MS segment. ZPH muons are seeded exclusively from MS-only muon candidates and are therefore orthogonal to the stdCB and LRT collections by construction and no additional overlap removal is required. All three containers are subsequently merged via a dedicated muon container-merging procedure.

Muon candidates are further required to satisfy identification working point (WP) criteria. Standard CB muons are selected with the `Tight` WP, LRT muons with the `Medium` WP, and ZPH muons with a custom set of requirements analogous to the `Tight` WP, as no dedicated WP exists for this reconstruction type. The exact selection criteria are listed in Table 7. Here, the q/p significance is defined as

$$\frac{|(q/p)_{\mathrm{ID}} - (q/p)_{\mathrm{ME}}|}{\sqrt{\sigma^2_{\mathrm{ID}} + \sigma^2_{\mathrm{ME}}}}, \tag{6.2}$$

where ME stands for MS-only extrapolated tracks. The ID/ME/CB momentum imbalance is defined as

$$\rho' = \frac{|p_{\mathrm{T}}^{\mathrm{ID}} - p_{\mathrm{T}}^{\mathrm{ME}}|}{p_{\mathrm{T}}^{\mathrm{CB}}}. \tag{6.3}$$

Table 7. Summary of key muon identification requirements applied in this analysis for each reconstruction category.

Requirement	`Tight` (stdCB)	`Medium` (LRT)	`Tight`-like (ZPH)
$\|\eta\|$	< 2.5	< 2.7	< 2.5
Num. precision layers*	≥ 2	$\geq 2^{\dagger}$	≥ 2
q/p significance	η/p_{T}-dependent	< 7	< 8
Combined fit χ^2/N_{DoF}	< 8	—	—
Momentum imbalance ρ'	η/p_{T}-dependent	—	—
Num. Pixel hits	≥ 1	—	$= 0$

* A precision layer requires ≥ 3 hits in a single MDT/CSC layer, and it is counted as a hole layer if ≥ 3 hits missing.
† For $|\eta| < 0.1$, at least one precision layer with no more than one precision hole is required instead.

Note that LRT muons provide no significant improvement in signal acceptance, as the signal muons are highly boosted and typically have $d_0 < 5$ mm, a regime in which the LRT track reconstruction is less effective than primary track reconstruction. This motivates the alternative reconstruction strategies of ZPH muons that relax the pixel hit requirement. As shown in Figure 44, the addition of ZPH muons increases signal acceptance by a factor of four in the $L_{xy} > 120$ mm region.

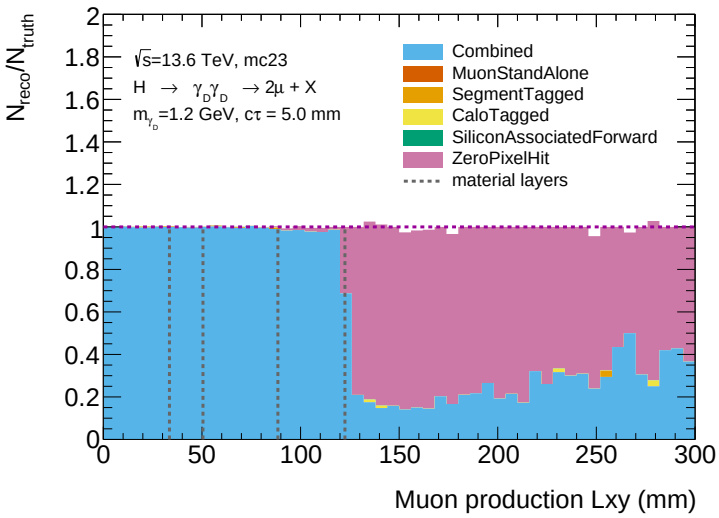


Figure 44. Muon reconstruction efficiency as a function of L_{xy} for different muon types. Events are selected to satisfy muon trigger requirements.

6.3.2 Secondary vertex reconstruction with muons

As discussed in Section 5.1.2.2, displaced vertices (DVs) within the ID are reconstructed with the Vertex Secondary Inclusive (VSI) algorithm. Several VSI vertex collections, corresponding to different configurations of the algorithm, are available for different analysis purposes. For this search, a new collection, `VSI_BoostedMuons`, is introduced to target the highly collimated muon pairs characteristic of boosted dark photon decays. Table 8 summarizes its key configuration differences with respect to `VSI_Leptons`, a collection previously developed for DV searches involving leptons.

Table 8. Key differences between `VSI_Leptons` and `VSI_BoostedMuons`.

Reconstruction step	**VSI_Leptons**	**VSI_BoostedMuons**
Track selection	ID tracks matched to an `Electron` or `Muon` object; tracks associated to the PV are vetoed	ID tracks matched to a `Muon` object; no track-to-PV association veto
2-track compatibility	≥ 1 track in pair has $\lvert d_0 \rvert > 1$ mm	2-track vertex distance from PV > 0.5 mm
Track attachment	$p_{\mathrm{T}} > 1000$ MeV	$p_{\mathrm{T}} > 500$ MeV

The VSI_BoostedMuons algorithm seeds vertex reconstruction from two-track combinations, requiring the tracks to be associated with muons and the seed vertex to be displaced at least 0.5 mm from the PV to suppress prompt backgrounds. Additional ID tracks with $p_{\mathrm{T}} \geq 500$ MeV may be attached to a vertex candidate at a distance to the DV ≥ 2 mm, provided that their inclusion improves the vertex fit quality. The resulting track multiplicity serves as a signal-background discriminant, as signal DVs are expected to contain exactly two tracks.

The vertexing efficiency is defined as the fraction of truth decay vertices within the fiducial region that are successfully matched to a reconstructed vertex. A reconstructed track is truth-matched to a truth particle if more than half of its detector hits are associated with that particle. A reconstructed vertex is truth-matched to a truth decay vertex if the majority of its total track momentum is carried by tracks truth-matched to particles originating from that decay.

Figure 45 shows the vertexing efficiency for a few representative signal points, comparing the VSI_Leptons configuration with the VSI_BoostedMuons configuration with and without the ZPH muons included. The low efficiency in VSI_Leptons at low radius is a consequence of the veto on tracks associated to the PV, as well as the 2-track d_0 requirement. The efficiency up to the last pixel layer is greatly enhanced in VSI_BoostedMuons. After the last pixel layer, the efficiency is recovered with the inclusion of the ZPH muons.

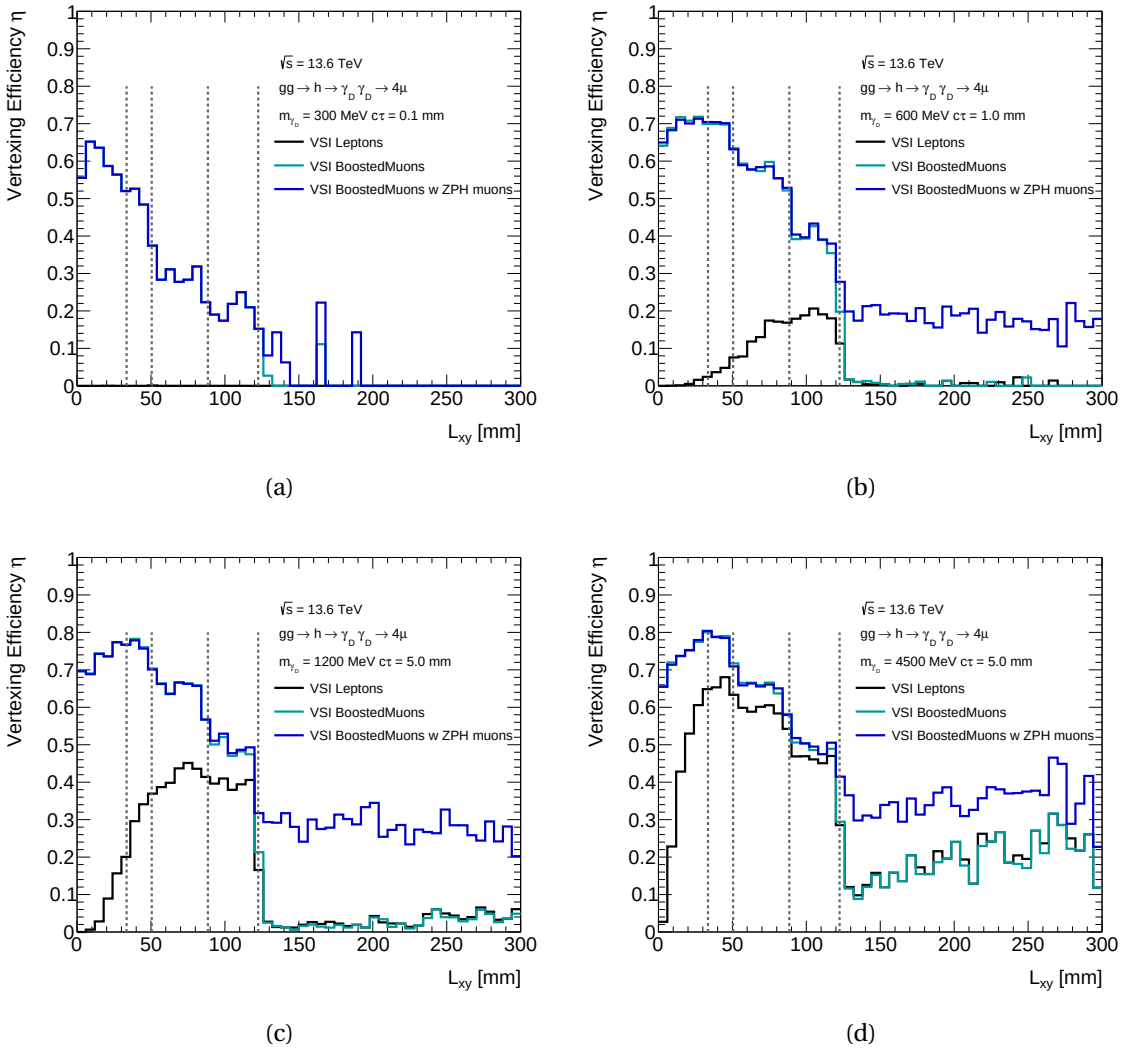


Figure 45. Vertexing efficiency for the four-muon final state as a function of L_{xy} for dark photon mass and lifetime (a) $m_{\gamma_D} = 300$ MeV, $c\tau = 0.1$ mm, (b) $m_{\gamma_D} = 600$ MeV, $c\tau = 1.0$ mm, (c) $m_{\gamma_D} = 1200$ MeV, $c\tau = 5.0$ mm, and (d) $m_{\gamma_D} = 4500$ MeV, $c\tau = 5.0$ mm.

6.3.3 Vertex pulling

When the opening angle between the two muons is small ($\Delta\phi < 0.005$), reconstructed vertices tend to shift toward the next pixel layer, a phenomenon referred to as *vertex pulling*. The reason for vertex pulling is that the vertex position uncertainty is larger in the direction of the track trajectory. As a result, collimated tracks often share pixel clusters, which can bias the vertex reconstruction toward that layer. The exact mechanism causing the vertex misreconstruction may be related to pixel cluster hit position determination in the tracking procedure, and is not yet fully understood. A study of the characterization of the vertex pulling effect is detailed in Section A.0.0.

Two cases of vertex pulling are identified:

1) DVs near the IP being pulled to Insertable B-Layer (IBL).

2) DVs produced in between two pixel layers being pulled to the further layer.

These two types of vertex pulling effects are illustrated using signal MC in Figure 46. Signal DVs from short-lifetime dark photons, which decay near the IP, are pulled to the IBL. The long-lived signal L_{xy} distributions also show distinct peaks at each material layer, closely resembling material interaction background. As a result, to preserve signal acceptance, the L_{xy} material map veto is not applied. Instead, material interaction background is characterized using MC simulation and data. The following selections on vertex L_{xy} significance and dimuon opening angle are specifically efficient in removing DVs from vertex pulling:

- $L_{xy}/\sigma(L_{xy}) \geq 3$
- $\Delta\phi_{\mu\mu} \geq 0.005$

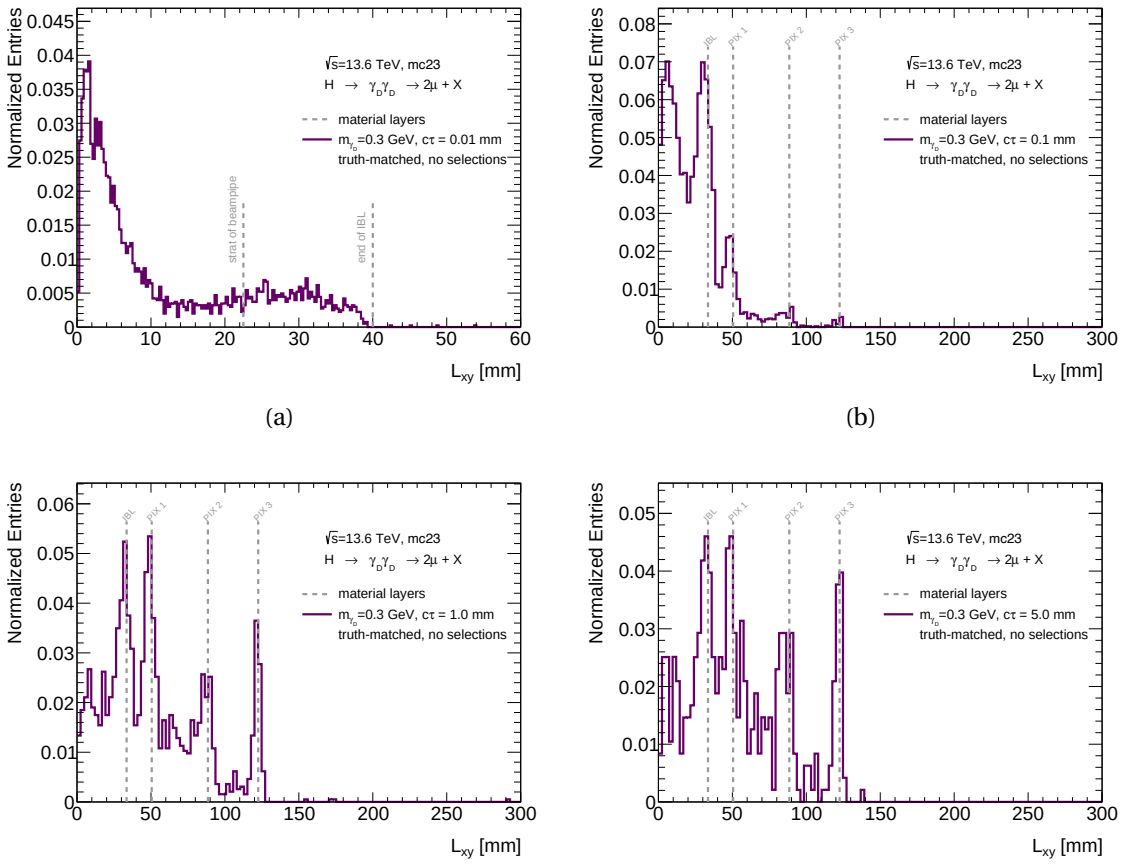


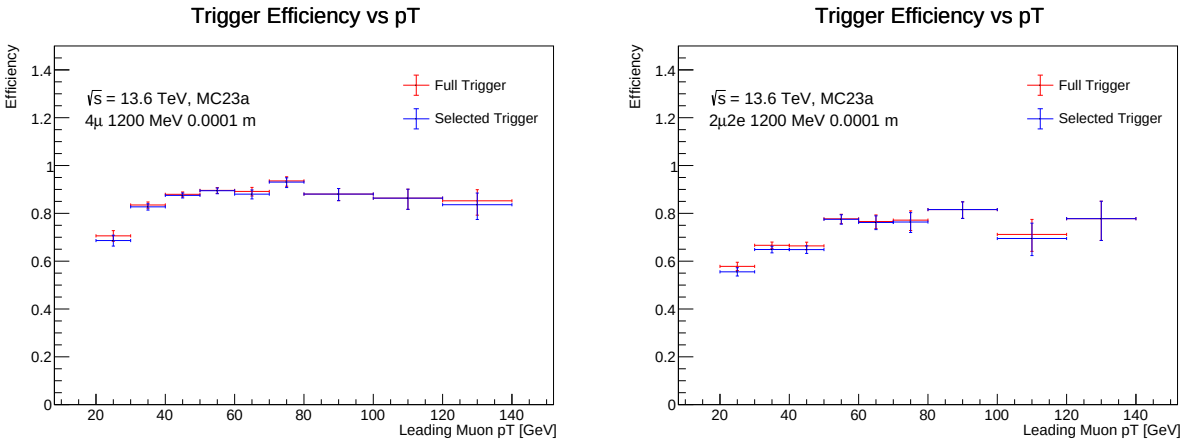
Figure 46. The L_{xy} distribution of HAHM $2\mu 2e$ signal for a 300 MeV dark photon with (a) $c\tau = 0.01$ mm, (b) $c\tau = 0.1$ mm, (c) $c\tau = 1$ mm, and (d) $c\tau = 5$ mm. In (a), the signal has a very short lifetime, and a fraction of DVs near the IP are being reconstructed at IBL. In (b), (c), and (d), the signal has a longer lifetime, therefore more DVs are being reconstructed at every pixel layer. The vertex pulling effects in signal MC are manifested in the peaks at the pixel layers in these distributions. No other selections have been applied to these distributions except for requiring a pair of truth-matched opposite-sign (OS) muons.

6.3.4 Triggers

The analysis uses single and dimuon triggers, combined with logical OR. The list of triggers used by year is given in Table 9. The most efficient combinations of two single-muon triggers and two dimuon triggers are selected. A comparison of the trigger efficiencies obtained using all unprescaled low-threshold single-muon and dimuon triggers and those obtained using the selected triggers listed in Table 9 is shown in Figure 47.

Table 9. Summary of chosen muon triggers by year

Year(s)	**Triggers**
2022–2024	**Single Muon:** HLT_mu24_ivarmedium_L1MU14FCH, HLT_mu50_L1MU14FCH **Dimuon:** HLT_2mu14_L12MU8F, HLT_mu22_mu8noL1_L1MU14FCH
2017–2018	**Single Muon:** HLT_mu26_ivarmedium, HLT_mu50 **Dimuon:** HLT_2mu14, HLT_mu22_mu8noL1
2016	**Single Muon:** HLT_mu26_ivarmedium, HLT_mu50 **Dimuon:** HLT_2mu14, HLT_mu22_mu8noL1
2015	**Single Muon:** HLT_mu20_iloose_L1MU15, HLT_mu50 **Dimuon:** HLT_2mu10, HLT_mu18_mu8noL1



(a) (b)

Figure 47. Comparison of trigger efficiency of full trigger chain (all low-threshold, unprescaled single and dimuon triggers) and selected triggers of an example 4-muon channel sample (left) and 2-muon channel sample (right)

6.4 Event selections

6.4.1 Baseline Selections

Events are preselected for analysis if they pass the trigger and contain at least two muons matched to a reconstructed DV. Preselected events for both the 2-muon and 4-muon channels must satisfy the baseline selections in Table 10.

Table 10. Baseline selections for the 2μ and 4μ channels.

- Pre-selections: the event satisfies the trigger requirement, contains at least one DV reconstructed by the VSI BoostedMuons algorithm (Section 6.3.2), and passes all selections required by the `LLP1` derivation.
- The event contains:
  - one displaced vertex (DV) in the 2μ channel, or
  - two displaced vertices in the 4μ channel,

  each fulfilling the selection criteria listed in Table 11.
- The event contains at least:
  - one vertexed muon matched to a single-muon trigger object in the 2μ channel, or
  - two vertexed muons matched to dimuon trigger objects in the 4μ channel,

  as defined in Table 9. Trigger-matched muon(s) must have a combined p_{T} exceeding the corresponding trigger threshold by 5%, and must satisfy the `Tight` isolation working point when matched to an isolated-muon trigger.

Figure 48 and Figure 49 show the 2-muon signal and heavy flavor MC kinematic distributions for events satisfying the baseline selections. HF background muons peak at lower transverse momentum than signal, motivating a tighter minimum p_T requirement. Examining the *close-by-corrected* muon isolation variables `ptvarcone30`/p_T and `topoetcone20`/p_T, signal muons are well isolated, whereas the background exhibits a broad tail of loosely isolated

Table 11. Common baseline selections for dimuon DVs.

Selection	Requirement
VSI configuration	BoostedMuons
Secondary vertex number of muon tracks	≥ 2
Number of dimuon DVs	≥ 1
Vertexed muon p_{T}	> 3 GeV
Vertexed muon $\lvert\eta\rvert$	< 2.5
dimuon invariant mass $m_{\mu\mu}$	< 20 GeV
DV transverse displacement L_{xy}	$\in [1(0.5), 300]$ mm for 2μ (4μ) channel

events. Transverse displacement provides further discrimination; the background is dominated by a large prompt peak that falls rapidly exponentially with L_{xy}, while the signal extends to larger displacement. Finally, the background rate is significantly enhanced at $m_{\mu\mu} \lesssim 5$ GeV, including known resonances, while at higher masses it is lower and more uniform.

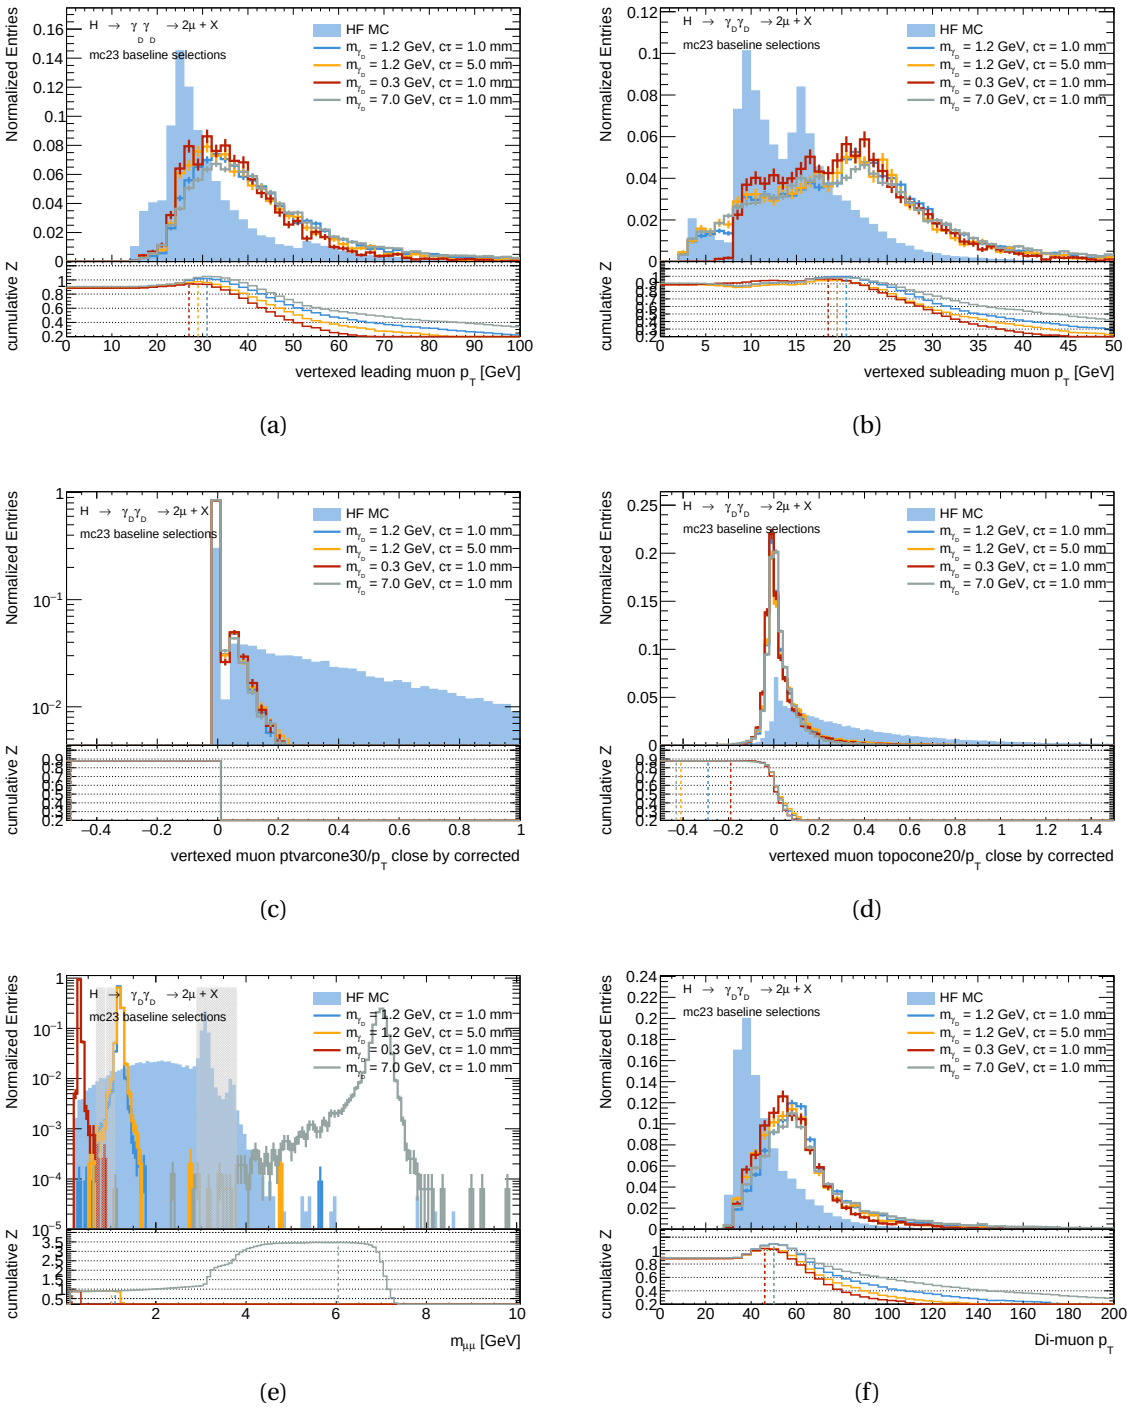
The same kinematic distributions are shown for the 4-muon channel in Figure 50, Figure 51 and **??**. Again, the background muons peak at lower transverse momentum. The 4-muon background shows a significant tail of less-isolated muons. The background dimuon invariant mass spectrum is more uniform than in the 2-muon channel, with less-prominent resonances. Finally, the background falls more rapidly in transverse displacement.

6.4.2 Signal Region Definition

After the baseline-selection, a tighter set of requirements is applied to select signal-like events. These selections are based on track and vertex quality, muon p_{T} and isolation, transverse displacement, and consistency of the DV topology with two-body decay. The full set of requirements is detailed in the following section.

6.4.2.1 Muon selections

Signal muons are expected to be boosted as they originate from a 125 GeV Higgs boson decaying through an intermediate low-mass resonance (dark photon). In contrast, many

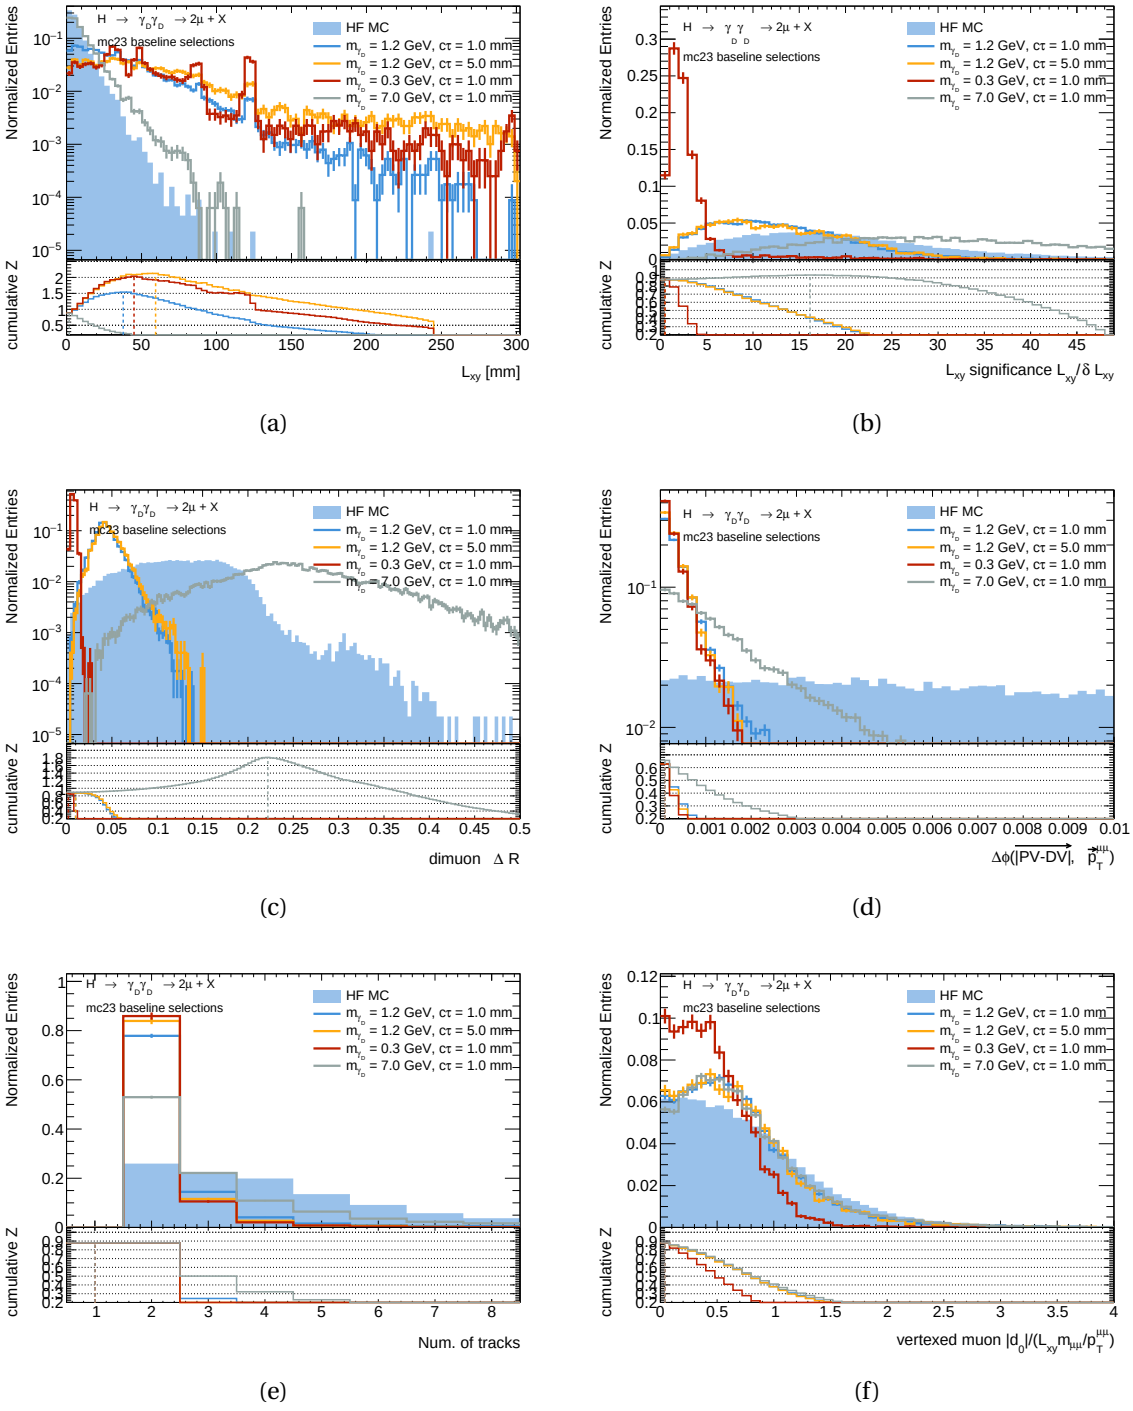


(a) (b)

(c) (d)

(e) (f)

Figure 48. Kinematic distributions of signal and heavy flavor simulation for the 2-muon channel after the baseline selections: muon p_{T}, isolation, dimuon invariant mass, and dimuon p_{T}.



(a) (b)

(c) (d)

(e) (f)

Figure 49. Kinematic distributions of signal and heavy flavor simulation for the 2-muon channel after the baseline selections: DV transverse displacement L_{xy}, L_{xy} significance, dimuon opening angle ΔR, dimuon pointing angle $\Delta\phi(\overrightarrow{DV} - \overrightarrow{PV}, \vec{p}_{\mu\mu})$, DV track multiplicity, and muon $d_0/(L_{\mathrm{xy}} m_{\mu\mu}/p_{\mathrm{T}}^{\mu\mu})$.

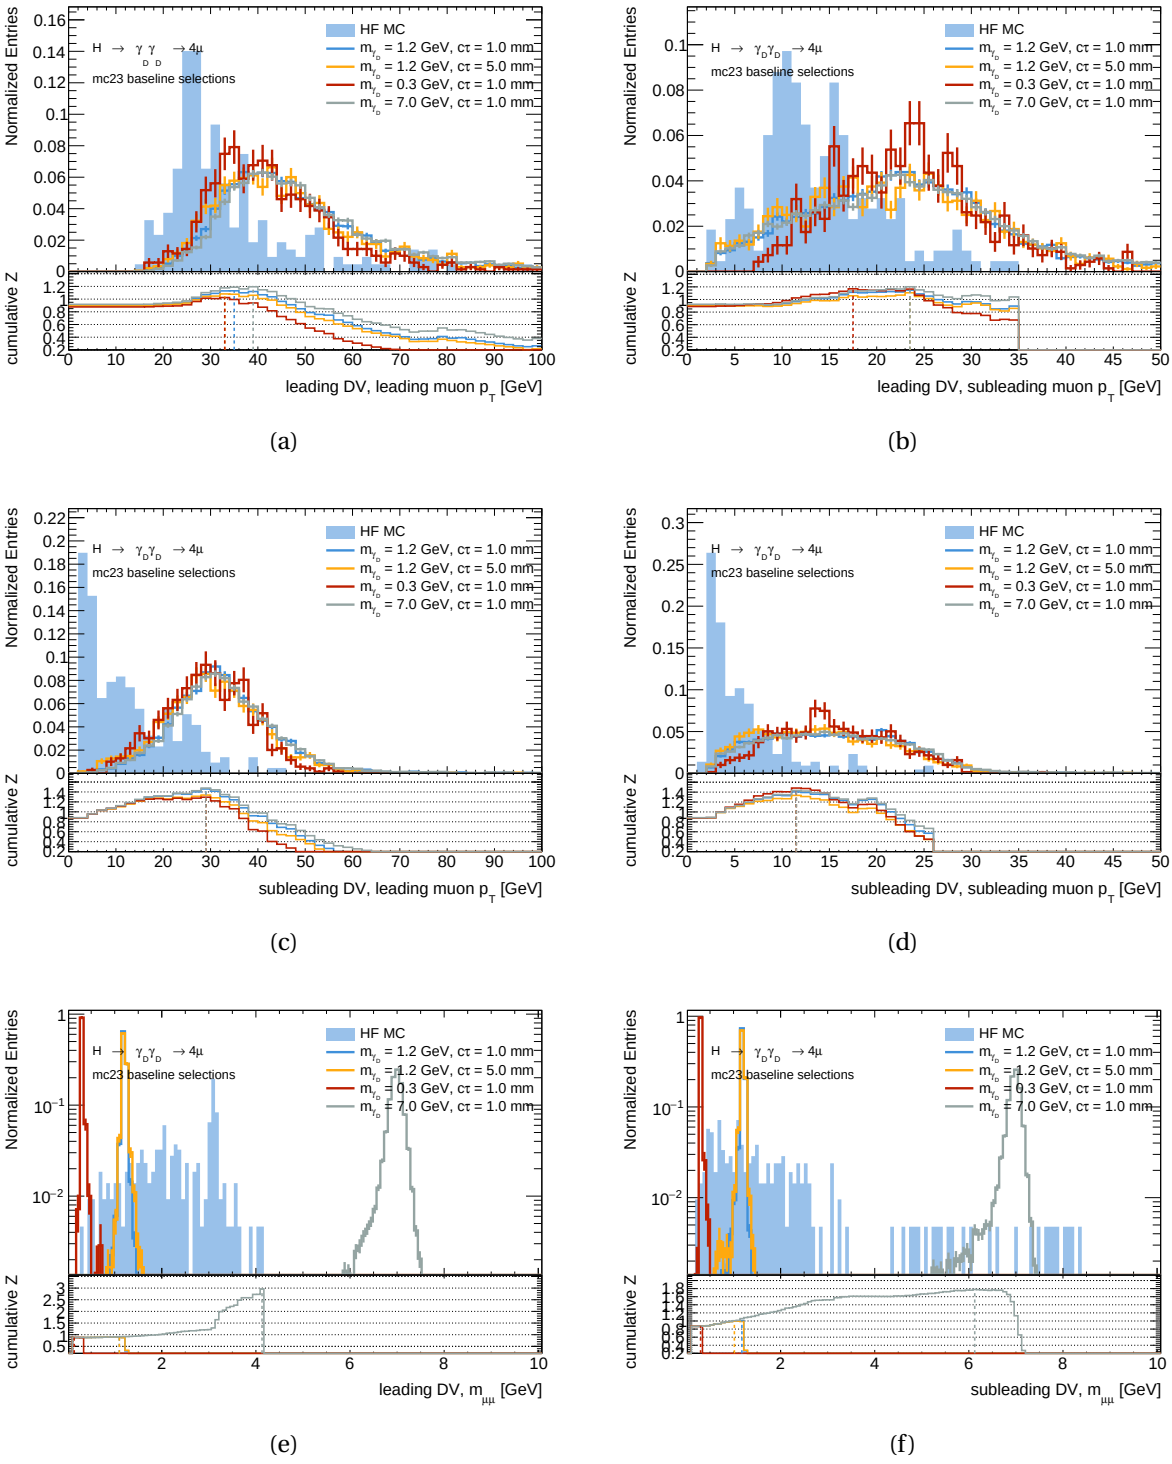
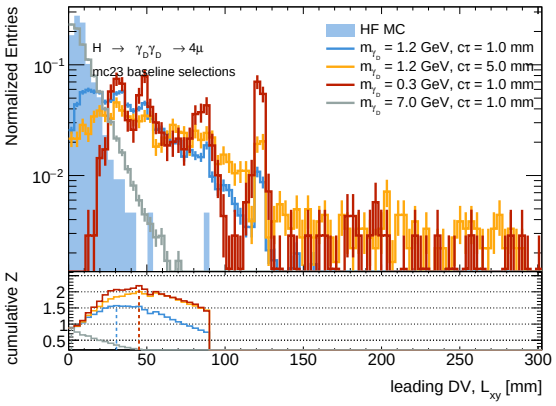
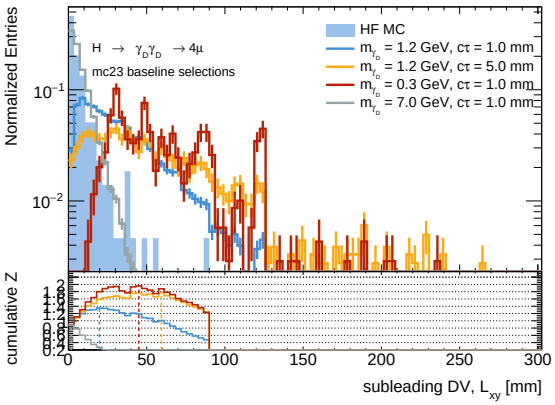


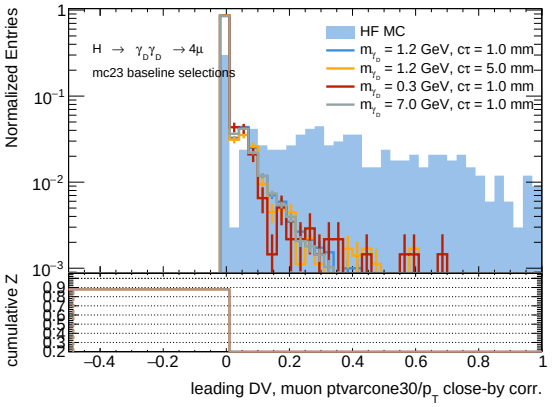
Figure 50. Kinematic distributions of signal and heavy flavor simulation for the 4-muon channel after the baseline selections: leading muon p_{T}, sub-leading muon p_{T}, and dimuon $m_{\mu\mu}$ for leading and sub-leading dimuons.



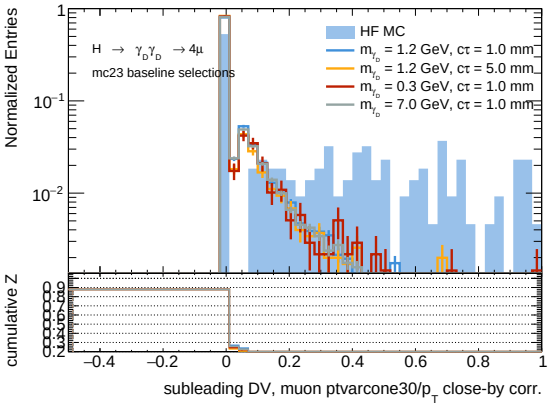
(a)



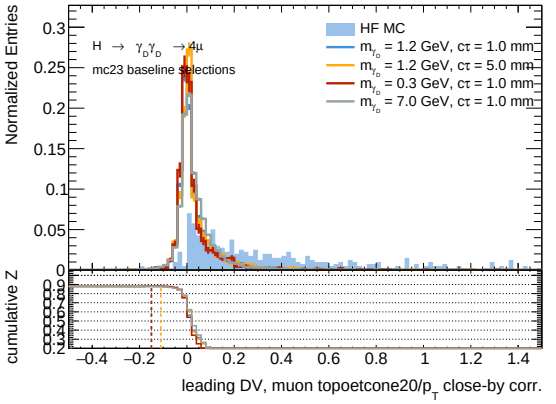
(b)



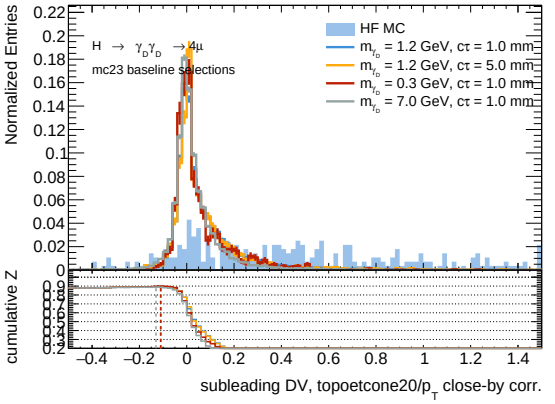
(c)



(d)



(e)



(f)

Figure 51. Kinematic distributions of signal and heavy flavor simulation for the 4-muon channel after the baseline selections: DV L_{xy}, muon track-based isolation, and muon calorimeter-based isolation for leading and sub-leading dimuons.

background sources, such as material interactions and in-flight decays, are not expected to be boosted and thus tend to produce muons with lower transverse momentum and a larger angular separation. To suppress such backgrounds, the leading muon on the DV is required to have a p_{T} of at least 15 GeV, and the sub-leading muon a minimum p_{T} of 5 GeV.

A large amount of background originates from meson decays in flight, where a combined muon is formed by incorrectly matching an ID track of a charged meson to an MS track of a muon from the meson decay. To suppress this background, muons are required to satisfy identification quality requirements outlined in Section 6.3. In addition, the combined muon track fit χ^2/DoF is required to be less than 2 for standard CB muons and LRT muons, and the muon ID track is required to have at least one TRT hit if muon $|\eta| < 2$.

Muon isolation requirements are applied to suppress background from HF decays and hadronic interactions with the detector material. Both track-based and calorimeter-based isolation are used. Track-based isolation is computed by summing the p_{T} of nearby tracks within a variable-sized isolation cone around the muon, denoted as $p_{\mathrm{T}}^{\mathrm{varcone30}}$. Calorimeter-based isolation is determined by summing the energy of topological clusters (topo-clusters) within a fixed-size cone around the muon, denoted as $E_{\mathrm{T}}^{\mathrm{topocone20}}$. Since signal muons tend to be collimated, a close-by correction is applied to the isolation variables to subtract the contribution from the nearby signal muon that may enter the isolation cone, ensuring an accurate estimation of the isolation. The *Loose* and *Tight* isolation working points, with close-by correction applied, are defined as:

$$\text{Loose:} \quad p_{\mathrm{T}}^{\mathrm{varcone30}}/p_{\mathrm{T}} < 0.15 \text{ and } E_{\mathrm{T}}^{\mathrm{topocone20}}/p_{\mathrm{T}} < 0.30$$

$$\text{Tight:} \quad p_{\mathrm{T}}^{\mathrm{varcone30}}/p_{\mathrm{T}} < 0.04 \text{ and } E_{\mathrm{T}}^{\mathrm{topocone20}}/p_{\mathrm{T}} < 0.15$$

The 2-muon signal region requires muons passing the *Tight* isolation working point, while the 4-muon channel uses the *Loose* isolation working point due to significantly less background.

6.4.2.2 Displaced vertex selections

Each displaced vertex (DV) is formed with at least 2 muon tracks, allowing additional ID tracks to be associated with it. The selections of signal DV make use of the kinematic features of a 2-body decay.

Vertex transverse displacement The transverse displacement of a DV is defined as

$$L_{\text{xy}} = |\overrightarrow{DV} - \overrightarrow{PV}|_{\text{xy}},$$

where $\overrightarrow{DV}$ and $\overrightarrow{PV}$ are 3D vectors from the global origin to the DV and PV positions, respectively. The significance of the transverse displacement of a DV is defined as

$$L_{\text{xy}} \text{ significance} = L_{\text{xy}}/\sigma(L_{\text{xy}}),$$

For the 2μ channel, DVs are required to have $L_{\text{xy}} \in [1,300]$ mm and L_{xy} significance ≥ 3 to reject vertices formed by prompt muons, including those affected by vertex pulling. The 4μ channel imposes a relaxed L_{xy} cut within [0.5,300] mm, with no cut on L_{xy} significance; this choice significantly improves signal efficiency without dramatically changing the already small expected background rate.

For vertices containing ZPH muons (see Section 6.3), an additional requirement of $L_{\text{xy}} \geq 110$ mm is imposed. Since muons are pair-produced in signal events, no ZPH muon is expected to be associated with vertices located within the last pixel layer.

Additional track veto The secondary vertex reconstruction algorithm, VSI, is configured to attach additional ID tracks to DVs with $L_{xy} \geq 2$ mm if it improves the vertex fit quality. Heavy flavor background events are expected to pick up more additional tracks due to cascading decays, while material interactions contribute via photon conversions and hadronic interactions; a selection on the number of additional tracks can therefore be used to reject such backgrounds.

The origins of the additional tracks were investigated using heavy flavor background and signal MC samples. From simulation studies, in-time additional tracks primarily consist of pions and kaons. However, the largest source of additional tracks is out-of-time pile-up tracks, identified by an unassigned MC-truth PDG ID. Due to the larger vertex fit uncertainty along the direction of flight, pile-up tracks are more likely to be back-to-back with the dimuon system in $\Delta\phi$. As a result, DVs at shorter displacement are more likely to accumulate pile-up tracks.

This correlation between L_{xy} and the number of additional tracks complicates the rejection of backgrounds that naturally have high track multiplicities, such as HF events, while maintaining signal acceptance, since DVs with lower L_{xy} are more likely to accumulate additional tracks from pile-up. Signals with shorter decay lengths, either due to shorter lifetimes or a smaller boost factor (i.e. larger mass), are more affected by a strict veto on DVs with additional tracks. Therefore, the following two types of additional track vetoes are applied.

1) Strict veto: reject DVs with track multiplicity > 2.

2) ΔR-dependent veto: reject DVs where the highest p_{T} additional track is within $\Delta R < 0.5$ of the dimuon system.

Studies of the ΔR distribution between the highest p_{T} additional track and the dimuon system, $\Delta R(p_{\text{max add. trk}}, p_{\mu\mu})$, using heavy flavor background MC, signal MC, and CR data show that in signal MC, additional tracks are mostly back-to-back with the dimuon system, peaking at $2.5 < \Delta R < 3.5$, whereas the heavy flavor background peaks at $\Delta R < 0.5$, with tracks largely identified as pions and kaons. This selection therefore targets DVs where additional tracks originate from physical processes rather than pile-up reconstruction effects, making it less sensitive to pile-up mis-modeling in MC.

While the ΔR-dependent veto is more physics-motivated, it was found to introduce additional background from prompt DY events in data. These events have low L_{xy} and a $m_{\mu\mu}$ peak at twice the muon mass threshold, impacting the sensitivity in the low-mass region of the

2μ channel ($m_{\mu\mu} < 5$ GeV). The strict additional track veto is therefore applied in this region instead.

For the 2μ-channel high-mass region ($m_{\mu\mu} > 5$ GeV) and the 4μ channel, the ΔR-dependent veto is employed, where it improves both signal efficiency and sensitivity.

Muon pointing angles For a particle with a finite lifetime decaying into two muons, the vertex formed by the two muons points toward the direction of the dimuon momentum. The ATLAS detector has different angular resolution in the transverse plane versus the z-direction; the angle between the direction of flight and the dimuon momentum in the x–y plane provides better signal-background separation power. The muon pointing angle in the x–y plane is defined as

$$\Delta\phi(\overrightarrow{DV} - \overrightarrow{PV}, \vec{p}_{\mu\mu}).$$

Requiring the dimuon momentum to point in the direction of the decaying particle removes backgrounds that do not originate from two-body decays, such as semi-leptonic decays of B mesons, material interactions, and accidental crossings of muons.

Muon opening angles The 2μ channel imposes a minimum opening angle requirement on the muon pair in the transverse plane, $\Delta\phi_{\mu\mu} \geq 0.005$, to remove prompt dimuons faking a displaced vertex due to vertex pulling. The $\Delta\phi_{\mu\mu}$ is calculated with respect to the DV. This selection also potentially removes a fraction of photon conversions to muons in the material layers.

Lifetime-scaled muon d_0 For vertices formed by a dimuon pair, the pseudo-proper lifetime of the LLP is defined as

$$\text{pseudo-}c\tau = L_{xy} m_{\mu\mu} / p_{\mathrm{T}}^{\mu\mu}.$$

Since the transverse impact parameter (d_0) of the muon is correlated with the LLP lifetime, a requirement on the lifetime-scaled d_0 is imposed to suppress backgrounds in which the observed displacement is not physically consistent with the LLP decay hypothesis: d_0/pseudo-$c\tau \geq 0.1$.

Cosmic muon veto To reject highly back-to-back muons, which is a distinct signature of cosmic muons, a requirement on the dimuon opening angle, $|\pi - \Delta\phi_{\mu\mu}| \geq 0.01$, is imposed. As discussed in Section 6.5.5, this effectively removes the cosmic background.

Orthogonality to 4μ channel To ensure that the same event is not analyzed in both the 2μ and 4μ channels, each event in the 2μ channel is checked to see if it passes the 4μ selection. If the event passes, it is vetoed in the 2μ channel.

6.4.2.3 2-muon Channel Signal Region

All of the selections above are applied in the 2-muon channel to reduce background from HF decays, material interactions, accidental crossings, and prompt DY events.

The full event selection in the 2-muon channel is shown in Table 12. The only difference in event selections between the low-mass region ($m_{\mu\mu} < 5$ GeV) and high-mass region ($m_{\mu\mu} \geq 5$ GeV) is the L_{xy} requirement and the additional track veto.

The L_{xy} regions are:

- Low-mass region ($m_{\mu\mu} < 5$ GeV): $L_{xy} \in [1, 300]$ mm.
- High-mass region ($m_{\mu\mu} \geq 5$ GeV): $L_{xy} \in [2, 300]$ mm.

This is because different background estimation methods are used in the two regions, and for the high-mass region, a tighter L_{xy} requirement is applied to remove prompt DY background.

The additional track vetoes used are:

– Low-mass region ($m_{\mu\mu} < 5$ GeV): $N_{\text{Tracks}} = 2$ (strict veto).

– High-mass region ($m_{\mu\mu} \geq 5$ GeV): $N_{\text{Tracks}} \geq 2$, and $\Delta R(p_{\text{add. trk}}, p_{\mu\mu}) > 0.5$ (ΔR-dependent veto).

Table 12. Signal region selection for the 2μ analysis channel. "DV" denotes the displaced vertex. "Muons" refers to the two muons forming the DV. χ^2_{ID} and χ^2_{CB} are for the inner-detector and combined tracks, respectively.

Event Selection 2μ-channel Signal Region		
DV	Number of tracks	$= 2$ (≥ 2) for $m_{\mu\mu} < 5$ (≥ 5 GeV)
	Number of muon tracks	$= 2$
	Sum of muon track charges	0 (OS pair)
	L_{xy}	$\in$ [1 (2), 300] mm for $m_{\mu\mu} < 5$ (≥ 5) GeV
	L_{xy} of DV with ZPH muon	≥ 110 mm
	$L_{xy}/\sigma(L_{xy})$	≥ 3
	χ^2/ndof	≤ 5
	$\Delta\phi(\overrightarrow{DV} - \overrightarrow{PV}, \vec{p}_{\mu\mu})$	≤ 0.005
	$\Delta\phi_{\mu\mu}$	≥ 0.005
	$\lvert\pi - \Delta\phi_{\mu\mu}\rvert$	≥ 0.01
	$\Delta R(p_{\text{add. trk}}, p_{\mu\mu})$	> 0.5 for $m_{\mu\mu} \geq 5$ GeV
Muons	Leading p_{T}	> 15 GeV
	Sub-leading p_{T}	> 5 GeV
	Identification	Tight (CB), custom Tight (ZPH), Medium (LRT)
	Isolation	`Tight_VarRad_CloseByCorr`
	χ^2_{ID}/ndof	< 2
	χ^2_{CB}/ndof	< 2
	TRT hits ($\lvert\eta\rvert < 2$)	> 0
	$\lvert d_0\rvert/(L_{xy} m_{\mu\mu}/p_{\text{T}}^{\mu\mu})$	≥ 0.1
Event	Orthogonal to 4μ channel	True

6.4.2.4 4-muon Channel Signal Region

The 4μ channel selects events containing two OS DVs with compatible dimuon masses and a four-muon invariant mass consistent with the Higgs boson mass, targeting the signal process $gg \to h \to \gamma_D\gamma_D \to 4\mu$, where each γ_D decays into a muon pair. The four-muon invariant mass is expected to peak near 125 GeV, and the signal region is defined as $114 < m_{4\mu} < 136$ GeV, corresponding to approximately 5σ of the signal width as determined by a Gaussian fit to the $m_{4\mu}$ distributions. Inverting this requirement yields an approximately signal-free control region for background studies in data.

The full event selection is summarized in Table 13. Relative to the 2μ channel, the four-muon mass constraint significantly reduces backgrounds, motivating the relaxation of the L_{xy}, $\Delta\phi(\overrightarrow{DV} - \overrightarrow{PV}, \vec{p}_{\mu\mu})$, and isolation requirements. The L_{xy} significance and $|d_0|/(L_{xy} m_{\mu\mu}/p_{\mathrm{T}}^{\mu\mu})$ selections are removed entirely. The ΔR-based additional track veto is retained, following the 2μ channel high-mass region.

Additionally, both dimuon systems are expected to be compatible with the dark photon mass m_{γ_D}, resulting in a small invariant mass difference $\Delta_{m_{\mu\mu}}$ between them. The vertex mass resolution degrades with increasing dimuon mass, as higher-mass dark photons produce higher-momentum muons with smaller curvature and therefore poorer momentum resolution, increasing the $\Delta_{m_{\mu\mu}}$ width at higher mass. This mass-dependent resolution is accounted for by defining two symmetric linear boundaries in the $m_{\mu\mu}^{\text{Leading}}$ versus $m_{\mu\mu}^{\text{Sub-leading}}$ plane, as illustrated in Figure 52. These boundaries define the linear selection on the two dimuon masses:

$$\begin{aligned} \textbf{upper boundary:} \quad & m_{\mu\mu}^{\text{Sub-leading}} \leq 1.1 \times m_{\mu\mu}^{\text{Leading}} + 0.15\ \text{GeV} \\ \textbf{lower boundary:} \quad & m_{\mu\mu}^{\text{Sub-leading}} \geq 0.91 \times m_{\mu\mu}^{\text{Leading}} - 0.15\ \text{GeV} \end{aligned} \tag{6.4}$$

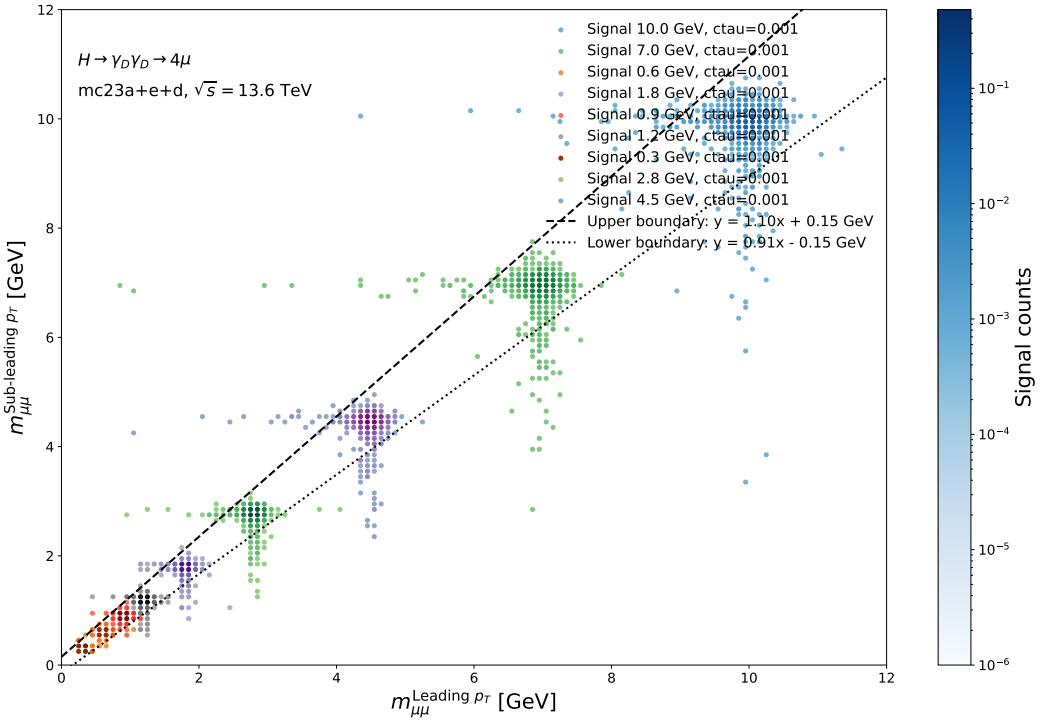


Figure 52. Signal region boundaries in the dimuon mass plane, defined by $m_{\mu\mu}^{\text{Leading}}$ and $m_{\mu\mu}^{\text{Subleading}}$.

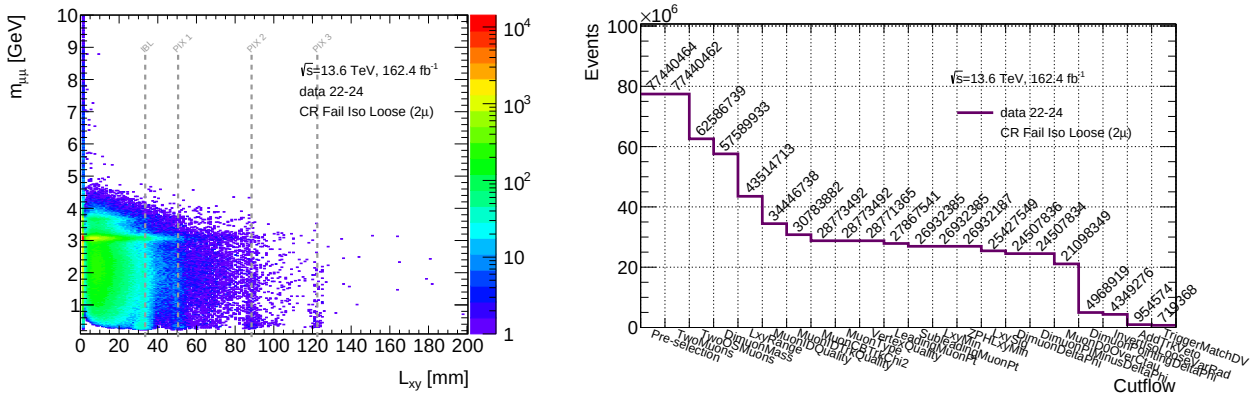
Table 13. Signal region selection for the 4μ analysis channel. "DV" denotes each displaced vertex. "Muons" refers to the two muons forming that DV. χ^2_{ID} and χ^2_{CB} are for the inner-detector and combined tracks, respectively.

Event Selection 4μ-channel Signal Region		
DV	Number of tracks	≥ 2
	Number of muon tracks	$= 2$
	Sum of muon track charges	0 (OS pair)
	L_{xy}	$\in [0.5, 300]$ mm
	L_{xy} of DV with ZPH muon	≥ 110 mm
	χ^2/ndof	≤ 5
	$\Delta\phi(\overrightarrow{DV} - \overrightarrow{PV}, \vec{p}_{\mu\mu})$	≤ 0.05
	$m_{\mu\mu}$	≤ 14 GeV
	$\lvert\pi - \Delta\phi_{\mu\mu}\rvert$	≥ 0.01
	$\Delta R(p_{\text{add. trk}}, p_{\mu\mu})$	> 0.5
Muons	Leading p_{T}	> 15 GeV
	Sub-leading p_{T}	> 5 GeV
	Identification	Tight (CB), custom Tight (ZPH), Medium (LRT)
	Isolation	`Loose_VarRad_CloseByCorr`
	χ^2_{ID}/ndof	< 2
	χ^2_{CB}/ndof	< 2
	TRT hits ($\lvert\eta\rvert < 2$)	> 0
Event	Number of displaced vertices	$= 2$
	$m_{\mu\mu}^{\text{Sub-leading}}$	$\leq 1.1\, m_{\mu\mu}^{\text{Leading}} + 0.15$ GeV $\geq 0.91\, m_{\mu\mu}^{\text{Leading}} - 0.15$ GeV
	$m_{4\mu}$	$\in [114, 136]$ GeV

6.4.3 Control region definitions

Several control regions are defined to estimate background contributions in the signal regions. These regions are constructed to be kinematically similar to the signal region, but remain orthogonal by inverting key selection criteria. This enables data-driven estimation of background rates and shapes while minimizing signal contamination.

Fail Loose Isolation In the Fail Isolation control region, events are selected where at least one DV muon fails the "Loose" isolation requirement (`Loose_VarRad_CloseByCorr`), while all other SR selections remain. Since the signal region applies Tight isolation in the 2μ channel and Loose isolation in the 4μ channel, inverting the Loose isolation requirement minimizes signal acceptance in the control region. Selecting non-isolated muons enhances contributions from heavy flavor decays or hadronic jets. Data distributions in the `Fail Iso Loose` region are shown in Figure 53.



(a) (b)

Figure 53. Data in `fail Isolation CR`. (a) dimuon invariant mass vs L_{xy} and (b) cutflow. The cutflow shows all selections applied in the `fail Isolation CR`, which are the same selection as SR except for inverting the isolation requirement.

Fail Pointing Angle The Fail Pointing Angle control region ($\Delta\phi(\overrightarrow{DV}-\overrightarrow{PV}, \vec{p}_{\mu\mu}) > 0.005$) selects DVs not originating from two-body decays. It is enriched in backgrounds from accidental

crossings and HF decays. It is used for background estimation and validation in the 2-muon channel high-mass region ($m_{\mu\mu} > 5$ GeV).

Inverted 4-muon Higgs Mass The 4μ invariant mass control region inverts the signal mass window 114–136 GeV described in Section 6.4.2.4. This window corresponds to approximately 5σ of the Gaussian-fitted signal width, evaluated at the signal mass with the worst resolution. By inverting this selection, the control region is constructed with negligible signal contamination.

Table 14. Summary of control regions

Control Region	Selection	Background
Fail Loose Isolation	$p_{\mathrm{T}}^{\mathrm{varcone30}}/p_{\mathrm{T}} \geq 0.15$ $E_{\mathrm{T}}^{\mathrm{topocone20}}/p_{\mathrm{T}} \geq 0.30$	Material interactions, Heavy flavor
Fail Pointing Angle	$\Delta\phi(\overrightarrow{DV} - \overrightarrow{PV}, \vec{p}_{\mu\mu}) \geq 0.005$	Heavy flavor
Three-track DVs	DV has exactly 3 attached tracks	Material Interactions
Inverted 4-muon Higgs Mass	$m_{4\mu} \notin [114, 136]$ GeV, $m_{4\mu} \geq 40$ GeV	Heavy flavor

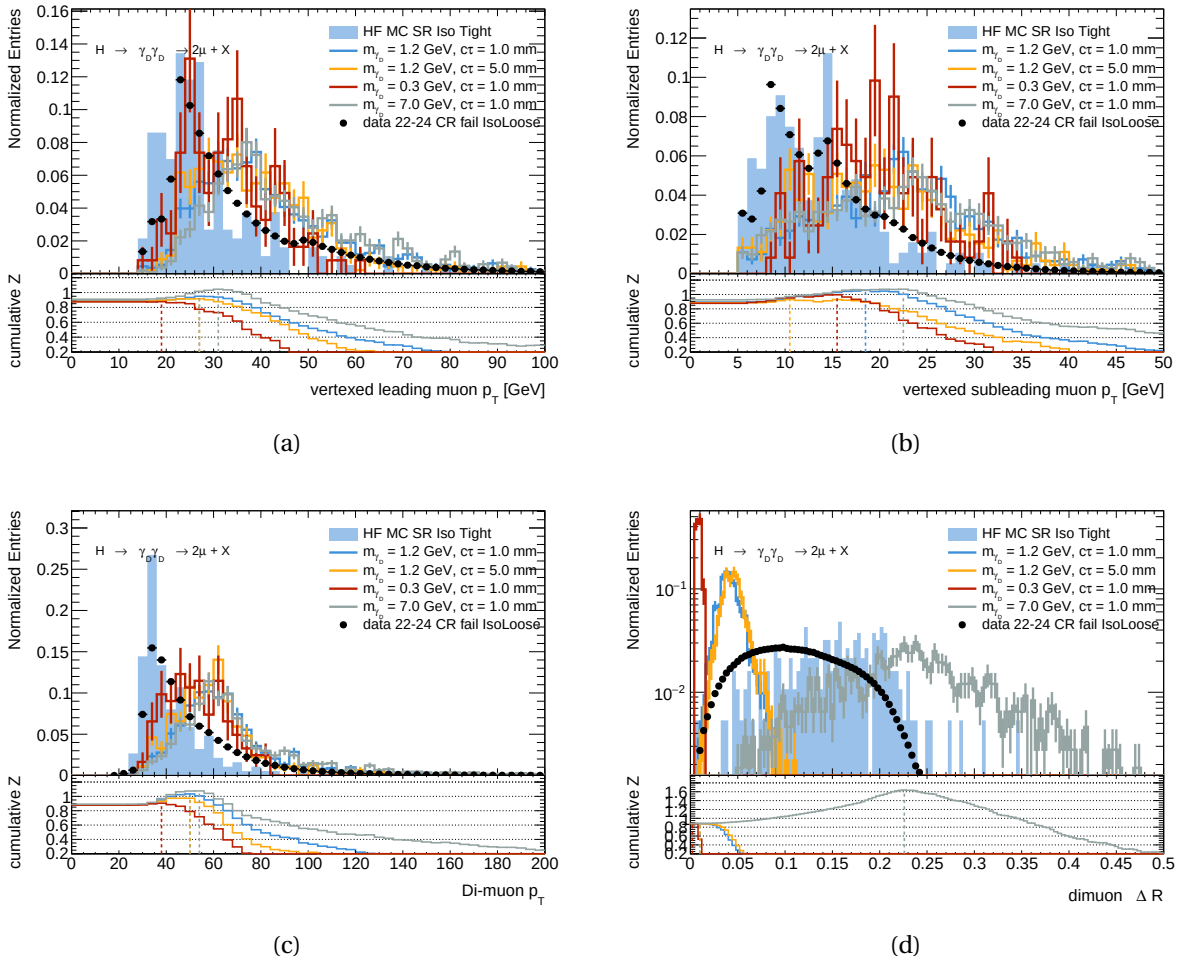
6.4.4 Signal vs background comparisons

Figures 54 and 55 compare signal, simulated HF background in SR, and partial Run 3 data in `Fail Loose Isolation` CR for the 2μ channel. The background is peaked at lower muon and dimuon p_T with significant overlap with the signal. In contrast, the signal dimuon pointing angle distributions are peaked at smaller values, while the background is approximately flat.

The $m_{\mu\mu}$ distributions show that the HF simulation describes data for $m_{\mu\mu} \lesssim 5$ GeV, but has no statistics at higher-mass. The overall shape of the $m_{\mu\mu}$ distribution clearly delineates regions of high and low background, motivating separate treatment in the analysis. Due to the large amount of HF background and a complex background composition in the $m_{\mu\mu} < 5$ GeV region, parametric functions are used to model signal and background and profile likelihood fits are performed in small, sliding $m_{\mu\mu}$ windows for best sensitivity. In the $m_{\mu\mu} > 5$ GeV region, where the background is relatively small with a simpler composition, a data-driven ABCD background

estimation combined with a cut-and-count analysis is used. Finally, the background extends far in L_{xy}, reaching beyond ~ 50 mm. Both regions in the 2-muon channel are further separated into multiple L_{xy} bins to optimize the sensitivity of long-lifetime signals.

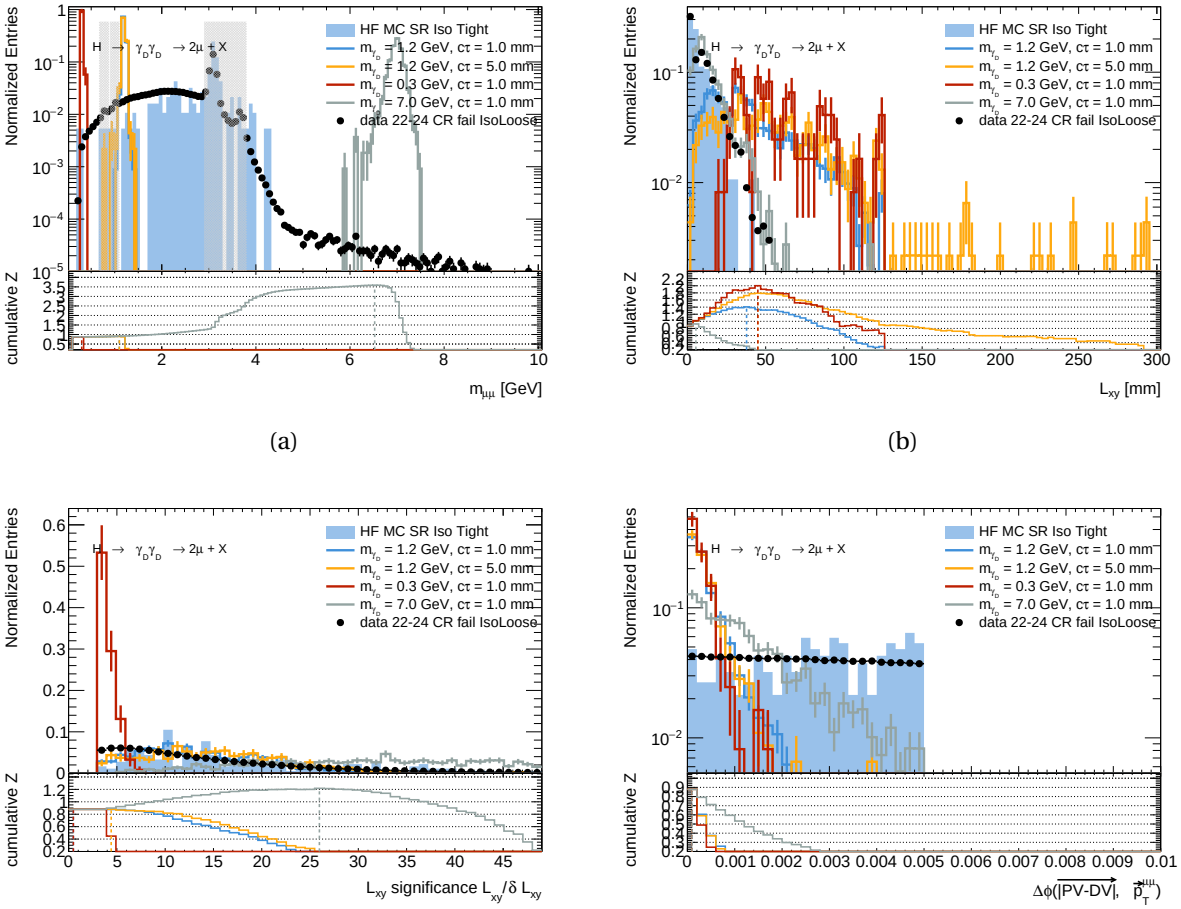
Figures 56 and 57 show a similar set of kinematics for the 4μ channel, but no simulated background events pass the selection. Instead, both the non-isolated control region data and the isolated control region outside of the 4μ invariant mass window around the Higgs mass are shown. The non-isolated background in data peaks at low p_T, significantly overlapping signal, but is flat in dimuon pointing angle. Since the overall background size is small, there is no need to bin the analysis according to L_{xy}. In contrast to the 2μ channel, there are no data events with $m_{\mu\mu}$ above the J/ψ peak in the non-isolated region. Since this channel is statistically limited, the search is performed in the 4μ mass window without separating the analysis into high- and low-$m_{\mu\mu}$ regions.



(a) (b)

(c) (d)

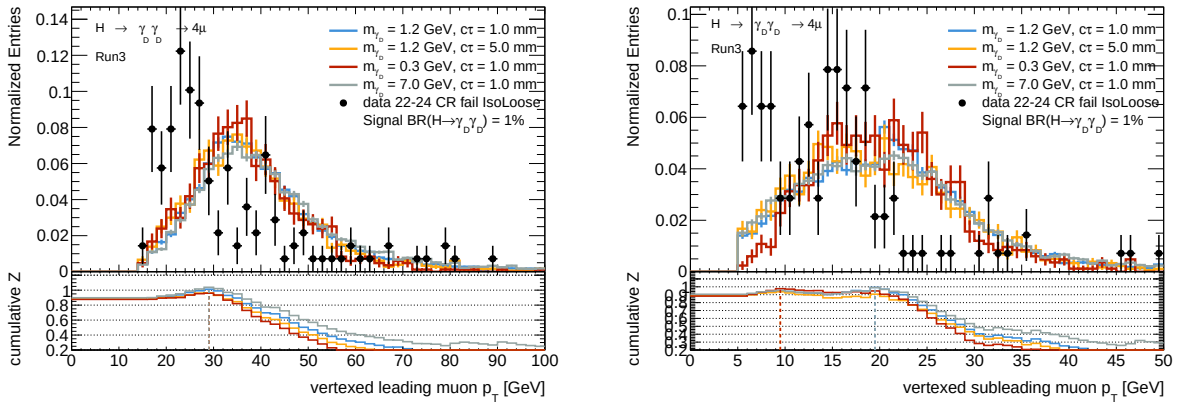
Figure 54. Kinematic distributions of 2-muon signal and heavy flavor MC under SR selections, and data 22–24 under CR Iso Loose selections: muon p_{T}, dimuon p_{T}, and opening angle.



(a) (b)

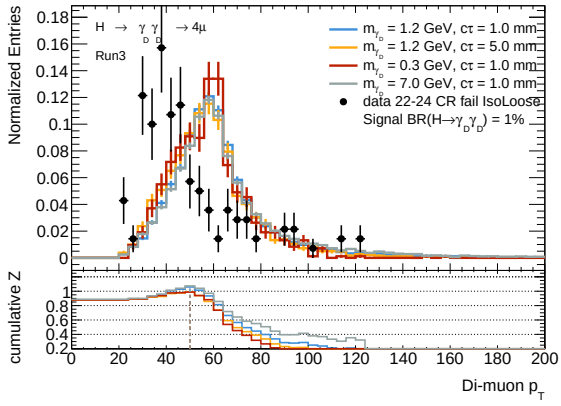
(c) (d)

Figure 55. Kinematic distributions of 2-muon signal and heavy flavor MC under SR selections, and data 22–24 under CR Iso Loose selections: dimuon invariant mass, transverse displacement, L_{xy} significance, and pointing angle.

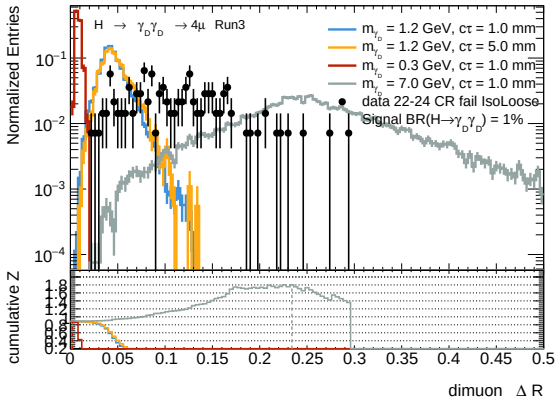


(a)

(b)



(c)



(d)

Figure 56. Kinematic distributions of 4-muon signal and heavy flavor MC under SR selections, and data 22–24 under CR Iso Loose selections: muon p_{T}, dimuon p_{T}, and opening angle.

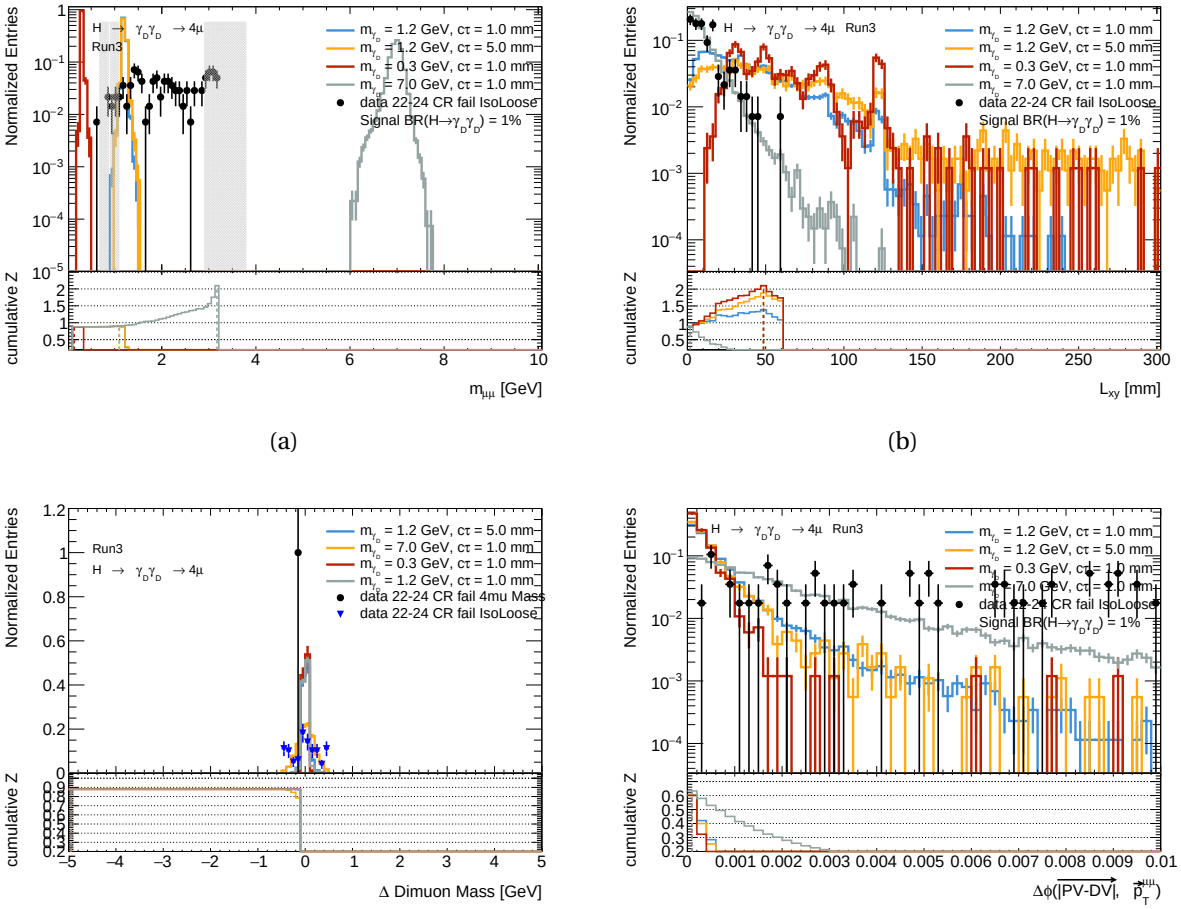


Figure 57. Kinematic distributions of 4-muon signal and heavy flavor MC under SR selections, and data 22–24 under CR Iso Loose selections: dimuon invariant mass, transverse displacement, dimuon mass difference, and pointing angle.

6.5 Background Composition

This section describes the background sources in more detail, how they contribute to the different signal regions, and their relative compositions. In some regions, multiple background processes contribute non-negligibly to the signal region. However, their $m_{\mu\mu}$ distributions are smooth and show no localized structures consistent with a narrow resonance signal. Consequently, these backgrounds can be modeled using closed-form continuous parametric functions. In other regions, the background composition is simpler and the overall yields are expected to be small. In these cases, a cut-and-count strategy is adopted, estimating the background using a data-driven ABCD method.

The 2μ low-mass region uses a parametric background model combined with a profile likelihood fit to extract the signal strength, due to its complex background composition and large HF background. Figure 58 shows the data distribution in the non-isolated CR in different L_{xy} windows.

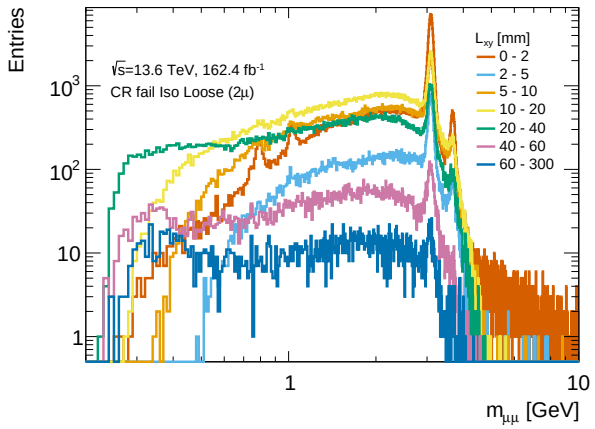


Figure 58. Distribution of the dimuon invariant mass $m_{\mu\mu}$ in different L_{xy} windows in 2022–2024 data in the Fail Loose Isolation CR.

The 2μ high-mass region has a simpler and relatively small background, and therefore uses an ABCD method to estimate the HF background, along with the uncorrelated crossing background estimates in Section 6.5.4, combined with a cut-and-count analysis.

The 4μ channel is a near-zero background analysis, with limited statistics in both the isolated and non-isolated control region data. The dominant background arises from events with two heavy flavor decays of similar $m_{\mu\mu}$. This channel therefore also uses a cut-and-count analysis, with the background estimated using a data-driven ABCD method.

6.5.1 Heavy flavor background

Muons from HF decays that form a secondary vertex are the dominant source of background, due to the large production cross section of $b\bar{b}$ and $c\bar{c}$ events and the sizable lifetime of B mesons, as illustrated in Figure 59.

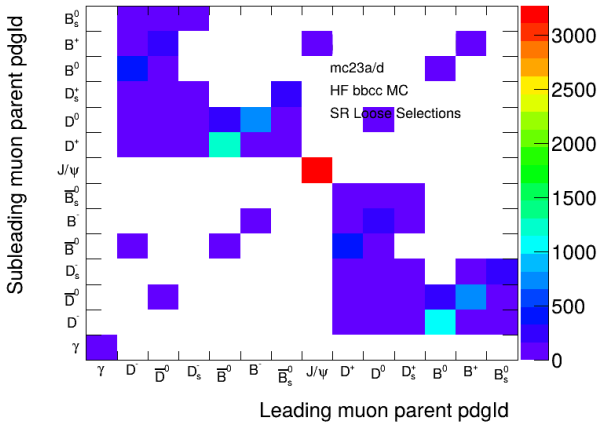
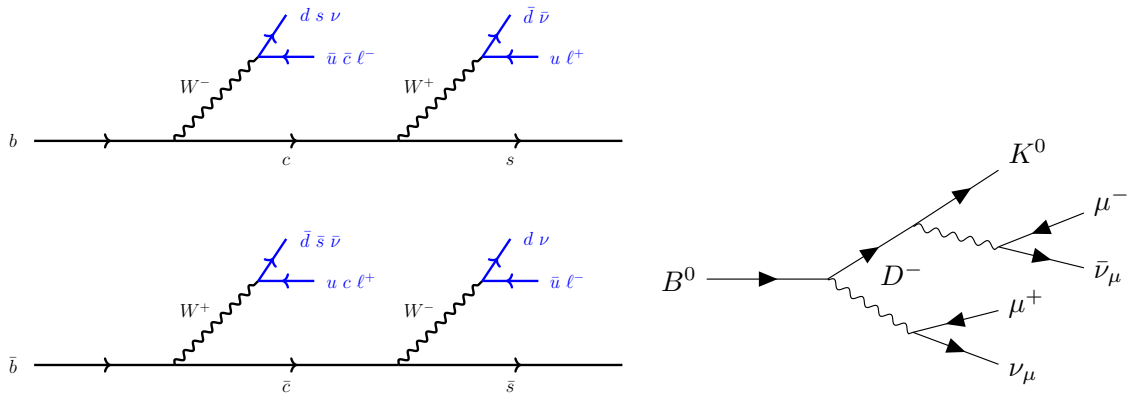


Figure 59. Parent particles of the leading and subleading muons from HF background simulation with a loose 2μ channel SR selection.

The large contribution from $J/\psi \to \mu\mu$ decays is suppressed in the 2μ channel by a mass veto covering the J/ψ mass window. The remaining non-resonant contribution mainly comes from cascade decays of B mesons. Figure 60 shows illustrative diagrams of HF background from semi-leptonic decays of a B meson and the decays of b quarks.

Figure 61 shows the $m_{\mu\mu}$ distributions in data and HF simulation, normalized to unity. The data are shown in the Fail Loose Isolation CR, while HF simulation is shown in both the CR

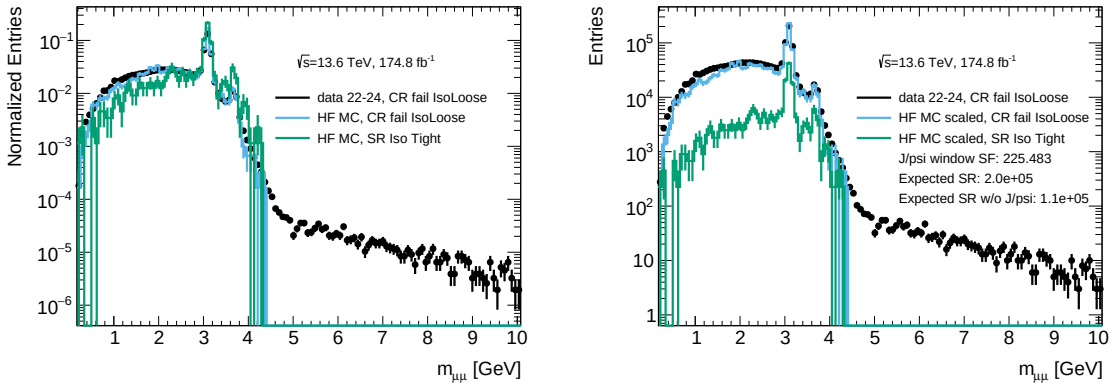


(a) (b)

Figure 60. Feynman diagrams for decays of b quarks, and a semi-leptonic decay of B^0 meson.

and the SR. A visible difference in the $m_{\mu\mu}$ shape is observed between the isolated SR and the non-isolated CR.

In Figure 61b, the CR simulation is scaled to CR data to match the event yield in the J/ψ mass window ([2.9, 3.2] GeV), exploiting the well-modeled fraction of B meson decays to J/ψ and semi-leptonic final states. The same scale factor is applied to the simulation in the SR to predict the HF background yield.



(a) (b)

Figure 61. The dimuon invariant mass distributions in data and HF simulation. Data distribution is shown in the fail-isolation CR, while HF MC is shown in both fail-isolation CR and pass-isolation SR. Figure (a) shows distributions normalized to unity. Figure (b) shows the distributions where MC in CR is scaled to have the same number of events as data under the J/ψ mass window, and MC in SR is scaled by the same factor as the CR.

However, a closure test, comparing the isolated J/ψ mass window data to the scaled HF simulation, shows that the simulation does not predict the exact background yield in the SR. In addition, HF contributions at $m_{\mu\mu} > 5$ GeV are present in data but are not modeled in the simulation.

For two muons from HF decays to produce an invariant mass larger than the B meson mass, each muon must originate from a different hadron and be vertexed by chance at large opening angle, a process far less probable than a dimuon pair from a single hadron decay. The simulation lacks sufficient statistics to model this background in the 2μ channel, and has negligible statistics for the 4μ channel. Consequently, a data-driven parametric background model is used for the 2μ channel $m_{\mu\mu} < 5$ GeV region, and a data-driven cut-and-count ABCD method is used for the $m_{\mu\mu} > 5$ GeV region and the 4μ channel.

6.5.2 Material interactions

The material interaction background comes from two processes: hadronic interactions of secondary particles with detector material, and photon conversions into muon pairs (the latter roughly four orders of magnitude less frequent than conversions to electrons). Both are heavily suppressed by the p_{T} and isolation requirements. While these events could be further reduced by vetoing DVs in material-dense regions of the detector, such a veto would reduce the signal acceptance significantly due to the vertex pulling effect (Section 6.3.3). Instead, the background is characterized across the $m_{\mu\mu}$–L_{xy} plane using simulation and data to assess its size relative to HF and guide the estimation strategy.

The properties of material interaction backgrounds are studied using the hadronic τ and $Z+\gamma$ MC samples described in Section 6.2, with the baseline selections of Section 6.4 applied. Additional truth-level selections are applied to isolate material interaction events in simulation, where muons are required to be produced in `Geant4` simulations where these processes are simulated, and for photon conversion events, muons are required to originate from the same

photon parent. The Geant4 simulation also simulates decay in-flights of prompt hadrons, such as kaon decays in-flight to muons, $K_S^0 \to \pi^+\pi^- \to \mu^+\mu^- 2\nu$. Therefore, a K_S^0 mass veto of $m_{\pi\pi} \notin [470, 530]$ MeV is applied to suppress this contribution, where $m_{\pi\pi}$ is the invariant mass of the two vertexed tracks reconstructed under the pion mass hypothesis.

Figure 62 shows hadronic interactions (top row) and photon conversions (bottom row). Both processes result in DVs concentrated at material layers with a leading muon p_{T} mostly below SR thresholds. Both processes peak near $2m_\mu$ and fall steeply above ~300 MeV, becoming negligible at higher $m_{\mu\mu}$. These kinematic properties ensure that the background rate is low in the signal region, below 1–2 GeV.

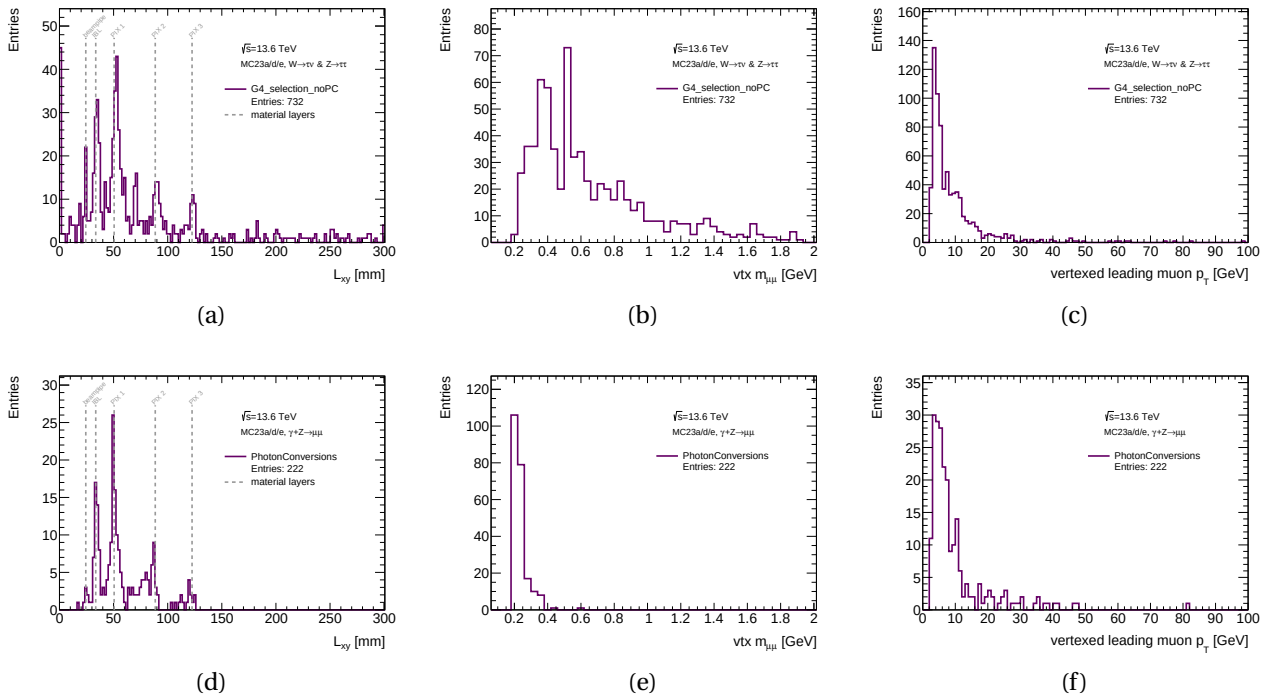


Figure 62. Kinematic distributions of material interaction backgrounds: L_{xy} (left), $m_{\mu\mu}$ (center), and leading muon p_{T} (right). The top row shows the hadronic τ sample ($W \to \tau\nu$ and $Z \to \tau\tau$) with Geant4 selections and photon conversions removed; the bottom row shows the $Z + \gamma$ sample with photon conversion selections applied.

Combining the simulation based study with studies in non-isolated data control regions, the $m_{\mu\mu}$ distribution of material interaction background is found to be smooth above ~300 MeV,

where it can be modeled with a parametric fit. Due to the resonant-like feature at the $2m_\mu$ threshold, the 2μ channel analysis does not consider signal with $m_{\mu\mu}$ below 400 MeV. For 400 MeV $< m_{\mu\mu} <$ 5 GeV, the material interaction background is estimated together with the dominant HF background using a parametric fit. In the 4μ channel, the low overall background rate allows an inclusive ABCD cut-and-count estimate at all $m_{\mu\mu}$, as it does not rely on its shape information. The K_S^0 mass veto is not applied in the SR, as muons from K_S^0 decays-in-flight are largely non-isolated with poor combined-track fit quality, making this contribution negligible after SR selections.

6.5.3 Prompt background

The dominant source of prompt background is low-mass DY production decaying to two muons, as illustrated in Figure 63. The muon pair is often boosted by ISR jets, so at low invariant mass the muons tend to be collimated.

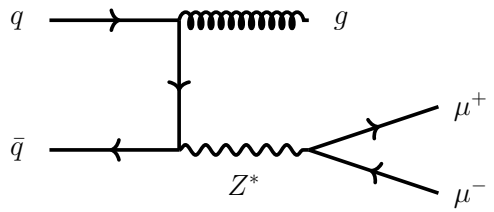


Figure 63. A Feynman diagram for Drell–Yan production decaying to two muons.

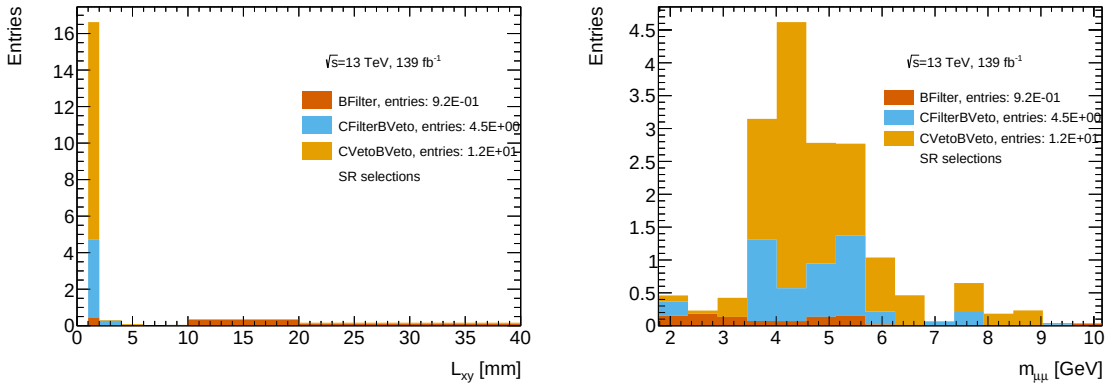
Prompt DY background enters the signal region through two mechanisms: (1) the vertex pulling effect causes prompt dimuons from low-mass DY to be reconstructed at the IBL, mimicking displaced vertices; and (2) prompt dimuons are reconstructed with a small apparent displacement due to finite vertex resolution.

Vertex pulling As discussed in Section 6.3.3, collimated muon pairs can bias the secondary vertex reconstruction, resulting in the vertex position being incorrectly assigned to a shared pixel cluster. A high-statistics single-particle-gun MC sample is used to simulate prompt

dimuon pairs and study the characteristics of pulled vertices. These events are further isolated in data by constructing a prompt-background-enriched CR, defined by inverting the L_{xy} significance requirement ($L_{xy}/\sigma(L_{xy}) < 3$). Requiring $\Delta\phi_{\mu\mu} > 0.005$ is found to be 98% efficient in removing pulled vertices. The residual prompt background due to vertex pulling is small. Since the $m_{\mu\mu}$ distribution from this background is smooth, it is included into the parametric fit background model described in Section 6.6.1.

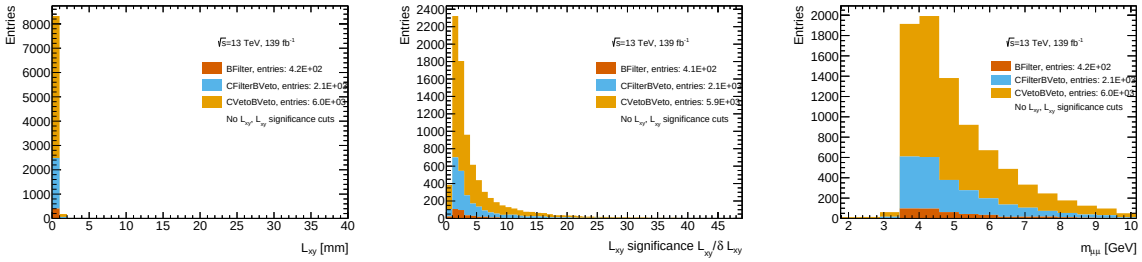
Prompt background without vertex pulling The remaining prompt background not due to vertex pulling is characterized using a low-mass DY simulation of $Z^* \to \mu\mu$ events with 3.6 GeV $< m_{\mu\mu} <$ 10 GeV. Figure 65 shows the SR L_{xy}, L_{xy} significance, and $m_{\mu\mu}$ distributions with the displacement requirements removed, illustrating the effectiveness of the L_{xy}-based selections at rejecting prompt background. Figure 64 shows the prompt DY background yield in the SR scaled to the full Run 2 luminosity.

The prompt background without vertex pulling is negligible above $L_{xy} > 2$ mm, with only ~16 events below $L_{xy} < 1$ mm. In the low-mass region ($m_{\mu\mu} < 5$ GeV), it is small relative to the dominant HF background and is included in the parametric fit. For the high-mass region ($m_{\mu\mu} > 5$ GeV), tightening the L_{xy} requirement to $L_{xy} > 2$ mm makes this background negligible.



(a) (b)

Figure 64. The L_{xy} and dimuon invariant mass distributions in low-mass DY MC with SR selections corresponding to the full Run 2 luminosity. Figure (a) shows the L_{xy} distributions, with an existing L_{xy} requirement at 1 mm. Figure (b) shows the dimuon invariant mass distribution of the selected events.



(a) (b) (c)

Figure 65. The L_{xy}, L_{xy} significance, and dimuon invariant mass distributions in low-mass DY MC with SR selections but removing L_{xy} and L_{xy} significance requirements.

6.5.4 Accidental crossings

The accidental crossing background consists of events with two uncorrelated muon tracks crossing and forming a displaced vertex, mostly from events where a prompt muon from the primary interaction is vertexed with a muon produced in a pile-up interaction in the same event. Since crossing probability is larger at larger opening angles, this background is found to be negligible in the low-mass region, or otherwise smooth in $m_{\mu\mu}$. In the high-mass region, the accidental crossing background is estimated with a data-driven ABCD method, as described in section 6.6.5.

6.5.5 Cosmic muon background

Cosmic muons are a potential source of background when a single cosmic muon is reconstructed as two back-to-back muons with $\Delta\phi \approx \pi$ and $\sum\eta \approx 0$, or when a cosmic muon crosses with a muon from the pp collision. The latter is included in the estimation for accidental crossings background, while the former is expected to be negligible after applying the full signal region selections. Among these selections, the muon pointing angle ensures a two-body decay topology, and a dedicated veto removes back-to-back muons with $|\pi - \Delta\phi| < 0.01$.

6.5.6 Summary of background composition

The background composition is summarized in Table 15, which provides a qualitative expectation for the size of each background in the signal region. The low-mass region ($m_{\mu\mu}$ < 5 GeV) is dominated by HF background. In the high-mass region ($m_{\mu\mu} \geq 5$ GeV), both HF and accidental crossing backgrounds are of similar size, while all other contributions are negligible. This composition motivates the data-driven estimation methods described in Section 6.6.

Table 15. Qualitative comparison of background contributions in the 2μ channel SR, estimated for the combined Run 2+3 luminosity of ~300 fb^{-1}.

Background	$m_{\mu\mu} < 5$ GeV	$m_{\mu\mu} \geq 5$ GeV	Reference
Heavy flavor	$\mathcal{O}(10^5)$	$\mathcal{O}(10)$	Section 6.5.1
Material Interactions	$\mathcal{O}(10^3)$	$\mathcal{O}(1)$	Section 6.5.2
Accidental Crossings	$\mathcal{O}(10^2)$	$\mathcal{O}(10)$	Section 6.5.4
Prompt DY with Vertex Pulling ($L_{xy} > 20$ mm)	$\mathcal{O}(10)$	Negligible	Section 6.5.3
Prompt DY without Vertex Pulling	$\mathcal{O}(10)$	Negligible for $L_{xy} > 2$ mm	Section 6.5.3
Decays in flight	$\mathcal{O}(1)$	Negligible	Section 6.5.2
Cosmics	Negligible	Negligible	Section 6.5.5

6.6 Background Estimation

Two data-driven strategies are used to estimate the background in the signal regions, chosen based on background composition and available statistics.

In the 2-muon low-mass region, where the large HF background is smoothly distributed in $m_{\mu\mu}$, a shape analysis is performed using profile likelihood fits over small, sliding mass windows. The signal and background components are each modeled with separate parametric functions. Control region data, normalized to the signal region via the observed J/ψ yield, serve as background templates for selecting the background functional form and assigning the associated modeling systematic uncertainty.

In the 2-muon high-mass region and the 4-muon channel, where backgrounds are small, a data-driven ABCD cut-and-count method is used. Two uncorrelated variables define four non-overlapping regions, where the background in the signal region A is estimated as:

$$N_A = \frac{N_B}{N_C} \times N_D. \tag{6.5}$$

When control region statistics are low, a Bayesian pseudo-experiment method is used to correctly propagate the statistical uncertainty. Gamma posteriors are placed on each control region yield, and an ensemble of ABCD predictions N'_A is constructed by sampling from these

posteriors. The median of the ensemble is reported as the central estimate, with the 16% and 84% quantiles defining asymmetric 1σ uncertainties.

6.6.1 2-muon channel low-mass region

The low-mass region ($m_{\mu\mu} < 5$ GeV) is dominated by HF background, whose $m_{\mu\mu}$ distribution is a smooth continuum punctuated by known resonances that must be masked.

Mass windows are adjusted when known SM resonances are present, in which case these mass regions are masked, as listed in Table 16. The sizes of the masked mass windows are determined by fitting the known resonances in data with a Double-Sided Crystal Ball (DSCB) function modeling the resonance signal plus a one-degree-of-freedom power-law function modeling the continuum HF background, as shown in Figure 66. The veto regions correspond approximately to $\pm 5\sigma$ of the DSCB fit. The details of these fit functions are explained later in this section.

Table 16. Known SM resonances.

SM resonance	Property	Mass Window
ρ meson	$u\bar{d}$ or $(u\bar{u}-d\bar{d})/\sqrt{2}$, ~ 775 MeV	[700, 900] MeV
ϕ meson	$s\bar{s}$, ~ 1,019 MeV	[900, 1100] MeV
J/ψ meson	$c\bar{c}$, ~ 3,097 MeV	[2800, 3400] MeV
$\psi(2S)$ meson	$c\bar{c}$, ~ 3,686 MeV	[3400, 4000] MeV

Since the dark photon signal is long-lived, and has a small width, the mass resolution is limited by detector reconstruction. Taking advantage of the narrow resonance feature, the background is fitted with a parametric function in small mass windows where the background is relatively smooth.

As a baseline strategy, each fit region is defined to be a mass window $m_{\text{sig}} \pm 5\sigma$, where σ is the signal mass resolution obtained from a fit to the Double-Sided Crystal Ball function, detailed in Section 6.6.2. Overlaps with the resonance veto regions are removed from the mass window,

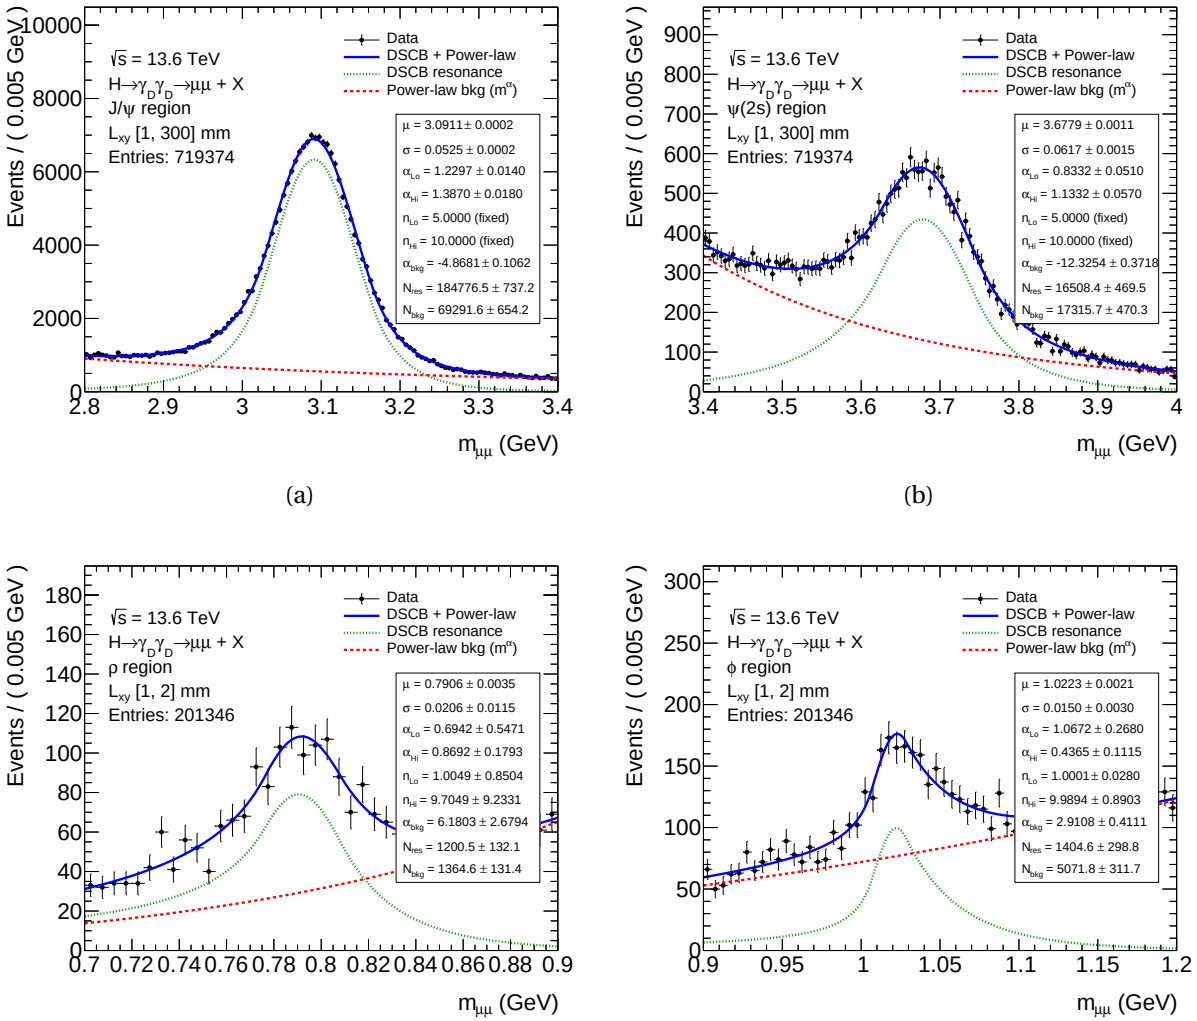


Figure 66. Fits to known SM resonances in the dimuon mass spectrum used to determine the masking windows. The dataset used is the `Fail Isolation Loose` CR data from 2022–2024. Shown are the J/ψ (a), $\psi(2S)$ (b), ρ (c), and ϕ (d) resonances. Each resonance window is fitted with a Double-Sided Crystal Ball signal function plus a power-law background (solid line), with signal and background components shown separately as dashed lines. For the ϕ and ρ resonances, the fits are only performed in the L_{xy} = 1–2 mm region to obtain the best mass resolution.

resulting in asymmetric background side-bands. Signal injection studies will be performed to assess the effectiveness of fitting in such regions, and remove signal points that cannot be constrained due to a lack of background side-bands.

Each signal mass window is further divided into multiple L_{xy} regions to maximize sensitivity. For each signal mass, a simultaneous fit over seven L_{xy} windows is performed, requiring the signal strength to be the same in all L_{xy} regions. To ensure the data statistics at each L_{xy} region, these regions are chosen empirically as:

$L_{xy} \in$ [1,2] mm, [2,5] mm, [5,10] mm, [10,20] mm, [20,40] mm, [40,60] mm, [60,300] mm.

Before unblinding, several CR datasets are used as background templates to define the background model and to estimate the expected sensitivity. The CRs provide a range of plausible background shapes, but an accurate normalization can only be obtained from SR data. As the best available approximation, the background templates are scaled to the expected SR yield using the normalization derived from the vetoed J/ψ mass window. The choice of background templates is described in detail in Section 6.6.3. The full procedure pre-unblinding is defined as follows:

1. Determine the functional form of the fit function.
2. Assign systematic uncertainties from background and signal modeling bias
3. Estimate the expected sensitivity using a background template and the chosen background model.

6.6.2 Signal Modeling

A Double-Sided Crystal Ball (DSCB) function is used to parametrize the dark photon signal, described in Equation (6.6):

$$f_{\text{sig}}(x, f, M, \sigma_{\text{CB}}, \alpha, n, \sigma_G) = \text{DSCB}(x, M_{\text{CB}}, \sigma_{\text{CB}}, \alpha_{\text{low}}, \alpha_{\text{high}}, n_{\text{low}}, n_{\text{high}}), \tag{6.6}$$

where x is the running dimuon invariant mass and other parameters are:

- M_{CB}: mean value of the Crystal Ball function,
- σ_{CB}: width of the Crystal Ball function,
- α_{low}: threshold for the left low-end tail power-law,
- α_{high}: threshold for the right high-end tail power-law,
- n_{low}: power in the left low-end tail power-law,
- n_{high}: power in the right high-end tail power-law.

The fit results of a few selected signal points of 1 mm $c\tau$ are shown in Figure 67. The size of the mass windows is chosen to be roughly $\pm$ 5–7 σ_{CB}. Since n_{low} and n_{high} are highly correlated with the α_{low} and α_{high} parameters, $n_{\mathrm{low}} = 5$ and $n_{\mathrm{high}} = 10$ are fixed for signal masses above 1 GeV in the fit for better convergence. In addition, the signal model is derived from fitting the $m_{\mu\mu}$ distribution in the full L_{xy} range, i.e. 1–300 mm. This is due to a lack of statistics in individual L_{xy} windows.

6.6.2.1 Interpolation of signal parametrization

Since signal samples are only generated at discrete mass and lifetime points, the DSCB parameters are interpolated across all masses and lifetimes. For each parameter, a smooth function of mass and lifetime is derived using fitted values from available samples. Figure 68 shows the interpolations of all signal parameters using linear functions. Since there are no meaningful variations of these parameters among different lifetimes, the signal parameters are chosen from the $c\tau = 1$ mm signal models for all lifetimes at a given signal mass. As a result, the mass window is defined as $m_{\mathrm{sig}} \pm 5\sigma$, with σ from the DSCB fits implemented as a smooth, mass-dependent function:

$$\sigma(m) = 0.015\, m + 0.005 \quad [\mathrm{GeV}]. \tag{6.7}$$

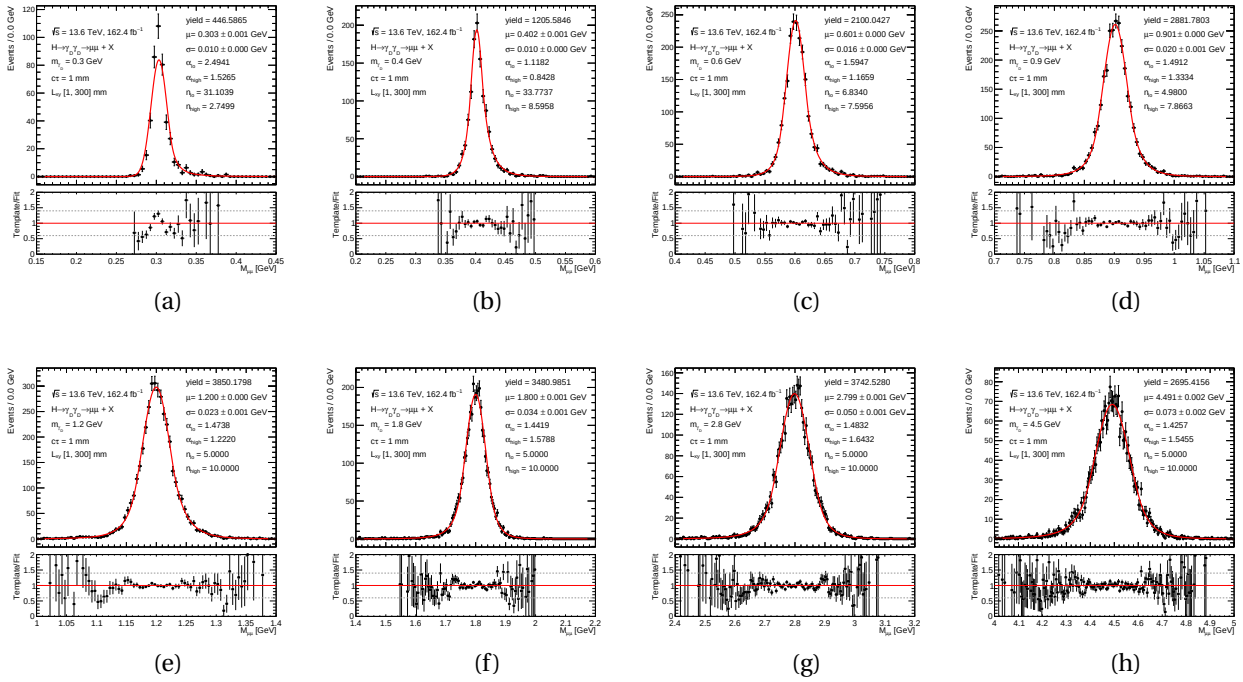


Figure 67. Signal MC fits with DSCB functions, inclusive in L_{xy}. The number of entries corresponds to data 22–24 luminosity with $BR(H \to A'A') = 1\%$.

A sliding-window scan over $m_{\text{sig}} \in [0.4, 5]$ GeV in 10 MeV steps is then performed, with resonance veto overlaps removed. The signal yields used in each fit region also follow interpolations of signal efficiency and the signal L_{xy} distribution across all mass points.

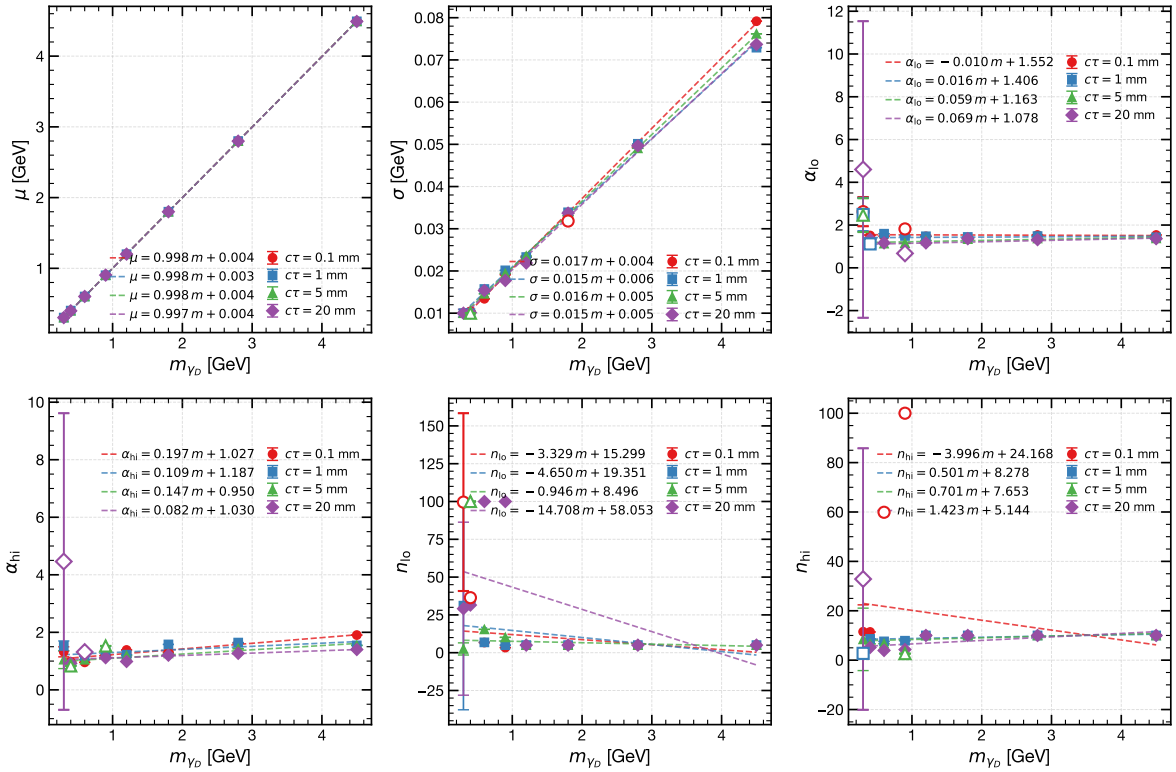


Figure 68. Double-Sided Crystal Ball (DSCB) shape parameters, mean (μ), width (σ), tail parameters (α_{lo}, α_{hi}) and power-law exponents (n_{lo}, n_{hi}), as a function of the dark photon signal mass m_{γ_D}, shown for four signal lifetimes $c\tau = 0.1, 1, 5, 20$ mm. The parameters are extracted from DSCB fits to the $m_{\mu\mu}$ distribution of simulated $H \to \gamma_D\gamma_D \to \mu\mu + X$ signal events at each mass point. Dashed lines show linear fits to the mass dependence, used to parameterize the signal shape as a function of mass. Points excluded from the linear fit by a 2σ iterative clipping procedure are shown as open markers.

6.6.2.2 Signal injection test

Signal injection tests characterize the modeling of the signal. The signal invariant mass shape is described with DSCB functions, as described in Section 6.6.2, and further interpolated

to different signal masses and lifetimes. The injection test evaluates the potential bias introduced by the signal modeling.

The procedure includes injecting a number of signal events (N_{inj}) into a smoothed background template, and comparing N_{inj} to the signal strength obtained from the signal-plus-background fit (N_{fit}). The largest difference between the injected and fitted signal yields for each signal point is assigned as a systematic uncertainty.

6.6.3 Background Templates

Several CR datasets are used as background templates to test the background modeling before unblinding. The CR background templates are normalized to the observed yield in the vetoed J/ψ mass window (2.9–3.2 GeV). A scale factor (SF) is derived as:

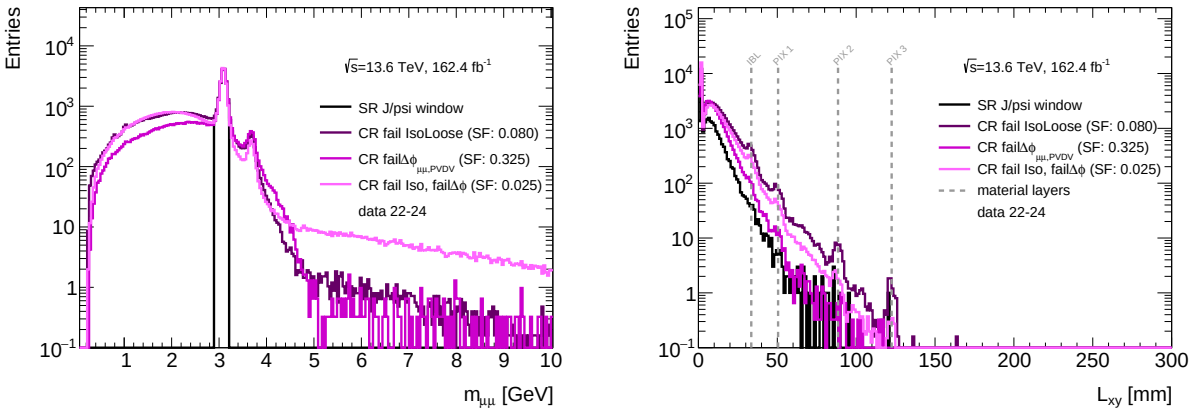
$$\text{SF} = \frac{N(\text{SR selection} | m_{\mu\mu} \in J/\psi \text{ window})}{N(\text{CR selection} | m_{\mu\mu} \in J/\psi \text{ window})}.$$

Figure 69 shows the dimuon invariant mass $m_{\mu\mu}$ and transverse decay length L_{xy} distributions for the SR and scaled CR datasets, where the CR is normalized to the SR in the J/ψ mass window. The SFs used to normalize the data CRs are listed in Table 17.

Table 17. Background templates from data control regions, and the corresponding scale factors (SFs) used to normalize the templates to the SR-like reference yield in the J/ψ peak.

Control Region Template	SF to J/ψ region
`Fail Loose Isolation`	0.0799
`Fail Pointing Angle`	0.3249
`Fail Loose Isolation, fail Pointing Angle`	0.0247

Since the background composition is expected to vary among the different CRs, both the total event yields and the $m_{\mu\mu}$ shapes differ across the three CRs after applying the normalization. This is expected, and does not affect the choice of background model, as the model is designed to be flexible enough to accommodate different shapes.



(a) (b)

Figure 69. Comparison of the dimuon invariant mass $m_{\mu\mu}$ (a) and transverse decay length L_{xy} (b) distributions between the SR J/ψ mass window (2.9–3.2 GeV) and scaled CR datasets, where the CR datasets are scaled to have the same number of events in the J/ψ mass window.

Within each mass window, the data are further divided in L_{xy}, starting from the seven nominal L_{xy} regions defined in Section 6.6.1. Based on the three J/ψ-scaled CR background templates, the following L_{xy} binning choice per signal-mass range is made in Table 18. Adjacent L_{xy} bins are combined at certain mass ranges so that every fit region retains $\gtrsim \mathcal{O}(10)$ expected events, the threshold below which the spurious-signal test of Section 6.6.4.3 is no longer informative due to template statistical fluctuations.

Table 18. L_{xy} regions used in each mass range. Adjacent default L_{xy} bins (Section 6.6.1) are merged where the event yield is expected to fall below $\mathcal{O}(10)$ events per $(m_{\mu\mu}, L_{xy})$ bin. Mass intervals overlapping the SM resonance vetoes (Table 16) are excluded.

m_{sig} range [GeV]	L_{xy} regions [mm]
$0.4 < m_{\text{sig}} < 0.8$	[0,2], [2,10], [10,20], [20,40], [40,300]
$1.0 < m_{\text{sig}} < 2.8$	[0,2], [2,5], [5,10], [10,20], [20,40], [40,60], [60,300]
$m_{\text{sig}} > 4.0$	[0,300]
$0.8 \leq m_{\text{sig}} \leq 1.0$, $2.8 \leq m_{\text{sig}} \leq 4.0$	sidebands of resonance vetoes (Table 16)

6.6.4 Background Modeling

The background in each fit region is modeled using a set of empirical functions. These functions include power laws ("Power") and exponentials with polynomials on index ("Epoly") as defined in Table 19. All of the tested functional forms are summarized in Table 20. All the coefficients are floating in the fit and determined by data.

Table 19. Definitions of empirical functions for the background modeling.

Function	Expression
PowerN	$m_{\mu\mu}^{(a_0+a_1 m_{\mu\mu}+a_2 m_{\mu\mu}^2+...+a_N m_{\mu\mu}^N)}$
EpolyN	$\exp(a_1 m_{\mu\mu}+a_2 m_{\mu\mu}^2+...+a_N m_{\mu\mu}^N)$

Table 20. List of tested empirical functional forms for the background modeling.

Function	Expression	nPar
Epoly1	$exp(a_1 m_{\mu\mu})$	1
Epoly2	$exp(a_1 m_{\mu\mu}+a_2 m_{\mu\mu}^2)$	2
Epoly3	$exp(a_1 m_{\mu\mu}+a_2 m_{\mu\mu}^2+a_3 m_{\mu\mu}^3)$	3
Epoly4	$exp(a_1 m_{\mu\mu}+a_2 m_{\mu\mu}^2+a_3 m_{\mu\mu}^3+a_4 m_{\mu\mu}^4)$	4
Power0	$pow(m_{\mu\mu}, a_0)$	1
Power1	$pow(m_{\mu\mu}, a_0+a_1 m_{\mu\mu})$	2
Power2	$pow(m_{\mu\mu}, a_0+a_1 m_{\mu\mu}+a_2 m_{\mu\mu}^2)$	3
Power3	$pow(m_{\mu\mu}, a_0+a_1 m_{\mu\mu}+a_2 m_{\mu\mu}^2+a_3 m_{\mu\mu}^3)$	4

6.6.4.1 Background model criteria

A background fit function is accepted if it satisfies the following criteria, evaluated using the following tests:

– **Goodness-of-fit χ^2 test**: The χ^2 probability of the background-only fit exceeds 1% on all background templates (see Section 6.6.4.2).

– **Spurious signal test on fit bias**: $S/\delta S < 50\%$, i.e. the spurious signal is less than 50% of the statistical uncertainty on the fitted signal yield, assuming background statistical errors corresponding to the Run 2 or Run 3 luminosity (see Section 6.6.4.3).

If multiple functions pass these criteria, the one with the fewest degrees of freedom is selected, or the one with the smallest spurious signal S if several remain. This strategy follows the approach outlined in Ref. [56].

6.6.4.2 Goodness of the background-only fit

The goodness-of-fit of the background-only fit is evaluated with a χ^2 test. The fit function used for background modeling must have a χ^2 probability greater than 1% for the background-only fits.

Examples of background-only fit results using selected fit functions are shown in Figures 70 to 72 for three representative mass windows (0.5, 1.8, and 4.5 GeV). The following background-only fits are performed in all L_{xy} windows using the `Fail Loose Isolation` CR dataset scaled to expected SR yields. These results show that the selected parametric functions are able to describe the background distribution in all fit regions.

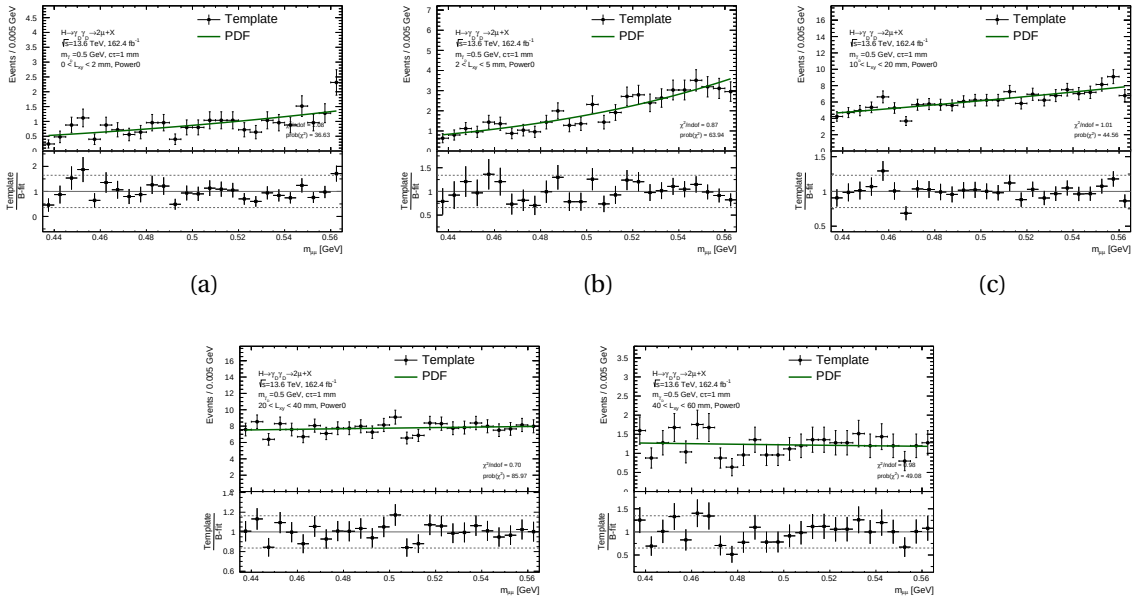
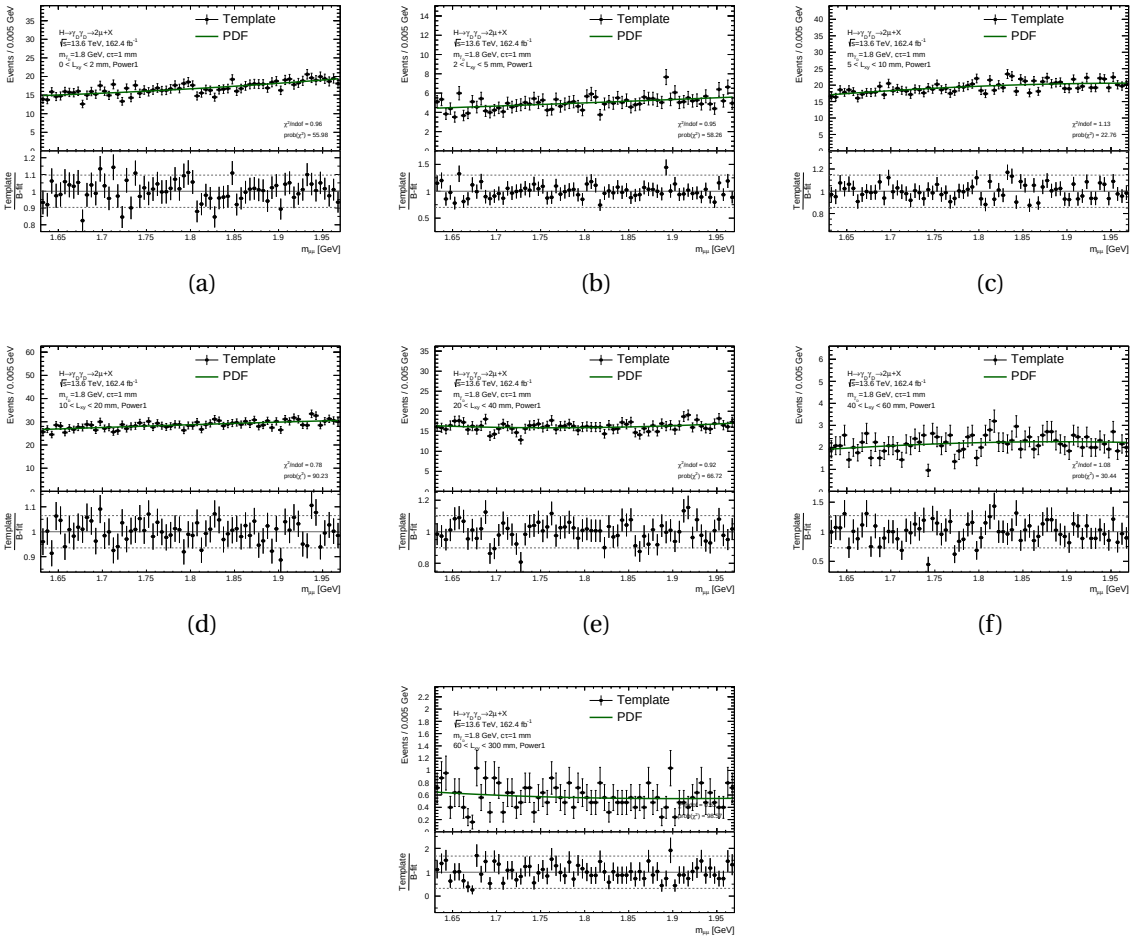


Figure 70. Background-only fits for the 0.5 GeV mass window using the `Power0` fit function, for the `Fail Isolation Loose` dataset. Shown are $L_{xy} \in [0,2]$ mm (a), $[2,10]$ mm (b), $[10,20]$ mm (c), $[20,40]$ mm (d), and $[40,300]$ mm (e).

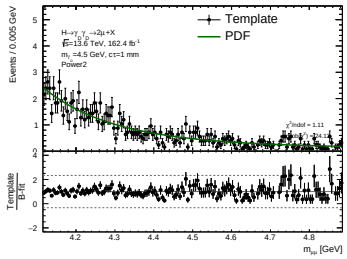


(a) (b) (c)

(d) (e) (f)

(g)

Figure 71. Background-only fits for the 1.8 GeV mass window using the `Power1` fit function, for the `Fail Isolation Loose` dataset. Shown are $L_{xy} \in [0,2]$ mm (a), [2,5] mm (b), [5,10] mm (c), [10,20] mm (d), [20,40] mm (e), [40,60] mm (f), and [60,300] mm (g).



(a)

Figure 72. Background-only fits for the 4.5 GeV mass window using the `Power2` fit function, for the `Fail Isolation Loose` dataset. Shown is $L_{xy} \in [0,300]$ mm (a).

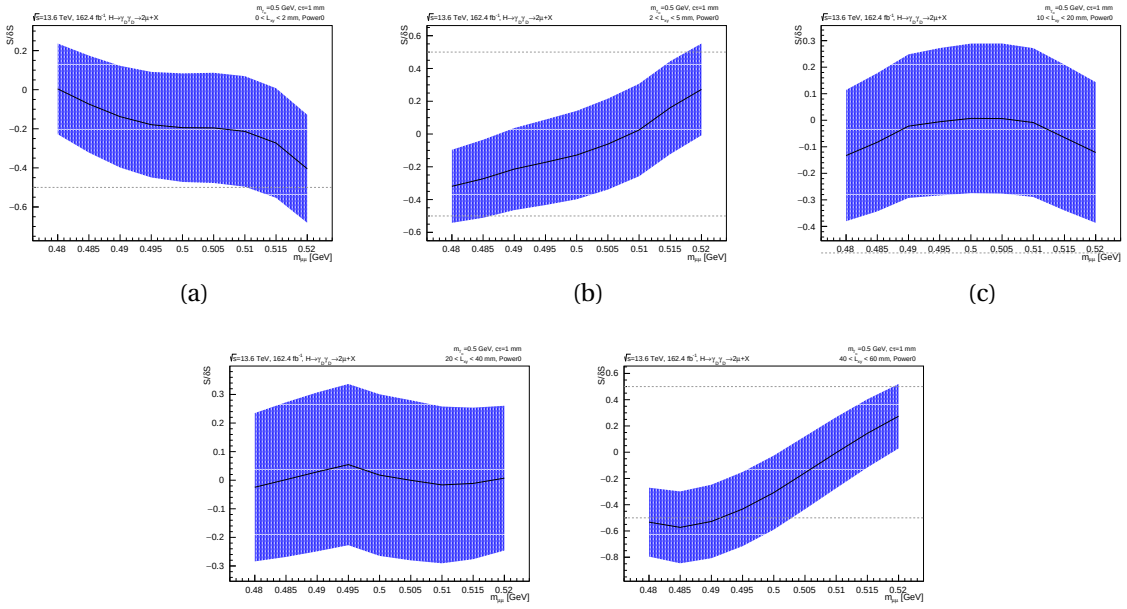
6.6.4.3 Spurious signal test

The choice of functional form for background modeling may introduce a bias on the fitted signal yield, known as *spurious signal* (SS). Spurious signal tests are performed by fitting a background+signal model to a background-only template, with the fitted signal yield defined as the SS size. To suppress bias due to statistical fluctuations of the template, its equivalent luminosity must be significantly larger than that of the data. Control region data are used as background templates here since no high-statistics background MC sample is available.

The spurious signal test serves two purposes: to identify the functional forms with the least background modeling bias, and to provide a systematic uncertainty on the signal yield due to background mis-modeling, as discussed in Section 6.7.1.1. For each candidate function, the spurious signal size S is extracted from the background-only template and compared relative to the statistical uncertainty of the fitted signal (δS), evaluated over the same mass windows around the signal mass. The δS is obtained assuming background statistical errors corresponding to the Run 2 (139 fb^{-1}) or Run 3 (162.4 fb^{-1}) luminosity, using the normalization specified in Section 6.6.3.

Candidates satisfying the criteria in Section 6.6.4.1 are further required to demonstrate flexibility, defined as the ability to satisfy the SS criterion across all CR templates and L_{xy} regions at a given mass window. Among these, the function with the fewest degrees of freedom is chosen.

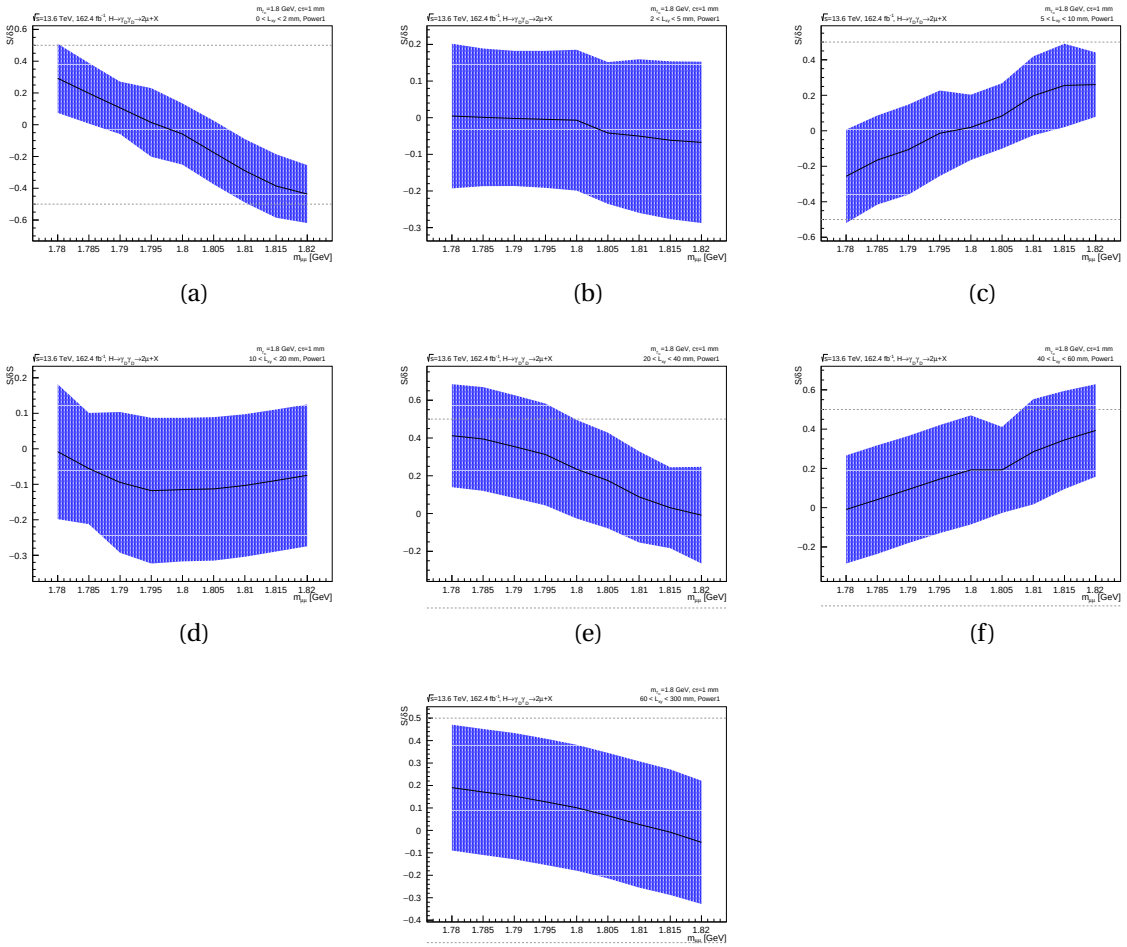
Spurious signal tests are performed for a sliding mass window scan from 0.4 to 5.7 GeV in 100 MeV steps, excluding known resonance windows as defined in Table 16. The $S/\delta S$ distributions for the `Fail Isolation Loose` template are shown in Figures 73 to 75 for representative signal masses of m_{γ_D} = 0.5, 1.8, and 4.5 GeV. While the $S/\delta S$ distributions are shown over a ±200 MeV scan window around the signal mass, the spurious signal uncertainty is determined based on the SS size at the signal mass, $S(m)$; the scan window is shown only to illustrate the distribution shape.



(a) (b) (c)

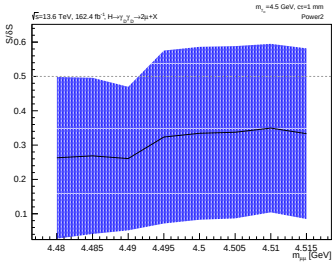
(d) (e)

Figure 73. $S/\delta S$ distributions for the 0.5 GeV mass window using the `Power0` fit function, for the `Fail Isolation Loose` dataset scaled to expected SR yields. Shown are $L_{xy} \in [0,2]$ mm (a), $[2,10]$ mm (b), $[10,20]$ mm (c), $[20,40]$ mm (d), and $[40,300]$ mm (e).



(a) (b) (c) (d) (e) (f) (g)

Figure 74. $S/\delta S$ distributions for the 1.8 GeV mass window using the `Power1` fit function, for the `Fail Isolation Loose` dataset scaled to expected SR yields. Shown are $L_{xy} \in [0,2]$ mm (a), $[2,5]$ mm (b), $[5,10]$ mm (c), $[10,20]$ mm (d), $[20,40]$ mm (e), $[40,60]$ mm (f), and $[60,300]$ mm (g).



(a)

Figure 75. $S/\delta S$ distribution for the 4.5 GeV mass window using the `Power2` fit function, for the `Fail Isolation Loose` dataset scaled to expected SR yields. Shown is $L_{xy} \in [0,300]$ mm (a).

Based on the results of the spurious signal test, the selected background PDFs are summarized in Table 21. A small number of fit regions fail the spurious signal criteria with the selected function. These failures are attributed to limited statistics of the background template, which results in higher spurious signal through statistical fluctuations rather than background mismodeling. The selected fit functions are therefore retained in these regions.

Table 21. Selected background PDF for each mass range in the 2-muon channel low-mass parametric fit.

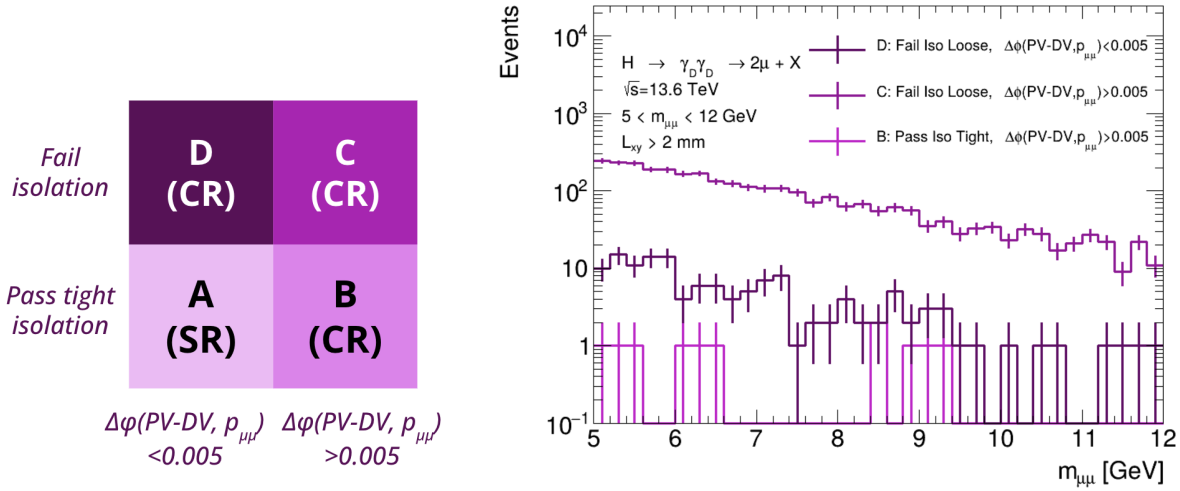
Mass range [GeV]	Selected PDF
[0.4, 0.6]	`Power0`
[1.2, 2.7]	`Power1`
[4.1, 5.0]	`Power2`
[5.1, 5.7]	`Power0`

6.6.5 2-muon channel high-mass region

The high-mass region ($m_{\mu\mu} > 5$ GeV) has two primary background contributions: accidental crossings of uncorrelated muon pairs, and HF decays where the two muons originate from different B mesons, producing a higher $m_{\mu\mu}$ than single-hadron decays. The ABCD plane uses muon isolation and the pointing angle $\Delta\phi(PV\text{–}DV, \vec{p}_{\mu\mu})$ as the two discriminants, as illustrated in Figure 76a.

The ABCD estimate (Equation (6.5)) is evaluated in a 2 GeV mass window around the targeted signal mass using the Bayesian pseudo-experiment method described above. The expected background yields in the signal region are shown in Table 22. The signal contamination is taken into account explicitly in the likelihood-based ABCD estimation, described in Section 6.8.2.

The ABCD method is validated in two validation regions: an intermediate isolation region, defined by events in which both muons fail the tight isolation requirement but satisfy a looser criterion, and a same-sign region, defined by requiring muon pairs carry the same charge. In



(a) (b)

Figure 76. The ABCD method used for background estimation in the high-mass region. Illustration of the ABCD plane defined by muon isolation and the pointing angle $\Delta\phi(PV–DV, \vec{p}_{\mu\mu})$ (a). Dimuon invariant mass distributions in data in control regions B, C, and D (b).

Table 22. Observed events in the control regions and estimated expected events in the signal region A for the Run 3 dataset corresponding to 174.8 fb^{-1}.

Region	Events	
	$6 < m_{\mu\mu} < 8$ GeV	$9 < m_{\mu\mu} < 11$ GeV
B: Pass Isolation, $\Delta\phi(PV-DV, p_{\mu\mu}) > 0.005$	3 ± 1.6	2 ± 1.4
C: Fail Isolation, $\Delta\phi(PV-DV, p_{\mu\mu}) > 0.005$	1166 ± 34	291 ± 17
D: Fail Isolation, $\Delta\phi(PV-DV, p_{\mu\mu}) < 0.005$	45 ± 6.7	11 ± 3.3
$\Rightarrow$ A: Pass Isolation, $\Delta\phi(PV-DV, p_{\mu\mu}) < 0.005$	$0.122^{+0.085}_{-0.057}$	$0.082^{+0.079}_{-0.045}$

both regions, fewer than one event is expected, and zero events are observed. The agreement between expectation and observation indicates closure, while being limited by the available statistical precision.

6.6.6 4-muon data-driven background estimation

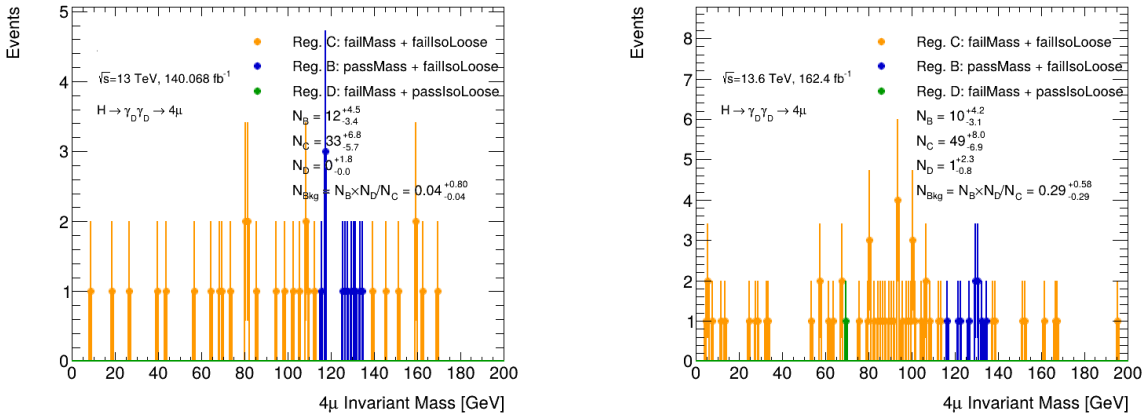
As in the 2-muon channel, the dominant background in the 4-muon channel comes from HF decays and is estimated using a data-driven ABCD method inclusive in L_{xy}. The ABCD plane is constructed from two uncorrelated variables: muon isolation and the $m_{4\mu}$ mass window. The four ABCD regions are defined by the pass/fail of each axis, with Region A (isolated muons, $m_{4\mu}$ in the Higgs window) as the signal region.

The two ABCD variables are shown to be uncorrelated by profiling the isolation variable in slices of $m_{4\mu}$ in data. Both $E_T^{\text{topocone20}}/p_T$ and $p_T^{\text{varcone30}}/p_T$ are flat for $m_{4\mu} > 40$ GeV, which validates the ABCD assumption in this region. A requirement of $m_{4\mu} > 40$ GeV is therefore imposed for the ABCD regions.

The background yield is estimated using the ABCD formula (Equation (6.5)), with the Bayesian pseudo-experiment method applied. The results are shown in Table 23 and Figure 77. Signal contamination in the CRs is explicitly accounted for in the statistical model.

Table 23. Background estimates in the 4μ signal region from the ABCD method with Bayesian pseudo-experiment sampling. N_{bkg} is defined as the median of the pseudo-experiment ensemble, and the uncertainties correspond to the 16% and 84% intervals.

Run	N_{bkg}
Run 3	$0.29^{+0.58}_{-0.29}$
Run 2	$0.04^{+0.80}_{-0.04}$



(a) Run 2: Observed event counts in the ABCD control regions.

(b) Run 3: Observed event counts in the ABCD control regions.

Figure 77. Summary of ABCD background estimates for the 4μ channel. Observed event counts in the ABCD control regions for Run 2 and Run 3.

This ABCD method is validated for closure using a validation region (VR) orthogonal to both the signal and control regions. The VR is constructed by inverting the linear cuts in the 2D plane of $m_{\mu\mu}^{\text{Leading } p_T}$ versus $m_{\mu\mu}^{\text{Sub-leading } p_T}$, as defined in Equation (6.4), and with relaxed event selections to increase statistics. In the SR-analogous VR, (114 GeV $\leq m_{4\mu} \leq$ 136 GeV), $N_{bkg} = 0.92^{+0.80}_{-0.53}$ is predicted, while zero events are observed.

6.7 Systematic Uncertainties

6.7.1 Uncertainties on background estimation methods

The background estimation in this analysis is fully data-driven. Systematic uncertainties are derived for each estimation method and are described in the subsections below.

6.7.1.1 Spurious signal uncertainty

In the low-mass parametric fit region, the systematic uncertainty on the signal yield due to background modeling bias is derived from the spurious signal test described in Section 6.6.1. It

is implemented as

$$N_{\text{signal}}(m) = \sigma_{\text{signal}} \mathscr{L}(\mathscr{A} \times \epsilon) + S_{\text{modeling}}(m)\, \theta_{\text{modeling}}, \tag{6.8}$$

where S_{modeling} is the spurious signal size derived from the background templates, and θ_{modeling} is a nuisance parameter constrained by an external Gaussian with mean 0 and width 1.

It is worth noting that the normalization of the background templates directly determines both the size of the spurious signal uncertainty and the reference statistical uncertainty. As discussed in Section 6.6.3, the normalization derived from the J/ψ mass window is not expected to be accurate, as confirmed by simulation. It is therefore important to understand how sensitive the spurious signal test results are to the choice of normalization.

Figure 78 illustrates this dependence for the $m_{\gamma_D} = 1.8$ GeV signal mass window and $L_{xy} \in [1,2]$ mm region, using the `Power0` function. As shown in Figure 78a, the absolute spurious signal size S grows with the template normalization, as expected from the larger statistical power of a higher-yield background. However, Figure 78b shows that the relative spurious signal size $S/N_{\text{background}}$ remains stable across normalizations, suggesting that the spurious signal scales proportionally with the background yield.

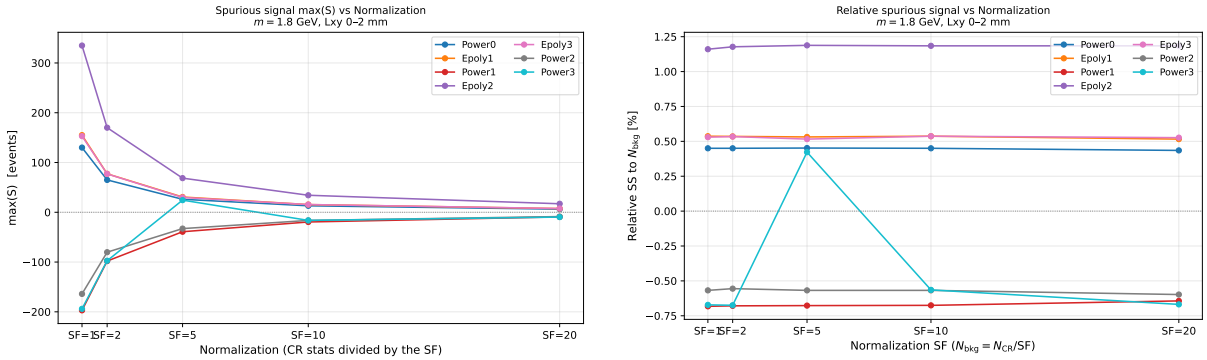
This motivates the use of a *relative spurious signal uncertainty*: rather than quoting an absolute spurious signal size that depends on the (uncertain) template normalization, the spurious signal is extracted relative to the background yield, with the background yield determined directly from the likelihood fit. The relative spurious signal size extracted from the SS test is defined as

$$k(m) = S(m)/N_{\text{background template}}, \tag{6.9}$$

such that the spurious signal uncertainty is implemented as

$$N_{\text{signal}}(m) = \sigma_{\text{signal}} \mathscr{L}(\mathscr{A} \times \epsilon) + k_{\text{modeling}}(m)\, N_{\text{background}}\, \theta_{\text{modeling}}, \tag{6.10}$$

where $N_{\text{background}}$ is the background yield from the fit, and its product with k gives the absolute spurious signal size entering as a systematic uncertainty.



(a) (b)

Figure 78. Spurious signal test results for the $m_{\gamma_D} = 1.8$ GeV signal mass window and $L_{xy} \in [1,2]$ mm region as a function of the template normalization, using the `Power0` function. (a) Absolute spurious signal size S as a function of normalization. (b) Relative spurious signal size $S/N_{\text{background}}$ as a function of normalization.

Preliminary results of the spurious signal uncertainty from all three CR templates are shown in Tables 24 to 26. Both absolute spurious signal size, $S(m)$, and the relative spurious signal size k are shown. While different L_{xy} window choices are used in different CR templates based on their statistics, the fit function choice is made such that it passes the selection criteria in all L_{xy} regions and for all templates.

The expected limits shown in Section 6.8 use $S(m)$ derived directly from the `Fail Isolation Loose` background template. In the final results, the most conservative spurious signal uncertainty across all background templates is used, together with the relative spurious signal size implementation $k(m)$.

Table 24. Spurious signal uncertainty derived from `Fail Isolation Loose` CR background template

m_{γ_D} [MeV]	L_{xy} [mm]	PDF	$S(m)$	N_{B}	$k(m) = S(m)/N_{\mathrm{B}}$
500	0–2	`Power0`	−0.757	23.3	−0.0325
500	2–10	`Power0`	−0.739	50.4	−0.0147
500	10–20	`Power0`	0.0742	161	4.598×10^{-4}
500	20–40	`Power0`	0.211	202	1.045×10^{-3}
500	40–300	`Power0`	−1.41	31.9	−0.0442
1800	0–2	`Power1`	−2.34	1.148×10^{3}	-2.039×10^{-3}
1800	2–5	`Power1`	−0.149	340	-4.379×10^{-4}
1800	5–10	`Power1`	0.889	1.322×10^{3}	6.727×10^{-4}
1800	10–20	`Power1`	−5.79	1.956×10^{3}	-2.96×10^{-3}
1800	20–40	`Power1`	6.77	1.103×10^{3}	6.139×10^{-3}
1800	40–60	`Power1`	1.98	148	0.0134
1800	60–300	`Power1`	0.506	38.8	0.013
4500	0–300	`Power2`	3.02	110	0.0274
5500	0–300	`Power0`	1.84	26.5	0.0695

Table 25. Spurious signal uncertainty derived from `Fail Isolation Loose && Fail Pointing Angle` CR background template

m_{γ_D} [MeV]	L_{xy} [mm]	PDF	$S(m)$	N_{B}	$k(m) = S(m)/N_{\mathrm{B}}$
500	0–2	`Power0`	2.13×10^{-14}	12.5	1.701×10^{-15}
500	2–10	`Power0`	−1.17	383	-3.059×10^{-3}
500	10–20	`Power0`	1.53	171	8.924×10^{-3}
500	20–40	`Power0`	−0.994	136	-7.282×10^{-3}
500	40–300	`Power0`	−0.93	15.8	−0.0589
1800	0–2	`Power1`	−4.99	1.945×10^{3}	-2.566×10^{-3}
1800	2–5	`Power1`	8.55	418	0.0205
1800	5–10	`Power1`	−5.16	1.386×10^{3}	-3.724×10^{-3}
1800	10–20	`Power1`	−7.21	1.843×10^{3}	-3.911×10^{-3}
1800	20–40	`Power1`	−1.99	896	-2.222×10^{-3}
1800	40–60	`Power1`	−1.64	105	−0.0156
1800	60–300	`Power1`	−0.296	20.2	−0.0147
4500	0–300	`Power2`	4.22	235	0.0179
5500	0–300	`Power0`	0.146	180	8.117×10^{-4}

Table 26. Spurious signal uncertainty derived from `Fail Pointing Angle` CR background template

m_{γ_D} [MeV]	L_{xy} [mm]	PDF	$S(m)$	N_{B}	$k(m) = S(m)/N_{\mathrm{B}}$
500	0–2	`Power0`	5.74	68.6	0.0837
500	2–10	`Power0`	−0.64	37.4	−0.0171
500	10–20	`Power0`	−12.1	74.4	−0.163
500	20–40	`Power0`	−6.42	39.3	−0.163
500	40–300	`Power0`	1.92	2.6	0.738
1800	0–2	`Power1`	6.06	774	7.835×10^{-3}
1800	2–5	`Power1`	−3.95	439	-8.996×10^{-3}
1800	5–10	`Power1`	−12.6	1.078×10^{3}	−0.0117
1800	10–20	`Power1`	0.0157	1.137×10^{3}	1.381×10^{-5}
1800	20–40	`Power1`	10.1	382	0.0264
1800	40–60	`Power1`	2.56	32.2	0.0796
1800	60–300	`Power1`	0.932	6.83	0.137
4500	0–300	`Power2`	15.7	269	0.0583
5500	0–300	`Power0`	2.13×10^{-14}	10.4	2.048×10^{-15}

6.7.1.2 ABCD estimation uncertainty

Uncertainties on background estimation using ABCD methods are due to correlations in ABCD variables and the coverage of all background compositions.

In the 2-muon channel high-mass region, a systematic uncertainty is taken from the maximum non-closure from performing the ABCD estimation with varying control region boundaries. As a result, a 25% systematic uncertainty is applied on all of the background predictions in the high-mass region.

In the 4-muon channel, systematics are assigned from a similar closure test where the 4-muon mass window is shifted and resized by ±2.2 GeV. The resulting systematic uncertainty is around 20%.

6.7.2 Uncertainties on the signal process

6.7.2.1 Experimental uncertainties

Luminosity The uncertainty in the integrated luminosity is 0.83% for the 2015–2018 dataset [57] and 1.9% for the 2022–2024 dataset [58]. These are derived following the methodology described in Ref. [57], using the LUCID-2 detector [59] as the primary luminosity measurement, complemented by Inner Detector and calorimeter information.

Pile-up reweighting Simulated events are reweighted to match the pile-up distribution observed in data. The associated uncertainty accounts for imperfect knowledge of the inelastic cross section and variations in data-taking conditions, and is propagated to the signal yield for all Run 2 and Run 3 samples.

Tracking and vertexing Uncertainties in track reconstruction and vertexing affect the signal efficiency. Standard tracking efficiency uncertainties are derived from $Z \to \mu\mu$ tag-and-probe measurements and include contributions from detector description, physics modeling, and pile-up. For Large Radius Tracking (LRT), efficiency uncertainties are estimated using $K^0_S \to \pi^+\pi^-$ tag-and-probe studies. Their applicability to highly displaced tracks is assessed and propagated to the signal yield. The secondary vertex reconstruction uncertainty is evaluated by smearing track parameters within their resolution and measuring the resulting variation in signal efficiency, covering potential mis-modeling of the vertex fit quality and transverse decay length resolution.

Muon-related uncertainties Muon reconstruction, calibration, and trigger effects are taken into account. Reconstruction, isolation, and momentum scale uncertainties for standard and LRT muons are applied following ATLAS recommendations. Zero-pixel-hit muons are primarily covered by standalone muon uncertainties, with an additional small contribution

assigned for track–segment matching efficiency. The modeling is validated in a control region enriched in such muons from material-induced resonances. Muon trigger efficiency scale factor uncertainties are propagated to the signal yield.

Primary vertex assignment Displaced signal can lead to mis-assignment of the primary vertex (PV), as displaced tracks may not contribute significantly to the $\sum p_T^2$ ranking. Since the PV position enters the computation of track-based isolation, the resulting uncertainty in isolation efficiency is evaluated from simulation and propagated to the signal yield.

6.7.2.2 Theoretical and modeling uncertainties

Theoretical uncertainties on the signal yield arise from parton distribution functions (PDFs), renormalization and factorization scale variations, and the strong coupling constant α_s. Scale uncertainties are evaluated by varying μ_R and μ_F independently by factors of 0.5 and 2.0, and taking the envelope of the resulting variations. PDF uncertainties are computed using the full eigenvector set of the nominal PDF (NNPDF3.0 NLO), while the α_s uncertainty is evaluated by comparing results obtained with $\alpha_s = 0.117$ and $\alpha_s = 0.119$.

6.8 Statistical Analysis and Expected Limits

Statistical analysis is performed separately in the 2-muon channel low-mass region, 2-muon channel high-mass region and 4-muon channel. Preliminary expected limits on the HAHM model parameters are presented here based on the background estimation results and signal simulations. Statistical uncertainties and background systematic uncertainties are included in the statistical analysis. The signal systematic uncertainties have not yet been fully evaluated at the time of writing, therefore placeholder values are used as a conservative estimation.

While the analysis procedure is identical for both the Run 2 and Run 3 datasets, only expected limits derived from the partial Run 3 dataset are shown, corresponding to an

integrated luminosity of 162.4 fb^{-1}. The full Run 2 and partial Run 3 results will be included in the full unblinded analysis.

6.8.1 2-muon channel: low-mass region

The 2-muon channel low-mass region takes a parametric fit approach for background estimation, as described in Section 6.6.1. The signal strength is extracted from a maximum likelihood fit to the reconstructed $m_{\mu\mu}$ distribution in sliding windows. Confidence intervals are constructed based on the profile-likelihood-ratio test statistic [60]. The systematic uncertainties listed in Section 6.7 are implemented in the fit as nuisance parameters constrained by additional Gaussian or log-normal likelihood terms.

In the likelihood, the signal yield N_S is parameterized in terms of the signal strength (μ) times Higgs to dark photon decay branching ratio as

$$N_S = \sigma_{\text{ggF}} \times \mu \times \text{BR}(H \to \gamma_D\gamma_D) \times 2 \times \text{BR}_{\mu\mu} \times (1 - \text{BR}_{\mu\mu}) \times \mathscr{L} \times \mathscr{A}, \tag{6.11}$$

where

- σ_{ggF} is the production cross section of the Higgs boson via gluon–gluon fusion (ggF),
- μ is the parameter of interest of the fit, $\text{BR}(H \to \gamma_D\gamma_D)$ is fixed at 1% for simplicity,
- $\text{BR}_{\mu\mu} = \text{BR}(\gamma_D \to \mu\mu)$ is the dark photon to muons decay branching ratio, which is determined entirely by dark photon mass,
- $\mathscr{L}$ is the integrated luminosity, and
- $\mathscr{A}$ the acceptance defined as the fraction of signal events being reconstructed and passing the event selections described in Section 6.4.

The best-fit value of the signal-strength parameter is used to set expected upper limits on the $\text{BR}(H \to \gamma_D\gamma_D)$ at a given dark photon mass and lifetime. Expected limits for several signal

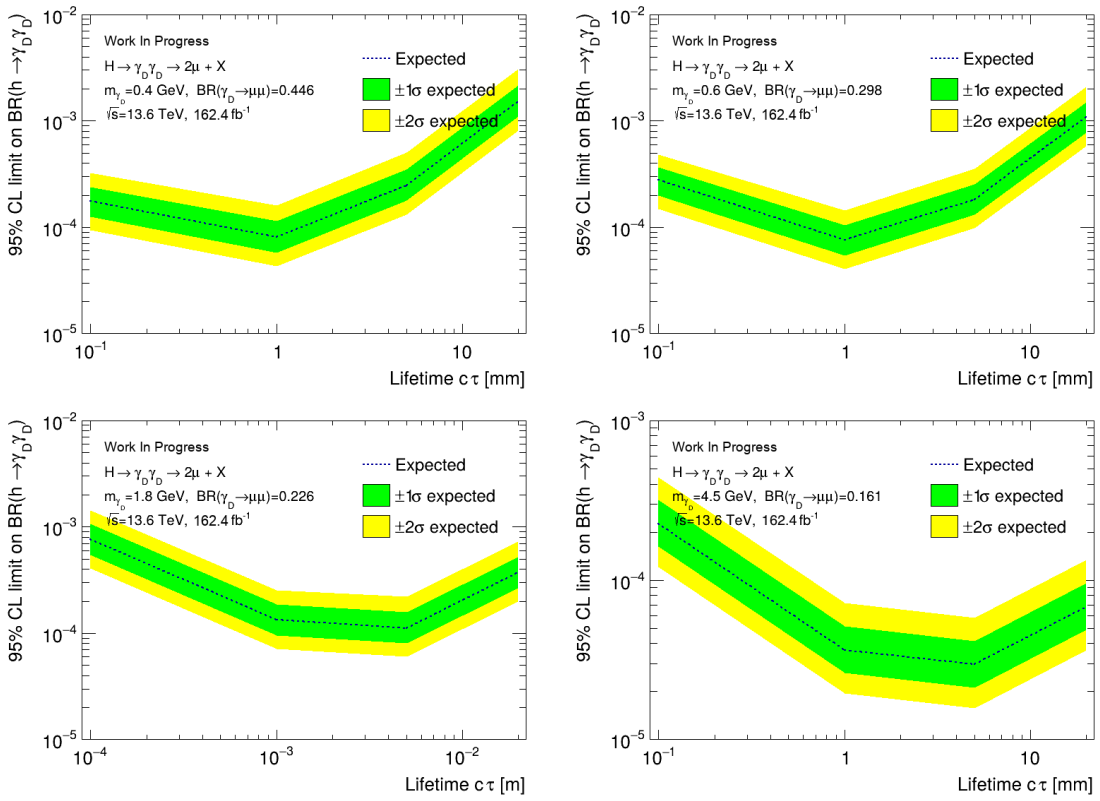


Figure 79. Expected 95% CL upper limits on the branching ratio BR($H \to \gamma_D\gamma_D$) as a function of the proper decay length $c\tau$, for dark photon masses m_{γ_D} = 0.4, 0.6, 1.8, and 4.5 GeV in the 2-muon channel low-mass region. The expected limits are derived from the partial Run 3 dataset ($\sqrt{s}$ = 13.6 TeV, 162.4 fb^{-1}) under the background-only hypothesis. The solid line shows the median expected limit, while the green and yellow bands indicate the $\pm1\sigma$ and $\pm2\sigma$ uncertainty intervals, respectively.

points are shown in Figure 79. An interpolation of signal models and yields as a function of signal mass is performed. Results from a sliding mass window fit over the full mass range for signal with $c\tau$ = 1 mm are shown in Figure 80. Mass windows overlapping with known SM resonances (ρ, ϕ, J/ψ, and $\psi(2S)$) are excluded, as defined in Table 16. The expected upper limits on BR($H \to \gamma_D\gamma_D$) are of the order of 10^{-3}–10^{-4}.

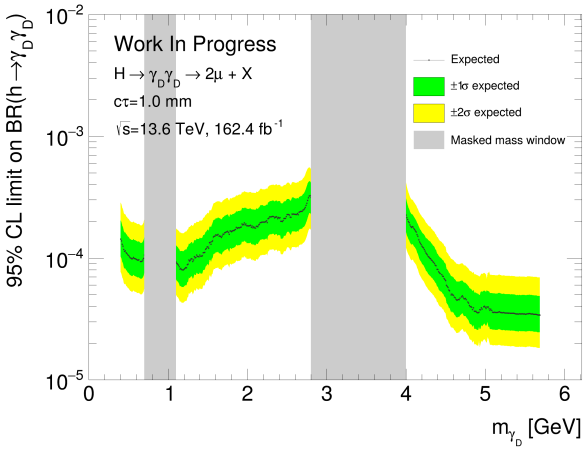


Figure 80. Expected 95% CL upper limits on the branching ratio BR($H \to \gamma_D\gamma_D$) as a function of the dark photon mass m_{γ_D} for a fixed proper decay length $c\tau = 1$ mm. The expected limits are obtained from a sliding mass window parametric-fit across signal masses using the partial Run 3 dataset ($\sqrt{s} = 13.6$ TeV, 162.4 fb^{-1}) under the background-only hypothesis. The solid line shows the median expected limit, while the green and yellow bands indicate the $\pm1\sigma$ and $\pm2\sigma$ uncertainty intervals, respectively. The gray shaded regions indicate excluded mass windows due to known SM resonances (ρ, ϕ, J/ψ, and $\psi(2S)$), as defined in Table 16.

6.8.2 2-muon channel: high-mass region

In the high-mass region of the 2-muon channel, a cut-and-count analysis is performed, with the background estimation described in Section 6.6.5. Expected 95% CL upper limits on BR($H \to \gamma_D\gamma_D$) are derived using a likelihood-based implementation of the ABCD method, where all four regions are modeled simultaneously as Poisson likelihoods.

Signal contamination in the control regions is accounted for via the corresponding signal yields. The background normalization is constrained by the ABCD relation and treated as a nuisance parameter. Limits are computed using the CL_s method with toy-based likelihood evaluation, rather than the asymptotic approximation, due to the low background statistics.

The statistical analysis is performed including the statistical uncertainty on the background estimates, a 25% systematic uncertainty on the background based on the variations of the ABCD plane, and a 20% placeholder systematic uncertainty on the signal. The resulting expected limits for the $m_{\gamma_D} = 7$ and 10 GeV mass points are shown in Figure 81.

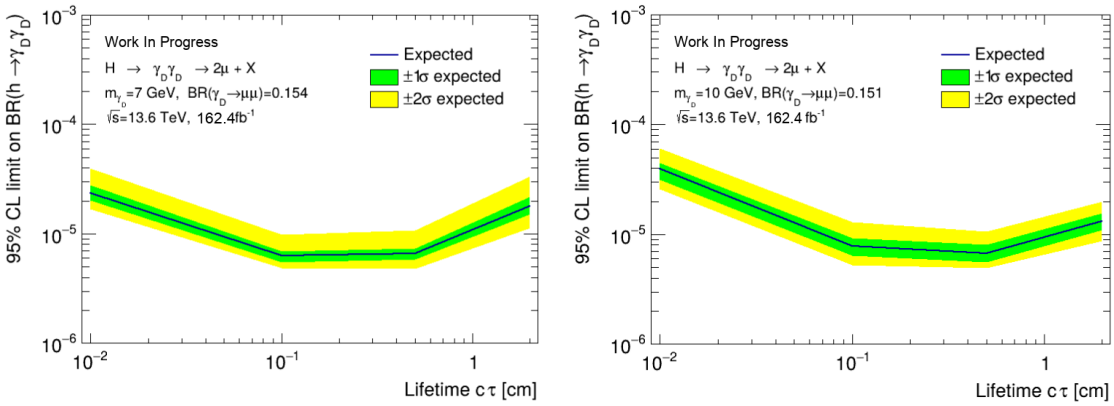
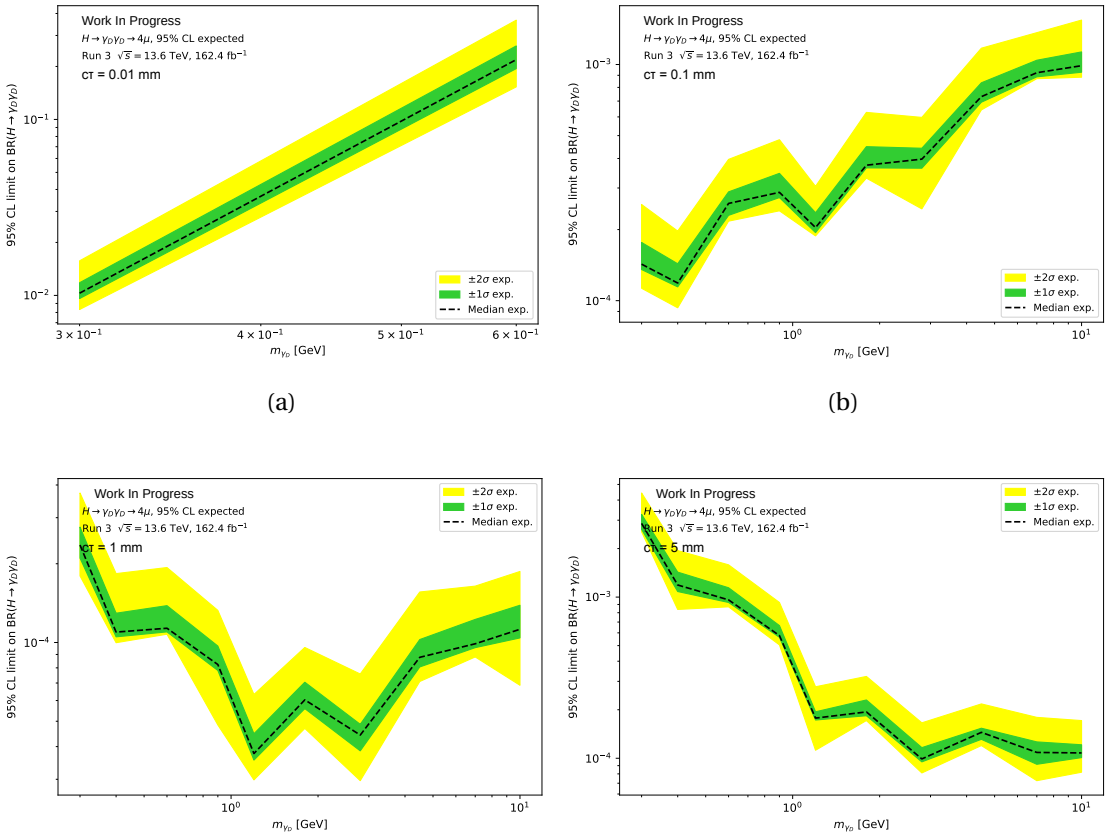


Figure 81. Expected 95% CL upper limits on the branching ratio BR($H \to \gamma_D\gamma_D$) as a function of the proper decay length $c\tau$, shown for dark photon masses $m_{\gamma_D} = 7$ GeV and $m_{\gamma_D} = 10$ GeV in the 2-muon channel high-mass region. The expected limits are derived from the partial Run 3 dataset ($\sqrt{s} = 13.6$ TeV, 162.4 fb^{-1}) under the background-only hypothesis. The solid line shows the median expected limit, while the green and yellow bands indicate the $\pm 1\sigma$ and $\pm 2\sigma$ uncertainty intervals, respectively.

6.8.3 4-muon channel

The 4-muon channel performs a similar cut-and-count analysis as in the 2-muon channel high-mass region. A single invariant mass window is used for all signal masses, as opposed to a sliding window used in the 2-muon channel. The expected upper limits on BR($H \to \gamma_D\gamma_D$) are calculated using the same CL_s method described in Section 6.8.2, which includes signal contamination in the likelihood fit, and uses toy-based likelihood estimation due to the low background statistics.

These preliminary limit estimates do not include any background systematic uncertainties, but do include a placeholder 10% systematic uncertainty on the signal yield. The expected 95% CL upper limits on BR($H \to \gamma_D\gamma_D$) as a function of m_{γ_D} for several fixed $c\tau$ values are shown in Figure 82. The strongest expected limit of BR($H \to \gamma_D\gamma_D$) $\lesssim 3.7 \times 10^{-5}$ is achieved at $m_{\gamma_D} = 1.2$ GeV and $c\tau = 1$ mm.



(a) (b)

(c) (d)

Figure 82. Expected 95% CL upper limits on the branching ratio BR($H \to \gamma_D\gamma_D$) as a function of the dark photon mass m_{γ_D}, shown for four fixed proper decay lengths $c\tau$ = 0.01, 0.1, 1, and 5 mm in the 4-muon channel. The expected limits are derived from the partial Run 3 dataset ($\sqrt{s}$ = 13.6 TeV, 162.4 fb^{-1}) under the background-only hypothesis. The solid line shows the median expected limit, while the green and yellow bands indicate the $\pm 1\sigma$ and $\pm 2\sigma$ uncertainty intervals, respectively.

Chapter 7

Conclusions and Outlook

This dissertation presents the complete analysis strategy and expected sensitivity of the first ATLAS search for low-mass, long-lived dark photons with displaced vertices in the ATLAS silicon tracker, targeting a theoretically motivated region of parameter space with small dark-photon masses or kinetic mixing strengths, which can be uniquely probed by collider experiments. The goal of this work is to expand the understanding of a potential dark sector via a vector boson mediator that interacts with the visible photon through kinetic mixing. This analysis primarily focuses on the scenario where the Higgs boson serves as a portal to the dark sector and decays to a pair of dark photons. The unique role of the Higgs boson makes it particularly sensitive to new physics, as it provides one of only three possible renormalizable interactions with any dark sector extension of the Standard Model.

The analysis targets dark-photon masses from 300 MeV to 10 GeV with kinetic mixing strengths of $\epsilon \sim 10^{-4}$–10^{-6}. The unique displaced dimuon vertex signature is expected to provide improved sensitivity in this region of parameter space compared to previous searches. The expected upper limits on Higgs to dark photon branching ratios reach $\mathrm{BR}(H \to \gamma_D\gamma_D) \sim 10^{-4}$–$10^{-5}$ in the most sensitive regions of m_{γ_D} and $c\tau$, exceeding the sensitivity of previous ATLAS searches targeting the same Higgs portal model via displaced leptons reconstructed in the calorimeter and Muon Spectrometer [61] and prompt leptons [62]. The expected sensitivity is also competitive with CMS results [63, 64], which target a similar signature with dedicated dimuon triggers. These expected limits are anticipated to closely reflect the final result after unblinding the data.

This analysis has established new reconstruction techniques for signatures involving displaced muons and secondary vertices from low-mass resonances in ATLAS. Novel

reconstruction features, most notably vertex pulling, have been characterized, and the dominant background contributions from heavy flavor decays and material interactions have been studied in detail for the first time in the context of low-mass, long-lived particle searches. These developments provide new insights and a technical foundation for future analyses targeting unconventional long-lived signatures, and the improved understanding of detector effects and backgrounds will benefit a broader class of physics analyses.

At the upcoming high-luminosity LHC, the ATLAS Inner Detector will be replaced by a nine-layer all-silicon Inner Tracker, which will provide higher track resolution and an extended reconstruction volume for displaced tracks and vertices. The luminosity upgrade will further provide an unprecedented amount of data, which opens up a new frontier for probing rare phenomena from physics beyond the Standard Model. In particular, a dark sector with feebly interacting mediators to the Standard Model is one of the most exciting avenues at this high-luminosity frontier.

APPENDIX 1

Characterization of Vertex Pulling Effect in ATLAS Secondary Vertex Reconstruction

Section 6.3.3 introduced *vertex pulling*, an effect in which reconstructed secondary vertex positions are biased toward the pixel layers when the vertexed muon tracks are highly collimated and share the same pixel cluster. The main impacts of this effect on the analysis are the material interaction background being indistinguishable from signal events reconstructed at material layers, and an additional source of low-mass Drell–Yan (DY) background when prompt dimuons are reconstructed near the IBL. This appendix studies the effect in detail, characterizing a strong correlation between vertex pulling and dimuon opening angle, which motivates selections that reduce its contribution to the background.

Two simulation samples are used to isolate dimuons affected by vertex pulling. The first is a low-lifetime, low-mass signal sample at $m_{\gamma_D} = 300$ MeV and $c\tau = 0.01$ mm, which provides a convenient way to isolate vertex pulling events. Due to its short lifetime, this signal is expected to be prompt-like, with dark photon decaying close to the primary vertex and following an exponential decay length distribution smeared by the vertex position resolution. The vertex pulling effect causes these prompt dimuons to be reconstructed at anomalously large L_{xy}, with a distribution biased toward the IBL. The low mass of 300 MeV additionally results in a higher fraction of collimated dimuons due to the larger boost, making it well-suited for this study.

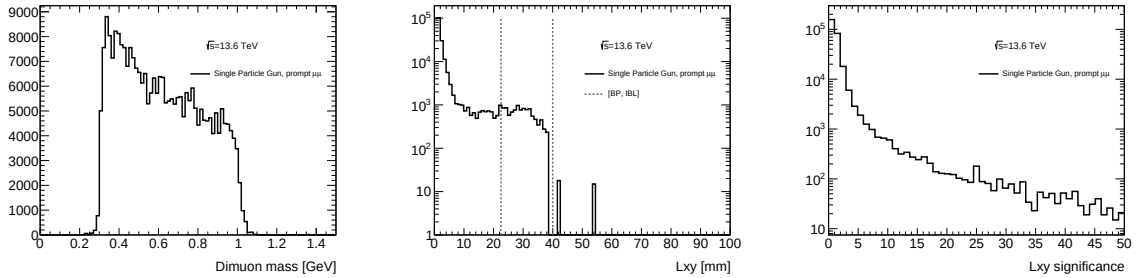
The second sample is a high-statistics single-particle-gun (SPG) Monte Carlo sample used to simulate prompt dimuons and study the properties of pulled vertices. Events are generated containing one BSM scalar particle and six charged pions. The scalar is produced with predefined kinematics, with its momentum and mass sampled within fixed ranges, and is forced to decay promptly into a pair of oppositely charged muons. The six additional pions

provide sufficient charged tracks to allow for primary vertex reconstruction. No pile-up is included in these events. This simulation sample is more inclusive in dimuon mass, covering a wider range of kinematic phase space than the signal sample.

The study is further validated using a control region enriched in prompt Drell–Yan events, which is used to confirm the vertex pulling characterization, evaluate the effect of selections designed to reduce this background, and estimate the residual background yield.

A.1 Characterization with the SPG simulation

Figure A.83 shows the $m_{\mu\mu}$, L_{xy}, and L_{xy} significance distributions for prompt dimuon events in the invariant mass range 0.2–1 GeV, simulated by the SPG sample. The L_{xy} distribution in Figure A.83b shows a second population at L_{xy} = 20–40 mm, which is a clear indication of vertex pulling since events are generated with dimuons produced at the interaction point ($L_{xy} = 0$).

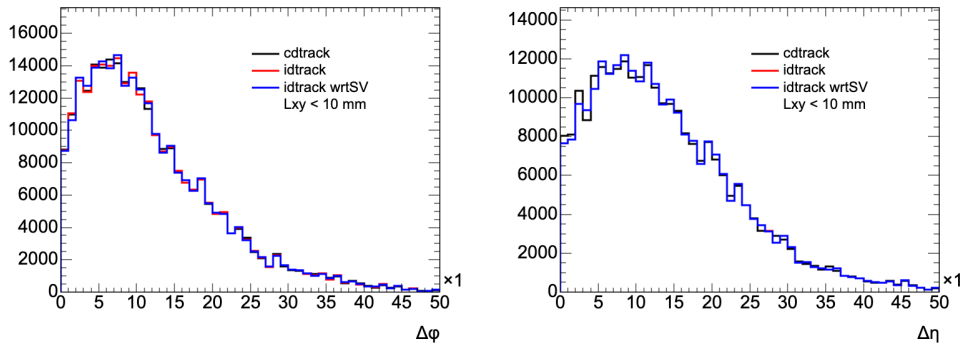


(a) (b) (c)

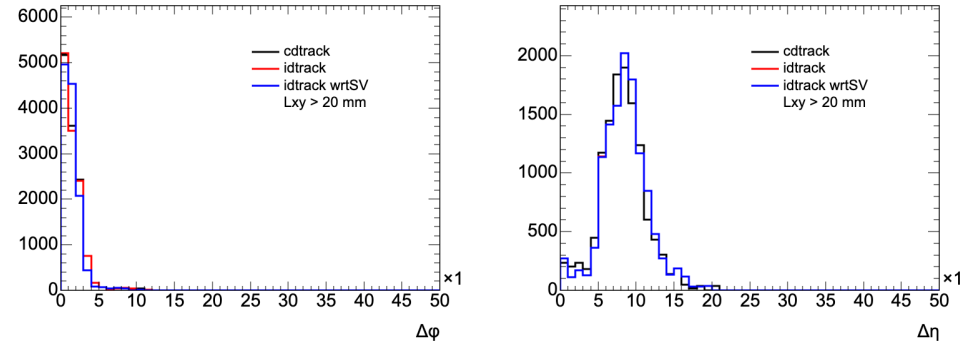
Figure A.83. The $m_{\mu\mu}$, L_{xy}, and L_{xy} significance distributions in the prompt-muon SPG sample.

The angular characteristics of pulled vertices are illustrated in Figure A.84, showing dimuon $\Delta\phi$ and $\Delta\eta$ for vertices at small ($L_{xy} < 10$ mm) and large ($L_{xy} > 20$ mm) transverse displacement. Prompt dimuons ($L_{xy} < 10$ mm) have a broad opening angle distribution, while pulled vertices ($L_{xy} > 20$ mm) have small $\Delta\phi$ and $\Delta\eta$, confirming that vertex pulling is strongly correlated with

track collimation. Figure A.85b additionally shows that dimuon $\Delta\phi$ is a better variable on which to apply selections for background rejection.



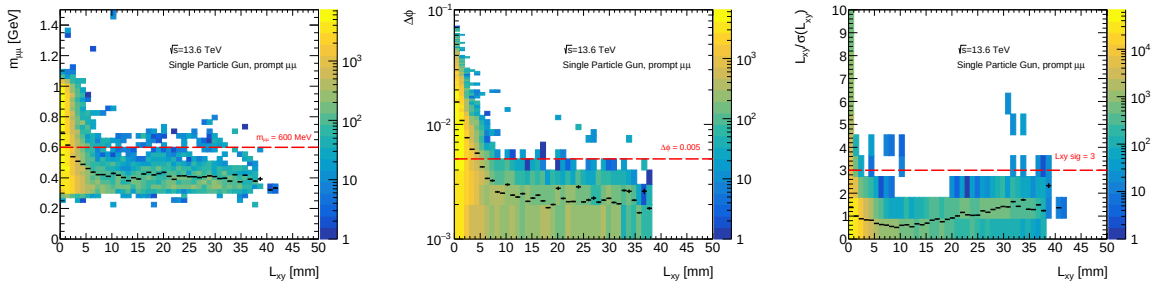
(a)



(b)

Figure A.84. dimuon $\Delta\eta$ and $\Delta\phi$ distributions in the prompt-muon SPG sample, for vertices reconstructed at (a) small displacement ($L_{xy} < 10$ mm) and (b) large displacement ($L_{xy} > 20$ mm). The large-displacement region is dominated by vertex pulling, and shows that the muon pairs are more collimated. The x-axis is scaled by 10^{-3}.

Figure A.85 shows the correlations between pulled secondary vertices and $m_{\mu\mu}$, dimuon $\Delta\phi$, and L_{xy} significance. Vertex pulling is strongly associated with small $\Delta\phi$ and low L_{xy} significance. Requiring $\Delta\phi > 0.005$ and L_{xy} significance > 3 therefore efficiently rejects pulled vertices from the prompt background.

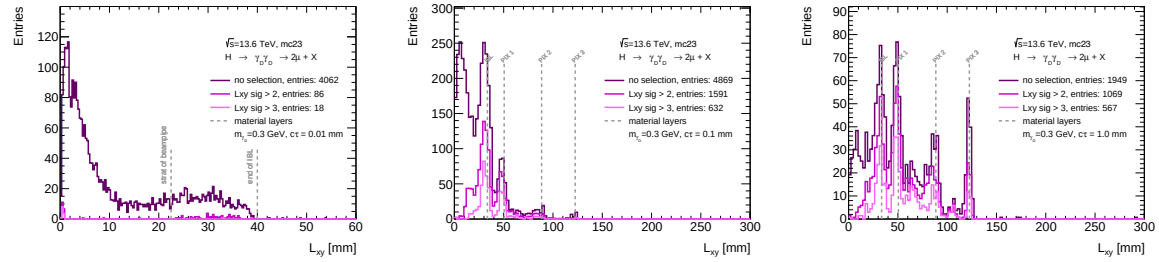


(a) (b) (c)

Figure A.85. The $m_{\mu\mu}$, dimuon $\Delta\phi$, and L_{xy} significance as a function of L_{xy} in the prompt-muon SPG sample.

A.2 Characterization with signal simulation

The impact of the L_{xy} significance selection on the signal is shown in Figure A.86 for the $m_{\gamma_D} = 0.3$ GeV signal, and in Figure A.87 for the $m_{\gamma_D} = 0.6$ and 1.2 GeV signals with distributions shown before and after applying L_{xy} significance > 2 and > 3 requirements. The vertex pulling effect is characterized as peaks in the L_{xy} distribution at the pixel layer positions. The L_{xy} significance requirement effectively suppresses prompt events (e.g. signal with short lifetimes) at low L_{xy}, while retaining most of the displaced signal events at larger L_{xy}.



(a) (b) (c)

Figure A.86. Reconstructed L_{xy} distributions for the $m_{\gamma_D} = 0.3$ GeV signal at $c\tau = 0.01$, 0.1, and 1 mm, shown before and after requiring L_{xy} significance > 2 and > 3. The dashed vertical lines label the pixel detector layers (IBL and PIX 1–3).

The impact of the dimuon $\Delta\phi$ selection is shown in Figure A.88 for $m_{\gamma_D} = 0.3$ GeV and in Figure A.89 for $m_{\gamma_D} = 0.6$ and 1.2 GeV, with distributions shown before and after applying $\Delta\phi >$

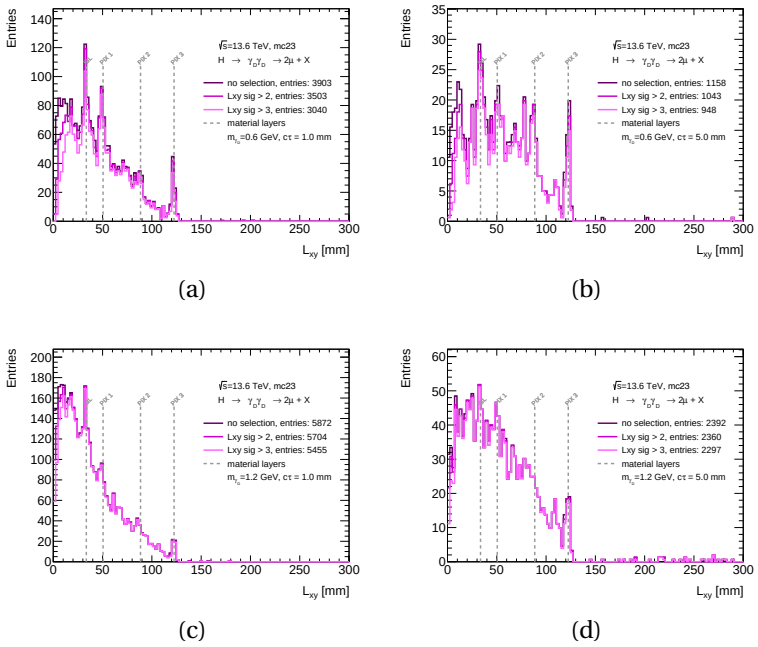


Figure A.87. Reconstructed L_{xy} distributions for the $m_{\gamma_D} = 0.6$ and 1.2 GeV signals at $c\tau = 1$ and 5 mm, shown before and after requiring L_{xy} significance > 2 and > 3. The dashed vertical lines label the pixel detector layers (IBL and PIX 1–3).

0.005. The requirement reduces the size of vertex pulling peaks more uniformly throughout all L_{xy}. The 300 MeV, $c\tau = 0.01$ mm signal, shown in Figure A.88a, illustrates the effectiveness of rejecting prompt dimuon events reconstructed at the IBL.

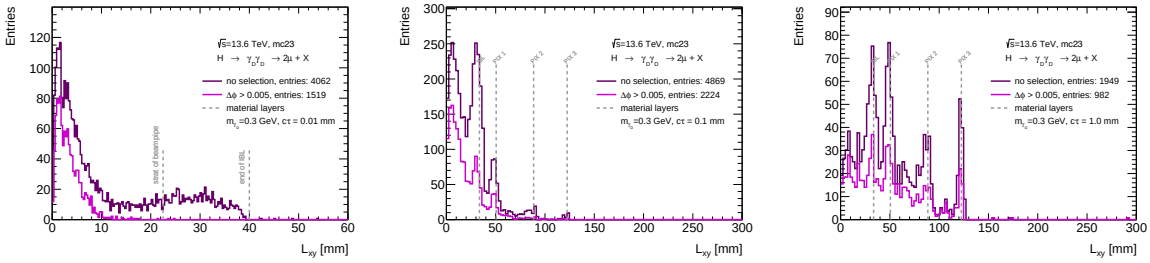


Figure A.88. Reconstructed L_{xy} distributions for the $m_{\gamma_D} = 0.3$ GeV signal at $c\tau = 0.01$, 0.1, and 1 mm, shown before and after requiring dimuon $\Delta\phi > 0.005$. The dashed vertical lines label the pixel detector layers (IBL and PIX 1–3).

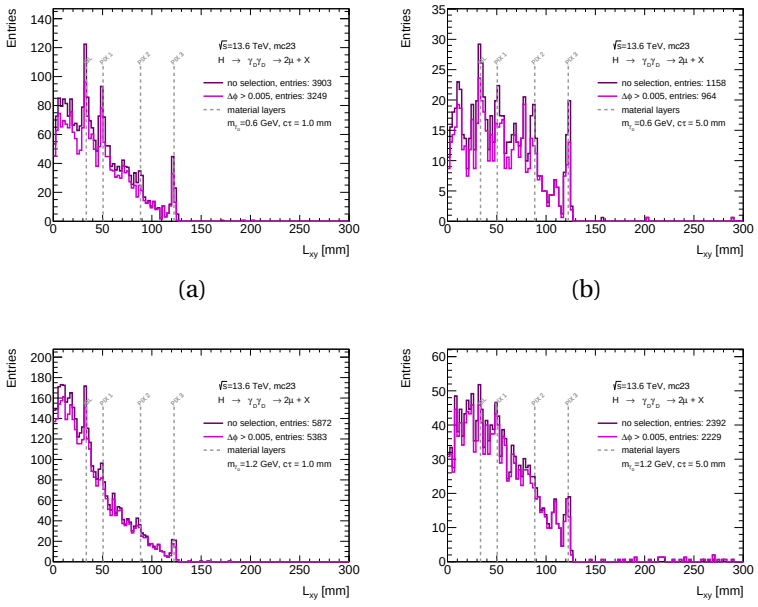


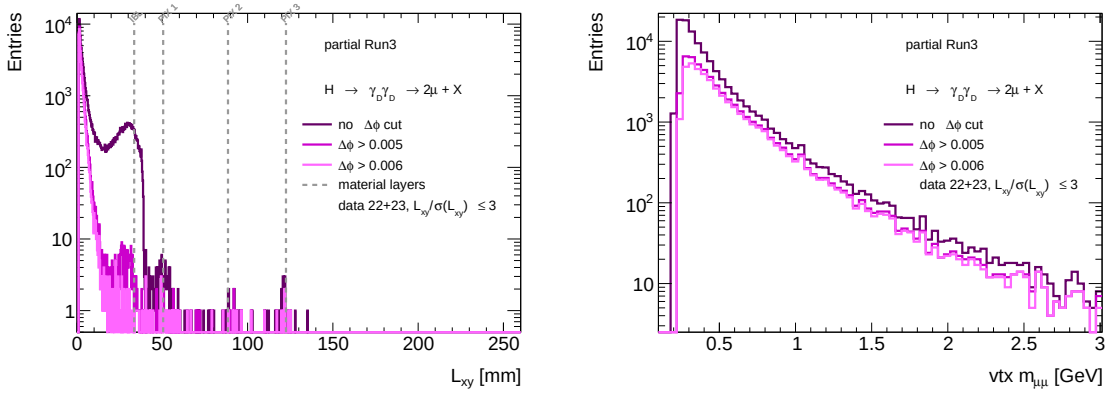
Figure A.89. Reconstructed L_{xy} distributions for the m_{γ_D} = 0.6 and 1.2 GeV signals at $c\tau$ = 1 and 5 mm, shown before and after requiring dimuon $\Delta\phi > 0.005$. The dashed vertical lines label the pixel detector layers (IBL and PIX 1–3).

A.3 Characterization with control region data

The efficiency of the $\Delta\phi$ selection is validated in data using the inverted L_{xy} significance control region. Figure A.90 shows the L_{xy} and $m_{\mu\mu}$ distributions after requiring $\Delta\phi$ to be larger than 0.005 and 0.006. As shown in Figure A.90a, a $\Delta\phi$ requirement of 0.005 is 98% efficient in removing DVs with L_{xy} larger than 20 mm. Figure A.90b shows that the invariant mass distribution falls smoothly from about twice the muon mass.

The residual prompt background from vertex pulling is estimated by extrapolating the number of vertex pulling background events from the `fail` L_{xy} `significance and fail dimuon` $\Delta\phi$ CR to the SR, using a scale factor derived from the SPG sample. Vertex pulling background events are defined as DVs with $L_{xy} > 20$ mm in the CR. The final estimate for the remaining pulled DVs in the SR data is approximately 4.8 events for 56 fb^{-1}, which is small compared to the dominant heavy flavor background in the relevant mass region. The low-

mass Drell–Yan contribution due to vertex pulling is included in the parametric background modeling in the 2-muon channel low-mass region, as discussed in Section 6.5.3.



(a) (b)

Figure A.90. Distributions of $m_{\mu\mu}$ and L_{xy} with $\Delta\phi$ required to be larger than 0.005 and 0.006 in the fail-L_{xy}-significance CR using 2022–2023 data.

APPENDIX 2

Analog Readback for the ATLAS Inner Tracker Upgrade Pixel Module Electrical Quality Control

For the up-coming high-luminosity LHC (HL-LHC), the ATLAS detector is undergoing significant upgrades to overcome the challenges of higher simultaneous collision rates and radiation levels. The ATLAS Inner Tracker (ITk) is a nine-layer all-silicon tracking detector for HL-LHC. This appendix documents the work on the development of the electrical quality control tests during the pre-production of the ITk Pixel modules, in particular the Analog Readback test.

The pre-production ITk pixel modules use the ITkPixV1.1 front-end (FE) chip, designed based on the RD53B readout chip [65]. A newer version ITkPixV2 is used for production. An electrical Quality Control (QC) procedure has been developed to ensure electrical performance of modules throughout production and integration, designed to test more than 10 thousand modules, across 25 sites, and throughout all stages. As part of the QC, the analog readback test measures and analyzes a large number of chip parameters and properties, including 72 internal biases, 4 different temperature sensors, and 42 Ring Oscillators for characterizing radiation damage to the FE.

B.1 The ATLAS Inner Tracker upgrade and Pixel Modules

The ATLAS ITk consists of five silicon pixel layers and four silicon strip layers, with the detector coverage extended out to $|\eta| < 4.0$. A schematic illustration of ITk is shown in Figure B.91, and a detailed layout is shown in Figure B.92a. The Transition Radiation Tracker from the current Inner Detector will be removed, replaced by nine silicon layers, as shown in Figure B.92b.

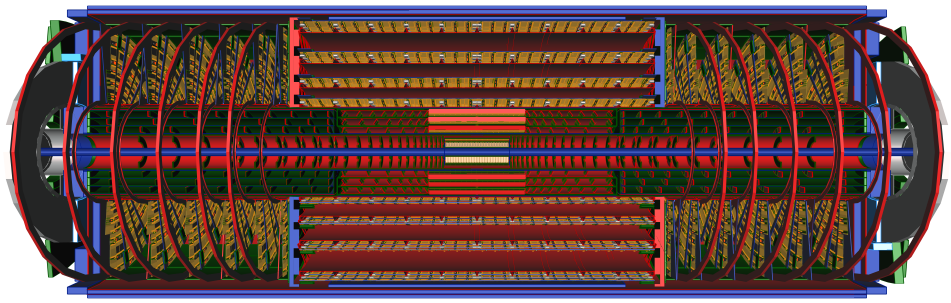
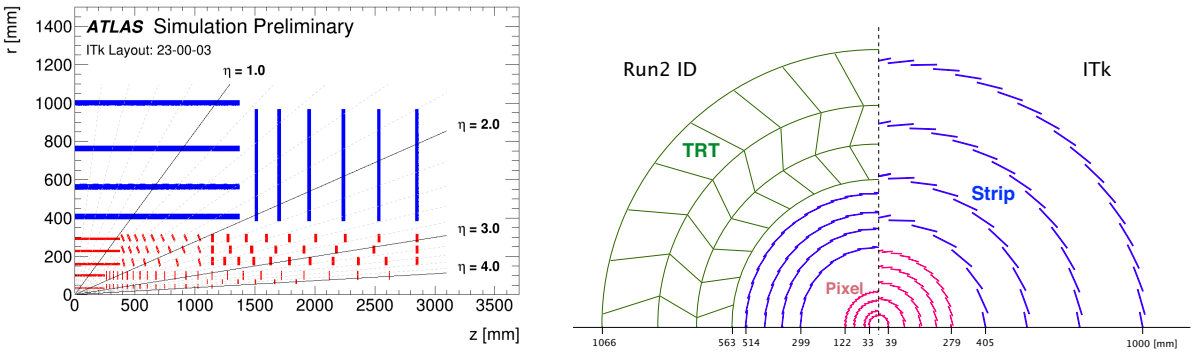


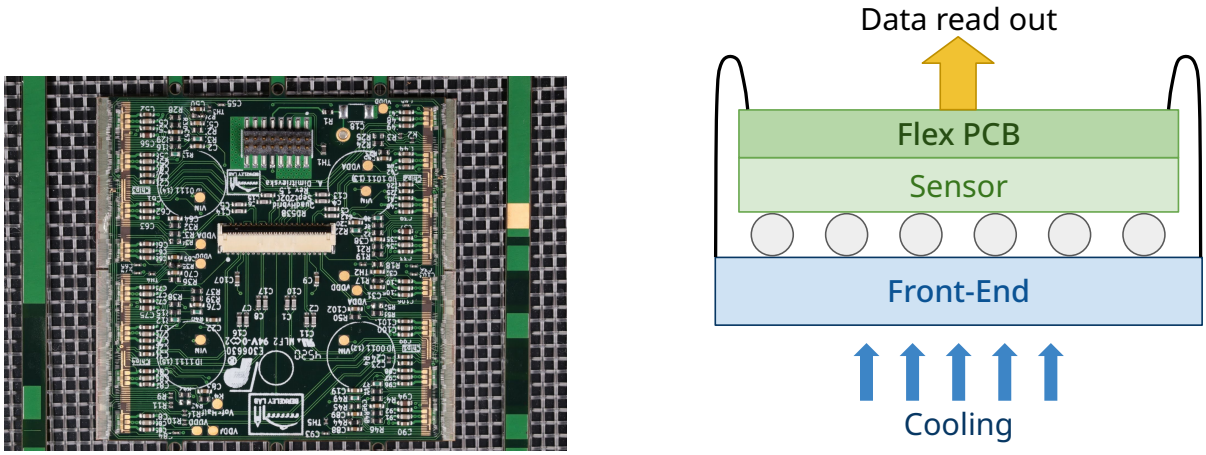
Figure B.91. A schematic illustration of the ATLAS ITk detector. Figure taken from [66].



(a) (b)

Figure B.92. Detailed layout of the ITk. Figure (a) shows one quadrant of the ITk in the r–z plane, with the pixel layers in red and the strip layers in blue [66]. Figure (b) shows a comparison between the current Run 2-Run 3 ATLAS Inner Detector, and ITk in the transverse plane [67].

The ITk pixel sensor pitch is 25×100 μm^2 in the barrel of the innermost layer, and 50×50 μm^2 in the rest of the pixel detector. Overall, there are a total of ~ 9.2k pixel modules and ~ 5 billion pixel channels, which is more than 50 times more pixels than the current Inner Detector. There are two types of ITk pixel modules: quad modules consisting of four FE chips, and triplet modules consisting of three FE chips. Each FE chip is of 384 pixel rows and 400 pixel columns, and of size 2×2 cm^2. Figure B.93 shows a picture and a schematic layout of a quad module, where 4 chips are bump-bonded to a single sensor, and are wire-bonded to a PCB.



(a) (b)

Figure B.93. An ITk Pixel quad module. Figure (a) shows a photograph of an assembled quad module, and Figure (b) shows a schematic layout, in which four front-end chips are bump-bonded to a single sensor and wire-bonded to the module PCB.

B.2 Analog readback in module QC

The module electrical QC defines the specifications, testing procedures and the tools to check the electrical performance of ITk Pixel modules during pre-production and production. The analog readback test measures all internal biases of the front-end chip in order to give access to the status of the chip during module QC, and after the modules are loaded onto the local support. These parameters include all internal voltages and currents and recordings of

the analog and digital supply voltages (VDDA and VDDD) and the ring-oscillator frequencies as a function of the trim setting. It also measures the temperature of the chip via temperature sensors (NTC, MOS, and Poly sensors) in order to check if the temperature sensors are functional, and provide a way to monitor the temperature of the chip throughout the testing process, and after the modules are loaded onto the local support. Different temperature sensors can also monitor the temperature at different parts of the chip.

B.2.1 Chip sensing and monitoring

The internal biases are read out through the chip monitoring block, shown in Figure B.94. The block includes an analog current multiplexer (IMUX), an analog voltage multiplexer (VMUX), an analog-to-digital converter (ADC), and dedicated VMUX, IMUX, and NTC pads. The current and voltage multiplexers feed the input of the ADC, and are available on the VMUX pad for external measurements. The current at IMUX can be measured as a voltage through VMUX across an external resistor R_{IMUX} to ground, and the same applies to the external NTC.

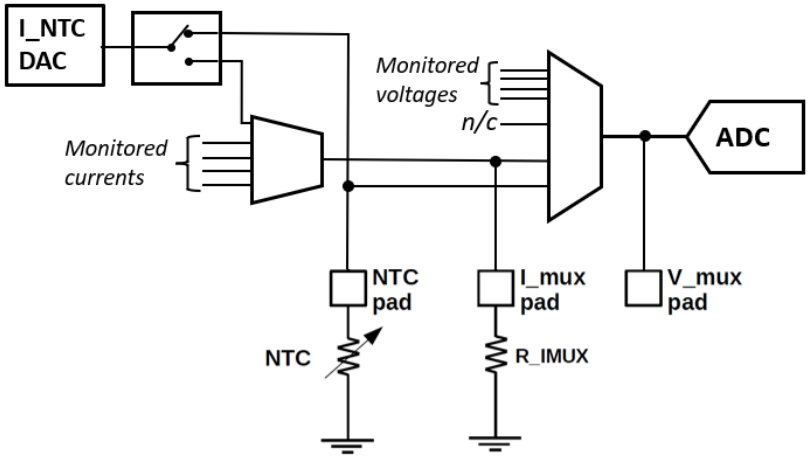


Figure B.94. Diagram of the RD53B FE chip monitoring block. The current multiplexer (IMUX) and analog voltage multiplexer (VMUX) feed the input of the analog-to-digital converter (ADC). Figure taken from [65].

The voltage multiplexer provides access to internal voltages such as VINA, VIND, VDDA, and VDDD, to the temperature sensors and to the radiation-monitoring sensor. The current multiplexer provides access to the internal bias for the phase-locked loop (PLL)-based Clock and Data Recovery circuit (CDR), the current mode logic (CML) data output driver, and the analog front-end (DIFF FE), as well as to the chip input and shunt currents. The full list of VMUX and IMUX channels is shown in Table B.28 and Table B.27, respectively. The individual components in the FE chip are shown in Figure B.95.

Table B.27. Current multiplexer (I_mux) assignments for the ITk Pixel FE chip [65].

Setting	**Selected Input**	**Setting**	**Selected Input**	**Setting**	**Selected Input**
0	IREF main ref. current	11	Capmeasure parasitic	22	DIFF FE Preamp Top-Left
1	CDR VCO main bias	12	DIFF FE Preamp Main array	23	DIFF FE VTH1 Right
2	CDR VCO buffer bias	13	DIFF FE PreComp	24	DIFF FE Preamp Top
3	CDR CP current	14	DIFF FE Comparator	25	DIFF FE Preamp Top-Right
4	CDR FD current	15	DIFF FE VTH2	26	not used
5	CDR buffer bias	16	DIFF FE VTH1 Main array	27	not used
6	CML driver tap 2 bias	17	DIFF FE LCC	28	Ana. input current/21000
7	CML driver tap 1 bias	18	DIFF FE Feedback	29	Ana. shunt current/21600
8	CML driver main bias	19	DIFF FE Preamp Left	30	Dig. input current/21000
9	NTC_pad current	20	DIFF FE VTH1 Left	31	Dig. shunt current/21600
10	Capmeasure circuit	21	DIFF FE Preamp Right	32–62	not used
				63	high Z

Table B.28. Voltage multiplexer (V_mux) assignments for the ITk Pixel FE chip [65].

Setting	**Selected Input**	**Setting**	**Selected Input**	**Setting**	**Selected Input**
0	Vref_ADC (GADC)	10	DIFF FE VTH1 Main array	31	Vref_CORE
1	I_mux pad voltage	11	DIFF FE VTH1 Left	32	Vref_PRE
2	NTC_pad voltage	12	DIFF FE VTH1 Right	33	VINA / 4
3	VCAL_DAC / 2	13	RADSENS Ana. SLDO	34	VDDA / 2
4	VDDA / 2 from capmeasure	14	TEMPSENS Ana. SLDO	35	VrefA
5	Poly TEMPSENS top	15	RADSENS Dig. SLDO	36	VOFS / 4
6	Poly TEMPSENS bottom	16	TEMPSENS Dig. SLDO	37	VIND / 4
7	VCAL_HI	17	RADSENS center	38	VDDD / 2
8	VCAL_MED	18	TEMPSENS center	39	VrefD
9	DIFF FE VTH2	19–30	Ana. GND	40–62	not used
				63	high Z

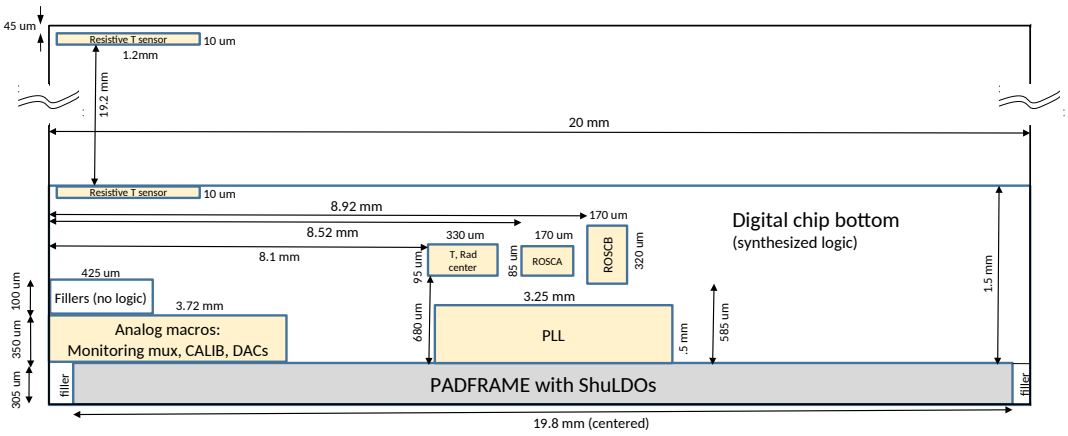


Figure B.95. Location of the main elements on the ITkPixV1.1 front-end chip, including the monitoring block, the analog and digital shunt-LDO regulators, the temperature and radiation sensors and the Ring Oscillators. Figure taken from [65].

B.3 Temperature measurement

Three types of temperature sensors are available in the ITkPixV1.1 chip: the resistor-based negative temperature coefficient (NTC) sensors, transistor-based metal-oxide semiconductor (MOS) sensors, and the polysilicon resistor temperature sensor. The analog readback test performs two types of measurement for the NTC sensor and one for the MOS sensors, while the test has not been developed for the polysilicon sensor.

The placement of the NTC sensors is shown in Figure B.96, where four external NTCs (one NTC per chip) are routed to their corresponding chip via wire bond pads. There is one Detector Control System (DCS)/Interlock External NTC for every module, which is connected externally via the power connector. The MOS sensors are inside the chip, with one at the center of the analog chip bottom, one at the analog SLDO, and one at the digital SLDO. The polysilicon sensors measure the temperature difference between the top and bottom edges of the chip, however, since the circuits to measure its resistance are not very precise in ITkPixV1.1, it is not useful for pre-production.

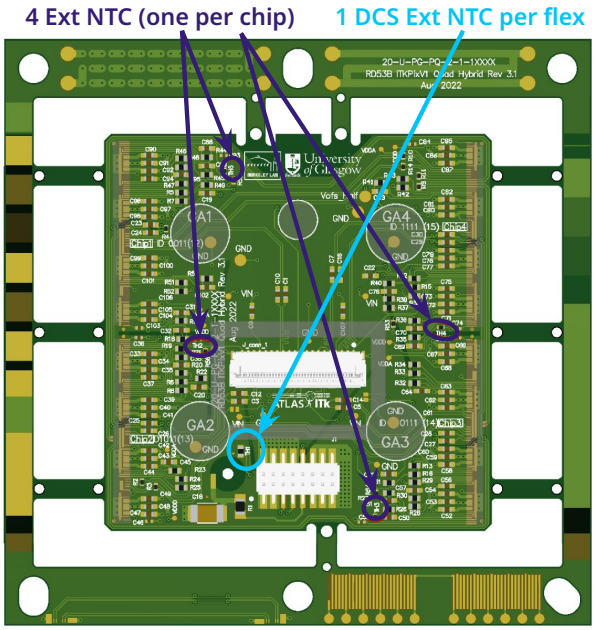


Figure B.96. Placement of the NTC temperature sensors on the ITkPixV1.1 front-end chip.

B.3.1 External NTC temperature sensor measurement

The temperature of the NTC sensors is measured via their resistance. For each external NTC sensor, the FE chip sources a known current through the NTC, which can be measured through IMUX. The voltage at the NTC pad is measured via VMUX. The NTC resistance is obtained from

$$V_{\mathrm{NTC}} = V_{\mathrm{mux2}} - \mathrm{GND}, \qquad I_{\mathrm{NTC}} = \frac{I_{\mathrm{mux9}} - \mathrm{GND}}{R_{\mathrm{IMUX}}}, \qquad R_{\mathrm{NTC}} = \frac{V_{\mathrm{NTC}}}{I_{\mathrm{NTC}}}, \tag{B.1}$$

which is converted to temperature using the Steinhart-Hart equation,

$$\frac{1}{T} = A + B \ln R_{\mathrm{NTC}} + C \left(\ln R_{\mathrm{NTC}}\right)^3, \tag{B.2}$$

where the coefficients A, B, and C are taken from the chip configuration. The external NTC is precise and easy to measure, but since it is attached to the flex, it only measures the external chip temperature.

B.3.2 DCS external NTC

The DCS external NTC is used by DCS monitoring and interlocking. During module QC, its resistance is measured by connecting an external multimeter to the corresponding pin on the power adapter board. The temperature is obtained using the Steinhart-Hart equation, and the precision is limited by the resolution of the multimeter and of the ADC.

B.3.3 MOS sensors

The MOS sensors are placed close to the shunt-LDO regulators, which are the main heat sources on-chip, which makes them more sensitive to the chip temperature. However, they are more complex to read out and change with radiation damage. The diagram of a MOS is shown in Figure B.97.

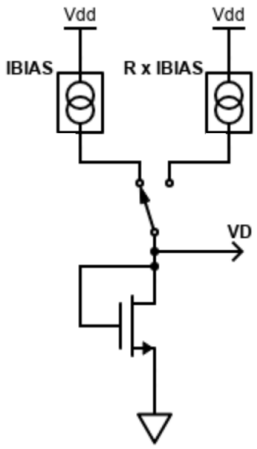


Figure B.97. Diagram of a MOS temperature sensor on the ITkPixV1.1 front-end chip. Figure taken from [65].

The voltage difference ΔV_D across a diode has a complementary-to-absolute-temperature variation:

$$\Delta V_D = V_D(R \times I_{\text{bias}}) - V_D(I_{\text{bias}}) = N_f \frac{k_B T}{q} \ln R, \tag{B.3}$$

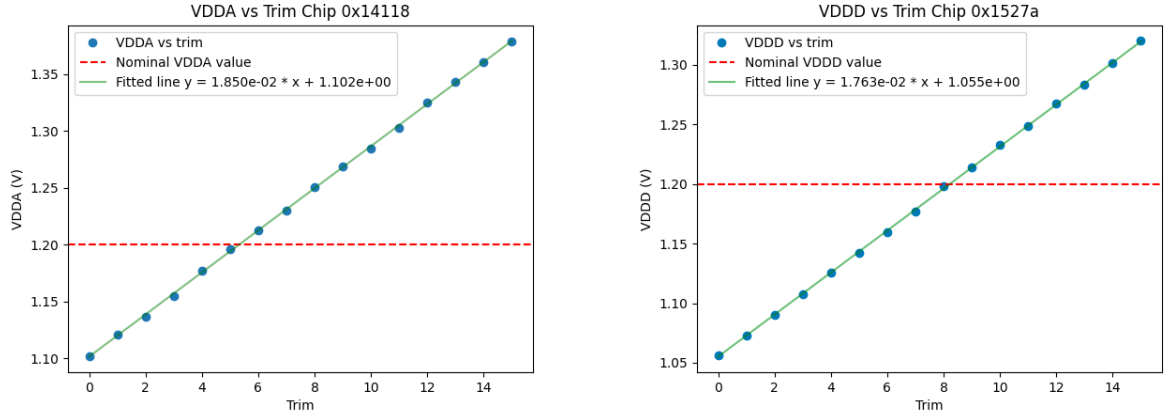
such that the absolute temperature can be measured from ΔV_D as

$$T = \frac{\Delta V_D \, q}{N_f \, k_B \ln R}, \tag{B.4}$$

where q is the electron charge, k_B the Boltzmann constant, R a fixed current ratio (nominally $R = 15$), and N_f the diode ideality factor taken from wafer probing. The MOS sensors are in principle very precise, while their precision is limited by the ideality factor N_f.

B.4 VDDA and VDDD measurement

The VDDA and VDDD regulator outputs should nominally produce approximately 1.2 V. These voltages are affected by the main reference current (Iref), which can be trimmed with wire bond pads when needed. The analog readback test measures the VDDA and VDDD as a function of trim settings, in order to find the trim value that gives VDDA and VDDD closest to 1.2 V. The chip configuration is then updated to use these trim settings. Example measurements of VDDA and VDDD as a function of trim are shown in Figure B.98.



(a) (b)

Figure B.98. Measured VDDA (a) and VDDD (b) as a function of the trim setting.

B.5 Ring Oscillator measurement

Ring Oscillators (ROSCs) are test circuits that measure the speed of logic cells. They characterize the radiation damage to different gates, with the oscillation frequency decreasing with increasing radiation damage. A diagram of a ring oscillator block is shown in Figure B.99. The ROSCs are located in two banks in the chip bottom, as shown in Figure B.95. Bank A has 8 ROSCs and Bank B has 34 ROSCs. It is important to measure them during module QC to make sure the chip has working ROSCs, as well as to collect raw measurements before radiation.

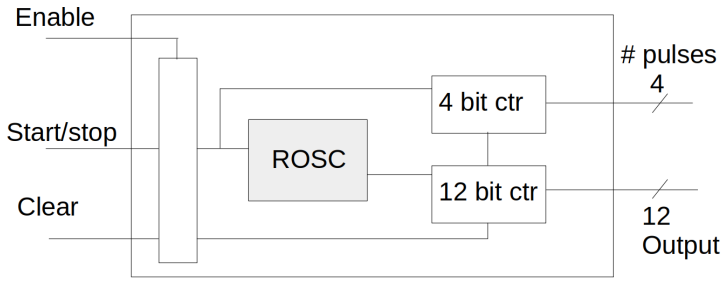


Figure B.99. Diagram of a ring oscillator block in the ITkPixV1.1 front-end chip. Figure taken from [65].

In analog readback, the ROSC oscillation frequency in MHz is measured as a function of VDDD under different trim values (and temperature). A linear relation is expected, where the slopes and offsets are determined for each ring oscillator. Example measurements are shown in Figure B.100.

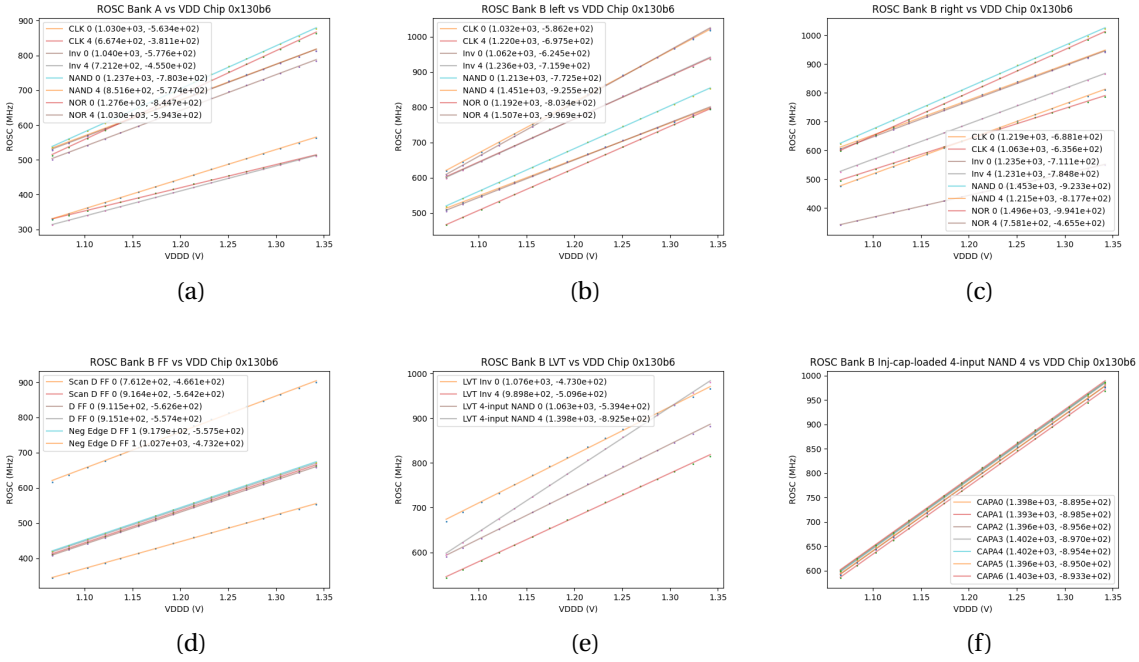


Figure B.100. Example ring oscillator measurements for the Bank A and Bank B ROSCs, showing the oscillation frequency as a function of VDDD. The figures correspond to Bank A (a), Bank B left (b), Bank B right (c), Bank B FF (d), Bank B LVT (e), and Bank B injection capacitor (f).

APPENDIX 3

HL-LHC ATLAS Inner Tracker Pixel Detector ToT-to-Charge Calibration and Tracking Performance

In the ATLAS Inner Tracker (ITk) pixel detector upgrade for the high-luminosity LHC, digitized charge information must be stored under extremely high hit rates. The precision of the ToT-to-charge calibration could directly impact cluster position resolution and tracking performance. This appendix summarizes a study aimed at understanding the impact of a pixel-level time-over-threshold (ToT) to charge calibration, as opposed to the front-end (FE)-level calibration currently used, on the cluster position resolution and overall tracking performance of the ITk. The ToT observable and its role in charge calibration and analog clustering are introduced in Section 5.1.1.

C.1 Time-over-threshold and charge reconstruction

In both the current ATLAS Pixel detector and the ITk upgrade, charge deposited in the pixel sensor is read out by the front-end electronics as a Time-over-Threshold (ToT) value when the preamplifier output exceeds an adjustable analog threshold. This threshold is tuned via charge injection to maintain a high signal-to-noise ratio. The ToT is a discrete, integer-valued quantity encoded in 4 bits, clocked at the 40 MHz LHC bunch-crossing frequency. The offline reconstruction then converts the digitized ToT back to charge, which is used as input to track reconstruction.

The charge reconstruction chain is illustrated in Figure C.101. For simulated events, the analog charge deposition is converted to ToT during the digitization (DIGI) step using the calibration function ToT = $f(Q)$ derived from the ITk pixel front-end response, providing a relatively accurate representation of the raw detector readout. The inverse conversion, Q =

f^{-1}(ToT), is then applied during reconstruction (RECO) step to both simulation and data, recovering the deposited charge from the digitized ToT values.

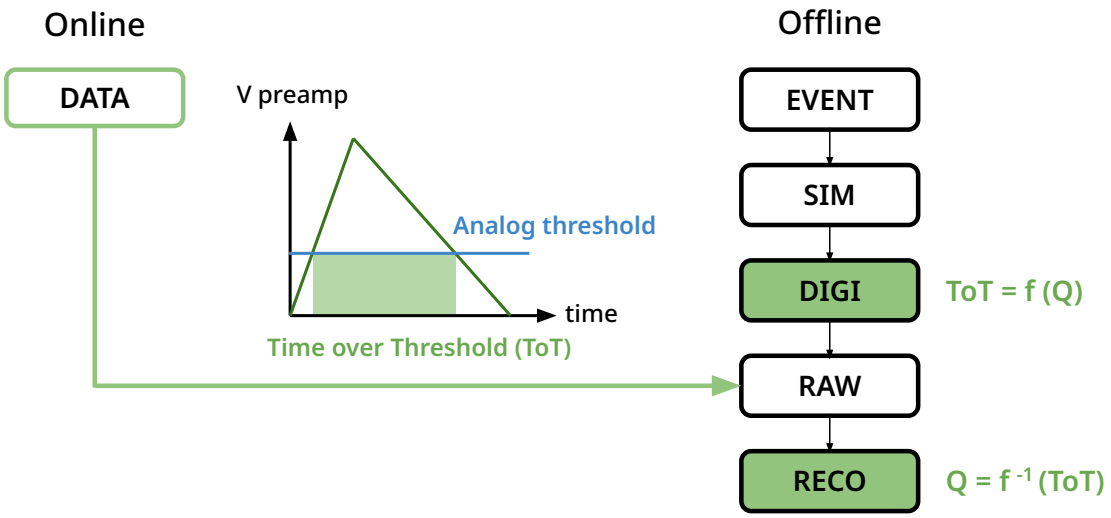


Figure C.101. A schematic illustration of the charge reconstruction chain at the ATLAS Pixel detector. The charge deposition in the pixel front-ends is measured by time-over-threshold (ToT), which has limited resolution. In the offline software, simulations (SIM) of events go through a digitization (DIGI) step to convert analog charge into digitized ToT based on the FE configuration. Both simulated and raw data events are then reconstructed (RECO), where discrete ToT values are converted back to charge.

The current ATLAS Pixel front-end is close to a linear FE, and a FE-level calibration is used, assuming all pixel channels in each FE have the same ToT–charge correlation, fitted by a polynomial. The ToT-to-charge conversion is then done using this polynomial fit. The same method has been proposed for the ITk Pixel detector, and a FE-level calibration is used for the ITk Pixel in the current simulation and reconstruction. In addition, the current Pixel FE has a higher ToT resolution, where the IBL has a 4-bit ToT and the other three pixel layers have an 8-bit ToT.

The ITk Pixel detector uses a differential front-end, which exhibits a non-linear ToT versus charge response. An important difference with respect to the current Pixel detector is that the ITk front-end lacks the ability to adjust the feedback current, which effectively controls the

charge-to-ToT conversion. As a result, unlike the current Pixel detector, the ITk Pixel front-end cannot be tuned to make all pixels behave uniformly.

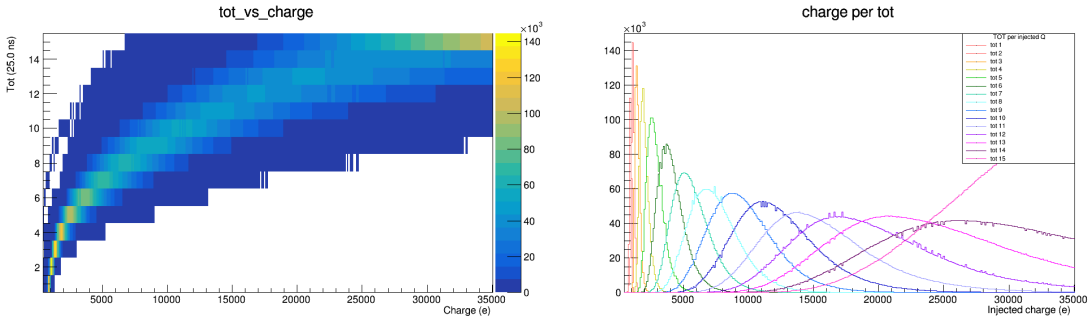
Because of the non-uniform per-pixel response, a FE-level calibration may not be optimal for the ITk Pixel detector. It is therefore worth investigating whether a pixel-level calibration can significantly improve the pixel cluster position resolution and, in turn, the overall tracking performance.

C.2 ToT-versus-charge response in the ITk Pixel front-end

The ToT versus charge response of the ITk pixel front-end is characterized using a single-chip module with a 3D sensor of pixel pitch 50×50 μm^2. The front-end is tuned to a threshold of 1000 e^-. A ToT scan is performed by injecting charges in the range 500–35000 e^- in steps of 100 e^-, from which the charge-to-ToT response is constructed across the full dynamic range of the front-end.

The resulting ToT-versus-charge distribution is shown in Figure C.102a, exhibiting a highly non-linear response. The charge per ToT value, shown in Figure C.102b, further demonstrates the saturation of resolution at higher injected charge values.

There exists an intrinsic charge resolution loss due to the limited ToT resolution. While each injected charge corresponds to a unique ToT value, the same ToT value can be mapped to a range of true deposited charges. The non-linear response, together with a large ToT dispersion, further implies that the reconstructed charge uncertainty varies significantly across the dynamic range, raising the question of whether a uniform FE-level calibration sufficiently captures pixel-to-pixel variations, or whether a pixel-level calibration is necessary.



(a) (b)

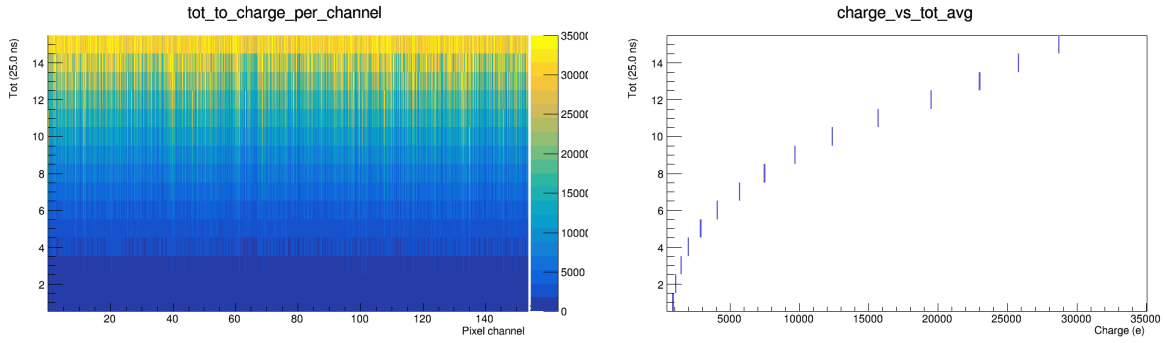
Figure C.102. ToT response of the ITk Pixel front-end as a function of injected charge, measured on a single-chip module with a 50×50 μm^2 3D sensor tuned to a threshold of 1000 e^-. Figure (a) shows the ToT as a function of injected charge, and Figure (b) shows the injected charge per ToT count.

C.3 Charge calibration at pixel-level vs at FE-level

From the ToT-versus-charge response, a pixel-level calibration look-up table is constructed for every channel in the front-end chip, mapping ToT values to reconstructed charge. Each FE consists of $400 \times 384 = 153600$ pixel channels, and the resulting per-channel calibration variation is shown in Figure C.103a. In comparison, a single FE-level calibration look-up table is derived by averaging the injected charge over all channels for each ToT value, shown in Figure C.103b. This averaged calibration is applied uniformly to every pixel in the chip, which does not account for the per-pixel variations.

To quantify the impact of the calibration granularity on the charge reconstruction resolution, a closure test is performed. A known injected charge Q_0 is compared with the reconstructed charge Q obtained from the corresponding ToT value. The closure test is performed for both the pixel-level and the FE-level (per-chip) calibrations using the look-up tables shown in Figure C.103. The difference between the reconstructed charge Q and the true injected charge Q_0 characterizes the charge reconstruction uncertainty and the resolution loss.

Figure C.104 shows the reconstructed charge Q as a function of the injected charge Q_0, and the corresponding relative residual $(Q - Q_0)/Q_0$, for the two types of calibrations. In the pixel-



(a) (b)

Figure C.103. ToT-to-charge calibration look-up tables for a single front-end chip. Figure (a) shows the pixel-level calibration look-up table, where the reconstructed charge on the z-axis is given for each pixel channel on the x-axis and ToT value on the y-axis. Figure (b) shows the FE-level calibration look-up table, obtained by averaging the injected charge over all channels for each ToT value.

level calibration, the finite width of the Q distribution at a fixed Q_0 comes from the loss of resolution due to the discretized ToT, and the relative residual is approximately ±20% across most of the dynamic range. In the FE-level calibration, a much larger charge resolution loss is observed caused by neglecting the pixel-level variations.

The width of the reconstructed-charge distribution at a given Q_0 directly quantifies the charge reconstruction uncertainty. Figure C.105 shows the RMS of the Q distribution as a function of Q_0 for the two types of calibrations. The pixel-level calibration shows a significantly smaller RMS as a function of Q_0 compared to the FE-level calibration. The fall-off of the two RMS distributions at high charge is an artifact of the saturation of ToT values.

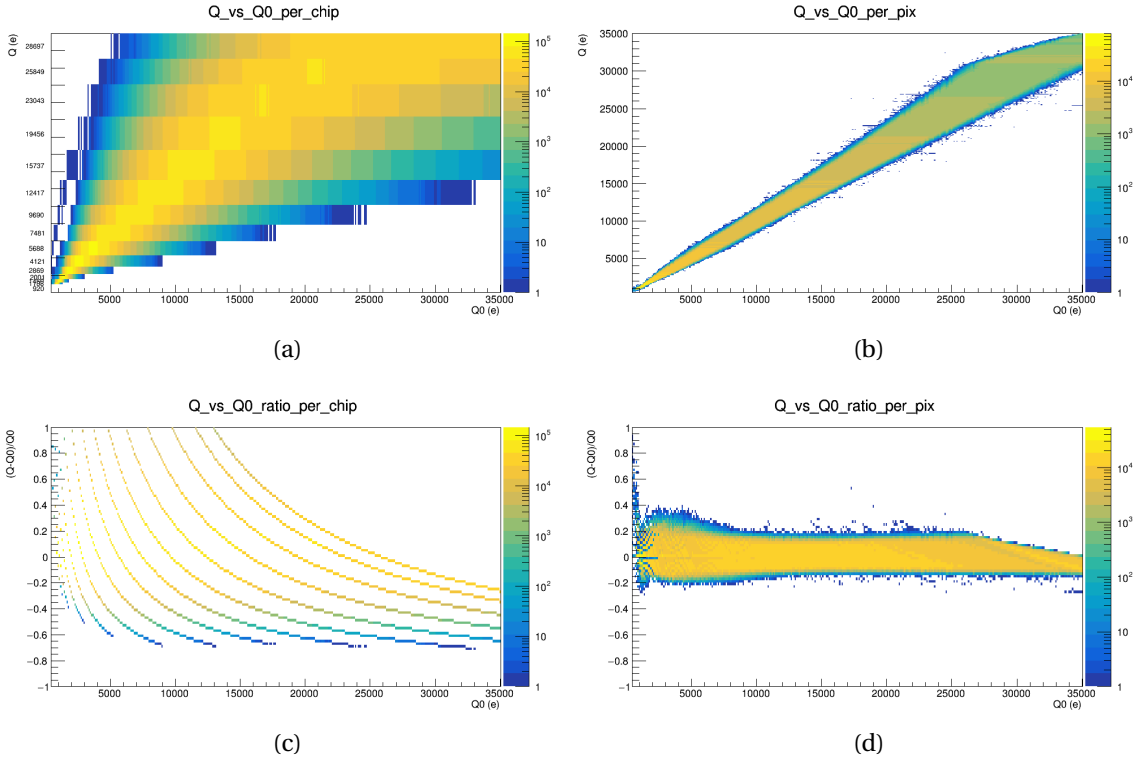


Figure C.104. Closure test comparing the FE-level (per-chip) and pixel-level calibrations. The reconstructed charge Q as a function of the injected charge Q_0 is shown for the FE-level (a) and pixel-level (b) calibrations at the top. The corresponding relative residuals $(Q - Q_0)/Q_0$ are shown for the FE-level (c) and pixel-level (d) calibrations at the bottom.

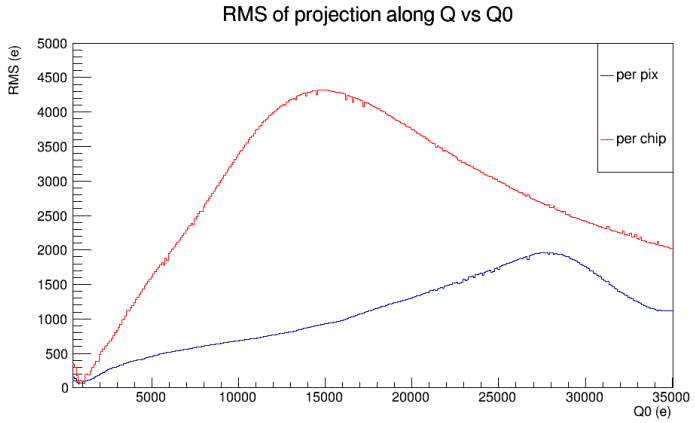


Figure C.105. RMS of the reconstructed-charge Q distribution as a function of the injected charge Q_0, for the pixel-level (per pix) and FE-level (per chip) calibrations.

C.4 Impact on tracking performance

Since charge information is used for analog clustering where the hit positions are found for track reconstruction, the charge calibration granularity may impact the overall tracking performance. This work further quantifies the impact of pixel-level charge calibration on track reconstruction relative to the current FE-level calibration approach.

Four configurations of digitization and reconstruction are compared, as outlined below:

1. **Original DIGI + Original RECO**: the baseline configuration using the current FE-level calibration for both digitization and reconstruction.

2. **Per-chip DIGI + Per-chip RECO**: a FE-level calibration applied in both steps using the look-up table derived in this study, as shown in Figure C.103b.

3. **Per-pixel DIGI + Per-pixel RECO**: a pixel-level calibration applied in both steps using the look-up table derived in this study, as shown in Figure C.103a.

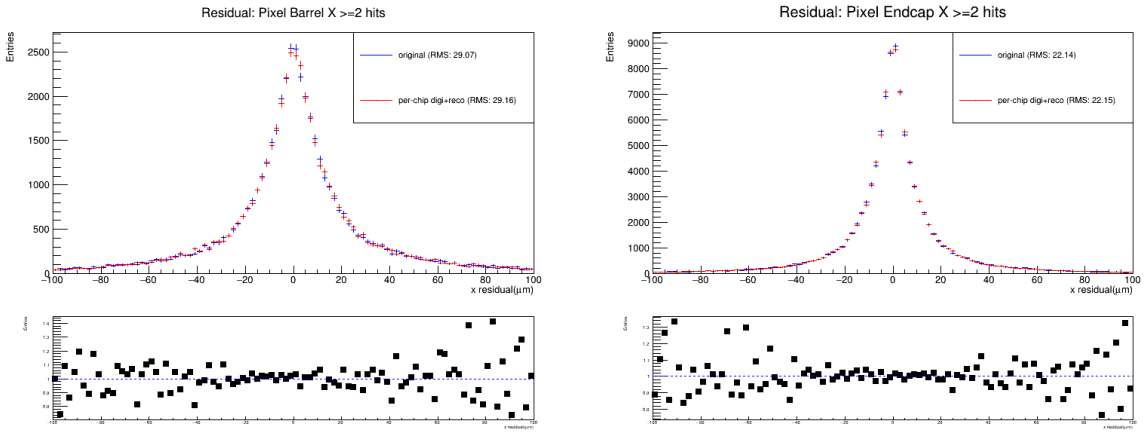
4. **Per-pixel DIGI + Per-chip RECO**: per-pixel calibration applied at digitization but only per-chip calibration applied at reconstruction, both using the look-up table derived in this study.

The configurations differ in the granularity at which the calibration function ToT = $f(Q)$ is applied during digitization (DIGI) and its inverse $Q = f^{-1}(\text{ToT})$ during reconstruction (RECO). Configuration 1 serves as a baseline reference for the original charge reconstruction. Configurations 2 and 3 apply the same calibration at both DIGI and RECO, and are therefore expected to give the same tracking performance as the baseline. Configuration 4 uses a per-pixel calibration for DIGI but a FE-level calibration for RECO, which could show a decrease in performance due to the charge resolution loss.

Tracking performance is evaluated using the hit position residual. Cluster position resolution and tracking efficiency can be used as complementary higher-level metrics, but are

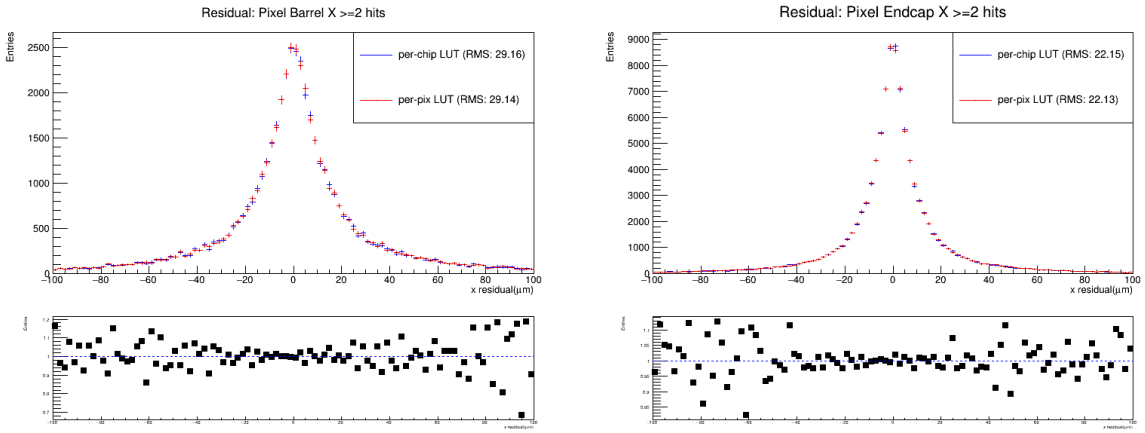
not shown here. The x-position residual in the pixel barrel and endcap is shown in Figures C.106 to C.108 for the three comparisons. The residual distributions are in good agreement across all configurations.

An explanation of this result may be that the hit position residuals shown here are the *Track–Cluster* residuals, defined as the difference in hit position between the cluster reconstruction and the track fit, which only reflects the track fit performance. To understand the true impact of the charge calibration granularity, a *Cluster–Truth* residual would be needed in future work.



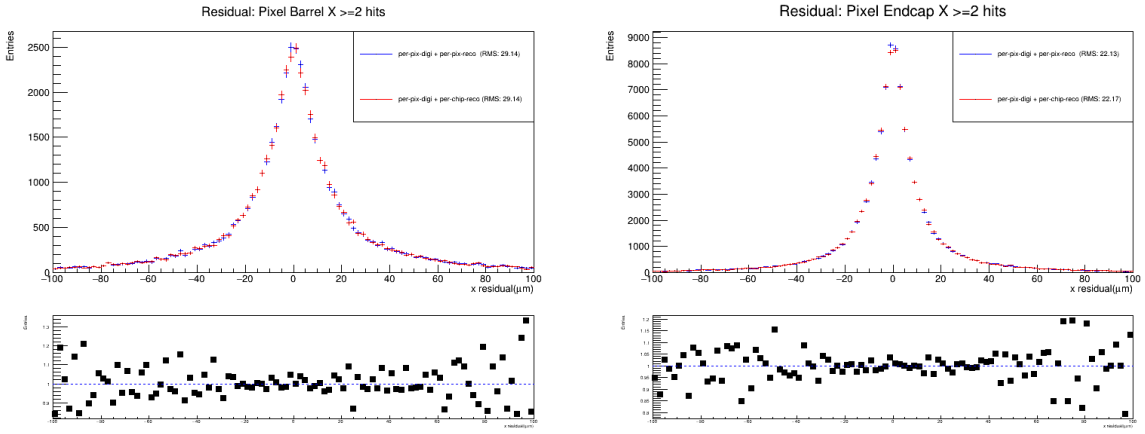
(a) (b)

Figure C.106. Pixel hit position residual in x, comparing the baseline (**Original DIGI + Original RECO**) and the per-chip calibration (**Per-chip DIGI + Per-chip RECO**).



(a) (b)

Figure C.107. Pixel hit position residual in x, comparing the per-chip (**Per-chip DIGI + Per-chip RECO**) and per-pixel (**Per-pixel DIGI + Per-pixel RECO**) calibrations.



(a) (b)

Figure C.108. Pixel hit position residual in x, comparing the matched per-pixel configuration (**Per-pixel DIGI + Per-pixel RECO**) and the mismatched configuration (**Per-pixel DIGI + Per-chip RECO**).

C.5 Challenges for future investigations

In summary, a potential loss in position resolution is difficult to isolate, as it is largely masked by several other effects in the reconstruction chain. A list of technical challenges is outlined below to help guide future investigations:

1) Analog clustering uses the charge information not only to determine the cluster position, but also to correct the hit position according to the track incident angle. This step is referred to as the *on-track* calibration.

2) During Run 2 and Run 3, the on-track calibration is performed with neural networks (NNs): a first NN determines whether a cluster originates from a single particle or is shared between multiple particles by predicting the number of hits in the cluster, after which a second NN, selected according to the predicted number of hits, calibrates the hit positions.

3) The same on-track calibration is planned for Run 4. However, the corresponding NNs have not yet been trained, so placeholder calibration constants are currently used for this step.

4) In ATHENA, only three clustering options are available: digital clustering, analog clustering with incident-angle correction, and the NN-based approach. Performing analog clustering *without* the incident-angle correction is not currently possible in ATHENA, although it may be feasible within ACTS.

APPENDIX 4

Non-resonant Anomaly Detection with Background Extrapolation

The work presented in this appendix was published in the Journal of High Energy Physics [68], in collaboration with Radha Mastandrea and Benjamin Nachman. It has been modified to fit the format of an appendix for this dissertation.

Complete anomaly detection strategies that are both signal sensitive and compatible with background estimation have largely focused on resonant signals. Non-resonant new physics scenarios are relatively under-explored and may arise from off-shell effects or final states with significant missing energy. In this paper, we extend a class of weakly supervised anomaly detection strategies developed for resonant physics to the non-resonant case. Machine learning models are trained to reweight, generate, or morph the background, extrapolated from a control region. A classifier is then trained in a signal region to distinguish the estimated background from the data. The new methods are demonstrated using a semi-visible jet signature as a benchmark signal model, and are shown to automatically identify the anomalous events without specifying the signal ahead of time.

D.1 Introduction

Despite the impressive predictability of the Standard Model (SM), it is a well-known fact that it does not account for all known phenomena and is plagued with a number of aesthetic issues. As such, the search for new, fundamental interactions is a central goal of particle physics across virtually every experiment. However, given that there is an uncountable number of new physics models, we simply do not have the bandwidth to test all possibilities, and it is certain that there are alternative hypotheses that we have not yet considered.

For this reason, a new paradigm has emerged to complement traditional, model-specific approaches. The new *anomaly detection* protocols seek to explore data with as little bias as possible in order to be broadly sensitive to many scenarios. Such protocols have been significantly advanced by modern machine learning (see [69–72] for a selection of machine learning reviews and anomaly detection challenges, and [73–170] for specific techniques). One powerful anomaly detection strategy ("weak supervision") is to isolate a region of phase space and compare data to a background-only reference. Signal model-agnostic evidence for new physics emerges when the data is not statistically consistent with the reference.

The core task of weakly supervised anomaly detection strategies is to construct the reference sample of background-only events. Existing strategies follow one of two approaches: (1) the reference sample is from (background-only) simulation and (2) the reference sample is learned from sidebands. The advantage of the simulation approach is that it puts few assumptions on the form of the signal. In particular, this approach can accommodate signals that are non-resonant, either due to the new physics being very massive or consisting of non-reconstructable (e.g. invisible) particles. The challenge with a simulation-based reference is that the simulation accuracy limits the anomaly detection sensitivity. In contrast, sideband methods posit that the signal, if it exists, is localized in at least one known dimension m. Focusing on a particular interval in m, a signal-sensitive region of phase space is then constructed using other features x, and regions in m away from the posited resonance are used to estimate the background $p_{\text{background}}(x|m)$. The benefit of sideband methods is that the reference is estimated directly from data, while the challenge is that not all signals are resonant.

We propose an approach to merge the advantages of the simulation-based and sideband-based cases to form a strategy that learns the reference from data without requiring a resonance. This method extends and modifies existing resonance approaches [87, 88, 102, 125, 153, 160, 171] to allow for extrapolation in m, as opposed to interpolation. Instead of the signal region being defined by an interval in a one-dimensional m, we now have a multidimensional m for

which the signal region is defined by thresholds[1] $m_0 > c_0, m_1 > c_1, ..., m_n > c_n$ for $m, c \in \mathbb{R}^n$. The conditional reference $p_{\text{background}}(x|m)$ is learned (directly or indirectly) from the complement of the signal region and then extrapolated to the signal region. If the machine learning models are smooth and the background remains qualitatively the same in the signal region (e.g. the functional form of the background distribution, particularly the dependence on the context variables, is continuous across the CR-SR boundary), there is reason to believe that the extrapolation can be accurate[2]. A new feature of the non-resonant case is that we need to additionally estimate $p_{\text{background}}(m)$. This is also needed in the resonant case, but it is less important (since $p(m)$ is relatively constant in the signal region) and thus approximate or simplified methods have been considered so far. We propose to estimate $p_{\text{background}}(x|m)$ using a likelihood-ratio method [87] and then combine this estimate with three proposals for estimating $p_{\text{background}}(m)$ based on reweighting [87], generative models [125], and template morphing [171].

The growing anomaly detection for fundamental interactions literature includes a number of proposals that can accommodate non-resonant new physics. In the weakly supervised case, simulation-based background estimates do not require resonant new physics [74, 76, 135, 145]. Nearly all data-based weakly supervised methods include non-resonant x features [87, 88, 102, 125, 153, 160, 167, 171]. In special cases, symmetries can be used to define the reference directly in data without requiring m to be resonant [131, 132, 143, 147]. Extrapolation of generative models has been considered for background estimation, but combined with classical model-specific search strategies [173]. Unsupervised methods like those based on autoencoders (e.g. Ref. [77, 78, 82] and many others) do not require resonances for achieving signal sensitivity in general. However, except in special cases, they often lack background estimation strategy [133],

[1] A non-rectangular signal region is also possible, but makes the accounting more difficult.

[2] We leave additional studies to impose minimalism in the extrapolation to future work via e.g. monotonicity [172].

or any guarantees of optimality. Our approach is unique because it combines signal sensitivity with background estimation[3] and is asymptotically optimal [88] in the limit of large datasets and effective machine learning.

Non-resonant signals can arise from many new physics models. For example, models of a strongly interacting dark sector can produce jets which contain both stable and unstable dark hadrons, predicting signatures of jets containing significant missing energy. Non-resonant signals are also characteristic of effective field theories, which can model new physics with masses beyond those accessible through on-shell production. ML-based anomaly detection can offer additional sensitivity to explore the vast signature space of non-resonant signals.

This paper is organized as follows. In Sec. D.2, we introduce the methodology we follow for our tail-based non-resonant anomaly detection procedure, providing details on the background extrapolation task. In Sec. D.3, we show a realistic physics example, considering jets emerging from a complex dark sector as our non-resonant signal. We conclude in Sec. D.4.

D.2 Methodology

D.2.1 Locating overdensities

The goal of a weakly-supervised anomaly search is to identify regions of phase space that are more represented by data (forming overdensities) than they are by the reference dataset. The search is performed in a signal region (SR) and the reference is estimated using information from a neighboring region.

In the resonant case, the SR is defined by an interval around a posited new particle mass $m \in [m_0 - \delta, m_0 + \delta]$. The neighboring region (called *sideband* or SB) is the region (or a subset of the region) outside of this interval. If the signal exists and is at mass m_0, we expect that most of

[3]The reference does not need to be used directly to estimate the background after making a cut on the anomaly score. One could use any classical non-resonant background estimation strategy like the ABCD method, possibly also with machine learning [174].

the signal is in the SR with little contamination in the SB. The reference background template is estimated using information from the SB. Since we expect the phase space to be smoothly varying with m and since δ is expected to be small relative to regions over which large changes occur in $p_{\text{background}}(x|m)$, we expect the background to be well-estimated by interpolating from the SB (see further details in Sec. D.2.2). Given samples from the background dataset[4], overdensities are identified with the Classification WithOut LAbels procedure (CWoLa) [73, 75, 175]. Events in data are all assigned the label "signal" and all events from the background template are labeled "background". An optimal classifier trained to distinguish these datasets will implicitly learn a function monotonically related to the signal-over-background likelihood ratio. Thresholding the classifier will then isolate an anomaly-like region of phase space.

In the non-resonant case, the setup is slightly different. Instead of having SB that surround the SR, the neighboring region is only on one side and is now called a *control region* or CR. However, we allow for m to be multidimensional – the more dimensions available to constrain $p_{\text{background}}(x|m)$, the better accuracy we might expect for reconstruction in the SR[5]. Resonances are in fact nearly a special case of this setup, since one could define e.g. $m' = 1/|m - m_0|$. As in the resonant case, we start by constructing background templates in the CR, and we expect most of the signal to be in the SR so that the CR can be used to estimate $p_{\text{background}}$ without bias. A schematic diagram of the non-resonant setup is presented in Fig. D.109. The challenge now is to extrapolate into the SR instead of interpolating as in the resonant case. Given the background template, the actual anomaly detection proceeds using the same CWoLa procedure described above. The main aspect that distinguishes different approaches is how the background templates are generated.

[4]There are variations on this where the density is used directly without samples [88], but the overall idea is the same.

[5]We validated this with simplified Gaussian samples.

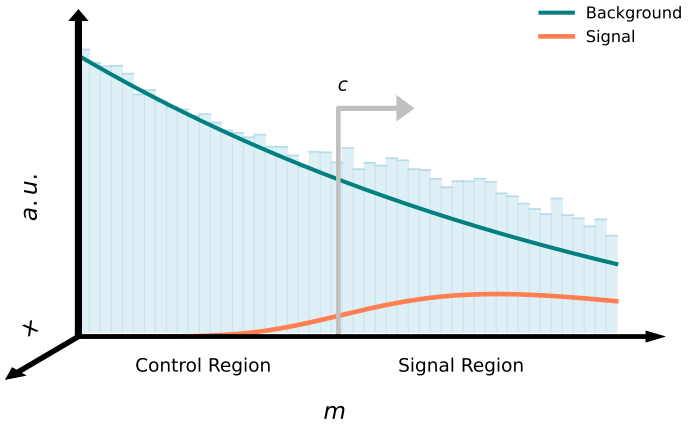


Figure D.109. A schematic of the setup for non-resonant anomaly detection. The signal region (SR) is defined by a one-sided cut on a context variable $m > c$. While the figure shows a one-dimensional context variable, in practice the context can be multidimensional. A number of other features x are used for the CWOLA anomaly detection classifier. It is also possible to include m in the classifier.

D.2.2 Constructing background templates

To construct samples of background-only events in the SR, we take the following steps:

1. **Estimate $p_{\text{data}}(x|m)$ in the CR.** We first learn the distribution of the chosen features for background-only events, conditioned on the context variables, in the CR. It is assumed that the conditioning on the context will allow the learned distribution to be extrapolated into the SR. For the interpolation case (i.e resonant anomaly detection), the analogous first step would be to learn $p_{\text{data}}(x|m)$ in the SB. Since we assume that the signal is mostly absent from the CR, $p_{\text{data}} \approx p_{\text{background}}$.

2. **Estimate $p_{\text{background}}(m)$ in the SR.** In order to sample from $p_{\text{background}}(x|m)$ in the SR, we need to be able to generate SR-like context m. In the interpolation case, this might be done by performing a parametric fit to a histogram of m in the SB and interpolating to the SR. This could also work in the extrapolation case, but (1) the shape may be non-trivial in the tails of m and (2) this approach does not scale well to more than one dimension.

To circumvent these challenges, we estimate $p_{\text{background}}(m)$ in the SR by reweighting context in the CR, $p_{\text{simulation}}(m)$, to match data and then extrapolating this function into the SR. A binary classifier parameterized in m [176, 177] is trained to distinguish simulated context from data context in the CR. The classifier output is then interpreted as a likelihood ratio $w(m)$ for estimating the background [87]. As in step 1, we assume that conditioning the network on the context will enable the learned weights to extrapolate into the SR.

3. **Sample $p_{\text{background}}(x|m)$ in the SR**. We first estimate $p_{\text{background}}(x|m)$ from the learned model in the CR, drawing context from $p_{\text{simulation}}(m)$, in the CR. The resulting events are then weighted by $w(m)$ to arrive at our estimate of $p_{\text{background}}(x,m)$ in the SR, following Eq. D.1.

$$\begin{aligned} p_{\text{background}}(x,m) &= p_{\text{background}}(x|m)\, p_{\text{background}}(m) \\ &= p_{\text{background}}(x|m)\, p_{\text{simulation}}(m) \frac{p_{\text{background}}(m)}{p_{\text{simulation}}(m)} \\ &= p_{\text{background}}(x|m)\, p_{\text{simulation}}(m)\, w(m) \end{aligned} \tag{D.1}$$

We consider three approaches for approximating $p_{\text{data}}(x|m)$ in the CR, based on three techniques previously developed for the resonance case: Simulation-Assisted Likelihood-Free Anomaly Detection (SALAD) [87], Classifying Anomalies through Outer Density Estimation (CATHODE) [125], and Flow-Enhanced Transportation for Anomaly detection (FETA) [171]. Our proposed methods are:

1. **Reweight** (SALAD-inspired). The parameterized, binary classifier used to approximate $p_{\text{background}}(m)$ is extended to include also x so that the likelihood ratio in both x and m is estimated at the same time.

2. **Generate** (CATHODE-inspired). A generative neural network (normalizing flow [178]) is trained to model data in the CR, conditioned on the context m. Given values of m, one can sample features x from the normalizing flow.

3. **Morph** (FETA-inspired). A normalizing flow-based model is trained to map samples in the CR from MC to data, conditioned on the context m. Like SALAD, FETA starts from a simulation, but instead of reweighting, it morphs the features directly [179, 180].

Note that $w(m)$ is estimated with SALAD in all three cases, so our new approaches are SALAD-{SALAD, CATHODE, FETA} hybrid methods.

We evaluate the performance of the Reweight, Generate, and Morph-constructed background samples against two benchmarks: (1) a fully supervised classifier trained to discriminate pure background from pure signal, and (2) an "idealized" classifier trained to discriminate pure background from true background + signal. The idealized classifier represents the actual best performance to which our background estimation methods should asymptote.

D.2.3 Network architectures and training specifications

All networks are implemented in PyTorch [181] and optimized with Adam [182]. For the Generate and Morph methods, we construct normalizing flow networks with the nflows package [183]. All models are trained with a train-validation split of 2/3–1/3 and are evaluated at the epoch of lowest validation loss. For the Reweighting, Generate, and Morph methods, architectures and hyperparameters were optimized through manual tuning such that in the case of zero signal injection, a binary classifier could not distinguish reconstructed background in the CR from data in the CR. This was necessary in order to keep the SR blinded.

The Reweighting method is implemented with a binary classifier, consisting of a dense neural network with 3 layers of 100 nodes each. We use a batch size of 512, a learning rate

of 10^{-3}, and train for up to 50 epochs with an early stopping patience of 5. The classifier that generates the context weights $w(m)$ is exactly the same as that for the Reweighting method. The Generate method is implemented with 4 layers of (Masked Affine Autoregressive Transform block of 128 hidden features, Reverse Permutation), where the transform block is built from a Masked Autoencoder for Distribution Estimation (MADE) architecture. Training is with a batch size of 216, a learning rate of 10^{-3}, and is carried out for up to 20 epochs with an early stopping patience of 5 epochs. For the Morph method, the base density flow uses the same architecture as the Generate method. The transfer flow is similar, but layers have only 64 hidden features, the batch size is reduced to 128, and the learning rate to 10^{-4}.

For the anomaly detection classifier networks, as well as the networks used for the closure tests (see Sec. D.3.2.1), we use a fully-connected architecture with 3 layers of 64 nodes each. We use a batch size of 512, a learning rate of 10^{-3}, and train for up to 50 epochs with an early stopping patience of 5 epochs. Networks are only trained on the non-context features. In order to ensure training stability, for a given set of generated samples, we crop the weights $w(m)$ by dropping all weights that are greater than 3 standard deviations above the mean. This has the effect of removing a very small number (usually on the order of a few tens, or $\sim 0.1\%$) of samples whose weights severely bias the resulting $p_{\text{background}}(x,m)$ distributions.

For the dataset, we consider a two-dimensional context m and a five-dimensional feature space x (to be explained in detail in Sec. D.3.1). The values of all seven variables are minmax-scaled to the range $(-2.5, 2.5)$ based on the minimum and maximum of the simulation (pure background) dataset before any network training is carried out. This was found to give more reliable flow training. When the number of signal events is small, the exact events that are injected into the "data" is important. Therefore, we rerun the procedures with 10 random signal injections and report the median and a 68-percentile spread. For each signal injection, we retrain the models with a different random initialization 20 times and average over these

models. We found (as in previous studies) that this ensembling helps with sensitivity and stability.

D.3 Application with dark QCD jets

This section illustrates the use of the above background extrapolation method in the case of finding Beyond the Standard Model (BSM) physics at the LHC. The physics model we are interested in is jets that contain both stable and unstable hadrons from a strongly interacting dark sector, or dark QCD. This type of signal comes from e.g. Hidden Valley models [184–187]. In particular, we assume there exist some new massive quarks charged under dark QCD that are promptly produced and go through QCD-like showering and hadronization to form "dark hadrons". It is possible that a fraction of these dark hadrons decay back to the SM hadrons, and the rest remain invisible. This results in the type of dark QCD signature called semi-visible jets [188–193]. Semi-visible jets are challenging to distinguish from SM QCD jets, but the unique signature from certain dark hadron mass scales and missing energy may be used for identification.

We assume that the dark quarks are produced through a mediator such as a massive gauge boson Z'. When the mediator is produced on-shell, we can reconstruct the invariant mass of the mediator from the semi-visible jets. However, when the invisible fraction of the jet is large, there will be a large amount of missing energy, resulting in a poor invariant mass resolution so that the reconstructed signal is no longer resonant. The invisible fraction of dark hadron decays is defined using the parameter, r_{inv}, the ratio of the number of stable dark hadrons over number of total dark hadrons.

For this study, we choose a 4 TeV Z' decay into a pair of dark quarks. There are 3 flavours of dark quarks, all with the same mass of $m_{q_{\text{D}}} = 250$ GeV. There are two types of dark hadrons, the dark pion and dark rho meson, whose masses are set to be the same as the confinement scale of the dark sector, $m_{\pi_{\text{D}}} = m_{\rho_{\text{D}}} = \Lambda_{\text{D}} = 500$ GeV. Note that this choice of Λ_{D} results in a small

number of dark mesons in the shower, and since the dark mesons are relatively heavy, the jet substructure comes from the dark meson decays and is less dependent on the hadronization process in the dark sector. A study of the effects from different mass parameters is presented in Appendix D.5. The AD procedure is insensitive to the simulation modeling, and therefore is complementary to methods that minimizes the hadronization uncertainties [194, 195].

D.3.1 Dataset specifications

Signal events are generated using the Hidden Valley module from PYTHIA 8.310 [196]. Background ("data" and "simulation") events are generated using MadGraph5 aMC@NLO [197] and showered using PYTHIA 8.310. All events use DELPHES [198] for detector simulation with the CMS card, and jets are clustered using the anti-k_t [50] algorithm with $R = 1.0$ using FASTJET [199]. For the background samples, the data and simulation events differ in the choices of renormalization and factorization scales (1 vs. 2, respectively) as well as the tuning (14 vs. 25). We show results using a different background simulation in Appendix D.6 to validate the robustness of the anomaly detection (AD) methods against larger discrepancies in data and simulation. A minimum of 200 GeV transverse momentum is required for the leading jet in the event. In Fig. D.110, we show the m_{jj} distribution for a few selections of r_{inv} values. To ensure the non-resonance of the signal, we set $r_{\text{inv}} = 1/3$ in this study.

We use the scalar sum of jet transverse momenta (H_{T}) and missing transverse momentum (MET) as the context variables to define a two-dimensional SR[6]. For a five-dimensional feature space, we use the dijet invariant mass m_{jj}, and the two-pronged and three-pronged N-subjetiness [200, 201] for both the leading and subleading jets. Distributions of context and feature variables for signal and background is shown in Fig. D.111. The SR is defined by the

[6]These were chosen since they are nearly independent. This is not strictly necessary for the method to work, but we found that it improves the extrapolation quality and the same features can be used for a classical matrix method / ABCD background estimate.

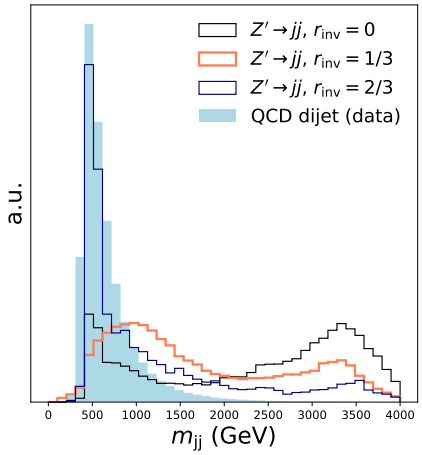


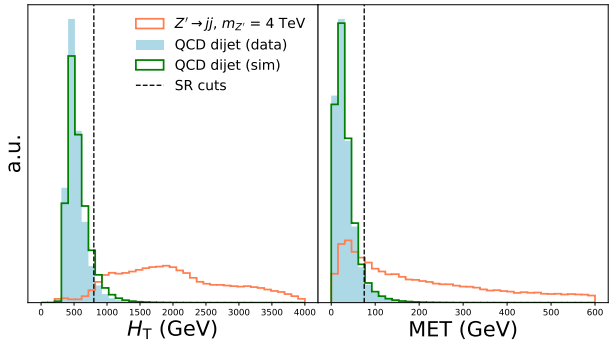
Figure D.110. A histogram of m_{jj}, corresponding to the reconstructed mass of the 4 TeV Z' particle, for a selection of choices within the Hidden Valley model with parameter r_{inv}. In this study, we set $r_{\text{inv}} = 1/3$, at which point the Z' signal is no longer clearly resonant.

cuts $H_{\text{T}} > 800$ GeV, and MET > 75 GeV, which ensures the signal-over-background ratio is much lower in the CR than that of the SR[7]. We summarize the number of generated events in Tab. D.29.

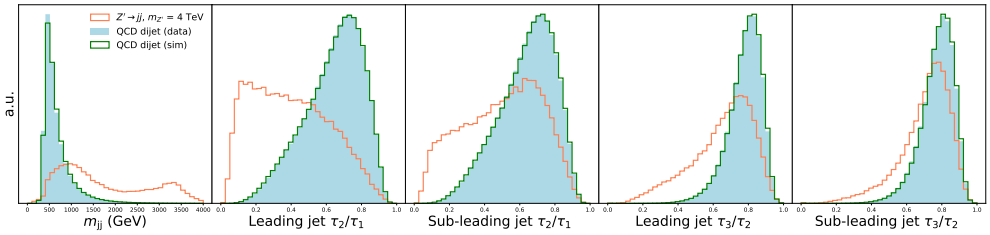
Table D.29. Breakdown of events generated for the Hidden Valley model. We found that large training sets in the CR were necessary for good performances of the Generate and Morph (the flow-based) methods, while the Reweight method worked well even with a smaller set. We did not explore oversampling in the SR, which may improve performance [125].

Purpose	Dataset	Number CR events	Number SR events
Training for generative models	Simulation (bkg.)	9,983k	126k
	Data (bkg.)	9,247k	72k
	Data (sig.)	14k	50k
Generated samples	Reweight Generate Morph	N/A	126k
Evaluation	Ideal AD samples	N/A	68k
	Fully supervised set	N/A	13k bkg., 30k sig.
	Test set	N/A	150k bkg., 20k sig.

[7] In addition to the selection of features, the definition of CR and SR will introduce some model dependence. We expect that these regions will work well for a wide range of physics model parameters, especially when the Z' mass is higher.



(a) Two-dimensional context variables. The SR is defined by the cuts $H_T > 800$ GeV, and MET > 75 GeV, which are shown on the plot.



(b) Five-dimensional feature variables.

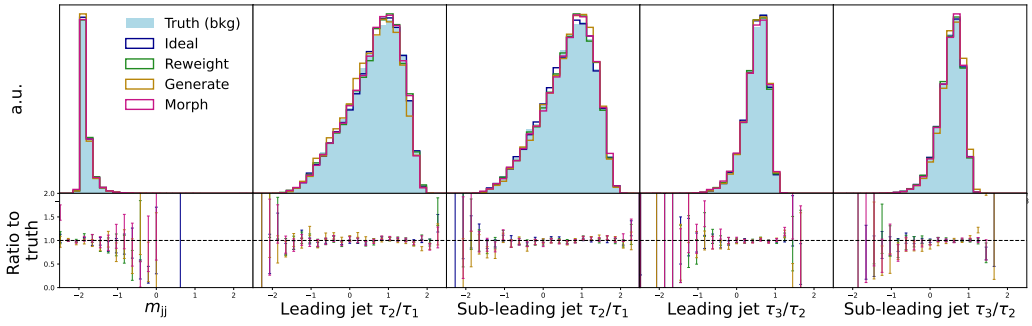
Figure D.111. Histograms of the observables used in the Hidden Valley non-resonant anomaly detection task, for background ("data" and "simulation") and signal events).

D.3.2 Results

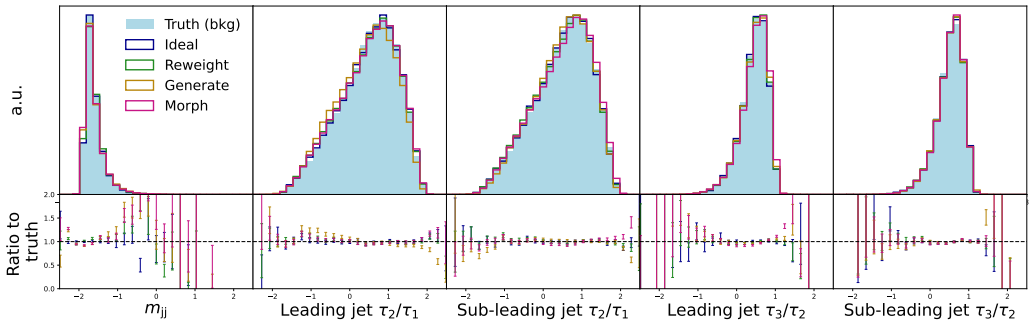
D.3.2.1 Closure tests

As a first test of extrapolation, we check that we are able to generate realistic background samples in both the CR and SR under zero-signal injection. In Fig. D.112, we show the five feature distributions as constructed by the Ideal, Reweight, Generate, and Morph methods in relation to truth, the zero-signal data. Feature reconstruction is excellent in the CR, and the ratios of the marginals to the truth are consistent with unity across virtually all of the domains of the observables. Reconstruction is good overall in the SR, although the reconstructed distributions (particularly for the Morph method) deviate visibly. We calculated the one-dimensional Wasserstein distances as a quantitative measurement of the closure test, characterizing how well the observables are reconstructed in the signal region and control region by Reweight, Generate and Morph methods. The average Wasserstein distances of Ideal, Reweight, Generate, and Morph across five feature variables are 0.0048, 0.0073, 0.0237, 0.0116 in CR, and 0.0058, 0.0130, 0.0293, 0.0202 in SR, respectively. Note that in order to avoid premature unblinding, we withhold a set of 10k simulation and 10k data events in the CR from network training and use this set to evaluate CR closure.

In Fig. D.113a, we show a more rigorous test by training a binary classifier to discriminate between the constructed samples and background-only data. We show the spread of ROC AUCs (receiver operating characteristic area-under-curves) across 10 binary classifier runs; for a classifier trained on two identical datasets, the ROC AUC is expected to be 0.5 for a perfect classifier, in the limit of infinite statistics and training time. Again, we see that there is generally excellent reconstruction in the CR and slightly worse reconstruction in the SR. In Fig. D.113b, we show the classifier rejection (1/false positive rate) curves for the discrimination task in the SR. For all of the methods, the rejection found in the zero-signal case is compatible (to within errorbars) with that of a random classifier.

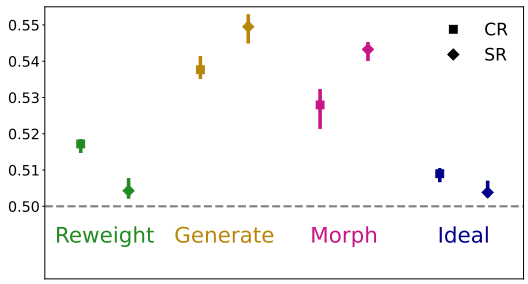


(a) Control region distributions.

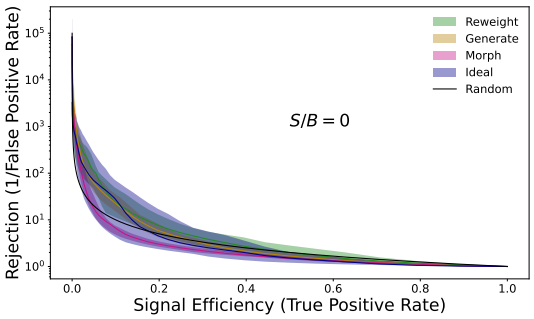


(b) Signal region distributions.

Figure D.112. Feature distributions as constructed by the Reweight, Generate, and Morph methods in the absence of signal. Errorbars in the ratio plots reflect Poisson statistics. Note that the values of the context and features have all been minmax-scaled to the range (-2.5, 2.5).



(a) ROC AUC spread in the CR and SR. The horizontal line is at 0.5, marking true random.



(b) Rejection in the SR.

Figure D.113. Results for a classifier tasked with discriminating extrapolated background samples from true background. The "Ideal" corresponds to an idealized classifier, which has been trained to discriminate two statistically identical sets of background, and thus represents a realistic random classifier. Errorbars and uncertainty bands show a 68-percentile spread across 20 independent classifier runs.

As stated in Sec. D.2.3, architectures and hyperparameters for the extrapolation networks were selected so as to make the spread of ROC AUCs in the CR consistent with random. There is a (seemingly unavoidable) systematic error associated with the extrapolation, such that Reweight, Morph, and Generate methods show a drop in similarity to true background in the SR. We believe that most of this drop can be attributed to the more general challenge of extrapolation for neural networks.

D.3.2.2 Signal detection tests

In this section, we carry out an anomaly detection test for the dark QCD physics model. The test is done by generating a signal-background discriminator using a weekly-supervised binary classifier following the CWoLa procedure, as explained in Sec. D.2. The classifier is trained to discriminate the background reference against data that consists of SR background with a known amount of signal injection. For the simulation, we scan over seven signal injections of signal (S) over background (B) of $S/B \in \{0, 3.1, 6.5, 9.3, 12, 15, 18\} \times 10^{-3}$, corresponding

to significances $S/\sqrt{B} \in \{0, 0.84, 1.75, 2.52, 3.32, 4.07, 4.99\}$, respectively[8]. In this section, we provide a number of metrics derived from this binary classifier.

We focus on metrics derived from the classifier SIC curves for the discrimination task in the SR. The SIC, defined as the true positive rate/$\sqrt{\text{false positive rate}}$, gauges the factor by which a signal's significance would improve by making a cut on the classifier score. We expect SIC $\gg 1$ for a well-performing classifier, while in the zero-signal case, the classifier would perform poorly.

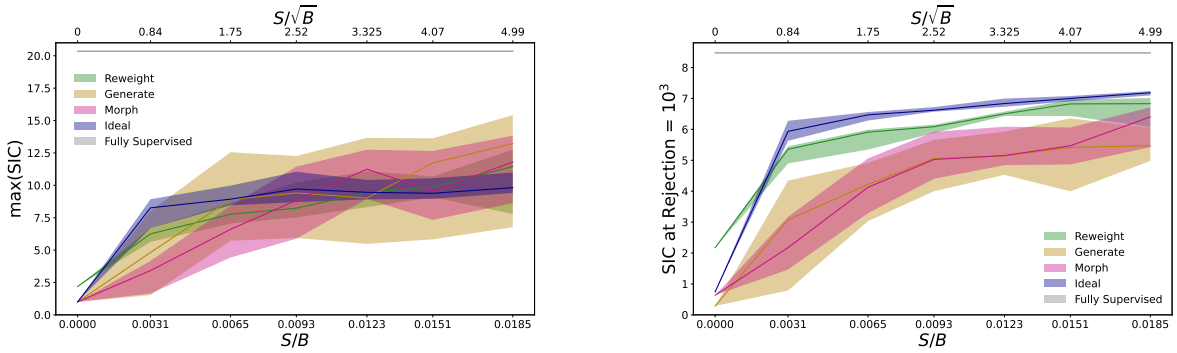
In Fig. D.114a, we show the maximum SICs across all signal efficiencies, corresponding to the largest achievable significance improvements. In Fig. D.114b, we show SICs evaluated at a fixed rejection (given by 1/false positive rate) of 10^3. The results are encouraging, demonstrating comparable performance to the idealized anomaly detector for all three methods. The Reweight method enhances a signal significance from below $1 (\sim 0.84)\sigma$ to the 5σ "discovery" threshold. The Generate and Morph methods perform less optimally with respect to the metric of SIC at a rejection of 10^3, but there is good evidence that the Generate and Morph methods could enhance a signal significance originally between 1–2σ to the discovery threshold. At zero signal injection, the higher SIC values for Reight is likely due to randomness, since there is nothing for the network to learn. What we see is that the mismodeling happens to be in a signal-like direction. It would be interesting for future studies to explore how these findings change for different signal and background models[9].

D.4 Conclusions

In this paper, we have extended background interpolation methods to the extrapolation case to enable weakly supervised, non-resonant anomaly detection (AD). Our strategy is

[8] Note that with our chosen context region bounds, there is a small amount of signal contamination in the CR. For the largest signal injection, the contamination is 4.3×10^{-3}%, from 384 events.

[9] See Appendix D.5, where we show how the detection task changes when using a different set of signal parameters.



(a) Maximum of the Significance Improvement Characteristic.

(b) Significance Improvement Characteristic evaluated at a rejection of 10^3.

Figure D.114. Metrics for an anomaly detection classifier tasked with discriminating extrapolated background samples from data in the SR. Uncertainty bands represent the 68-percentile spread over 10 different signal injections, all of which have been score-averaged over 20 independent classifier runs.

based on the construction of a realistic set of background events through extrapolation from a background enriched control region (CR) into a signal region (SR). We proposed three extrapolation approaches based on reweighting, generation, and morphing which parallel the SALAD, CATHODE, and FETA methods, respectively. Studies in the resonant case have shown complementarity between these methods [162] which may be especially important in the extrapolation case.

The proposed methods are tested on a realistic BSM signal consisting of dark QCD jets. We found that in the zero-signal case, there was good closure between the reconstructed samples and the background data, showing that our background templates were well-modeled and unbiased. The agreement between predicted background and the true background in CR implies our methods are optimized for interpolation, and deviation in the SR might come from the intrinsic difficulty of extrapolation. Finally, we have also shown the performances of all the methods in a realistic AD task, demonstrating their abilities to elevate non-significant signal injection to the threshold of possible detection. While we only showed the performance for a

single physics signal model, we stress that the protocol was optimized on only the background, and the sensitivity should extend to a broad class of new physics scenarios.

Our study, while encouraging, brings to light some of the many challenges associated with non-resonant AD. Future work is necessary in order to develop more robust non-resonant AD. Possible avenues may involve using networks or data preprocessing techniques that are especially well-suited for extrapolation. It is also worth testing the learning of $w(m)$ function given multiple different simulation sets that are qualitatively different from the data, in order to verify the effectiveness of the Reweight method[10].

With background extrapolation methods, the range of possible signals that could be detected using an overdensity analysis would be greatly expanded. Not only are we able to explore a larger parameter phase space of known resonance signals that was previously not reconstructable, but we are also able to access non-resonant off-shell effects from heavy particles. It is likely that no one method will achieve the best sensitivity to all scenarios, so a set of techniques are required to achieve broad coverage.

Data and code availability

The physics datasets and parameter cards used for all simulations have been made available on Zenodo at `https://zenodo.org/records/10154213`. The analysis code is available at `https://github.com/hep-lbdl/non-resonant-AD-extrapolation`.

D.5 Comparing different signal parameters

In this section, we compare the performance of three proposed non-resonant AD methods (Reweight, Generate, and Morph) on different signal parameters of the dark QCD model. We

[10] See Appendix D.6 for an example of a modified simulation set.

choose the following two sets of parameters, both with $r_{\text{inv}} = 1/3$ in addition to the 4 TeV Z' signal used in Sec. D.3:

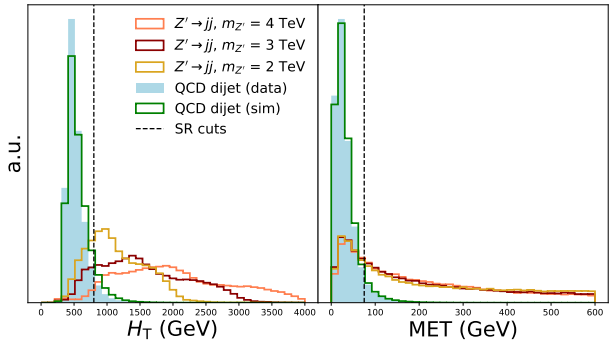
- $m_{Z'} = 2$ TeV, $m_{\pi_{\text{D}}} = m_{\rho_{\text{D}}} = \Lambda_{\text{D}} = 200$ GeV, and $m_{q_{\text{D}}} = 100$ GeV.
- $m_{Z'} = 3$ TeV, $m_{\pi_{\text{D}}} = m_{\rho_{\text{D}}} = \Lambda_{\text{D}} = 300$ GeV, and $m_{q_{\text{D}}} = 150$ GeV.

These two sets of parameters, denoted as the 2 TeV and 3 TeV signal, respectively, are chosen to have different m_{jj} and H_{T} distributions from the 4 TeV signal, with more signal located towards the bulk of the background. The choice of $m_{\pi_{\text{D}}}$, $m_{\rho_{\text{D}}}$, Λ_{D} and $m_{q_{\text{D}}}$ values gives a similar two-pronged and three-pronged jet substructures as the 4 TeV signal in Sec. D.3.

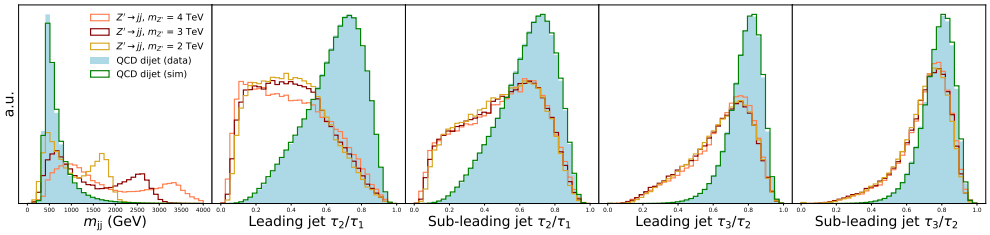
Fig. D.115 shows the distributions of context and feature variables of three sets of signal parameters used in this study. We apply the same SR definition with $H_{\text{T}} > 800$ GeV and MET > 75 GeV. Note that the distributions of the context variable H_T and the feature variables m_{jj} for the 2 TeV signal look much more similar to those of background, so we expect the discrimination task to be harder for lower signal masses.

In Figs. D.116 and D.117, we show the results for an anomaly detection classifier working on a range of signal detections for the 2 TeV and 3 TeV signals, respectively. (These plots are analogous to Fig. D.114 for the 4 TeV signal.) As expected, the performances of all AD methods show low sensitivity to the 2 TeV signal, with Reweight achieving a similar significance improvement as the idealized classifier. Performance is better across the board for the 3 TeV signal, although the qualitative features are the same as for the 2 TeV signal.

For both 2 TeV and 3 TeV signals, the gap in performance between the idealized classifier and the fully supervised classifier is much larger than it is for the 4 TeV signal. This difference might be ascribed to the fact that the lower-mass signals have a greater resemblance to the background process in variables such as H_T and m_{jj}. Additionally, the amount of signal contamination in the CR is higher for the lower mass Z' signals: for the 2 TeV Z', the highest signal injection has a CR contamination of 0.01% (1257 events), and for the 3 TeV Z', the highest

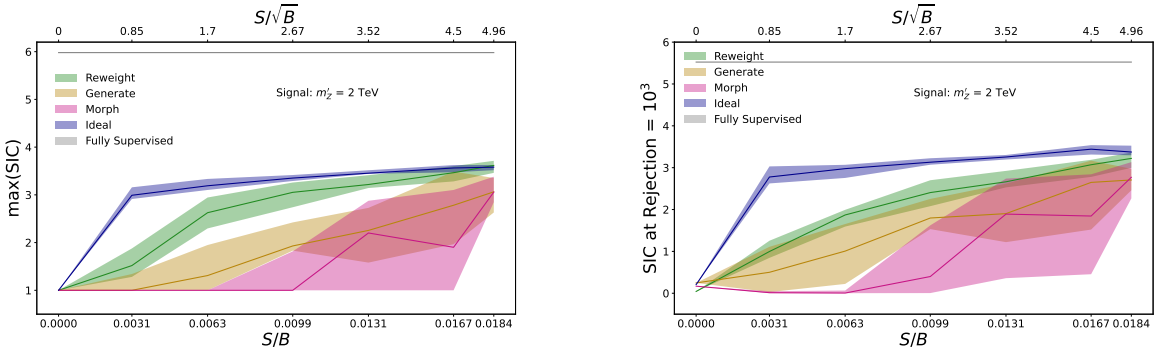


(a) Two-dimensional context variables



(b) Five-dimensional feature variables

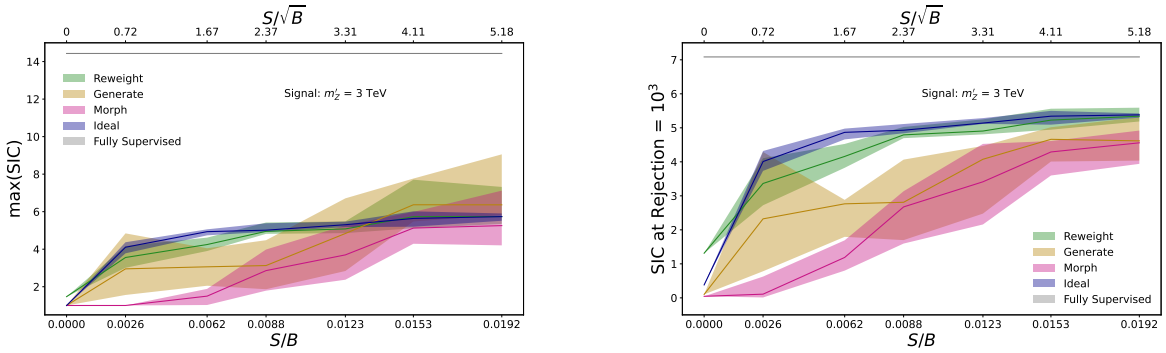
Figure D.115. Histograms of the observables used in the Hidden Valley non-resonant anomaly detection task, for background ("data" and "simulation") and signal events).



(a) Maximum of the Significance Improvement Characteristic.

(b) Significance Improvement Characteristic evaluated at a rejection of 10^3.

Figure D.116. Metrics for an anomaly detection classifier tasked with discriminating extrapolated background samples from data in the SR. The background is QCD dijet events. The signal is a 2 TeV Z' producing two dark QCD jets with $2m_{q_D} = m_{\pi_D} = m_{\rho_D} = \Lambda_D = 200$ GeV. Uncertainty bands represent the 68-percentile spread over 10 different signal injections, all of which have been score-averaged over 20 independent classifier runs.



(a) Maximum of the Significance Improvement Characteristic.

(b) Significance Improvement Characteristic evaluated at a rejection of 10^3.

Figure D.117. Metrics for an anomaly detection classifier tasked with discriminating extrapolated background samples from data in the SR. The background is QCD dijet events. The signal is a 3 TeV Z' producing two dark QCD jets with $2m_{q_\mathrm{D}} = m_{\pi_\mathrm{D}} = m_{\rho_\mathrm{D}} = \Lambda_\mathrm{D} = 300$ GeV. Uncertainty bands represent the 68-percentile spread over 10 different signal injections, all of which have been score-averaged over 20 independent classifier runs.

contamination is 0.005% (551 events). A significant contamination could bias the extrapolated background distributions.

However, there is still sufficient evidence that the Reweight method could reliably enhance the signal significance from 1σ to the discovery threshold for all considered signal parameters. In contrast, the Generate and Morph methods appear to be more sensitive to the signal model, while still effective at picking up on moderate signal significances originating at $\sim 3\sigma$. It is worth exploring in which scenario the Reweight method becomes less effective than Generate and Morph in the future.

D.6 Modifying the simulation

In the main text of this work, the distributions of simulated background are similar to the data distributions in most variables. This feature could decrease the difficulty of the reweighting procedure in all three methods. In this section, we investigate the effects of using a “morphed”

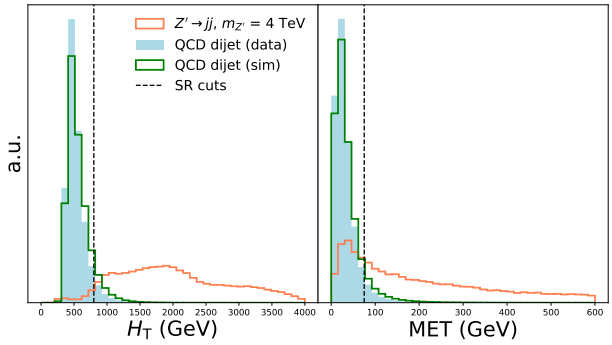
simulation set that might better represent the expected discrepancy between simulation and actual collider data.

We introduce a hand-tuned modification to the simulation, both to the features and context, defined in Eq. D.2:

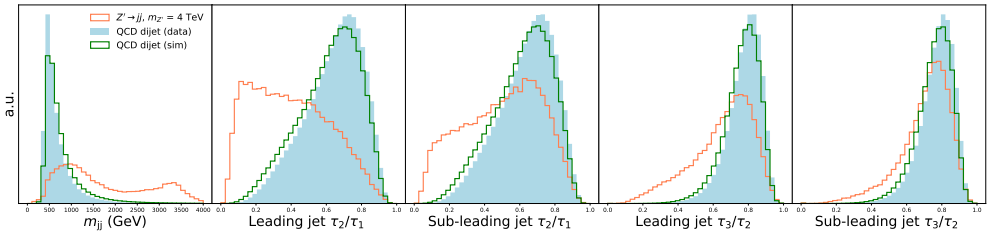
$$\begin{aligned} H_T: &\quad x \to x \\ \text{MET}: &\quad x \to x\left(1+\frac{x}{500}\right) \\ m_{jj}: &\quad x \to x\left(1+\frac{x}{6000}\right) \\ \tau_i: &\quad x \to x^{1.1}. \end{aligned} \tag{D.2}$$

Through this modification, the distributions of simulation are slightly morphed to be different from data, as shown in Fig. D.118.

In Fig. D.119, we compare the results for an anomaly detection classifier working on a range of signal detections for the modified simulation, using the same 4 TeV signal as was done in the main text. We make the comparison on the basis of the SIC at a rejection of 10^3; in Fig. D.119a, we show the results for the unmodified simulation, and in Fig. D.119b, we show the results from modified simulation. The performances are compatible, both quantitatively and qualitatively. In particular, the signal injections at which each method reaches the detection threshold are almost equal. The main difference is that the uncertainty bands are wider for the morphed simulation. This is logical given that when the simulation is more different from the "truth" background, the networks must learn a less trivial function to map from signal to background.

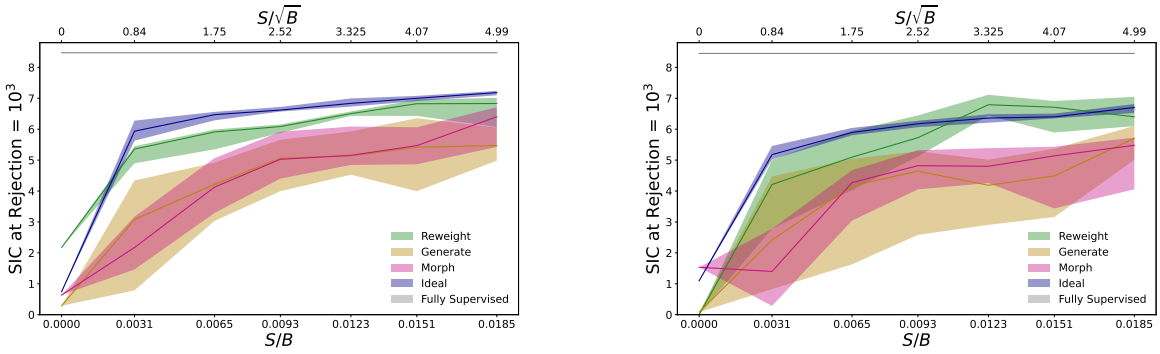


(a) Two-dimensional context variables



(b) Five-dimensional feature variables

Figure D.118. Histograms of the observables used in the Hidden Valley non-resonant anomaly detection task, for background ("data" and "simulation") and signal events). The QCD dijet simulation has been modified to appear different from QCD dijet data.



(a) Metric for unmodified simulation (identical to Fig. D.114b).

(b) Metric for modified simulation

Figure D.119. SIC at a rejection of 10^3 for an anomaly detection classifier tasked with discriminating extrapolated background samples from data in the SR. Uncertainty bands represent the 68-percentile spread over 10 different signal injections, all of which have been score-averaged over 20 independent classifier runs.

[1] G. Bertone and D. Hooper. 'History of dark matter'. In: *Rev. Mod. Phys.* 90.4 (2018). DOI: 10.1103/revmodphys.90.045002. URL: https://doi.org/10.1103%2Frevmodphys.90.045002.

[2] A. Borriello and P. Salucci. 'The dark matter distribution in disc galaxies'. In: *Mon. Notices Royal Astron. Soc.* 323.2 (2001), pp. 285–292. DOI: 10.1046/j.1365-8711.2001.04077.x. URL: https://doi.org/10.1046%2Fj.1365-8711.2001.04077.x.

[3] E. W. Kolb and M. S. Turner. *The Early Universe*. Vol. 69. 1990. ISBN: 978-0-201-62674-2. DOI: 10.1201/9780429492860.

[4] Matt Graham, Christopher Hearty and Mike Williams. 'Searches for Dark Photons at Accelerators'. In: *Annual Review of Nuclear and Particle Science* 71.1 (2021), pp. 37–58. DOI: 10.1146/annurev-nucl-110320-051823. URL: https://doi.org/10.1146%2Fannurev-nucl-110320-051823.

[5] David Curtin et al. 'Illuminating dark photons with high-energy colliders'. In: *J. High Energy Phys.* 2015.2 (2015). DOI: 10.1007/jhep02(2015)157. URL: https://doi.org/10.1007%2Fjhep02%282015%29157.

[6] P. Ilten et al. 'Serendipity in dark photon searches'. In: *J. High Energy Phys.* 2018.6 (2018). DOI: 10.1007/jhep06(2018)004. URL: https://doi.org/10.1007%2Fjhep06%282018%29004.

[7] Andrea Caputo and Rouven Essig. 'The Dark Photon: a 2026 Perspective'. In: Mar. 2026. arXiv: 2603.08430 [hep-ph].

[8] Georges Aad et al. 'Search for light neutral particles decaying promptly into collimated pairs of electrons or muons in pp collisions at $\sqrt{s}$ = 13 TeV with the ATLAS detector'. In: *Eur. Phys. J. C* 85.3 (2025), p. 335. DOI: 10.1140/epjc/s10052-025-13916-5. arXiv: 2407.09168 [hep-ex].

[9] Georges Aad et al. 'Search for light long-lived neutral particles that decay to collimated pairs of leptons or light hadrons in pp collisions at $\sqrt{s}$ = 13 TeV with the ATLAS detector'. In: *JHEP* 06 (2023), p. 153. DOI: 10.1007/JHEP06(2023)153. arXiv: 2206.12181 [hep-ex].

[10] CMS Collaboration. 'Search for long-lived particles decaying into muon pairs in proton-proton collisions at $\sqrt{s}$= 13 TeV collected with a dedicated high-rate data stream'. In: *J. High Energy Phys.* 2022.4 (2022), p. 62. ISSN: 1029-8479. DOI: 10.1007/JHEP04(2022)062. URL: https://doi.org/10.1007/JHEP04(2022)062.

[11] Carl D. Anderson. *Cloud Chamber Photo of Tracks of a Positron*. Caltech Images Collection. Item 1.2-7, Caltech Archives and Special Collections, https://collections.archives.caltech.edu/repositories/2/archival_objects/105304. Aug. 1932.

[12] Ernest O. Lawrence. 'Method and Apparatus for the Acceleration of Ions'. US1948384A. U.S. Patent No. 1948384, https://patents.google.com/patent/US1948384A/en. Feb. 1934.

[13] V.D. Barger and R.J.N. Phillips. *Collider Physics*. Frontiers in Physics. Avalon Publishing, 1996. ISBN: 9780201149456. URL: https://books.google.com/books?id=UFcutAEACAAJ.

[14] J. J. Aubert et al. 'Experimental Observation of a Heavy Particle *J*'. In: *Phys. Rev. Lett.* 33 (23 Dec. 1974), pp. 1404–1406. DOI: 10.1103/PhysRevLett.33.1404. URL: https://link.aps.org/doi/10.1103/PhysRevLett.33.1404.

[15] J. E. Augustin et al. 'Discovery of a Narrow Resonance in e^+e^- Annihilation'. In: *Phys. Rev. Lett.* 33 (1974), pp. 1406–1408. DOI: 10.1103/PhysRevLett.33.1406.

[16] S. Navas et al. 'Review of particle physics'. In: *Phys. Rev. D* 110.3 (2024), p. 030001. DOI: 10.1103/PhysRevD.110.030001.

[17] Lyndon Evans and Philip Bryant. 'LHC Machine'. In: *JINST* 3 (2008), S08001. DOI: 10.1088/1748-0221/3/08/S08001.

[18] Ewa Lopienska. 'The CERN accelerator complex, layout in 2022'. In: (2022). General Photo. URL: https://cds.cern.ch/record/2800984.

[19] Morad Aaboud et al. 'Comparison between simulated and observed LHC beam backgrounds in the ATLAS experiment at $E_{\text{beam}} = 4$ TeV'. In: *JINST* 13.12 (2018), P12006. DOI: 10.1088/1748-0221/13/12/P12006. arXiv: 1810.04450 [hep-ex].

[20] 'LHC Design Report Vol.1: The LHC Main Ring'. In: (June 2004). Ed. by Oliver S. Bruning et al. DOI: 10.5170/CERN-2004-003-V-1.

[21] *Superconductivity and the LHC: the early days*. CERN Courier, vol. 51, pp. 21–27. Nov. 2011. URL: https://cerncourier.com/a/superconductivity-and-the-lhc-the-early-days-2/.

[22] ATLAS Collaboration. *Luminosity Public Results Run 3*. ATLAS Public TWiki. Accessed: April 2026. 2024. URL: https://twiki.cern.ch/twiki/bin/view/AtlasPublic/LuminosityPublicResultsRun3.

[23] I. Zurbano Fernandez et al. 'High-Luminosity Large Hadron Collider (HL-LHC): Technical design report'. In: 10/2020 (Dec. 2020). Ed. by I. Béjar Alonso et al. DOI: 10.23731/CYRM-2020-0010.

[24] E. Bravin. 'Transverse emittance'. In: *CAS - CERN Accelerator School: Beam Instrumentation*. May 2020. DOI: 10.48550/arXiv.2005.05770. arXiv: 2005.05770 [physics.acc-ph].

[25] The ATLAS Collaboration. 'The ATLAS Experiment at the CERN Large Hadron Collider'. In: *J. Instrum.* 3.08 (2008), S08003. DOI: 10.1088/1748-0221/3/08/S08003. URL: https://dx.doi.org/10.1088/1748-0221/3/08/S08003.

[26] Georges Aad et al. 'The ATLAS experiment at the CERN Large Hadron Collider: a description of the detector configuration for Run 3'. In: *JINST* 19.05 (2024), P05063. DOI: 10.1088/1748-0221/19/05/P05063. arXiv: 2305.16623 [physics.ins-det].

[27] ATLAS Collaboration. 'The ATLAS Experiment at the CERN Large Hadron Collider'. In: *JINST* 3 (2008), S08003. DOI: 10.1088/1748-0221/3/08/S08003.

[28] H.H.J. ten Kate. 'The ATLAS superconducting magnet system at the Large Hadron Collider'. In: *Physica C* 468 (2008), pp. 2137–2142. DOI: 10.1016/j.physc.2008.05.146.

[29] ATLAS Collaboration. *ATLAS Insertable B-Layer: Technical Design Report*. ATLAS-TDR-19; CERN-LHCC-2010-013. 2010. URL: https://cds.cern.ch/record/1291633. Addendum: ATLAS-TDR-19-ADD-1; CERN-LHCC-2012-009. 2012. URL: https://cds.cern.ch/record/1451888.

[30] *Study of the mechanical stability of the ATLAS Insertable B-Layer*. Tech. rep. Geneva: CERN, 2015. URL: https://cds.cern.ch/record/2022587.

[31] Claudia Marcelloni De Oliveira. 'IBL installation into the inner detector of the ATLAS Experiment side C'. General Photo. 2014. URL: https://cds.cern.ch/record/1702006.

[32] V. Zivkovic et al. 'The FE-I4 pixel readout system-on-chip resubmission for the insertable B-Layer project'. In: *JINST* 7 (2012), p. C02050. DOI: 10.1088/1748-0221/7/02/C02050.

[33] 'ATLAS inner detector: Technical design report. Vol. 1'. In: (Apr. 1997).

[34] ATLAS Collaboration. *ATLAS LAr Calorimeter Phase-II Upgrade: Technical Design Report*. ATLAS-TDR-027; CERN-LHCC-2017-018. 2017. URL: https://cds.cern.ch/record/2285582.

[35] ATLAS Collaboration. *ATLAS New Small Wheel: Technical Design Report*. ATLAS-TDR-020; CERN-LHCC-2013-006. 2013. URL: https://cds.cern.ch/record/1552862.

[36] 'Standard Model Summary Plots June 2024'. In: (2024).

[37] ATLAS Collaboration. 'Study of the material of the ATLAS inner detector for Run 2 of the LHC'. In: *JINST* 12 (2017), P12009. DOI: 10.1088/1748-0221/12/12/P12009. arXiv: 1707.02826 [hep-ex].

[38] Georges Aad et al. 'Track and Vertex Reconstruction with the ATLAS Inner Detector'. In: (May 2026). arXiv: 2605.07585 [physics.ins-det].

[39] R. Fruhwirth. 'Application of Kalman filtering to track and vertex fitting'. In: *Nucl. Instrum. Meth. A* 262 (1987), pp. 444–450. DOI: 10.1016/0168-9002(87)90887-4.

[40] ATLAS Collaboration. 'Performance of the reconstruction of large impact parameter tracks in the inner detector of ATLAS'. In: *Eur. Phys. J. C* 83 (2023), p. 1081. DOI: 10.1140/epjc/s10052-023-12024-6. arXiv: 2304.12867 [hep-ex].

[41] R Frühwirth and W Waltenberger. *Adaptive Multi-vertex fitting*. Tech. rep. Geneva: CERN, 2005. DOI: 10.5170/CERN-2005-002.280. URL: https://cds.cern.ch/record/803519.

[42] *Performance of vertex reconstruction algorithms for detection of new long-lived particle decays within the ATLAS inner detector*. Tech. rep. All figures including auxiliary figures are available at https://atlas.web.cern.ch/Atlas/GROUPS/PHYSICS/PUBNOTES/ATL-PHYS-PUB-2019-013. Geneva: CERN, 2019. URL: https://cds.cern.ch/record/2669425.

[43] ATLAS Collaboration. 'Muon reconstruction and identification efficiency in ATLAS using the full Run 2 pp collision data set at $\sqrt{s} = 13$ TeV'. In: *Eur. Phys. J. C* 81 (2021), p. 578. DOI: 10.1140/epjc/s10052-021-09233-2. arXiv: 2012.00578 [hep-ex].

[44] *ZeroPixelHit Muons Reconstruction*. https://gitlab.cern.ch/atlas/athena/-/blob/main/PhysicsAnalysis/DerivationFramework/DerivationFrameworkLLP/src/RecoverZeroPixelHitMuons.cxx. RecoverZeroPixelHitMuons.

[45] ATLAS Collaboration. 'Topological cell clustering in the ATLAS calorimeters and its performance in LHC Run 1'. In: *Eur. Phys. J. C* 77 (2017), p. 490. DOI: 10.1140/epjc/s10052-017-5004-5. arXiv: 1603.02934 [hep-ex].

[46] ATLAS Collaboration. *Improved electron reconstruction in ATLAS using the Gaussian Sum Filter-based model for bremsstrahlung*. ATLAS-CONF-2012-047. 2012. URL: https://cds.cern.ch/record/1449796.

[47] ATLAS Collaboration. 'Electron and photon performance measurements with the ATLAS detector using the 2015–2017 LHC proton–proton collision data'. In: *JINST* 14 (2019), P12006. DOI: 10.1088/1748-0221/14/12/P12006. arXiv: 1908.00005 [hep-ex].

[48] ATLAS Collaboration. 'Electron reconstruction and identification in the ATLAS experiment using the 2015 and 2016 LHC proton–proton collision data at $\sqrt{s} = 13\,\text{TeV}$'. In: *Eur. Phys. J. C* 79 (2019), p. 639. DOI: 10.1140/epjc/s10052-019-7140-6. arXiv: 1902.04655 [physics.ins-det].

[49] ATLAS Collaboration. 'Jet reconstruction and performance using particle flow with the ATLAS Detector'. In: *Eur. Phys. J. C* 77 (2017), p. 466. DOI: 10.1140/epjc/s10052-017-5031-2. arXiv: 1703.10485 [hep-ex].

[50] Matteo Cacciari, Gavin P. Salam and Gregory Soyez. 'The anti-k_t jet clustering algorithm'. In: *JHEP* 04 (2008), p. 063. DOI: 10.1088/1126-6708/2008/04/063. arXiv: 0802.1189 [hep-ph].

[51] ATLAS Collaboration. 'Reconstruction of hadronic decay products of tau leptons with the ATLAS experiment'. In: *Eur. Phys. J. C* 76 (2016), p. 295. DOI: 10.1140/epjc/s10052-016-4110-0. arXiv: 1512.05955 [hep-ex].

[52] ATLAS Collaboration. 'The performance of missing transverse momentum reconstruction and its significance with the ATLAS detector using $140\,\text{fb}^{-1}$ of $\sqrt{s} = 13\,\text{TeV}$ pp collisions'. In: (2024). arXiv: 2402.05858 [hep-ex].

[53] ATLAS Collaboration. 'ATLAS Event Display: W boson decays to a muon and a neutrino'. General Photo. 2023. URL: https://cds.cern.ch/record/2859307.

[54] ATLAS Collaboration. 'Search for light long-lived neutral particles that decay to collimated pairs of leptons or light hadrons in pp collisions at $\sqrt{s} = 13$ TeV with the ATLAS detector'. In: *J. High Energy Phys.* 6 (2023). DOI: 10.1007/jhep06(2023)153. URL: https://doi.org/10.1007%2Fjhep06%282023%29153.

[55] CMS Collaboration. 'Search for a Narrow Resonance Lighter than 200 GeV Decaying to a Pair of Muons in Proton-Proton Collisions at $\sqrt{s} = 13$ TeV'. In: *Phys. Rev. Lett.* 124 (13 2020), p. 131802. DOI: 10.1103/PhysRevLett.124.131802. URL: https://link.aps.org/doi/10.1103/PhysRevLett.124.131802.

[56] *Measurement of the $H \to \gamma\gamma$ fiducial cross section with 31.4 fb^{-1} of 13.6 TeV proton-proton collision data with the ATLAS detector*. Tech. rep. Geneva: CERN, 2022. URL: https://cds.cern.ch/record/2842926.

[57] ATLAS Collaboration. 'Luminosity determination in pp collisions at $\sqrt{s} = 13$ TeV using the ATLAS detector at the LHC'. In: *Eur. Phys. J. C* 83 (2023), p. 982. DOI: 10.1140/epjc/s10052-023-11747-w. arXiv: 2212.09379 [hep-ex].

[58] *Preliminary luminosity calibration of the ATLAS 13.6 TeV data recorded in 2024 and combination with the 2022 and 2023 measurements*. Tech. rep. Geneva: CERN, 2025. URL: https://cds.cern.ch/record/2948582.

[59] G. Avoni et al. 'The new LUCID-2 detector for luminosity measurement and monitoring in ATLAS'. In: *JINST* 13 (2018), P07017. DOI: 10.1088/1748-0221/13/07/P07017.

[60] Glen Cowan et al. 'Asymptotic formulae for likelihood-based tests of new physics'. In: *Eur. Phys. J. C* 71 (2011). [Erratum: Eur. Phys. J. C **73** (2013) 2501], p. 1554. DOI: 10.1140/epjc/s10052-011-1554-0. arXiv: 1007.1727 [physics.data-an].

[61] ATLAS Collaboration. 'Search for light long-lived neutral particles that decay to collimated pairs of leptons or light hadrons in pp collisions at $\sqrt{s} = 13$ TeV with the ATLAS detector'. In: *JHEP* 06 (2023), p. 153. DOI: 10.1007/JHEP06(2023)153. arXiv: 2206.12181 [hep-ex].

[62] Georges Aad et al. 'Search for light neutral particles decaying promptly into collimated pairs of electrons or muons in pp collisions at $\sqrt{s} = 13$ TeV with the ATLAS detector'. In: *Eur. Phys. J. C* 85.3 (2025), p. 335. DOI: 10.1140/epjc/s10052-025-13916-5. arXiv: 2407.09168 [hep-ex].

[63] Armen Tumasyan et al. 'Search for long-lived particles decaying into muon pairs in proton-proton collisions at $\sqrt{s}$ = 13 TeV collected with a dedicated high-rate data stream'. In: *JHEP* 04 (2022), p. 062. DOI: 10.1007/JHEP04(2022)062. arXiv: 2112.13769 [hep-ex].

[64] *Search for long-lived particles decaying into muons in proton-proton collisions at* $\sqrt{s}$ = 13.6 TeV *using the CMS scouting data sets*. Tech. rep. Geneva: CERN, 2025. URL: https://cds.cern.ch/record/2938078.

[65] Maurice Garcia-Sciveres, Flavio Loddo and Jorgen Christiansen. *RD53B Manual*. Tech. rep. Geneva: CERN, 2019. URL: https://cds.cern.ch/record/2665301.

[66] ATLAS Collaboration. *Expected tracking and related performance with the updated ATLAS Inner Tracker layout at the High-Luminosity LHC*. ATL-PHYS-PUB-2021-024. 2021. URL: https://cds.cern.ch/record/2776651.

[67] ATLAS Collaboration. *Sensitivity of the ATLAS experiment to long-lived particles with a displaced vertex and* E_T^{miss} *signature at the HL-LHC*. ATL-PHYS-PUB-2018-033. 2018. URL: https://cds.cern.ch/record/2647992.

[68] Kehang Bai, Radha Mastandrea and Benjamin Nachman. 'Non-resonant anomaly detection with background extrapolation'. In: *JHEP* 04 (2024), p. 059. DOI: 10.1007/JHEP04(2024)059. arXiv: 2311.12924 [hep-ph].

[69] Gregor Kasieczka et al. 'The LHC Olympics 2020: A Community Challenge for Anomaly Detection in High Energy Physics'. In: (Jan. 2021). arXiv: 2101.08320 [hep-ph].

[70] T. Aarrestad et al. 'The Dark Machines Anomaly Score Challenge: Benchmark Data and Model Independent Event Classification for the Large Hadron Collider'. In: (May 2021). arXiv: 2105.14027 [hep-ph].

[71] Georgia Karagiorgi et al. 'Machine Learning in the Search for New Fundamental Physics'. In: (Dec. 2021). arXiv: 2112.03769 [hep-ph].

[72] Matthew Feickert and Benjamin Nachman. 'A Living Review of Machine Learning for Particle Physics'. In: (Feb. 2021). arXiv: 2102.02770 [hep-ph].

[73] Jack H. Collins, Kiel Howe and Benjamin Nachman. 'Anomaly Detection for Resonant New Physics with Machine Learning'. In: *Phys. Rev. Lett.* 121.24 (2018), p. 241803. DOI: 10.1103/PhysRevLett.121.241803. arXiv: 1805.02664 [hep-ph].

[74] Raffaele Tito D'Agnolo and Andrea Wulzer. 'Learning New Physics from a Machine'. In: *Phys. Rev.* D99.1 (2019), p. 015014. DOI: 10.1103/PhysRevD.99.015014. arXiv: 1806.02350 [hep-ph].

[75] Jack H. Collins, Kiel Howe and Benjamin Nachman. 'Extending the search for new resonances with machine learning'. In: *Phys. Rev.* D99.1 (2019), p. 014038. DOI: 10.1103/PhysRevD.99.014038. arXiv: 1902.02634 [hep-ph].

[76] Raffaele Tito D'Agnolo et al. 'Learning Multivariate New Physics'. In: (2019). DOI: 10.1140/epjc/s10052-021-08853-y. arXiv: 1912.12155 [hep-ph].

[77] Marco Farina, Yuichiro Nakai and David Shih. 'Searching for New Physics with Deep Autoencoders'. In: (2018). DOI: 10.1103/PhysRevD.101.075021. arXiv: 1808.08992 [hep-ph].

[78] Theo Heimel et al. 'QCD or What?' In: *SciPost Phys.* 6.3 (2019), p. 030. DOI: 10.21468/SciPostPhys.6.3.030. arXiv: 1808.08979 [hep-ph].

[79] Tuhin S. Roy and Aravind H. Vijay. 'A robust anomaly finder based on autoencoder'. In: (2019). arXiv: 1903.02032 [hep-ph].

[80] Olmo Cerri et al. 'Variational Autoencoders for New Physics Mining at the Large Hadron Collider'. In: *JHEP* 05 (2019), p. 036. DOI: 10.1007/JHEP05(2019)036. arXiv: 1811.10276 [hep-ex].

[81] Andrew Blance, Michael Spannowsky and Philip Waite. 'Adversarially-trained autoencoders for robust unsupervised new physics searches'. In: *JHEP* 10 (2019), p. 047. DOI: 10.1007/JHEP10(2019)047. arXiv: 1905.10384 [hep-ph].

[82] Jan Hajer et al. 'Novelty Detection Meets Collider Physics'. In: (2018). DOI: 10.1103/PhysRevD.101.076015. arXiv: 1807.10261 [hep-ph].

[83] Andrea De Simone and Thomas Jacques. 'Guiding New Physics Searches with Unsupervised Learning'. In: *Eur. Phys. J.* C79.4 (2019), p. 289. DOI: 10.1140/epjc/s10052-019-6787-3. arXiv: 1807.06038 [hep-ph].

[84] Anna Mullin et al. 'Does SUSY have friends? A new approach for LHC event analysis'. In: (2019). DOI: 10.1007/JHEP02(2021)160. arXiv: 1912.10625 [hep-ph].

[85] Giovanna Menardi Alessandro Casa. 'Nonparametric semisupervised classification for signal detection in high energy physics'. In: (2019). arXiv: 1809.02977 [hep-ex].

[86] Barry M. Dillon, Darius A. Faroughy and Jernej F. Kamenik. 'Uncovering latent jet substructure'. In: *Phys. Rev.* D100.5 (2019), p. 056002. DOI: 10.1103/PhysRevD.100.056002. arXiv: 1904.04200 [hep-ph].

[87] Anders Andreassen, Benjamin Nachman and David Shih. 'Simulation Assisted Likelihood-free Anomaly Detection'. In: *Phys. Rev. D* 101.9 (2020), p. 095004. DOI: 10.1103/PhysRevD.101.095004. arXiv: 2001.05001 [hep-ph].

[88] Benjamin Nachman and David Shih. 'Anomaly Detection with Density Estimation'. In: *Phys. Rev. D* 101 (2020), p. 075042. DOI: 10.1103/PhysRevD.101.075042. arXiv: 2001.04990 [hep-ph].

[89] J. A. Aguilar-Saavedra, Jack H. Collins and Rashmish K. Mishra. 'A generic anti-QCD jet tagger'. In: *JHEP* 11 (2017), p. 163. DOI: 10.1007/JHEP11(2017)163. arXiv: 1709.01087 [hep-ph].

[90] M. Romão Crispim et al. 'Transferability of Deep Learning Models in Searches for New Physics at Colliders'. In: *Phys. Rev. D* 101.3 (2020), p. 035042. DOI: 10.1103/PhysRevD.101.035042. arXiv: 1912.04220 [hep-ph].

[91] M. Crispim Romao et al. 'Use of a Generalized Energy Mover's Distance in the Search for Rare Phenomena at Colliders'. In: (Apr. 2020). DOI: 10.1140/epjc/s10052-021-08891-6. arXiv: 2004.09360 [hep-ph].

[92] Oliver Knapp et al. 'Adversarially Learned Anomaly Detection on CMS Open Data: re-discovering the top quark'. In: (2020). DOI: 10.1140/epjp/s13360-021-01109-4. arXiv: 2005.01598 [hep-ex].

[93] ATLAS Collaboration. 'Dijet resonance search with weak supervision using 13 TeV pp collisions in the ATLAS detector'. In: (2020). DOI: 10.1103/PhysRevLett.125.131801. arXiv: 2005.02983 [hep-ex].

[94] Barry M. Dillon et al. 'Learning the latent structure of collider events'. In: (May 2020). DOI: 10.1007/JHEP10(2020)206. arXiv: 2005.12319 [hep-ph].

[95] M. Crispim Romao, N.F. Castro and R. Pedro. 'Finding New Physics without learning about it: Anomaly Detection as a tool for Searches at Colliders'. In: (June 2020). DOI: 10.1140/epjc/s10052-020-08807-w. arXiv: 2006.05432 [hep-ph].

[96] Oz Amram and Cristina Mantilla Suarez. 'Tag N' Train: A Technique to Train Improved Classifiers on Unlabeled Data'. In: (Feb. 2020). DOI: 10.1007/JHEP01(2021)153. arXiv: 2002.12376 [hep-ph].

[97] Taoli Cheng et al. 'Variational Autoencoders for Anomalous Jet Tagging'. In: (July 2020). arXiv: 2007.01850 [hep-ph].

[98] Charanjit K. Khosa and Veronica Sanz. 'Anomaly Awareness'. In: (July 2020). arXiv: 2007.14462 [cs.LG].

[99] Punnathat Thaprasop et al. 'Unsupervised Outlier Detection in Heavy-Ion Collisions'. In: (July 2020). arXiv: 2007.15830 [hep-ex].

[100] Stephon Alexander et al. 'Decoding Dark Matter Substructure without Supervision'. In: (Aug. 2020). arXiv: 2008.12731 [astro-ph.CO].

[101] J. A. Aguilar-Saavedra, F. R. Joaquim and J. F. Seabra. 'Mass Unspecific Supervised Tagging (MUST) for boosted jets'. In: (2020). DOI: 10.1007/JHEP03(2021)012. arXiv: 2008.12792 [hep-ph].

[102] Kees Benkendorfer, Luc Le Pottier and Benjamin Nachman. 'Simulation-Assisted Decorrelation for Resonant Anomaly Detection'. In: (Sept. 2020). arXiv: 2009.02205 [hep-ph].

[103] Adrian Alan Pol and Victor Berger and Gianluca Cerminara and Cecile Germain and Maurizio Pierini. 'Anomaly Detection With Conditional Variational Autoencoders'. In: (2020). arXiv: 2010.05531 [cs.LG].

[104] Vinicius Mikuni and Florencia Canelli. 'Unsupervised clustering for collider physics'. In: (Sept. 2020). arXiv: 2010.07106 [physics.data-an].

[105] Melissa van Beekveld et al. 'Combining outlier analysis algorithms to identify new physics at the LHC'. In: (Oct. 2020). arXiv: 2010.07940 [hep-ph].

[106] Sang Eon Park et al. 'Quasi Anomalous Knowledge: Searching for new physics with embedded knowledge'. In: (Nov. 2020). arXiv: 2011.03550 [hep-ph].

[107] Darius A. Faroughy. 'Uncovering hidden patterns in collider events with Bayesian probabilistic models'. In: (Dec. 2020). arXiv: 2012.08579 [hep-ph].

[108] George Stein, Uros Seljak and Biwei Dai. 'Unsupervised in-distribution anomaly detection of new physics through conditional density estimation'. In: (Dec. 2020). arXiv: 2012.11638 [cs.LG].

[109] Purvasha Chakravarti et al. 'Model-Independent Detection of New Physics Signals Using Interpretable Semi-Supervised Classifier Tests'. In: (Feb. 2021). arXiv: 2102.07679 [stat.AP].

[110] Joshua Batson et al. 'Topological Obstructions to Autoencoding'. In: (Feb. 2021). arXiv: 2102.08380 [hep-ph].

[111] Andrew Blance and Michael Spannowsky. 'Unsupervised Event Classification with Graphs on Classical and Photonic Quantum Computers'. In: (Mar. 2021). arXiv: 2103.03897 [hep-ph].

[112] Blaž Bortolato et al. 'Bump Hunting in Latent Space'. In: (Mar. 2021). arXiv: 2103.06595 [hep-ph].

[113] Jack H. Collins et al. 'Comparing Weak- and Unsupervised Methods for Resonant Anomaly Detection'. In: (Apr. 2021). arXiv: 2104.02092 [hep-ph].

[114] Barry M. Dillon et al. ‘Better Latent Spaces for Better Autoencoders’. In: (Apr. 2021). arXiv: 2104.08291 [hep-ph].

[115] Thorben Finke et al. ‘Autoencoders for unsupervised anomaly detection in high energy physics’. In: (Apr. 2021). arXiv: 2104.09051 [physics.data-an].

[116] David Shih et al. ‘Via Machinae: Searching for Stellar Streams using Unsupervised Machine Learning’. In: (Apr. 2021). arXiv: 2104.12789 [astro-ph.GA].

[117] Oliver Atkinson et al. ‘Anomaly detection with Convolutional Graph Neural Networks’. In: (May 2021). arXiv: 2105.07988 [hep-ph].

[118] Alan Kahn et al. ‘Anomalous Jet Identification via Sequence Modeling’. In: (May 2021). arXiv: 2105.09274 [hep-ph].

[119] Tommaso Dorigo et al. ‘RanBox: Anomaly Detection in the Copula Space’. In: (June 2021). arXiv: 2106.05747 [physics.data-an].

[120] Sascha Caron, Luc Hendriks and Rob Verheyen. ‘Rare and Different: Anomaly Scores from a combination of likelihood and out-of-distribution models to detect new physics at the LHC’. In: (June 2021). arXiv: 2106.10164 [hep-ph].

[121] Ekaterina Govorkova et al. ‘LHC physics dataset for unsupervised New Physics detection at 40 MHz’. In: (July 2021). arXiv: 2107.02157 [physics.data-an].

[122] Gregor Kasieczka, Benjamin Nachman and David Shih. ‘New Methods and Datasets for Group Anomaly Detection From Fundamental Physics’. In: July 2021. arXiv: 2107.02821 [stat.ML].

[123] Sergey Volkovich, Federico De Vito Halevy and Shikma Bressler. ‘The Data-Directed Paradigm for BSM searches’. In: (July 2021). arXiv: 2107.11573 [hep-ex].

[124] Ekaterina Govorkova et al. ‘Autoencoders on FPGAs for real-time, unsupervised new physics detection at 40 MHz at the Large Hadron Collider’. In: (Aug. 2021). arXiv: 2108.03986 [physics.ins-det].

[125] Anna Hallin et al. ‘Classifying Anomalies THrough Outer Density Estimation (CATHODE)’. In: (Sept. 2021). arXiv: 2109.00546 [hep-ph].

[126] Bryan Ostdiek. ‘Deep Set Auto Encoders for Anomaly Detection in Particle Physics’. In: (Sept. 2021). arXiv: 2109.01695 [hep-ph].

[127] Katherine Fraser et al. ‘Challenges for Unsupervised Anomaly Detection in Particle Physics’. In: (Oct. 2021). arXiv: 2110.06948 [cs.LG].

[128] Pratik Jawahar et al. ‘Improving Variational Autoencoders for New Physics Detection at the LHC with Normalizing Flows’. In: (Oct. 2021). arXiv: 2110.08508 [hep-ph].

[129] Juan Herrero-Garcia, Riley Patrick and Andre Scaffidi. ‘Signal-agnostic dark matter searches in direct detection data with machine learning’. In: (Oct. 2021). arXiv: 2110.12248 [hep-ph].

[130] J. A. Aguilar-Saavedra. ‘Anomaly detection from mass unspecific jet tagging’. In: (Nov. 2021). arXiv: 2111.02647 [hep-ph].

[131] Rupert Tombs and Christopher G. Lester. ‘A method to challenge symmetries in data with self-supervised learning’. In: (Nov. 2021). arXiv: 2111.05442 [hep-ph].

[132] Christopher G. Lester and Rupert Tombs. ‘Stressed GANs snag desserts, a.k.a Spotting Symmetry Violation with Symmetric Functions’. In: (Oct. 2021). arXiv: 2111.00616 [hep-ph].

[133] Vinicius Mikuni, Benjamin Nachman and David Shih. ‘Online-compatible Unsupervised Non-resonant Anomaly Detection’. In: (Nov. 2021). arXiv: 2111.06417 [cs.LG].

[134] S. V. Chekanov and W. Hopkins. ‘Event-based anomaly detection for new physics searches at the LHC using machine learning’. In: (Nov. 2021). arXiv: 2111.12119 [hep-ph].

[135] Raffaele Tito d’Agnolo et al. ‘Learning New Physics from an Imperfect Machine’. In: (Nov. 2021). arXiv: 2111.13633 [hep-ph].

[136] Florencia Canelli et al. ‘Autoencoders for Semivisible Jet Detection’. In: (Dec. 2021). arXiv: 2112.02864 [hep-ph].

[137] Vishal S. Ngairangbam, Michael Spannowsky and Michihisa Takeuchi. ‘Anomaly detection in high-energy physics using a quantum autoencoder’. In: (Dec. 2021). arXiv: 2112.04958 [hep-ph].

[138] Layne Bradshaw, Spencer Chang and Bryan Ostdiek. ‘Creating Simple, Interpretable Anomaly Detectors for New Physics in Jet Substructure’. In: (Mar. 2022). arXiv: 2203.01343 [hep-ph].

[139] J. A. Aguilar-Saavedra. ‘Taming modeling uncertainties with Mass Unspecific Supervised Tagging’. In: (Jan. 2022). arXiv: 2201.11143 [hep-ph].

[140] Thorsten Buss et al. ‘What’s Anomalous in LHC Jets?’ In: (Feb. 2022). arXiv: 2202.00686 [hep-ph].

[141] Sulaiman Alvi, Christian Bauer and Benjamin Nachman. ‘Quantum Anomaly Detection for Collider Physics’. In: (June 2022). arXiv: 2206.08391 [hep-ph].

[142] Barry M. Dillon, Radha Mastandrea and Benjamin Nachman. ‘Self-supervised Anomaly Detection for New Physics’. In: (May 2022). arXiv: 2205.10380 [hep-ph].

[143] Mattias Birman et al. ‘Data-directed search for new physics based on symmetries of the SM’. In: *Eur. Phys. J. C* 82.6 (2022), p. 508. DOI: 10.1140/epjc/s10052-022-10454-2. arXiv: 2203.07529 [hep-ph].

[144] John Andrew Raine et al. ‘CURTAINs for your Sliding Window: Constructing Unobserved Regions by Transforming Adjacent Intervals’. In: (Mar. 2022). arXiv: 2203.09470 [hep-ph].

[145] Marco Letizia et al. ‘Learning new physics efficiently with nonparametric methods’. In: (Apr. 2022). arXiv: 2204.02317 [hep-ph].

[146] C. Fanelli, J. Giroux and Z. Papandreou. '"Flux+Mutability": A Conditional Generative Approach to One-Class Classification and Anomaly Detection'. In: (Apr. 2022). arXiv: 2204.08609 [cs.LG].

[147] Thorben Finke et al. 'Boosting mono-jet searches with model-agnostic machine learning'. In: *JHEP* 08 (2022), p. 015. DOI: 10.1007/JHEP08(2022)015. arXiv: 2204.11889 [hep-ph].

[148] Rob Verheyen. 'Event Generation and Density Estimation with Surjective Normalizing Flows'. In: (May 2022). arXiv: 2205.01697 [hep-ph].

[149] Barry M. Dillon et al. 'A Normalized Autoencoder for LHC Triggers'. In: (June 2022). arXiv: 2206.14225 [hep-ph].

[150] Sascha Caron, Roberto Ruiz de Austri and Zhongyi Zhang. 'Mixture-of-theories Training: Can We Find New Physics and Anomalies Better by Mixing Physical Theories?' In: (July 2022). arXiv: 2207.07631 [hep-ph].

[151] Sang Eon Park, Philip Harris and Bryan Ostdiek. 'Neural Embedding: Learning the Embedding of the Manifold of Physics Data'. In: (Aug. 2022). arXiv: 2208.05484 [hep-ph].

[152] Jernej F. Kamenik and Manuel Szewc. 'Null Hypothesis Test for Anomaly Detection'. In: (Oct. 2022). arXiv: 2210.02226 [hep-ph].

[153] Anna Hallin et al. 'Resonant anomaly detection without background sculpting'. In: (Oct. 2022). arXiv: 2210.14924 [hep-ph].

[154] Gregor Kasieczka et al. 'Anomaly Detection under Coordinate Transformations'. In: (Sept. 2022). arXiv: 2209.06225 [hep-ph].

[155] Jack Y. Araz and Michael Spannowsky. 'Quantum-probabilistic Hamiltonian learning for generative modelling & anomaly detection'. In: (Nov. 2022). arXiv: 2211.03803 [quant-ph].

[156] Radha Mastandrea and Benjamin Nachman. 'Efficiently Moving Instead of Reweighting Collider Events with Machine Learning'. In: *36th Conference on Neural Information Processing Systems*. Dec. 2022. arXiv: 2212.06155 [hep-ph].

[157] Julian Schuhmacher et al. 'Unravelling physics beyond the standard model with classical and quantum anomaly detection'. In: (Jan. 2023). arXiv: 2301.10787 [hep-ex].

[158] Stephen Roche et al. 'Nanosecond anomaly detection with decision trees for high energy physics and real-time application to exotic Higgs decays'. In: (Apr. 2023). arXiv: 2304.03836 [hep-ex].

[159] Tobias Golling et al. 'The Mass-ive Issue: Anomaly Detection in Jet Physics'. In: *34th Conference on Neural Information Processing Systems*. Mar. 2023. arXiv: 2303.14134 [hep-ph].

[160] Debajyoti Sengupta et al. 'CURTAINs Flows For Flows: Constructing Unobserved Regions with Maximum Likelihood Estimation'. In: (May 2023). arXiv: 2305.04646 [hep-ph].

[161] Vinicius Mikuni and Benjamin Nachman. 'High-dimensional and Permutation Invariant Anomaly Detection'. In: (June 2023). arXiv: 2306.03933 [hep-ph].

[162] Tobias Golling et al. 'The Interplay of Machine Learning–based Resonant Anomaly Detection Methods'. In: (July 2023). arXiv: 2307.11157 [hep-ph].

[163] Louis Vaslin, Vincent Barra and Julien Donini. 'GAN-AE : An anomaly detection algorithm for New Physics search in LHC data'. In: (May 2023). arXiv: 2305.15179 [hep-ex].

[164] ATLAS Collaboration. 'Anomaly detection search for new resonances decaying into a Higgs boson and a generic new particle X in hadronic final states using $\sqrt{s} = 13$ TeV pp collisions with the ATLAS detector'. In: (June 2023). arXiv: 2306.03637 [hep-ex].

[165] Sergei V. Chekanov and Rui Zhang. 'Boosting sensitivity to new physics with unsupervised anomaly detection in dijet resonance search'. In: (Aug. 2023). arXiv: 2308.02671 [hep-ex].

[166] D. Abadjiev et al. 'Autoencoder-based Anomaly Detection System for Online Data Quality Monitoring of the CMS Electromagnetic Calorimeter'. In: (Sept. 2023). arXiv: 2309.10157 [physics.ins-det].

[167] Gerrit Bickendorf et al. 'Combining Resonant and Tail-based Anomaly Detection'. In: (Sept. 2023). arXiv: 2309.12918 [hep-ph].

[168] Thorben Finke et al. 'Back To The Roots: Tree-Based Algorithms for Weakly Supervised Anomaly Detection'. In: (Sept. 2023). arXiv: 2309.13111 [hep-ph].

[169] Erik Buhmann et al. 'Full Phase Space Resonant Anomaly Detection'. In: (Oct. 2023). arXiv: 2310.06897 [hep-ph].

[170] Marat Freytsis, Maxim Perelstein and Yik Chuen San. 'Anomaly Detection in Presence of Irrelevant Features'. In: (Oct. 2023). arXiv: 2310.13057 [hep-ph].

[171] Tobias Golling et al. 'Flow-enhanced transportation for anomaly detection'. In: *Phys. Rev. D* 107.9 (2023), p. 096025. DOI: 10.1103/PhysRevD.107.096025. arXiv: 2212.11285 [hep-ph].

[172] Ouail Kitouni, Niklas Nolte and Mike Williams. 'Robust and Provably Monotonic Networks'. In: *35th Conference on Neural Information Processing Systems*. Nov. 2021. arXiv: 2112.00038 [cs.LG].

[173] Joshua Lin, Wahid Bhimji and Benjamin Nachman. 'Machine Learning Templates for QCD Factorization in the Search for Physics Beyond the Standard Model'. In: *JHEP* 05 (2019), p. 181. DOI: 10.1007/JHEP05(2019)181. arXiv: 1903.02556 [hep-ph].

[174] Gregor Kasieczka et al. 'ABCDisCo: Automating the ABCD Method with Machine Learning'. In: (July 2020). DOI: 10.1103/PhysRevD.103.035021. arXiv: 2007.14400 [hep-ph].

[175] Eric M. Metodiev, Benjamin Nachman and Jesse Thaler. 'Classification without labels: Learning from mixed samples in high energy physics'. In: *JHEP* 10 (2017), p. 174. DOI: 10.1007/JHEP10(2017)174. arXiv: 1708.02949 [hep-ph].

[176] Kyle Cranmer, Juan Pavez and Gilles Louppe. 'Approximating Likelihood Ratios with Calibrated Discriminative Classifiers'. In: (June 2015). arXiv: 1506.02169 [stat.AP].

[177] Pierre Baldi et al. 'Parameterized neural networks for high-energy physics'. In: *Eur. Phys. J.* C76.5 (2016), p. 235. DOI: 10.1140/epjc/s10052-016-4099-4. arXiv: 1601.07913 [hep-ex].

[178] Danilo Rezende and Shakir Mohamed. 'Variational Inference with Normalizing Flows'. In: *Proceedings of the 32nd International Conference on Machine Learning* 37 (July 2015). Ed. by Francis Bach and David Blei, pp. 1530–1538. URL: http://proceedings.mlr.press/v37/rezende15.html.

[179] Tobias Golling et al. 'Flows for Flows: Morphing one Dataset into another with Maximum Likelihood Estimation'. In: (Sept. 2023). arXiv: 2309.06472 [hep-ph].

[180] Samuel Bright-Thonney et al. 'Chained Quantile Morphing with Normalizing Flows'. In: (Sept. 2023). arXiv: 2309.15912 [hep-ph].

[181] Adam Paszke et al. 'PyTorch: An Imperative Style, High-Performance Deep Learning Library'. In: *Advances in Neural Information Processing Systems 32*. Curran Associates, Inc., 2019, pp. 8024–8035. URL: http://papers.neurips.cc/paper/9015-pytorch-an-imperative-style-high-performance-deep-learning-library.pdf.

[182] Diederik P. Kingma and Jimmy Ba. *Adam: A Method for Stochastic Optimization*. 2017. arXiv: 1412.6980 [cs.LG].

[183] Conor Durkan et al. *nflows: normalizing flows in PyTorch*. Version v0.14. Nov. 2020. DOI: 10.5281/zenodo.4296287. URL: https://doi.org/10.5281/zenodo.4296287.

[184] Matthew J. Strassler and Kathryn M. Zurek. 'Echoes of a hidden valley at hadron colliders'. In: *Phys. Lett. B* 651 (2007), pp. 374–379. DOI: 10.1016/j.physletb.2007.06.055. arXiv: hep-ph/0604261.

[185] Lisa Carloni and Torbjorn Sjostrand. 'Visible Effects of Invisible Hidden Valley Radiation'. In: *JHEP* 09 (2010), p. 105. DOI: 10.1007/JHEP09(2010)105. arXiv: 1006.2911 [hep-ph].

[186] Lisa Carloni, Johan Rathsman and Torbjorn Sjostrand. 'Discerning Secluded Sector gauge structures'. In: *JHEP* 04 (2011), p. 091. DOI: 10.1007/JHEP04(2011)091. arXiv: 1102.3795 [hep-ph].

[187] Simon Knapen, Jessie Shelton and Dong Xu. 'Perturbative benchmark models for a dark shower search program'. In: *Phys. Rev. D* 103.11 (2021), p. 115013. DOI: 10.1103/PhysRevD.103.115013. arXiv: 2103.01238 [hep-ph].

[188] Timothy Cohen, Mariangela Lisanti and Hou Keong Lou. 'Semivisible Jets: Dark Matter Undercover at the LHC'. In: *Physical Review Letters* 115.17 (Oct. 2015). ISSN: 1079-7114. DOI: 10.1103/physrevlett.115.171804. arXiv: 1503.00009. URL: http://dx.doi.org/10.1103/PhysRevLett.115.171804.

[189] Timothy Cohen et al. 'LHC searches for dark sector showers'. In: *Journal of High Energy Physics* 2017.11 (Nov. 2017). ISSN: 1029-8479. DOI: 10.1007/jhep11(2017)196. arXiv: 1707.05326. URL: http://dx.doi.org/10.1007/JHEP11(2017)196.

[190] Hugues Beauchesne et al. 'Collider phenomenology of Hidden Valley mediators of spin 0 or 1/2 with semivisible jets'. In: *Journal of High Energy Physics* 2018.8 (Aug. 2018). ISSN: 1029-8479. DOI: 10.1007/jhep08(2018)030. arXiv: 1712.07160. URL: http://dx.doi.org/10.1007/JHEP08(2018)030.

[191] Elias Bernreuther et al. 'Casting a graph net to catch dark showers'. In: *SciPost Physics* 10.2 (Feb. 2021). ISSN: 2542-4653. DOI: 10.21468/scipostphys.10.2.046. arXiv: 2006.08639. URL: http://dx.doi.org/10.21468/SciPostPhys.10.2.046.

[192] 'Search for resonant production of strongly coupled dark matter in proton-proton collisions at 13 TeV'. In: *Journal of High Energy Physics* 2022.6 (June 2022). ISSN: 1029-8479. DOI: 10.1007/jhep06(2022)156. arXiv: 2112.11125. URL: http://dx.doi.org/10.1007/JHEP06(2022)156.

[193] 'Search for non-resonant production of semi-visible jets using Run 2 data in ATLAS'. In: (2023). arXiv: 2305.18037 [hep-ex].

[194] Timothy Cohen, Joel Doss and Marat Freytsis. 'Jet substructure from dark sector showers'. In: *Journal of High Energy Physics* 2020.9 (Sept. 2020). ISSN: 1029-8479. DOI: 10.1007/jhep09(2020)118. arXiv: 2004.00631. URL: http://dx.doi.org/10.1007/JHEP09(2020)118.

[195] Timothy Cohen, Jennifer Roloff and Christiane Scherb. 'Dark sector showers in the Lund jet plane'. In: *Physical Review D* 108.3 (Aug. 2023). ISSN: 2470-0029. DOI: 10.1103/physrevd.108.l031501. arXiv: 2301.07732. URL: http://dx.doi.org/10.1103/PhysRevD.108.L031501.

[196] Christian Bierlich et al. *A comprehensive guide to the physics and usage of PYTHIA 8.3*. 2022. arXiv: 2203.11601 [hep-ph].

[197] J. Alwall et al. 'The automated computation of tree-level and next-to-leading order differential cross sections, and their matching to parton shower simulations'. In: *Journal of High Energy Physics* 2014.7 (July 2014). DOI: 10.1007/jhep07(2014)079. URL: https://doi.org/10.1007%2Fjhep07%282014%29079.

[198] J. de Favereau et al. 'DELPHES 3, A modular framework for fast simulation of a generic collider experiment'. In: *JHEP* 02 (2014), p. 057. DOI: 10.1007/JHEP02(2014)057. arXiv: 1307.6346 [hep-ex].

[199] Matteo Cacciari, Gavin P. Salam and Gregory Soyez. 'FastJet User Manual'. In: *Eur. Phys. J. C* 72 (2012), p. 1896. DOI: 10.1140/epjc/s10052-012-1896-2. arXiv: 1111.6097 [hep-ph].

[200] Jesse Thaler and Ken Van Tilburg. 'Maximizing Boosted Top Identification by Minimizing N-subjettiness'. In: *JHEP* 02 (2012), p. 093. DOI: 10.1007/JHEP02(2012)093. arXiv: 1108.2701 [hep-ph].

[201] Jesse Thaler and Ken Van Tilburg. 'Identifying Boosted Objects with N-subjettiness'. In: *JHEP* 03 (2011), p. 015. DOI: 10.1007/JHEP03(2011)015. arXiv: 1011.2268 [hep-ph].