Search for dark gauge bosons decaying into displaced lepton-jets in proton–proton collisions at $\sqrt{s} = 13$ TeV with the ATLAS detector

by

Miriam Diamond

A thesis submitted in conformity with the requirements
for the degree of Doctor of Philosophy
Graduate Department of Physics
University of Toronto

Abstract

Search for dark gauge bosons decaying into displaced lepton-jets in proton–proton collisions at $\sqrt{s} = 13$ TeV with the ATLAS detector

Miriam Diamond
Doctor of Philosophy
Graduate Department of Physics
University of Toronto
2017

The dark photon (A'), the gauge boson carrier of a hypothetical new force, has been proposed in a wide range of Beyond the Standard Model (BSM) theories, and could serve as our window to an entire dark sector. A massive A' could decay back to the Standard Model (SM) with a significant branching fraction, through kinetic mixing with the SM photon. If this A' can be produced from decays of a dark scalar that mixes with the SM Higgs boson, collider searches involving leptonic final states provide promising discovery prospects with rich phenomenology. This work presents the results of a search for dark photons in the mass range $0.2 \lesssim m_{A'} \lesssim 10$ GeV decaying into collimated jets of light leptons and mesons, so-called "lepton-jets". It employs 3.57 fb^{-1} of data from proton-proton collisions at a centre-of-mass energy of $\sqrt{s} = 13$ TeV, collected during 2015 with the ATLAS detector at the LHC. No deviations from SM expectations are observed. Limits on benchmark models predicting Higgs boson decays to A's are derived as a function of the A' lifetime; limits are also established in the parameter space of $m_{A'}$ vs. kinetic mixing parameter ϵ . These extend the limits obtained in a similar search previously performed during Run 1 of the LHC, to include dark photon masses $2 \lesssim m_{A'} \lesssim 10$ GeV and to cover higher ϵ values for $0.2 \lesssim m_{A'} \lesssim 2$ GeV, and are complementary to various other ATLAS A' searches. As data-taking continues at the LHC, the reach of lepton-jet analyses will continue to expand in model coverage and in parameter space.

To my parents, who have been there for me through all the collapses of the quantum wavefunctions.

Acknowledgements

The first acknowledgement goes to my supervisor, William Trischuk, who provided me with the guidance and cookies required to carry out this work. I would also like to thank the entire University of Toronto ATLAS group for their support and camaraderie throughout my doctoral studies, especially professors Peter Krieger, Bob Orr, Pierre Savard, Pekka Sinervo, and Richard Teuscher. Special thanks to theorist Bob Holdom for welcoming my experimentalist visits to his office. High-fives go out to the Displaced LJs analysis team: Antonio Policicchio, Monica Verducci, Stefano Giagu, Marco Schioppa, Robert Les, Michela Del Gaudio, Marianna Testa, and in memory of the late Guido Ciapetti. The mentorship of Art McDonald has been greatly important to me in my particle physics adventures. Finally, I am eternally grateful to my cardiologist Blair Grubb for making my life as I know it possible, and to NASA flight surgeon Smith Johnston for helping me get off the ground.

Contents

List of Tables

List of Figures

List of Appendices

Chapter 1

Introduction

The search for dark matter (DM) is amongst the most exciting frontiers in the quest for Beyond the Standard Model (BSM) physics. Having recently completed the SM with the discovery of the 125 GeV Higgs boson, general-purpose detectors at the world's most powerful particle accelerator, the Large Hadron Collider (LHC) at CERN, are aspiring to now shed light on DM. Run 2 of the LHC, with pp collisions at $\sqrt{s} = 13$ TeV, began in early 2015 and offers exciting new discovery potential with respect to the Run 1 ($\sqrt{s} = 7 - 8$ TeV) results.

Overwhelming evidence for the existence of DM comes from the realm of astrophysics, including baryogenesis, large-scale structure formation, cosmic microwave background radiation properties, galaxy cluster collisions, galactic rotation curves, and gravitational lensing. The most recent estimates from PLANCK [1] indicate that DM makes up a quarter of our universe, five times as much as SM material. All this evidence, however, is based on the gravitational interactions of the DM. While the DM is commonly assumed to be composed of BSM particles – rather than non-luminous conglomerations of SM material, or even a non-existent artefact of modified gravity – the hypothesized particles' properties and non-gravitational interactions remain shrouded in mystery.

The leading DM paradigm over the past several years has been that of the Weakly Interacting Massive Particle (WIMP), constituting "cold" dark matter. According to this paradigm, DM is composed of particles with masses near the weak scale, moving non-relativistically, with small couplings to SM particles. Rich worldwide programs, falling into three main categories, have been searching for signs of WIMPs. The first category, direct detection experiments, waits for WIMPs to collide with nuclei in large-volume, low-background detectors, and seeks to observe the resulting nuclear recoils. The second category, indirect detection, searches the sky for signatures (such as electrons, positrons, and gamma rays) of WIMP self-annihilations in the cosmos. The third category, collider searches, seeks to produce and detect WIMPs from high-energy collisions of SM particles in the laboratory. The three categories are complementary: in probing different energy scales, they may in fact probe different processes involving SM-DM interactions (see Figure 1.1). Meanwhile, some WIMP models are accessible through all three categories, but with each category probing a different area of parameter space.

While great strides have been made in extending the reach of all three categories of searches, no definitive WIMP discovery has yet been made. Run 2 of the LHC has ruled out large swaths of parameter space in minimal WIMP models with masses ranging from the weak scale up to a few TeV. Meanwhile, several direct and indirect detection anomalies have accumulated, which are difficult to explain with the minimal WIMP models. This has spurred interest in a wide variety of DM models that depart from the WIMP paradigm altogether, but also in "dark sector" models where the WIMPs are part of an entire dark particle zoo. Given the intricate structure of the SM, which describes only a fifth of the matter density of the universe, it is hardly unreasonable to expect a dark sector with its own intricate structure. This structure may well include new "dark forces", such as a version of electromagnetism carried by a dark gauge boson, acting between dark sector particles. While the presence of new long-range forces in the universe may seem surprising, the weakness of the interactions between the SM and the dark sector implies that such dark forces may easily have gone undetected in all experiments so far.

Faced with a vast sea of possible DM candidates (see Figure 1.2), where should we now focus our searches? One approach is to "look for the portal". In short, all our attempts to detect DM other than through its gravitational interactions assume that some non-gravitational coupling(s) – in other words, portal(s) – between the DM and the SM actually exist. Since we must begin our searches from our known base, the SM, it is eminently reasonable to target what is most accessible to us – the portal(s) rather than the DM itself – as our first step. Fortunately for us, there is a long history in the BSM literature, not necessarily tied to the BSM context, of studying possibilities for adding extra interactions to the SM. The viable extra interactions, now cast as portals, are tightly restricted by the symmetries of the SM. They fall into 4 categories: hypercharge, Higgs, neutrino, and axion. General BSM searches in these categories become re-cast as DM searches by adding the constraints from astrophysical observables.

Two of these portals, hypercharge and Higgs, are relevant to models where a massive dark gauge boson mixes with the SM photon. Such models provide rich phenomenology, with varying experimental signatures in different regions of parameter space. Collider experiments are able to probe certain signatures, and therefore certain regions, that are inaccessible to direct and indirect detection experiments. One of these signatures consists of "lepton-jets" (LJs), where highly-boosted, fairly low-mass dark gauge bosons decay to pairs of highly-collimated leptons. Depending on the dark gauge boson lifetime, the LJs can be prompt, i.e. produced at the primary vertex of a collision event, or displaced. ATLAS (A Toroidal LHC Apparatus [2]) and CMS (Compact Muon Solenoid [3]), the two large multi-purpose detectors at the LHC, conducted LJ searches in Run 1. This work documents a Run 2 search for displaced LJs in ATLAS, using 3.57 fb^{-1} of pp collision events recorded in 2015. Improvements over the Run 1 search [4], particularly in trigger efficiency, combine with the increased Higgs boson production cross-section at higher $\sqrt{s}$ to give the Run 2 search sensitivity to dark gauge bosons over a larger swath of parameter space.

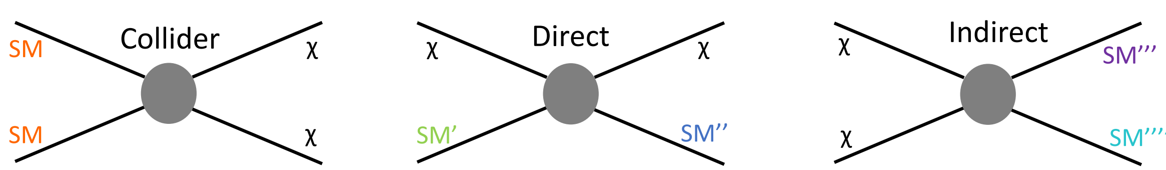

Figure 1.1: Schematic representation of the DM processes probed by collider, direct detection, and indirect detection experiments. The particle species represented by SM through SM"" may or may not be the same.

Given that LHC data has ruled out large swaths of parameter space in traditionally-favoured WIMP and supersymmetric models, collider search programs are increasingly turning their attention to "exotic"

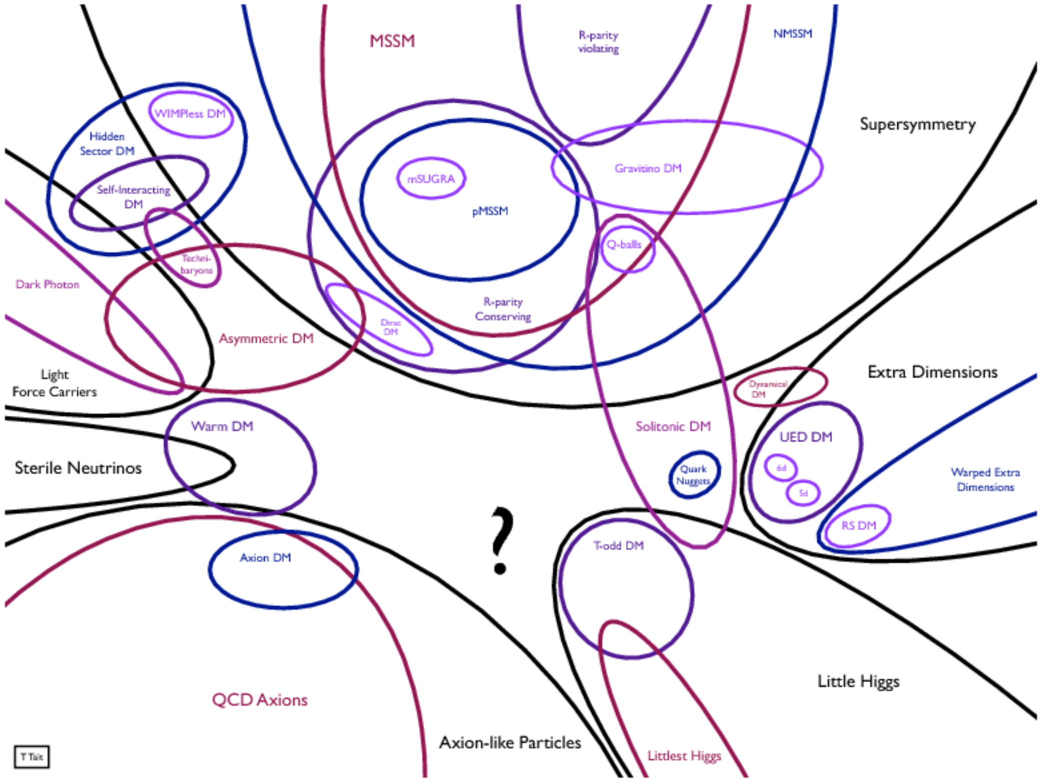

Figure 1.2: A graphical representation of the sea of DM candidates currently circulating in the literature, with no particular quantities on the axes. Dark photons appear in violet, midway up the picture at the left edge. [5]

signatures in hopes of making a discovery. LJ analyses, requiring innovative triggers and non-standard object reconstruction, are on the front lines of this shift, and are complementary in parameter space to other ATLAS exotic searches. The results presented here solidify the place of LJ searches in the ever-expanding ATLAS dark gauge boson search program, and are in part available in ATLAS conference note [6]. Chapters 2 and 3 of this work discuss the theory and collider phenomenology of dark gauge boson models. Chapter 4 summarizes the relevant workings of the ATLAS detector. Chapter 5 describes the methodology and analysis flow of the displaced LJ search, entirely blinded (i.e. without examining the signal region in the 2015 data). Chapter 6 gives the (unblinded) results, and their interpretation as limits in model parameter space. Description of the systematic uncertainties is reserved for Chapter 7. Finally, Chapter 8 discusses future extensions for ATLAS Run 2 LJ searches.

Chapter 2

Dark Gauge Boson Models

This chapter begins by summarizing the theory framework, as well as the notation, for the Hypercharge and Higgs portals. It goes on to discuss models constructed "top-down" (i.e. from overarching theoretical motivations), and then models constructed "bottom-up" (i.e. from DM phenomenological motivations), that incorporate both these portals. The final section outlines the primary motivations for dark gauge boson searches at the LHC in particular.

2.1 Hypercharge Portal

There is a long history in the literature of theories that extend the SM by adding a new U(1) gauge symmetry, hereafter referred to as U(1)'. These models generally assume the SM fields are all fundamentally uncharged under U(1)', but note that the occurrence of "kinetic mixing" between the SM U(1) and the BSM U(1)' gives SM fermions charged under U(1) a small effective charge under U(1)' as well. The U(1)' can occur as the only BSM gauge symmetry, i.e. in SU(2) × U(1) × U(1)', or in conjunction with BSM non-Abelian gauge groups G_n. Going all the way back to the first widely-cited U(1)' models, the mass of the vector boson mediator of the U(1)' can range from zero (as in the "paraphoton" of [7]) to heavier than the Z (as in the "second Z" of [8]). The lighter the mediator, the longer range a force it conveys; a massless mediator implies an infinite-range force similar to electromagnetism. The presence of a new long-range force in nature may seem surprising, but because it would act only weakly on SM fermions, it may indeed have gone unnoticed in all experiments so far.

This work will focus on masses between the MeV and the weak scale, rather than on the massless (paraphoton) case or on the heavier-than-Z case. Notation in the literature has become somewhat inconsistent over the years; in accordance with the most common convention, this work will refer to the U(1)' field as A'_μ, its mediator as a "dark gauge boson" A', the fairly low-mass (1 MeV $\lesssim m_{A'} \lesssim$ 10 GeV) case as the "dark photon" γ_{d} , and the fairly massive (10 GeV $\lesssim m_{A'} \lesssim m_Z$) case as the "dark Z" Z_{D} .

Assuming the A' is massive and that all SM fields are uncharged under U(1)', the relevant gauge terms in the effective Lagrangian [9] are:

$$\mathcal{L} \subset \frac{-1}{4}\hat{B}_{\mu\nu}\hat{B}^{\mu\nu} - \frac{1}{4}\hat{A'}_{\mu\nu}\hat{A'}^{\mu\nu} + \frac{\epsilon}{2\cos\theta_W}\hat{A'}_{\mu\nu}\hat{B}^{\mu\nu} + \frac{1}{4}m_0^2\hat{A'}^{\mu}\hat{A'}_{\mu} \tag{2.1}$$

with the SM U(1)$_Y$ field strength $\hat{B}_{\mu\nu} = \partial_\mu \hat{B}_\nu - \partial_\nu \hat{B}_\mu$, the U(1)' field strength $\hat{A'}_{\mu\nu} = \partial_\mu \hat{A'}_\nu - \partial_\nu \hat{A'}_\mu$, the Weinberg mixing angle θ_W, the dark gauge boson mass term m_0 , and the kinetic mixing parameter ϵ. The kinetic mixing, i.e. the $\hat{A'}_{\mu\nu}\hat{B}^{\mu\nu}$ operator, is what gives fermions charged under U(1) a small effective charge under U(1)' as well.

The kinetic terms can be diagonalized via the field redefinition:

$$\begin{pmatrix} A'_0 \\ B \end{pmatrix} = \begin{pmatrix} \sqrt{1-\frac{\epsilon^2}{\cos^2\theta_W}} & 0 \\ \frac{-\epsilon}{\cos\theta_W} & 1 \end{pmatrix} \begin{pmatrix} \hat{A}' \\ \hat{B} \end{pmatrix}.$$

After electroweak symmetry breaking, with this field redefinition the mass matrix can be written in the basis $(A^\mu, {Z_0}^\mu, {A'_0}^\mu)$ as:

$$\mathcal{M}^2 = (m_0^2/\delta^2) \begin{pmatrix} 0 & 0 & 0 \\ 0 & 1 & -\eta\sin\theta_W \\ 0 & -\eta\sin\theta_W & \eta^2\sin^2\theta_W + \delta^2 \end{pmatrix}$$

with $\eta = \frac{\epsilon}{\cos\theta_W\sqrt{1-\frac{\epsilon^2}{\cos^2\theta_W}}}$, and $\delta = \frac{4m_0}{(g^2+g_D^2)v^2}$ where v is the Higgs vacuum expectation value (vev) and g_D is the U(1)' gauge coupling. While the SM photon remains massless (since electromagnetism remains unbroken), the other two fields mix, with the mass eigenstates Z and A' derived via:

$$\begin{pmatrix} Z \\ A' \end{pmatrix} = \begin{pmatrix} \cos\alpha & \sin\alpha \\ -\sin\alpha & \cos\alpha \end{pmatrix} \begin{pmatrix} Z_0 \\ A'_0 \end{pmatrix}$$

with the mixing angle

$$\tan\alpha = \frac{1-\eta^2\sin^2\theta_W - \delta^2 - \text{Sign}(1-\delta^2)\sqrt{4\eta^2\sin^2\theta_W + (1-\eta^2\sin^2\theta_W-\delta^2)^2}}{2\eta\sin\theta_W}.$$

The masses are then:

$$m_{Z,A'}^2 = \frac{m_0^2}{2\delta^2}\left(1+\delta^2+\eta^2\sin^2\theta_W \pm \text{Sign}(1-\delta^2)\sqrt{(1+\delta^2+\eta^2\sin^2\theta_W)^2 - 4\delta^2}\right).$$

In the limit of small mixing ($\epsilon \ll 1$ and $\delta \ll 1$), the mass formulas reduce to:

$$\begin{aligned} m_{A'}^2 &\simeq m_0^2(1-\eta^2\tan^2\theta_W) \\ m_Z^2 &\simeq (m_0^2/\delta^2)(1+\eta^2\tan^2\theta_W). \end{aligned} \quad (2.2)$$

The $m_{A'}$ comes out to be less than m_Z for a wide range of η and δ.

This paradigm can be expanded by also allowing for kinetic mixing between the U(1)' and the SM weak neutral current $\hat{J}_\mu^{NC}$, adding to the Lagrangian an extra term

$$\epsilon_Z \hat{A'}_\mu \hat{J}_\mu^{NC}$$

where ϵ_Z is the weak kinetic mixing parameter. This creates what is known as "mass matrix mixing" between the A' and the Z. However, built into ϵ_Z is a factor of $(m_{A'}/m_Z)^2$ with respect to ϵ. For small

$m_{A'}$ (i.e. for the $\gamma_{\rm d}$ in this work), the suppression of ϵ_Z means the mass matrix mixing can be neglected.

In the case of massless A', the U(1)' must be kept unbroken, and kinetic mixing can induce tiny charges under the SM U(1) for the BSM particles charged under U(1)'. These BSM particles then manifest as "millicharged particles", i.e. possessing a small fraction of the electron's electromagnetic charge [7]. As long as the A' has mass, however, the BSM fields charged under U(1)' do not acquire electromagnetic millicharge.

In modern terms, the kinetic mixing has become known as the "hypercharge portal" or "vector portal" to any new sector consisting of BSM particles charged only under U(1)'. The "portal" terminology is appropriate to emphasize that the kinetic mixing is the dark gauge boson's only renormalizable interaction with the SM.

2.2 Higgs Portal

A further class of models, known as Hidden Abelian Higgs [11], arises if the U(1)' is spontaneously broken by a Higgs mechanism to give the A' its mass. It requires adding to the theory a SM-singlet "dark scalar" field S, with U(1)' charge q_S. Then s_0, the uneaten singlet-Higgs-like component of S, in general has a renormalizable coupling to the SM-like Higgs boson h_0. Interest in the possibility of mixing between h_0 and s_0, known as the "Higgs portal", has been steadily rising ever since the discovery of the 125 GeV SM-like Higgs boson.

A Higgs mechanism is not necessarily the only way to generate a mass for the A'. In string models (see Section 2.3), alternatives include the Stuckelberg mechanism or a non-zero Fayet-Iliopoulos term. Outside the realm of string models, however, a massive A' is often assumed to be accompanied by at least one dark Higgs boson. Meanwhile, the dark Higgs sector may be more complicated than a single scalar, analogous to extended SM Higgs sectors (two-Higgs-doublet, etc). This work will assume only one dark Higgs boson.

In electroweak symmetry breaking in Hidden Abelian Higgs models, not only does the SM Higgs doublet H acquire a vev $< H >= (0, v/\sqrt{2})$, but the dark singlet S also acquires a vev: $< S >= v_S/\sqrt{2}$. Then the general renormalizable potential V_0 for the scalar fields [9] is

$$V_0(H,S) = -\mu^2|H|^2 + \lambda|H|^4 - \mu_S|S|^4 + \kappa|S|^2|H|^2 \tag{2.3}$$

where κ is the Higgs mixing parameter, $\mu^2 = v^2\lambda + \kappa v_S^2/2$, and $\mu_s^2 = v_S^2\lambda + \kappa v^2/2$. Expanding in small fluctuations h_0 and s_0 about the vacuum, the Higgs mass matrix in the (h_0, s_0) basis is:

$$\mathcal{M}^2 = \begin{pmatrix} 2v^2\lambda & v\ v_S\kappa \\ v\ v_S\kappa & 2v_S^2\lambda^S \end{pmatrix}.$$

The mass eigenstates h and s are then defined via:

$$\begin{pmatrix} h \\ s \end{pmatrix} = \begin{pmatrix} \cos\theta_h & \sin\theta_h \\ -\sin\theta_h & \cos\theta_h \end{pmatrix} \begin{pmatrix} s_0 \\ h_0 \end{pmatrix}$$

with mixing angle

$$\tan\theta_h = \frac{v^2\lambda - v_S^2\lambda_S - \mathrm{Sign}(v^2\lambda - v_S^2\lambda_S)\sqrt{v^4\lambda^2 + v_S^4\lambda_S^2 + v^2v_S^2(\kappa^2 - 2\lambda\lambda_S)}}{v\ v_S\kappa}.$$

The Higgs boson masses are then

$$m_{h,s}^2 = v^2\lambda + v_S^2\lambda_S \pm \mathrm{Sign}(v^2\lambda - v_S^2\lambda_S)\sqrt{v^4\lambda^2 + v_S^4\lambda_S^2 + v^2v_S^2(\kappa^2 - 2\lambda\lambda_S)}$$

and the dark photon mass term is

$$m_0 = g_D q_S v_S.$$

For small θ_h, the h is dominantly SM-Higgs-like and the s is dominantly singlet-Higgs-like. Then these mass formulas reduce to:

$$\begin{aligned} m_h^2 &\simeq 2\lambda v^2 + 2s_h^2(\lambda v^2 - \lambda_S v_S^2) \\ m_s^2 &\simeq 2\lambda_S v_S^2 + 2s_h^2(\lambda v^2 - \lambda_S v_S^2) \end{aligned} \quad (2.4)$$

with $s_h = \frac{\kappa v\ v_S}{2v_S^2\lambda_S - v^2\lambda} \simeq \tan\theta_h$. It is generally assumed that the h is the 125 GeV boson discovered at the LHC.

2.3 Hidden Valley, SUSY, and String Models

In modern terms, "hidden sector" scenarios are defined as models in which a sector of SM-singlet particles, interacting amongst themselves, is weakly coupled to the SM. This coupling may be generated by fields at or above the TeV scale, appearing in irrelevant operators (operators whose high dimensionality suppresses them by the mass scale or powers of the mass scale). In such cases, in the limit that the mass scale goes to ∞, the hidden sector becomes completely divorced from the SM. Alternatively, the coupling may be generated by marginal operators (such that the couplings are independent of the mass scale), with very small coupling constants.

A particular sub-class of these models, known as "hidden valleys", has recently been garnering increasing attention in the literature. In hidden valleys, the lowest one or more of the hidden particles is forced to decay to SM particles rather than to other hidden particles, due to a mass gap or symmetry [12]. The increasing popularity of "hidden valley" models is in part due to their appearance in various supersymmetric (SUSY) constructions that evade the ever-more-stringent experimental constraints, and in various string theory constructions. But it also stems in part from pure optimism, as the decays of hidden particles to SM particles give us hope of observing smoking-gun signatures in multiple different types of experiments, including hadron colliders, lower-energy e^+e^- colliders, and high-intensity fixed-target experiments.

In hidden valley models, some "portal", i.e. neutral particle or particle/anti-particle pair, couples both to SM operators and to operators made from valley fields. The hypercharge portal and the Higgs portal both fit this description, and the A' can serve as the lowest valley particle. The range of theoretically viable valleys is huge; given the complexity of the SM, there is no well-motivated reason to expect the valley to be minimal. This huge range can in fact be seen as a strength of the model from a "naturalness" point of view, as it implies a lack of need for fine-tuning of the hidden valley. Non-minimal hidden valley models often predict the production of multiple A's in cascade decays of heavier valley particles. Some also feature hadronization and showering of valley particles, with hidden hadrons of widely varying lifetimes (in analogy to QCD hadrons) producing A's in clusters, e.g. in [13]. The coupling strength between the A' and the valley particles is controlled by the dark coupling constant α_d; moderate to high α_d causes valley particles to radiate copious numbers of A's.

Low-mass hidden sectors can be easily incorporated into SUSY models, by forcing the hidden sector to softly break SUSY and then communicate SUSY-breaking to the visible sector. Such models have attracted interest because a low-mass hidden Higgs boson can ameliorate the little hierarchy problem, and SUSY particles can evade existing experimental constraints by decaying into the hidden sector [14]. SUSY makes the "portal" much more complicated, e.g. with kinetic mixing involving a dark bino and gauginos. That being said, the salient point in this context is that the lightest visible superpartner decays into a cascade of hidden sector particles, producing copious numbers of dark gauge bosons along the way.

Finally, string models often incorporate hidden sectors and include gauge groups containing one or more U(1) factors. In fact, in compactifications of Type II strings to four space-time dimensions, hidden gauge factors are a common feature in the low-energy effective field theory. The interactions of the corresponding hidden gauge bosons with visible particles are loop-suppressed, leading to weak coupling. For example, a stack of N coincident branes generically gives rise to a U(N) gauge group that decomposes into SU(N) × U(1). The SM may correspond to one such brane-stack, and the hidden sector to another such brane-stack. Open strings, each with one end stuck to the SM brane-stack and the other end stuck to the hidden sector brane-stack, would then generate kinetic mixing between the SM U(1) and the hidden U(1)' [15]. In other examples, the U(1)' lives on a spacetime-filling D7 brane, or on anti D3-branes [16].

2.4 Dark Matter Models

The previous sections have outlined models built "top-down", motivated from the realm of BSM theory. Recently, however, several models built "bottom-up", from the realm of phenomenology, have used dark gauge bosons to address specific indirect and direct detection observations that may constitute hints of DM.

2.4.1 Indirect and Direct Detection Anomalies

The motivation for indirect detection experiments is that DM particles randomly colliding in the cosmos could annihilate with each other to produce e^+e^- pairs, which would in turn annihilate to γ's. Earth-based or space-based detectors and telescopes can then hope to spot the excess $e^+/e^-/\gamma$'s coming from outer space. DM annihilation would presumably occur more frequently in areas with higher DM density, such as "halos" surrounding large collections of visible matter, as determined by gravitational clustering. The fluxes, spectra, and directional origin of the $e^+/e^-/\gamma$'s would therefore be dependent on the DM density profile, as well as on the particular nature of the DM. The galactic centre is an especially interesting place to search for DM signatures, since the density of the DM halo should be high there. Reliably identifying $e^+/e^-/\gamma$ excesses over astrophysical and cosmic backgrounds presents a host of challenges, and various poorly-understood exotic astrophysical phenomena may mimic DM indirect detection signatures. Nonetheless, great technical strides have been made in indirect detection experiments over the past several years.

The most relevant indirect detection anomalies observed to date are:

- An excess in e^+ fraction in $\sim 10-50$ GeV cosmic rays. Initially seen by HEAT (High Energy Antimatter Telescope, a balloon-borne instrument) and AMS-01 (Alpha Magnetic Spectrometer, on the International Space Station), and then seen by PAMELA (Payload for Antimatter Matter Exploration and Light-nuclei Astrophysics, a satellite-borne instrument) [17]. Possibly supported by Fermi Gamma-Ray Space Telescope [18], and possibly related to an excess in $e^+ + e^-$ at 500-800 GeV seen by ATIC (Advanced Thin Ionization Calorimeter, a balloon-borne instrument) [19];

- A haze in the cosmic microwave background near the galactic centre. Seen by WMAP (Wilkinson Microwave Anisotropy Probe, a satellite-borne instrument), possibly compatible with a 1-3 GeV gamma-ray excess in this area seen by Fermi Gamma-Ray Space Telescope [20] [21];

- A 511 keV emission line from the galactic centre, compatible with the annihilation of positrons at energies of a few MeV. Seen by EGRET (Energetic Gamma Ray Experiment Telescope, on-board the Compton Gamma Ray Observatory) and by an imaging spectrometer on-board INTEGRAL (the European Space Agency observatory) [22], as well as in observations dating back to the 1970s by various balloon-borne experiments.

Meanwhile, direct detection experiments wait for DM particles to collide with nuclei in large-volume, low-background detectors, and seek to observe the nuclear recoils from the collisions. A wide variety of detector materials and designs have been employed to target different DM-nucleus interaction types and recoil energies. While almost all direct detection experiments have reported results consistent with expected SM backgrounds, the DAMA/LIBRA experiment [23] has reported a longstanding seasonally-oscillating signal. Two other experiments, CDMS [24] and CoGeNT [25], have also reported hints of signals, which may or may not be consistent with the DAMA signal depending upon a variety of assumptions [26].

Feasible non-DM explanations do exist for each of the above anomalies. The natural question to ask, though, is whether the dominant DM paradigm over the past several years – that of "cold DM" in the form of WIMPs – can simultaneously explain most or all of these anomalies. Simple "collisionless" WIMP models, i.e. with no BSM forces acting between WIMPs, have great difficulty accounting for multiple indirect detection excesses when all the details are worked out. Even more difficult for them is generating a signal in DAMA (and/or CDMS and/or CoGeNT) but not in other direct detection experiments. Enter models in which the role of the A' is to convey a dark force amongst heavier hidden sector particles that are the WIMPs. In direct detection experiments, the A' can significantly increase the DM-nucleus scattering cross section, or alter the kinematics of the scattering. In indirect detection, the A' can significantly enhance the DM annihilation rates at low velocities. A comprehensive review of such models, and their relevance to each of the above anomalies, can be found in [26]; a brief summary will be given in the next section.

There are some models in which the A' itself is the DM; these models generally require departing from the cold DM paradigm entirely. For example, for $m_{A'} \sim 0.1$ MeV and $\epsilon \sim 10^{-12}$, the A' may be a "lukewarm" DM candidate [27]. However, such models have difficulty explaining some or all of the above anomalies, and the A' mass and kinetic mixing fall far below the ranges considered in this work.

2.4.2 Secluded DM

According to the "thermal relic" paradigm, the WIMPs were in thermal equilibrium with the SM in the early universe. As the universe expanded and cooled, the WIMPs experienced "freeze-out", i.e. dropped out of thermal equilibrium with the SM. The time – or equivalently, temperature – of freeze-out depends upon the thermally-averaged annihilation cross-section times velocity ($< \sigma \times v >$) of the WIMPs. The amount of DM present in the universe at freeze-out gives the "relic density", which has presumably persisted to this day. The relic density is therefore known from measurements of the current energy density of DM in the universe, $\Omega_{DM} h^2 \approx 0.12$ with h the Hubble constant [1]. This relic density is achieved for $< \sigma \times v > \approx 3 \times 10^{-26}$ cm^3/s; the fact that this corresponds approximately to the expectation for an electrically-neutral particle with weak-scale mass interacting via the weak force is known as the "WIMP miracle". There is some flexibility in the relic density constraints, in that too large a value of $< \sigma \times v >$, leading to too small a relic density, doesn't necessarily rule out a DM candidate. It is always possible that the candidate in question is just one of multiple DM species, where other species make up the rest of the relic density. Too small a value of $< \sigma \times v >$, leading to too high a relic density, does rule out a candidate, as it "over-closes" the universe with too much DM. A WIMP model can also be ruled out if it interferes with "Big Bang nucleosynthesis", i.e. if it changes the amounts of hydrogen and helium formed in the early universe.

Attempts to fit simple collisionless thermal WIMP models to the e^+ excess typically yield a $< \sigma \times v >$ one to three orders of magnitude too large. In addition, the WIMPs would have to display an overwhelming preference for annihilation to hard leptons over annihilation to hadrons or soft leptons. All the obvious SM candidates for annihilation channels – W, Z, Higgs, or heavy quarks – fail in this regard. What is needed, therefore, is a BSM way to enhance the DM annihilation cross-section in the galactic halo without affecting the thermal relic abundance, and to ensure the WIMP annihilation preferentially produces hard leptons. One such category of thermal relic WIMP models is known as "secluded DM", in which thermal relic WIMPs ψ annihilate to metastable mediators with delayed decays back to SM states [28]. The A' is an example of a possible metastable mediator, where the WIMPs are charged under U(1)'. The s is generally also included in the model, such that spontaneous symmetry breaking acts on the U(1)' near the weak scale. The two primary mechanisms responsible for WIMP annihilation are then:

$$\begin{aligned} \psi + \overline{\psi} &\rightarrow A'^{*} \rightarrow Z^{*}/\gamma^{*} \rightarrow \text{SM states} \\ \psi + \overline{\psi} &\rightarrow 2A' \rightarrow \text{SM states} \end{aligned} \tag{2.5}$$

where * denotes virtuality.

ψ remains in equilibrium with the SM in the early universe by interacting with the SM thermal bath via the A'. The parameters of the model are constrained by the requirements that the A' and s must decay before the start of nucleosynthesis, A' decays must remain in thermal equilibrium, and the correct WIMP relic abundance must be achieved by the ψ. The first two constraints translate into limits on the lifetime of the A', which depends upon ϵ and κ. The third constraint translates into limits on the $< \sigma \times v >$ of the ψ.

There are two regimes for this type of model: $m_\psi < m_{A'}$ and $m_\psi > m_{A'}$. In the first regime, the on-shell process of Equation 2.5 is kinematically disallowed, and so fairly large mixing is required in order to ensure sufficient ψ annihilation cross-section, and the model is subject to numerous experimental constraints. In the second regime, the on-shell channel opens up, and the mixing can be taken small enough to evade these constraints. This makes the second regime more promising from the point of view of compatibility with experiment. In either case, there is an interesting region of parameter space in the WIMP mass range of $\mathcal{O}(\text{MeV}) > 2m_e$, as the WIMP annihilations could then produce the 511 keV photons and/or the positron excess while evading other experimental constraints. In the case of $2m_e \lesssim m_{A'} \lesssim m_\psi \alpha_d$, there is a "recombinant annihilation" process $\psi + \overline{\psi} \to \psi\text{onium} + A'$. This recombinant process would be significantly enhanced in the galactic DM halo relative to the on-shell process of Equation 2.5 which would have dominated at freeze-out. The ψonium would eventually decay to A's, and then to SM states with light leptons heavily favoured – possibly accounting for the excess positron flux. Extensive discussions of the model constraints and potential explanatory power are given in [28] and [29].

2.4.3 Inelastic and Exciting DM

A further category of models (e.g. [30], [31] [32]) unites secluded DM with two more theories, Inelastic Dark Matter (iDM) and Exciting Dark Matter (XDM), under a single explanatory framework of SUSY + hidden gauge symmetries.

The XDM model [33] features a dark sector that includes two or more nearly-degenerate WIMP states. The higher-energy state serves as an "excited" state, such that inelastic scattering between two WIMPs could raise one or both of them from the lower state. The inelastic scattering would be mediated by a A'. Assuming an energy splitting $\mathcal{O}(\text{MeV})$ greater than $2m_e$, subsequent WIMP decay back to the lower state would likely be accompanied by the emission of a e^+e^- pair that could then produce the 511 keV photons.

The iDM model [34] features a dark sector that includes a massive particle charged under a spontaneously-broken dark gauge symmetry. The dark particle's symmetry-breaking mass is assumed to be much smaller than its symmetry-preserving mass, splitting its mass eigenstates $\chi_1,\ \chi_2$ such that its interactions with itself and with the SM are primarily off-diagonal (i.e. between different mass eigenstates). DM annihilation would therefore occur primarily through "co-annihilation", i.e. χ_1 with χ_2. This model has 3 important BSM parameters besides $m_{A'}$ and ϵ: $m_{\chi_1},\ m_{\chi_2} - m_{\chi_1},\ \alpha_d$. The heavier eigenstate χ_2 may have been abundant in the early universe, facilitating efficient co-annihilation, but depleted post-freezeout. Then post-freezeout, the average velocities of DM particles in halos would be insufficient to up-scatter χ_1's into χ_2's, and the relative lack of χ_2's would shut down the co-annihilation. Most direct detection signatures would thereby be suppressed, as they rely on annihilation. Meanwhile, the DAMA excess could potentially be explained with mass splitting $\mathcal{O}(100\text{ keV})$, by interactions between a χ_1 or χ_2 with a nucleus. The explanatory power and experimental signatures of each of iDM and XDM are discussed extensively in [35].

The frameworks of [30] and [31] propose a multiplet of dark sector states at masses $\mathcal{O}(500\text{-}800)$ GeV, and a non-Abelian A'. Splittings in the multiplet of $\mathcal{O}(100\text{ keV})$ create iDM effects; other splittings in the multiplet of $\mathcal{O}(\text{MeV})$, along with inelastic scattering mediation by the non-Abelian A', create XDM effects. Furthermore, the WIMPs in the multiplet would annihilate into leptons at a level well above what would be expected from a traditional thermal relic, due to a Sommerfeld enhancement in the annihilation cross-section brought about by the A'. This enhanced annihilation could explain the positron excess, and create a flux of e^+e^- pairs and gamma rays from the galactic centre. The electrons would synchroton-radiate in the galactic magnetic field to produce the WMAP haze, while the gamma rays would account for the Fermi observation. The A' would have to be low-mass, with Compton wavelength $\gtrsim 1$ GeV^{-1}, in order to produce the necessary enhancement. The non-Abelian nature of the A' also requires, in general, multiple dark Higgs bosons. The model can fruitfully be developed further by incorporating SUSY, as making the dark sector SUSY-breaking naturally puts the A' at the GeV scale and the WIMP mass splittings at the MeV scale. This leads to "SuperDark" models that look much like the hidden valley SUSY models; kinetic mixing serves as the leading interaction between the dark sector and the (N)MSSM. The enhanced gamma ray production implications of such SUSY-based models are explored in [36]. The framework of [32] offers a simplification of the above, by employing only an Abelian U(1)' and a single dark Higgs boson in conjunction with SUSY to generate the necessary WIMP mass splittings and GeV-scale U(1)' symmetry-breaking.

2.5 Motivations for Dark Gauge Boson Searches at the LHC

Given the above models, dark gauge bosons earn a place in the LHC search programme from a variety of perspectives.

With the discovery of a SM-like Higgs boson now a few years old, shedding light on DM has taken centre stage in the LHC's discovery quest. A host of WIMP searches at LHC (most of which target a single SM object recoiling against missing transverse energy) have so far yielded no discovery, joining the ranks of direct and indirect detection experiments in setting ever more stringent limits. Meanwhile, models for the particle content of the dark sector have become ever more complicated, with trickier ways to evade bounds from traditional WIMP searches. The situation therefore calls for the consideration of other, complementary approaches. Whenever we assume there is hope of detecting DM through any means other than its gravitational influence, we are automatically assuming that the DM has some non-gravitational coupling to the SM – in other words, that some dark portal does exist. Setting out from the realm of the SM, there are great advantages to focusing on the portal as our nearest target, rather than on the rest of the dark sector. This frees us from many of the contingencies of the unknown dark sector particle content and mass hierarchies, giving us a high degree of model independence and discovery potential. In non-minimal hidden valley models, for example, the A' may very well be the best candidate for collider detection amongst the hidden zoo, especially in cases where long decay chains yield high A' multiplicity.

The A' has the additional allure of showing up not only in various "top-down" theoretical frameworks (such as extended SUSY scenarios), but also in various 'bottom-up" DM models. This gives the A' a place alongside traditional SUSY searches, as well as in less-conventional "exotics" DM searches.

Post-Higgs boson discovery, the LHC is aiming not only for the next big discovery, but also for precision measurements of the SM-like Higgs boson. A major aspect of this programme is the determination of $h \to$ SM and $h \to$ invisible BRs. If the Higgs portal is present, the A' could make a noticeable contribution to the Higgs boson decays. Not only would we take advantage of the Higgs boson to extend our A' discovery potential, but we would need to take the A' into account in order to understand the Higgs boson's behaviour.

Chapter 3

Collider Phenomenology of Dark Gauge Bosons

This chapter begins by summarizing the A' production and decay channels relevant in collider searches. It then defines the scope and role of LJ searches in the wider A' search program in ATLAS. Next comes a high-level sweep over the parameter space of A' models, outlining the current experimental constraints and delineating the region to be explored in this work. This is followed by a detailed description of A' branching ratios and lifetimes. Finally, the benchmark models are introduced that will be employed throughout the analysis.

3.1 Production and Decay Channels

Assuming the A' is massive enough to decay into leptons and light hadrons f, the following A' channels are most relevant at colliders:

- Drell-Yan $pp \to A' \to f^+f^-$
- Singly-exotic Higgs boson decay $h \to Z^*A' \to 4f$
- Doubly-exotic Higgs boson decay $h \to 2A' \to 4f$
- Hidden Valley Higgs boson decay $h \to$ Dark States $\to nA' \to 2nf + X$ (where n is the A' multiplicity, and X may or may not include collider-stable hidden states manifesting as missing transverse energy)
- Dark SUSY $pp \to$ SUSY States $\to nA' \to 2nf + X$

The first three processes (Figure 3.1) can be considered on their own, or as building blocks in the hidden valley and dark SUSY models.

If only the hypercharge portal is present, with no Higgs portal nor SUSY, the Drell-Yan channel is most sensitive, with the singly-exotic Higgs channel offering slightly weaker sensitivity. The Drell-Yan channel is covered by general ATLAS dilepton searches, i.e. bump-hunts in the $e^+e^-/\mu^+\mu^-$ spectrum. At $m_{A'} \lesssim 10$ GeV, Drell-Yan limits are not competitive with existing limits from other types of experiments. Limits for $m_{A'} \gtrsim 160$ GeV were obtained in [37] by re-casting a dilepton resonance search in 20 fb^{-1}

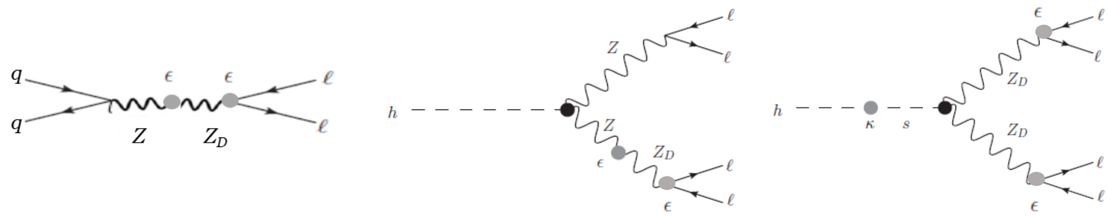

Figure 3.1: The three basic A' channels most relevant at colliders: Drell-Yan (left), singly-exotic Higgs boson decay (centre), and doubly-exotic Higgs boson decay (right).

of 8 TeV ATLAS data, which provided slightly better sensitivity than electroweak precision (EWP) experiments. The EWP constraints in this mass range, based on global fits to measurements made at LEP, Tevatron, and LHC, are documented in [9]. Projections for 10 GeV $\lesssim m_{A'} \lesssim$ 80 GeV, for hypothetical specialized analyses of the 8 TeV data, were estimated in [38] and show good potential for again exceeding the sensitivity of EWP. Meanwhile, singly-exotic Higgs boson decay was the subject of a dedicated ATLAS search in [39], hunting for a resonance in the $Z^* \to e^+e^-/\mu^+\mu^-$ spectrum using invariant masses of final-state lepton combinations as additional handles. Limits for $15 < m_{A'} < 55$ GeV, established using 20 fb^{-1} of 8 TeV data, were not competitive with EWP.

This work will focus on cases where the Higgs portal is also present, as in Hidden Abelian Higgs and various Hidden Valley models. Then the doubly-exotic channel gains the upper hand, providing a huge boost in sensitivity. This is because the $h \to 2A'$ or $h \to$ Dark States can proceed through just one Higgs mixing (κ) rather than two kinetic mixings (ϵ^2). Note that one kinetic mixing must still occur in $A' \to f^+f^-$, maintaining a handle on ϵ. Such cases can be further subdivided into:

- Heavy (Z_{D}) case: $m_{Z_{\mathrm{D}}} \gtrsim 10$ GeV
- Light (γ_{d}) case: $m_{\gamma_{\mathrm{d}}} \lesssim 10$ GeV

The important difference between these two cases is that at LHC, the γ_{d} come out highly boosted (due to their small mass), so that the final-state decay products f^+f^- are highly collimated. Meanwhile, the Z_{D} decay products come out with enough angular separation to be considered "non-collimated", i.e. they fall within the domain of the standard ATLAS reconstruction systems. In the Z_{D} case, the analysis in [39] placed bounds on the signal strength and on κ. The γ_{d} case, where f is restricted to electrons, muons, or light hadrons (which will collectively be referred to as l), provides the motivation for LJ searches and forms the basis for this work.

The above discussion illustrates how the A' search in ATLAS must be divided into smaller searches, by mass range. In some cases, the searches must be further sub-divided according to the A' lifetime, since decays in different parts of the detector produce markedly different experimental signatures. In all cases, if the A' is so long-lived as to be considered collider-stable, it can be detected only in missing transverse energy (MET) analyses. Even within the realm of LJ searches, a range of final-state topologies come into play. For example, in hidden sector and dark SUSY models, moderate to high α_d can lead to the emission of γ_{d}'s from other dark particles. Specifically, for a low-mass γ_{d} and a low-mass dark particle χ, emission of the γ_{d} receives a strong enhancement in the direction collinear to the χ. As discussed in [40], this “radiating DM” phenomenon is analogous to the emission of collinear photons in QED, or the emergence of parton showers in QCD. It leads to multiple γ_{d}’s emerging along the same line-of-flight, such that their decay products all form a single LJ. This is distinct from hidden sector and dark SUSY scenarios where multiple γ_{d}’s are produced from different parents in complicated multi-pronged decay chains; in such cases, each γ_{d} forms a distinct LJ well-separated from the others. Meanwhile, the width of an individual LJ can be affected not only by the number of γ_{d}’s that decayed to form it, but also by the boosts (and therefore the masses) of the γ_{d}’s , and by the details of interactions within the dark sector.

3.2 Scope and Role of LJ Searches

LJs as an exotic signature were initially introduced in the context of A' searches, and this is still the context in which the signature is most heavily probed. However, LJs also show up in the final-state predictions of other BSM models, such as right-handed neutrinos and protophobic/axial vector bosons. While these other models do not constitute the focus of this work, they are briefly discussed at the end of this section for the sake of completeness.

3.2.1 A' Searches

Figure 3.2 illustrates the placement of LJ searches in the overall scheme of ATLAS A' analyses. The LJ searches, divided into Prompt and Displaced, cover zero- to medium-lifetime γ_{d} and fall within the domain of the Unusual Exotic Signatures and Higgs Decays (UEH) Group. At $\sqrt{s} = 8$ TeV, a Displaced search was performed in [4], and a Prompt search in [41]. A Prompt search at $\sqrt{s} = 13$ TeV, using 2015-1016 data, is currently underway. The LJ searches can be considered the low-mass limit of non-collimated Z_{D} searches, which also fall within the UEH group domain. The zero- to short-lifetime Z_{D} cases (both singly- and doubly-exotic Higgs boson decays) were covered in Run 1 [39]; the medium-lifetime Z_{D} case (doubly-exotic Higgs boson decay to displaced muons only) was not covered in Run 1, and the Run 2 analysis is still underway. For completeness, it should be mentioned that CMS performed a search for Prompt LJs composed only of muons at 8 TeV [42].

Collider-stable lifetimes and Drell-Yan searches both fall outside the UEH Group domain, as ATLAS has large separate groups for MET and dilepton resonance searches, respectively. The main reason for this separation is that UEH is essentially a collection of specialized model-driven searches. In contrast, MET searches encompass a huge diversity of BSM models, including all “mono-X” scenarios in which particle(s) that escape ATLAS undetected recoil against some SM object X (jet, photon, weak gauge boson, Higgs boson, etc.) The MET searches hence have their own dedicated groups, which are driven by generic signatures rather than by specific models. Similarly, the dilepton resonance searches encompass all models that include a new BSM particle with a e^+e^- decay channel. They hence have their own dedicated group, with specialized “bump-hunting” statistical frameworks to analyze the e^+e^- invariant mass distributions. If a signal is found by a MET or dilepton resonance search, the channel(s) in question would be compatible with a host of other BSM models besides dark photons. The signal, while a discovery of BSM physics, would be at best a hint of dark photons in particular.

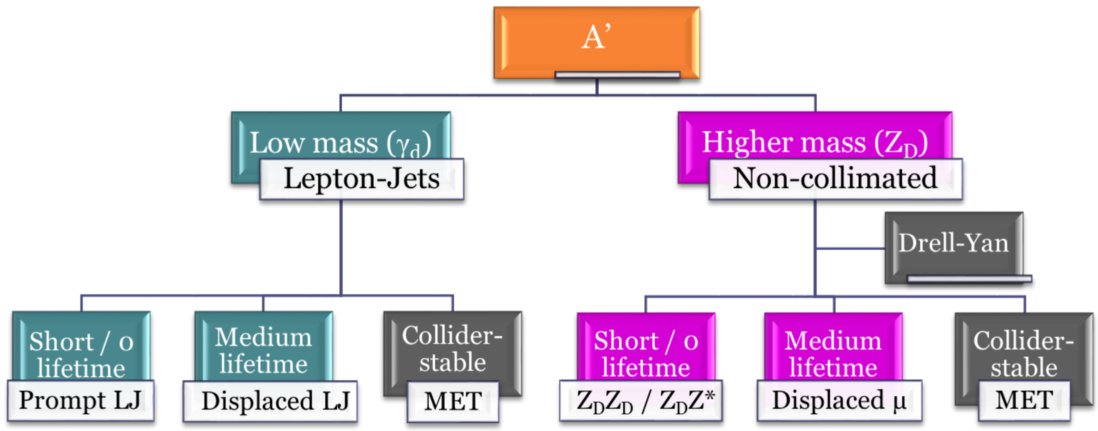

Figure 3.2: The overall scheme of ATLAS A' analyses. Green: LJ searches. Violet: Other UEH Group A' searches. Gray: Searches outside the UEH Group.

3.2.2 Other Exotic Searches

Recently, discussion of models of MeV-scale bosons different from the A' was spurred by an anomaly observed in the M1 transition from an 18.15 MeV isoscalar magnetic dipole state to the ground state ($J^\pi = 1^+,\ T = 0 \to J^\pi = 0^+,\ T = 0$) in ^{8}Be. Specifically, the authors of [43] measured the angular correlation (opening angle) of the electron and positron emitted in the transition. In transitions by standard internal pair creation (IPC), QED predicts a rapid and smoothly falling distribution of opening angles. Significant deviations were seen from the IPC predictions at large opening angles, with up to 6.8σ significance. The authors of [43] noted that if the transition takes place by emission of a short-lived neutral particle decaying into an electron-positron pair (rather than by standard IPC), this would create a sharp peak at such large opening angles. The angular correlations resulting from a hypothetical boson X were simulated for various m_X values, and a χ^2 analysis performed on the data to extract the best-fit m_X. The result of the analysis is $m_X c^2 = 16.70 \pm 0.35$ (stat) MeV. Interestingly, for a closely-related transition from an 17.6 MeV isovector state ($J^\pi = 0^+,\ T = 0$) to the ground state, the authors of [43] observed no significant deviations from IPC predictions in the angular correlations. To accommodate this observation while evading all other existing experimental constraints, other authors showed that an X with purely vector couplings to quarks would have to be protophobic (coupling much more weakly to protons than to neutrons). An X with primarily axial couplings to quarks, however, need not be protophobic [44]. Neither a protophobic vector boson nor an axial vector boson is exactly the A'. UV-complete models with a protophobic vector boson are difficult to construct. In the axial case, however, it is possible to construct models similar to the hypercharge portal and Higgs portal models of Chapter 2. An example is a $U(1)_{RH}$ gauge theory with couplings to right-handed fermions, kinetic mixing between hypercharge and the $U(1)_{RH}$, and two Higgs doublets charged under the $U(1)_{RH}$ [44]. The salient point is that in a detector such as ATLAS, an A' with mass m_X decaying to electron pairs may mimic the LJ signature of the X decaying to electron pairs.

A Displaced LJ signature may also arise in the neutrino Minimum Standard Model, where three right-handed neutrinos ($N_1,\ N_2,\ N_3$) are added to the SM. The lightest of these, N_1, would constitute stable

DM, while the N_2 and N_3 would have masses $\sim 0.2 - 100$ GeV. In order to evade current experimental bounds, the mixings between the right-handed and SM neutrinos would have to be very small, leading to long lifetimes for the right-handed neutrinos. Processes involving N production from W bosons, followed by N decay to SM products, have particularly promising LHC discovery potential [45]. The N may be sufficiently boosted to produce a Displaced LJ if it decays to leptons (Figure 3.3).

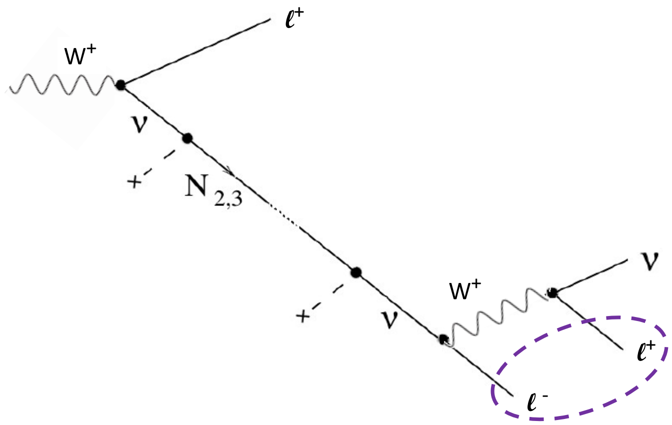

Figure 3.3: The production of a right-handed neutrino N_i, and its decay to SM leptons. The N may be sufficiently boosted to produce a LJ, illustrated in purple.

3.3 Parameter Space

Dark photon parameter space in the literature is often expressed in the two-dimensional plane (ϵ, $m_{A'}$), where only the presence of the hypercharge portal is assumed. Observations from several different types of experiments have been employed to establish limits in various regions of this parameter space. This work will focus on the range 10 MeV $\lesssim m_{\gamma_{\rm d}} \lesssim$ 10 GeV, as outlined in the previous section. As shown in Figure 3.4, the higher ϵ range ($\gtrsim 10^{-3}$) has been dominated by B-factories and by electron and muon magnetic moment measurements, while the low ϵ ($\lesssim 10^{-7}$) range is accessible only through supernova-cooling observations. The middle range ($10^{-3} \lesssim \epsilon \lesssim 10^{-7}$) has been dominated by fixed-target beam-dump experiments. However, these beam-dumps are sensitive only at the lower end of the $m_{A'}$ range, leaving a gap at higher $m_{A'}$ waiting to be explored by colliders. A third dimension in parameter space must then be added, to account for the Higgs mixing that plays such an important role in the collider search sensitivity.

A brief summary of the existing experimental constraints shown in Figure 3.4 is as follows:

- a_μ, a_e: The A' would contribute to the anomalous magnetic moment of the electron and of the muon, through a one-loop diagram [46]. This contribution is proportional to ϵ^2 for a wide range of $m_{A'}$, ruling out $\epsilon \gtrsim 10^{-3}$ for $m_{A'} \lesssim 0.1$ GeV. Currently, there is a discrepancy of $\sim 3.4\sigma$ between Brookhaven precision measurements of a_μ and the theoretical prediction, which could actually be explained by the A' in a narrow band of parameter space with ϵ just smaller than the excluded range [47]. However, this favoured range is largely disfavoured by other experiments.

- APEX (APrime Experiment): Dedicated fixed-target beam dump experiment at Jefferson Laboratory, with a 1.131 and then a 2.262 GeV electron beam. Designed to detect e^+e^- pairs produced from A' decay, where the A' would be radiated from electrons scattering off the target nuclei (bremsstrahlung). This requires searching for a narrow resonance in the e^+e^- invariant mass spectrum. [48]

- MAMI (A1 Collaboration at Mainz Microtron): A similar experimental design to APEX, but with an 885 MeV electron beam. [49]

- Non-dedicated beam dump experiments (Orsay, E774, E141, E137, U70): High-intensity, tightly-collimated pulsed electron beams incident on metal targets, with calorimeters and/or spectrometers located tens of metres to a few hundred metres downstream. [50–54]

- Long-baseline neutrino experiments (LSND, CHARM): Essentially employed as beam dump experiments in this context. [55, 56]

- B-factories / e^+e^- colliders (BaBar): The A' could be produced at e^+e^- colliders via $e^+e^- \to A'\gamma$, $e^+e^- \to A'\ s_{\rm d}$ where $s_{\rm d}$ is a dark scalar, or $e^+e^- \to V \to PA'$ where V is a vector meson and P is a pseudoscalar meson. In the latter case, the pseudoscalar could be easily tagged by one of its main decays, while the A' could be detected if it decays into a pair of SM leptons. [57, 58]

- Kaon experiments (NA48/2 at CERN Super Proton Synchrotron, KLOE): Searches for $\pi^0 \to \gamma A'$ followed by prompt $A' \to e^-e^+$, amongst samples of fully-reconstructed $\pi^0 \to \gamma e^-e^+$ decay candidates. [59–61]

- Nuclear physics (HADES): Searches for $A' \to e^-e^+$ in inclusive dielectron spectra measured in p(3.5 GeV) + p, Nb reactions and in Ar (1.756 GeV/u) + KCl reactions. [62]

- Supernova SN1987A neutrino observations: A' production would contribute to energy loss from the hot core of a collapsing supernova, if the A' is sufficiently long-lived to travel a few km before decaying. [63, 64]

While a sub-MeV (and even massless) A' is theoretically possible, the MeV-to-GeV range for $m_{A'}$ is generally favoured by DM models, and by naturalness considerations in SUSY models. In some string models, $m_{A'}$ and ϵ are typically both proportional to the string scale, and are therefore related to each other. Outside of these models, however, $m_{A'}$ and ϵ are essentially independent parameters, and theory gives rather little guidance regarding the likely order-of-magnitude of ϵ. This is why experiments must "look everywhere" along this direction of parameter space. Theoretically there is no clear minimum for ϵ, while the maximum could be all the way up at $\mathcal{O}(1)$ since the dimension-four hypercharge portal operator is unsuppressed by any mass scale. Even if the value of ϵ is assumed to go to 0 at very high energies, ϵ at lower energies can be generated by perturbative or non-perturbative effects, and loops of various orders – involving heavy messengers charged under both $U(1)_Y$ and $U(1)'$ – can drive it to anywhere from 10^{-2} to 10^{-8} [65]. However, the range $10^{-3} \lesssim \epsilon \lesssim 10^{-4}$ is particularly interesting from a "naturalness" point of view in SUSY models, as it arises from integrating out heavy fields charged under both the SM and dark sector. When supersymmetric kinetic mixing is taken into account, this ϵ also naturally generates the GeV scale in the dark sector, which is promising for WIMPs [26]. Meanwhile, certain string models can drive ϵ lower by several orders of magnitude. The range $\epsilon \lesssim 10^{-7}$ has the advantage of easily allowing WIMPs to escape direct detection bounds, without the need for iDM or XDM [65].

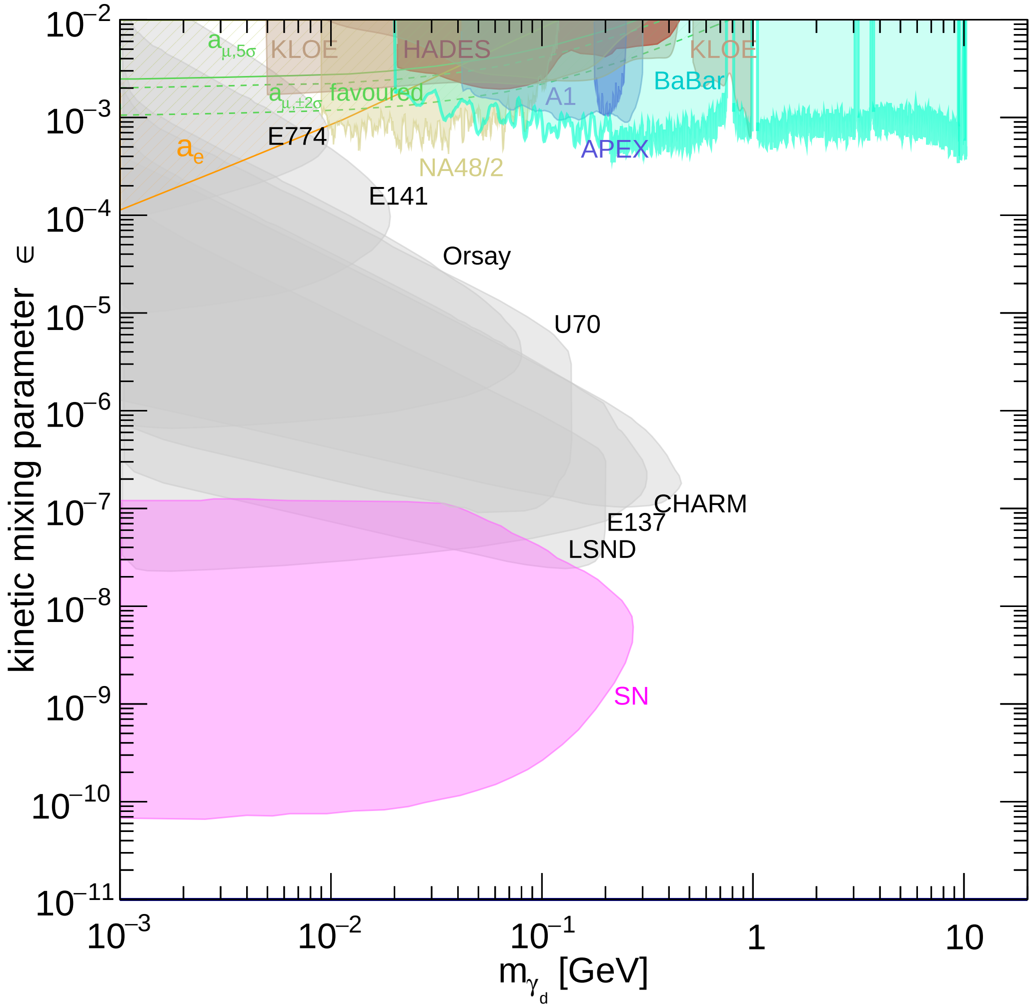

Figure 3.4: Existing constraints in A' parameter space, zoomed-in to the area of interest in this work: B-factories (teal), beam-dump and long-baseline neutrino experiments (gray), kaon experiments (light brown), nuclear physics experiments (dark brown), supernova observations (violet), electron magnetic moment (orange dashed line), muon magnetic moment (green dashed lines), dedicated A'-search fixed-target experiments (blue). The empty region to the right of the beam-dump experiments is the target of LJ searches in this work.

3.4 Branching Ratios and Lifetimes

The partial width, to lowest order in κ, for the doubly-exotic Higgs boson decay proceeding via Higgs mixing is [9]:

$$\Gamma(h \to 2A') = \frac{{\kappa'}^2 v^2}{32\pi m_h^5}\sqrt{1-\frac{4m_{A'}^2}{m_h^2}}\left[(m_h^2+2m_{A'}^2)^2-8m_{A'}^2(m_h^2-m_{A'}^2)\right] \tag{3.1}$$

with $\kappa' = \frac{\kappa m_h^2}{|m_h^2 - m_s^2|}$.

The $A'l^+l^-$ coupling is [9]:

$$g_{A'll} = \frac{g}{\cos\theta}\left[-\sin\alpha(t_3\cos^2\theta - Y\sin^2\theta) + Y\eta\cos\alpha\sin\theta\right] \tag{3.2}$$

where t_3 is the weak isospin of the fermion and Y is its hypercharge. To lowest order, this yields decay widths [9]

$$\Gamma(A' \to l^+l^-) = \frac{1}{24\pi m_{A'}}\sqrt{1-\frac{4m_l^2}{m_{A'}^2}}\left[m_{A'}^2(g_L^2+g_R^2) - m_l^2(-6g_Lg_R+g_L^2+g_R^2)\right] \tag{3.3}$$

where L, R refer to left- and right-handed components of the fermion. This tree-level approximation must then be corrected for threshold effects, hadronic resonances, and higher-order QCD. The resulting branching ratios for the A' are shown in Figure 3.5. Combining these with Equation 3.1 yields the BR($h \to 2A' \to 4l$) shown in Figure 3.6. Meanwhile, the resulting proper lifetime τ for the A' is shown in Figure 3.7; the salient point is that the lifetime depends upon both mass and ϵ. The following formula, used throughout the Displaced LJs analysis, gives a sufficient approximation:

$$\gamma c\tau \simeq \left(\frac{10^{-4}}{\epsilon}\right)^2 \left(\frac{100\mathrm{MeV}}{m_{A'}}\right)^2 \mathrm{mm} \tag{3.4}$$

Exploring longer lifetimes provides access to lower ϵ, thereby motivating searches for Displaced LJs. The lower limit on the ϵ range that can be probed by displaced LJ searches is imposed primarily by the lifetime at which the A' becomes effectively collider-stable, i.e. decays outside the detector volume. In ATLAS, it is convenient to define collider stability as $\gamma c\tau \gtrsim 10$ m. For γ $\mathcal{O}(10-100)$, which is typical in LHC events, this translates by Equation 3.4 to a lower limit on ϵ of $\mathcal{O}$ (10^{-7}) for $m_{A'} \sim 1$ GeV. Note that the value of ϵ does not affect the BRs, only the A' lifetime.

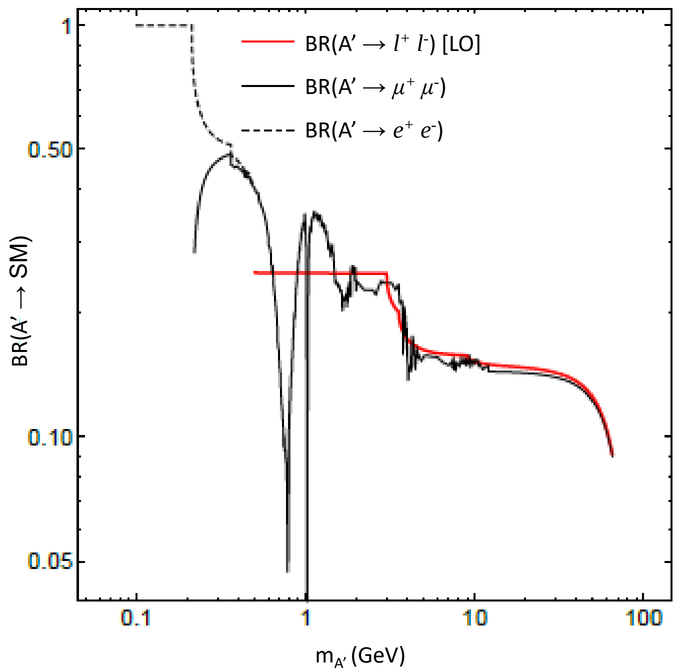

Figure 3.5: Branching ratios of the A' as a function of its mass in the simplest hypercharge portal model, from [9]. Red: leading-order estimate of the leptonic BR of the A'. Black lines: estimates of the electron and muon BRs, incorporating QCD corrections up to 3 loops.

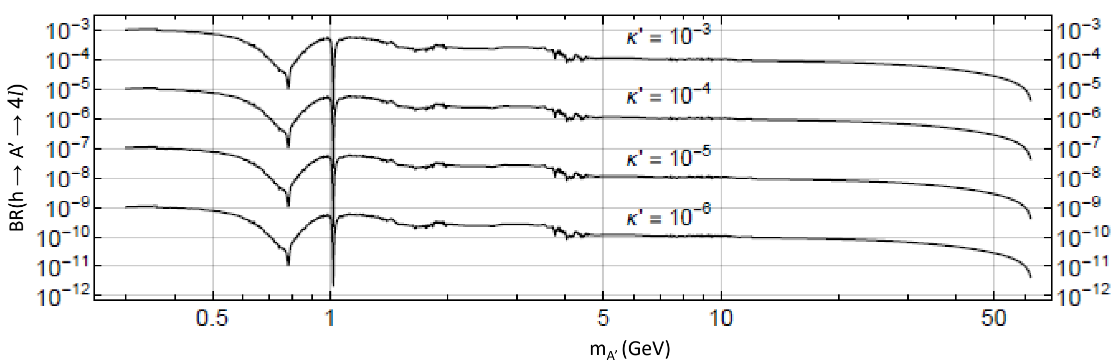

Figure 3.6: BR($h \to 2A' \to 4l$) as a function of the A' mass and of κ' in the simplest Hidden Abelian Higgs model, from [9].

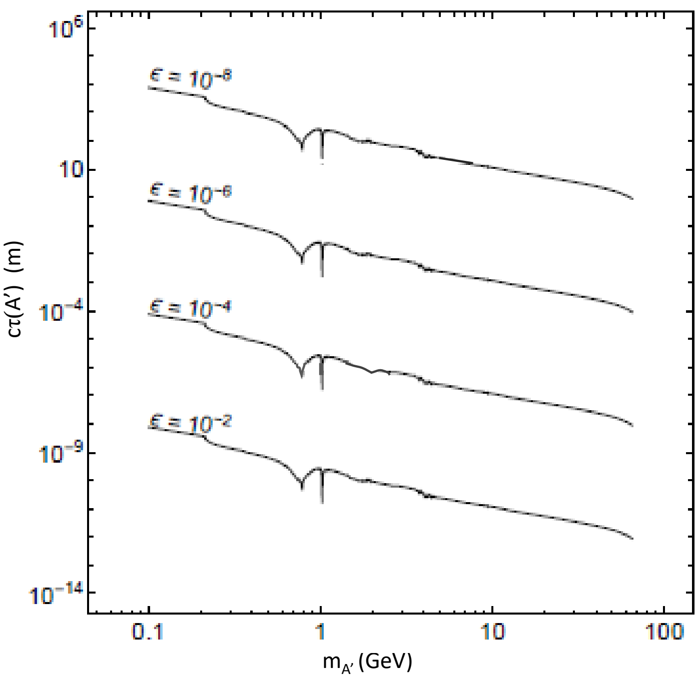

Figure 3.7: Proper lifetime of the A' as a function of its mass and of ϵ in the simplest hypercharge portal model, from [9].

3.5 Benchmark Models

There has been lively debate in the experimental and phenomenology communities, regarding benchmark model-based vs model-independent analysis approaches. The first approach aims to establish limits that are as stringent as possible for a specific model; analysis cuts are hence finely-tuned for the kinematic features of that particular model. Advanced optimization and multivariate analysis tools, such as boosted decision trees, are often employed, leading to cuts that are based on highly effective but opaque black-box discriminants. While this can improve the limits by several percent, it makes it difficult to "re-cast" the results, i.e. to use them for placing limits on other models with similar but non-identical experimental signatures. The latter approach aims for ease of re-casting, at the expense of limit stringency on any one particular model. It promotes the use of only basic features of the experimental signature, along with the production of extensive acceptance and efficiency tables as a function of a few key kinematic variables. The Displaced LJs analysis employs aspects of both approaches, with a fairly model-independent stage followed by limit-setting on benchmark models that are chosen to be representative of a broad class of models.

3.5.1 Simplified LJs

In a LJ, the e, μ, π multiplicity depends on the number of $\gamma_{\rm d}$'s that merge into a single lepton jet, and on the mass of the $\gamma_{\rm d}$. Other properties of the LJs strongly depend on the unknown structure of the hidden sector and its couplings to the visible sector. For instance, the LJ shape is controlled, in part, by the typical boost of the hidden particles, which in turn is determined by the mass ratio of the decaying hidden sector particle and the produced visible sector one.

The model-independent stages of the LJ analysis adopt a simplified approach, with a generic definition of LJs containing one or two $\gamma_{\rm d}$'s (see Figure 3.8). In the latter case, the two $\gamma_{\rm d}$'s are assumed to come from a parent dark scalar $s_{\rm d}$.

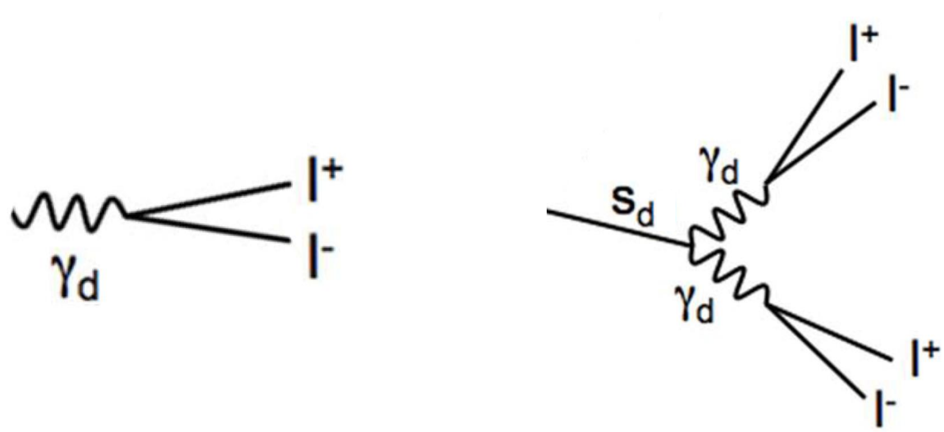

Figure 3.8: Left: LJ produced by the decay of one $\gamma_{\rm d}$. Right: LJ produced by the decay of one hidden scalar $s_{\rm d}$ into two $\gamma_{\rm d}$'s.

3.5.2 FRVZ Models

For limit-setting purposes, the two benchmark models used in this analysis are Falkowsky-Ruderman-Volansky-Zupan (FRVZ) models [14, 66] where a Higgs boson H decays to pairs of hidden fermions. The Higgs boson may be either SM-like (125 GeV), or a heavier BSM variety. In the first benchmark model (Figure 3.9, left) the hidden fermion decays to a γ_d and a lighter hidden fermion, the HLSP (Hidden Lightest Stable Particle). In the second model (Figure 3.9, right) the hidden fermion f_{d_2} decays to an HLSP and a hidden scalar s_d that in turn decays to pairs of γ_d's. For brevity, these will be referenced as the $2\gamma_\mathrm{d}$ model and the $4\gamma_\mathrm{d}$ model, respectively. The benchmark models are simplified by imposing the restriction that only muon/electron/pion pairs are produced from the γ_d's, although other light hadrons and/or tau leptons could theoretically be produced at sufficiently high γ_d masses. In general, hidden sector radiation can produce extra γ_d's , but the number produced is model-dependent because the number of radiated γ_d's is proportional to the size of the dark gauge coupling α_d. Hidden radiation is not included in the FRVZ benchmark models; this corresponds to assuming $\alpha_d \lesssim 0.01$.

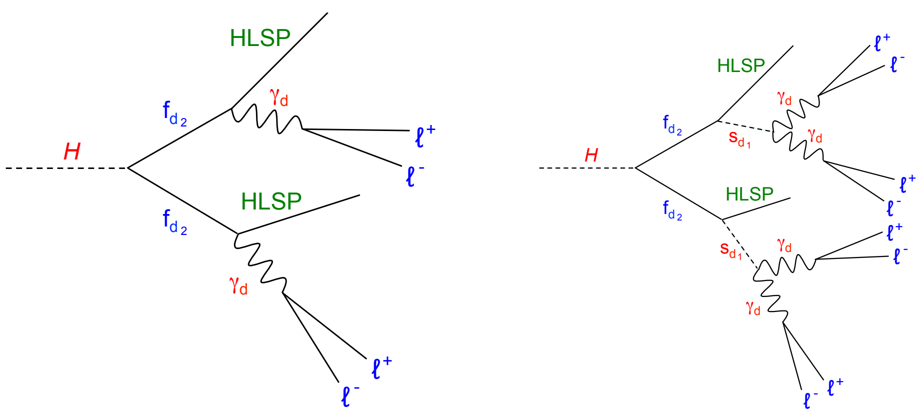

Figure 3.9: The two FRVZ models used as benchmarks in the analysis. Left: the $2\gamma_\mathrm{d}$ model. Right: the $4\gamma_\mathrm{d}$ model. Reproduced as in [6].

These models provide several benefits as benchmarks:

- Sensitivity to both κ and ϵ, since $H \to f_{d_2}\ f_{d_2}$ proceeds via the Higgs portal while the γ_d decay proceeds via the hypercharge portal.
- Flexibility in the amount of $E_\mathrm{T}^\mathrm{miss}$ produced, depending upon the kinematics of the HLSPs (controlled by the dark sector mass spectrum). In the limit of no energy carried away by the HLSPs, the $2\gamma_\mathrm{d}$ model has essentially the same experimental signature as the doubly-exotic Higgs boson decay in Figure 3.1.
- At least two LJs in the final state.
- Not explicitly dependent upon SUSY, but can be interpreted as a simplified version of various SUSY models.

These models were used as benchmarks in the Run 1 analysis, assuming the Higgs boson to be SM-like (125 GeV). The Run 2 analysis adds versions of the benchmarks with a heavy BSM Higgs boson (800 GeV), for which the narrow-width approximation is employed. In all cases, the Higgs boson is assumed to be produced via gluon-gluon fusion, as this is the leading SM Higgs boson production process at both Run 1 and Run 2 energies (Figure 3.10). The $\sqrt{s}$ increase from 8 TeV in Run 1 to 13 TeV in Run 2 boosts the gg production cross-section of the SM Higgs boson from 18.5 to 43.9 pb ([67]), helping the analysis sensitivity. The gg production cross-section for the heavy BSM Higgs boson is left as an unknown in the analysis interpretation, as it depends upon the couplings of the BSM Higgs boson which could vary greatly from those of a SM-like Higgs.

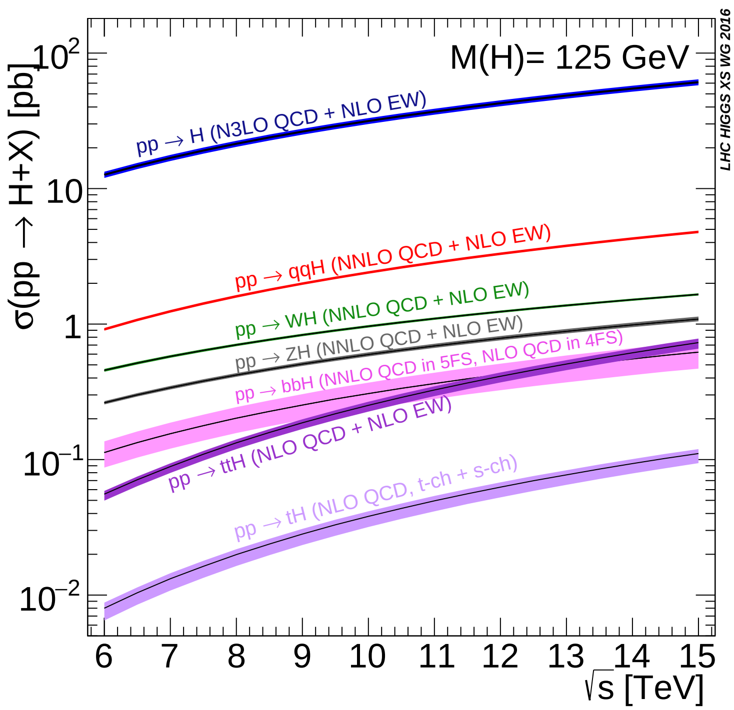

Figure 3.10: SM cross-sections for the 125 GeV Higgs boson production processes, as a function of $\sqrt{s}$. [67]

3.6 Summary

In the presence of both the Higgs Portal and the Hypercharge Portal, searches for A's in decay chains beginning with h are sensitive to both ϵ and κ. These searches in ATLAS are divided up according to mass; the Displaced LJs analysis focuses on the production of two or more fairly low-mass A' (denoted $\gamma_{\rm d}$), i.e. $0.2 \lesssim m_{\gamma_{\rm d}} \lesssim 10$ GeV. It further focuses on the case $10^{-3} \lesssim \epsilon \lesssim 10^{-7}$, which leads to $\gamma_{\rm d}$ lifetime long enough to yield displaced decay products, but not long enough to usually escape the detector. Accurate parameterizations for the lifetimes and branching ratios of the $\gamma_{\rm d}$ to electrons, muons, and light hadrons are available in this regime. While a wide range of other experiments – from beam-dumps to B-factories to supernova observations – have already established limits in $(m_{A'},\ \epsilon)$ parameter space, they do not reach the region probed by the Displaced LJs analysis. The LJ limits do depend upon an extra parameter, namely κ, that the other experimental limits do not. The LJ shapes, and number and species of constituents, are highly model-dependent, necessitating a simplified approach with a generic definition of LJs. As in the 8 TeV version of the Displaced LJs analysis, two exotic Higgs boson decay "FRVZ" processes, one yielding $2\gamma_{\rm d}$ and the other yielding $4\gamma_{\rm d}$, are then used as benchmarks for limit-setting. The 13 TeV analysis documented in this work benefits from the higher Higgs boson production cross-section, and explores the decay of a heavy BSM Higgs boson as well as the SM Higgs boson.

Chapter 4

The ATLAS Detector

ATLAS is a multi-purpose detector [2], primarily studying *pp* collisions, located at one of four interaction points along the 26.7 km LHC tunnel. The main ATLAS sub-detectors (Figure 4.1) consist of an inner tracking system (ID) surrounded by electromagnetic and hadronic calorimeters (ECal and HCal) which are in turn surrounded by a muon spectrometer (MS). As an understanding of the sub-detector volumes is important for the LJ analysis, the R extent of each main sub-detector is given in Table 4.1.

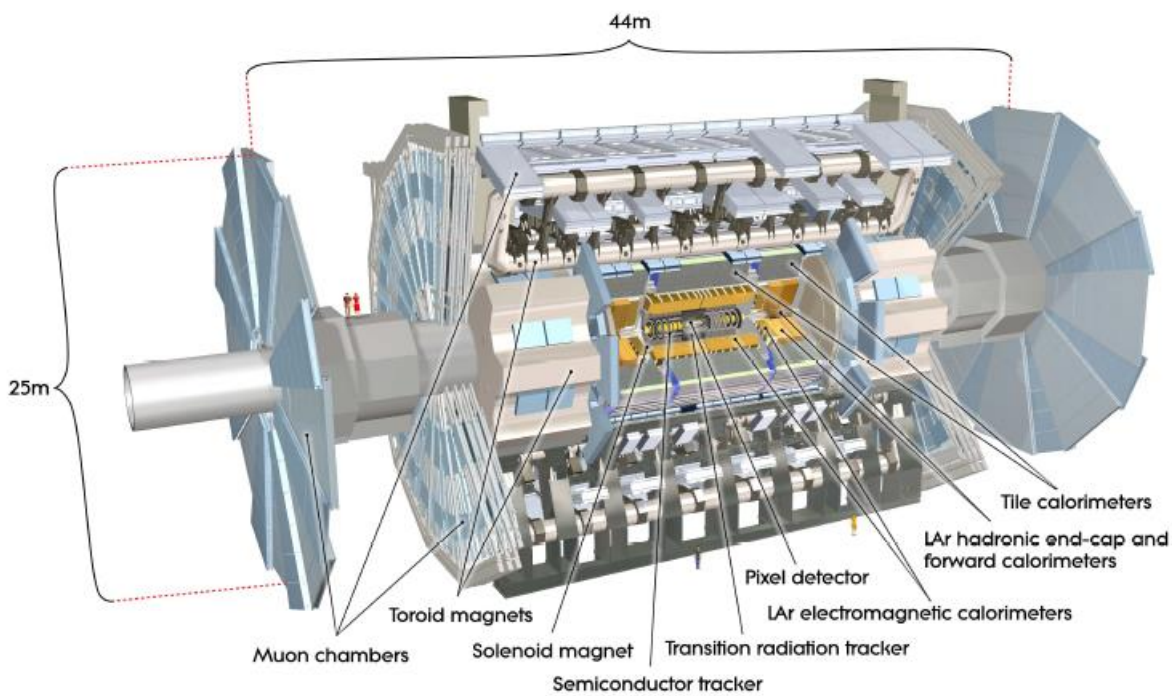

Figure 4.1: Overall layout of the ATLAS detector. [2]

ATLAS employs a hybrid system of superconducting magnets [68], consisting of one solenoid, one barrel toroid, and two end-cap toroids. The solenoid, aligned on the beam axis, provides a 2 T axial magnetic field for bending charged tracks in the ID. The toroids provide fields of approximately 0.5 T with a bending integral of $\simeq$ 2.5 Tm in the barrel, and 1 T with a bending integral $\simeq$ 6 Tm in the

Sub-detector	R Extent (m)
Inner Detector	0.03 - 1.40
Electromagnetic Calorimeter	1.40 - 2.28
Hadronic Calorimeter	2.28 - 4.25
Muon Spectrometer	4.35 - 10.57

Table 4.1: The R extent of each main sub-detector.

end-caps, for bending muon tracks in the MS.

For the pp running conditions in the LHC during 2015 data-taking, the key parameters relevant to ATLAS analyses are summarized in Table 4.2.

Parameter	Design Value	Typical 2015 Value
Beam energy (TeV)	7.0	6.5
Average protons/bunch (10^{11})	1.15	1.1
Bunch spacing (ns)	25	25
Maximum peak luminosity (10^{34} cm^{-2} s^{-1})	1.0	0.51
Total integrated luminosity (fb^{-1})	-	4.2
ATLAS recorded luminosity (fb^{-1})	-	3.9

Table 4.2: Key parameters describing the LHC running conditions for 2015 pp data-taking, compared to their design values.

4.1 Coordinate System

Throughout this work, the standard ATLAS coordinate system will be employed, with the origin at the interaction point (IP):

- The z-axis is aligned with the beam pipe.
- In Cartesian coordinates, the x-axis points from the IP to the centre of the LHC ring, with the y-axis pointing upward (right-handed).
- In cylindrical coordinates, (R, ϕ) describes the transverse plane, with ϕ the azimuthal angle around the beam pipe.
- In polar coordinates, the polar angle is θ.
- Transverse momentum is defined as $p_{\mathrm{T}} = p \sin\theta$, and transverse energy as $E_{\mathrm{T}} = E \sin\theta$.
- Pseudorapidity is defined as $\eta = -\ln\tan(\theta/2) = \tanh^{-1}(p_z/p)$.
- Opening angle between two objects is defined as $\Delta R = \sqrt{(\Delta\eta)^2 + (\Delta\phi)^2}$.

The use of η is often preferred over θ because differences in η are Lorentz-invariant under z boosts, in the limit of massless particles. For massive particles, it is rapidity $y = \frac{1}{2}\ln\frac{E+p_z}{E-p_z}$ rather than η that is Lorentz-invariant under z boosts, but in high-energy collisions, $m \ll p$ and so η converges to y. A $\Delta\eta$ measurement is therefore independent of the z boosts of the rest frames of the particles involved. Since the particles involved are partons, and each parton carries some unknown fraction of the p_z of the proton to which it belongs, the partons' rest frames generally have different z boosts. There is the additional advantage that the η distributions of final-state particles are approximately uniform in a variety of simple collision processes.

Particle tracks in ATLAS are often described in terms of their transverse impact parameter (d_0) and longitudinal impact parameter (z_0). The impact parameter represents a distance of the point of closest approach of the track to a reference (see Figure 4.2). For ID tracks, this reference is always a line parallel to the z axis and passing through the transverse beamspot position. For MS tracks, this reference is event-dependent: specifically, a line parallel to the z axis and passing through the primary vertex of the event. Tracks are also often described by their p/q, a signed quantity representing momentum over charge.

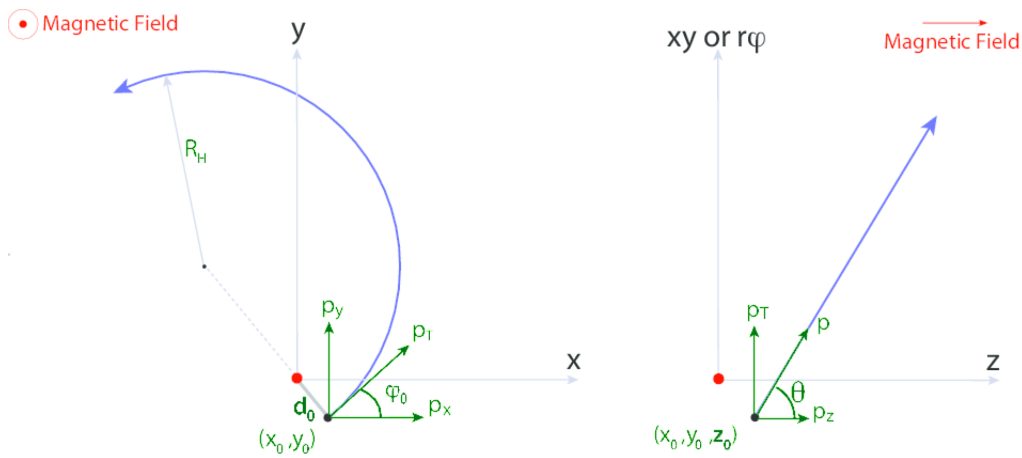

Figure 4.2: Impact parameters, transverse (d_0) and longitudinal (z_0), where the red dot represents the reference. [69]

4.2 Inner Detector and Tracking

The ID is composed of three sub-systems, as shown in Figure 4.3: Pixel, SemiConductor Tracker (SCT), and Transition Radiation Tracker (TRT). The Pixel and SCT constitute the precision tracking detectors, covering $|\eta| < 2.5$; the Pixels achieve the highest granularity. The TRT provides coarser track-following further from the IP, up to $|\eta| = 2.0$.

The three concentric Pixel barrel layers [70] present in Run 1 have been supplemented by a fourth innermost layer, the Insertable B-Layer [71], for Run 2. The three Pixel layers in each end-cap are arranged as disks perpendicular to the beam axis. Each layer is composed of silicon sensors, where each pixel is 50×400 μm^2 (50×250 μm^2 in the Insertable B-Layer) in $R - \phi \times z$, with intrinsic accuracy of 10 μm ($R - \phi$), 115 μm (z). The standard ATLAS tracking algorithms require any reconstructed ID track to have at least one hit in the Pixel sub-system.

The SCT is composed of 8 strip layers [72], with pitch ~ 0.08 mm. The layers are arranged in pairs, where each pair measures both $R - \phi$ and z by combining the measurements from its two members. The barrel layers are composed of small-angle (40 mrad) stereo strips, where one member of each pair runs azimuthally. The intrinsic accuracies per module are 10 μm ($R - \phi$) and 115 μm (z). In the end-cap, there are 4 layers of strips running radially, each paired with a layer of stereo strips at an angle of 40 mrad. The intrinsic accuracies per module are 17 μm ($R-\phi$) and 580 μm (z).

The TRT is composed of straw tubes [73], each 4 mm in diameter. The barrel straws are 144 cm long, arranged parallel to the beam axis, while the end-cap straws are 37 cm long, arranged radially. Unlike the Pixels and SCT, the TRT provides only $R-\phi$ information; the intrinsic accuracy per straw is 130 μm. A track originating from the IP typically generates 36 TRT hits.

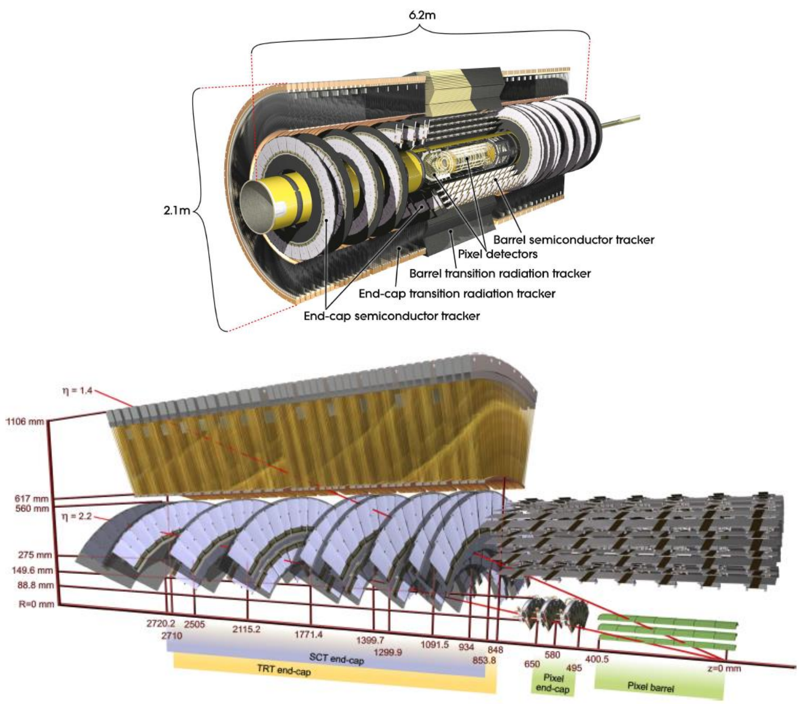

Figure 4.3: Top: Overall layout of the inner detector. Bottom: Path traversed by two charged tracks (red lines). [2]

4.3 Calorimeters

The layout of the calorimeters is shown in Figure 4.4. The calorimeter granularity and depth segmentation vary with $|\eta|$, depending upon the measurement requirements of the physics processes of interest and upon the radiation environment.

4.3.1 Electromagnetic Calorimeter

The electromagnetic calorimeter (ECal) is a liquid argon (LAr) detector [74] with accordion-shaped lead absorber plates and kapton electrodes. It is composed of a barrel component ($|\eta| < 1.475$) and two end-cap components ($1.375 < |\eta| < 3.2$). The barrel is made from two identical half-barrels with a small gap between them at $z = 0$, and is segmented in depth into three layers for $|\eta| < 1.35$ and two layers in the remainder. Each end-cap is made from two coaxial wheels, one covering $1.375 < |\eta| < 2.5$ and the other covering the region $2.5 < |\eta| < 3.2$. They are segmented in depth into three layers for $1.5 < |\eta| < 2.5$ and two layers in the remainder. In the region of $|\eta| < 1.8$, an additional LAr layer serves

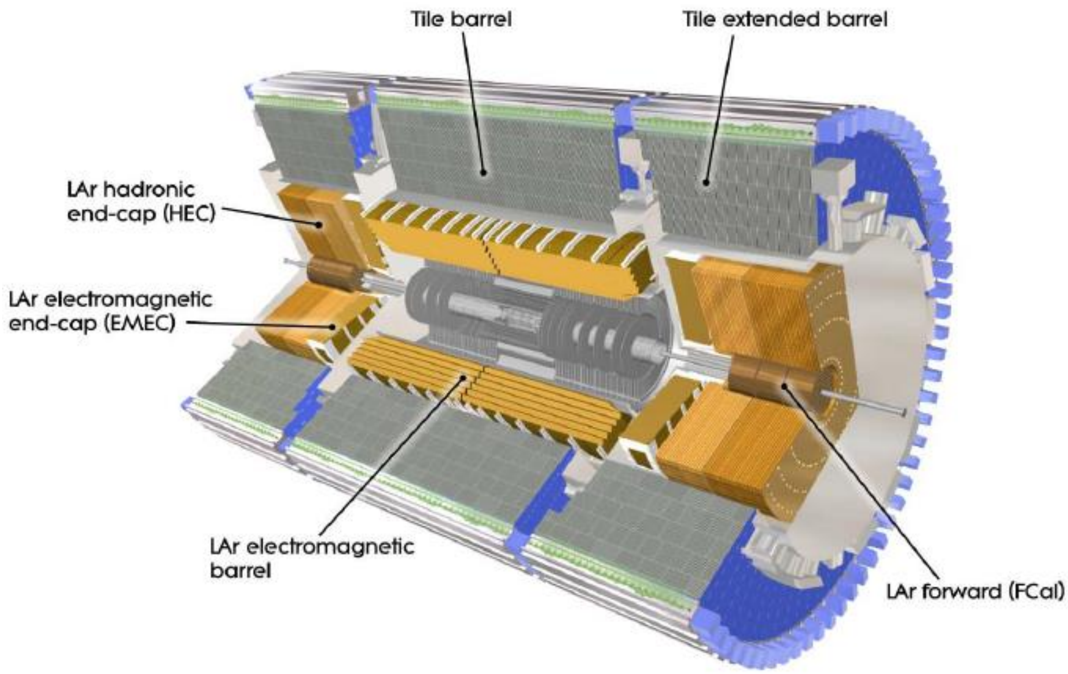

Figure 4.4: Layout of the calorimeters. [2]

as a presampler to correct for the energy lost by electrons and photons upstream of the calorimeter. The $|\eta|$ granularity varies from 0.00325 to 0.1, and the ϕ granularity from 0.025 to 0.1. The finest granularity is found nearest the IP, to accomplish precision measurements of electrons and photons. The ECal is > 22 radiation lengths thick in the barrel, and > 24 in the end-caps.

In addition, the inner module of the Forward Calorimeter (FCal) [75] is optimized for EM measurements. It is also a LAr module, but with copper absorbers. It consists of a metal matrix, with longitudinal channels containing concentric rods and tubes that serve as the electrodes. The LAr is contained in the very small gaps between rod and tube.

4.3.2 Hadronic Calorimeter

The tile calorimeter is a sampling calorimeter [76] with scintillating tiles as the active material and steel as the absorber. Its barrel and two extended barrels are divided by depth into 3 layers and azimuthally into modules, with the barrel covering $|\eta| < 1.0$ and the extended barrels $0.8 < |\eta| < 1.7$. The $\eta \times \phi$ granularity is 0.1×0.1, except in the outer layer where it is 0.2×0.1. Its layer thicknesses are 1.5, 4.1 and 1.8 interaction lengths in the barrel, and 1.5, 2.6 and 3.3 in the extended barrel.

The hadronic end-cap calorimeter (HEC) [74] uses LAr as the sampling material, held in gaps between copper plate absorbers. It is composed of two wheels per end-cap, the inner at $1.5 < |\eta| < 2.5$ and the outer at $2.5 < |\eta| < 3.2$, each divided by depth into 2 layers and azimuthally into 32 wedge-shaped modules. The inner layers have granularity 0.1×0.1 in $\eta \times \phi$, while the outer layers have granularity 0.2×0.2.

There is a transition region or "gap" between the barrel and end-cap (see Figure 4.5), filled with

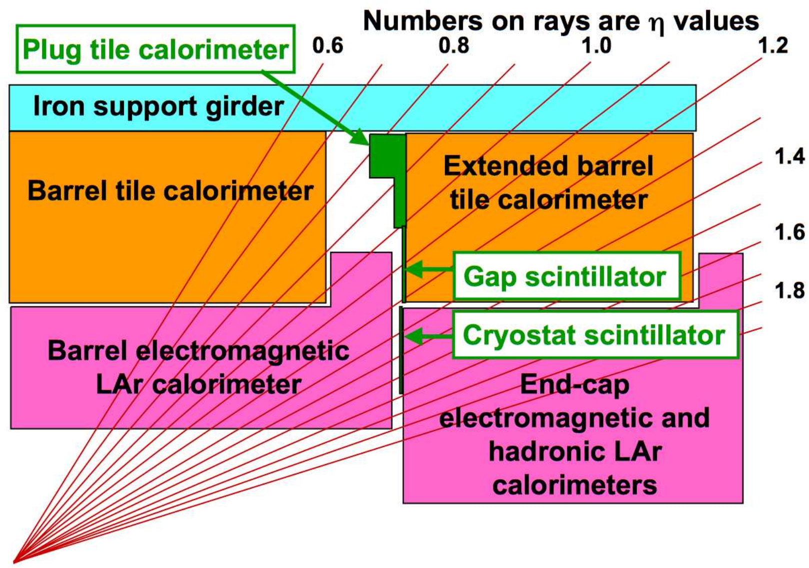

Figure 4.5: Schematic of the transition region between the barrel and end-cap. [2]

cables and services for the ID as well as power supplies and services for the ECal barrel, where part of the structure has been removed from the tile calorimeter. Part of this region is covered by an extra tile calorimeter module, the "plug tile calorimeter". In the region $1.0 < |\eta| < 1.2$, two "tile gap scintillators" with independent readout are attached to the tile module surface, providing signals which can be used to correct for energy losses in the inactive material in the gap. In addition, in the region $1.2 < |\eta| < 1.6$, two "cryostat scintillators" penetrate the region between the outer walls of the barrel and end-cap cryostats, providing signals which can be used to correct for energy losses in the outer wall of the barrel cryostat and in the ID services. The tile gap and cryostat scintillators will play a role in jet selection in the Displaced LJs analysis (see Section 5.5.6).

The two outer modules of the FCal [75] are optimized for hadronic measurements. They are similar in structure to the inner module, but with tungsten absorbers. Specifically, while the tubes in the electrodes are still copper, the rods are tungsten, and the matrix filling the space between electrodes is composed of tungsten slugs. These modules are optimized for a high absorption length; the tungsten was chosen for its ability to provide containment, and minimise the lateral spread, of hadronic showers.

4.3.3 Topological Calorimeter Clusters

Calorimeter towers can be used directly for jet reconstruction. However, the calorimeter jets employed in the Displaced LJ analysis are instead built from topological calorimeter clusters. These "topo-clusters" are groups of calorimeter cells that are intended to follow the shower development [77], taking full of advantage of the fine calorimeter segmentation and efficiently suppressing calorimeter noise. The topo-cluster formation algorithm makes extensive use of the signal-to-noise ratio for each cell, where the noise is estimated as the absolute value of the energy deposited in that cell divided by the RMS of the energy distribution measured in events triggered at random bunch crossings. The algorithm starts a cluster with a seed cell, whose signal-to-noise ratio must exceed a given threshold. Neighbouring cells that have a signal-to-noise ratio above a second (lower) threshold are added iteratively to the cluster. Then, all cells neighbouring the cluster, regardless of signal-to-noise ratio, are added to it. Finally, a splitting step is performed to separate showers from close-by particles, using local energy maxima in the cluster as seeds for a new iteration of the clustering algorithm. The energy of the final topo-cluster is defined as the energy sum of all the constituent cells. The reconstructed direction of the topo-cluster is defined as the average, weighted by absolute cell energy, of the pseudorapidities and azimuthal angles of the constituent cells.

Position-dependent time-of-flight corrections are made to the timing information associated with calorimeter cell signals. Specifically, the particle that made the deposition in the calorimeter cell is assumed to originate from the IP and travel at the speed of light; the theoretical time-of-flight to the cell position is calculated based on the point's distance from the IP, and this theoretical time-of-flight is subtracted from the actual time recorded.

4.3.4 Calorimeter Noise Discrimination

Significant coherent or sporadic noise in the LAr calorimeters can lead to the reconstruction of fake energy deposits. The characteristic ionization signal shape is used to eliminate these fake deposits. Specifically, the quadratic difference Q_{cell}^{LAr} between the measured and simulated pulse shapes in a given cell is calculated online using four measured samples of the pulse shape. A high Q_{cell}^{LAr} value indicates poor signal shape, i.e. a high probability of noise. The Q_{cell}^{LAr} is eventually used in various jet-level quantities in the Displaced LJs analysis, to distinguish real jets from jets reconstructed from fake energy deposits.

4.4 Muon Spectrometer

The Muon Spectrometer [78] is designed for two main functions: triggering on muons, and measuring their momenta. The momentum measurements are accomplished by two types of precision-tracking chambers, Monitored Drift Tubes (MDTs) and Cathode-Strip Chambers (CSCs). The CSCs cover $2.0 < |\eta| < 2.7$ in the end-caps; the remaining end-cap and barrel coverage in the range $|\eta| < 2.7$ is provided by the MDTs. Meanwhile, the trigger capabilities are provided by Resistive Plate Chambers (RPCs) and Thin Gap Chambers (TGCs). The RPCs cover the barrel ($|\eta| < 1.05$) while the TGCs cover the end-caps ($1.05 < |\eta| < 2.4$). The combination of an RPC and MDT is referred to as a "station".

The basic performance parameters of each of these 4 types of sub-systems are summarized in Table 4.3. The layout of the MS is illustrated in Figure 4.6.

Chamber	RMS, z (mm)	RMS, R (mm)	RMS, ϕ (mm)	RMS, time (ns)	n/track
MDT	0.035	–	–	–	20 (barrel and end-cap)
CSC	–	0.040	5	7	4 (end-cap)
RPC	10	–	10	1.5	6 (barrel)
TGC	–	2 - 6	3 - 7	4	9 (end-cap)

Table 4.3: The basic performance parameters of each of the 4 MS sub-systems. RMS provides a measure of the chamber resolution, while n/track refers to the number of measurements generated per track. [78]

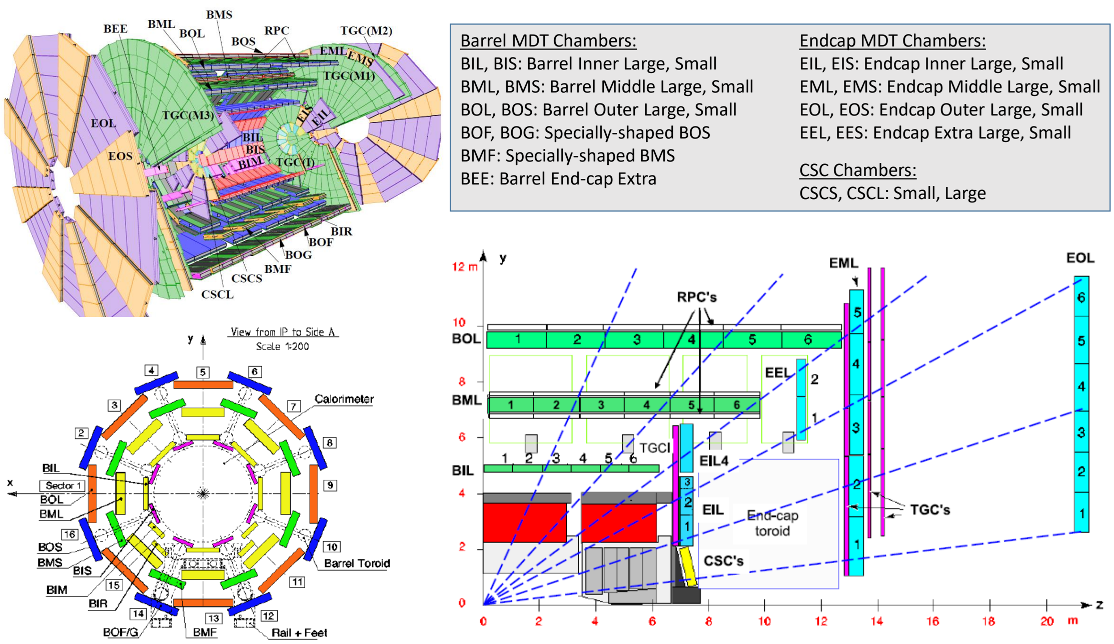

Figure 4.6: Top: overall configuration of the MS. Bottom left: cross-section of the MS barrel in the transverse plane (i.e. perpendicular to the beam line). Bottom right: cross-section of the MS in the plane containing the beam axis; muons would travel along the dashed blue lines in the absence of magnetic field, but are bent in this plane by the toroidal magnetic field. Illustrations modified from [2].

4.4.1 Precision-Tracking Chambers

The precision-tracking chambers are arranged in three-layer octants, where each octant of each layer is sub-divided into a "small" and a "large" sector. These two sectors differ slightly in lateral extent and overlap to a small degree in ϕ. The barrel layers form concentric shells at $R \simeq 5, 7.5, 10$ m. The end-cap layers form wheels at $|z| \simeq 7.4, 14, 21.5$ m. In addition, End-cap Extra chambers are located in an intermediate ring at $|z| \simeq 10.8$ m, and Barrel End-cap Extra chambers are located at the outer circumference of the end-cap cryostat. The chamber sizes increase, along with the layer dimensions, with distance from the IP. The specific shapes and dimensions are designed to optimize solid angle coverage, although some gaps in chamber coverage are unavoidable. A gap at $|\eta| \leq 0.08$ in the large sectors and $|\eta| \leq 0.04$ in the small sectors allows room for services. Additional gaps in two of the octants accommodate the "feet" of the ATLAS support structure, necessitating specially-shaped chambers in these octants. The air-core magnet design minimises the amount of material – approximately 1.3 radiation lengths per octant in most of the barrel – that muons traverse in the MS. The material budget consists mainly of the muon chambers themselves, toroid coils, vacuum vessels, and supports for the chambers and magnets.

The innermost layer in each end-cap is composed of CSCs, because the high particle fluxes and muon track densities in these regions require higher rate capability, finer time resolution, and lower neutron sensitivity. The other layers of the end-caps, and all the barrel layers, are composed of MDTs.

Monitored Drift Tubes

Each MDT chamber consists of 3 or 4 layers of drift tubes (Figure 4.7). The tubes contain pressurized Ar/CO_2 gas, with a central W-Rh wire held at high voltage to collect the electrons from ionisation. Signal transmission to the electronics is at one end of the tube, opposite from the connection to the High Voltage supply. A radial electric field is generated in the cylindrical geometry; since the coordinate of a track is determined by the radial distance from the nearest wire, the measurement accuracy depends only weakly upon the track's angle of incidence with respect to the chamber plane. Only the arrival time of the track portion closest to the wire is actually used for measuring the track coordinate, and an adjustable dead-time in the front-end of the readout chain ignores the rest of the pulse train. An online calibration system compensates for distortions in the electric field caused by positive ions, to avoid degradation in spatial resolution at high counting rates.

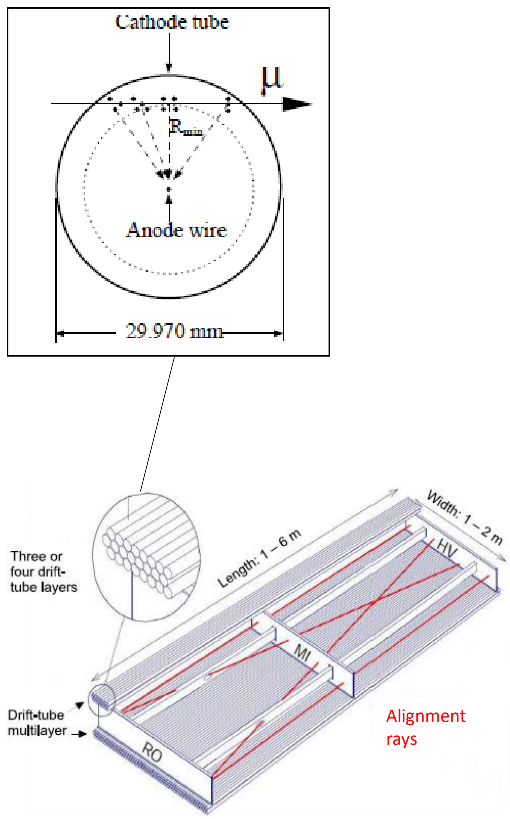

Figure 4.7: Schematic of an MDT chamber, with a close-up view of an individual tube. Illustration modified from [2].

The MDT readout employs a combination of on- and off-chamber electronics. On-chamber, there is one Amplifier / Shaper / Discriminator (ASD) chip per 8 tubes. Measurement of the pulse height of the ASD signal is performed by an Analogue-to-Digital Converter and encoded as part of the ADS output, which allows for pulse height dependent corrections to be implemented. The ASD outputs are sent to Time-to-Digital Converters, which store arrival times of leading and trailing edges. The ASDs and converters are programmable via shift registers, allowing for interventions such as changes of dead-time settings and deactivation of noisy channels. Groups of ASDs and converters are controlled by a local processor, which interacts with the off-chamber electronics and sends data to the MDT Readout Driver.

Cathode-Strip Chambers

Each CSC is a 4-plane multiwire proportional chamber with segmented cathodes (Figure 4.8). The anode wires are operated at high voltage in Ar/CO_2 gas. The wires are all oriented in the radial direction; cathode strips parallel and perpendicular to the wire allow for measurement of both R and ϕ. The track position is determined by interpolating between the induced charges on neighbouring cathode strips. The cathode segmentation is therefore designed to sample the induced charge distribution as precisely as possible, without introducing too many readout channels. Every readout strip is connected to an amplifier, with two intermediate floating strips located between every two neighbouring readout strips. The floating strips capacitively couple the induced signal to the readout strips, thereby providing additional charge interpolation and reducing the differential non-linearity of position measurements. The spatial resolution is then determined primarily by the signal-to-noise ratio, mainly limited by the electronic noise of the pre-amplifiers, and by the readout pitch. The fine timing resolution is enabled by low electron drift times, while low neutron sensitivity is achieved by keeping the gas volume low.

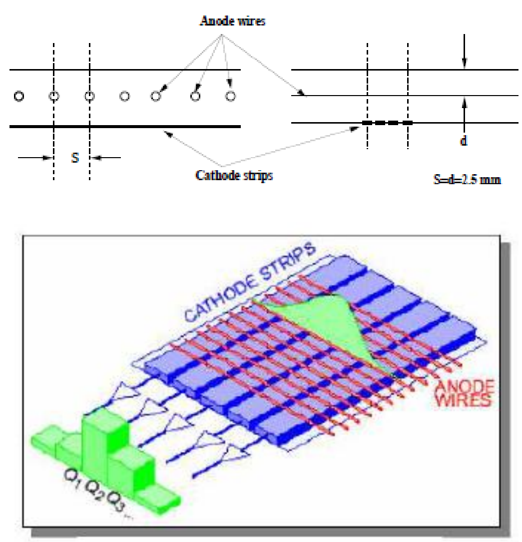

Figure 4.8: Schematic diagram of a CSC chamber. [2]

The CSC readout system is composed mostly of off-chamber electronics, as the high radiation levels pose a challenge for on-chamber components. On-chamber, for every 24 channels, there is one ASM-1 chip for amplifying and shaping the signals. Groups of ASM-1's are connected to ASM-2 boards that sample the pulse train and store the train in a capacitor array until a readout is requested. A calibration system, assisted by the high stability of the amplifiers and the low leakage of the capacitor storage system, ensures the necessary intercalibration between nearby channels. The ASM-2's digitize and multiplex the data, and send it to the off-chamber Readout Driver.

4.4.2 Trigger Chambers

The main role of the trigger chambers is to provide fast and coarse tracking and transverse momentum information. The trigger looks for fast coincidences, separately in η and in ϕ, between the three last layers along a muon track. The deviation from straightness of this path provides a measure of the track curvature, which is employed as a criterion for the track to pass predefined momentum trigger thresholds. The trigger chambers also provide ϕ measurements, which are later used to complement MDT measurements (see Section 4.4.3). Similarly to the calorimeters (see Section 4.3.3), the trigger chambers employ position-dependent time-of-fight corrections. The corrected time distribution for muons from the IP is therefore centred at 0 for each chamber.

The RPCs form 3 concentric cylindrical layers in the barrel. In the end-caps, 7 layers of TGCs complement the middle layer of MDTs ($1.05 \leq |\eta| \leq 1.92$), while 2 layers of TGCs complement the inner layer of MDTs ($1.92 \leq |\eta| \leq 2.4$). The end-caps provide a much more challenging trigger environment than the barrel, motivating the use of TGCs rather than RPCs in the end-caps. One aspect of the end-cap challenge is the radiation level, which is up to 10 times higher than in the barrel. Another aspect is the lower ratio of p_{T} to p in the end-caps (located outside the magnetic field) than in the barrel, necessitating finer and $|\eta|$-dependent granularity in the end-caps in order to maintain the same p_{T} resolution.

To avoid acceptance losses at low momenta, the trigger chambers have regions of overlap amongst each other. Algorithms in the trigger systems provide special treatment for these overlap regions to eliminate double-counting of tracks leading to fake dimuon triggers. One region, however, is often excluded from the trigger algorithms to avoid high fake rates: the transition region between the barrel and end-cap toroids, $1.30 \leq |\eta| \leq 1.65$, where there are strong magnetic field inhomogeneities.

Resistive Plate Chambers

The RPCs are gaseous parallel electrode-plate detectors, with no wires (Figure 4.9). Each chamber is composed of two contiguous units; each unit is in turn made of two independent detector layers, i.e. gas volumes, of $C_2H_2F_4$/Iso-$C_4H_{10}SF_6$. Each volume is bounded by two plastic laminate plates, separated by insulating spacers with a uniform electric field. Metallic strips mounted on the outer faces of the plates read out the signal via capacitive coupling. Strips in one direction read out ϕ, while strips in the orthogonal direction read out η. The chambers are operated in avalanche mode, providing rate-independent time resolution and low signal jitter. Each readout strip is connected to a front-end GaAs circuit consisting of a shaping amplifier and a comparator, where the voltages defining the comparator thresholds are provided by external digital-to-analogue converters. The strips are soldered directly to printed boards such that every two ϕ strips are OR-ed to form a single logical strip with the same time jitter.

The outer RPC layer, with its large lever arm, selects 9-35 GeV tracks for the high-p_{T} triggers. The two inner chambers are responsible for 6-9 GeV tracks for the low-p_{T} triggers.

Thin Gap Chambers

The TGCs are multi-wire proportional chambers (Figure 4.10), similar in operating principle to the CSCs. The wire groups measure R, while segmented readout strips on the cathode planes measure ϕ. The anode wires are operated at high voltage in n-pentane/CO_2 gas. The granularity of the chambers

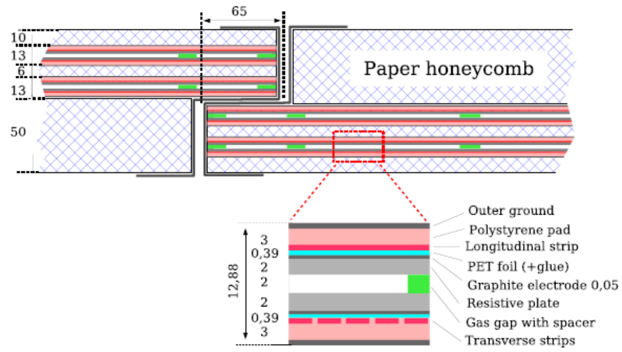

Figure 4.9: Schematic diagram of an RPC chamber. [2]

varies with $|\eta|$ in order to maintain the required momentum resolution. Staggered alignment of the wire groupings in consecutive layers, and of the radial strips, optimises the position resolution. The fine time resolution is achieved by high electric field around the wires and small wire-to-wire distances. Signals arrive within 25 ns, with azimuthal granularity of 2-3 mrad as seen from the IP. The signals are amplified in on-chamber electronics, corrected for time-of-flight and synchronized to the beam-crossing frequency, and then sent to the off-chamber electronics.

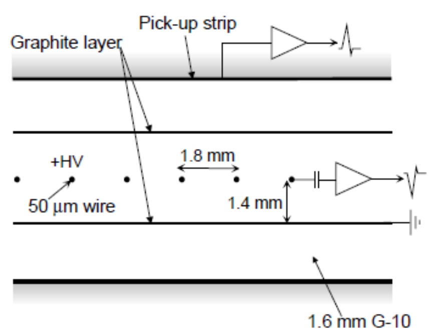

Figure 4.10: Schematic diagram of a TGC. [2]

The 7 layers of the outer TGC ring are arranged in a triplet + 2 doublets, where the triplet helps cope with false coincidences from background hits. These layers are azimuthally segmented into sectors, with each sector further azimuthally sub-divided into modules composed of up to 4 chambers.

4.4.3 Muon Tracking Reconstruction

The following sections document the process of muon track reconstruction [79]. The algorithms and performance of the muon trigger systems will be discussed in Section 4.5, while analysis-level muon quality selections will be discussed in Section 5.2.2. Muon tracks are first constructed "standalone" from the MS system only, and then combined with ID tracks and/or calorimeter information if possible. Standalone muon reconstruction issues will play a significant role in the Displaced LJ analysis, in Section

5.2.2.

Standalone Reconstruction

In the reconstruction of standalone tracks, the first stage consists of searching for hit patterns inside each chamber to form track "segments". The η track coordinate is determined from the precision-tracking chambers. The CSCs measure ϕ as well, but the MDTs do not. Therefore, hits in the MDTs are then matched in η with hits in the trigger chambers; since the trigger chambers also measure ϕ, the ϕ coordinate for tracks passing through the MDTs is taken from the trigger chambers. This method works unambiguously if there is only one track per event present in any given MDT/trigger chamber pair. For uncorrelated tracks, this is almost always the case, given the low average occupancy of the chambers. For correlated close-by muon tracks, such as from the muonic decay of a low-mass particle, there may be two or more hits from different tracks in the same chamber, creating ambiguity in the combination of $|\eta|$ and ϕ hits. The easiest solution is to push the ambiguity-resolving to the next stage of track reconstruction, where muon track candidates are combined with ID tracks. However, if no ID tracks are available, such as in displaced decays of long-lived particles, mis-reconstruction may occur. Note that within the CSCs, multitrack ambiguities are resolved more easily, providing good two-track resolution. This is accomplished by correlating the charges collected – specifically, the pulse heights – on the two planes of orthogonal strips. Due to the broad pulse height distribution in the chambers and the high resolution of the analogue-to-digital converters, different tracks almost always have different pulse height values. In the rare cases of near-equality, information from the other chamber layers is used to make the η/ϕ assignment of the tracks.

Once the segments have been built, muon track candidates are constructed from combinations of segments in different layers. The combinatorial search algorithm first uses segments from the middle layers, where more trigger hits are available, as seeds. It then extends the seeding to the other layers. The segments are selected based on hit-multiplicity and fit-quality criteria, and matched based on their relative positions and angles. A minimum of two matching segments are required to build a track, except in the barrel-endcap transition region where a single high-quality segment is sufficient.

Any given segment may belong to multiple track candidates at this stage. The next stage is overlap removal, a set of algorithms responsible for deciding the best assignment of the segment. In order to maintain high efficiency for close-by muons, it may allow a segment to be shared between two tracks. In fact, tracks with segments in three different layers are kept even if they are identical in the two innermost layers, as long as they share no hits in the outermost layer. A global χ^2 fit is applied to the hits associated with each track candidate, and a candidate is accepted if the fit satisfies certain selection criteria. Otherwise, hits responsible for large contributions to the χ^2 are removed, and the fitting is repeated. A hit-recovery procedure is also performed, and the fitting repeated if additional hits consistent with the candidate trajectory are found. Performance of the overlap removal with regard to close-together muons will play a significant role in the Displaced LJ analysis in Section 6.1.1.

For standalone reconstruction, the driving performance goal is $\sim 10\%$ transverse momentum resolution for 1 TeV tracks, which is equivalent to ~ 50 μm resolution on the measurement of a 500 μm sagitta. The goal for the range of accessible muon momenta is approximately 3 GeV to 3 TeV. The momentum resolution in the range of $\sim 30 - 200$ GeV, the range relevant for the Displaced LJ analysis, is actually dominated by multiple scattering effects rather than by the precision of the MDTs.

Combined Reconstruction

The MS tracks are incorporated into one of three main muon object types, depending on which other ATLAS subdetectors provide compatible information.

- Extrapolated: No additional information. The track parameters are re-calculated at the IP, taking into account the estimated energy loss of the muon in the calorimeters. This entails applying a loose requirement on compatibility with the IP, which may be problematic for highly displaced muon tracks. In addition, the muon is required to traverse at least 2 MS layers (3 in the forward region). This type is intended primarily for muons in the region $2.5 < |\eta| < 2.7$, which is not covered by the ID.

- Combined: Puts the entire MS track together with a track independently reconstructed in the ID. A global refit is applied using the hits from both the ID and MS, during which the algorithm may add or remove MS hits to improve the fit quality. Most muons are reconstructed "outside-in", where an MS track is extrapolated inward and matched to an ID track. "Inside-out" reconstruction, where an ID track is extrapolated outward and matched to an MS track, is also performed.

- Segment-tagged: Puts at least one local track segment in the MDT or CSC together with a track independently reconstructed in the ID. This type is intended for muons that cross only one MS layer, due either to low p_{T} or to falling in regions with reduced MS acceptance.

There is an additional muon type, Calorimeter-tagged, that does not use MS information at all. For this type, an ID track is matched to a calorimeter energy deposit compatible with a minimum-ionizing particle. It has the lowest purity of all the muon types, and is intended to cover a gap region in MS acceptance ($|\eta| < 0.1$, $15 < p_{\mathrm{T}} < 100$ GeV).

4.5 Trigger System

The Trigger and Data Acquisition system [80] is composed of a hardware-based first-level trigger (L1) and a software-based high-level trigger (HLT). L1 receives input from several detector sub-systems, including the calorimeter trigger system (L1Calo) and muon trigger system (L1Muon). Events accepted by L1 are buffered in the Read-Out System, and organized into Region-of-Interest (RoI) information before being processed by the HLT. The HLT employs the RoI information for regional reconstruction in various trigger algorithms. Events accepted by the HLT are then transferred to local storage, in different Streams based on the event type: minimum-bias, cosmic, calibration, etc.

Due to the increase in energy from Run 1, the Run 2 trigger rates were on average 2.0 to 2.5 times larger for the same luminosity with the same trigger criteria. The decrease in bunch-spacing from Run 1 further increased most muon trigger rates, due to more out-of-time pile-up (see Section 4.6). Various un-prescaled trigger thresholds therefore had to be raised from their Run 1 values. However, the HLT was upgraded from Run 1, when it had been composed of separate Level-2 and Event Filter farms. The single homogeneous farm, along with upgrades to the Read-Out System, data storage, and data collection network, allowed for various new triggers to be implemented. One such new trigger, in particular, plays an important role in the Displaced LJs analysis (see Section 5.4).

4.5.1 L1 Trigger

The L1Calo trigger is composed of a preprocessor, followed by two further processors, Cluster and Jet/Energy-sum, in parallel. Common Merger Modules then carry out the logic required to produce global trigger results. The preprocessor digitizes and calibrates the analogue signals from the calorimeters, reducing the granularity by summing the cells into trigger towers of size 0.1×0.1 in $\eta \times \phi$. It then synchronises the timing of all inputs to compensate for different times-of-flight and signal path-lengths, assigns signals to the correct bunch-crossing, carries out pedestal subtraction, applies a noise threshold, and performs a final E_{T} calibration. The Cluster processor identifies electron/photon and tau candidates, with programmable E_{T} and isolation criteria. It defines an RoI as a 2×2 tower cluster in the ECal in which the E_{T} of at least one nearest-neighbour tower pair exceeds a given threshold. Isolation-veto thresholds can be set using the surrounding towers, as well as using the towers directly behind in the HCal. The tau algorithm differs from the electron/photon algorithm in that it favours narrow hadronic objects. The Jet/Energy-sum turns jet trigger elements (0.2×0.2 sums in $\eta \times \phi$) into jet objects, using sliding windows of 0.2×0.2 to 0.4×0.4, and produces global sums of scalar and missing transverse momentum. It defines an RoI as 2×2 to 8×8 tower clusters for which the summed EM and hadronic E_{T} exceeds a given threshold. The electron/photon/tau triggers extend to $|\eta| < 2.5$, and the jet triggers to $|\eta| < 3.2$. The energy in L1 towers is calibrated at the "electromagnetic energy scale", which accurately reconstructs the energy deposited in an EM shower, but underestimates the energy deposited by hadrons. This becomes especially important for triggers that later make use of the ratio of hadronic to EM energy deposition, as is the case for one of the Displaced LJs triggers (see Section 5.4).

The L1Muon trigger searches for coincidences of hits in two or three trigger stations within a road from the IP through the MS. The width of the road scales inversely with the p_{T} threshold to be applied, with a total of 6 thresholds operating concurrently. In the barrel, the low-p_{T} trigger requires 3-out-of-4 coincidence of the two inner RPC layers; the high-p_{T} trigger looks for an additional coincidence with the outer RPC layer. In the end-cap, 3-out-of-4 coincidence is required in the doublet pair planes, 2-out-of-3 coincidence for the triplet wire planes, and 1-out-of-2 hits for the triplet strip planes. The triplet coincidence requirements are based on the fact that in each triplet, signals from three wireplanes but only two strip-planes are read out for the trigger. The search algorithm is performed in both the ϕ and η projections to reduce fake triggers. RoI's are defined as approximately 0.1×0.1 in $\eta \times \phi$. The output of the L1Muon trigger provides information on up to two muon-track candidates per muon trigger sector, including the RoI position and the p_{T} threshold passed for each candidate.

4.5.2 HLT Trigger

The HLT performs two versions of muon reconstruction: a "fast" stage specifically for the trigger, followed by a "precision" stage similar to the offline process. Since the offline process was detailed in Section 4.4.3, only the fast stage will be discussed here. The fast reconstruction refines each L1 muon candidate by including MDT data in the given RoI. A track fit is performed using the MDT drift times and positions, and a p_{T} value is assigned using lookup tables. This creates an "MS-only" muon track, which is then extrapolated back to the IP using the same track extrapolator that is employed offline. If the extrapolation is compatible with an ID track, a combined muon candidate is formed with refined track parameter resolution. The majority of muon triggers employ combined muon candidates, but some specialized triggers, as in the Displaced LJs analysis, employ MS-only tracks.

The HLT also performs fast and precision track triggering. The fast stage consists of trigger-specific pattern recognition algorithms. The precision stage is similar to the offline version, but seeded by tracks and space-points identified by the fast stage, and typically configured to run only within the given RoI's.

The HLT employs a series of reconstruction algorithms to convert calorimeter readouts into cell objects and cluster objects. These objects serve as input to the reconstruction of electron, photon, tau, and jet candidates, as well as $E_{\mathrm{T}}^{\mathrm{miss}}$ and muon isolation. The HLT has access to the full detector granularity, greatly improving the energy and position measurements with respect to L1. The reconstruction proceeds in two stages: unpacking and clustering. The unpacking of the calorimeter data can be performed only for the given RoI's, which is useful for well-separated subjects such as electrons, photons, muons, and taus. Alternatively, it can be performed for for the full calorimeter, which is necessary for jets and for global event quantities such as $E_{\mathrm{T}}^{\mathrm{miss}}$. The unpacked data is then converted into a collection of cells and passed to a clustering algorithm. Two different clustering algorithms can be used to reconstruct the clusters of deposited energy: fast, which uses a sliding-window, and topo-clustering, which is similar in performance to the offline reconstruction.

4.6 Luminosity and Pileup

The preferred sub-detector for luminosity measurements in 2015 data-taking was the Beam Conditions Monitor (BCM) [81], also used for generating beam-abort signals and studying beam-induced backgrounds. It consists of one station on either side of, and close to, the IP. Each station is composed of 4 modules, with 2 sensors each, distributed symmetrically in ϕ, at $|z| = 184$ cm, $|\eta| = 4.2$. To tolerate the high radiation doses, the sensors are diamond rather than silicon. FPGA electronics digitize the signals and determine in-time and out-of-time multi-module coincidences. The BCM luminosity measurements are cross-checked by measurements from another sub-detector, LUCID (LUminosity measurement using Cerenkov Integrating Detector). It is located much further from the IP, at $z \approx 17$ m [2].

Pileup, i.e. simultaneous pp collisions in the same event, plays an important role in all Run 2 analyses, even more than in Run 1. "In-time" pileup refers to the number of interactions that occur in the same bunch crossing. The number of reconstructed vertices in a given event is generally indicative of the amount of in-time pileup; it increases with instantaneous luminosity. "Out-of-time" pileup refers to overlapping signals in the detector from neighbouring bunch crossings. Smaller in-train bunch separation leads to more out-of-time pileup. The average number of interactions per bunch crossing is denoted $< \mu >$. The smallest chunk of data in ATLAS in which luminosity is made available for analyses is known as a luminosity block ("lumiblock"), which is composed of $\sim 10^9$ bunch crossings. However, the pileup often varies significantly between bunch crossings within a given lumiblock. $< \mu >$ can be taken as an average over all bunch crossings in one particular lumiblock, or over multiple lumiblocks. Most analyses, including the Displaced LJ analysis, employ the former rather than the latter measure.

4.7 Key Points for Displaced LJ Analysis

The Muon Spectrometer plays the key role in the reconstruction of muonic Displaced LJs, making the details of its components and measurement capabilities highly relevant in the design of the analysis. Since electron and pion Displaced LJs manifest as jets in the calorimeters, the properties of the EM and Hadronic Calorimeters are important primarily from the point of view of jet reconstruction. The Inner

Detector is relevant in the analysis only for vetoing prompt objects and for quantifying LJ isolation. The workings and limitations of the trigger systems are determining factors in the analysis sensitivity, in view of the importance of innovative triggers for Displaced LJs. Luminosity is employed for normalizing MC samples and calculating limits; pileup considerations enter into the systematic uncertainties and the optimization of certain analysis variables.

Chapter 5

Displaced Lepton-Jet Search Description

This chapter gives the step-by-step methodology of the LJ analysis, entirely blinded (i.e. without examining the signal region in 2015 data) and without revealing any actual results. It can hence be used as a recipe for independently replicating the analysis, without bias from foreknowledge of results. The first order of business is to gain some physical intuition about the signal, using Monte Carlo (MC) of the benchmark processes. With this intuition in mind, the analysis defines and constructs LJ candidates, of a few different classes, using pre-selected physics objects from events that pass a set of triggers. It then proceeds through two stages of selection, which involve imposing sets of requirements ("cuts"): first at the level of individual LJ candidates, and then at the level of entire events. The first stage is MC-driven, while the second stage is data-driven. Descriptions are provided of the main backgrounds against which the cuts are designed to discriminate, and of the methodologies for optimizing the cut values. Relevant details of the software framework for the analysis are also provided along the way. For reference, a summary of the analysis flow appears in Section 5.7.

The blinding scheme dictated that the analysis code had to be run on all MC up to the end of the LJ level selection, with the values of these cuts fully determined, before applying any code to data. The rest of the code had to be set up so as to dynamically determine the values of the data-driven cuts and generate the final event yields all in one go. This part of the code was tested on $\sim 10\%$ percent of the data, in order to check that it worked properly, before it was applied to the entire data set. Approval from the UEH subgroup was required before un-blinding. While the systematic uncertainties were largely determined with the analysis still blinded, discussions of the systematics are postponed to Chapter 7, in the interests of readability and flow of this document.

5.1 LJ Monte Carlo

The first step in the analysis is to characterize the ATLAS detector response to the $\gamma_{\rm d}$ decay products as a function of various key observables:

- Particle species
- $c\tau$ of $\gamma_{\rm d}$

- ΔR of LJ constituents
- p_{T} of γ_{d} and of LJ constituents
- η of γ_{d}

5.1.1 LJ Gun Monte Carlo Tool

A specialized MC tool known as the "LJ Gun" was developed, which simulates LJs with one or two γ_{d}'s as in Figure 3.8. The LJ Gun decays the γ_{d}'s into electron, muon, and charged pion pairs according to the γ_{d} mass, as in Section 3.4. Two different polarization states, longitudinal and transverse, can be selected for the γ_{d}. Pile-up, i.e. multiple interactions in the same bunch crossing (see Section 4.6), are included in the simulation. The search for Displaced LJs necessitates the exploration of the sensitivity of multiple different sub-detectors – ECal, HCal, and MS – to the γ_{d} decay products. Consequently, the γ_{d} lifetime is chosen such that, taking into account the boost factor of the γ_{d}, 80% of the decays occur inside a cylinder of 8 m radius and 14 m of half-length. All γ_{d}'s that decay outside this fiducial volume are re-generated until the decay point lies in the fiducial volume. The LJ Gun thereby guarantees that all the relevant regions of the detector are equally populated.

The LJ Gun was originally developed by members of the Run 1 Displaced LJs analysis team, prior to my joining the analysis effort. Substantial changes were made to the core ATLAS MC framework for Run 2; I was responsible for upgrading and maintaining the LJ Gun to allow it to run within the Run 2 framework, and for helping validate the simulated samples it produced. For LJs with one γ_{d}, samples of 100000 events were generated with γ_{d} masses of 0.05, 0.15, 0.40, 0.90 and 1.5 GeV, in both possible polarization states. These masses were specifically chosen to cover a large range of BRs of the γ_{d}. The ϕ, η, and p_{T} of the γ_{d}'s were generated according to a uniform distribution in the ranges 5 to 100 GeV, -2.5 to 2.5, and -3.14 to 3.14 rad, respectively. For LJs with two γ_{d}'s, samples of 100000 events have been generated with s_{d} masses of 1.0, 2.0, 5.0 and 10 GeV. For each mass of the s_{d}, all the kinematically allowed γ_{d} masses in the one-γ_{d} set were generated, with the same η, ϕ, and p_{T} distributions.

5.1.2 Benchmark Model Monte Carlo

The set of parameters used to generate the MC samples for the FRVZ benchmark models (see Section 3.5.2) are listed in Table 5.1. The Higgs boson is generated through the gluon-gluon fusion production mechanism, the dominant process for a low-mass Higgs boson. The gluon-gluon fusion Higgs boson production cross section in pp collisions at $\sqrt{s} = 13$ TeV, estimated at the next-to-next-to-leading order [67], is $\sigma = 43.92$ pb for $m_H = 125$ GeV. The masses of the HLSP and f_{d_2} are chosen to be light relative to the Higgs boson mass, and far from the kinematic threshold at $m_{\mathrm{HLSP}} + m_{\gamma_{\mathrm{d}}} = m_{f_{\mathrm{d}_2}}$. For the chosen γ_{d} mass (0.4 GeV), the γ_{d} decay BRs are set to their expected values according to Section 3.4: 0.45 e^+e^-, 0.45 $\mu^+\mu^-$, 0.10 $\pi^+\pi^-$. The γ_{d} proper lifetime $c\tau$ is set such that a large fraction of the decays occurs inside the sensitive ATLAS detector volume, i.e. up to ~ 8 m in R and ~ 14 m in z. The detection efficiency can then be estimated for a range of γ_{d} mean lifetimes through re-weighting of the generated samples; the $c\tau$ must of course be converted to a lab-frame decay distance taking into account the large boost of the γ_{d}.

Samples for each of the two benchmark processes were produced separately, with 300000 events each. Additional samples with a high-mass BSM Higgs boson (800 GeV) rather than a SM Higgs boson, but the same decay chain, were also simulated. A custom model in MadGraph5 v2.2.3 [82] for parton-level simulation, linked to the PYTHIA8 generator [83] v8.186 for showering and hadronization, was used for the truth-level MC. I was responsible for creating and maintaining the MadGraph5 model, and ensuring proper linkage with the rest of the simulation chain.

m_H [GeV]	$m_{f_{d_2}}$ [GeV]	m_{HLSP} [GeV]	$m_{s_{d_1}}$ [GeV]	m_{γ_d} [GeV]	$c\tau_{\gamma_d}$ [mm]
125	5.0	2.0	-	0.4	47.0
125	5.0	2.0	2.0	0.4	82.40
800	5.0	2.0	-	0.4	11.76
800	5.0	2.0	2.0	0.4	21.04

Table 5.1: Parameters used for the MC simulation of FRVZ events, for $m_H = 125$ and $m_H = 800$ GeV. First and third rows are for the $2\gamma_d$ final state (first benchmark model); second and fourth rows are for the $4\gamma_d$ final state (second benchmark model). Reproduced as in [6].

5.1.3 Software Framework

The standard ATLAS MC Simulation Data Flow [84] is shown in Figure 5.1. The LJ Gun creates EVNT files directly; in the case of the FRVZ samples, the EVNT files are created by MadGraph5 + Pythia8. The MC flow is more complicated than for data, which simply proceeds from Raw byte-stream format to ESD. The Raw byte-stream represents the event data produced from the HLT, roughly analogous to RDO.

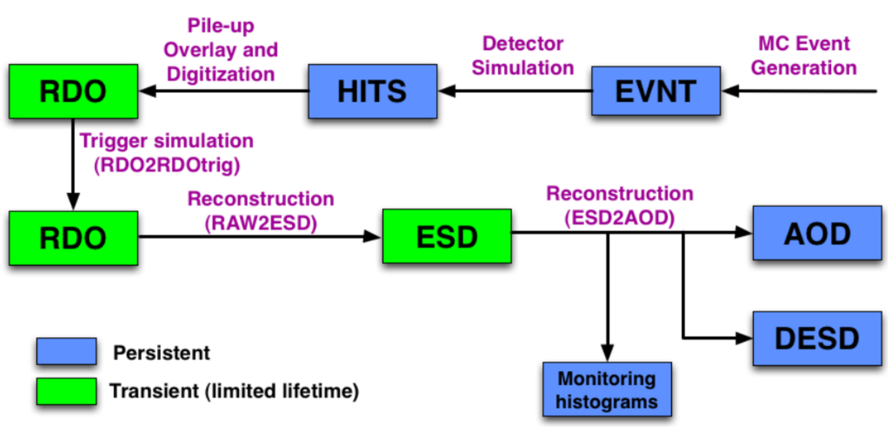

EVNT (Event Data): contains the truth event information as produced by MC generators	**ESD (Event Summary Data):** contains the detailed output of the detector reconstruction
HITS (Hits Data): contains simulated energy deposits in active detector volumes and related particle information	**AOD (Analysis Object Data):** contains a summary of each reconstructed event, with sufficient information for common analyses
RDO (Raw Data Object): Representation of the byte-stream information (i.e. RAW data format) in simulation	**DESD (Derived Event Summary Data):** contains detailed information for targeted events, for specialized analyses

Figure 5.1: The standard ATLAS MC Simulation Data Flow [85].

5.2 Searching for LJ Candidates

With the help of the MC samples, the next steps in the analysis consist of:

- Defining classes of LJs.
- Performing pre-selection on the classes of reconstructed physics objects that may become LJ constituents.
- Building LJ candidates from the preselected physics objects.

5.2.1 LJ Classification

At the detector level, a γ_d decaying to a muon pair before the first MS trigger chambers is identified by two muons in the MS, while a γ_d decaying to an electron/pion pair before or inside the calorimeters is seen as one or two calorimeter jets. Hence, a cluster of only muons and no jets in a narrow cone is the signature of an LJ with all γ_d's decaying to muon pairs. A cluster of two muons and one or two jets is typical of an LJ with one γ_d decaying into a muon pair and one γ_d decaying into an electron/pion pair. An LJ with one or two γ_d's , both decaying to electron/pion pairs, results in one or more jets.

LJs are defined and classified according to the muon and/or jet content, found within a cone of opening ΔR, and satisfying specialized selection criteria. If at least two muons and no jets are found in the cone, the LJ is classified as Type0. Otherwise, if there are at least two muons and only one jet in the cone, the LJ is of Type1. The remaining jets satisfying the selection criteria, and with no muons in their cones, are defined as Type2 LJs. LJs containing only one γ_d can contribute only to Type0 and Type2; a Type1 LJ must contain at least two γ_d's . The classification of LJs is therefore:

- Type0: LJs with all γ_d's decaying to muons. This type selects γ_d decays beyond the pixel detector, up to the first trigger plane of the MS.
- Type1: LJs with one γ_d decaying to a muon pair and one γ_d decaying to an electron/pion pair. The range of decay distances targeted by this type extends from the last ID pixel layer up to the end of the HCal for γ_d decaying into an electron/pion pair, and from the last ID pixel layer up to the first trigger plane of the MS for the γ_d decays to muons.
- Type2: LJs with all γ_d's decaying to electron/pion pairs in the HCal. The displaced decay requirement is crucial, as it greatly reduces the otherwise-overwhelming background due to QCD multi-jet production.

Figure 5.2 schematically shows the LJ classification.

5.2.2 Muon Pre-Selection

For Type0 and Type1 LJs, "displaced" muons are defined as muons whose tracks begin beyond the last Pixel layer of the ID. The standard ATLAS muon reconstruction always attempts to match MS tracks to ID tracks (see Section 4.4.3) to form Combined muons. However, the ID track reconstruction requires at least one hit in the pixel layers. Therefore, a displaced muon has no ID track, and is reconstructed only as an Extrapolated muon using MS information. For the purposes of this analysis, a "Non-Combined"

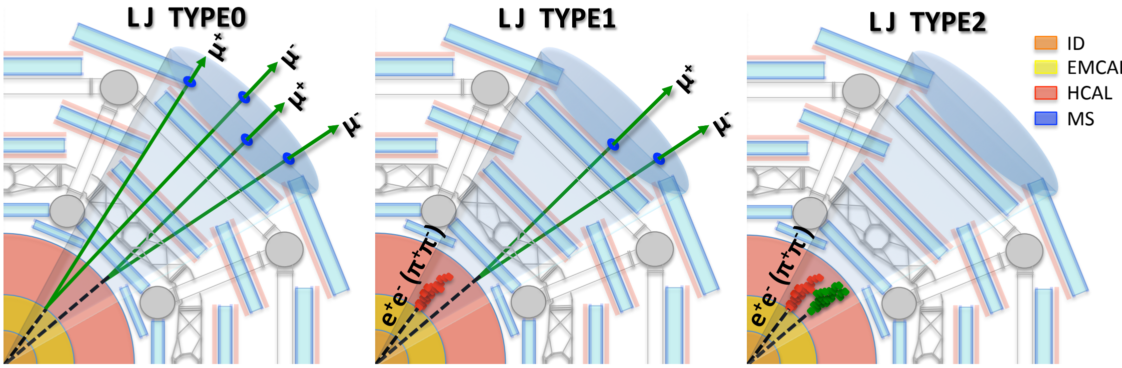

Figure 5.2: Schematic picture of the LJ classification according to the γ_d decay final states. Left: Type0 LJ, only muons. Centre: Type1 LJ, muons and jets. Right: Type2 LJ, only jets. Reproduced as in [6].

muon is defined as an Extrapolated muon that does not match any Combined muon. By requiring muon constituents of Type0 and Type1 LJs to be Non-Combined, prompt muons from the IP are discarded.

The following quality criteria for the selection of the muons follow the recommendations of the ATLAS Muon Combined Performance group, except requiring a minimum of two MDT stations instead of the standard three. This reduction allows for the recovery of decays that occur after the innermost layer of muon chambers but before the first MS trigger chamber. (See Section 4.4 for a description of the components of the MS.)

- Barrel: MS track with hits in at least two MDT stations (number of hits ≥ 10) and at least one RPC ϕ-hit;
- Endcap: MS track with ≥ 1 hit in the TGC stations, and at least one of the following cases:
  - ≥ 10 hits in the MDT stations
  - ≥ 3 hits in the CSC chambers and ≥ 3 hits in the EML/EMS/EOL/EOS MDT stations.

The muon search is limited to $|\eta| < 2.4$, corresponding to the ID and muon trigger coverage. Tracks in the $|\eta|$ interval 1.0 to 1.1 are rejected, to avoid a transition region of the MS between barrel and endcap.

5.2.3 Electron and Pion Pre-Selection

Calorimeter jets for Type1 and Type2 LJs are built from topoclusters (see Section 4.3.3) using the anti-k_t algorithm [86], with radius parameter $R = 0.4$. This algorithm has become the default for use in ATLAS because it is quite fast, and provides infrared and colinear safety.

The displaced decay of a γ_d in the HCal produces a jet with unusually low electromagnetic fraction f_{EM}, defined as the ratio of the energy deposited in the EM calorimeter to the total jet energy, since the γ_d deposits no energy as it traverses part of the HCal. The standard minimum f_{EM} criterion is therefore removed from the standard ATLAS "loose cleaning" jet quality selection criteria [87] that the calorimeter jets are required to satisfy. In order to eliminate jets reconstructed from fake energy deposits in the calorimeters, the remaining criteria consist of removing jets that meet any of the following conditions:

- $f_{HEC} > 0.5, \ |f_Q^{HEC}| > 0.5, \ <Q> \ > 0.8$

- $|E_{neg}| > 50$ GeV
- $f_{EM} \ > 0.95,\ f_Q^{LAr} > 0.8,\ < Q > \ > 0.8, |\eta| < 2.8$
- $f_{max} > 0.99,\ |\eta| < 2$

where the various discriminating variables are defined as follows, with Q_{cell}^{LAr} defined in Section 4.3.4:

- $< Q >$: Average jet quality, the energy-squared weighted average of the Q_{cell}^{LAr} of the cells in the jet, normalized to between 0 and 1.
- f_Q^{LAr}: Fraction of the energy in the LAr calorimeter cells of a jet that have $Q_{cell}^{LAr} > 4000$, i.e. poor signal shape quality.
- f_Q^{HEC}: Fraction of the energy in the HEC calorimeter cells of a jet that have $Q_{cell}^{LAr} > 4000$, i.e. poor signal shape quality.
- $|E_{neg}|$: Sum of the energy of all cells in the jet with negative energy. A small amount of negative energy can be present in a good jet due to electronic and pile-up noise; large amounts of negative energy are generally due to noisy calorimeter cells.
- f_{max}: The maximum fraction of the total jet energy deposited in any single calorimeter layer.

This selection is performed by a standard ATLAS analysis tool.

A jet energy calibration procedure is then performed by another set of standard ATLAS analysis tools, as per the procedure described in [88]. The important aspects include:

- Area-based pile-up subtraction, along with residual pile-up corrections
- p_{T} and η-dependent energy corrections, derived from nominal 2015 MC simulation
- Corrections to reduce the dependence of the jet energy measurement on the longitudinal and transverse structure of the jets
- Corrections for jets that are not fully contained in the calorimeters.

Finally, an additional cut $p_{\mathrm{T}} \geq 20$ GeV is then imposed on the calibrated jets.

5.2.4 Construction of LJ Candidates

A dedicated analysis tool was created expressly for the purposes of identifying LJ candidates in samples. Its modular structure and configurable properties allow it to be used by other analyses featuring LJs (such as Prompt LJs), not just by Displaced LJs. It constructs analysis-level LJ objects, whose constituents are required to pass the above pre-selection criteria. It also calculates, and provides easy access to, a number of kinematical and compositional properties of the LJ candidates. For example, the overall LJ momentum is obtained from the vector sum over all momenta of the LJ constituents. I was responsible for the creation, maintenance, and validation of this tool.

Type0 and Type1

Type0 and Type1 LJ candidates are built using a simple clustering algorithm that combines all the muons lying within a cone of fixed size in (η, ϕ) space. The algorithm is seeded by the highest p_{T} muon. If at least two muons and no jets are found in the cone, the LJ is accepted as Type0; if at least two muons and at least one jet are found, the LJ is taken as Type1. The search is then repeated with all unassociated muons, until no muon seeds are left. The size of the search cone for the above algorithm is determined based on the distribution of the maximum ΔR between the final-state muons in the MC samples of the benchmark models (Figure 5.3). The cone size is chosen as $\Delta R = 0.5$, to contain all the γ_{d} decay products in nearly all cases (or 99.8% in the $4\gamma_{\mathrm{d}}$ $m_H = 800$ GeV case).

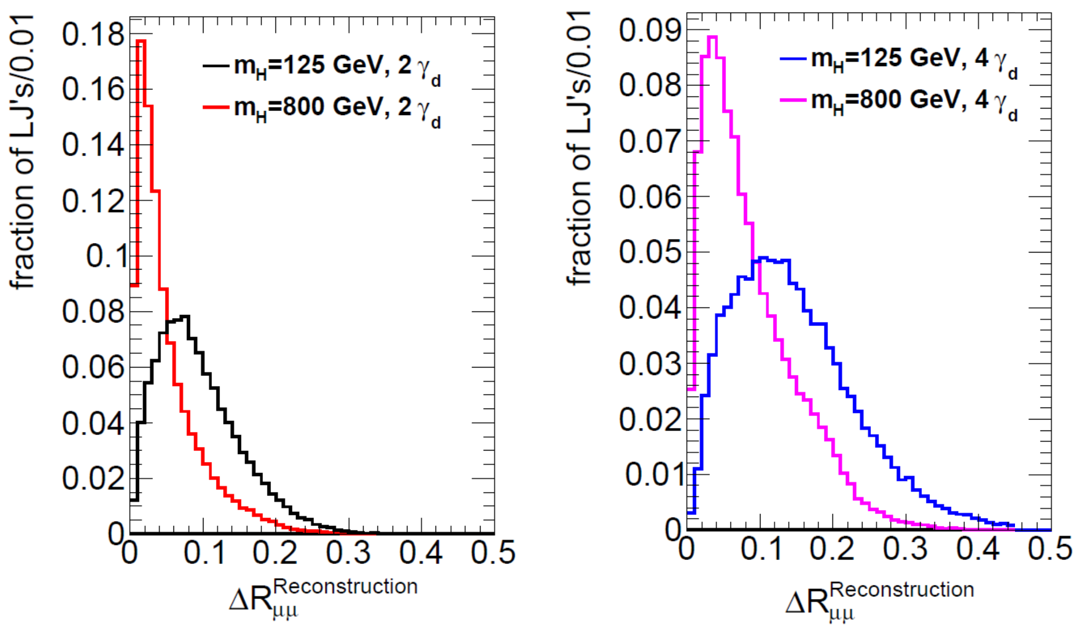

Figure 5.3: Distribution of the maximum ΔR between the final-state muons, evaluated in MC samples of the benchmark models, one sample with SM Higgs boson and one with a BSM heavy Higgs boson. Left: benchmark process with $2\gamma_{\mathrm{d}}$. Right: benchmark process with $4\gamma_{\mathrm{d}}$.

Type2

A Type2 LJ candidate requires only one calorimeter jet, with no muons within a cone of $\Delta R = 0.4$. Table 5.2 shows the fraction of γ_{d}'s decaying into electron/pion pairs that yield a reconstructed jet, in MC. Even in the case of $s_{\mathrm{d}} \to 2(\gamma_{\mathrm{d}} \to e^+e^-/\pi^+\pi^-)$, only a single jet – rather than multiple jets – is reconstructed from the final-state electrons/pions.

5.3 Software Framework

The search for LJ candidates begins with xAOD files (see Section 5.1.3). These files are then put through a centralized reduction framework, which makes a much smaller file (DxAOD) from the xAOD. The framework accomplishes this by removing entire events that do not pass the analysis triggers,

Process	m_H [GeV]	Reconstructed Jets per LJ	Fraction of LJs
$\gamma_{\rm d} \to e^+e^-/\pi^+\pi^-$	125	0	0.66
		1	0.34
	800	0	0.46
		1	0.54
$s_{\rm d} \to 2(\gamma_{\rm d} \to e^+e^-/\pi^+\pi^-)$	125	0	0.70
		1	0.30
	800	0	0.36
		1	0.64

Table 5.2: Reconstructed jets per LJ in MC, for $\gamma_{\rm d} \to e^+e^-/\pi^+\pi^-$ and for $s_{\rm d} \to 2(\gamma_{\rm d} \to e^+e^-/\pi^+\pi^-)$.

removing certain types of information or collections of data objects that are irrelevant for the analysis, and removing individual data objects irrelevant for the analysis. A “Derivation”, which specifies exactly which of the above operations is performed, is defined for the Displaced LJs analysis. I was responsible for the definition of the Derivation, and liaising with the reduction framework specialists to ensure the proper production of the DxAODs. Along with the DxAOD files, configuration files are also generated to accompany MC datasets, to allow the MC events to be re-weighted for pileup conditions. MC samples are generally produced in advance of the actual data-taking period, at which time only an estimate of the pileup conditions can be put into the simulation. The MC samples then need to be re-weighted during analysis, in order to match the pileup conditions actually found in the data.

The ATLAS software is structured such that analysis on the DxAOD files may be performed either using the proprietary ATLAS Athena (C++ based) framework, or using ROOT. A number of standard ATLAS tools, to accomplish standard tasks such as performing pileup re-weighting and applying quality criteria and calibrations, exist under both frameworks. The Displaced LJ analysis employs the second option. The result is a file of ROOT n-tuples and histograms. The analysis flow through the software is shown in Figure 5.4.

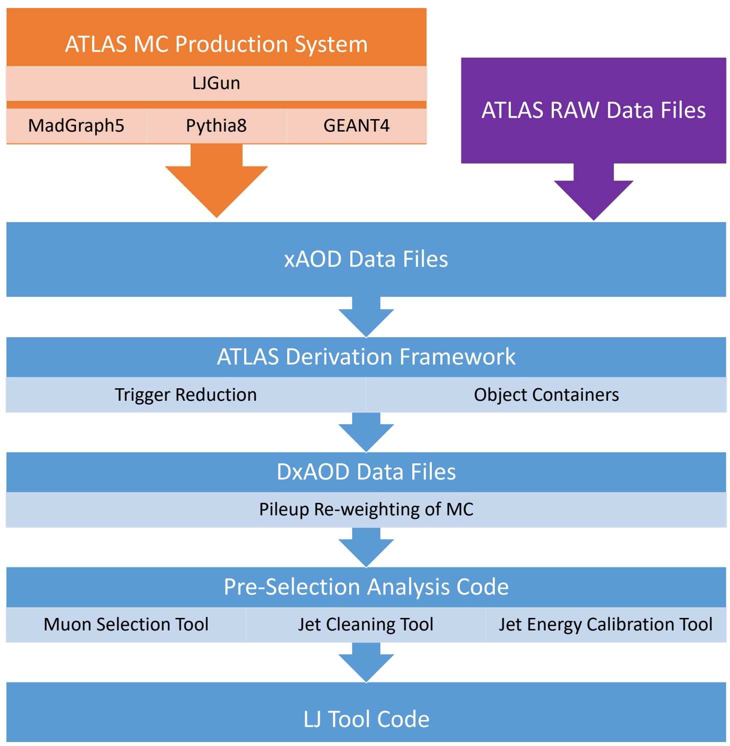

Figure 5.4: The analysis flow. Orange: pathway for MC only. Violet: pathway for data only. Blue: common pathway. Lighter-coloured rectangles indicate the use of sub-programs or tools. Modified from [85].

5.4 Triggers

Having clearly defined what constitutes an LJ candidate, it is vital to consider triggers to select events with such candidates. At this point, it is also necessary to begin considering the main sources of background to the LJ signal. The background with the largest cross-section, dwarfing the signal cross-section by several orders of magnitude, is QCD multi-jet production. Discrimination against this background must begin at the trigger level. Other background sources are cosmic ray muons that cross the detector in time coincidence with a bunch-crossing interaction, beam halo muons (beam-induced background, abbreviated BIB), and cosmic ray muons giving hard bremsstrahlung in the HCal. Since these backgrounds have much smaller rates than QCD multi-jet production, discrimination against them at the trigger level is less important, and can be left to a greater extent to later analysis stages.

The majority of the standard ATLAS triggers [80] are optimized for selecting objects originating from the IP, and are very inefficient in selecting events with displaced decay vertices. Therefore the analysis uses the logical OR of the following dedicated triggers, which do not favour objects originating from the IP:

- Tri-muon MS-only : Selects events with at least three MS-only muon tracks with $p_{\mathrm{T}} \geq 6$ GeV.
- Narrow-Scan: Selects events with two MS-only muon tracks in a narrow cone.
- CaloRatio: Selects events with an isolated calorimeter jet of low EM fraction.

As the use of these innovative triggers plays a large role in the overall sensitivity of the analysis, the workings of the triggers are discussed in detail below.

5.4.1 Trigger Descriptions

The Tri-Muon MS-only trigger [89] is seeded at L1 by a cluster of three muon ROIs in a $\Delta R = 0.4$ cone. At HLT, if any jet with $p_{\mathrm{T}} \geq 2$ GeV is found within a $\Delta R = 0.5$ cone, the event is rejected. This strong rejection is not ideal from a signal acceptance point of view, but is necessary for discrimination against the QCD background; without the rejection, the trigger rate becomes too high for the trigger system to tolerate.

The Narrow-Scan trigger was introduced for the 2015 data-taking, and adopts a specialised and novel approach for a wide range of signal models featuring highly collimated muons. The Narrow-Scan algorithm begins by requiring at least one L1 trigger muon object, unlike other multi-muon triggers which usually require at least two L1 trigger muon objects. Requiring at least two L1 objects leads to severe signal efficiency losses for collimated muons, due to the limited granularity at L1 which results in fewer reconstructed L1 muon objects than truth muons, and due to matching ambiguities between the L1 and HLT muon objects. To compensate for the high rate from only one L1 muon object, the HLT searches for a second muon without requiring it to match an L1 muon object. To limit the online resources consumption, the HLT search is limited to a $\Delta R = 0.5$ cone around the L1 muon. Neither of the HLT muons is required to have a matching ID track; in fact, the second muon is explicitly required to be unmatched to an ID track. The first muon must have $p_{\mathrm{T}} \geq 20$ GeV, while the second muon must have $p_{\mathrm{T}} \geq 6$ GeV.

The CaloRatio trigger [89] is seeded at L1 by a tau candidate with ≥ 60 GeV. A tau seed is used rather than a jet seed because the tau seed uses a narrower calorimeter region. An RoI of calorimeter cells is built around the L1 seed, and anti-k_t jet reconstruction with $R = 0.4$ performed in the RoI. The reconstructed jet must have $|\eta| \leq 2.5$ to ensure ID tracks can be matched to it, as well as $E_{\mathrm{T}} \geq 30$ GeV. A cut on the calorimeter energy ratio is then applied, requiring $f_{EM} \leq 0.0.063$, and the surviving jets are used to determine RoIs for track reconstruction. Precision track triggering (see Section 4.5.2) is run in each of these regions; any jet found to have a track with $p_{\mathrm{T}} \geq 2$ GeV within $\Delta R \leq 0.2$ of its axis is rejected. Finally, a beam halo rejection algorithm is applied (see Section 5.5.5) to reject fake jets from beam halo muons.

5.4.2 Trigger Efficiencies

The resulting trigger efficiencies on the FRVZ models, defined as the ratio between the number of triggered events and the total number of MC generated ones, are shown in Table 5.3. Efficiencies for the case of a BSM 800 GeV Higgs boson, as well as for the SM-like 125 GeV Higgs boson, are included. The Narrow-Scan trigger, which was not yet available for the Run 1 analysis, provides the majority of the trigger efficiency in Run 2. As shown in the Table, trigger efficiencies are a factor of $\sim 2-3$ higher in Run 2 than in Run 1.

Trigger	$2\gamma_{\mathrm{d}}$ $m_H = 125$ GeV	$2\gamma_{\mathrm{d}}$ $m_H = 800$ GeV	$4\gamma_{\mathrm{d}}$ $m_H = 125$ GeV	$4\gamma_{\mathrm{d}}$ $m_H = 800$ GeV
Tri-muon MS-only	2.0	2.4	4.9	7.8
Narrow-Scan	10.6	23.0	8.3	38.4
Calo-Ratio	0.3	9.7	0.1	7.4
OR of all	11.9	32.0	11.8	44.8
Run 1	4.6	-	6.2	-

Table 5.3: Trigger efficiencies, in percent, for MC samples of the FRVZ benchmark processes. The final row shows the overall trigger efficiency from the Run 1 Displaced LJ search, which had no $m_H = 800$ GeV benchmark, for comparison. The other rows are reproduced as in [6].

5.5 LJ Selection

Once LJ candidates are identified in events that pass the triggers, two stages of selection are performed:

- Selection on the individual LJ candidates.
- Event-level selection on events that have at least two LJs passing the selection criteria.

This section discusses the first stage; Section 5.6 discusses the second stage.

In defining the selection criteria, the goal is to achieve an optimal signal-to-background ratio, while keeping enough signal candidates to maintain reasonable discovery or limit-setting potential. As mentioned in the previous section, the main sources of background to the LJ signal are cosmic ray muons that cross the detector in time coincidence with a bunch-crossing interaction, BIB, cosmic ray muons giving hard bremsstrahlung in the HCal, and QCD multi-jet production. A secondary goal is to maintain as much model independence as possible, up until the limit-setting stage. For example, invariant mass cuts on the LJs and $E_{\mathrm{T}}^{\mathrm{miss}}$ cuts were deliberately avoided, as they would strongly depend upon the dark sector mass hierarchy, as were boosted decision trees. Deciding on the criteria, and optimizing the cut values, required the consensus and combined efforts of all members of the Displaced LJs analysis group.

5.5.1 Data Samples for Background Studies

The nominal LHC configuration for pp collisions contains 3564 bunch crossings per revolution, but not all bunches are actually filled with protons. "Empty" bunch crossings are unfilled, and furthermore have at least 5 unfilled bunches on both sides. These allow for the study of cosmic ray events. The same triggers used by the analysis in the filled bunches were also active during the 2015 data-taking in the empty bunch crossings. The Tri-Muon and Narrow-Scan events in the empty bunch crossings were recorded in the "CosmicMuons" data stream, while the Cal-Ratio events were recorded in the "CosmicCalo" stream. The average ratio between filled and empty bunch crossings was $\sim 2:1$, but the LJ analysis triggers were heavily pre-scaled (by a factor of ~ 50) in the empty bunch crossings. This yielded an "Empty" data sample that has small statistics, but still useful for studying the cosmic backgrounds.

Meanwhile, a di-jet control sample consisting of ~ 1.0 fb^{-1} of the 2015 data is used for studying QCD multi-jet background. Events in this sample are required to pass at least one of the ATLAS dedicated single-jet triggers, and contain two jets passing the standard "good medium" quality criteria [87] with $p_{\mathrm{T}} > 20$ GeV, $|\eta| < 2.5$. These two jets must be separated by $\Delta\phi > 2.5$, and the average of their p_{T} must exceed 0.4 of the p_{T} of any other jet in the event.

5.5.2 Muon Impact Parameter: Cosmic Background

A potentially troublesome source of background for Type0 and Type1 LJs consists of cosmic ray muons that cross the detector coincidentally in time with a filled bunch-crossing interaction. To reduce this background, a cut on the muon track impact parameter $|z_0| < 280$ mm is used, where $|z_0|$ is defined as the minimum distance in the z coordinate of the MS extrapolated track to the main interaction vertex. This cut is loose enough to pass most of the displaced muons, while still eliminating many of the cosmics.

In order to determine the optimized cut value, a search for Type0 LJ candidates (with the muon pre-selections discussed in Section 5.2.2 and the LJ candidate construction discussed in Section 5.2.4) was performed on the Empty data sample. It was also performed on the $4\gamma_{\mathrm{d}}$ FRVZ benchmark MC sample with $m_H = 125$ GeV; of all the FRVZ samples, this one exhibits the broadest $|z_0|$ distribution, due to the lower LJ boost. The $|z_0|$ of the muon constituents are shown in Figure 5.5; the cut $|z_0| < 280$ mm removes only 5% of the signal while eliminating 96% of the cosmic background.

The potential for the use of an alternative cut, based on muon timing, for the elimination of cosmic ray muon background is discussed in the Appendix.

5.5.3 Jet Timing: Cosmic Background

Energy deposits from hard bremsstrahlung of cosmic ray muons in the HCal can be reconstructed as jets with low f_{EM}, creating a background to the Type1 and Type2 LJ selections. The discriminating variable used to remove such cosmic jets is the jet timing Δt_{Calo}, defined as the energy-weighted mean time difference between the bunch-crossing time and the time of energy deposition in the calorimeter cells. The time of the energy deposition in the calorimeter cells is automatically corrected by presumed time-of-flight, assuming the particle that made the deposition originated from the IP and travelled at the speed of light (see Section 4.3.3). In order to determine the optimized cut value, a search for Type1 and Type2 LJ candidates – with the pre-selections discussed in Sections 5.2.2 and 5.2.3 and the LJ candidate construction discussed in Section 5.2.4 – was performed on the Empty data sample. It was also performed on the $2\gamma_{\mathrm{d}}$ FRVZ benchmark MC sample.

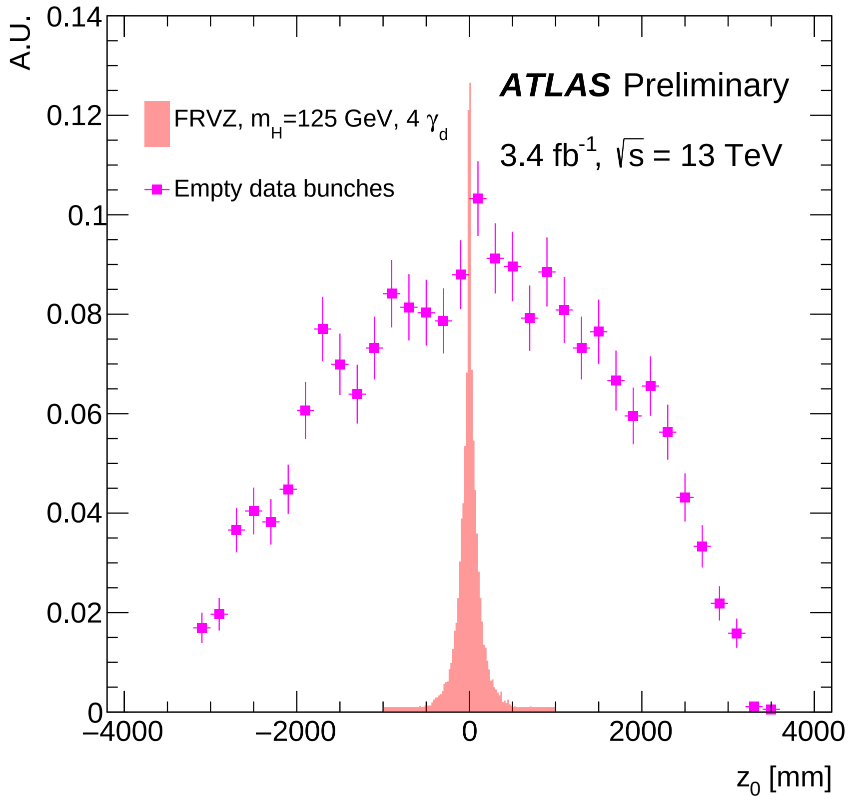

Figure 5.5: Impact parameter $|z_0|$ for the muon constituents of Type0 LJs in empty bunch-crossings (cosmic ray muons) and in the $m_H = 125$ GeV $4\gamma_\mathrm{d}$ FRVZ MC sample. Reproduced as in [6].

Figure 5.6 shows that cosmic rays tend to have a wide Δt_Calo distribution, while the signal MC Type1 and Type2 jets have a narrow peak at the bunch-crossing time ($\Delta t_\mathrm{Calo} = 0$). The width of the cosmic distribution is controlled by the time window associated with each bunch crossing, which is in turn controlled by the interval between bunch crossings. The FRVZ distribution is centred near 0, but somewhat biased toward positive values, as a result of the time-of-flight correction. Removing jets with Δt_Calo outside the interval between -4 ns and +4 ns causes $< 5\%$ loss of signal.

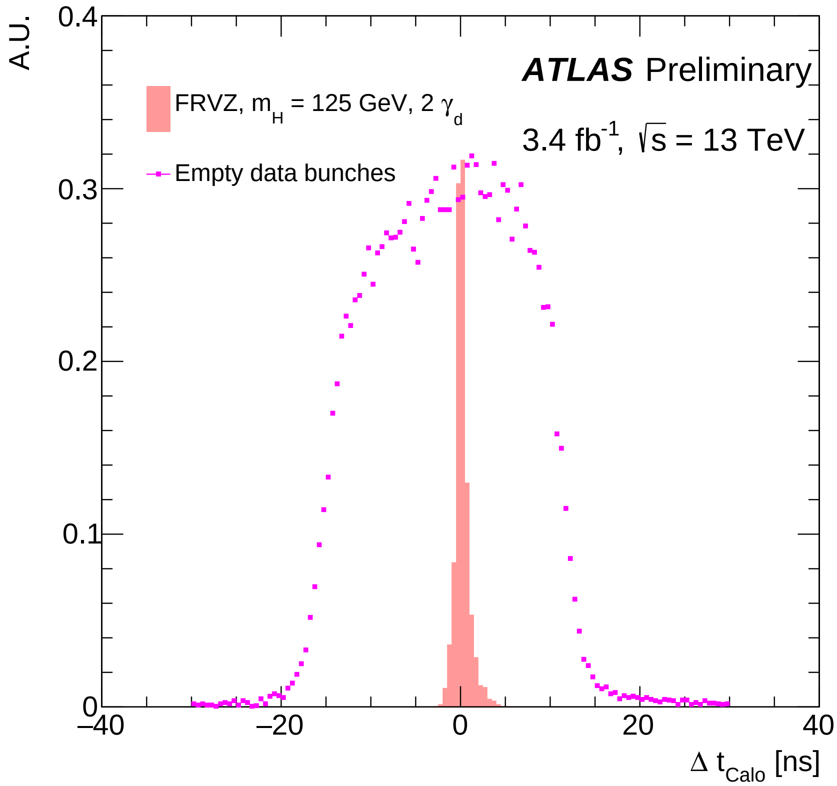

Figure 5.6: Jet timing Δt_{Calo} for the jet constituents of Type1 and Type2 LJ candidates in empty bunch-crossings, i.e. cosmic rays, and in the m_H = 125 GeV $2\gamma_{\mathrm{d}}$ FRVZ MC sample. Reproduced as in [6].

5.5.4 EM Fraction, Width, and JVT: Jet Background

The main background source for Type2 LJs is the QCD production of multi-jet events. As already discussed, Type2 LJs have unusually low f_{EM}, which can serve as a discriminating variable. Two additional discriminating variables are employed to reduce the multijet background: jet width (W) and Jet Vertex Tagging (JVT).

The jet width W is defined as the average distance of a jet constituent to the jet direction, weighted by the p_{T} of the constituents:

$$W = \frac{\sum_i \Delta R^i \cdot p_{\mathrm{T}}^i}{\sum_i p_{\mathrm{T}}^i} \tag{5.1}$$

where $\Delta\mathrm{R}^{\mathrm{i}} = \sqrt{(\Delta\phi_{\mathrm{i}})^2 + (\Delta\eta_{\mathrm{i}})^2}$ is the distance between the jet axis and the i^{th} jet constituent, and p_{T}^i is the constituent p_{T} with respect to the beam axis.

The JVT is designed to reject pile-up jets while retaining jets coming from the primary vertex (PV) of the event. The PV is defined as the vertex whose associated tracks have the largest Σp_{T}^2. Due to their displacement, the Type2 LJs mimic pile-up jets in this regard, so the JVT is applied in the opposite manner: to reject jets coming from the PV while retaining Type2 LJs. JVT is calculated by a standard ATLAS analysis tool [90] that matches ID tracks to the jet, identifies which of these tracks originate from the PV, and then constructs a likelihood for the jet to be associated with the PV. The likelihood uses the scalar sum of the p_{T} of all the tracks matched to the jet, the scalar sum of the p_{T} of the matched tracks that originate from the PV, the calibrated jet p_T, and corrections to account for pile-up dependence. Displaced jets, i.e. jets whose matched tracks do not originate from the PV, are assigned a low likelihood and hence a low JVT. Jets that do not have any matched tracks are assigned an arbitrary negative JVT value.

Cut Optimization

To study the QCD multi-jet background and to optimize the cuts on the above discriminating variables, a Tag and Probe method is used with the dijet control sample. The leading jet in the event is considered the "tag" jet, while the sub-leading jet is used as the "probe". Only events that have both a tag and a probe jet, with $E_{\mathrm{T}} \geq 20$ GeV and $\eta \leq 2.4$, are considered. An additional cut of $|\Delta\phi| \geq 2.5$ between the tag and the probe is also applied, to select back-to-back topologies similar to those seen in FRVZ events.

To optimize the cut values, the discriminating variables are ordered according to their separation power, SP, defined as [91]:

$$SP = \sqrt{\frac{1}{2}\int \frac{(p_S(x) - p_B(x))^2}{(p_S(x) + p_B(x))} dx}$$

where $p_S(x)$ and $p_B(x)$ are the signal and background probability density functions, respectively, of a given discriminating variable x.

The cut on the first (most discriminating) variable is then chosen to maximize the significance S, defined as [92]

$$S = \sqrt{2 \times ((N_S + N_B) \times \log(1 + N_S/N_B) - N_S)} \tag{5.2}$$

where N_S and N_B are the numbers of signal and background jets with the first variable less than a given cut. This significance measure is based upon the likelihood for signal strength μ,

$$L(\mu) = \frac{(\mu N_S + N_B)^N}{N!} \exp -(\mu N_S + N_B)$$

where $N = N_S + N_B$ is the total number of events. The test statistic for discovery is then $-2\ln[L(0)/L(\mu)]$; the nominal signal hypothesis ($\mu = 1$) is taken in the "asymptotic" limit of large N_B.

Then, the distributions of the other variables are re-calculated after applying the cut on the first one, and the procedure is repeated until the cuts for all variables are obtained.

According to the above procedure, the variables are ordered with their optimal cuts as:

1. $f_{EM} < 0.1$
2. $W < 0.058$
3. JVT < 0.56

In Figures 5.7, 5.8 and 5.9, the distributions of the discriminating variables are shown for MC FRVZ signal jets, and for jets from the di-jet data control sample. In each of these Figures, the corresponding significance variable S is also shown. The cut value is imposed where S reaches a maximum. In the case of JVT, S reaches a maximum plateau, with small fluctuations on the plateau; these fluctuations are purely statistical. To avoid dependence on the statistical fluctuations, the cut is imposed at the low edge of the plateau rather than strictly at the maximum.

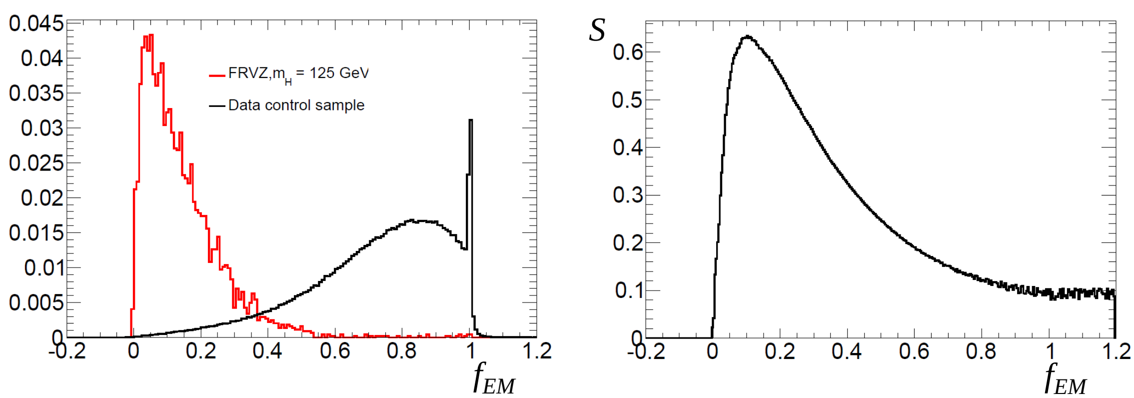

Figure 5.7: Left: f_{EM} for MC FRVZ signal jets (red), and for jets from the di-jet data control sample (black). Right: the significance variable S defined in Equation 5.2.

MVA Optimization

As a cross-check to the above optimization procedure, I employed the Rectangular Cut Optimization method for MultiVariate Analysis in the TMVA (version 4.2.0) package [91] for ROOT. Input to the TMVA was the FRVZ $2\gamma_{\rm d}$ $m_H = 125$ GeV MC sample, and a subset of approximately 50000 events from the di-jet data control sample. Monte Carlo sampling was used for the multivariate analysis fitter method. The JVT, W, and f_{EM} were the only input variables, along with the MC event weights. The TMVA output yields the same ranking of the variables by separation power, and very similar cut values for optimal significance (Table 5.4).

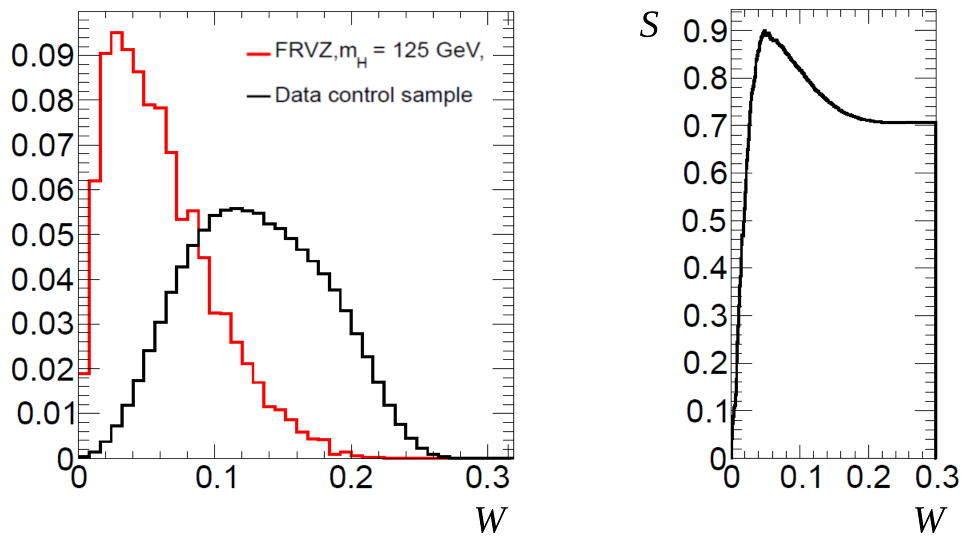

Figure 5.8: Left: W for MC FRVZ signal jets (red), and for jets from the di-jet data control sample (black). Right: the significance variable S defined in Equation 5.2.

Variable	Separation Power	Cut Value
EM Fraction	0.81	0.11
Jet Width	0.40	0.05
JVT	0.34	0.56

Table 5.4: TMVA ranking of jet discriminating variables by separation power, and cut values for optimal significance.

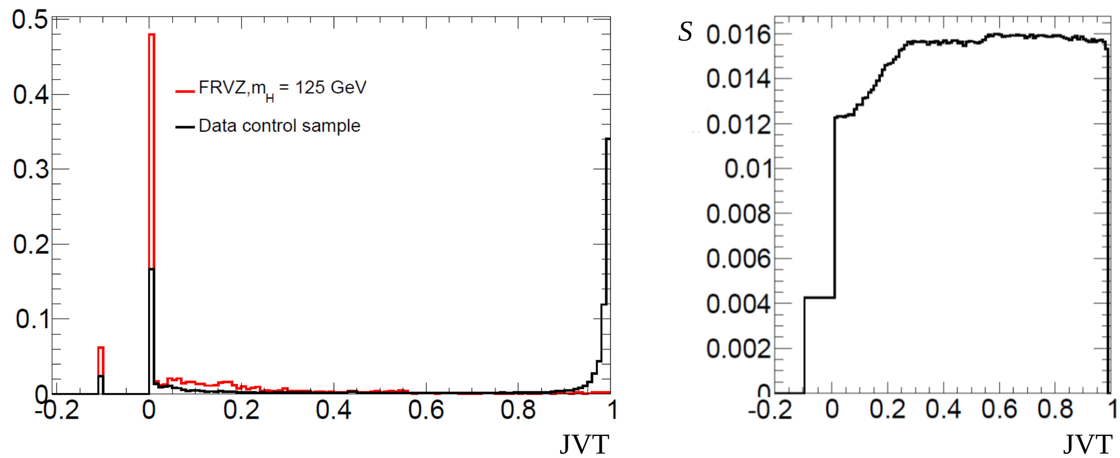

Figure 5.9: Left: JVT for MC FRVZ signal jets (red) and for jets from the di-jet data control sample (black). Right: the significance variable S defined in Equation 5.2.

5.5.5 Beam Induced Background: Jet Background

A potential source of fake Type2 LJs is Beam Induced Background (BIB), i.e. a catastrophic energy loss of a high energy beam-background muon in the material of the calorimeters, manifesting itself as a high-energy calorimeter deposit that is reconstructed as a jet. If this bremsstrahlung occurs in the HCal, the resulting jet has a very low f_{EM} and therefore resembles a Type2 LJ. If the bremsstrahlung is not hard enough to stop the muon entirely, the muon keeps travelling into the opposite side of the MS from which it began. A clear example of this type of event from actual data is shown in Figure 5.10.

The CalRatio trigger (see Section 5.4), employs an implementation of the "Two-Sided No-Time Method", developed by the Non-Collision Backgrounds group, in order to tag and eliminate such fake Type2 LJs. I acted as a liaison between the Exotics group and the Non-Collision Backgrounds group, creating and validating a tool to implement this method at analysis level. The method identifies a BIB fake jet using the muon segments, parallel to the beam pipe and matched to the jet in ϕ, that generally accompany it. In the upper pane of Figure 5.10, the muon segments on either side of the MS are shown in the insets, while the large energy deposit in an HCal cell corresponding to the fake jet is shown in the main display. In the LJ case, since the problematic jet fake must strike the calorimeters, the relevant muon segments show up in the endcap MDTs. The lower pane of Figure 5.10 shows the energy deposits in the HCal and in the endcap MDTs all lining up parallel to the z axis.

For each jet, the method searches through all the jet constituents to find high-energy "candidate BIB clusters". It also searches through all the muon segments in that event, to find "candidate BIB segments". It then applies matching criteria between each candidate BIB cluster and each candidate BIB segment. If a candidate BIB cluster matches at least one candidate BIB segment on both sides of the detector, the jet to which the cluster belongs is tagged as BIB. The exact criteria for candidate selection and matching, given in Table 5.5, come directly from the NCB group's recommendations.

Candidate BIB cluster	Calorimeter jet High-energy calo cluster	Constructed from high-energy calo cluster ECal energy + HCal energy > 10 GeV
Candidate BIB segment	In MDTs Along beampipe	Module is EIL or EIS $(\theta_{global\ position} - \theta_{momentum}) < 10°$
Segment-cluster matching	Close in ϕ Compatible in r	$(\phi_{cluster} - \phi_{segment}) > 4°$ $\lvert r_{cluster} - r_{segment} \rvert < 800$ mm

Table 5.5: Criteria employed in the analysis implementation of the Two-Sided No-Timing Method for tagging fake Type2 LJs from BIB. See Section 4.4 for a description of the MDTs and the EIL and EIS modules in the MS.

5.5.6 The HCal transition region: Jet Background

A source of background jets with artificially low f_{EM} consists of jets reconstructed in the transition region between the barrel and end-cap cryostat, where the HCal consists of only tile-gap or cryostat scintillators (see Section 4.3.2). For jets with $f_{EM} < 0.2$ in the di-jet control sample, Figure 5.11 (blue distribution) shows two sharp peaks. These correspond to jets reconstructed mostly from the small energy released in the tile gap and cryostat scintillators, which are therefore significantly affected by noise contributions. They are removed by requiring the fraction of energy deposited in the HCal region $1.0 < |\eta| < 1.6$, corresponding to the tile gap and cryostat scintillators, to be less than 10% of the total jet energy (red distribution).

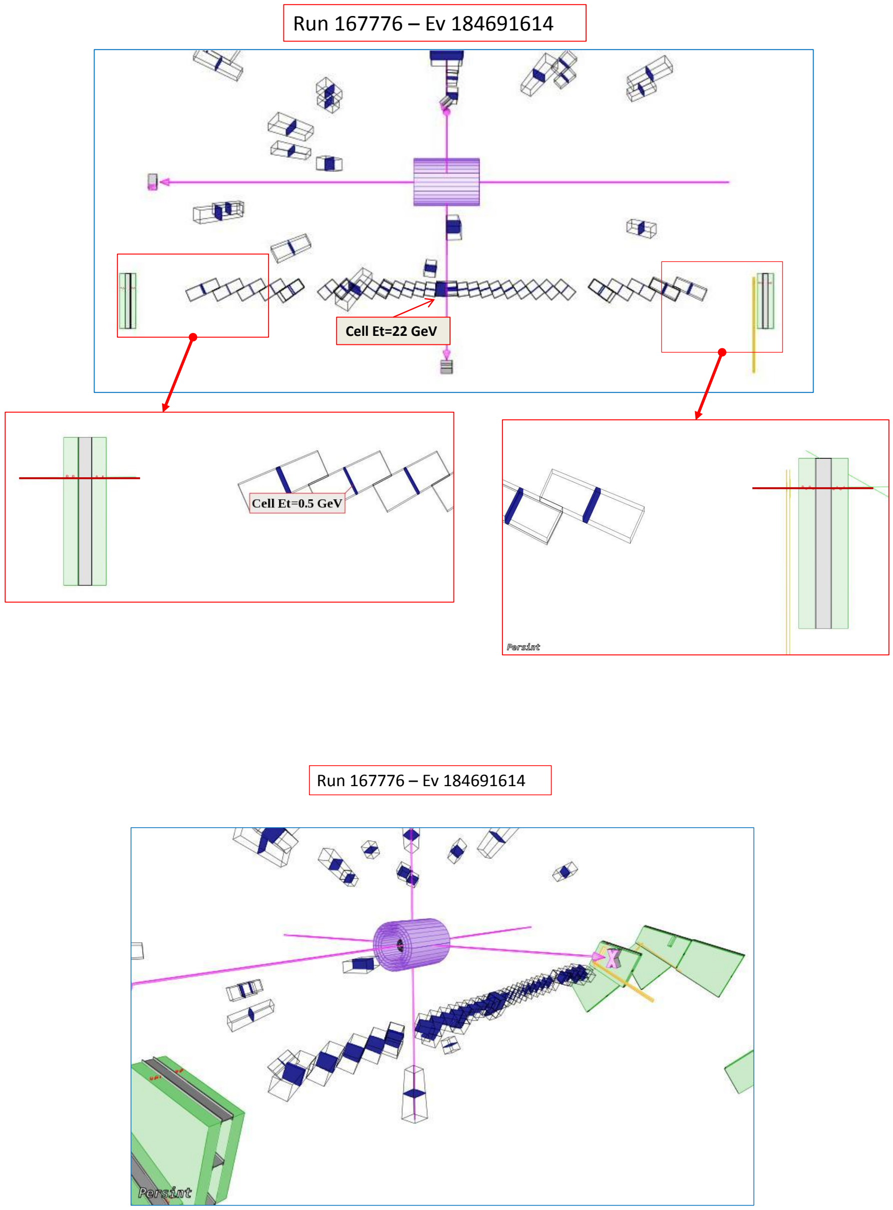

Figure 5.10: Event display from the Persint visualization program [93], illustrating an example of BIB mimicking a Type2 LJ. The display shows only detector modules with significant energy deposits (dark blue) in one of their elements. Top: energy deposits are seen in the forward MDT chambers (insets), together with a large energy deposit in an HCal cell (highlighted in the central display). Bottom: Rotated view, showing the energy deposits in the HCal and in the MDT chambers forming a diagonal line of dark blue rectangles (parallel to the z axis) in the lower part of the display.

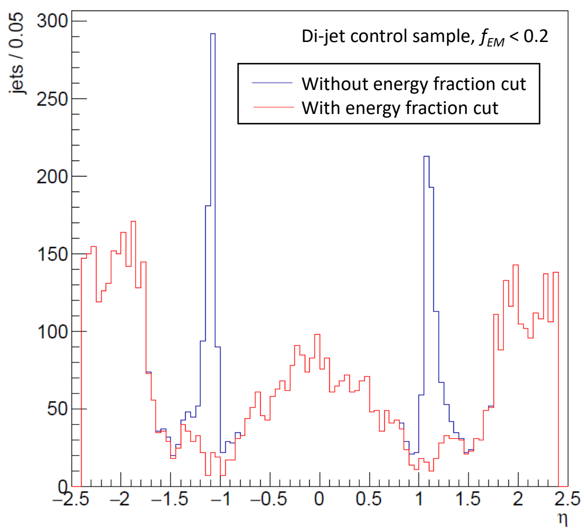

Figure 5.11: η distribution of jets in the di-jet control sample with $f_{EM} < 0.2$, without (blue) and with (red) the requirement that the energy fraction released in the tile gap and cryostat scintillators must be less than 0.1.

5.6 Event Selection

The analysis selects only events with at least two LJs that pass all the above criteria. Cuts on two additional quantities, the LJ isolation in the ID and the $|\Delta\phi|$ between the two LJs, are then imposed at event level. The cut values are obtained through a data-driven technique that also produces the residual background estimates. This section provides a description of these event-level quantities, and then of the data-driven technique. The actual results of the data-driven technique will be shown in Section 6.3.

5.6.1 Isolation in Inner Detector

The displaced LJs are expected to be highly isolated in the ID. Therefore the multi-jet background can be significantly reduced by requiring track isolation around the LJ direction in the ID. The ID isolation variable $\Sigma p_{\mathrm{T}}^{\mathrm{ID}}$ is defined as the sum of the transverse momenta of all ID tracks with $p_{\mathrm{T}} > 500$ MeV, matched to the PV of the event, within a cone of $\Delta R = 0.5$ around the LJ. The requirement of matching to the PV helps reduce the pileup dependence of $\Sigma p_{\mathrm{T}}^{\mathrm{ID}}$. The identification of the primary vertex is performed by a standard ATLAS analysis "vertexing tool", the selection of ID tracks is performed with the help of a standard "track selection tool", and the $\Sigma p_{\mathrm{T}}^{\mathrm{ID}}$ is calculated using a standard "isolation tool".

It is expected that the $\Sigma p_{\mathrm{T}}^{\mathrm{ID}}$ of displaced muonic LJs should behave in the same way as would the $\Sigma p_{\mathrm{T}}^{\mathrm{ID}}$ of prompt muonic LJs, if in the latter case the p_{T} contributions of ID tracks of muon constituents of the LJs are subtracted out. In order to validate that the $\Sigma p_{\mathrm{T}}^{\mathrm{ID}}$ of displaced muonic LJs is behaving properly, it is useful to checking whether this expectation in fact holds. For this purpose, a search for LJs was performed on samples of $Z \to \mu\mu$ events in MC and in 2015 data [1] , and the $\Sigma p_{\mathrm{T}}^{\mathrm{ID}}$ calculated

[1] Triggers: loose-quality muon, $p_{\mathrm{T}} > 20$ GeV OR standard muon, $p_{\mathrm{T}} > 50$ GeV. Selection: exactly one opposite-sign lepton pair, invariant mass 66 to 116 GeV, each muon $p_{\mathrm{T}} > 25$ GeV and $|\eta| < 2.4$, each muon passing standard quality criteria with additional calorimeter isolation and ID isolation criteria.

for these “ersatz prompt” LJs. To study the behaviour of $\Sigma p_{\mathrm{T}}^{\mathrm{ID}}$ in background events, a LJ search with the pre-selections discussed in Sections 5.2.2 and 5.2.3 was performed on the di-jet control sample and on the Empty sample, and the $\Sigma p_{\mathrm{T}}^{\mathrm{ID}}$ calculated for these background LJs. As shown in Figure 5.12, the $\Sigma p_{\mathrm{T}}^{\mathrm{ID}}$ agreement between the ersatz prompt LJs and the Type0 LJs from the FRVZ benchmark samples is indeed fairly good, which enables the analysis to use the $Z \to \mu\mu$ distributions for validating the signal modelling in this variable. As also shown in the Figure, the signal distribution indeed differs drastically from the di-jet distribution, although not so drastically from the Empty distribution. This indicates the $\Sigma p_{\mathrm{T}}^{\mathrm{ID}}$ is useful as a discriminating variable against the multi-jet background, although not against the cosmic background.

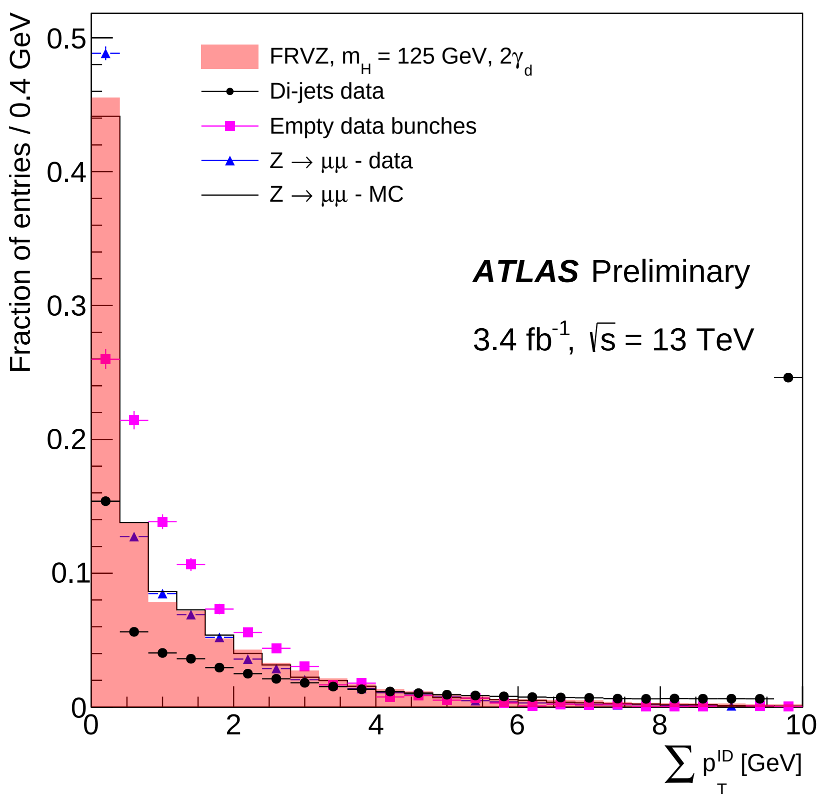

Figure 5.12: The distributions of $\Sigma p_{\mathrm{T}}^{\mathrm{ID}}$ for the di-jet data control sample (black circles), $Z \to \mu\mu$ data control sample (blue triangles), and Empty sample (violet squares). All distributions are normalized to unit area; the overflow events for the di-jet data control sample ($p_{\mathrm{T}} > 10$ GeV) are filled in the last interval. Reproduced as in [6].

As an event-level variable for the analysis, $\max(\Sigma p_{\mathrm{T}}^{\mathrm{ID}})$ is used, where max indicates the highest $\Sigma p_{\mathrm{T}}^{\mathrm{ID}}$ from amongst the LJs of the event. The criterion $\max(\Sigma p_{\mathrm{T}}^{\mathrm{ID}}) < 4.5$ GeV, with the cut value determined by the data-driven optimization performed in Section 6.3, greatly reduces the multijet background while having little effect on the LJ signal selection efficiency.

5.6.2 LJ separation in the azimuthal plane

In the FRVZ models, the two LJs in each event are expected to be back-to-back in the azimuthal (ϕ) plane. This can be seen from the distribution of $|\Delta\phi|_{\rm LJ}$, i.e. $|\Delta\phi|$ between the two reconstructed LJs, in the FRVZ $2\gamma_{\rm d}$, $m_H = 125$ GeV MC sample (Figure 5.13). The distributions from the di-jet control sample and the Empty sample are also shown in the Figure. A cut $|\Delta\phi|_{\rm LJ} > 0.63$ rad, with the cut value determined by the data-driven optimization performed in Section 6.3, significantly reduces the multijet and cosmic backgrounds without significantly affecting the signal.

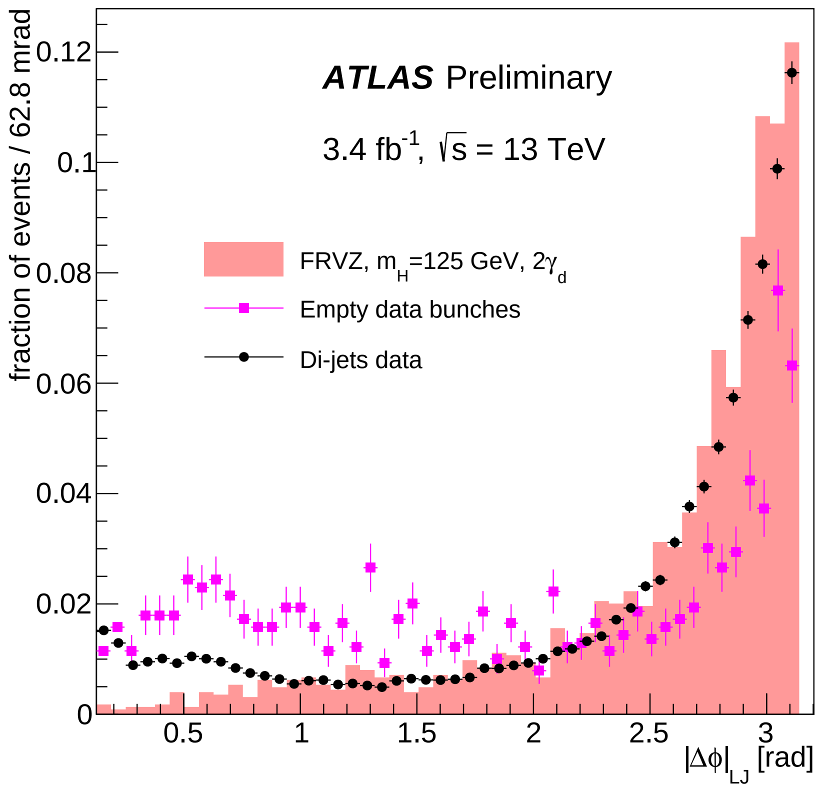

Figure 5.13: $\Delta\phi$ between the two reconstructed LJs from the FRVZ $2\gamma_{\rm d}$, $m_H = 125$ GeV MC sample. For comparison, the distributions from the Empty sample and from the di-jet control sample are also shown. The uncertainties are statistical only. Reproduced as in [6].

5.6.3 Data-Driven Cut Optimization

The backgrounds that remain after the LJ selection are QCD multi-jet production and cosmics. In order to optimize the cuts on the two event-level quantities $\max(\Sigma p_{\rm T}^{\rm ID})$ and $|\Delta\phi|_{\rm LJ}$ against these backgrounds, a data driven "ABCD" method is used. This is a simplified matrix method that serves as a simultaneous background estimation and signal hypothesis test in signal and control regions. It is implemented according to the ATLAS Statistics Forum ABCD recommendations. A full RooFit workspace was developed to perform both the cut optimizations and the background estimations.

The events distribution after the LJ selection stage is factorized in the ($\max(\Sigma p_{\mathrm{T}}^{\mathrm{ID}})$, $|\Delta\phi|_{\mathrm{LJ}}$) plane. The cuts on these two variables subdivide the plane into four rectangular regions, A B C D. The signal region is A; the others are control regions. Neglecting the possibility of signal contamination in the control regions regions, and assuming the two variables are uncorrelated, the number of background events in the signal region can be evaluated as:

$$n_A = n_D \times \frac{n_B}{n_C} \tag{5.3}$$

Statistical Details

The likelihood function used for the simultaneous fit is the product of the four likelihoods for the counting experiments in the A, B, C and D regions:

$$L(n_A, n_B, n_C, n_D|\mu, \theta_\mu) = \prod_{i=A,B,C,D} \frac{e^{-N_i} N_i^{n_i}}{n_i!};$$

μ is the parameter of interest, defined as the signal strength, i.e. the ratio between the observed and expected number of signal events in the A region. θ_μ indicate the set of nuisance parameters used to describe systematic uncertainties and parameters describing the different backgrounds. Then:

$$\begin{aligned}
N_A &= \mu N^S_{A,\,expected} + N^B_A; \\
N_B &= \mu\epsilon_B + N^B_B \tau_B; \\
N_C &= \mu\epsilon_C + N^B_C \tau_B \tau_D; \\
N_D &= \mu\epsilon_D + N^B_D \tau_D
\end{aligned}$$

with ϵ_i signal scaling factors estimated from simulation used to describe small leakages of the signal into the background control regions, and τ_B and τ_D nuisance parameters. N^S denotes number of signal events and N^B denotes number of background events; the number of background events in the signal region is related to the number of events in the control regions via the ABCD method ansatz A = BD/C.

The full likelihood described above is multiplied with additional Gaussian probability distribution functions to constrain the nuisance parameters. Best estimates of the nuisance parameters and of the parameter of interest are obtained using profile likelihood fits, and p-values and signal limits are derived using the CL_s procedure as described in [94]. This is all implemented according to the ATLAS Statistics Forum frequentist limit recommendations.

5.7 Summary of Selection Criteria

The complete set of selection criteria, composed of the LJ selection and event-level selection, is given in Table 5.6. For Type0 and Type1 LJs, the requirement that each muon constituent must be Non-Combined (see Section 5.2.2) has been moved from the pre-selection stage to the end of the LJ selection stage, for the purposes of increasing the statistics throughout the selection so that systematic uncertainties can be more accurately determined. Since the data-driven optimization of the cut values in the event selection stage is un-blinded, the actual cut values are left to the next chapter.

LJ Type	Criterion	Description
Type0 / 1	Impact parameter	Requires $\|z_0\| < 280$ mm for each muon constituent
Type1 / 2	Jet timing	Removes jets outside $\pm$ 4 ns time window
Type2	HCal gap	Requires energy deposited in tile gap and cryostat scintillators $\leq$ 10% of the jet energy
Type2	EM fraction	Requires jet $f_{EM} < 0.1$
Type2	Jet width	Requires jet $W < 0.058$
Type2	Jet Vertex Tagger	Requires JVT $\leq$ 0.558
Type2	BIB	Uses "Two-Sided No Time Method" to remove jets from beam halo
Type0 / 1	Veto CB muons	Require muons to be Non-Combined
Event-level	ID isolation	See Section 6.3
Event-level	Azimuthal separation	See Section 6.3

Table 5.6: Summary of the LJ selection and event-level selection criteria, reproduced as in [6].

Chapter 6

Displaced Lepton-Jet Search Results

This chapter describes the results of the LJ search, conducted according to the methodology set out in the previous chapter. The first product of the analysis is a set of LJ detection efficiencies, before the application of selection cuts. These efficiencies can be used by phenomenologists to "re-cast" this search, i.e. to test any model that predicts displaced LJs regardless of whether the rest of the model's properties resemble those of the benchmarks. Next come the results of the first stage of selection, which is at the level of individual LJ candidates. Then come the data-driven optimization and the results of the event-level selection, yielding model-independent expected and observed signal event counts. Finally, the event counts are interpreted as limits on the benchmark models, and as constraints in the more general dark photon parameter space. Discussion of systematic uncertainties is postponed to Chapter 7.

6.1 LJ Detection Efficiencies

There has been increasing demand from phenomenologists in recent years to publish detection efficiency results, to permit re-casting of searches (as discussed in Section 3.5). Given that experimentalists cannot be expected to cover the plethora of new models constantly emerging in the literature, this demand is well-justified. Therefore, the first product of the Displaced LJ search is an evaluation of acceptance times reconstruction efficiency as a function of decay distance for each type of LJ. Designated $A \times Eff$, this is defined as the ratio between the number of reconstructed LJs and the total number of truth $\gamma_{\rm d}$'s decaying to the appropriate products – muons for Type0 LJs, electrons or pions for Type 2 – in a given decay distance bin. The number of truth $\gamma_{\rm d}$'s is determined from the FRVZ MC. The number of reconstructed LJs is taken after the pre-selection, with no trigger cuts applied; in the case of Type2 LJs, the f_{EM} and jet width cuts are also applied. The $A \times Eff$ gives a measure of the ATLAS detector's ability to reconstruct displaced collimated pair(s) of muons or electrons/pions, independently of trigger limitations, and as independently as possible of specific analysis cuts. The two additional cuts for Type2 LJs are necessary in order to select displaced, rather than prompt, decays.

6.1.1 Type0 LJs

The Type0 $A \times Eff$ is shown in Figure 6.1, as a function of the $\gamma_{\rm d}$ transverse decay distance L_{xy} in the barrel ($|\eta| \leq 0.9$) and of the longitudinal decay distance L_z in the endcap, for the $m_H = 125$ GeV and the $m_H = 800$ GeV $2\gamma_{\rm d}$ FRVZ samples. In both cases, at $\lesssim 1$ m, efficiency is poor because the LJs are not sufficiently displaced. Starting at ~ 3 m, $A \times Eff$ decreases with increasing distance from the primary vertex, due to the decrease of the separation between the two muons in the innermost MDT stations. The algorithm for reconstructing particle tracks in the MS includes “overlap removal” to avoid track duplication of the same muon, and this often rejects one of the two muons if the two candidate tracks share segments in one chamber (see Section 4.4.3). The increase in $A \times Eff$ immediately after the innermost MS chambers, ~ 4.5 m in the barrel and ~ 7.5 m in the endcaps, is due to the separation of the two muons under the influence of the MS magnetic field, which increases the probability of both muons being reconstructed successfully in the face of overlap removal. The efficiency drops to zero after the second layer of MDT chambers in the barrel region (third layer in the endcap) because there are no longer enough MDT hits to reconstruct a muon track. The $A \times Eff$ for the 800 GeV sample is systematically lower than for the 125 GeV sample, due to the higher boost of the γ_{d} and the consequent decrease in the separation of the muon tracks.

The Type0 $A \times Eff$ is also shown in Figure 6.1 as a function of the p_{T} of the γ_{d}. Below ~ 40 GeV, the $A \times Eff$ rises with p_{T} as more of the γ_{d} decay products emerge with enough p_{T} to be reconstructed. Above ~ 40 GeV, the $A \times Eff$ drops with p_{T} , as the decay products become more collimated and hence more difficult to reconstruct. Figure 6.2 shows two-dimensional plots of the Type0 $A \times Eff$, as a function of $(p_{\mathrm{T}},\ L_{xy})$ and of $(p_{\mathrm{T}},\ L_z)$.

6.1.2 Type2 LJs

The Type2 $A \times Eff$ is shown in Figure 6.3 as a function of L_{xy} and of L_z, with the cuts on the jet width and EM fraction ($W \leq 0.058$, $f_{EM} \leq 0.1$) applied to the reconstructed LJs after the preselection. Due to the EM fraction cut, most of the LJs reconstructed correspond to decays in the HCal. In the barrel region of the ECal, the $A \times Eff$ increases as the decay point approaches the HCal; decays past the HCal are not reconstructed at all. The $A \times Eff$ for the 800 GeV sample is higher than for the 125 GeV sample because the decay products are produced with higher p_{T} , and hence are more likely to be reconstructed as a jet. The Type2 $A \times Eff$ is also shown in Figure 6.3 as a function of the p_{T} of the γ_{d}; it demonstrates a similar p_{T} dependence as the Type0. Figure 6.4 shows two-dimensional plots of the Type2 $A \times Eff$, as a function of $(p_{\mathrm{T}},\ L_{xy})$ and of $(p_{\mathrm{T}},\ L_z)$.

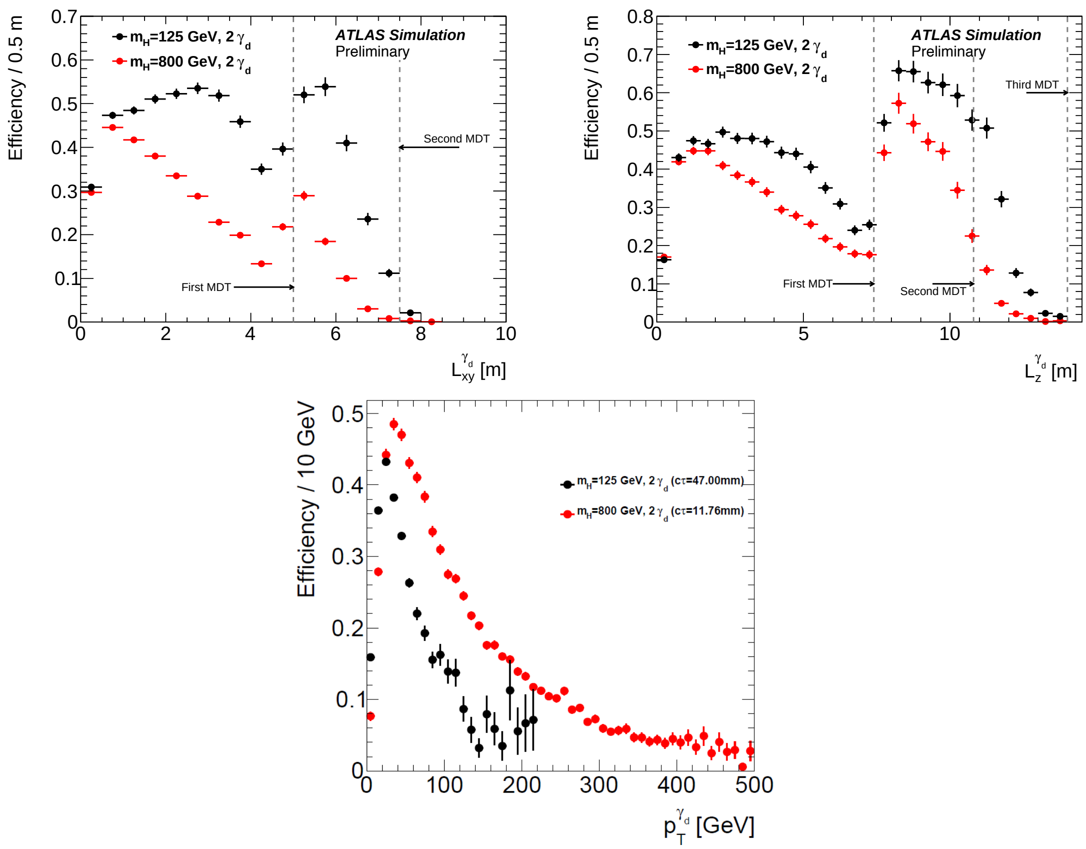

Figure 6.1: Top left: $A \times Eff$ in the barrel ($|\eta| \leq 0.9$) for Type0 LJs, as a function of the transverse decay distance L_{xy} of the γ_{d}. Reproduced as in [6]. Top right: $A \times Eff$ in the endcaps, as a function of the longitudinal decay distance L_z, where the dashed grey lines display the approximate positions of the MDT chambers. Reproduced as in [6]. Bottom: $A \times Eff$ throughout the detector, as a function of the p_{T} of the γ_{d}. Evaluated in MC samples of the $2\gamma_{\mathrm{d}}$ FRVZ benchmark model. The uncertainties are statistical only.

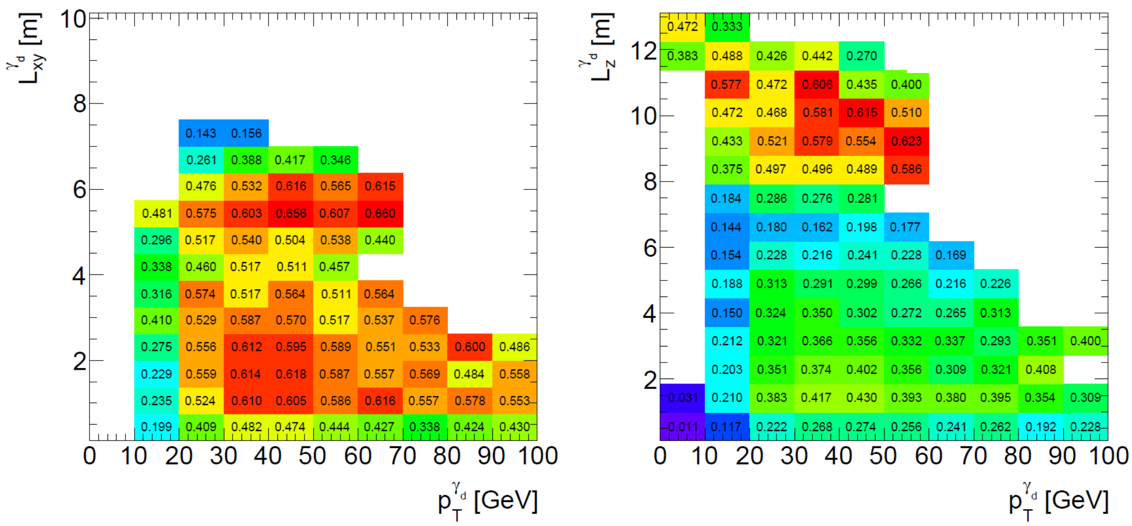

Figure 6.2: Left: $A \times Eff$ in the barrel ($|\eta| \leq 0.9$) for Type0 LJs, as a function of the transverse decay distance L_{xy} and of the p_{T} of the γ_{d}. Right: $A \times Eff$ in the endcaps, as a function of the longitudinal decay distance L_z and of the p_{T} of the γ_{d}. In the white areas, $A \times Eff = 0$, or the statistics are too low to calculate a meaningful $A \times Eff$.

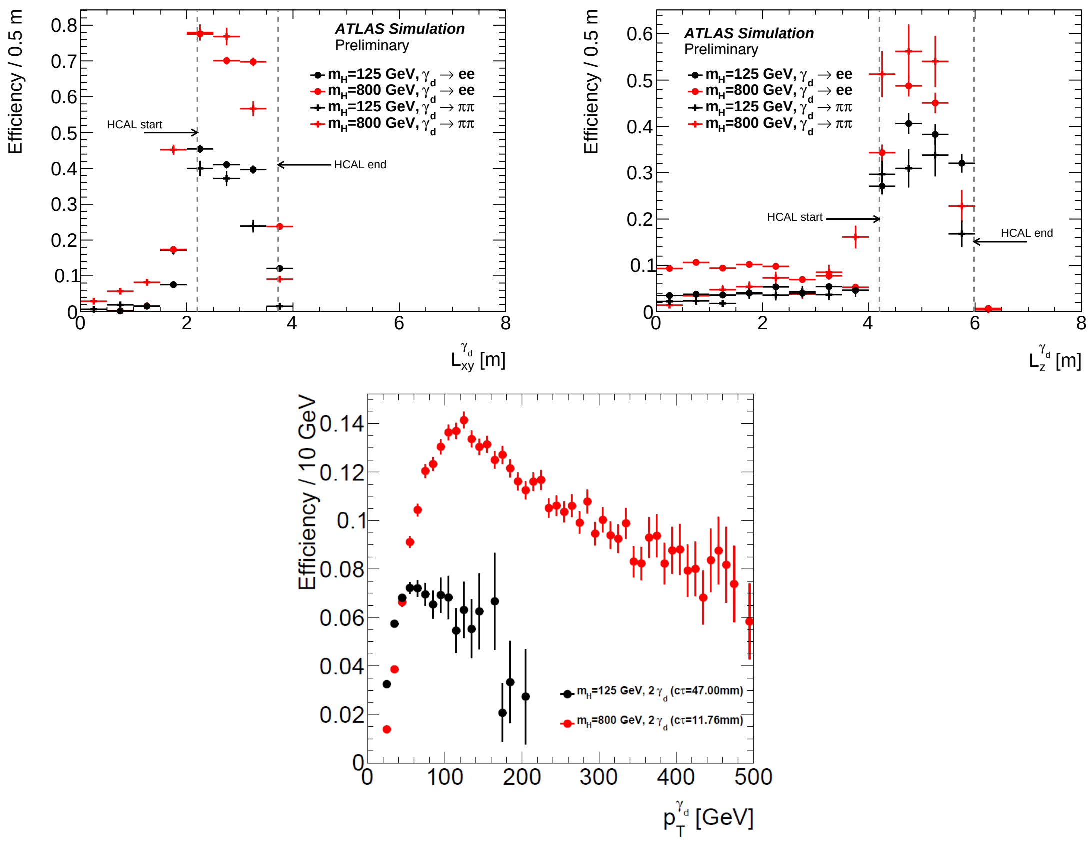

Figure 6.3: Top left: $A \times Eff$ in the barrel ($|\eta| \leq 0.9$) for Type2 LJs, as a function of the transverse decay distance L_{xy} of the γ_d. Reproduced as in [6]. Top right: $A \times Eff$ in the endcaps, as a function of the longitudinal decay distance L_z, where the grey dashed lines display the boundaries of the HCal. Reproduced as in [6]. Bottom: $A \times Eff$ throughout the detector, as a function of the p_T of the γ_d. The uncertainties are statistical only.

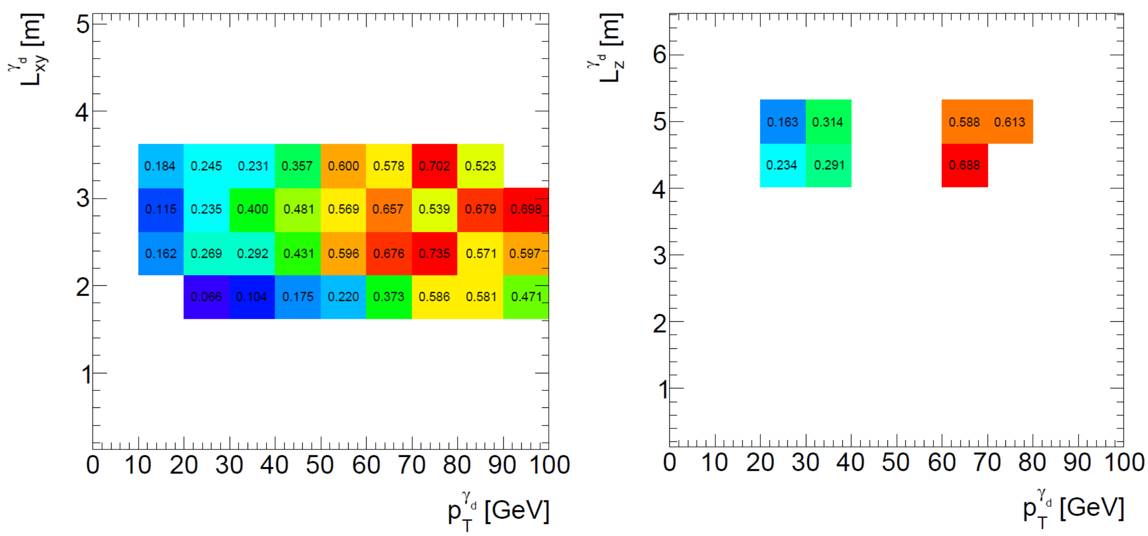

Figure 6.4: Left: $A \times Eff$ in the barrel ($|\eta| \leq 0.9$) for Type2 LJs, as a function of the transverse decay distance L_{xy} and of the p_{T} of the γ_{d} . Right: $A \times Eff$ in the endcaps, as a function of the longitudinal decay distance L_z and of the p_{T} of the γ_{d}. In the white areas, $A \times Eff = 0$, or the statistics are too low to calculate a meaningful $A \times Eff$.

6.2 LJ Selection: Cut Flow

This section summarizes the results of the cut-flow up to the end of the LJ selection described in Section 5.5; the next section will discuss the event-level selection. Given the high volume of MC and data involved, running the analysis code on all the samples required the combined efforts of all members of the Displaced LJs analysis team. In order to properly compare MC to data, the total number of events in each MC sample is rescaled to the production cross-section for the sample processes multiplied by the total integrated luminosity in the 2015 data sample. For the 125 GeV Higgs boson, the SM gg fusion production cross section (43.87 pb at $\sqrt{s}$ = 13 TeV), as given in [67], is always assumed. The uncertainties on event numbers in this section are statistical only.

6.2.1 FRVZ Monte Carlo

The results of the cut flow applied to the $2\gamma_\mathrm{d}$ and $4\gamma_\mathrm{d}$ FRVZ MC samples, for $m_H = 125$ GeV, are shown in Table 6.1. Assuming a 10% BR of the 125 GeV Higgs boson to the hidden sector, 138 ± 3 events in the $2\gamma_\mathrm{d}$ model and 117 ± 2 events in the $4\gamma_\mathrm{d}$ model would be expected to survive the LJ selection stage.

Cut	$2\gamma_\mathrm{d}$ FRVZ MC	4 γ_d FRVZ MC
Trigger selection	1858 ± 10	1836 ± 10
Preselection	382 ± 5	427 ± 5
z_0	334 ± 4	302 ± 4
Jet timing	261 ± 4	237 ± 4
HCal gap	226 ± 3	211 ± 3
EM fraction	153 ± 3	141 ± 3
Jet width	146 ± 3	134 ± 3
JVT	145 ± 3	133 ± 3
BIB removal	145 ± 3	133 ± 3
CB muon veto	138 ± 3	117 ± 2

Table 6.1: Number of events surviving each stage of the LJ selection on the two $m_H = 125$ GeV FRVZ MC samples. Assumes BR($H \to \gamma_\mathrm{d}$'s)= 0.1. Numbers, and associated uncertainties, scaled to the integrated luminosity of the data sample.

The results of the cut flow applied to the FRVZ models with $m_H = 800$ GeV are shown in Table 6.2, assuming a cross section of 1.0 pb for each of the two benchmark processes. This σ was not chosen according to any particular physics motivation, but rather to allow easy re-scaling to any predicted values. Then 70 ± 1 events in the $2\gamma_\mathrm{d}$ model and 155 ± 2 events in the $4\gamma_\mathrm{d}$ model would be expected to survive the LJ selection stage.

6.2.2 SM Background Monte Carlo

Additional potential backgrounds to the Displaced LJ signal, apart from cosmics, BIB, and QCD multijet production, include all the SM processes which lead to real prompt muons with or without hadronic jets in the final state. To verify that these additional potential backgrounds do not actually make significant contributions, the trigger, pre-selection, and LJ selection criteria were applied to official ATLAS MC samples of the following SM processes:

- single boson (W/Z) + jets

Cut	$2\gamma_d$ FRVZ MC	$4\ \gamma_d$ FRVZ MC
Trigger selection	1192 ± 4	1631 ± 5
Preselection	234 ± 2	467 ± 3
z_0	218 ± 2	399 ± 2
Jet timing	176 ± 2	338 ± 2
HCal gap	146 ±1	303 ± 2
EM fraction	85 ± 1	203 ±2
Jet width	83 ± 1	198 ± 2
JVT	81 ± 1	196 ± 2
BIB removal	81 ± 1	196 ± 2
CB muon veto	70 ± 1	155 ± 2

Table 6.2: Number of events surviving each stage of the LJ selection on the two $m_H = 800$ GeV FRVZ MC samples. Assumes a cross section of 1.0 pb for each of the two benchmark processes. Numbers, and associated uncertainties, scaled to the integrated luminosity of the data sample.

- single-top
- $t\bar{t}$
- Other heavy-flavour ($b\bar{b}$, $c\bar{c}$)
- diboson ($WW,\ WZ,\ ZZ$)

Pile-up is included in all these samples. No events survive the LJ selection; as an example, the cut-flow on the $t\bar{t}$ sample is shown in Table 6.3.

Cut	Number of Events
Preselection	33400
z_0	32300
Jet timing	24683
HCal gap	19046
EM fraction	1903
Jet width	1726
JVT	1443
BIB removal	1443
CB muon veto	0

Table 6.3: Results of the LJ selection on the $t\bar{t}$ MC sample. Numbers scaled to the integrated luminosity of the data sample.

6.2.3 2015 Data

Data used for this analysis were collected during the entire 2015 data-taking period, including only runs in which all the ATLAS subdetectors were running at nominal conditions as defined in the standard ATLAS Good Run List. The total integrated luminosity corresponds to 3.57 fb^{-1} with a 2.1% uncertainty, calculated as per the standard methodology detailed in [95]. This calculation relies upon a preliminary calibration of the luminosity scale using a pair of xy beam-separation scans performed in June 2015. In this data, 1.68×10^7 events pass the trigger selection and have a valid primary vertex. The presence of a valid PV is irrelevant to the presence of final-state LJs, but must be required at the analysis level in order for quantities such as $\Sigma p_{\text{T}}^{\text{ID}}$ to be valid.

The results of the cut-flow through the LJ selection on this data sample are shown from Tables 6.4 to 6.8 for each channel, i.e. each possible final-state pair combination of LJ types. The list of cuts for each LJ type is given in Table 5.6; the list of cuts for each channel is simply the combination of the cuts for the relevant types.

Cut	Number of Events, Type0-Type0
Preselection	3375
z_0 cut	1949
Veto CB muons	13

Table 6.4: Result of the cut flow, through the LJ selection, on the 2015 data sample for the Type0-Type0 channel.

Cut	Number of Events, Type0-Type1	Number of Events, Type1-Type1
Preselection	6491	12436
z_0	4852	9024
Jet timing	4563	8498
Veto CB muons	1	1

Table 6.5: Result of the cut flow, through the LJ selection, on the 2015 data sample for the Type0-Type1 and Type1-Type1 channels.

Cut	Number of Events, Type0-Type2	Number of Events, Type1-Type2
Preselection	326208	822330
z_0	277424	719698
Jet timing	176163	481729
HCal gap	110969	332028
EM fraction	5154	12958
Jet width	2732	7086
JVT	863	1978
BIB removal	862	1975
Veto CB muons	61	8

Table 6.6: Result of the cut flow, through the LJ selection, on the 2015 data sample for the Type0-Type2 and Type1-Type2 channels.

Cut	Number of Events, Type2-Type2
Preselection	412993
Jet timing	129628
HCal gap	53630
EM fraction	1358
Jet width	686
JVT	347
BIB removal	335

Table 6.7: Result of the cut flow, through the LJ selection, on the 2015 data sample for the Type2-Type2 channel.

Cut	Number of Events, all Types
Preselection	1583830
z_0	1425940
Jet timing	800594
HCal gap	509701
EM fraction	32544
Jet width	23578
JVT	16262
BIB removal	16246
Veto CB muons	419

Table 6.8: Result of the cut flow, through the LJ selection, on the 2015 data sample for all channels combined.

6.3 Data-Driven Cut Optimization for Event-Level Selection, and Background Estimation

The "ABCD method", introduced in Section 5.6.3, was performed on the event-level variables $\max(\Sigma p_{\mathrm{T}}^{\mathrm{ID}})$ and $|\Delta\phi|_{\mathrm{LJ}}$ to optimize the cut values and estimate the residual backgrounds. Figure 6.5 shows the ABCD plane for the 2015 data after the LJ selection stage. The lines dissecting the plane correspond to the results of the cut optimization defining the signal region: $\max(\Sigma p_{\mathrm{T}}^{\mathrm{ID}}) \leq 4.5$ GeV, and $|\Delta\phi|_{\mathrm{LJ}} \geq 0.628$ rad.

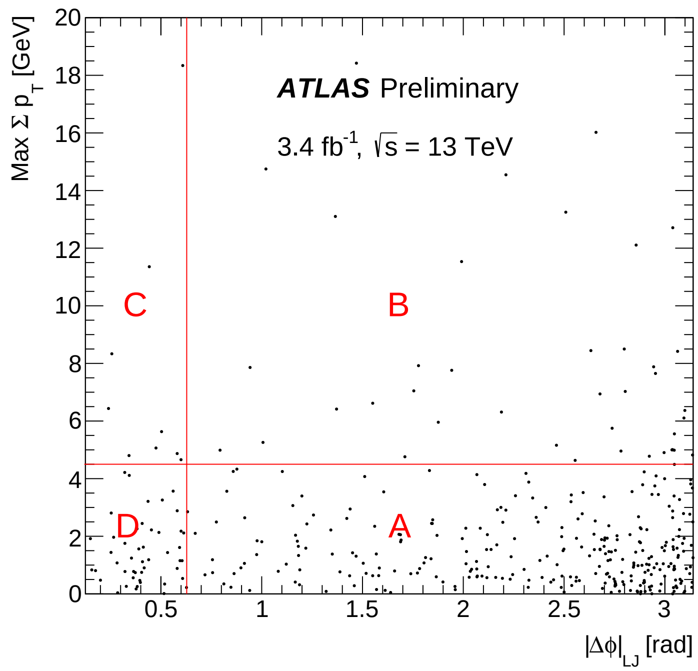

Figure 6.5: The ABCD method in the ($\max(\Sigma p_{\mathrm{T}}^{\mathrm{ID}})$, $|\Delta\phi|_{\mathrm{LJ}}$) plane, in 2015 data after the LJ selection stage. The lines dissecting the plane show the results of the cut optimization defining the signal region A: $\max(\Sigma p_{\mathrm{T}}^{\mathrm{ID}}) \leq 4.5$ GeV, and $|\Delta\phi|_{\mathrm{LJ}} \geq 0.628$ rad. Reproduced as in [6].

6.4 Final Results

6.4.1 Model-Independent

Table 6.9 shows the number of events in the 2015 data sample surviving the LJ selection and event-level selection, as well as the residual backgrounds estimated using the data-driven ABCD method, for each channel. Table 6.10 shows the total observed number of events in the signal region, and the expected residual background, for all channels combined. Since the background for the Type2–Type2 channel is expected to be much higher than for the other channels, the results are also shown with this channel excluded, and for this channel only.

Channel	LJ Selection	Event-Level Selection	Residual Background
Type0 - Type0	13	9	2 ± 3
Type0 - Type1	1	0	0 ± 0.5
Type0 - Type2	61	37	22 ± 7
Type1 - Type1	1	0	$(0.3 \pm 8.1) \times 10^{-5}$
Type1 - Type2	8	0	0 ± 0.4
Type2 - Type2	335	239	241 ± 90

Table 6.9: Number of events surviving the LJ selection and event-level selection on the 2015 data sample, as well as residual backgrounds estimated using the data-driven ABCD method, for the various LJ pair types.

Category	Observed events	Expected background
All channels	285	231 ± 12 (stat) $\pm$ 62 (syst)
Type2-Type2 channel excluded	46	31.8 ± 3.8 (stat) $\pm$ 8.6 (syst)
Type2-Type2 channel only	239	241 ± 41(stat) $\pm$ 65(syst)

Table 6.10: Results of the ABCD method compared with the observed number of events in data, after the application of all selection criteria. Reproduced as in [6].

6.4.2 Benchmark Models

Table 6.11 shows the number of signal MC events surviving after the preselection, LJ selection, and event-level selection, for the $m_H = 125$ GeV $2\gamma_\mathrm{d}$ and $4\gamma_\mathrm{d}$ FRVZ models. These results assume BR($H \rightarrow \gamma_\mathrm{d}$'s) = 0.1. Table 6.12 gives analogous results for $m_H = 800$ GeV, assuming a cross section of 1.0 pb for each signal process. This σ was not chosen according to any particular physics motivation, but rather to allow easy re-scaling to any predicted values. In both Tables, the number of events has been rescaled to the 3.57 fb^{-1} total integrated luminosity.

LJ pair type	$2\gamma_\mathrm{d}$ FRVZ Preselection	$2\gamma_\mathrm{d}$ FRVZ LJ Selection	$2\gamma_\mathrm{d}$ FRVZ Event-Level	$4\gamma_\mathrm{d}$ FRVZ Preselection	$4\gamma_\mathrm{d}$ FRVZ LJ Selection	$4\gamma_\mathrm{d}$ FRVZ Event-Level
Type0-Type0	176 ± 3	97 ± 2	80 ± 2	211 ± 3	81 ± 2	70 ± 2
Type0-Type1	-	-	-	72 ± 2	27 ± 1	16.0 ± 0.8
Type0-Type2	201 ± 3	39 ± 1	31 ± 1	164 ± 3	10.0 ± 0.6	9.0 ± 0.6
Type1-Type1	-	-	-	4.0 ± 0.4	1.0 ± 0.2	1.0 ± 0.2
Type1-Type2	-	-	-	26 ± 1	0.40 ± 0.01	0.01 ± 0.01
Type2-Type2	20 ± 1	2.0 ± 0.5	2.0 ± 0.5	13 ± 0.8	0.34 ± 0.10	0.34 ± 0.10

Table 6.11: Number of LJ pairs after the pre-selection, LJ selection, and event-level selection, on the $m_H = 125$ GeV FRVZ models. Type1 LJs can only be produced in the $4\gamma_\mathrm{d}$ model.

LJ pair type	$2\gamma_\mathrm{d}$ FRVZ Preselection	$2\gamma_\mathrm{d}$ FRVZ LJ Selection	$2\gamma_\mathrm{d}$ FRVZ Event-Level	$4\gamma_\mathrm{d}$ FRVZ Preselection	$4\gamma_\mathrm{d}$ FRVZ LJ Selection	$4\gamma_\mathrm{d}$ FRVZ Event-Level
Type0-Type0	47 ± 0.7	22.0 ± 0.3	16.0 ± 0.2	127 ± 1	66± 1	56.0 ± 0.8
Type0-Type1	-	-	-	91 ± 1	35.0 ± 0.5	17.0 ± 0.4
Type0-Type2	117 ± 1.2	36.0 ± 0.6	27.0 ± 0.5	181 ± 2	36.0 ± 0.6	31.0 ± 0.6
Type1-Type1	-	-	-	16.0 ± 0.1	4.7 ± 0.1	1.4 ± 0.1
Type1-Type2	-	-	-	64 ± 1	8.5 ± 0.2	3.3 ± 0.2
Type2-Type2	65 ±0.9	12.0 ± 0.3	10.0 ± 0.3	36.0 ± 0.7	4.0 ± 0.21	3.2 ± 0.2

Table 6.12: Number of LJ pairs after the pre-selection, LJ selection, and event-level selection, on the $m_H = 800$ GeV FRVZ models. Type1 LJs can only be produced in the $4\gamma_\mathrm{d}$ model.

Under the same assumptions, Table 6.13 shows the total expected number of signal events surviving the event-level selection, for the $2\gamma_\mathrm{d}$ and $4\gamma_\mathrm{d}$ FRVZ models. With $m_H = 125$ GeV, the expected signal is 113 ± 2 events for the $2\gamma_\mathrm{d}$ model, and 96 ± 2 events for the $4\gamma_\mathrm{d}$ model. With $m_H = 800$ GeV, the expected signal is 53 ± 0.6 events for the $2\gamma_\mathrm{d}$ model, and 112 ± 1 events for the $4\gamma_\mathrm{d}$ model. Since the background for the Type2-Type2 channel is expected to be much higher than for the other channels, the results are also shown with this channel excluded, and for this channel only.

Category	$m_H = 125$ GeV $2\gamma_\mathrm{d}$	$m_H = 125$ GeV $4\gamma_\mathrm{d}$	$m_H = 800$ GeV $2\gamma_\mathrm{d}$	$m_H = 800$ GeV $4\gamma_\mathrm{d}$
All events	113 ± 2	96 ± 2	53.0 ± 0.6	112 ± 1
Type2-Type2 excluded	111 ± 2	96 ± 2	43.0 ± 0.5	109 ± 1
Type2-Type2 only	2.0 ± 0.5	0.34 ± 0.10	10.0 ± 0.3	3.2 ± 0.2

Table 6.13: Total expected number of signal events surviving the event-level selection, for each benchmark model. Reproduced as in [6].

6.5 Limits in Benchmark Models

In this section, upper limits are set on the Higgs boson branching fraction to γ_d's , as a function of the γ_d mean lifetime, in the FRVZ benchmark models. Given the relatively low signal efficiency and high background in the Type-Type2 channel, it has been excluded from the limit-setting. The CL_s method [94] is employed for the limit-setting, where the signal region is populated from the data-driven background estimate and from the appropriate signal hypothesis. It is executed separately for the $2\gamma_\text{d}$ and $4\gamma_\text{d}$ models, and as a function of the γ_d proper lifetime.

6.5.1 Pseudo-Experiments

The method begins with a pseudo-experiment process, which proceeds as follows:

1. A benchmark MC sample with γ_d lifetime $c\tau_0$ is chosen as the "reference sample".
2. The $A \times Eff$ for the γ_d after the event-level selection is evaluated as a function of the γ_d decay length (L_xy in the barrel, L_z in the endcap) and of its p_T, in the MC reference sample.
3. A large number of MC "pseudo-experiments" (indexed by j) with different proper lifetimes, $0.5 \leq c\tau_j \leq 5000$ mm, are set up. A pseudo-experiment j consists of taking the p_T of each γ_d from the reference sample, but then re-assigning the decay length of the γ_d using an exponential distribution whose parameter is $c\tau_j$.
4. Each pseudo-experiment event is weighted by the $A \times Eff$ of its γ_d's , based on their decay length and p_T values. The $A \times Eff\ (L,\ p_T)$ is obtained from Step 2.
5. The number of events in each pseudo-experiment j is rescaled by the ratio of the integrated $A \times Eff$ at $c\tau_j$ to the integrated $A \times Eff$ at $c\tau_0$.

The ratios in the final step are plotted in Figure 6.6. For each of a few values of $c\tau_j$, a full MC simulation was performed, from which the integrated $A \times Eff$ was extracted. As shown in Figure 6.6, each of these $A \times Eff$ values was then compared to the result of the pseudo-experiment at the given $c\tau_j$, in order to validate the implementation of the pseudo-experiment process.

6.5.2 Limit-Setting

The event numbers in the rescaled pseudo-experiments, together with the expected background and observed numbers of events in region A, are employed to obtain upper limits at the 95% confidence level (CL) on signal strength, μ_{95}, for the $2\gamma_\text{d}$ and for the $4\gamma_\text{d}$ model. The μ_{95} can be expressed as

$$\mu_{95} = \frac{N_{95}}{N_{mod}} = \frac{(\sigma \times \text{BR})_{95}}{(\sigma \times \text{BR})_{mod}}$$

where N_{95} is the upper limit at 95% CL on the number of signal events, $(\sigma\times\ \text{BR})_{95}$ is the upper limit at 95% CL on the signal cross-section times branching ratio, N_{mod} is the number of signal events predicted from the benchmark model, and $(\sigma\times\ \text{BR})_{mod}$ is the signal cross-section times branching ratio predicted from the model. The use of 95% for the CL is in accordance with the standard for ATLAS results.

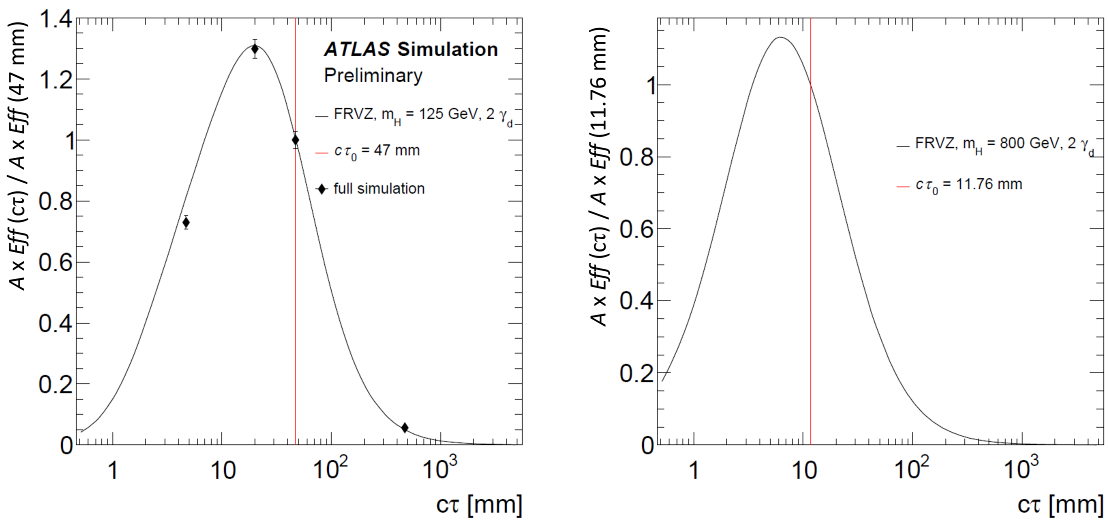

Figure 6.6: Ratio of the integrated $A \times Eff$ at a given $c\tau$, obtained using the pseudo-experiment process, to the integrated $A \times Eff$ at $c\tau_0$ (red line) of the reference MC FRVZ sample. Left: For the $m_H = 125$ GeV $2\gamma_\mathrm{d}$ sample, $c\tau_0 = 47$ mm. The markers show the $A \times Eff$ ratio found using full simulation MC with $c\tau = 4.7,\ 20,\ 47,\ 470$ mm, which indicate good agreement with the pseudo-experiment process. Reproduced as in [6]. Right: For the $m_H = 800$ GeV $2\gamma_\mathrm{d}$ sample, $c\tau_0 = 11.76$ mm.

For the 125 GeV Higgs boson, N_{95} is given as a function of the Higgs boson BR to γ_d's by

$$N_{95}(\mathrm{BR}(H \to \gamma_\mathrm{d}'\mathrm{s})) = \mu_{95} \times \sigma_{ggF} \times \mathrm{BR}(H \to \gamma_\mathrm{d}'\mathrm{s}) \times L$$

where σ_{ggF} is the SM cross-section for Higgs boson production by gluon-gluon fusion, and L the luminosity. N_{mod} is

$$N_{mod}(\mathrm{BR}(H \to \gamma_\mathrm{d}'\mathrm{s})) = \sigma_{ggF} \times \mathrm{BR}(H \to \gamma_\mathrm{d}'\mathrm{s}) \times (A \times Eff) \times \mathrm{BR}(\gamma_\mathrm{d}) \times L$$

where $\mathrm{BR}(\gamma_\mathrm{d})$ is the theoretical BR of the γ_d to all final states included in the search. For the 800 GeV Higgs boson, whose production cross-section σ_{800} is unknown, N_{95} is given as a function of this production cross-section times the Higgs boson BR to γ_d's :

$$N_{95}(\sigma_{800} \times \mathrm{BR}(H \to \gamma_\mathrm{d}'\mathrm{s})) = \mu_{95} \times \sigma_{800} \times \mathrm{BR}(H \to \gamma_\mathrm{d}'\mathrm{s}) \times L$$

.

The exclusion limits on $\sigma\times$ BR for the $m_H = 125$ GeV $2\gamma_\mathrm{d}$ and $4\gamma_\mathrm{d}$ FRVZ models, as a function of the γ_d $c\tau$, are shown in Figure 6.7. The exclusion limits for the $m_H = 800$ GeV models are shown in Figure 6.8. In both Figures, the limits do not use the Type2–Type2 channel, and the "expected" limits are generated using Asimov data sets. Table 6.14 shows the γ_d $c\tau$ ranges that are excluded. The $m_H = 125$ GeV exclusions assume $\mathrm{BR}(H \to \gamma_\mathrm{d}\text{'s}) = 0.1$, corresponding to the lower dashed horizontal line in each pane of Figure 6.7. For example, the lower end of the $2\gamma_\mathrm{d}$ exclusion range, 2.2 mm, corresponds to the $c\tau$ value of the first point in the left pane of Figure 6.7 where the lower dashed line meets the solid "observed limit" curve. The upper end of the exclusion range, 111.3 mm, corresponds to the $c\tau$ value of the second point where the dashed line meets the curve. All $c\tau$ values between these two points are excluded. Similarly, the $m_H = 800$ GeV exclusions in Table 6.14 assume a $\sigma\times$ BR of 5 pb, corresponding to the dashed horizontal line in each pane of Figure 6.8.

The essential form of the limit calculations can be understood by examining one parameter point in one model. For example, for the $2\gamma_{\rm d}$ $m_H = 125$ GeV model, the observed number of events in region A yields $\mu_{95}^{obs} = 0.0016$ by the CL_s method. Meanwhile, at the parameter point $c\tau = 2.2$ mm, $A\times Eff = 0.0023$ (Figure 6.6). Since Type2–Type2 LJs are not included, the following final states are possible, with BRs determined as in Section 3.4:

- $\mu\mu$ $\mu\mu$: $\text{BR}(\gamma_{\rm d} \to \mu\mu)^2 = 0.45^2 = 0.20$
- $ee\ \mu\mu\ +\ \mu\mu\ ee$: $2 \times \text{BR}(\gamma_{\rm d} \to \mu\mu) \times \text{BR}(\gamma_{\rm d} \to ee) = 2 \times 0.45 \times 0.47 = 0.42$
- $\pi\pi\ \mu\mu\ +\ \mu\mu\ \pi\pi$: $2 \times \text{BR}(\gamma_{\rm d} \to \mu\mu) \times \text{BR}(\gamma_{\rm d} \to \pi\pi) = 2 \times 0.45 \times 0.08 = 0.073$.

This yields $A\times Eff \times \text{BR}(\gamma_{\rm d}) = 0.0023 \times (0.20 + 0.42 + 0.073) = 0.0016$. Since this equals the observed μ_{95}, this parameter point is exactly on the border of being ruled out. Parameter points with a higher $A\times Eff \times \text{BR}(\gamma_{\rm d})$ are ruled out. Given $\sigma_{ggF} = 43870$ fb, $\text{BR}(H \to \gamma_{\rm d}\text{'s}) = 0.1$, and $L = 3.57$ fb^{-1},

$$N_{mod} = 0.1 \times 43870 \text{ fb} \times 3.57 \text{ fb}^{-1} = 15661$$

and observed

$$N_{95} = 0.0016 \times 0.1 \times 43870 \text{ fb} \times 3.57 \text{ fb}^{-1} = 24.87.$$

The justification for not including the Type2–Type2 channel can be seen intuitively, without detailed calculations, by examining $S/\sqrt{B}$ values where B is the expected background in region A. For example, for the $2\gamma_{\rm d}$ $m_H = 125$ GeV model, the Type2–Type2 channel has $S = 2$ (Table 6.13) and $B = 241$ (Table 6.10), yielding $S/\sqrt{B} = 0.13$. Meanwhile, for all other channels combined, $S = 111$, $B = 32$, yielding $S/\sqrt{B} = 20$. Therefore, although the Type2–Type2 channel accounts for $\sim 40\%$ of the $\gamma_{\rm d}$ decays according to the theoretical BRs, including this channel would obviously drag down the $S/\sqrt{B}$ terribly, and would not contribute to the limits produced.

Table 6.14 includes the $\gamma_{\rm d}$ $c\tau$ exclusions from the 8 TeV Displaced LJs search for $m_H = 125$ GeV [4], for purposes of comparison. Notwithstanding the much greater luminosity of the Run 1 data sample, 20.3 fb^{-1} rather than 3.57 fb^{-1}, the Run 2 analysis extends the Run 1 limits at low $c\tau$, mainly thanks to the improved trigger efficiency. The increase in $\sqrt{s}$ also plays a role, by increasing Higgs boson production cross-section but also by further boosting the $\gamma_{\rm d}$'s. The greater $\gamma_{\rm d}$ boost shifts the limits lower in proper $\gamma_{\rm d}$ lifetime, since a lower lifetime results in the same lab-frame decay length. While beneficial at low $c\tau$, this reduces the search sensitivity at high $c\tau$ compared to the 8 TeV analysis.

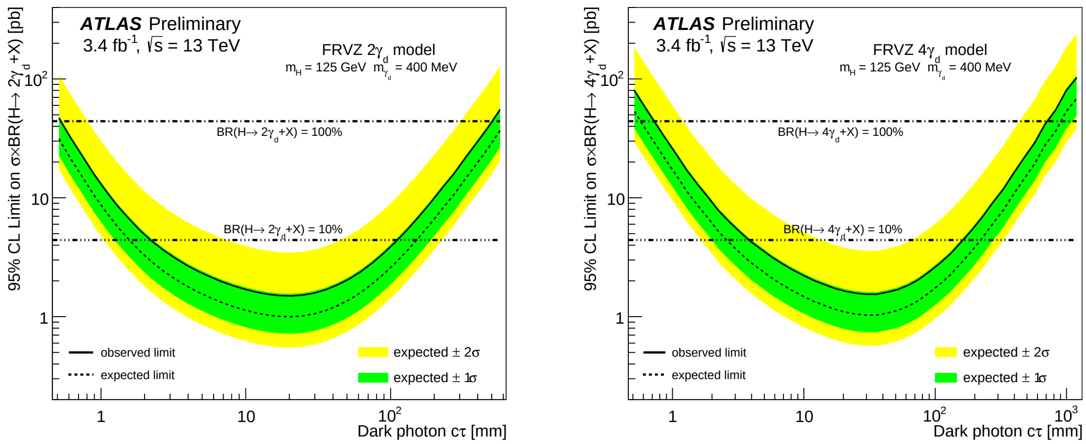

Figure 6.7: Left: The 95% upper limits on $\sigma\times$ BR for the $m_H = 125$ GeV $2\gamma_\mathrm{d}$ FRVZ model, as a function of the γ_d $c\tau$. The expected limit is shown as the dashed curve, and the observed limit as the solid curve. The horizontal lines correspond to $\sigma\times$ BR for two particular values of the BR. Right: analogous limits for the $4\gamma_\mathrm{d}$ FRVZ model. Reproduced as in [6].

FRVZ model	Excluded $c\tau$ [mm] $m_H = 125$ GeV, BR 10%	Excluded $c\tau$ [mm] $m_H = 800$ GeV, $\sigma\times$ BR $= 5$ pb
$2\gamma_\mathrm{d}$	$2.2 \leq c\tau \leq 111.3$	$0.6 \leq c\tau \leq 63$
Run 1	$14 \leq c\tau \leq 140$	
$4\gamma_\mathrm{d}$	$3.8 \leq c\tau \leq 163.0$	$0.8 \leq c\tau \leq 186$
Run 1	$15 \leq c\tau \leq 260$	

Table 6.14: Ranges of γ_d $c\tau$ excluded at 95% CL in the FRVZ benchmark models, with Type2-Type2 events excluded. Limits from the Run 1 Displaced LJs analysis, in which there was no $m_H = 800$ GeV benchmark, are included for comparison; the Run 2 results are reproduced as in [6].

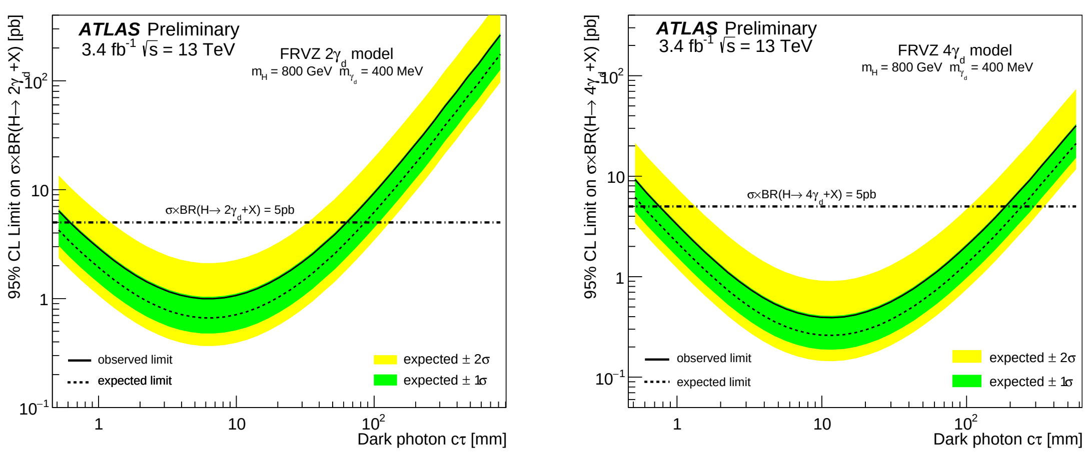

Figure 6.8: Left: The 95% upper limits on $\sigma\times$ BR for $m_H = 800$ GeV $2\gamma_\mathrm{d}$ (left) FRVZ model, as a function of the γ_d $c\tau$. The expected limit is shown as the dashed curve, and the observed limit as the solid curve. The horizontal lines correspond to $\sigma\times$ BR $= 5$ pb. Right: analogous limits for the $4\gamma_\mathrm{d}$ FRVZ model. Reproduced as in [6].

6.6 Limits in Dark Photon Parameter Space

The $c\tau$ exclusion ranges for the FRVZ models can be converted into exclusion contours in the plane of $(m_{A'},\ \epsilon)$ for hypercharge portal models, via Equation 3.4. While the machinery to perform this conversion was developed by other members of the analysis team for the Run 1 analysis, I adapted it for the Run 2 results. As discussed in Section 3.3, the Higgs portal must also be present for these limits to apply; unlike the other experimental limits established in this parameter space, they depend upon the Higgs boson branching fraction to the hidden sector, or equivalently the Higgs mixing parameter κ. For different values of this branching fraction (0.05, 0.10, 0.20) of the 125 GeV Higgs boson, the resulting 90% CL exclusion regions for $H \to 2\gamma_{\mathrm{d}} + X$ are shown in Figure 6.9, with the Type2–Type2 channel again excluded. These exclusion regions can be combined with those from the 8 TeV Displaced LJ search [4], yielding Figure 6.10. The 8 TeV limits contribute additional low-ϵ coverage, but fall short of the 13 TeV limits in high-ϵ coverage, for $m_{\gamma_{\mathrm{d}}} \lesssim 2$ GeV; above this mass, the 13 TeV contours stand alone. The combined Displaced contours can then be overlaid with the results of the Prompt LJs search at $\sqrt{s} = 8$ TeV [41], demonstrating the complementarity of the analyses (Figure 6.11).

The left edge of the Displaced LJ contours in parameter space is at $\sim 2m_{\mu}$, corresponding to the kinematic limit below which the Type0 LJs are lost. The right edge, at $m_{\gamma_{\mathrm{d}}} = 10.5$ GeV, is actually determined by lack of signal MC and theoretical BR(γ_{d}) predictions above this mass. The opening angle of the γ_{d} decay products would be expected to rise rapidly at higher masses, leading to fewer reconstructed LJs and landing in the domain of non-collimated searches. The extent of the reconstruction efficiency loss cannot be estimated without full signal MC, which was not generated due to limited availability of computational resources for simulation. In addition, the theoretical BR(γ_{d}) predictions for BR($\gamma_{\mathrm{d}} \to$ hadrons) at higher masses become difficult to generate due to hadronic resonances; no usable parameterization is currently available. The 8 TeV contours were cut off at $m_{\gamma_{\mathrm{d}}} = 2$ GeV due to even more restricted MC availability.

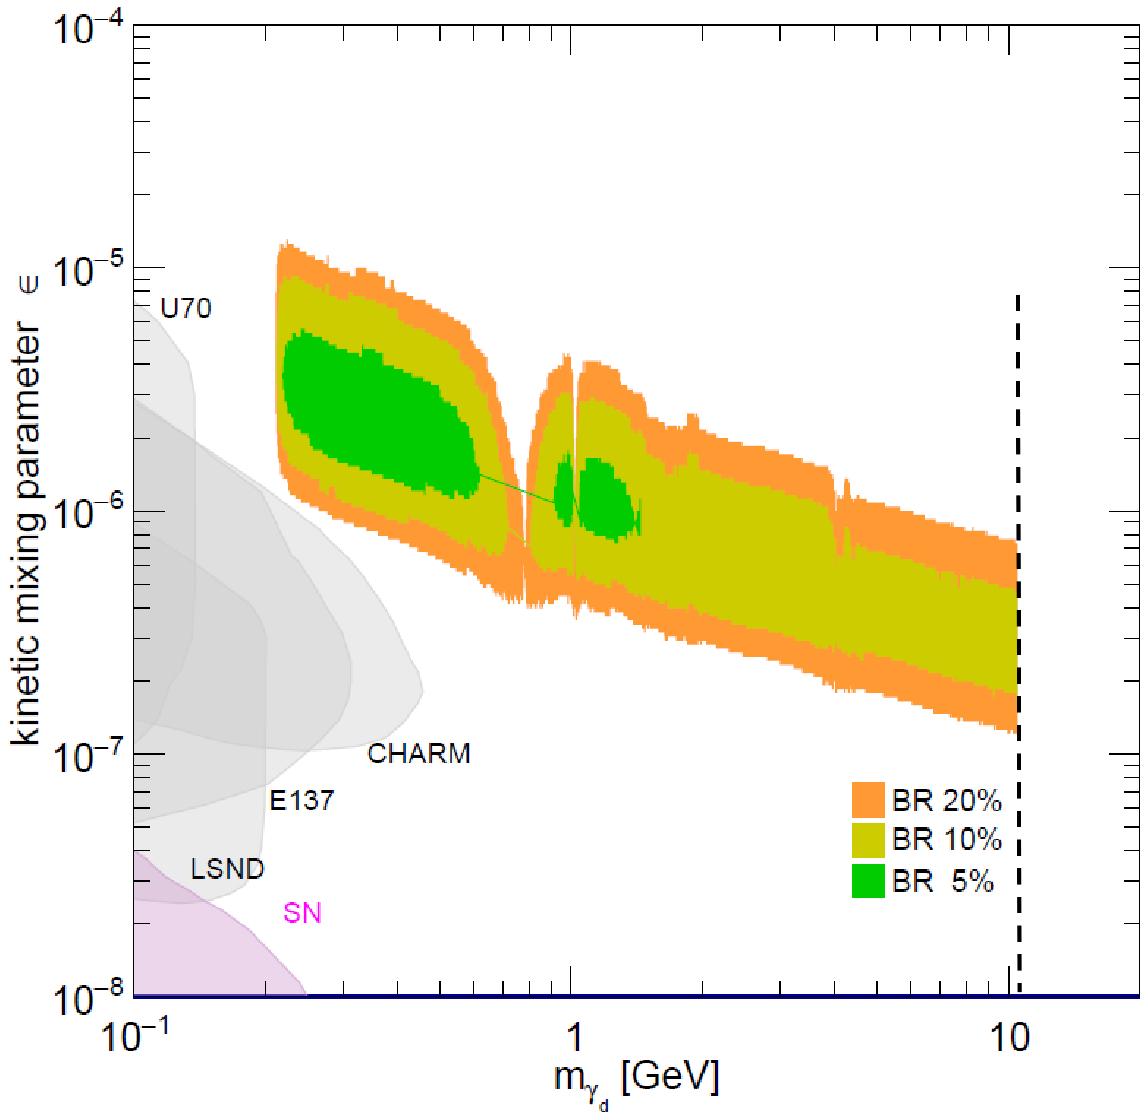

Figure 6.9: 90% CL exclusion regions for $H \to 2\gamma_{\mathrm{d}} + X$. The contours are cut off at $m_{A'} \approx 10.5$ GeV (dashed vertical line) due to limitations in MC availability and BR predictions, as discussed in the text.

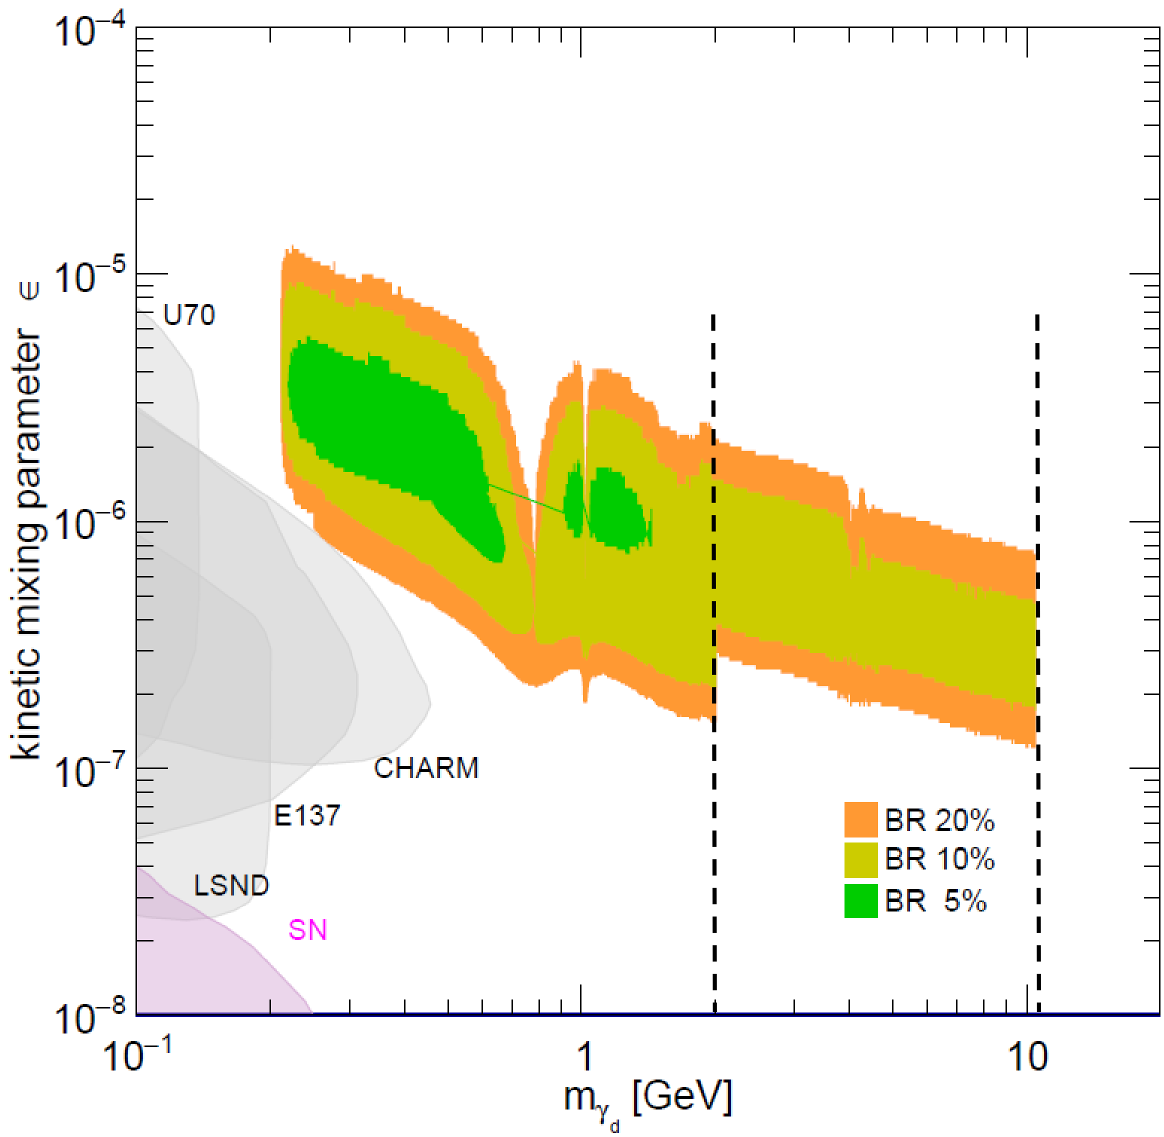

Figure 6.10: 90% CL exclusion regions for $H \to 2\gamma_\mathrm{d} + X$, combined with those from the $\sqrt{s} = 8$ TeV Displaced LJs analysis. The 13 TeV analysis is cut off at $m_{A'} \approx 10.5$ GeV, while the 8 TeV analysis is cut off at $m_{A'} \approx 2$ GeV (dashed vertical lines), due to limitations in MC availability and BR predictions as discussed in the text.

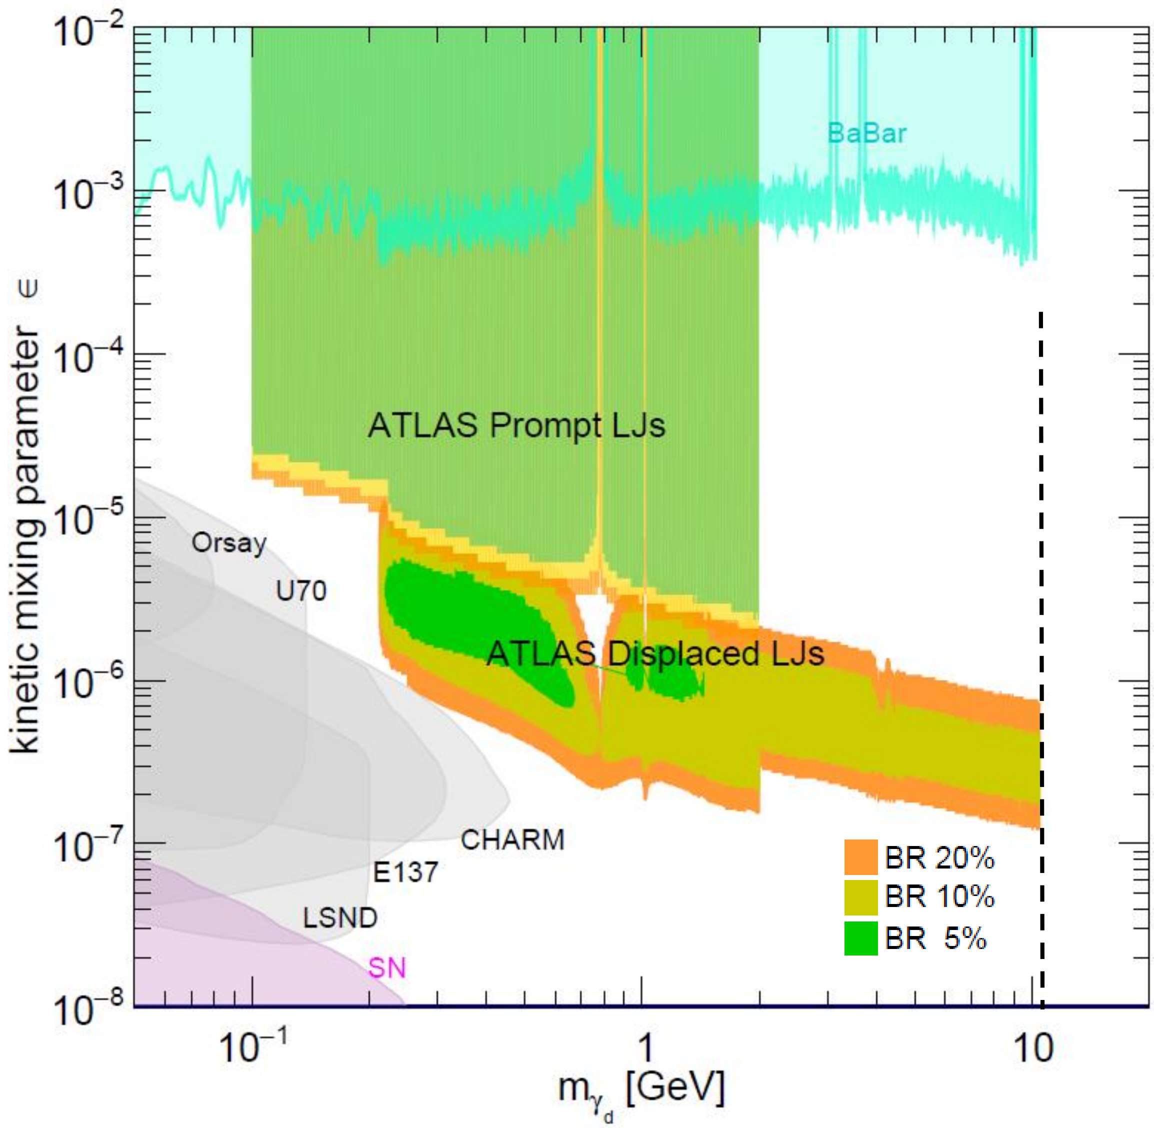

Figure 6.11: 90% CL exclusion regions for $H \to 2\gamma_{\mathrm{d}} + X$, combined with those from the $\sqrt{s} = 8$ TeV Displaced LJs analysis, overlaid with the results of the Prompt LJs search at $\sqrt{s} = 8$ TeV. The Displaced contours are cut off at $m_{A'} \approx 10.5$ GeV (dashed vertical line) due to limitations in MC availability and BR predictions, as discussed in the text.

Chapter 7

Systematic Uncertainties

Several sources of systematic uncertainty ("systematics") in the analysis are included as input to the final results presented in Section 6.4. One source, the overall normalization uncertainty of the integrated luminosity, has a standard centrally-determined ATLAS value for 2015 data, calculated as per the methodology detailed in [95]. The other sources required dedicated investigations, sometimes in consultation with experts outside the Displaced LJs analysis team whom I sought out and liaised with, and are discussed in detail in the upcoming sections of this chapter. The first is multi-jet and cosmic background estimation, where the ABCD data-driven method (see Section 5.6.3) makes two major contributions to the systematics: one from possible residual correlations between the two variables employed in the method, and the other from possible non-closure. The second is reconstruction efficiency for pairs of highly-collimated non-combined muons, evaluated using a Tag-and-Probe method applied to $J/\psi \to \mu\mu$ 2015 data and MC samples. Next come the trigger efficiencies: the systematic on the Narrow-Scan efficiency is evaluated using a Tag-and-Probe method applied to $J/\psi \to \mu\mu$ 2015 data and MC samples, while for the systematics for the Tri-Muon and Cal-Ratio triggers are taken from analyses of 2012 data. Finally, the presence of pile-up affects the efficiency of the $\max(\Sigma p_{\mathrm{T}}^{\mathrm{ID}})$ cut, contributing a systematic that is evaluated by computing the $\Sigma p_{\mathrm{T}}^{\mathrm{ID}}$ of $Z \to \mu\mu$ events in 2015 data as a function of the number of interaction vertices in the event. When conducting searches for exotic signatures, involving non-standard analysis strategies, such dedicated systematics investigations are a necessary price to be paid.

For the sake of completeness, small sources of systematic uncertainty ($\lesssim 1\%$), while not discussed in detail, are briefly stated at the end of this chapter. The uncertainty on the Higgs boson production cross-section is not included in the analysis systematics, following the ATLAS convention for all Run 2 analyses involving Higgs boson decays.

7.1 Multijet and Cosmics Background

The total systematic uncertainty on the multijet and cosmics background obtained from the ABCD method is rather high, 27%. It is the sum in quadrature of two contributions outlined below, as defined by the statistics expert on the Displaced LJs team, Stefano Giagu, in consultation with the ATLAS Statistics Forum.

7.1.1 Correlation between ABCD Variables

The ABCD method relies on the assumption that the variables $|\Delta\phi|_{\mathrm{LJ}}$ and $\max(\Sigma p_{\mathrm{T}}^{\mathrm{ID}})$ are highly uncorrelated. This assumption is verified in Figure 7.1, where a linear correlation would appear as an overall slope in the black points. The correlation coefficient turns out to be 0.02, invisible in the Figure. However, this does not allow us to entirely neglect the possibility of correlation; there would be uncertainty on this coefficient even if it were to come out as 0, and in the interests of being conservative, a systematic uncertainty associated with the overall effect of correlation on the ABCD method must be assigned. This uncertainty was determined by creating a gap between the signal region and the control regions, discarding all events that fall in the gap, and re-running the ABCD method. The signal region was kept fixed, while the size of the gap was varied based on the expected resolution for the two variables. The expected resolution for $|\Delta\phi|_{\mathrm{LJ}}$, ~ 0.1 rad, was estimated by comparing the LJ direction at the MC generator level with the reconstructed direction. The expected resolution for $\max(\Sigma p_{\mathrm{T}}^{\mathrm{ID}})$, 0.25 GeV, was estimated from the variable's distribution in a sample of $Z \to \mu\mu$ events in 2015 data (see Section 7.2). A bootstrap resampling method was then applied to subtract out the effects of statistical fluctuation, leaving only the systematic uncertainty which amounts to 15%.

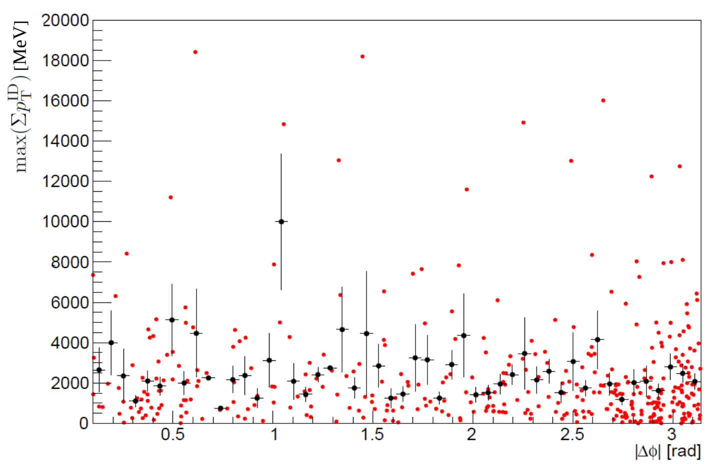

Figure 7.1: In 2015 data after the LJ selection stage, the correlation factor between $\max(\Sigma p_{\mathrm{T}}^{\mathrm{ID}})$ and $|\Delta\phi|_{\mathrm{LJ}}$ is 0.02. The points with error bars represent the profile histogram of $\max(\Sigma p_{\mathrm{T}}^{\mathrm{ID}})$ as a function of $|\Delta\phi|_{\mathrm{LJ}}$.

7.1.2 Non-Closure

The second contribution to the systematic uncertainty from the ABCD method comes from potential non-closure, i.e. Equation 5.3 failing to hold for each background source. The ABCD method is designed to always attain closure if there is only one source of background. In the case of the LJ analysis, however, there are two distinct backgrounds still in play, multijets and cosmics. Therefore, the closure of the method for each background source must be evaluated by examining a sample that is highly dominated by only one of the background sources. For this purpose, a dedicated version of the LJ selection was created with the cuts on z_0 (for Type0 and Type1 LJs) and on the jet timing (for Type1 and Type2 LJs) inverted, and run on the 2015 data sample. The resulting sample should be highly enriched in cosmics. Then, yet another version was created by instead loosening the requirement on f_{EM} to 0.6 and running over the di-jet data sample; with the higher statistics, the sample should be highly dominated by multijets. The results of the ABCD method on these samples are summarized in Table 7.1.

Dedicated Sample	Expected n_A	Observed n_A
Multijet	169 ± 72	141 ± 12
Cosmic	1257 ± 361	1082 ± 33

Table 7.1: Results of the ABCD method on two dedicated samples.

The two samples were then combined, with a scale factor derived by fitting the sum of the two $|\Delta\phi|_{\mathrm{LJ}}$ distributions to the nominal 2015 data. The resulting number of events in the signal region for the combined sample is 329 ± 75, compared with the observed 270 ± 27 in the nominal selection. From all this, a systematic uncertainty of $(329 - 270)/270 \times 100\% = 22\%$ is assigned.

7.2 Reconstruction Efficiency for Highly-Collimated Non-Combined Muon Pairs

To estimate the systematics on the $\gamma_{\mathrm{d}} \to \mu\mu$ reconstruction efficiency, standard ATLAS muon systematics cannot be used, because their validity has not been demonstrated for the case of highly-collimated non-combined muon pairs. Therefore, the reconstruction efficiency is measured in $J/\psi \to \mu\mu$ events, using the Tag-and-Probe method described in detail in [96], as a function of the ΔR between the two muons and of muon p_{T}. $J/\psi \to \mu\mu$ events are employed because the di-muon system born from very boosted J/ψ decay, i.e. γ >100, shows spatial collimation similar to the case of muonic LJs. While there are various SM di-muon systems that could satisfy this criterion, the J/ψ is the only one that has reasonably low background, and high enough cross section to guarantee sufficient statistics. High statistics are required in order to obtain an efficiency measurement in each bin of p_T and ΔR between the two muons.

The ATLAS combined muon reconstruction efficiency is the product of the efficiency in the ID, the efficiency in the MS, and the matching efficiency between the ID and MS tracks (see Section 4.4.3 for a description of the muon reconstruction process and muon types). In its standard usage [96], the Tag-and-Probe method is designed to be sensitive to either the ID efficiency or the MS and matching efficiency. $Z \to \mu\mu$ decays are selected by requiring two oppositely-charged muon tracks, with invariant mass near m_Z. One of the muon tracks must correspond to a Combined muon; this muon track is the "tag". If the ID efficiency is to be measured, the other track ("probe") must correspond to a MS track, and the efficiency is the fraction of probes which can be associated to an ID track. If the MS and matching efficiency is to be measured, the other track (probe) must be an ID track, and efficiency is the fraction of ID probes which can be associated to a Combined or Segment-tagged muon. In its modified usage for the Displaced LJs analysis, only the MS efficiency is to be measured. Therefore the probe is an ID track, but the efficiency is the fraction of ID probes which can be associated to an MS track.

The muon reconstruction efficiency in a 2015 data sample dominated by $J/\psi \to \mu\mu$ events are compared to those in MC, and the difference taken as systematic error. The comparisons were performed primarily by Marco Schioppa, the muon expert in the Displaced LJs analysis group. The $J/\psi \to \mu\mu$ MC events are centrally produced by ATLAS. Meanwhile, the data sample corresponds to an integrated luminosity of 3.5 fb^{-1} , collected using a combination of J/ψ triggers with minimum muon p_{T} thresholds ranging from 6 to 18 GeV. The selection then requires the muon that fired the trigger to have $p_T > 4$ GeV, pass the standard quality criteria for combined muons, and fall in the range $|\eta| < 2.4$ corresponding to the acceptance of the ID. This triggering muon represents the tag. The second muon of the decay, reconstructed in the ID with opposite charge to the tag, represents the probe. The probe must have $p_T > 4$ GeV, and the invariant mass of the tag and probe must satisfy $2700 < m_{tag+probe} < 3500$ MeV. This search is repeated for all the tag and probe pairs found in the event; the pair whose $m_{tag+probe}$ is closest to $m_{J/\psi}$ is chosen. The background contamination is approximately halved by applying a cut on the difference of the z_0 impact parameter between the tag and the probe: $\Delta z_0 = |z_0^{tag} - z_0^{probe}| < 5$ mm. This cut keeps $\sim 97\%$ of the signal, with ~ 450000 probes passing, but the signal-to-background ratio for the probes is only $\sim 1/80$ at this stage.

Tag-and-probe pairs are grouped into two independent samples: pairs in which the probe is matched to an MS track (matched sample) and pairs in which the probe is not (unmatched sample). Typically the unmatched sample contains a small number of signal events, dominated by the background continuum. The efficiency is derived by a simultaneous log-likelihood fit of the two statistically independent distributions of $m_{tag+probe}$:

$$\begin{aligned} f_{matched}(m) &= N_{tot} Eff_\mu S(m;\theta_S) + B(m;\theta_B) \\ f_{unmatched}(m) &= N_{tot}(1 - Eff_\mu) S(m;\theta_S) + B(m;\theta_B) \end{aligned} \tag{7.1}$$

where $S(m;\theta_S)$ and $B(m;\theta_B)$ are the signal and background models with parameter sets θ_S and θ_B, respectively. The parameters θ_S are forced to be the same in the matched and unmatched samples, while the θ_B parameters are left to vary independently. The efficiency Eff_μ and the total number of events N_{tot} are extracted from the fitting procedure. A Crystal Ball function [97] is used for the signal, and a second-order polynomial function for the background models. The fits are performed in bins of p_{T}, η and ϕ for the probe, and in bins of ΔR between tag and probe.

The difference between the MC and the signal fit to data is shown in Figure 7.2). In the interval $0.02 \leq \Delta R \leq 0.06$, where the LJ decay products are expected to fall, it amounts to 15%.

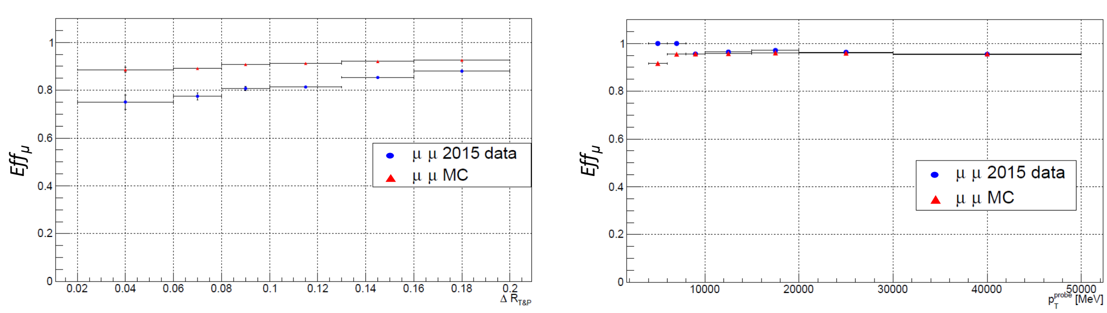

Figure 7.2: Muon probe reconstruction efficiency measured using tag-and-probe pairs from $J/\psi \to \mu^+\mu^-$ events, for both data and signal MC. Left: as a function of the ΔR between tag and probe. Right: as a function of the probe p_{T}.

7.3 Trigger Efficiencies

The systematic uncertainty on the efficiency of each of the triggers employed in the analysis is evaluated separately. The largest of these systematics is on the CalRatio trigger.

7.3.1 CalRatio

The systematic uncertainty on the CalRatio trigger is assumed to be the same as in 2012 data [98]. There, the performance of the trigger in data was compared with the performance in MC, studying each main trigger requirement individually: minimum jet E_{T} , maximum f_{EM} , and maximum $\Sigma p_{\mathrm{T}}^{\mathrm{ID}}$. These comparisons were performed by a dedicated "CalRatio" team in the UEH subgroup. 2012 MC dijet samples ($\sqrt{s}$ = 8 TeV) were used along with a 2012 dijet sample, where events in both samples were required to pass a single-jet ($p_{\mathrm{T}} > 55$ GeV) trigger. The Tag-and-Probe method was employed, where the triggering jet served as the tag and the jet with the next-highest p_{T} was used as the probe. Both tag and probe had to satisfy $|\eta| < 2.5$, with $|\Delta\Phi| > 2.0$ between tag and probe. Using only the probes, data and MC "turn-on curves" – plots of efficiency as a function of jet E_{T}, f_{EM}, and $\Sigma p_{\mathrm{T}}^{\mathrm{ID}}$ – were constructed for each of the three trigger requirements above. For jet E_{T} and $\Sigma p_{\mathrm{T}}^{\mathrm{ID}}$, the discrepancies between the data curve and MC curve are negligible. The f_{EM} curve is shown in Figure 7.3. The difference between data and MC is taken at the trigger working point, maximum $\log(f_{EM}) = 1.2$, where it is 11%.

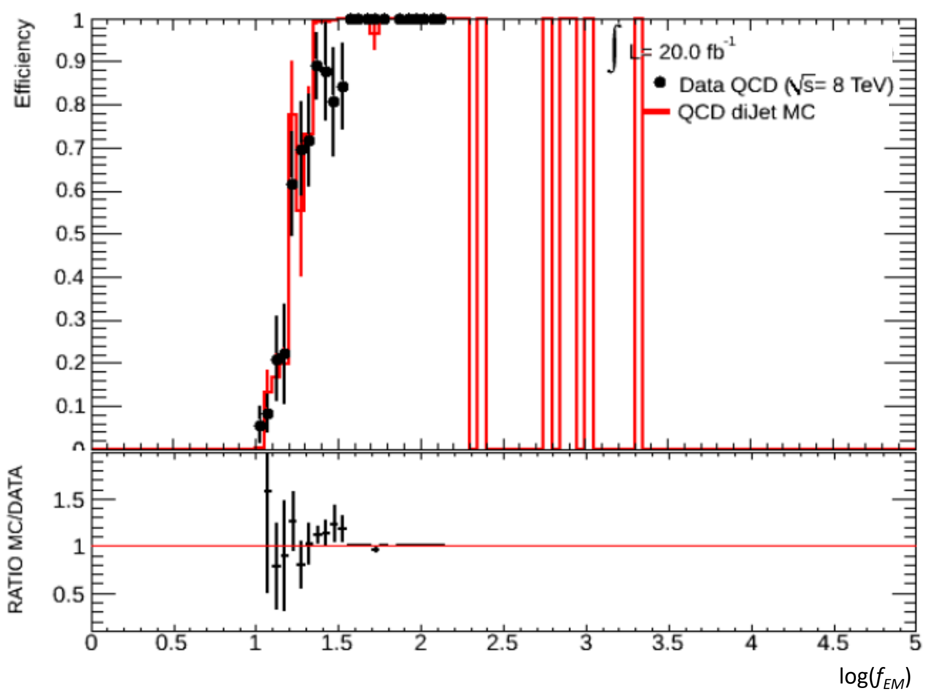

Figure 7.3: CalRatio turn-on curves from [98], for 2012 dijet MC and data, for f_{EM} . The difference between data and MC is 11% at the trigger working point for maximum $\log(f_{EM})$, 1.2.

7.3.2 Narrow-Scan

This systematic is evaluated using the Tag-and-Probe method with $J/\psi \to \mu\mu$ events, as documented in Section 7.2. In this case, however, the tag muon must have $p_T > 20$ GeV and fire the Narrow-Scan trigger, and the probe must appear in the MS with $\Delta R < 0.5$ with respect to the tag. The narrow-scan trigger efficiency, as a function of the probe p_{T} and of the ΔR, is shown in Figure 7.4 for both the $J/\psi \to \mu\mu$ data and MC. The difference between data and MC in the range most relevant for LJs, $0.02 \leq \Delta R \leq 0.06$, amounts to 6%.

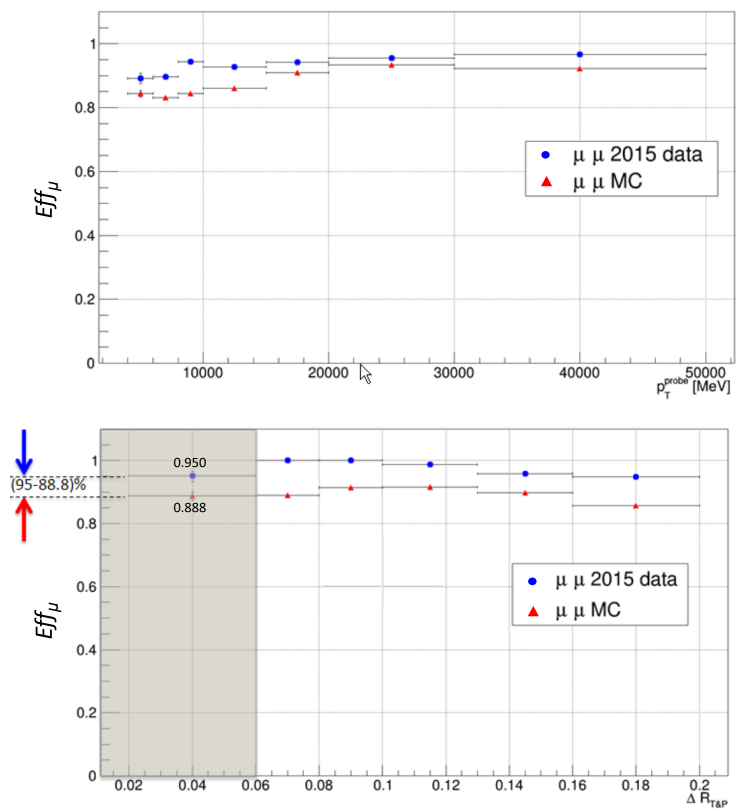

Figure 7.4: Narrow-scan trigger efficiency, measured using probes from $J/\psi \to \mu^+\mu^-$ events, for both data and MC. Top: as a function of the probe p_{T}. Bottom: as a function of the ΔR between tag and probe, where the range of interest for LJs (≤ 0.06) is shaded, and the discrepancy between efficiency in this range in data (0.950) and MC (0.888) is highlighted.

7.3.3 Tri-Muon

The systematic uncertainty on the Tri-muon trigger efficiency is assumed to be the same as in the 2012 data [4]. There, it was assumed to be dominated by the systematic uncertainty on the efficiency of the "2mu6" part of the trigger chain, i.e. requiring 2 muons with $p_{\mathrm{T}} > 6$ GeV. I was not involved in the assessments of the 2mu6 efficiency, as they were performed by the Run 1 Displaced LJs team prior to my joining the analysis. The 2mu6 efficiency uncertainty was evaluated using the Tag-and-Probe method on 2012 $J/\psi \to \mu\mu$ MC and data. The selection criteria for the Tag-and-Probe were similar to those in Section 7.2. However, only the matched sample was employed, and the associated MS tracks were used as the probes. The trigger efficiency corresponds to the probability for a probe to pass the $p_{\mathrm{T}} > 6$ GeV muon trigger. The efficiency as a function of ΔR between the tag and probe agrees well between 2012 data and MC (see Figure 7.5). The small discrepancy between data and MC yields a systematic uncertainty of 5.8%.

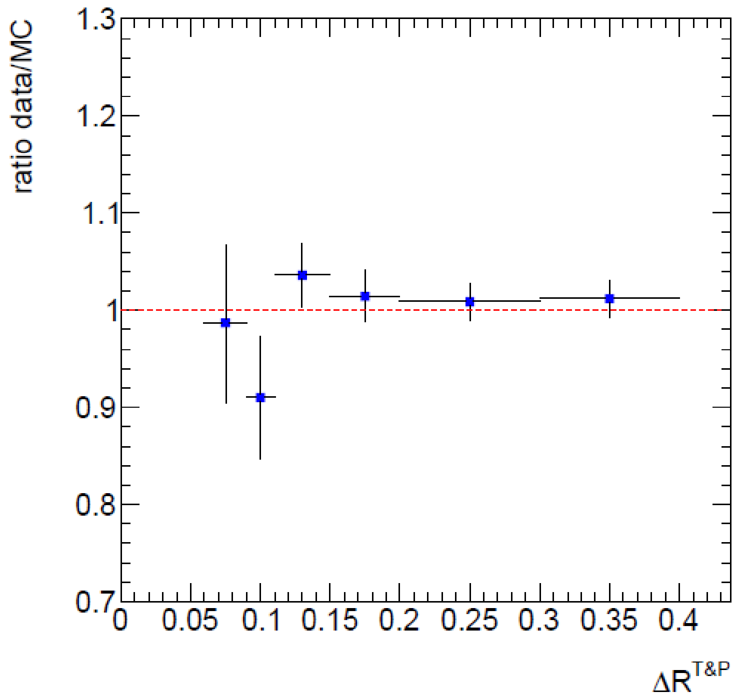

Figure 7.5: Ratio between 2mu6 trigger efficiency in 2012 $J/\psi \to \mu\mu$ data and MC, as a function of ΔR between the tag and probe.

7.4 Efficiency of ID Isolation Cut

The efficiency of the $\Sigma p_{\mathrm{T}}^{\mathrm{ID}}$ cut is defined here as the number of muons with $\Sigma p_{\mathrm{T}}^{\mathrm{ID}}$ not exceeding a given value, divided by the total number of muons. This efficiency was evaluated in $Z \to \mu\mu$ data and MC (see Section 5.6.1), where the data was divided into seven sub-samples of events with increasing numbers of interaction vertices. As shown in Figure 7.6, the efficiency as a function of event $\Sigma p_{\mathrm{T}}^{\mathrm{ID}}$ shows no significant difference between data and MC. The systematic uncertainty is taken as the maximum variation of the efficiency at 4.5 GeV, the value of the max($\Sigma p_{\mathrm{T}}^{\mathrm{ID}}$) analysis cut. It amounts to 5.1%.

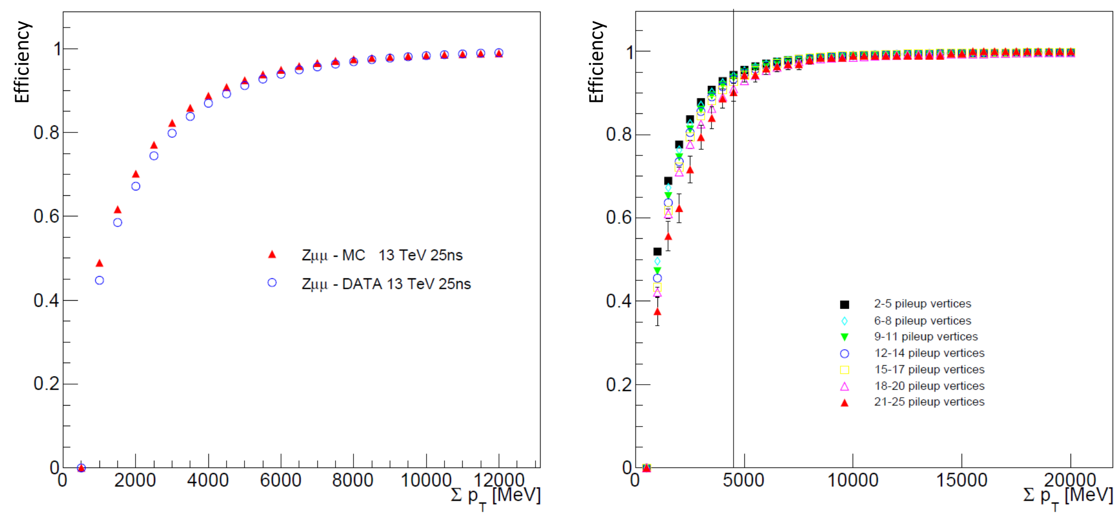

Figure 7.6: Left: Efficiency of the $\Sigma p_{\mathrm{T}}^{\mathrm{ID}}$ cut as a function of muon $\Sigma p_{\mathrm{T}}^{\mathrm{ID}}$, in a $Z \to \mu\mu$ data sample and in MC. Right: In data, for various different ranges of number of interaction vertices per event.

7.5 Small Systematics

The standard ATLAS method for evaluating Jet Energy Scale (JES) uncertainty components consists of performing an in situ jet pseudorapidity intercalibration analysis, as per [88]. For the Type1 and Type2 LJs, this analysis had to be repeated specifically for low f_{EM} jets. Results in the lowest f_{EM} range, namely 0 to 0.05 – corresponding to the LJ signal region – yield an uncertainty below 1%.

The reconstructed p_{T} of the γ_{d} often differs somewhat from the MC generator-level p_{T} value. This difference propagates to the reconstructed γ_{d} decay products, which affects the efficiency of all cuts (e.g. quality criteria) that have p_{T} dependence. Based on data vs MC studies, the resulting systematic uncertainty comes out to below 1%.

7.6 Summary of Systematics

Table 7.2 summarizes all systematics considered and their corresponding values. Adding them in quadrature yields a total of $\sim$ 34%. While the uncertainties on the event yields in Table 6.10 have a greater contribution from systematics than from statistics, the systematic uncertainties could in fact be brought down with higher statistics. This is particularly evident in the largest source of systematic uncertainty, namely the background estimation. A larger sample of events in empty bunch-crossings would allow the cosmic background to be much more comprehensively characterized, leading to its reliable estimation by means other than the ABCD method. Then only the multijet background estimation would be performed using ABCD, bringing down the non-closure systematic. This would visibly narrow the uncertainty bands in Figure 6.7 and 6.8, but would not visibly alter the exclusion contours in Section 6.6.

Systematic Uncertainty	Values (in %)
Multi-jet and cosmic background estimation	27
Single γ_{d} reconstruction efficiency	15
Trigger: CalRatio	11
Trigger: Narrow Scan	6.0
Trigger: Tri-muon	5.8
Effect of pile-up on $\Sigma p_{\mathrm{T}}^{\mathrm{ID}}$	5.1
Luminosity	2.1
Jet energy scale	≤ 1
p_{T} of reconstructed γ_{d}	≤ 1

Table 7.2: Summary of the systematic uncertainties, reproduced as in [6].

Chapter 8

Future Extensions

This chapter explores future extensions planned or possible for ATLAS Run 2 LJ analyses. Sensitivity will gradually improve with luminosity as more data is gathered; a wealth of 2016 data is already waiting to be examined (see Figure 8.1). Taking full advantage of this data will make luminosity just one of several potential sources of improvement, as described in this chapter. The first two sections discuss particular analysis enhancements that would improve the limits in the benchmark models: reconstruction of tracks that begin partway through the Pixel subsystem of the ID, and inclusion of more final-state topologies. The next section shows the scale of the overall improvement that could be attained through a synthesis of these enhancements. The final section discusses other DM models that could also be constrained through enhanced LJ searches.

8.1 Large-Radius Tracking

The gap in sensitivity between the Prompt and Displaced analyses (Section 6.6) at $\epsilon \sim 5 \times 10^{-6}$ to 1×10^{-5} occurs because neither analysis deals with tracks that begin partway through the Pixel subsystem of the ID. Standard ATLAS tracking assumes that all interesting charged physics are produced at the IP; it is tuned to reconstruct tracks that point back to the IP with very high efficiency and low fake rates, at the expense of displaced tracks. Several cuts throughout the tracking process – on impact parameters, minimum number of hits in the innermost ID layers, etc. – discriminate against fake tracks and reduce processing time, but also prevent most displaced tracks from being reconstructed. The Large-Radius Tracking (LRT) initiative, with participants from the ID Tracking Combined Performance group as well as from various interested Exotics and SUSY analyses, was conceived in order to extend the existing ATLAS offline tracking framework to recover such displaced tracks. I served as the Exotics Liaison on this initiative since its inception. A preliminary version of LRT has been implemented, with validation plots produced by Andreas Sogaard, but still needs to be tuned and optimized.

8.1.1 Standard ATLAS Tracking

In order to understand LRT, it is necessary to first briefly discuss the standard ATLAS offline tracking framework [69, 99]. The standard offline tracking proceeds in two passes: first “inside-out”, and then “outside-in”.

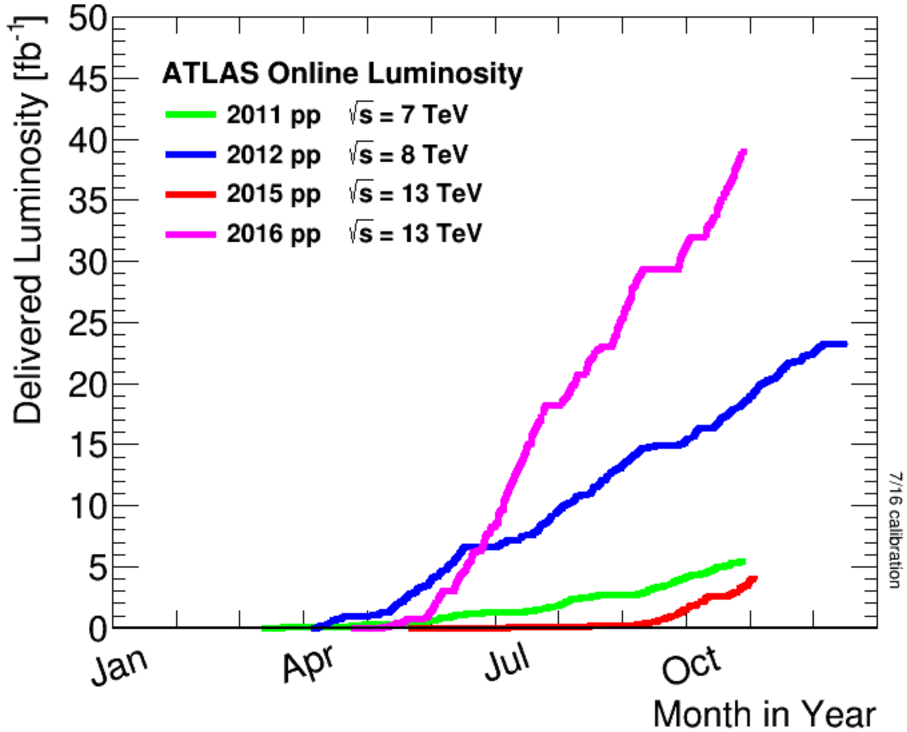

Figure 8.1: Cumulative luminosity of *pp* collisions versus day delivered to ATLAS. 2016 provided ~ 8 times more Run 2 data than 2015, and about twice the cumulative luminosity of Run 1.

The inside-out process is summarized in Figure 8.2. It begins with a data preparation stage, where clusters of Pixel and SCT hits are converted into 3-dimensional "SpacePoints". The next stage, candidate creation, builds "seeds" from combinations of 3 SpacePoints: either 3 Pixel, 2 Pixel + 1 SCT, or 3 SCT. In the case of seeds built from SCT hits, a beam spot constraint is applied, which implicitly assumes the seeds correspond to tracks that point back to the IP. Each seed provides enough directional information to build a track "road" through the rest of the Pixel and SCT sub-detectors, along which to search for more hits. Combinatorial track finding algorithms follow the roads to assign additional hits to seeds, producing track candidates. The third stage, candidate selection, separates the valid tracks from the fake tracks (quasi-random collections of hits and mis-matched track segments) amongst the candidates. This is done with the help of a track fitter that assigns a χ^2 to each candidate, and a scoring system that assigns a quality value to each candidate according to the χ^2 along with number of other criteria. An "ambiguity processor" sorts through candidates that share one or more hits, and decides which of the hits should actually belong to which of the tracks. A high number of clusters unshared with any other candidates helps a candidate pass the quality cuts; clusters that are shared with other candidates can hinder a candidate. A candidate is also hindered by "holes", i.e. absences of hits on Pixel or SCT layers along the path of the track between the first and last hits, especially if the holes are in the Pixel layers. Candidates that are selected become "Si-only tracks". The final stage is TRT extension, where a road search in the TRT is performed along the direction of each Si-only track. If TRT hits are added to an Si-only track, the result is an "Extended track"; another ambiguity processor and scoring system then decides between the Extended and the Si-only track.

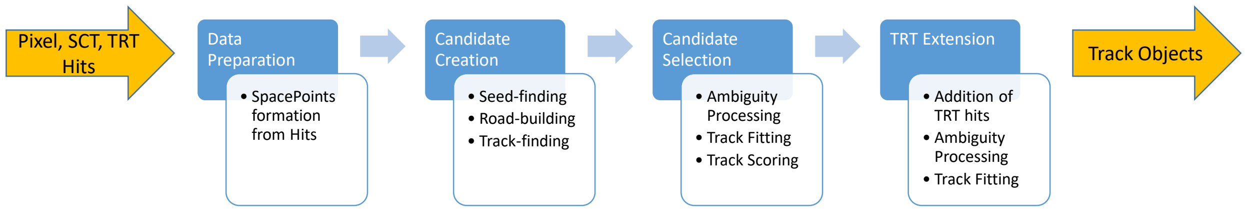

Figure 8.2: Summary of the inside-out tracking process.

The outside-in process is similar, but begins by forming seeds in the TRT and propagating them to the SCT and Pixel sub-detectors.

8.1.2 Modifications for LRT

LRT applies a third tracking pass, which is similar to the inside-out pass but only uses hits not already assigned to tracks in the previous passes. Restricting the tracking process to these "leftover" hits greatly reduces the combinatorics, which allows several of the constraints throughout the process to be loosened without making the pass computationally intractable. Specifically, the beam-spot constraint is removed in the candidate creation stage, and several cuts are loosened in the candidate selection stage. The LRT cut values are compared to the standard cut values in Table 8.1. The most important differences consist of the loosened impact parameter cuts (see Section 4.1 for the definition of d_0, z_0). For certain cuts, tracks likely to be from "secondaries", i.e. particles produced from interactions of other particles with the detector material, are treated differently than tracks from "primaries". Since the third pass does require considerable CPU time, approximately twice the sum of the first two passes, LRT is designed to be run on pre-selected sub-sets of events likely to contain displaced tracks.

Cut	Value in Standard Tracking	Value in LRT
max d_0, primaries [mm]	10	300
max d_0, secondaries [mm]	100	300
max z_0 [mm]	320	1500
max $\|\eta\|$	2.7	5
min p_{T} for primaries [MeV]	100	500
min p_{T} for secondaries [MeV]	400	500
max p_{T} [TeV]	none	1.0
min Pixel+SCT unshared clusters	6	5
max Pixel+SCT shared clusters	1	2
max holes on track	3	2
max Pixel holes on track	2	1

Table 8.1: Summary of Standard Tracking vs Large-Radius Tracking (LRT) cuts.

8.1.3 Preliminary Indicators of LRT Performance

The preliminary version of LRT has been tested on a "long-lived Z'" MC sample, consisting of $pp \to Z' \to \mu^+\mu^-$ events where the Z' is a 250 GeV BSM gauge boson with proper lifetime $c\tau = 500$ mm. The sample includes pile-up and full truth information. One important indicator of the LRT performance is the tracking efficiency in reconstructing the "signal" tracks, i.e. the muon tracks from the displaced Z' decay. In this respect, LRT significantly outperforms the standard tracking for d_0 up to ~ 200 mm, as shown in Figure 8.3. Other important factors are the fake rate, i.e. the rate of reconstructing tracks that don't actually correspond to any truth particle, and the reconstruction quality, i.e. how closely a reconstructed track matches a particle's truth track. One indicator of these factors is given by the "track truth-matching probability", which is determined by the fraction of clusters on the track that correspond to a single truth particle. Tracks with low truth-matching probability are usually fakes; tracks with higher (but still < 1) truth-matching probability usually correspond to truth particles but are poorly-reconstructed. As shown in Figure 8.4, LRT produces more low-probability tracks than standard tracking, with spikes at $\sim 20 - 30\%$ probability. Reduction of the fake rate is one of the main goals of the ongoing tuning.

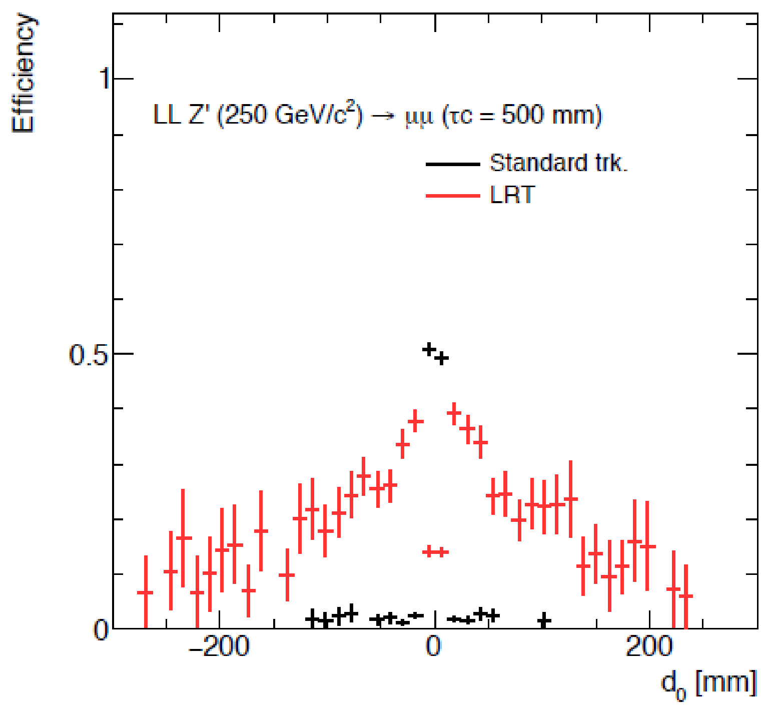

Figure 8.3: Track reconstruction efficiency for LRT and standard tracking, for a $pp \to Z' \to \mu^+\mu^-$ MC sample where the Z' has proper lifetime $c\tau = 500$ mm. The efficiency is calculated only for the signal tracks, i.e. the muon tracks from the displaced Z' decay.

Ongoing studies to tune the LRT modifications include:

- Reconstruction quality: residuals and pulls of all track parameters (d_0, z_0, θ, ϕ, q/p) in bins of truth-matching probability.
- Data/MC comparison: verifying that the MC accurately models the LRT performance in data.
- TRT extension: verifying that the extension algorithms work efficiently and accurately in LRT, and optimizing track selection based on TRT outliers and hits.

Several Run 2 Exotics and SUSY analyses are already testing the preliminary version of LRT, in conjunction with other upgrades to their displaced vertexing and displaced lepton identification codes.

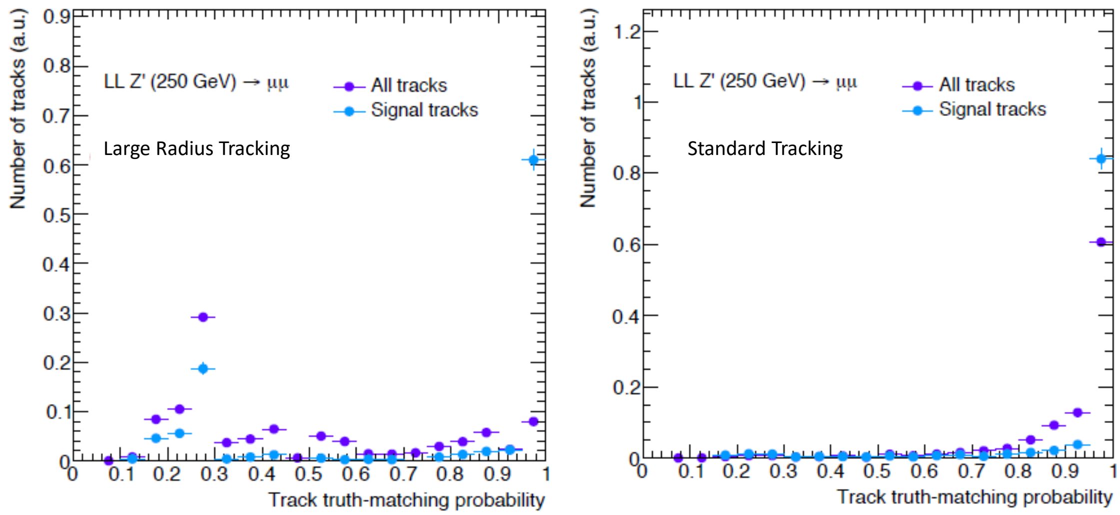

Figure 8.4: Track truth-matching probability for LRT (left) and standard tracking (right), for a $pp \to Z' \to \mu^+\mu^-$ MC sample where the Z' has proper lifetime $c\tau = 500$ mm. "Signal tracks" refer to the muon tracks from the displaced Z' decay.

LRT will be included in the ATLAS tracking software releases that will be used for reprocessed 2015-2016 data, and onward.

8.2 Additional Final-State Topologies

The Displaced LJ searches have so far required at least two LJs per signal event. Searching for only one final-state LJ, "mono-LJ searches", would require novel methods of reducing the backgrounds in order to achieve a sufficient signal-to-background ratio, but would open up a range of new possibilities. One such possibility consists of searching for cases where one of the two γ_d's escapes the detector without yet decaying, manifesting simply as E_T^{miss} . Such a search could provide sensitivity to even greater $c\tau$, corresponding to average decay length of ~ 1 m, and therefore sensitivity to even lower ϵ values.

Another enhancement that could be implemented, possibly in conjunction with the mono-LJ search, consists of searching for W/Z associated production of the Higgs boson that begins the FRVZ decay chain. Although the associated production cross-section ($\sigma_{WH} = 1.38$ pb, $\sigma_{ZH} = 0.87$ pb) is much lower than the gluon-gluon fusion production cross-section ($\sigma_{gg} = 43.92$ pb) for the 125 GeV Higgs boson at $\sqrt{s} = 13$ TeV [67], the ability to trigger on hard leptons and/or jets from the W/Z decay could substantially increase the trigger efficiency. Simulation of the $4\gamma_d$ FRVZ process with associated production could pose difficulties in MadGraph5 because of the high number of final-state particles, 8 in the LJs + 2 from the W/Z decay, especially if the combinatorics of all the different possible W/Z decay products are taken into account. The $4\gamma_d$ FRVZ process with gluon-gluon fusion, yielding 8 final-state particles in the LJs, already pushes the limit of the MadGraph5 simulation capabilities. Limiting the possible W/Z decay products in the simulation, and/or use of another MC generator, may therefore be necessary.

8.3 Combined Analysis

Plans are underway to combine 13 TeV Prompt LJ, Displaced LJ, displaced non-collimated muons, and $Z_D Z_D$ analyses (discussed in Section 3.2.1) to obtain a unified bound on the doubly-exotic Higgs boson decay shown in Figure 3.1. This process was used as the benchmark for the 8 TeV version of the $Z_D Z_D$ analysis [39]; since it is essentially the limit of the 2 $\gamma_{\rm d}$ FRVZ process with no $E_{\rm T}^{\rm miss}$ carried away by the lightest stable hidden particle, LJ results can easily be applied to it. No 13 TeV searches have yet been completed for the Prompt LJ or $Z_D Z_D$ analyses, but they are well underway, and estimates of their reach can be made based on the 8 TeV versions. The displaced non-collimated muons analysis would be new for Run 2, and would essentially serve as a longer-lifetime extension of the $Z_D Z_D$ analysis. Under mildly optimistic assumptions, the displaced non-collimated muon analysis would completely bridge the gap between the LJ and $Z_D Z_D$ searches in $m_{A'}$. Assuming some or all of the extensions discussed in the previous sections are applied to the LJ searches, the gap in ϵ between the Prompt and LJ analyses would be closed, and the ϵ reach would be extended to lower values.

The overall reach of LHC searches for the doubly-exotic Higgs boson decay was originally estimated by the phenomenologist authors of [9], as shown in Figure 8.5 for 300 fb^{-1} of 14 TeV data. The above combined analysis, although at $\sqrt{s}$ of only 13 TeV, would aim to cover most if not all of this estimated reach by the end of Run 2. As can be seen in the Figure, the low-mass reach of the LJ searches has already exceeded the phenomenologists' expectations.

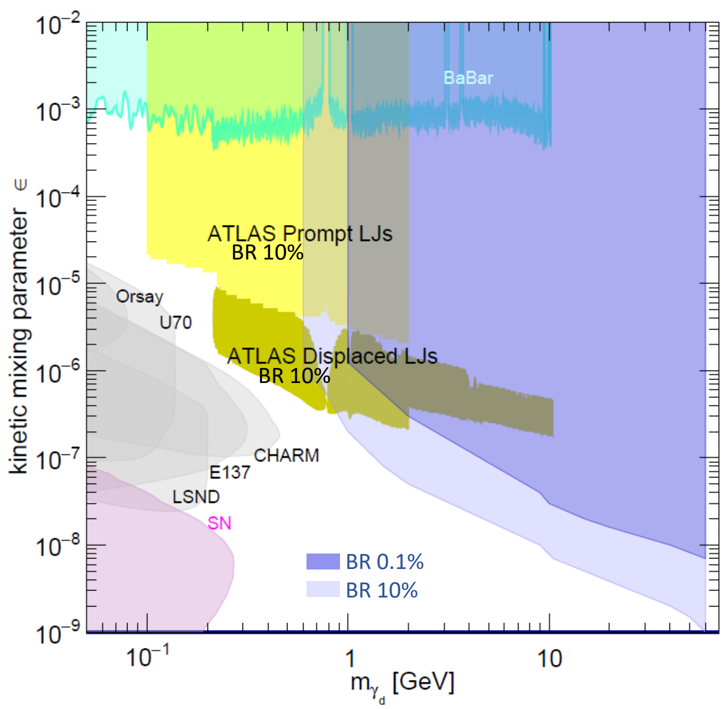

Figure 8.5: Projected overall reach of LHC searches for doubly-exotic Higgs boson decay, estimated by phenomenologists in [9] for 300 fb^{-1} of 14 TeV data. The blue contours correspond to two different values of BR($h \to Z_{\rm D} Z_{\rm D}$). The Displaced LJ (8 TeV + 13 TeV) and Prompt LJ (8 TeV) contours are for BR($h \to Z_{\rm D} Z_{\rm D}$) = 10%, the same BR as the lighter blue contour. Constraints from various other experiments (cyan, gray, violet) are the same as those discussed in Section 3.3.

8.4 Other Benchmark Models

There are a number of other models that could be analyzed as benchmarks. One such possibility is a SUSY hidden sector model (see Section 2.3), whose experimental signature is very similar to the FRVZ models except that it also includes soft final-state jets (Figure 8.6). This model was employed as a benchmark in the 8 TeV Prompt LJs analysis [41]; preliminary studies would have to be performed to determine whether the soft final-state jets could provide a useful additional handle for the Displaced LJs analysis.

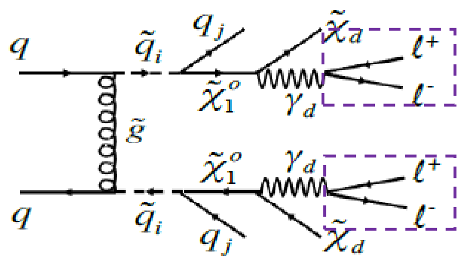

Figure 8.6: SUSY hidden sector model, where the $\tilde{\chi}_d$ (collider-stable dark SUSY particles) manifest as $E_{\mathrm{T}}^{\mathrm{miss}}$ and the q_j (SM quarks of flavour j) manifest as final-state soft jets. Purple boxes indicate LJs.

Another possibility is Drell-Yan A' production (see Section 3.1), where the A' decay products are sufficiently collimated to form an LJ. This would be significantly more challenging, as it would require a mono-LJ search. It would also likely require triggering on a high-p_{T} jet from initial state radiation.

Three more possibilities, Inelastic Dark Matter (introduced in Section 2.4.3), Radiating DM (Section 3.1), and Protophobic Boson models (Section 3.2.2), are discussed in further detail below.

8.4.1 Inelastic Dark Matter

In iDM models, LHC energies would typically be large enough to produce both the lighter and heavier DM mass eigenstates χ_1 and χ_2. Since the dominant DM coupling to SM states is through the off-diagonal interaction, both eigenstates would be produced simultaneously. Certain iDM processes can generate events with a single $A' \to$ LJ in the final state. Furthermore, in promising areas of parameter space in such models, the LJ comes out displaced.

Benchmark Process

The mass hierarchies of particular interest feature $m_{A'} > m_{\chi_1} + m_{\chi_2}$, such that the decay $A' \to \chi_1\chi_2$ occurs on-shell. With such a mass hierarchy, this is not a secluded DM scenario. The decay $\chi_2 \to \chi_1 A'^* \to \chi_1 l^+ l^-$ can occur via an offshell A'. Particularly interesting at the LHC are processes in which the original A' is produced in association with a hard SM object such as a high p_{T} jet. The hard SM object is important because it provides something to trigger on. An example of such a process (Figure 8.7), as suggested in [34], is:

$$pp \to j + A' \to j\chi_1\chi_2 \to jl^+l^-\chi_1\chi_1 \tag{8.1}$$

For $m_{A'}$ at the GeV scale, the A' would be highly boosted when recoiling off the jet. Then the decay products $l^+l^-\chi_1\chi_1$ would come out highly collimated; since the χ_1 is presumed to be collider-stable,

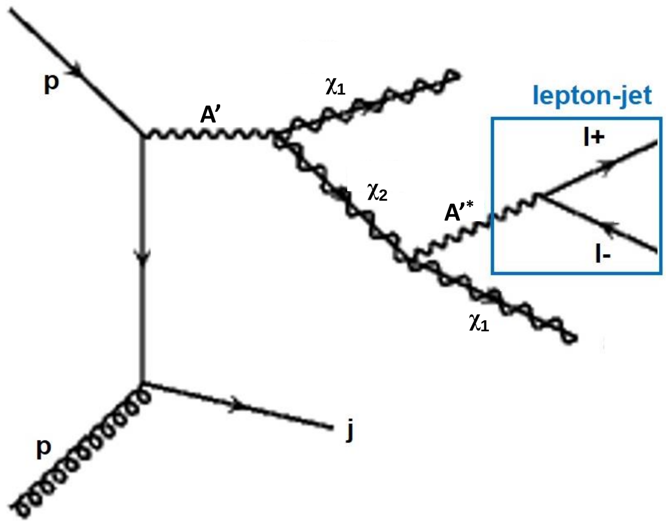

Figure 8.7: The iDM process of interest, where j is a high p_{T} jet.

the resulting experimental signature consists of $E_{\mathrm{T}}^{\mathrm{miss}}$ aligned with a LJ. The lifetimes of the χ_2 and A' determine the displacement of the LJ.

Discovery Potential

The LHC discovery potential for such processes is estimated in detail in [34]. The area of parameter space defined by χ_1, χ_2 masses in the range of hundreds of MeV to tens of GeV, with mass splitting $\Delta = m_{\chi_2} - m_{\chi_1}$ of order $0.01 - 0.1\ m_{\chi_1}$, is currently poorly constrained by direct and indirect detection experiments; constraints are mainly from B-factories, LEP, and $(g-2)_\mu$. In this area, values of DM-SM couplings that typically yield the appropriate thermal relic density make the χ_2 and/or A' long-lived. In the limit of $\Delta m \ll m_{\chi_1} \lesssim m_{A'}$, the DM annihilation rate, and therefore the relic density, depends mostly on m_{χ_1} and the dimensionless interaction strength $y = \epsilon^2 \alpha_D \left(\frac{m_{\chi_1}}{m_{A'}}\right)^4$. Parameter space in this model is therefore typically represented in the $(m_{\chi_1},\ y)$ plane. Experimental bounds, however, often vary with α_D, Δ, and/or $m_{\chi_1}/m_{A'}$ individually, which makes it somewhat complicated to express the bounds in this parameter space. Examples from [34] of existing experimental bounds and estimated Displaced LJ discovery potential are shown in Figure 8.8 for several different values of various parameters. The estimates are intended as "ballpark" figures only; since they do not include modelling of data-driven backgrounds, they simply apply reasonable estimates of trigger and LJ detection efficiencies, and then require 10 surviving signal events with 300 fb^{-1} of 13 TeV data. The salient point is that a Displaced LJ search could cover areas of parameter space compatible with the relic density that are untouched by other experiments. The new sensitivity provided by such a search generally increases with Δ and decreases with $m_{\chi_1}/m_{A'}$. Most of the new parameter space has m_{χ_1} of the order of tens of GeV, with $m_{A'}$ larger by a factor of $\sim 2 - 10$. This puts the A' mass in the territory of current non-collimated Z_{D} searches, rather than the current LJ γ_{d} searches.

Search Challenges

A Displaced LJ-based search for this iDM model poses two main experimental challenges, besides the presence of only one LJ in the final state: the relatively high Z_{D} mass, and the low p_{T} of the LJ constituents.

I employed the custom MadGraph5 model that I built for the FRVZ processes (see Section 5.1.2) to model the iDM process, with:

- The FRVZ HLSP used as the iDM χ_1
- The FRVZ f_{d_2} used as the iDM χ_2
- The HLSP f_{d_2} Z_{D} coupling ($\chi_1\chi_2Z_{\mathrm{D}}$ coupling) set to be non-zero.

Truth-level events were then produced using MadGraph5+PYTHIA8 for ~ 20 points in iDM parameter space, and compared to the MC results in [34] with some help from the authors. Key kinematical distributions for 4 of those parameter points are shown in Figure 8.9; these points were subsequently put through the ATLAS MC production system to create official MC xAOD files. The Z_{D} $c\tau$ in each sample is tuned for a mean lab-frame decay distance of about 4.5 m, such that $80-90\%$ of decays occur within the ATLAS volume; the χ_2 $c\tau$ is tuned for a lab-frame decay distance of about 8.0 cm, such that most decays occur within the ID. The MC demonstrates that given m_{χ_1} values of tens of GeV, only certain combinations of parameters yield small enough opening angle of the offshell Z_{D} decay products (i.e. $\Delta R \lesssim 0.4$) to form LJs. In order to achieve optimal coverage of parameter space, a complementary non-collimated Z_{D} search will be needed.

The low p_{T} of the Z_{D} decay products will greatly reduce the efficiencies of the triggers used in Section 5.4, since a large fraction of the decay products will fail to pass the minimum trigger p_{T} thresholds. Therefore, the analysis will have to trigger on the hard jet instead; the estimates in [34] assume the use of a $p_{\mathrm{T}} > 120$ GeV jet trigger. A significant fraction of the decay products will still fail the analysis-level quality criteria, or will fail to be reconstructed into xAOD physics objects altogether – for example, muon objects are only reconstructed down to 6 GeV. The p_{T} distributions of the A' LJ parent and of the sub-leading constituent of muonic LJs are shown in Figure 8.9.

Given that searching for only one final-state LJ will require additional discriminating variables, a minimum $E_{\mathrm{T}}^{\mathrm{miss}}$ cut could be imposed, as well as a maximum cut on on the $|\Delta\phi|$ between the reconstructed $E_{\mathrm{T}}^{\mathrm{miss}}$ direction and the LJ direction.

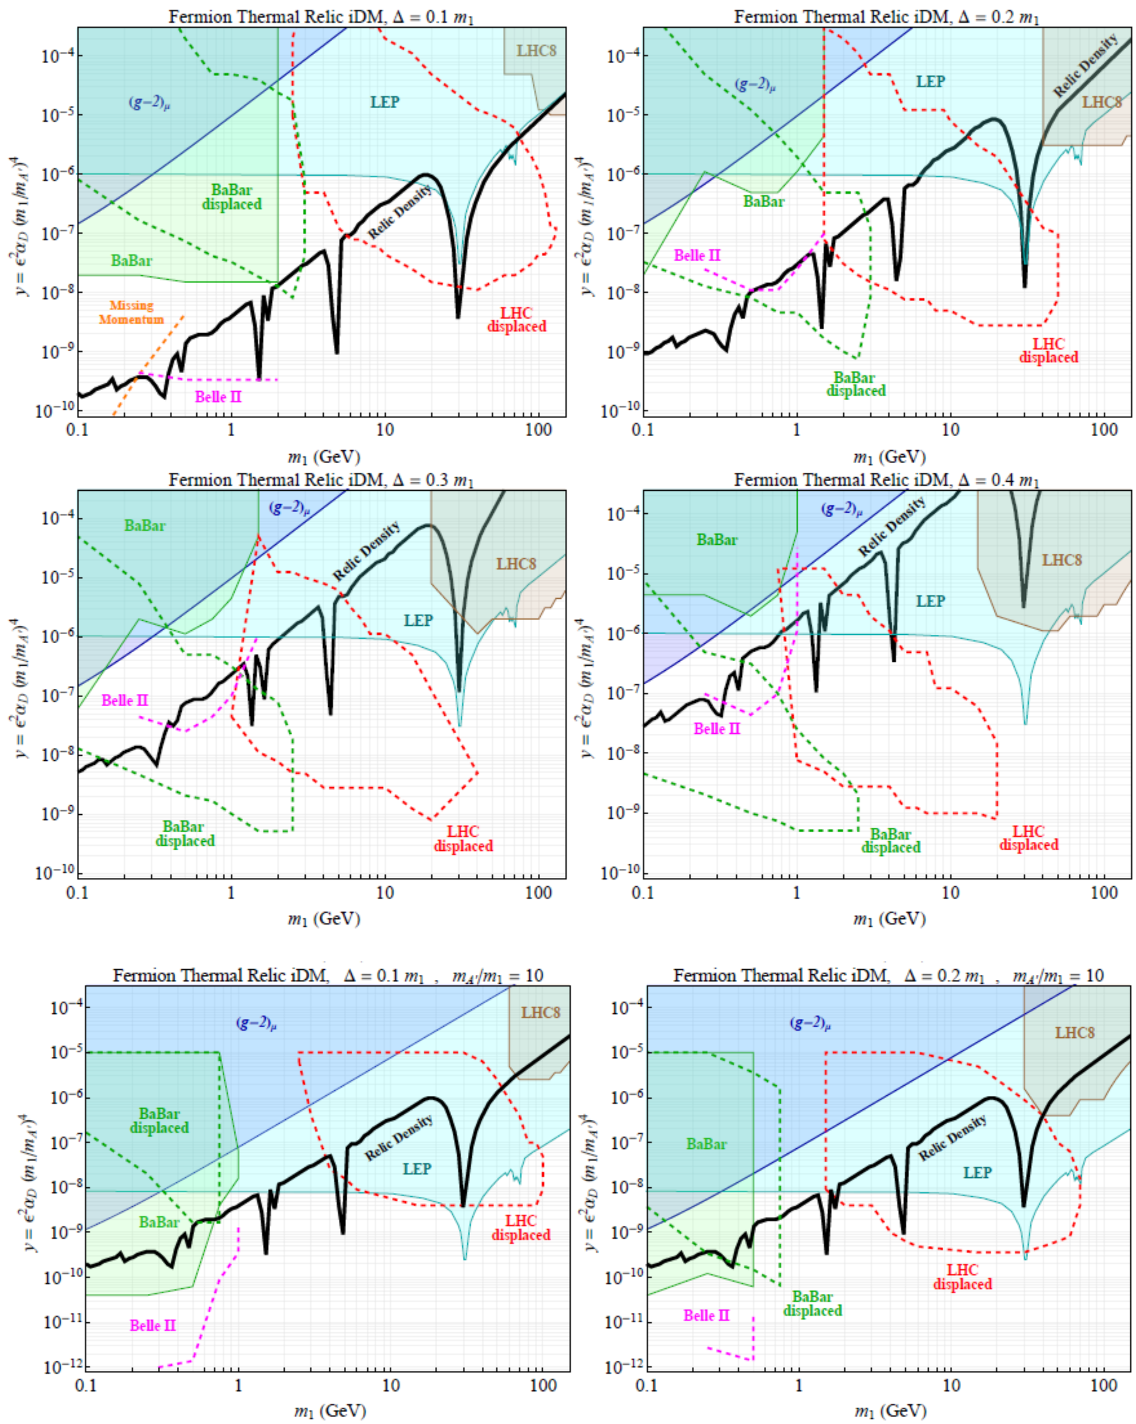

Figure 8.8: Top and middle panels: $m_1 = m_{\chi_1}$, $\alpha_D = 0.1$, $m_{A'}/m_{\chi_1} = 3$. Bottom panel: $m_1 = m_{\chi_1}$, $\alpha_D = 0.1$, $m_{A'}/m_{\chi_1} = 10$. Black line: thermal relic density target. Solid filled areas: ruled out by existing experimental constraints (BaBar $\gamma+E_{\mathrm{T}}^{\mathrm{miss}}$, LEP, $(g-2)_\mu$, re-cast LHC results at 8 TeV). Red dashed line: projections for a Displaced LJ search with 300 fb^{-1} of 13 TeV data. Green dashed line: projections for BaBar $\gamma+E_{\mathrm{T}}^{\mathrm{miss}}$ + displaced lepton searches. Violet line: projections for Belle II $\gamma+E_{\mathrm{T}}^{\mathrm{miss}}$ search. Orange dashed line: projections for a hypothetical fixed-target missing-momentum experiment proposed by the phenomenologist authors. [34]

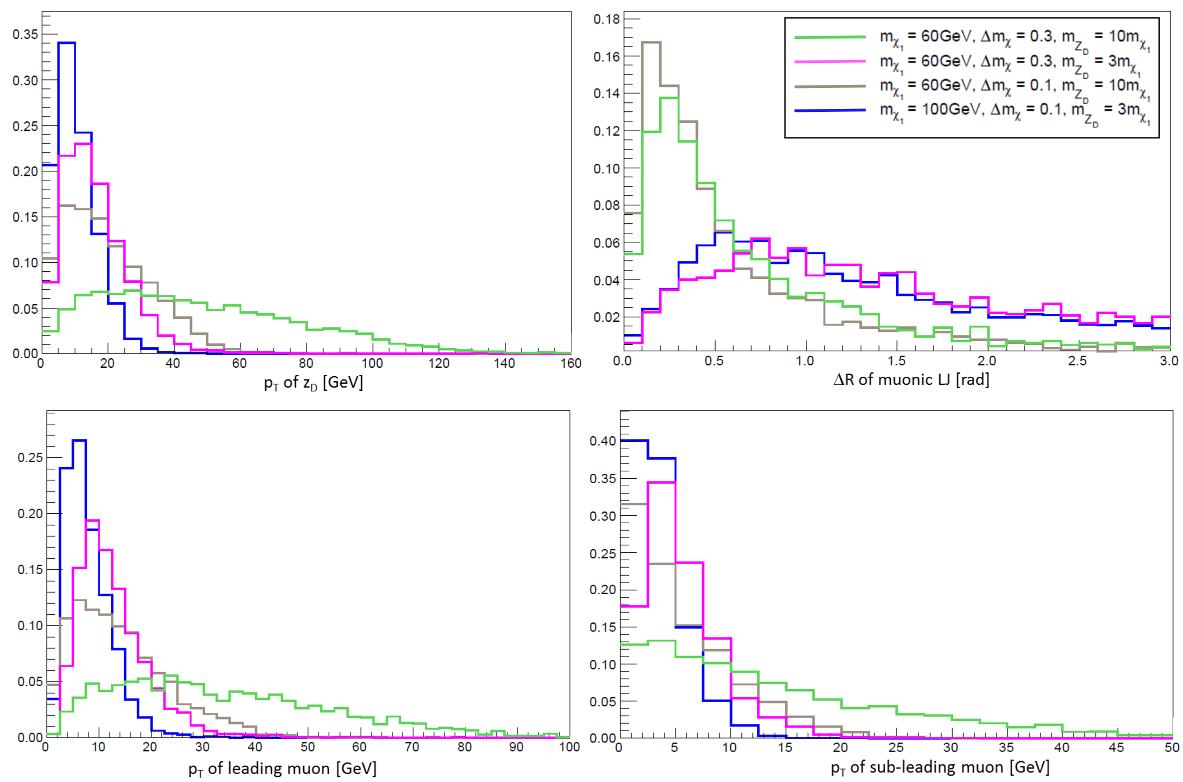

Figure 8.9: Truth-level MC of the iDM process, produced using MadGraph5+PYTHIA8, for various parameter points. Top left: the p_{T} distribution of the offshell Z_{D}. Top right: the opening angle of its muon decay products. Bottom: the p_{T} distribution of its leading and sub-leading muon decay products.

8.4.2 Radiating DM

In radiating DM models, multiple soft γ_{d}'s emitted from the same χ emerge along the same line-of-flight, such that their decay products all form a single LJ. A moderate dark coupling constant, $\alpha_d \sim 0.2$, results in a few γ_{d}'s being radiated for each χ pair. In [40], the dynamics of this dark radiation is explored in detail, both analytically and numerically. The dynamics do not depend upon the DM production mechanism, and are therefore the same in any model that pair-produces the χ at non-negligible rates from pp collisions (Figure 8.10). A re-casting of both the 8 TeV Displaced LJs search and the 7 TeV Prompt LJs search was performed by the authors of [40], along with sensitivity projections for 100 fb^{-1} of 13 TeV data (Figure 8.11). These 13 TeV projections use muonic LJs only, because of the authors' lack of access to the detailed detector simulations that would be necessary to estimate the multijet backgrounds for other LJ types, and assume the same trigger and reconstruction efficiencies as at 7 and 8 TeV. An updated re-casting using the 13 TeV Displaced search and the 8 TeV Displaced search has not been performed. In my discussions with the authors, we agreed that detailed optimization based on 13 TeV MC could yield somewhat better Run 2 limits than using the simple re-casting, and that including Type1 and Type2 LJs would presumably yield substantial improvements.

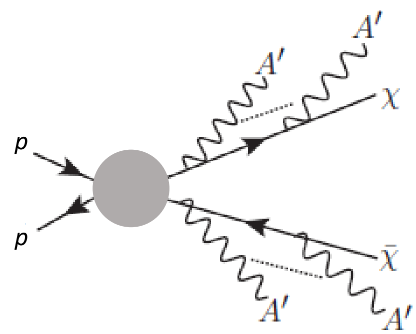

Figure 8.10: Radiating DM in [40], valid for any model that pair-produces the χ at non-negligible rates from pp collisions.

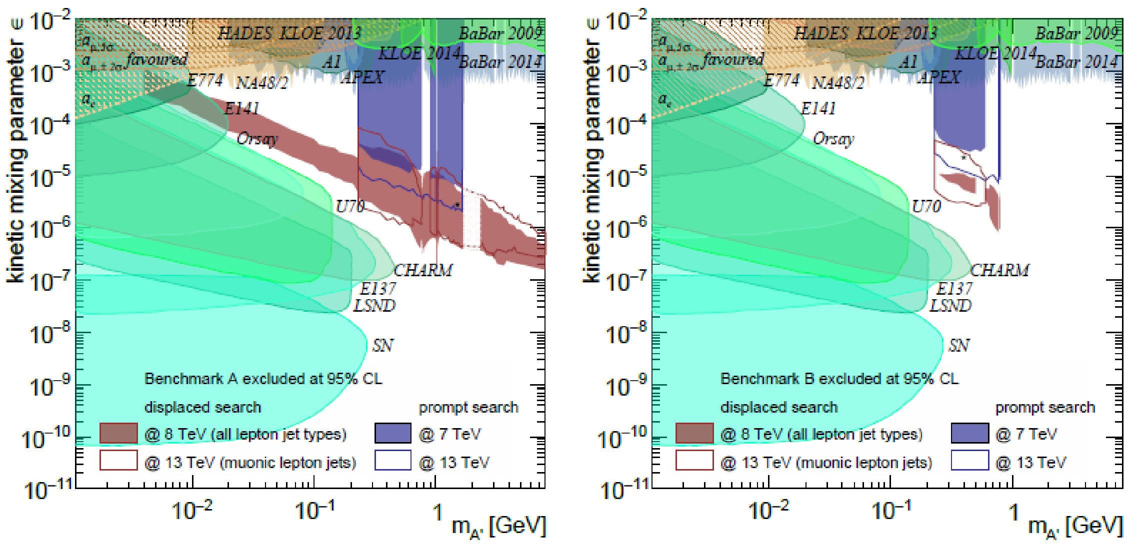

Figure 8.11: Bounds in A' parameter space in radiating DM models ($\alpha_d = 0.2$), obtained in [40] by recasting ATLAS 7 TeV and 8 TeV LJ search results. The 13 TeV contours, which are projections for 100 fb^{-1}, include only Type0 LJs and obtained without detailed detector simulations. Left: $m_\chi = 4$ GeV. Right: $m_\chi = 0.4$ GeV.

8.4.3 Protophobic / Axial Boson Models

As discussed in Section 3.2.2, an A' with mass m_X decaying to electron pairs may mimic the LJ signature of a protophobic/axial boson X decaying to electron pairs. It has been suggested that an anomaly recently observed in transitions of in ^{8}Be could be explained by $m_X \approx 17$ MeV. While the limits in Section 6.5.2 were obtained by excluding the Type2-Type2 LJ channel, the use of only the Type2 channel can place limits on such X boson models, assuming the X can somehow be produced from Higgs boson decays. The Type2-Type2 limits should therefore be translatable into limits on the X, if the $A \times Eff$ for the γ_d at m_X could be reliably determined and turned out to be non-negligible. An example is shown in Figure 8.12, for the case where an 800 GeV BSM Higgs boson decays to the X boson, if hypothetically the $A \times Eff$ at m_X were the same as the $A \times Eff$ at ~ 400 MeV. Unfortunately, in actuality $m_X = 16.7$ MeV is significantly below the mass range probed in the Displaced LJ search (see Section 6.5.2); specific full MC simulations would have to be performed for $m_X = 16.7$ MeV to determine whether the LJ searches have any reasonable sensitivity.

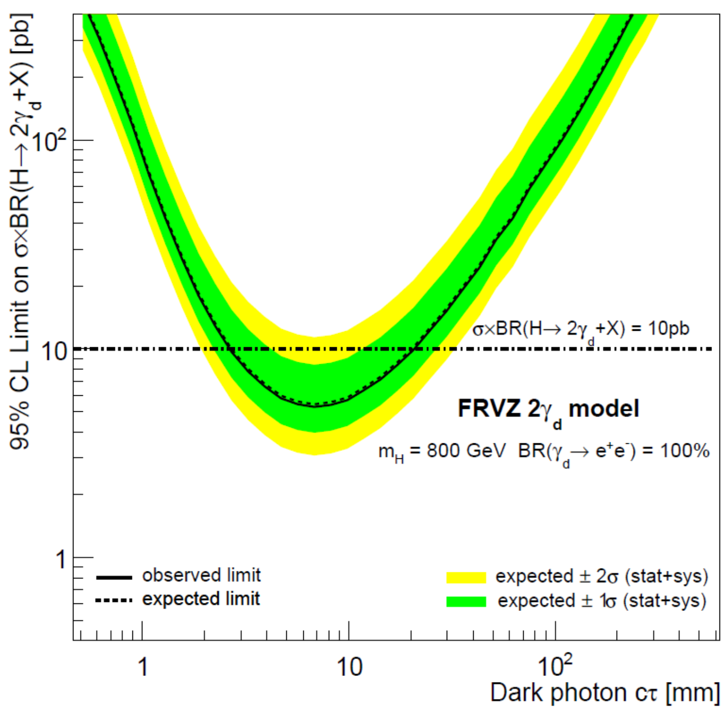

Figure 8.12: The 95% upper limits on $\sigma\times$ BR as a function of the γ_d $c\tau$, for the $m_H = 800$ GeV $2\gamma_\mathrm{d}$ FRVZ model with $m_{\gamma_\mathrm{d}} = 16.7$ MeV. Assumes γ_d decay to electron pairs (Type2-Type2 LJs) only. The expected limit is shown as the dashed curve, and the observed limit as the almost-identical solid curve. The horizontal lines correspond to $\sigma\times$ BR $= 10$ pb. The excluded range is $2.7 \leq c\tau \leq 21$ mm.

8.5 Summary

Increased luminosity is just one of several potential sources of improvement for future ATLAS LJ analyses. One avenue is to strengthen the limits on the current benchmark models. The gap between the Prompt and Displaced limits, at $\epsilon \sim 5 \times 10^{-6}$ to 1×10^{-5}, could be closed by reconstructing tracks that begin partway through the Pixel subsystem of the ID. In addition, more signal events could be detected by including in the search more final-state topologies that can arise in the models, such as a single LJ or associated Higgs boson production. With these improvements, combining the LJ searches with other ATLAS A' searches would yield coverage of an impressively large swath of parameter space. Another avenue is to expand the LJ analyses to other DM models. Preliminary studies indicate that Inelastic DM, Radiating DM, and protophobic/axial boson models are promising candidates.

Chapter 9

Conclusions

In the ongoing quest for answers to the DM mystery, a dark gauge boson A' appears in various promising models, both "top-down" (theoretically motivated from the realm of BSM physics) and "bottom-up" (phenomenologically motivated from astrophysical, direct and indirect detection observations). Mixing between the A' and the SM photon, controlled by the kinetic mixing parameter ϵ, would serve as a portal between the SM and the dark sector. Particularly interesting for the LHC are models in which the A' gets its mass from a dark Higgs boson, which mixes with the SM Higgs boson to create a Higgs portal as well. For A' masses on the order of hundreds of MeV to a few GeV, i.e. the γ_{d} case, highly-boosted γ_{d}'s emerging from LHC pp collisions and then decaying back to the SM would give rise to a distinctive experimental signature. This signature consists of lepton-jets, i.e. collimated pairs of light leptons and mesons, which would be displaced in the case of long-lived γ_{d}'s .

This work has presented the results of a search for displaced LJs, in pp collisions at $\sqrt{s} = 13$ TeV in a sample of 3.57 fb^{-1} collected during 2015 with the ATLAS detector at the LHC. Highlights are shown in the ATLAS conference note [6]; no deviations from Standard Model expectations are observed. Limits on benchmark models predicting Higgs boson decays to long-lived γ_{d}'s , which in turn produce displaced LJs, are derived as a function of the γ_{d} proper lifetime $c\tau$. The branching ratio of the 125 GeV Higgs boson to γ_{d}'s is found to be below 10% (at 95% CL) in the range $2.2 < c\tau < 111.3$ mm in a $H \to 2\gamma_{\mathrm{d}}$ + X benchmark model, and in the range $3.8 < c\tau < 163$ mm in a $H \to 4\gamma_{\mathrm{d}}$ + X model. These extend the limits established in a similar search previously performed during Run 1 of the LHC (20.3 fb^{-1} of $\sqrt{s} = 8$ TeV data), largely thanks to enhancements in the trigger efficiency and to the $\sqrt{s}$ increase. In addition, constraints on a heavy BSM Higgs boson are considered in these benchmark models for the first time. The cross section times branching ratio of an 800 GeV Higgs boson to γ_{d}'s is found to be below 5 pb (at 95% CL) in the range $0.6 < c\tau < 63$ mm in the $H \to 2\gamma_{\mathrm{d}}$ + X model, and in the range $0.8 < c\tau < 186$ mm in the $H \to 4\gamma_{\mathrm{d}}$ + X model. The constraints on the benchmark models are then translated into the $(m_{A'}, \epsilon)$ parameter space that has become standard in the literature for dark photon models. Constraints from the Displaced LJ analysis are complementary in ϵ to those from Prompt LJ analyses. The ATLAS constraints, although dependent upon an extra parameter representing the Higgs boson branching ratio to A's, cover a region in this parameter space that is untouched by other types of experiments.

The analysis strikes a balance between the model-independent and benchmark-model approaches, by defining and selecting generic Displaced LJ objects and then setting limits in the context of broadlyrepresentative benchmark models. The acceptance times efficiency plots produced for the different types of Displaced LJs, as a function of γ_d p_T and decay length, can be employed by phenomenologists to place ballpark limits on nearly any model featuring final-state Displaced LJs. The analysis cuts that follow emphasize accessible experimental quantities rather than opaque multivariable discriminants, further increasing the utility of the results to phenomenologists. The successful "re-casting" of the Run 1 results into a Radiating DM framework already demonstrates the value of this approach. Meanwhile, the FRVZ models chosen as benchmarks for limit-setting are representative of a broad class of models that include both the hypercharge and Higgs portals. While some requirements on the final state (e.g. at least two well-separated Displaced LJs) must be enforced in order to reduce the backgrounds, the analysis cuts are largely independent of final-state variables that depend on the exact dark sector mass parameters of the models, with no invariant mass or $E_\mathrm{T}^\mathrm{miss}$ cuts.

As Run 2 data-taking continues at the LHC, the discovery or limit-setting potential of LJ searches will continue to improve with increasing luminosity. Several other planned enhancements to the LJ searches – including the incorporation of muon timing information, large-radius inner detector track reconstruction, mono-LJ analyses, and W/Z associated production analyses – have also been put forward. The contours in parameter space presented in this work could be expanded by performing additional MC at $m_{A'} > 10$ GeV and obtaining a usable parameterization of γ_d BRs in this mass range. Then, unifying the enhanced LJ searches with complementary ATLAS searches at even higher masses would provide complete coverage of a much larger swath of parameter space.

Looking at the long-term picture, it can be reasonably expected that the limits presented in this work will soon be surpassed. The main goal of this work, in fact, is to lay the groundwork for these future developments. Therefore, the primary takeaways from this work should not be the exact values of the limits themselves, but rather the techniques that allow Displaced LJ searches to be conducted in the challenging Run 2 experimental conditions, and the potential applications of these searches. In the post-Higgs boson discovery era, a sea of BSM physics possibilities is circulating in the literature, while vast previously-favoured regions of WIMP parameter space have been ruled out by LHC data. Searches must now look in more unusual places for the next discovery, entailing a shift toward exotic signatures, innovative triggers, and non-standard reconstruction strategies. LJ analyses are on the front line of this shift, with a bright future in the ATLAS search program and solid niche in the global DM quest.

Appendix A

This appendix explores the possibility of using timing measurements from the RPCs in the MS (Section 4.4.2) to construct a discriminating variable against cosmic ray backgrounds. Such a discriminating variable could be used in addition to, or in place of, the muon impact parameter $|z_0|$, and was employed in the $\sqrt{s} = 8$ TeV Displaced LJs analysis. The investigation of its possible usage for Run 2 data was a joint effort between myself and the Run 1 timing experts in the analysis group, Antonio Policicchio and Monica Verducci. The RPC timing information was not fully available in the xAOD files used in this analysis for the 2015 data, but will eventually be available after a planned central re-processing. Eliminating the cosmic ray background at the LJ selection stage, prior to the event-level selection, would leave only multijet background for the data-driven ABCD method. This would greatly reduce the systematic uncertainty associated with the method, particularly the non-closure component (described in Section 7.1). It would also likely reduce the final background contamination in the signal region, improving the overall analysis sensitivity.

RPC Time Measurements

The time recorded by an RPC chamber for a muon hit is automatically corrected by the presumed time-of-flight of the muon, assuming the muon originates from the IP (see Section 4.4.2). The corrected time distribution for muons from the IP is therefore centred at 0 for each RPC chamber. Cosmic ray muons, however, are generally downward-going. In the upper hemisphere of the detector, $\phi > 0$, they travel the wrong way with respect to the time-of-flight assumption (see Figure A.1). The corrected time on each chamber therefore comes out negative.

It was found that a more robust method than using the RPC times on individual chambers is to use the difference in RPC times between chambers. For this purpose, the RPC is divided into layers as follows:

- Layer 1: BMS, $L_{xy} \leq 7.75$ m
- Layer 2: BML, $7.75 \leq L_{xy} \leq 9.00$ m
- Layer 3: BOS + BOL, $L_{xy} \geq 9.00$ m

The quantities $t_3 - t_1,\ t_3 - t_2,\ t_2 - t_1$, where t_x is the corrected time for layer x, are then candidates for discriminating variables. Since a muon often generates multiple nearby hits in a given RPC chamber, t can be assigned either as the earliest of these hit times, or as the mean of a Gaussian fit to these hit times.

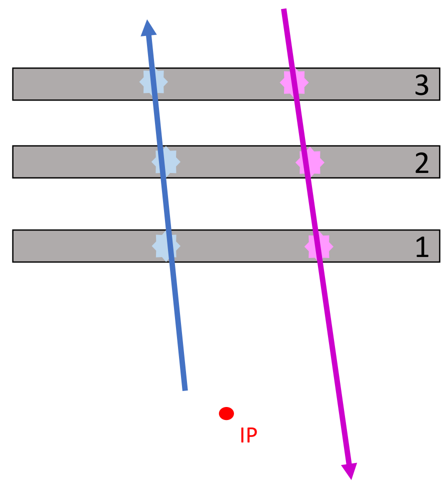

Figure A.1: Displaced muon (blue) vs cosmic ray muon (purple) travelling through three stations in the upper hemisphere of the RPC. The cosmic ray muon is going the wrong way, i.e. toward the IP, with respect to the time-of-flight assumption.

Preliminary Tests

Distributions of the quantities $t_3 - t_1,\ t_3 - t_2,\ t_2 - t_1$ (referenced collectively here as ΔT quantities) in the upper detector hemisphere were constructed using data from 3 different streams of a single 13 TeV run (Figure A.2):

- Express stream, containing minimum-bias events in filled bunch crossings
- CosmicMuons stream, containing events passing L1Muon-seeded triggers in empty bunch crossings
- CosmicCalo stream, containing events passing L1Calo-seeded triggers in empty bunch crossings.

The cosmic ray muons constituting backgrounds to the Displaced LJ search would mostly closely resemble the muons in the CosmicMuons and CosmicCalo samples. The timing distributions of the signal muons would resemble those of the muons in the Express sample. Even though the signal muons are displaced from the IP, while the Express muons generally originate from the IP, this makes only a small difference in the timing distribution. The average values of each of the ΔT quantities for $\phi > 0$ show significant separation between CosmicMuons/CosmicCalo and Express (Table A.1), indicating good prospects for using one of these quantities as a discriminating variable. For example, requiring $t_3 - t_1 > -10$ for $\phi > 0$ would reject $\gtrsim 97\%$ of cosmic muons with a signal efficiency of $\sim 99\%$.

Stream	$t_3 - t_1$ [ns]	$t_3 - t_2$ [ns]	$t_2 - t_1$ [ns]
CosmicMuons	-19	-15	-8
CosmicCalo	-21	-17	-7
Express	0	-1	-2

Table A.1: Average values for each of the ΔT quantities – where each t value is taken as the mean of a Gaussian fit to the muon's hit times in an RPC chamber – for $\phi > 0$ in the Express, CosmicMuons and CosmicCalo streams.

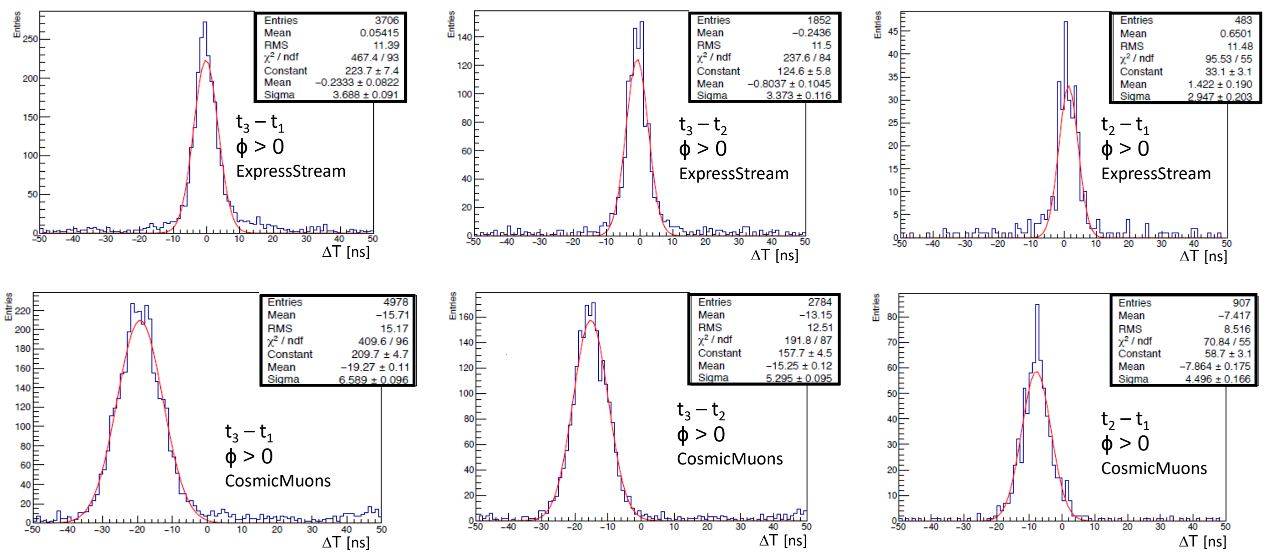

Figure A.2: Distributions of each of the ΔT quantities – where each t value is taken as the mean of a Gaussian fit to the muon's hit times in an RPC chamber – for $\phi > 0$ in the Express, CosmicMuons and CosmicCalo streams.

Bibliography

[1] P. Ade et al. Planck 2013 results. XVI. Cosmological parameters. *Astron. Astrophys.*, 571:A16, 2014, 1303.5076.

[2] The ATLAS Collaboration. The ATLAS Experiment at the CERN Large Hadron Collider. *Journal of Instrumentation*, 3(08):S08003, 2008.

[3] S. Chatrchyan et al. The CMS experiment at the CERN LHC. *JINST*, 3:S08004, 2008.

[4] G. Aad et al. Search for long-lived neutral particles decaying into lepton jets in proton-proton collisions at $\sqrt{s} = 8$ TeV with the ATLAS detector. *JHEP*, 11:088, 2014, 1409.0746.

[5] M. A. Deliyergiyev. Recent Progress in Search for Dark Sector Signatures. 2015, 1510.06927.

[6] Search for long-lived neutral particles decaying into displaced lepton jets in proton–proton collisions at $\sqrt{s} = 13$ TeV with the ATLAS detector. Technical Report ATLAS-CONF-2016-042, CERN, Geneva, Aug 2016.

[7] B. Holdom. Two U(1)'s and ϵ Charge Shifts. *Phys. Lett.*, 166B(2):196, 1986.

[8] P. Galison and A. Manohar. Two Z's or Not Two Z's? *Phys. Lett.*, 136B(4):279, 1984.

[9] D. Curtin, R. Essig, S. Gori, and J. Shelton. Illuminating Dark Photons with High-Energy Colliders. *JHEP*, 02:157, 2015, 1412.0018.

[10] B. Dobrescu and I. Mocioiu. Spin-dependent macroscopic forces from new particle exchange. *JHEP*, 11:005, 2006, hep-ph/0605342.

[11] J. Wells. How to Find a Hidden World at the Large Hadron Collider. 2008, 0803.1243.

[12] D. Curtin et al. Exotic decays of the 125 GeV Higgs boson. *Phys. Rev.*, D90(7):075004, 2014, 1312.4992.

[13] M. Strassler and K. Zurek. Echoes of a hidden valley at hadron colliders. *Phys. Lett.*, B651:374–379, 2007, hep-ph/0604261.

[14] A. Falkowski, J. Ruderman, T. Volansky, and J. Zupan. Hidden Higgs Decaying to Lepton Jets. *JHEP*, 05:077, 2010, 1002.2952.

[15] J. Kumar and J. Wells. CERN LHC and ILC probes of hidden-sector gauge bosons. *Phys. Rev.*, D74:115017, 2006, hep-ph/0606183.

[16] M. Goodsell, J. Jaeckel, J. Redondo, and A. Ringwald. Naturally Light Hidden Photons in LARGE Volume String Compactifications. *JHEP*, 11:027, 2009, 0909.0515.

[17] O. Adriani et al. The PAMELA Mission: Heralding a new era in precision cosmic ray physics. *Physics Reports*, 544(4):323 – 370, 2014.

[18] A. Abdo et al. Measurement of the Cosmic Ray e+ plus e- spectrum from 20 GeV to 1 TeV with the Fermi Large Area Telescope. *Phys. Rev. Lett.*, 102:181101, 2009, 0905.0025.

[19] J. Chang et al. An excess of cosmic ray electrons at energies of 300 - 800 GeV. *Nature*, 456:362–365, 2008.

[20] G. Dobler, D. Finkbeiner, I. Cholis, T. Slatyer, and N. Weiner. The Fermi Haze: A Gamma-ray Counterpart to the Microwave Haze. *ApJ*, 717:825–842, July 2010, 0910.4583.

[21] T. Daylan, D. Finkbeiner, D. Hooper, T. Linden, S. Portillo, N. Rodd, and T. Slatyer. The characterization of the gamma-ray signal from the central Milky Way: A case for annihilating dark matter. *Phys. Dark Univ.*, 12:1–23, 2016, 1402.6703.

[22] G. Weidenspointner et al. The sky distribution of 511-keV positron annihilation line emission as measured with INTEGRAL/SPI. *ESA Spec. Publ.*, 622:25, 2007, astro-ph/0702621.

[23] R. Bernabei et al. First results from DAMA/LIBRA and the combined results with DAMA/NaI. *Eur. Phys. J.*, C56:333–355, 2008, 0804.2741.

[24] R. Agnese et al. Silicon Detector Dark Matter Results from the Final Exposure of CDMS II. *Phys. Rev. Lett.*, 111(25):251301, 2013, 1304.4279.

[25] C. Aalseth et al. CoGeNT: A Search for Low-Mass Dark Matter using p-type Point Contact Germanium Detectors. *Phys. Rev.*, D88:012002, 2013, 1208.5737.

[26] J. Ruderman and T. Volansky. Decaying into the Hidden Sector. *JHEP*, 02:024, 2010, 0908.1570.

[27] J. Redondo and M. Postma. Massive hidden photons as lukewarm dark matter. *JCAP*, 0902:005, 2009, 0811.0326.

[28] M. Pospelov, A. Ritz, and M. Voloshin. Secluded WIMP Dark Matter. *Phys. Lett.*, B662:53–61, 2008, 0711.4866.

[29] M. Pospelov and A. Ritz. Astrophysical Signatures of Secluded Dark Matter. *Phys. Lett.*, B671:391–397, 2009, 0810.1502.

[30] N. Arkani-Hamed and N. Weiner. LHC Signals for a SuperUnified Theory of Dark Matter. *JHEP*, 12:104, 2008, 0810.0714.

[31] N. Arkani-Hamed, D. Finkbeiner, T. Slatyer, and N. Weiner. A Theory of Dark Matter. *Phys. Rev.*, D79:015014, 2009, 0810.0713.

[32] C. Cheung, J. Ruderman, L. Wang, and I. Yavin. Kinetic Mixing as the Origin of Light Dark Scales. *Phys. Rev.*, D80:035008, 2009, 0902.3246.

[33] D. Finkbeiner and N. Weiner. Exciting Dark Matter and the INTEGRAL/SPI 511 keV signal. *Phys. Rev.*, D76:083519, 2007, astro-ph/0702587.

[34] E. Izaguirre, G. Krnjaic, and B. Shuve. Discovering Inelastic Thermal-Relic Dark Matter at Colliders. *Phys. Rev.*, D93(6):063523, 2016, 1508.03050.

[35] D. Finkbeiner, T. Slatyer, N. Weiner, and I. Yavin. PAMELA, DAMA, INTEGRAL and Signatures of Metastable Excited WIMPs. *JCAP*, 0909:037, 2009, 0903.1037.

[36] J. Ruderman and T. Volansky. Searching for Smoking Gun Signatures of Decaying Dark Matter. 2009, 0907.4373.

[37] J. Cline, G. Dupuis, Z. Liu, and W. Xue. The windows for kinetically mixed Z'-mediated dark matter and the galactic center gamma ray excess. *JHEP*, 08:131, 2014, 1405.7691.

[38] I. Hoenig, G. Samach, and D. Tucker-Smith. Searching for dilepton resonances below the Z mass at the LHC. *Phys. Rev.*, D90(7):075016, 2014, 1408.1075.

[39] G. Aad et al. Search for new light gauge bosons in Higgs boson decays to four-lepton final states in pp collisions at $\sqrt{s} = 8$ TeV with the ATLAS detector at the LHC. *Phys. Rev. D*, 92:092001, Nov 2015.

[40] M. Buschmann, J. Kopp, J. Liu, and P. Machado. Lepton Jets from Radiating Dark Matter. *JHEP*, 07:045, 2015, 1505.07459.

[41] G. Aad et al. A search for prompt lepton-jets in pp collisions at $\sqrt{s} = 8$ TeV with the ATLAS detector. *JHEP*, 02:062, 2016, 1511.05542.

[42] V. Khachatryan et al. A search for pair production of new light bosons decaying into muons. *Phys. Lett.*, B752:146–168, 2016, 1506.00424.

[43] A. Krasznahorkay et al. Observation of Anomalous Internal Pair Creation in Be8 : A Possible Indication of a Light, Neutral Boson. *Phys. Rev. Lett.*, 116(4):042501, 2016, 1504.01527.

[44] J. Kozaczuk, D. Morrissey, and S. Stroberg. Light Axial Vectors, Nuclear Transitions, and the ^{8}Be Anomaly. 2016, 1612.01525.

[45] J. Helo, M. Hirsch, and S. Kovalenko. Heavy neutrino searches at the LHC with displaced vertices. *Phys. Rev.*, D89:073005, 2014, 1312.2900.

[46] M. Endo, K. Hamaguchi, and G. Mishima. Constraints on Hidden Photon Models from Electron g-2 and Hydrogen Spectroscopy. *Phys. Rev.*, D86:095029, 2012, 1209.2558.

[47] H. Davoudiasl, H. Lee, and W. Marciano. Dark Side of Higgs Diphoton Decays and Muon g-2. *Phys. Rev.*, D86:095009, 2012, 1208.2973.

[48] S. Abrahamyan et al. Search for a New Gauge Boson in Electron-Nucleus Fixed-Target Scattering by the APEX Experiment. *Phys. Rev. Lett.*, 107:191804, Nov 2011.

[49] H. Merkel et al. Search for Light Gauge Bosons of the Dark Sector at the Mainz Microtron. *Phys. Rev. Lett.*, 106:251802, Jun 2011.

[50] M. Davier and H. Ngoc. An unambiguous search for a light higgs boson. *Physics Letters B*, 229(1):150 – 155, 1989.

[51] J. Blumlein and J. Brunner. New Exclusion Limits for Dark Gauge Forces from Beam-Dump Data. *Phys. Lett.*, B701:155–159, 2011, 1104.2747.

[52] J. Blumlein and J. Brunner. New Exclusion Limits on Dark Gauge Forces from Proton Bremsstrahlung in Beam-Dump Data. *Phys. Lett.*, B731:320–326, 2014, 1311.3870.

[53] E. Riordan et al. Search for short-lived axions in an electron-beam-dump experiment. *Phys. Rev. Lett.*, 59:755–758, Aug 1987.

[54] A. Bross, M. Crisler, S. Pordes, J. Volk, S. Errede, and J. Wrbanek. Search for short-lived particles produced in an electron beam dump. *Phys. Rev. Lett.*, 67:2942–2945, Nov 1991.

[55] S. N. Gninenko. Constraints on sub-GeV hidden sector gauge bosons from a search for heavy neutrino decays. *Phys. Lett.*, B713:244–248, 2012, 1204.3583.

[56] R. Essig, R. Harnik, J. Kaplan, and N. Toro. Discovering new light states at neutrino experiments. *Phys. Rev. D*, 82:113008, Dec 2010.

[57] J. Lees et al. Search for a Dark Photon in e^+e^- Collisions at BaBar. *Phys. Rev. Lett.*, 113(20):201801, 2014, 1406.2980.

[58] B. Aubert et al. Search for Dimuon Decays of a Light Scalar Boson in Radiative Transitions $\Upsilon \to \gamma A^0$. *Phys. Rev. Lett.*, 103:081803, Aug 2009.

[59] J. Batley et al. Search for the dark photon in π^0 decays. *Phys. Lett.*, B746:178–185, 2015, 1504.00607.

[60] D. Babusci et al. Limit on the production of a light vector gauge boson in ϕ meson decays with the KLOE detector. *Physics Letters B*, 720(1 - 3):111 – 115, 2013.

[61] F. Archilli et al. Search for a vector gauge boson in ϕ meson decays with the KLOE detector. *Physics Letters B*, 706(4-5):251 – 255, 2012.

[62] G. Agakishiev et al. Searching a Dark Photon with HADES. *Phys. Lett.*, B731:265–271, 2014, 1311.0216.

[63] J. Dent, F. Ferrer, and L. Krauss. Constraints on Light Hidden Sector Gauge Bosons from Supernova Cooling. 2012, 1201.2683.

[64] H. Dreiner, J. Fortin, C. Hanhart, and L. Ubaldi. Supernova constraints on MeV dark sectors from e^+e^- annihilations. *Phys. Rev.*, D89(10):105015, 2014, 1310.3826.

[65] R. Essig et al. Working Group Report: New Light Weakly Coupled Particles. In *Proceedings, Community Summer Study 2013: Snowmass on the Mississippi (CSS2013): Minneapolis, MN, USA, July 29-August 6, 2013*, 2013, 1311.0029.

[66] A. Falkowski, J. Ruderman, T. Volansky, and J. Zupan. Discovering Higgs Decays to Lepton Jets at Hadron Colliders. *Phys. Rev. Lett.*, 105:241801, 2010, 1007.3496.

[67] D. de Florian et al. Handbook of LHC Higgs Cross Sections: 4. Deciphering the Nature of the Higgs Sector. 2016, 1610.07922.

[68] A. Yamamoto et al. The ATLAS central solenoid. *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment*, 584(1):53 – 74, 2008.

[69] P. Akesson et al. ATLAS Tracking Event Data Model. Technical Report ATL-SOFT-PUB-2006-004. ATL-COM-SOFT-2006-005. CERN-ATL-COM-SOFT-2006-005, CERN, Geneva, Jul 2006.

[70] G. Aad et al. ATLAS pixel detector electronics and sensors. *Journal of Instrumentation*, 3(07):P07007, 2008.

[71] M. Capeans et al. ATLAS Insertable B-Layer Technical Design Report. Technical Report CERN-LHCC-2010-013. ATLAS-TDR-19, Sep 2010.

[72] A. Ahmad et al. The silicon microstrip sensors of the ATLAS semiconductor tracker. *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment*, 578(1):98 – 118, 2007.

[73] The ATLAS TRT Collaboration. The ATLAS Transition Radiation Tracker (TRT) proportional drift tube: design and performance. *Journal of Instrumentation*, 3(02):P02013, 2008.

[74] *ATLAS liquid-argon calorimeter: Technical Design Report.* Technical Design Report ATLAS. CERN, Geneva, 1996.

[75] A. Artamonov et al. The ATLAS Forward Calorimeter. *Journal of Instrumentation*, 3(02):P02010, 2008.

[76] J. Abdallah et al. Design, Construction and Installation of the ATLAS Hadronic Barrel Scintillator-Tile Calorimeter. Technical Report ATL-TILECAL-PUB-2008-001. ATL-COM-TILECAL-2007-019, CERN, Geneva, Nov 2007.

[77] G. Aad et al. Jet energy measurement with the ATLAS detector in proton-proton collisions at $\sqrt{s} = 7$ TeV. *Eur. Phys. J.*, C73(3):2304, 2013, 1112.6426.

[78] S. Palestini. The muon spectrometer of the ATLAS experiment. *Nuclear Physics B - Proceedings Supplements*, 125:337 – 345, 2003.

[79] G. Aad et al. Muon reconstruction performance of the ATLAS detector in proton proton collision data at $\sqrt{s}$ =13 TeV. *Eur. Phys. J.*, C76(5):292, 2016, 1603.05598.

[80] M. Aaboud et al. Performance of the ATLAS Trigger System in 2015. 2016, 1611.09661.

[81] V. Cindro et al. The ATLAS Beam Conditions Monitor. *Journal of Instrumentation*, 3(02):P02004, 2008.

[82] J. Alwall, M. Herquet, F. Maltoni, O. Mattelaer, and T. Stelzer. MadGraph 5 : Going Beyond. *JHEP*, 06:128, 2011, 1106.0522.

[83] T. Sjostrand, S. Mrenna, and P. Skands. A Brief Introduction to PYTHIA 8.1. *Comput. Phys. Commun.*, 178:852–867, 2008, 0710.3820.

[84] G. Aad et al. The ATLAS Simulation Infrastructure. *Eur. Phys. J.*, C70:823–874, 2010, 1005.4568.

[85] I. Bird et al. Update of the Computing Models of the WLCG and the LHC Experiments. Technical Report CERN-LHCC-2014-014. LCG-TDR-002, Apr 2014.

[86] M. Cacciari, G. Salam, and G. Soyez. The Anti-k(t) jet clustering algorithm. *JHEP*, 04:063, 2008, 0802.1189.

[87] Selection of jets produced in 13TeV proton-proton collisions with the ATLAS detector. Technical Report ATLAS-CONF-2015-029, CERN, Geneva, Jul 2015.

[88] Jet Calibration and Systematic Uncertainties for Jets Reconstructed in the ATLAS Detector at $\sqrt{s} = 13$ TeV. Technical Report ATL-PHYS-PUB-2015-015, CERN, Geneva, Jul 2015.

[89] G. Aad et al. Triggers for displaced decays of long-lived neutral particles in the ATLAS detector. *JINST*, 8:P07015, 2013, 1305.2284.

[90] G. Aad. Performance of pile-up mitigation techniques for jets in pp collisions at $\sqrt{s} = 8$ TeV using the ATLAS detector. *Eur. Phys. J.*, C76(11):581, 2016, 1510.03823.

[91] A. Hocker et al. TMVA - Toolkit for Multivariate Data Analysis. *PoS*, ACAT:040, 2007, physics/0703039.

[92] G. Cowan, K. Cranmer, E. Gross, and O. Vitells. Asymptotic formulae for likelihood-based tests of new physics. *Eur. Phys. J.*, C71:1554, 2011, 1007.1727. [Erratum: Eur. Phys. J.C73,2501(2013)].

[93] D Pomarde and M Virchaux. The Persint visualization program for the ATLAS experiment. *eConf*, C0303241(cs.GR/0305057):MOLT009. 9 p, May 2003.

[94] A. Read. Presentation of search results: The CL(s) technique. *J. Phys.*, G28:2693–2704, 2002. [,11(2002)].

[95] ATLAS LUCID Collaboration. Measurement of the luminosity with the new LUCID-2 detector in 2015. Technical Report ATL-COM-FWD-2016-008, CERN, Geneva, May 2016.

[96] Muon reconstruction efficiency in reprocessed 2010 LHC proton-proton collision data recorded with the ATLAS detector. Technical Report ATLAS-CONF-2011-063, CERN, Geneva, Apr 2011.

[97] J. Gaiser. *Charmonium Spectroscopy From Radiative Decays of the J/ψ and ψ'*. PhD thesis, SLAC, 1982.

[98] Search for pair produced long-lived neutral particles decaying in the ATLAS hadronic calorimeter in pp collisions at $\sqrt{s} = 8$ TeV. Technical Report ATLAS-CONF-2014-041, CERN, Geneva, Jul 2014.

[99] T. Cornelissen et al. Concepts, Design and Implementation of the ATLAS New Tracking (NEWT). Technical Report ATL-SOFT-PUB-2007-007. ATL-COM-SOFT-2007-002, CERN, Geneva, Mar 2007.