

SEARCHES FOR SUB-GEV DARK MATTER WITH NOVA AND LDMX AND
PERFORMANCE STUDIES OF THE COSMIC RAY VETO FOR MU2E

Tyler Gregory Horoho
Kokomo, Indiana

Bachelor of Science, Indiana University Bloomington, 2018

A Thesis submitted to the Graduate Faculty
of the University of Virginia in Candidacy for the Degree of
Doctor of Philosophy
Department of Physics

University of Virginia
June 27, 2025

Edmond Craig Dukes, Chair
Robert Craig Group
Julian Heeck
Bradley Johnson

Searches for sub-GeV Dark Matter with NOvA and LDMX and Performance Studies of the Cosmic Ray Veto for Mu2e

Tyler Gregory Horoho

(ABSTRACT)

The constituents of dark matter are still unknown, and the viable possibilities span a very large mass range. Specific scenarios for the origin of dark matter sharpen the focus to within about an MeV to 100 TeV. Most of the stable constituents of known matter have masses in the lower range, and a thermal origin for dark matter works in a simple and predictive manner in this mass range as well. Non-gravitational couplings of dark matter to Standard Model particles would make it possible to produce dark matter at accelerators. Through 120 GeV proton-nucleus collisions, NOvA is effectively a beam dump experiment for producing and detecting dark matter. This thesis presents a search for dark matter – electron scattering in the NOvA near detector from sub-GeV dark matter produced at the target. A future search for sub-GeV dark matter with the Light Dark Matter eXperiment (LDMX) is discussed, along with a study of LDMX’s expected sensitivity to visibly decaying dark photons. Additionally, performance studies of the Mu2e cosmic ray veto (CRV) are presented. A calibration procedure for the CRV counters is discussed, along with tests of the efficiency, position resolution, and light yield decline of the CRV counters.

Dedication

To my sun, my moon, and all my stars, Grace.

Acknowledgments

First and most importantly, my wife, Grace, for her unwavering support over the past five years and beyond. She has been completely understanding whenever coursework consumed my personal time or when my research demanded that we move from the city we love to the Chicago suburbs. I cannot envision a more considerate and kind partner to accompany me in my physics journey.

My advisor, Craig Dukes, for his commitment to support my career goals and for providing copious opportunities for me to grow and succeed as a scientist. In addition, I want to thank Craig Group for his mentorship on a significant portion of the work discussed in this thesis.

My fellow graduate students, particularly Sid Ajith, Hunter Presley, Diego Ibarra Hoyos, Misho Sokhashvili, Ben Wade, and Jason Pioquinto, for their friendship. I am grateful for the time spent working together over Zoom in our first year to overcome the qual and all of the trips to go rock climbing, kayaking, and hiking.

My mentors and colleagues on LDMX, Mu2e, and NOvA for fostering a welcoming environment and for their effort in my development as a scientist. In particular, I want to thank Matt Solt, Yuri Oksuzian, Ralf Ehrlich, and Alex Himmel for their invaluable mentorship. I also want to thank those whom I worked alongside in advancing the work presented in this thesis: Jessica Pascadlo, Kieran Wall, Muve, Barnali Brahma, Josh Greaves, and Yongyi Wu.

This is not an exhaustive list of all the people who left a positive impact on my throughout the course of my PhD, and I apologize for any omissions.

Contents

List of Figures	xii
List of Tables	xviii
1 Introduction and Motivation	1
1.1 Overwhelming Evidence for Dark Matter	2
1.1.1 Galaxy Rotation Curves	3
1.1.2 Baryon Acoustic Oscillations	4
1.1.3 Weak Gravitational Lensing	8
1.1.4 Multi-messenger Astronomy Constraints	10
1.2 Thermal Relic Dark Matter	12
1.2.1 The WIMP Miracle?	14
1.2.2 Light Dark Matter	16
1.2.2.1 Empirical Motivations for Vector Portal Dark Matter	20
1.2.2.2 Detection Methods	21
I Search for Light Dark Matter with NOvA	27
2 The NOvA Experiment	28

2.1	The NuMI Beam	28
2.2	The Near Detector	31
2.3	Event Reconstruction	36
3	Dark Matter Analysis	40
3.1	Dark Matter Production and Detection	41
3.1.1	Production by NuMI	41
3.1.2	Detection by the Near Detector	42
3.2	Strategy for Dark Matter Search	44
3.3	Data Samples	48
3.3.1	Neutrino Simulation	49
3.3.1.1	Beam Simulation	50
3.3.1.2	Neutrino Interactions with GENIE	51
3.3.1.3	Detector Response	52
3.3.1.4	Dedicated Samples of ν - e and MEC Events	52
3.3.2	Dark Matter Simulation	54
3.3.2.1	Simulation of Production and Scattering	54
3.3.2.2	Generic Single-Electrons Samples	61
3.3.2.3	Signal Template Generation by Sample Re-weight	63
3.4	Event Selection	67

3.4.1	Single Electron Event Selection	67
3.4.2	Overlay Effect on Event Reconstruction and Selection	71
3.5	Systematic Uncertainties	75
3.5.1	Detector Response Uncertainties	75
3.5.2	Flux Uncertainties	76
3.5.3	Cross Section Uncertainties	77
3.6	Analysis Results	79
3.7	Discussion	84
II	Visibly Decaying Dark Photons with LDMX	87
4	The Light Dark Matter eXperiment	88
4.1	The Beamline	90
4.2	The Detector	91
4.2.1	Tracking System	91
4.2.2	Electromagnetic Calorimeter	93
4.2.3	Hadronic Calorimeter	95
4.3	Simulation Framework	97
4.4	Missing-Momentum Search	99
5	Visibly Decaying Dark Photons in LDMX	105

5.1	Simulated Event Samples	107
5.1.1	Visible A' Simulation	107
5.1.2	Background Simulation	107
5.1.3	HCal Reconstruction	109
5.2	Event Selection	110
5.2.1	HCal Selection Cuts	111
5.2.1.1	The Visibles Boosted Decision Tree	112
5.2.2	Backgrounds	117
5.2.2.1	Photo-nuclear backgrounds	118
5.2.3	Muon Pair Production	119
5.2.4	Total Background	119
5.2.5	Full Dataset Prospects	119
5.3	Signal Rates and Sensitivity Estimates	121
5.3.1	A' Efficiency	121
5.3.2	A' Rates	122
5.3.3	Sensitivity Estimates	124
5.4	Discussion	127
III	Performance Studies of the Mu2e Cosmic Ray Veto	131
6	The Mu2e Experiment	132

6.1	The Mu2e Experimental Apparatus	133
6.1.1	Muon beamline	133
6.1.2	The Detectors	134
6.2	The Cosmic Ray Veto	136
6.2.1	The Cosmic Ray Veto Detector Design	137
6.2.2	Readout of the Cosmic Ray Veto	140
7	Performance studies of the Mu2e CRV	143
7.1	Cosmic Ray Test Stand	143
7.1.1	Trigger Paddles	145
7.2	Data Processing	147
7.2.1	Calibration	147
7.2.1.1	Calibration at a Stable Temperature	147
7.2.1.2	Temperature-Dependent Calibration	149
7.2.1.3	Obtaining the Reference Calibration	151
7.2.2	Reconstructing the Light Yield	152
7.3	Efficiency Studies	153
7.3.1	Event Selection	154
7.3.2	Measurement of the Track Stub Efficiency	156
7.3.3	Monte Carlo Studies	159

7.3.4	Efficiency without Defining Modules	163
7.4	Position Resolution Study	165
7.5	Aging of CRV Modules	170
7.6	Extrapolation to Overall CRV Efficiency	175
7.7	Discussion	176
	Bibliography	178
	Appendices	196

List of Figures

1.1	The Standard Model of particle physics.	3
1.2	The rotation curve of NGC 6503.	5
1.3	The Cosmic Microwave Background.	6
1.4	The power spectrum of the Cosmic Microwave Background.	7
1.5	Image of the Bullet cluster.	9
1.6	The mechanism of thermal freeze-out.	13
1.7	A summary of experimental sensitivities to the dark matter-nucleon cross section for GeV-scale dark matter.	16
1.8	A one-loop kinetic mixing process between the A' and photon, with massive fermions Ψ that carry both electric charge and $U_D(1)$ charge.	18
1.9	Comparison of sub-GeV thermal relic targets in direct detection searches and accelerator-based searches.	22
1.10	Dark photon production mechanisms.	23
1.11	A cartoon showing dark matter production (left) and re-scattering (right).	24
2.1	The Fermilab accelerator complex.	29
2.2	The NuMI beam line.	31
2.3	The neutrino energy in various off-axis scenarios.	32

2.4	The expected dark matter flux in various off-axis scenarios.	32
2.5	Illustration of a NOvA ND cell with wavelength-shifting fiber.	33
2.6	A 3D view of the NOvA detector.	34
2.7	An example NOvA event display showing multiple interactions in a single event.	37
2.8	The result of the vertex finding algorithm.	38
2.9	Result of the Fuzzyk prong finding algorithm.	39
3.1	The signal and background shapes of the $E\theta^2$ variable.	46
3.2	Example of how the 90% confidence level exclusion value for the like- likelihood ratio was determined.	48
3.3	The ratio of pion flux between MIPP and BdNMC.	57
3.4	The NOvA dark matter flux correction procedure.	58
3.5	The expected number of dark matter events at 1×10^{21} POT.	60
3.6	The “generic” dark matter energy distribution.	62
3.7	Signal energy and angle correlation.	62
3.8	Example ratio histograms used in the re-weighting procedure.	64
3.9	The simulation and re-weighting procedure of the single electron events.	65
3.10	A comparison of the $E\theta^2$ spectra between re-weighted generic samples and the “true” distributions.	66
3.11	The discriminator values for the NuoneID and Epi0ID neural networks.	70

3.12	Event selection efficiency as a function of $E\theta^2$	72
3.13	The single-electron event reconstruction and selection efficiency on non-overlay and overlay samples.	74
3.14	The best-fit prediction compared to data.	81
3.15	The nominal prediction compared to data in E and θ	82
3.16	NOvA's 90% CL exclusion limit compared to other experiments. . . .	83
4.1	A plan view of the LDMX detector.	89
4.2	A diagram of the LDMX beamline.	91
4.3	The magnetic field strength as a function of position along the beamline.	92
4.4	A diagram of the LDMX ECal.	94
4.5	ECal response to MIPs and electromagnetic showers.	95
4.6	Isometric view of the LDMX HCal.	97
4.7	The relative rates of background processes in LDMX, along with the veto handles used to reject them.	99
4.8	Transverse momentum (p_T) distribution of the recoil electron for trig- gered events.	101
4.9	Distribution of BDT discriminator values and the maximum number of photo-electrons in the HCal for 1 GeV signal and PN background events.	103
4.10	The projected sensitivity of LDMX to sub-GeV dark matter.	104

5.1	The LDMX apparatus with a visibles sketch overlaid.	106
5.2	Six of the twelve BDT features.	113
5.3	The other six BDT features.	114
5.4	Visibles BDT ROC curve (left) and discriminator value output (right).	116
5.5	The FOM as a function of the discriminator value on the Visibles BDT for four test masses.	118
5.6	Final signal efficiency as a function of the A' decay position after all cuts. The shaded region represents the coverage of the HCal.	122
5.7	Signal efficiency through each step of the selection cuts.	123
5.8	The fraction of beam energy carried away by the A'	125
5.9	The angle of the A' momentum with respect to the beam axis.	125
5.10	LDMX's expected sensitivity to a visibly decaying A' signal.	126
6.1	Isometric drawing of the Mu2e apparatus.	133
6.2	Mu2e's signal window.	135
6.3	The Mu2e cosmic ray veto detector.	137
6.4	A polystyrene scintillator extrusion with two cavities for wavelength- shifting fibers, coated with a TiO_2 -polystyrene mixture.	138
6.5	Exploded view of the manifold at the end of a di-counter showing the fiber guide bar, SiPM mounting block, SiPM carrier boards, SiPMs, and counter motherboard.	139

6.6	Architecture of the CRV readout for a “normal” module read out on both ends.	140
7.1	The cosmic-ray test stand at the Wideband hall at Fermilab.	144
7.2	A diagram of a trigger paddle used in the cosmic-ray test stand. . . .	145
7.3	Example ADC waveform from the two SiPM channels of one end of a CRV counter, 127 ADC samples each.	146
7.4	The pedestal-subtracted dark count distribution from a single SiPM displays distinct peaks that correspond to one, two, and three photoelectrons.	148
7.5	A flow chart of the temperature-dependent calibration procedure of the CRV.	149
7.6	Typical SiPM photoelectron distribution example from cosmic-ray muons, fitted with a convoluted Gauss-Landau function.	154
7.7	Distributions of muon and non-muon events for the three parameters used in event selection cuts.	156
7.8	The single layer inefficiency of a CRV-T module as a function of the PE threshold.	157
7.9	The measured 3/4 inefficiency compared to the extrapolated single-layer efficiency.	158
7.10	A cartoon illustrating two muon tracks passing through a module. . .	159
7.11	Comparison of the light yield in simulation and data.	161

7.12	The directly measured 3/4 inefficiency and the extrapolated 3/4 inefficiency from the single layer efficiency for MC and data.	162
7.13	Event displays for simulated events with similar hit patterns to data events that fail the 3/4 veto requirement.	164
7.14	Comparison of the single-layer inefficiency with and without cleanup cuts.	166
7.15	The time offset between AFE chips in FEBs on opposite ends of the module.	168
7.16	The reconstructed position of cosmic-ray events from TOA differences.	170
7.17	An example of the SPE calibration constant for a “bad” channel over time.	172
7.18	Light yield over time of an example SiPM. Each point corresponds to a sub-run.	174
7.19	Distribution of measured aging in all counters with 1.8 mm WLS fibers.	174
7.20	The difference in the measured aging rate of two SiPMs in the same counter.	175

List of Tables

3.1	Composition of NOvA ND.	59
3.2	Summary of NOvA dark matter samples.	67
3.3	NOvA selection cut efficiency for backgrounds.	71
3.4	NOvA selection cut efficiency for the signal.	71
3.5	Summary of how 1σ systematic uncertainties affect the normalization of ν - N and ν - e backgrounds and χ - e signal events.	78
3.6	Summary of the predicted number of background events.	80
3.7	A summary of the largest systematic pulls in the dark matter best fit and neutrino-only fit.	82
4.1	Summary of selection cut effect on signal and backgrounds in the LDMX missing-momentum search.	103
5.1	Summary of filters and weights used in the background simulations. .	109
5.2	Number of background and signal events used for training the Visibles BDT.	116
5.3	Summary of selection cut effect on signal and backgrounds in the LDMX visibles search.	120
6.1	Expected backgrounds in Mu2e Run I.	136

7.1	Temperature correction parameters for the Mu2e CRV calibration. . .	150
-----	---	-----

Chapter 1

Introduction and Motivation

The Standard Model of particle physics is our most complete and accurate picture describing the composition of the Universe and interactions within it at the most fundamental level. Decades of discovery for particles predicted by the Standard Model have been a great triumph for particle physics, culminated by the discovery of the Higgs boson in 2012 [1, 2]. The particle content of the Standard Model, shown in Figure 1.1, contains six quarks and six leptons that compose the matter of everyday life, four gauge bosons that are responsible for three of the four fundamental forces of nature (electromagnetism, the weak nuclear interaction, and the strong nuclear interaction), and the Higgs boson, which provides a mechanism for particles to obtain mass.

However, there are several empirical observations that cannot be addressed within the framework of the Standard Model:

- Observations in astronomy indicate that the Universe contains a large amount of “dark matter,” which has not been directly observed and so far has only been inferred through its gravitational influence on visible matter. A more comprehensive discussion of the evidence for dark matter and explanations for it is given below.
- The Universe is dominated by matter over antimatter, which is possible in

the Standard Model through processes that violate charge-parity (CP) symmetry. However, the amount of CP violation predicted by the Standard Model is not sufficient to explain the observed matter-antimatter asymmetry of the Universe [3].

- Neutrinos have been observed to oscillate in their flavor as they propagate, and therefore have mass [4, 5]. Neutrinos do not acquire mass through the same mechanism as other fermions in the Standard Model and should consequently be massless.

Many theories beyond the Standard Model (BSM) can explain the above phenomena, and predictions in these theories are within reach of current- and next-generation particle physics experiments.

1.1 Overwhelming Evidence for Dark Matter

There is an overwhelming amount of evidence for the existence of dark matter. Over the last roughly 100 years, astronomers have continued to observe anomalies in data that point in the direction of new physics, and yet particle physicists have come up empty for an explanation. Evidence for the existence of dark matter is found on multiple distance scales in the Universe and through multiple detection mechanisms.

The first hint for the existence of dark matter came in 1933, when Fritz Zwicky used the virial theorem to calculate the total mass of the Coma Cluster from the observed velocities of nebulae in the cluster [6]. Zwicky found that the total mass to keep the cluster gravitationally bound must be significantly larger than the observed luminous matter, which he attributed to “dark matter.” Zwicky’s work was largely unnoticed




Figure 1.1: The Standard Model of particle physics is a group of the known elementary particles, consisting of six quarks, six leptons, four gauge bosons, and the Higgs boson.

for decades, and it was not until the 1970s that more evidence for dark matter became too apparent to ignore.

1.1.1 Galaxy Rotation Curves

In the late 1970s, Vera Rubin, with Kent Ford and Norbert Thonnard, observed that the orbital velocity of stars and diffuse gas in the outer galactic disk of spiral galaxies was too large to be accounted for by ordinary matter¹ alone [7]. Along with luminous matter from stars, Rubin included measurements from diffuse gas,

¹“ordinary matter” = matter understood within the context of the Standard Model of particle physics.

which came from 21 cm light emitted by the spin-flip transition of electrons in neutral hydrogen. This allowed the orbital velocity measurement to extend well beyond where luminous matter is emitted from the galactic disk. Figure 1.2 shows the galaxy rotation curve for the spiral galaxy NGC 6503 and how the distribution of visible and dark matter would affect the observed rotation speeds.² Outside of the stellar disk of the galaxy, the orbital velocity is expected to fall off as $v^2 = GM/r$ from Newtonian dynamics. Measurements show instead that the rotational velocity remains constant or slightly increases at large radii. From the Newtonian dynamics formula above, a constant rotational velocity would be satisfied if the mass distribution increased linearly with radius, meaning there is a large amount of mass—comparable to the amount of mass seen in the luminous matter of the galaxy—lies outside of the stellar disk, called the dark matter “halo.”

1.1.2 Baryon Acoustic Oscillations

Galaxy rotation curves indicate that dark matter halos have approximately 3–10 times more mass than the visible matter in galaxies. More precise measurements on the total amount of dark matter in the Universe relative to the amount of baryonic matter³ are measured from anisotropies of the Cosmic Microwave Background (CMB) radiation [9].

The CMB is the thermal “afterglow” of the Big Bang. Conditions were too hot in the early universe for stable atoms to form, and the Universe was filled with an opaque

²Zwicky also noted that observations of the rotational velocity of spiral galaxies by van Maanen were too high to keep the galaxy bound by the amount of luminous matter [6].

³Technically, “baryonic matter” should only constitute protons and neutrons, but I will use this term to refer to protons, neutrons, and electrons in the context of early Universe physics.




Figure 1.2: The rotation curve of NGC 6503 (points). The best fit is shown as the solid curve, along with the contribution from the stellar disk (dashed curve), diffuse gas (dotted curve), and the dark matter halo (dot dashed curve). From Ref. [8].

plasma.⁴ When the Universe expanded and cooled enough for stable atoms to form, the thermal radiation of the Universe could travel freely, which can now be detected as a blackbody spectrum peaking in the microwave frequency due to cosmological redshift. The CMB radiation is almost perfectly isotropic; it has a mean temperature of 2.7 K and root mean square variations of around $100 \mu\text{K}$ [9]. Variations away from the mean temperature in the CMB are shown in Figure 1.3.

These small perturbations away from perfect isotropy are due to fluctuations in the

⁴Thomson scattering of light by charged particles constituting the plasma is why the Universe was opaque.




Figure 1.3: The Cosmic Microwave Background has small temperature fluctuations due to over/under-dense regions in the early Universe.

density of baryonic matter caused by acoustic density waves in the primordial plasma of the early Universe, known as baryon acoustic oscillations (BAO). These acoustic waves are caused by the interplay between the gravitational attraction of all matter and the repulsion of baryonic matter from outward exerted radiation pressure in the over-dense regions. Dark matter stays at the center of these over-dense regions since it has a very small or no coupling to photons, so it does not oscillate like the baryons. Therefore, by looking at the size and frequency of baryon acoustic oscillations, which is embedded in the density fluctuations of the CMB, a measurement of both the total matter content and the baryonic matter content in the Universe can be made.

The total matter density Ω_m and the baryonic matter density Ω_b are determined as parameters in a fit to the observed power spectrum of the CMB, shown in Figure 1.4. The Planck Collaboration measured the total matter density of the Universe to be $\Omega_m h^2 = 0.1430 \pm 0.0011$, where h is the dimensionless Hubble parameter $h = H_0/(100 \text{ km/s/Mpc}) \approx 0.674$ [9]. Of this, the contributions from dark matter $\Omega_\chi h^2$ and baryonic matter $\Omega_b h^2$ are $\Omega_\chi h^2 = 0.1200 \pm 0.0012$ and $\Omega_b h^2 = 0.02237 \pm 0.00015$.

This means there is about five times as much dark matter in the Universe as baryonic matter according to CMB measurements.




Figure 1.4: The power spectrum of the Cosmic Microwave Background. The red dots represent measurements from the Planck Collaboration [9], and the blue curve represents the best-fit model, which assumes the presence of cold dark matter and a cosmological constant (Λ CDM).

The observed amount of baryonic matter from the CMB is in agreement with predictions from Big Bang Nucleosynthesis (BBN), which describes the production of hydrogen, deuterium, helium, and lithium in the first few minutes after the Big Bang [10]. Measurements of the abundance of light elements in the Universe agree remarkably well with the inferred baryon density from baryon acoustic oscillations [11]. This provides an independent confirmation that a large amount of non-baryonic matter is needed to explain the total matter content measured from the CMB power spectrum.

1.1.3 Weak Gravitational Lensing

The abundance of dark matter in the Universe has remained relatively stable on timescales that we can observe. The CMB provides a snapshot of the dark matter abundance at around 380,000 years after the Big Bang, when the Universe was cool enough for neutral hydrogen to form. The abundance of dark matter billions of years after the Big Bang is gauged through weak gravitational lensing of the large-scale structure of the Universe [12]. Gravitational lensing is the deflection of light from distant objects through the gravitational warping of spacetime caused by massive objects in the foreground. Surveys of galaxy clusters and the gravitational lensing caused by these galaxy clusters show that there is not nearly enough luminous matter in the clusters to account for the amount of lensing observed [12]. The amount of dark matter in the large-scale structure required to match the observed lensing is consistent with measurements on the CMB: about five times more dark matter than baryonic matter. These measurements also support the idea that dark matter played a key role in the large-scale structure formation of the Universe. Lensing measurements showed a filamentary structure to the distribution of mass in the Universe, and galaxies were distributed throughout these filaments. This structure mapped from gravitational lensing is also consistent with mapped catalogs of stars, galaxies, and quasars. Simulations of the evolution of the Universe to form this structure require the presence of dark matter [13].

Alternative explanations to dark matter, such as large-scale modifications to gravity [14], are strongly disfavored by weak gravitational lensing observations of the mergers of galaxy clusters [15, 16]. During the collision of galaxy clusters, visible matter in the clusters interacted and slowed, whereas the majority of the mass continued largely unaffected. The distribution of visible matter in the clusters is mapped

by looking at visible light and X-rays—indicative of star formation in the extreme environment of a galaxy cluster merger. The distribution of mass is mapped from the observation of weak gravitational lensing in the cluster. Figure 1.5 shows an image of the Bullet cluster and the distribution of both visible matter and mass. The discrepancy between the distribution of visible matter and the distribution of mass in the Bullet cluster cannot be explained by a large-scale modification to gravity (without dark matter needing to be introduced anyway) and provides strong evidence for a dark matter explanation to all of the above astronomical observations.




Figure 1.5: Image of the Bullet cluster. The blue contours show the projected mass density from gravitational lensing. The magenta contours show the X-ray brightness map measured by the Chandra X-ray Observatory. From Ref. [16].

The Bullet cluster observation also places stringent constraints on dark matter self-interactions. Because the dark matter component of the cluster merger remains spatially separated from the hot gas (which experiences significant drag and slows down), the dark matter must be nearly collisionless. Observations of the Bullet cluster place a limit on the self-interaction cross section per unit mass of dark matter to be less than around $1 \text{ cm}^2/\text{g}$ [17].

Along the same lines, simulations of collisionless dark matter fail to account for some observations on the scale of galaxies and small groups of galaxies. The density of collisionless dark matter in galaxies is predicted to increase sharply toward the galactic center, but observations of dwarf galaxies and low-surface-brightness galaxies show flat central density profiles [18]. Additionally, simulations predict that the Milky Way’s dark matter halo should contain several massive subhalos with satellite galaxies. These subhalos are “too big to fail” at forming stars, meaning they would inevitably have become visible galaxies. However, the number of observed bright satellites around the Milky Way is fewer than what simulations predict with collisionless dark matter [19]. A possible solution to these two problems is if dark matter has a small, velocity-dependent self-interaction [20].⁵

1.1.4 Multi-messenger Astronomy Constraints

To further emphasize that all of the observed phenomena are due to some yet-to-be discovered particle and not the misunderstood nature of gravity [22], multi-messenger astronomy measurements of neutron star mergers strongly constrain alternative theories of gravity. Neutron star mergers are a source of gravitational radiation from the inspiraling of massive objects and electromagnetic radiation from the resulting kilonova of the merger. Modifications to gravity that explain dark matter phenomena predict that there would be a significant delay (by as much as several hundred days) between the observed gravitational wave signal and the electromagnetic signal, but the two signals are observed nearly simultaneously, as expected from general relativity [23].

⁵A velocity-independent cross section for dark matter is strongly disfavored, as an excess of dark matter annihilation signals would be observed [21].

While there is a large amount of evidence for dark matter through a variety of indirect measurements at multiple cosmic distance scales, the true nature of dark matter remains a mystery. These indirect measurements also leave the possible identity of dark matter largely unconstrained; the lightest viable candidates are about 10^{-25} times the mass of an electron and the heaviest viable candidates are several solar masses⁶, or about 10^{62} times the mass of an electron. This is about 90 orders of magnitude to search. What we do know about dark matter is summarized as follows:

1. Its presence has been observed indirectly through gravitational effects, but never directly by “seeing” it. Therefore, dark matter must have no interaction or a very feeble interaction with photons, which implies it has little or no electric (or magnetic) charge. It also does not strongly self-interact.
2. There is about five times as much dark matter in the Universe as baryonic matter.
3. Dark matter abundance has remained constant over time (or at least since recombination), so dark matter is stable or has a lifetime far greater than the age of the Universe.
4. Evidence for the existence of dark matter is observed across the whole sky. With this and the cosmological principle⁷, one can assume dark matter exists everywhere, including throughout the Milky Way and around Earth.
5. Since dark matter seemingly played a large role in large-scale structure formation in the Universe, it must be non-relativistic.

⁶Although not completely ruled out, primordial black holes are very tightly constrained and thus are not favored [24, 25].

⁷The cosmological principle is the notion that the Universe is homogeneous and isotropic, at least when viewed on large enough distance scales.

1.2 Thermal Relic Dark Matter

Exploring Point 2 above, the fact that there is roughly the same amount of dark matter as baryonic matter (at least within the same order of magnitude) is a compelling reason for one to assume that there may be an interaction between dark matter and ordinary matter that connects their origin in the early Universe. Under this assumption—that dark matter and ordinary matter were in thermal equilibrium with each other in the early Universe through some interaction—the mass range for potential dark matter candidates becomes much more tightly constrained, between the sub-MeV to ~ 100 TeV range [26]. In the early Universe, when the temperature was much higher than the dark matter mass and the density was high, creation and annihilation processes were in thermal equilibrium. As the Universe expanded and the temperature fell below the dark matter mass, creation processes became inefficient and the dark matter number density began to annihilate away. At some later time, the annihilation rate became lower than the Hubble expansion rate and annihilation processes also became inefficient. At this point, the dark matter abundance was “frozen out” to a stable value, which is observed today. A diagram of this freeze-out mechanism is shown in Figure 1.6.

Starting from the Boltzmann equation, the observed abundance of dark matter Ω_χ is used to determine the necessary thermally averaged annihilation cross section $\langle\sigma v\rangle$ for dark matter by equating the annihilation rate to the Hubble expansion rate. This yields a model-independent relation of [28]

$$\Omega_\chi h^2 \simeq \frac{3 \times 10^{-27} \text{cm}^3 \text{s}^{-1}}{\langle\sigma v\rangle}, \quad (1.1)$$

meaning that based on abundance measurements from the CMB, $\langle\sigma v\rangle \simeq 3 \times 10^{-26} \text{cm}^3 \text{s}^{-1}$.




Figure 1.6: The mechanism of thermal freeze-out. The solid line is the density of dark matter that remains in thermal equilibrium as the universe expands and cools. The dashed line represents the freeze-out density for different $\langle\sigma v\rangle$ scenarios. The red region corresponds to a time in the universe when the production and annihilation rates are equal. The purple region corresponds to a time when the production rate is smaller than the annihilation rate, but both still regularly occur. The blue region corresponds to a time when the universe's expansion has overwhelmed the annihilation of dark matter. From Ref. [27].

As mentioned above, the assumption of dark matter being a freeze-out thermal relic constrains the dark matter mass to be in the sub-MeV to ~ 100 TeV range. The upper limit comes from unitarity constraints on the partial-wave expansion of dark matter annihilation [29]. The lower mass limit (~ 1 MeV) comes from the fact that if dark matter is too light, it would be relativistic at freeze-out and could not adequately

contribute to the observed large-scale structure formation [9, 12].

1.2.1 The WIMP Miracle?

The co-moving annihilation cross section above is what one would expect for a new particle that interacts via the Standard Model weak interaction in the GeV to TeV mass range [30], called the weakly interacting massive particle (WIMP). WIMPs are naturally predicted in supersymmetric extensions to the Standard Model, with some of the lightest particles predicted in supersymmetry being viable dark matter candidates [31]. Supersymmetry postulates a symmetry between integer spin particles and half-integer spin particles, which implies the existence of partner particles for all known particles with different spin properties. Supersymmetry could also explain why the Higgs mass is so much smaller than the Planck mass, which was one of the original motivations for the theory [32]. If evidence for supersymmetry was found through collider experiments, which probe electroweak-scale energies, this would be a strong clue toward a WIMP explanation for dark matter. For these reasons, WIMPs have historically been a popular dark matter candidate and one garnering the attention of many experimental searches.

At this time, there is currently no evidence for conventional supersymmetry from experiments at the Large Hadron Collider, which is the only accelerator complex operating at energies that can probe predictions from supersymmetry [33]. In addition, searches for the direct detection of WIMPs through their scattering on the nuclei of noble liquids have not yet found a positive signature [34–37]. The sensitivities of direct detection WIMP searches are rapidly approaching the “neutrino fog,” which are dark matter-nucleon coupling strengths where events originating from neutrino

interactions will impose a significant background—complicating the viability of this search method [38]. Figure 1.7 shows the sensitivity of direct detection experiments for GeV-scale dark matter, with the neutrino fog shown in dark blue. The sudden dip in the neutrino fog around 10 GeV is due to a transition between solar neutrinos as a dominant background (below 10 GeV) and atmospheric neutrinos as a dominant background (above 100 GeV). In this transition region the major source of background neutrinos comes from diffuse supernova neutrinos [39]. The loss in sensitivity of many of these experiments around this same mass is due to a different mechanism; below 20 MeV, nuclear recoils from dark matter start to become too soft to resolve in detectors. Current world-leading sensitivities for GeV-scale dark matter-nucleon cross sections are already within an order of magnitude of the neutrino fog.

The annihilation cross section is related to the dark matter mass m_χ and the mediator mass m_Φ as follows [41]:

$$\langle\sigma v\rangle\propto\frac{m_\chi^2}{m_\Phi^4}. \quad (1.2)$$

For a WIMP candidate, Φ is the Z boson. For smaller WIMP masses, $\langle\sigma v\rangle$ becomes smaller, which increases the expected abundance (since $\Omega_\chi \propto 1/\langle\sigma v\rangle$). Based on the relic abundance of dark matter measured by the CMB, calculations by Lee and Weinberg show that dark matter particles interacting via the weak interaction would lead to an overproduction of dark matter from what is observed if the dark matter mass is less than about 2 GeV [42]. With WIMPs not being a viable model below 2 GeV, a new interaction needs to be introduced to search for sub-GeV dark matter candidates. This circumvents the Lee-Weinberg bound by introducing a comparably light mediator through which light dark matter particles could annihilate.




Figure 1.7: A summary of experimental sensitivities to the dark matter-nucleon cross section for GeV-scale dark matter. The shaded gray is the excluded parameter space and the dashed lines show the expected 90% confidence level exclusion sensitivities of various experiments. Note this figure is from 2022 and recent results from LZ have excluded much of the parameter space from their expected sensitivity [36]. The neutrino fog is shown in the blue contours. From Ref. [40].

1.2.2 Light Dark Matter

This thesis will explore a model known as vector portal dark matter [41, 43], meaning that the “portal” interaction between the dark sector and visible sector is mediated by a new vector boson. The minimal vector portal model consists of a dark photon (A'), which is a new dark $U_D(1)$ gauge boson that kinetically mixes [44] with the Standard

Model photon, and a dark matter particle χ carrying $U_D(1)$ charge g_D that serves as the dark matter candidate. The $U_D(1)$ symmetry is broken, so the A' acquires a mass $m_{A'}$ through the Higgs mechanism. This symmetry breaking could be (but is not necessarily) governed by the same mechanisms as in electroweak symmetry breaking [45]. This produces the following Lagrangian:

$$\mathcal{L} = \mathcal{L}_{\text{SM}} - \frac{1}{4}F'_{\mu\nu}F'^{\mu\nu} + \frac{1}{2}m_{A'}^2 A'^\mu A'_\mu - \frac{\varepsilon}{2}F'_{\mu\nu}F^{\mu\nu} + \mathcal{L}_\chi, \quad (1.3)$$

where $\mathcal{L}_{\text{SM}}$ is the Standard Model Lagrangian, $F^{\mu\nu}$ is the electromagnetic field strength tensor, $F'_{\mu\nu} = \partial_\mu A'_\nu - \partial_\nu A'_\mu$ is the dark field strength tensor, ε is the kinetic mixing parameter between the dark photon and the Standard Model photon, and $\mathcal{L}_\chi$ is the dark matter Lagrangian, which depends on the nature of the dark matter:

$$\mathcal{L}_\chi = \begin{cases} ig_D A'_\mu (\chi^* \partial^\mu \chi - \chi \partial^\mu \chi^*) + m_\chi \chi^* \chi + \frac{1}{2} \partial_\mu \chi^* \partial^\mu \chi & \text{(Scalar } \chi) \\ \frac{1}{2} g_D A'_\mu \bar{\chi} \gamma^\mu \gamma^5 \chi + \frac{1}{2} m_\chi \bar{\chi} \chi + \frac{1}{2} \partial_\mu \bar{\chi} \partial^\mu \chi & \text{(Majorana fermion } \chi) \\ ig_D A'_\mu \bar{\chi}_1 \gamma^\mu \chi_2 + \frac{1}{2} m_1 \bar{\chi}_1 \chi_1 + \frac{1}{2} m_2 \bar{\chi}_2 \chi_2 + i \bar{\chi}_1 \gamma^\mu \partial_\mu \chi_2 & \text{(Pseudo-Dirac fermion } \chi) \end{cases} \quad (1.4)$$

The pseudo-Dirac fermions $\chi_{1,2}$ are mass eigenstates of a Dirac pair of two-component Weyl fermions (η, ξ) that have opposite $U_D(1)$ charge and a $U_D(1)$ conserving (breaking) Dirac (Majorana) mass term m_D (m_M), with $m_{1,2} \simeq m_D \mp m_M$ [46].

Something that has been swept under the rug above is that kinetic mixing between the dark photon and Standard Model photon necessitates the existence of massive particles Ψ, Ψ' that couple to both gauge fields. An example of this kinetic mixing at the one-loop level is shown in Figure 1.8. These particles could have masses far above the energy scale of the symmetry breaking, but the kinetic mixing persists at much lower energy scales. This one-loop process generates kinetic mixing with strength [44,

[45]:

$$\varepsilon = \frac{eg_D}{16\pi^2} \log\left(\frac{m_{\Psi'}}{m_{\Psi}}\right) \sim 10^{-5} - 10^{-3}, \quad (1.5)$$

where e is the electromagnetic coupling constant. A similar relation follows if the theory only contains one massive particle carrying both charges, with $m_{\Psi'}$ replaced by the theory's UV cutoff energy Λ . Grand Unified Theories motivate two-loop kinetic mixing in the predicted range $\varepsilon \sim 10^{-6} - 10^{-3}$ [45].




Figure 1.8: A one-loop kinetic mixing process between the A' and photon, with massive fermions Ψ that carry both electric charge and $U_D(1)$ charge.

The minimal vector portal model has four free parameters: $m_\chi, m_{A'}, \varepsilon$, and $\alpha_D \equiv g_D^2/4\pi$. To keep the theory perturbative, both ε and α_D must be less than unity. An upper bound of $\alpha_D \lesssim 0.5$ is motivated by the running of α_D and keeping the theory perturbative up to the electroweak scale [47]. The annihilation cross section at the tree level is [48]

$$\sigma v = \frac{8\pi}{3} \frac{\alpha \varepsilon^2 \alpha_D m_\chi^2 v^2}{(m_{A'}^2 - 4m_\chi^2)^2}, \quad (1.6)$$

where α is the electromagnetic fine-structure constant.

This thesis will explore two different mass hierarchy scenarios in the model, which substantively change the physical processes to search for. The first scenario is if $m_{A'} > m_\chi$, specifically $m_{A'} = 3m_\chi$. In this scenario, the A' will almost always decay on-shell to a pair of dark matter particles, given that $\alpha_D \gg \alpha \varepsilon$. The second scenario

is if $m_{A'} < m_\chi$. In this case, it is kinematically forbidden for an A' to decay on-shell to dark matter, so the preferred decay channel will be to Standard Model particles like an e^+e^- pair, provided that $m_{A'} > 2m_e$. An A' could still decay to dark matter off-shell, but this is suppressed by a $1/m_{A'}^4$ propagator factor. This mass hierarchy scenario is less interesting from the perspective of a simple dark matter model that predicts the correct relic abundance, but it is still interesting from the perspective of a search for a new long-lived particle that provides evidence of new physics. Near $m_{A'} \approx 2m_\chi$, dark matter annihilation is resonantly enhanced, allowing for sufficient annihilation rates for smaller couplings between dark matter and the Standard Model [49]. The annihilation cross section also becomes insensitive to α_D near this resonance. Because smaller couplings are required in this resonant regime, it is generally harder to probe thermal relic values in experiments. But the mass assumption $m_{A'} = 3m_\chi$ is sufficiently far from the resonant enhancement and is also small enough that thermal relic parameter values are not overstated in reach plots. This mass assumption is the same as that made by other measurements, which allows for easier comparison of results [50–55].

For use later in this thesis, Eq. (1.6) is re-written for the mass scenario $m_{A'} = 3m_\chi$ as

$$\sigma v = \frac{8\pi}{3} \frac{\alpha v^2 Y}{m_\chi^2}, \quad (1.7)$$

where the dimensionless parameter Y , defined as

$$Y \equiv \varepsilon^2 \alpha_D \left(\frac{m_\chi}{m_{A'}} \right)^4, \quad (1.8)$$

characterizes the interaction strength between dark matter and Standard Model particles. Note the term $(m_\chi/m_{A'})^4$. As stated above, the assumption $m_{A'} = 3m_\chi$ is

small enough that thermal relic parameter values are not overstated in reach plots, which are usually plotted in Y vs m_χ . Larger mass ratios will essentially lower the expected values for Y without changing the physically interesting parameter, the coupling strength between the dark sector and the visible sector ε .

1.2.2.1 Empirical Motivations for Vector Portal Dark Matter

There are a few empirical observations that have historically motivated the dark photon and vector portal dark matter. Although a dark sector explanation for these observations is now disfavored, the dark photon still remains a viable model to facilitate interactions between dark matter and the Standard Model. The first observation was the 511 keV positron excess in the galactic center measured by the INTEGRAL telescope [56]. It was suggested that dark matter annihilation ($\chi\chi^\dagger \rightarrow e^+e^-$) could be the cause. However, this explanation is now disfavored due to constraints from BBN and the CMB. The process $\chi\chi^\dagger \rightarrow e^+e^-$ would inject energy into the early Universe's plasma, disrupting BBN and causing disagreement from what is observed. In addition, the required dark matter annihilation cross section for the observed positron excess is much higher than what is allowed by CMB constraints.

Another observation motivating the dark photon was the measurement of the magnetic moment of the muon (muon $g - 2$). Measurements at Brookhaven National Laboratory showed a discrepancy by over 3σ between the observed muon $g - 2$ and theoretical predictions [57]. Follow-up measurements of the muon $g - 2$ at Fermilab further strained the discrepancy to more than 5σ [58]. To explain the discrepancy between the prediction and measurement, a contribution to the muon $g - 2$ by the dark photon was suggested [59]. However, the dark photon explanation for the muon $g - 2$ anomaly is now disfavored for the following reasons. The preferred parame-

ter space for dark photons to explain the anomaly has been probed by several other experiments at 90% confidence level with no positive signature seen [60–62]. Also, the muon $g - 2$ anomaly is now seemingly less anomalous, as new calculations of the hadronic vacuum polarization contribution to the muon $g - 2$ from lattice QCD have lessened the tension between measurement and theoretical predictions [63].

Visibly decaying dark photons have been proposed as an explanation for the 6.8σ excess of events in the invariant mass distribution of e^+e^- pairs produced in the nuclear transitions of excited $^8\text{Be}^*$ to its ground state via internal pair production [64]. This proposed protophobic boson would have a mass of 16.7 MeV and a coupling to electrons between 2×10^{-4} and 1.4×10^{-3} . However, a dark photon explanation is also now disfavored, as much of the parameter space preferred by this new boson has been ruled out by NA48 and NA64 [61, 62].

1.2.2.2 Detection Methods

This thesis will present an accelerator-based search for light dark matter, which complements direct detection searches [65, 66]. Accelerator production mimics the relativistic kinematics of the early universe; in these conditions, the production and scattering of dark matter are not suppressed by velocity or spin effects. Velocity- and spin-suppression can pose a significant challenge for direct detection searches compared to accelerator searches, as seen in Figure 1.9.

The discovery of a new particle that has properties consistent with a dark matter candidate is not necessarily confirmation of the discovery of the dark matter permeating the Universe. For this confirmation, a signal would need to be observed in direct detection experiments that matches the expectation from this new particle.




Figure 1.9: Comparison of sub-GeV thermal relic targets in direct detection searches and accelerator-based searches. The thermal relic target region for direct detection searches spans around ten orders of magnitude, whereas the same region in accelerator-based searches is condensed to about two orders of magnitude. From Ref. [67].

Accelerator-based searches can characterize the particle properties, more easily probe parameter space, and provide hints to direct detection experiments on where to hone their search, and direct detection experiments provide evidence of the cosmological abundance of the particle, making the two search methods complementary.

Due to the kinetic mixing between the photon and dark photon, dark photons can be produced in any interaction that would produce photons in the final state. The dominant production mechanism depends on both the experimental setup and the mass of the dark photon. In proton-beam fixed-target experiments, which typically

produce a shower of hadrons, the decay of neutral mesons (such as π^0, η) will be the dominant production mechanism if the dark photon is less massive than the meson. In electron-beam fixed-target experiments or for very massive dark photons with a proton beam, bremsstrahlung will be the dominant production mechanism. These production mechanisms are illustrated in Figure 1.10.




Figure 1.10: Dark photon production via the decay of neutral mesons (left) and bremsstrahlung (right).

There are many ways of identifying the production of dark matter in fixed-target experiments. One way is by looking for a signature of dark matter scattering in a detector downstream of the target, typically known as a “beam dump” search. Beam dump searches rely on both the production and re-scattering of dark matter, as seen in Figure 1.11. The number of expected signal events in a beam dump experiment scales as

$$N_{\text{signal}} \propto N_{\text{beam}} \cdot P_{\text{prod}} \cdot P_{\text{acc}} \cdot P_{\text{det}}, \quad (1.9)$$

where N_{beam} is the number of beam particles on target, P_{prod} is the probability of dark matter production, P_{acc} is the probability the dark matter particle passes through the detector due to angular acceptance, and P_{det} is the probability of detecting a dark matter re-scattering signal. Beam dump searches can exploit the versatility of detectors that may not necessarily be designed to search for dark matter re-scattering. For example, most of the current constraints on sub-GeV dark matter are set by

experiments nominally designed to detect neutrinos [51, 55, 68].




Figure 1.11: A cartoon showing dark matter production (left) and re-scattering (right).

Another way of identifying the production of dark matter in fixed-target experiments is through a missing-momentum/energy signature. In this type of search, only the production of dark matter is required (the left diagram of Figure 1.11). In a missing-momentum/energy search, the number of signal events scales as

$$N_{\text{signal}} \propto N_{\text{beam}} \cdot P_{\text{prod}} \cdot P_{\text{miss}}, \quad (1.10)$$

where P_{miss} is the probability the dark matter escapes the detector without interacting. Since the cross sections of dark matter production and dark matter-nucleon/electron scattering are small due to the necessarily feeble coupling strengths, P_{prod} and P_{det} are small while P_{miss} will be close to 1. Therefore, typically several orders of magnitude fewer beam particles on target are required to generate the same number of signal events in a missing-momentum/energy search compared to a beam dump search. However, the technical requirements for a missing-momentum experiment pose their own challenges. A hermetic detector is required to ensure Standard Model processes are not mistaken for a signature of dark matter production. The detector must also be radiation-tolerant, since a high fluence of particles and radiation

will be produced in the beam interactions with the target, and hermeticity requires the detector to be placed near the interaction point.

This thesis will describe a beam dump search for light dark matter with the NOvA experiment, which provides world-leading constraints to MeV-scale vector portal dark matter. A future, missing-momentum search for light dark matter with the Light Dark Matter eXperiment (LDMX) will also be discussed, along with a study to determine LDMX's expected sensitivity to visibly decaying dark photon signatures ($A' \rightarrow e^+e^-$), probing the mass hierarchy scenario $m_{A'} < 2m_\chi$. Finally, on a tangential note, performance studies of the cosmic ray veto (CRV) system for Mu2e are presented. The LDMX hadronic calorimeter is based on the design of the Mu2e CRV, so the performance studies of the CRV presented here will inform the expected detector performance of the LDMX hadronic calorimeter. Mu2e will search for the neutrino-less conversion of a muon to an electron, which is predicted to occur in many BSM theories.

Part I

Search for Light Dark Matter with NOvA

Chapter 2

The NOvA Experiment

The NuMI Off-axis ν_e Appearance experiment, or NOvA (the ν is stylized as a ν), is a long-baseline neutrino oscillation experiment based at the Fermi National Accelerator Laboratory (Fermilab) and Ash River, Minnesota. NOvA is designed to observe $\nu_\mu(\bar{\nu}_\mu) \rightarrow \nu_e(\bar{\nu}_e)$ appearance and $\nu_\mu(\bar{\nu}_\mu) \rightarrow \nu_\mu(\bar{\nu}_\mu)$ disappearance using two detectors: a near detector (ND) situated 1 km downstream of the neutrino production target at Fermilab to measure the flux of neutrinos, and a far detector (FD) situated 810 km downstream of the neutrino target in Ash River to observe ν_e appearance and ν_μ disappearance. The analysis in this thesis is focused on a dark matter search in the ND, so only details about the ND are provided here. However, both the ND and FD are functionally identical, i.e., they use similar materials, designs, and electronics.

2.1 The NuMI Beam

NOvA sources its neutrinos from the Neutrinos at the Main Injector (NuMI) beam at Fermilab [69]. A diagram of the accelerator complex at Fermilab is shown in Figure 2.1. A beam of 120 GeV protons is accelerated and directed onto a graphite target. Charged hadrons produced in this collision—mainly pions and kaons—can then be focused according to their charge using magnetic focusing horns. These charged hadrons then decay into the desired beam of neutrinos or antineutrinos,

depending on the charge sign focused by the focusing horns.




Figure 2.1: Diagram of the Fermilab accelerator complex [69].

The protons originate from a 35 keV H^- source, where a radio-frequency quadrupole accelerates the ions to 750 keV into the linear accelerator (Linac) [70]. The Linac then accelerates the ions to 400 MeV before sending them through a carbon foil to strip two electrons from the ions, leaving only a proton. The protons are then injected into the Booster ring, where they are accelerated to 8 GeV and grouped into 84 bunches at 52.8 MHz (called a “batch”) [70]. These Booster batches are then injected into the Recycler and Main Injector rings.

The Recycler and Main Injector are two synchrotrons with a circumference 7 times

larger than the Booster, allowing them to store and accelerate 6 Booster batches. The Recycler acts as a staging area and transitions batches from the Booster to the Main Injector. In the Recycler, proton batches are “slip-stacked”, where two batches are combined into one to increase the proton intensity delivered by the Main Injector [70]. Slip-stacking is achieved by injecting two batches one after the other and slowing down the first batch until the two overlap. After batches from the Booster are combined into six double batches, they are delivered to the Main Injector. The Main Injector accelerates the batches to 120 GeV before they are delivered to the NuMI target.

The beam of 120 GeV protons from Fermilab’s Main Injector is directed onto a graphite target. A diagram of the NuMI beam line is shown in Figure 2.2. Interactions of the proton beam in the target produce hadrons, primarily pions and kaons. Charged hadrons produced in the interactions are focused toward the beam axis by two magnetic horns. The direction of current in the magnetic horns selects for charge sign in the hadrons, and this consequently selects for either neutrinos or antineutrinos. Forward horn current (FHC) selects positively charged particles, which primarily decay to neutrinos. Reverse horn current (RHC) selects negatively charged particles, which primarily decay to antineutrinos. After the focusing horns, the charged particles enter a 675 m long decay pipe, where they decay primarily into muons (antimuons) and muon-antineutrinos (muon-neutrinos). A hadron absorber downstream of the decay tunnel removes the remaining protons and other hadrons from the beam. Muons are absorbed by the subsequent earth shield, while neutrinos continue through to the NOvA ND and beyond.




Figure 2.2: Diagram of the NuMI beam line showing how the neutrino beam is selected with the focusing horns.

2.2 The Near Detector

The NOvA ND is situated 14.6 mrad off-axis from the NuMI beam line and is 100 m underground at Fermilab. The ND's placement 100 m underground provides the benefit of a significantly reduced rate of cosmic rays. This off-axis placement is to ensure a narrow range of neutrino energies for the 810 km baseline of the oscillation analysis, which also benefits this analysis by restricting the neutrino energies. The distribution of neutrino energies expected in the ND, as a function of the off-axis placement of the detector, is shown in Figure 2.3. Although not optimized for a dark matter search, NOvA's off-axis location provides a near-optimal flux of potential light dark matter particles. Figure 2.4 shows the dark matter flux for a range of light dark matter masses as a function of the off-axis location, with NOvA's angular coverage highlighted in green.

The NOvA ND is a segmented tracking calorimeter constructed from extruded polyvinyl chloride (PVC) cells filled with liquid scintillator. Each cell is 3.9 cm wide, 6.6 cm deep (in the beam direction), and 4 m long. The liquid scintillator is a blend of 95% mineral oil and 5% pseudocumene with trace concentrations of wavelength-shifting fluors [71]. The composition of the detector by mass is 67% carbon, 16% chlorine, 11%




Figure 2.3: The neutrino energy distributions observed by the ND for different off-axis configurations from NuMI.




Figure 2.4: The expected flux of light dark matter over various off-axis positions. The angular coverage for NOvA is shaded in green and the angular coverage for MINERvA, an on-axis detector, is shaded in blue.

hydrogen, 3% titanium, 3% oxygen, and other trace elements. Each PVC cell has a wavelength-shifting (WLS) fiber to collect light, shift its wavelength from blue (400–450 nm) to green (490–550 nm), and deliver it to an avalanche photodiode (APD).

The PVC cells are coated with titanium dioxide to increase reflectivity and allow more light to be captured by the WLS fiber. An illustration of a NOvA ND cell is shown in Figure 2.5.




Figure 2.5: Illustration of a NOvA ND cell with wavelength-shifting fiber.

Each plane of the ND is made up of 96 cells. Sequential planes alternate in their vertical or horizontal orientation so that three-dimensional position reconstruction can be done. A schematic of this 3D reconstruction, with the top and side detector views, is shown in Figure 2.6. The fully active portion of the ND consists of 192 planes for a total depth of 12.9 m in the beam direction. At the end of the ND, there are 11 additional planes of PVC interlaced between 10 cm steel plates between PVC planes to better contain muons produced in the detector (called the “muon catcher”). The muon catcher is only two thirds the height of the main detector, and the analysis in Chapter 3 did not utilize the muon catcher in any way.

Being a low Z calorimeter, the ND is optimized for distinguishing electron showers from hadronic showers and muons. As electrons travel through the ND, they immediately induce electromagnetic (EM) showers that begin to spread out. The (mostly carbon) materials of the ND result in the detector having a radiation length¹ of ~ 36 cm (~ 5.5 planes) and a Molière radius² of ~ 11 cm (2–3 cell widths). A π^0 produced in a neutrino-nucleus interaction will immediately decay into two photons

¹A radiation length is the mean distance an electron will travel before losing $\sim 63\%$ of its energy, or 77% of the mean free path of a high energy photon will travel before undergoing pair production.

²The Molière radius is the radius of a cylinder containing on average 90% of an electromagnetic shower’s energy.




Figure 2.6: An example neutrino interaction in the NOvA detector. The left plot shows a 3D view of the detector, where the alternating orientation of the PVC cells can be seen. The right plot shows the two plane views of the detector for each orientation of the cells.

due to its short lifetime. These photons will undergo pair production and leave an electromagnetic shower in the detector, making them difficult to distinguish from electrons. However, photons are more likely than electrons to travel several planes before pair-producing because of the long mean free path of the detector. This separates them from the point of interaction, allowing for easier distinction from electron showers. Other final-state particles can easily be distinguished from electrons as well; muons will leave long, straight tracks in the detector and hadrons will generally leave much longer, wider showers due to the hadronic nature of their interactions.

The two ends of each WLS fiber are optically connected to a single APD cell. Each APD collects the light from both ends of 32 fibers. The APDs have an 85% quantum

efficiency for the 490–550 nm light delivered by the fibers. Light is converted to an electrical signal through the photoelectric effect, and the electrical signal is amplified by 150-fold. Electrical signals are recorded as “hits” when the signal exceeds the analog to digital conversion threshold. This threshold is individually determined for each APD pixel based on the amount of thermal electronic noise. Thermal noise is reduced by keeping the APDs cooled to -15°C by a thermoelectric cooler.

Each APD is connected to a custom front-end board (FEB) that performs the amplification, integration, shaping, digitization, and packaging of signals from the APD. The heart of the FEB is an application-specific integrated circuit (ASIC). Each ASIC serves 32 input channels (for each pixel of the APD) that are each amplified, integrated, shaped, “multiplexed” (multiple channels are grouped together) 8-to-1, and sent to an analog-to-digital converter (ADC). The multiplexing occurs at 16 MHz. After the ADC, digitized signals (converted to a 12 bit format) are sent to a field-programmable gate array (FPGA), which decides whether to keep an event or drop it. Data is then sent to a data concentrator module (DCM), a custom single-board computer. The DCM takes the input of up to 64 FEBs and time-orders input signals corresponding to $50\mu\text{s}$ windows, or “microslices”. These microslices are then further organized into 5 ms “millislices”, which are then sent to a farm of servers. There, the data is stored on a buffer node, awaiting triggering. Data on the buffer nodes are stored for ~ 20 minutes before being deleted. When a trigger is received, data on the buffer nodes is inspected for millislices that have a timestamp in coincidence with the trigger. If that condition is met, the buffer nodes send the triggered millislices to a data logger that writes the data to disk. Note that the trigger used in the analysis described in Chapter 3 is simply the $9.5\mu\text{s}$ proton delay window on the NuMI target, each occurring about 2 s apart.

2.3 Event Reconstruction

To make the hits recorded by the data acquisition system useful for physics analysis, event reconstruction is necessary. The event reconstruction chain is a series of algorithms that takes in hits, groups them together based on their proximity in space and time, classifies individual particles within the event, and calculates kinematic variables associated with the event. These algorithms are applied to both real data and simulated data. This allows simulation tools to be validated and for data to be understood based on truth-level information available in simulation.

In a given time window of the trigger, there may be multiple interactions in the detector, as shown in Figure 2.7. These interactions are happening almost simultaneously, so separating the many different interactions requires both temporal and spatial information to “slice” the event. Figure 2.7 also shows the coordinate system for the ND, where the origin is at the center of the front face of the detector.

Hits nearby in space and time are grouped into a “slice” using a clustering algorithm called the Time Density Slicer (TDSlicer). This process is done separately for each detector view (xz plane and yz plane). Cluster centroids are identified by finding local maxima in the density of hits using a density clustering algorithm [72]. From the centroids, slices are built up by adding nearby hits within a minimum distance from other hits in the slice [73]. The xz and yz views of the slices are then combined, along with time, into full 4D slices by averaging the z -position and time values between possible pairs of view slices.

For each slice, an interaction vertex is identified. The vertex reconstruction algorithm uses a modified Hough transform with a Gaussian-smeared voting scheme [74]. This pattern recognition algorithm takes pairs of pixels in each detector view as an input,




Figure 2.7: An example ND event display showing multiple interactions within a single event. The top pane shows the xz plane of the detector and the bottom pane shows the yz plane of the detector. The beam direction is from left to right. The color scale shows energy deposition.

and fits a line in polar coordinates³ between the points. The polar coordinates for each pair of points are mapped in (ρ, θ) space, where peaks will arise from particle trajectories in a slice. The highest peak in (ρ, θ) space is taken as the first Hough line, and hits lying along the line are removed from consideration in the next iteration of the algorithm. The process is repeated with the remaining hits until they have all been assigned to Hough lines. The intersections of Hough lines are used to seed an elastic arms algorithm [75]. This algorithm searches for an optimal vertex by finding a set of lines, or “arms”, that describe the particle tracks in the slice. The optimal arms

³Polar coordinates (ρ, θ) are used instead of $(x/y, z)$ in order to parameterize vertical trajectories.

and vertex is the set that minimizes an energy cost function based on the distance and the association factor between hits and arms [76]. Hits are penalized for not being associated with any arms, and arms are penalized if their first hit is far from the candidate vertex. An example result of this vertex finding algorithm is shown in Figure 2.8.




Figure 2.8: The result of finding Hough lines, shown in yellow, and the candidate vertex, shown as a red cross, for an example slice.

From the newly defined interaction vertex, hits in a slice are grouped together into particle “prongs,” which are clusters of hits with a defined starting point and direction. This process, hereinafter called Fuzzyk, uses a probabilistic clustering algorithm [77] based on the Fuzzy k-means technique [78]. Fuzzyk allows hits to belong to more than one cluster—hence “fuzzy”—and is probabilistic because a hit does not have to possess a membership probability of one. This allows isolated hits to be treated as noise rather than being associated with a particle cluster. In each detector view, the Fuzzyk algorithm maps out the angular hit distribution from the vertex. The angle

with the greatest hit density is used as the first prong trajectory, and hits are assigned membership probability to the prong [76]. Hits with less than 1% membership in that prong are used in the next iteration of finding a prong. The last step is to match the 2D prongs between the two detector views to make a 3D prong. Any prongs that can't be matched between the detector views remain as 2D prongs. Figure 2.9 shows the result of the Fuzzyk algorithm on an example slice. In the analysis described in Chapter 3, selection cuts are made on these prongs to isolate interactions that more likely occurred from the interaction of dark matter with the detector.




Figure 2.9: Result of the Fuzzyk prong finding algorithm. This is the same example slice shown in Figure 2.8.

Chapter 3

Dark Matter Analysis

This analysis is the first search by NOvA for sub-GeV dark matter produced in the NuMI beam and scattered through their non-gravitational interactions with electrons in the NOvA near detector. The presence of a dark matter signal would manifest as an excess of single-electron shower events that is kinematically consistent with dark matter-electron scattering. The vector portal is chosen as the benchmark model [79], described in Section 1.2.2, to look for single forward-scattered electrons produced by light dark matter. For comparison to other experimental limits and the parameters expected from cosmic observations, the assumption $m_{A'} = 3m_\chi, \alpha_D = 0.5$ is made on the model. This assumption makes the parameter $Y \equiv \alpha_D \varepsilon^2 (m_\chi/m_{A'})^4$ a function of only the kinetic mixing ε . To search for a dark matter signal in (m_χ, Y) space, a raster scan is performed over a wide range of possible dark matter masses. At each mass, a profile likelihood fit is performed, considering the neutrino background and systematic uncertainties.

This chapter presents the analysis in detail, including a description of the simulation samples used in the fit, the cuts applied to select for single-electron events, and the results of the fit to 2.55×10^{21} protons on target (POT) of data recorded by the NOvA ND.

3.1 Dark Matter Production and Detection

3.1.1 Production by NuMI

In the interaction of the proton beam with the target, dark photons could be produced by rare decays of neutral mesons, such as π^0 and η , or via proton bremsstrahlung, as shown in Figure 1.10. The dark photons decay promptly into dark matter particles. This would essentially contaminate NOvA’s neutrino beam with dark matter.

At small values of $m_{A'}$, particularly below m_{π^0} , neutral meson (ϕ) decays provide the dominant production channels for sufficiently light mediators and light dark matter particles χ :

$$\phi \rightarrow \gamma + A' \rightarrow \gamma + \chi\chi^\dagger. \quad (3.1)$$

The total number of dark matter particles from decays of such mesons can be written as

$$N_\chi = 2N_{\text{POT}}N_{\phi/\text{POT}}\text{Br}(\phi \rightarrow \chi\chi^\dagger\gamma), \quad (3.2)$$

where the factor of two takes into account the production of the $\chi\chi^\dagger$ pair, N_{POT} is the number of POT, $N_{\phi/\text{POT}}$ is the number of mesons per POT, and $\text{Br}(\phi \rightarrow \chi\chi^\dagger\gamma)$ is the branching ratio of the meson to a photon and pair of dark matter particles. There are approximately 0.5 π^0 produced per POT [80]. For η , this analysis assumes that the production rate is around 1/30 of that for π^0 [81]. This is a conservative assumption for the beam energy of NuMI [82, 83]. Note that this estimate of the flux is slightly underestimated, as the kinematically suppressed secondary and tertiary mesons produced in the target and also the off-shell three-body decays of mesons are ignored. Mesons produced in the beam dump downstream of the target and focusing horns are also ignored.

For higher mass dark photons ($m_{A'} > 120$ MeV), an additional production channel—proton bremsstrahlung—becomes important:

$$p + N \rightarrow p + N + A' \rightarrow p + N + \chi\chi^\dagger. \quad (3.3)$$

This channel could provide a significant dark matter source with very little angular spread:

$$N_\chi = 2 \frac{N_{\text{POT}}}{\sigma(pN)} \sigma(pN \rightarrow pN\chi\chi^\dagger), \quad (3.4)$$

where $\sigma(pN) \sim 40$ mb is the total proton-proton cross section for a beam of 120 GeV.

3.1.2 Detection by the Near Detector

A fraction of this dark matter beam enters the ND where it could interact with Standard Model particles through elastic scattering on a nucleon or electron, as shown in Figure 1.11. This analysis only considers dark matter-electron scattering, and not dark matter-nucleus scattering, which was done for the following reasons. The electron channel for dark matter scattering is not suppressed at the same level as it is for neutrino scattering [84], especially compared to the nuclear channel when $m_{A'}$ is small. This leads to relatively smaller backgrounds in an electron-channel search. In addition, dark matter-electron elastic scattering is a purely leptonic process in the visible sector, making it much less uncertain, especially compared to dark matter-nucleus scattering, which must account for nuclear effects in the signal production. Furthermore, the signal signature in the electron channel is clean and easily identifiable in the NOvA ND: a single, forward-going electron. For these reasons, this analysis focused on dark matter-electron scattering, and it is unlikely that any further sensitivity could be gained by exploring the dark matter-nucleus interaction channel.

The kinematics of dark matter-electron (χ - e) elastic scattering are described by relativistic elastic scattering. In the laboratory frame, (initial) electrons in the detector are considered to be at rest. The four momenta of the initial-state (p_χ^0, p_e^0) and final-state (p'_χ, p'_e) particles are

$$p_\chi^0 = (E_\chi, \mathbf{p}_\chi), \quad p_e^0 = (m_e, 0), \quad p'_\chi = (E'_\chi, \mathbf{p}'_\chi), \quad p'_e = (E_e, \mathbf{p}_e). \quad (3.5)$$

The Mandelstam variables can be obtained as

$$\begin{cases} s = (p_\chi^0 + p_e^0)^2 = m_\chi^2 + m_e^2 + 2m_e E_\chi, \\ t = (p_e^0 - p'_e)^2 = 2m_e^2 - 2m_e E_e, \\ u = (p_\chi^0 - p'_\chi)^2 = m_\chi^2 + m_e^2 - 2E_\chi E_e + 2(E_\chi^2 - m_\chi^2)^{1/2}(E_e^2 - m_e^2)^{1/2} \cos \theta_e, \end{cases} \quad (3.6)$$

where θ_e is the angle of the outgoing electron with respect to the incoming particle.

Using $s + t + u = 2m_\chi + 2m_e$, and assuming $E_\chi, E_e \gg m_e$, gives the kinematic relation

$$\cos \theta_e = \beta^{-1} \left[1 - \frac{m_e(1 - E_e/E_\chi)}{E_\chi} \right], \quad (3.7)$$

where $\beta \equiv \sqrt{1 - (m_\chi/E_\chi)^2}$. For high incoming particle energies of several GeV, and in the small angle approximation, the outgoing electron is very forward and

$$E_e \theta_e^2 = 2m_e \beta (1 - E_e/E_\chi) + 2E_e (1 - \beta), \quad (3.8)$$

hereinafter referred to as $E\theta^2$, is small.

It is worth pointing out here that neutrino-electron (ν - e) elastic scattering has the same two-body kinematics as dark matter-electron scattering, and is a purely leptonic

process. Replacing $E_\chi \rightarrow E_\nu$ and $m_\chi \rightarrow m_\nu$, and by ignoring the negligible mass of the neutrino, $\beta \rightarrow 1$ and

$$E\theta^2 = 2m_e(1 - E_e/E_\nu). \quad (3.9)$$

Since $m_\nu \ll m_\chi$, $E\theta^2$ will be more forward for neutrino-electron scattering than dark matter scattering. This will allow for distinction between the $E\theta^2$ distribution of ν - e scattering and χ - e scattering.

3.2 Strategy for Dark Matter Search

This kinematic property of the scattered electron with small $E_e\theta_e^2$ shows the “forwardness” or boost in the beam direction and can be used as an important signal selection criterion in the data analysis. This work searches for light dark matter signals within $E\theta^2$ space.

The main background consists of true electron showers from neutrino-electron elastic scattering. Because of similar two-body collision kinematics, the ν - e processes also leave a single, forward-going electron shower but with $E\theta^2$ peaking closer to zero than for dark matter-electron scattering. For example, in the NOvA ν - e cross section analysis, the typical region of interest is $E\theta^2 < 0.005 \text{ GeV Rad}^2$. The dashed purple distribution in Figure 3.1 shows the shape of the $E\theta^2$ distribution for ν - e scattering, which is sharply peaked in the first bin.

Besides the irreducible background from true electron showers produced by neutrino interactions, other showers produced by neutrino-nucleus interactions can occur and occasionally look like an electron shower. However, a significant fraction of this background can be suppressed with well-defined event selection criteria, which will be

discussed in Section 3.4. This background is expected to be approximately flat across the $E\theta^2$ values of interest to this analysis, shown as the dot-dashed pink distribution in Figure 3.1. To properly constrain this background in the fit, a region of higher $E\theta^2$ outside the signal window is included, which is $0.02 < E\theta^2 < 0.048 \text{ GeV Rad}^2$.

Unlike neutrinos, the dark matter particle of interest has a much higher mass than an electron, and the energy of the dark matter particles ranges from a few GeV up to several tens of GeV. The recoil energy of the electrons scattered off dark matter particles, depending on the mass and kinetic energy of the dark matter particle, is also higher than that from neutrinos. The ratio of the dark matter signal signature to the neutrino background increases in $E\theta^2$ space, peaking somewhere between $0.005 < E\theta^2 < 0.015 \text{ GeV Rad}^2$ depending on the dark matter mass. This can be seen in the solid red, yellow, and blue distributions in Figure 3.1, which show the distribution shape of a 1, 10, and 100 MeV dark matter signal, respectively. Most signal events will have $E\theta^2 < 0.02 \text{ GeV Rad}^2$.

To probe for a dark matter signal in the (m_χ, Y) parameter space, a raster scan [85] is used over a wide range of dark matter masses. The raster scan method performs a fit in Y at various fixed values of m_χ , as opposed to performing a fit in both m_χ and Y simultaneously. At each mass, a Poisson negative log-likelihood ($-2 \ln \mathcal{L}$) fit to the data is performed in $E\theta^2$ space from 0 to 0.048 GeV Rad^2 . The best-fit parameters of the fit minimize $-2 \ln \mathcal{L}$. The mass range for this search is from 1 to 200 MeV.

A binned template is used for the prediction, which consists of three components and are obtained from Monte Carlo (MC) simulation:

- χ - e signal template, denoted as $\mathcal{S}_{\chi-e}$, which is corrected by a factor $\mathcal{R}_{m_\chi}$ that depends on the dark matter mass m_χ and will be discussed further in Section 3.3.2,




Figure 3.1: The signal and background shapes of the $E\theta^2$ variable. Arbitrary scales are used for the normalization to show the shape difference clearly.

- ν - e background template, denoted as $\mathcal{B}_{\nu-e}$, and
- neutrino-nucleus interaction background template, denoted as $\mathcal{B}_{\nu-N}$.

The linear combination of the three components constitutes the prediction $\mathcal{P}_{m_\chi}$:

$$\mathcal{P}_{m_\chi}(a) = a \cdot \mathcal{S}_{\chi-e} \otimes \mathcal{R}_{m_\chi} + \mathcal{B}_{\nu-e} + \mathcal{B}_{\nu-N}, \quad (3.10)$$

where the parameter a is the normalization of the signal template. The nominal

normalization of the background templates $\mathcal{B}_{\nu-e}$ and $\mathcal{B}_{\nu-N}$ are set to their Standard Model predictions (from the results of the simulations described in Section 3.3.1). The $\nu-e$ background is well-understood, with a cross section uncertainty below 1%, and the $\nu-N$ background's normalization is understood within cross section uncertainties, which can be as high as 15%. The normalization of the signal template, a , can be converted to the physical parameter of interest Y , which will be explained in Section 3.3.2.

For each value of m_χ , a best-fit value is determined for a by minimizing $-2 \ln \mathcal{L}$. The parameter a is constrained to be non-negative in the fit. Systematic uncertainties are implemented as nuisance parameters in the fit, with penalty terms assigned equal to the square of the number of standard deviations by which they vary from their nominal value. The systematics included in the fit and their impact on the sensitivity are discussed in Section 3.5.

To determine NOvA's sensitivity to a dark matter signal, the 90% confidence-level exclusion upper limit on Y is determined from the best-fit at each mass using the profile likelihood ratio as a test statistic and using the profile construction method [86]. Because this search is probing near a physical boundary (small values of Y , which must take on non-negative values), Wilks' theorem may not provide accurate confidence intervals for exclusion. Instead, after fitting to data, Poisson-fluctuated pseudo-experiments are generated around the best-fit point to determine the distribution of profile likelihood ratios. From this distribution, the profile likelihood ratio that is larger than that of 90% of pseudo-experiments is established. In the asymptotic limit that the experiment is probing away from a physical boundary, the profile construction method reproduces a $\Delta\chi^2$ distribution, which would require a profile likelihood ratio of 2.71 to exclude at the 90% confidence level. However, the required value for

the 90% exclusion in this study is between 1.61 and 2.14, depending on the mass being tested. An example of how the 90% exclusion critical value was determined for an example mass ($m_\chi = 1$ MeV) is shown in Figure 3.2.




Figure 3.2: The distribution of profile likelihood ratios from 1000 Poisson-fluctuated pseudo-experiments around the best fit. The cumulative fraction of pseudo-experiments is shown in purple, and 90% of pseudo-experiments have a likelihood ratio value below 2.10.

3.3 Data Samples

Data used in this analysis consists of 2.55×10^{21} POT recorded by the NOvA ND between 2014 and 2020. This is NOvA’s entire dataset for NuMI beam operation with FHC.

Monte Carlo (MC) simulations are used to study the signal and background. Signal events in this analysis are single electron showers produced by dark matter-electron scattering. Background events include an irreducible background from the single

electron showers produced by neutrino-electron scattering and a reducible background from neutrino-nucleus interactions.

For MC studies, three samples are used to estimate the background:

1. a 5.54×10^{21} POT inclusive NOvA near detector MC sample (used to estimate the background, except for the two processes in the following two items),
2. a 1.35×10^{23} POT sample of ν - e interactions, and
3. a 2.0×10^{24} POT sample of ν_e charged-current (CC) meson exchange current (MEC) interactions.

The latter two samples are necessary to account for low statistics in these processes in the inclusive MC sample.

3.3.1 Neutrino Simulation

The NOvA MC simulation chain uses a series of simulated processes from various software packages. There are three key steps in the simulation process, where the output of each is the input of the next:

1. Beam simulation: the package G4Numi is used to simulate hadron production, transport, and decay at the NuMI target. This generates NOvA's neutrino beam. The flux is corrected using the Package to Predict the FluX (PPFX) [87].
2. Neutrino interactions: Using GENIE [88, 89], neutrino interactions with electrons and nucleons in the NOvA near detector are simulated.

3. Final-state particle propagation and detector response: The final-state particle information from the GENIE simulation is input to Geant4 [90–92], which propagates the particles through the detector and simulates detector response (the production of scintillation and Cherenkov light and any secondary particles). Geant4 then simulates the transport of these photons produced in the detector through the WLS fibers to the APDs.

The end result of this simulation chain are files with raw hits that mimic real detector data. These files are then fed through the same reconstruction algorithm that is used on real data.

3.3.1.1 Beam Simulation

The expected flux of neutrinos is simulated using the G4NuMI package, a Geant4-based Monte Carlo simulation tool that is used by other Fermilab-based neutrino experiments. This package includes specifications for the NuMI beamline, including the proton beam, the geometry and composition of the target, and the magnetic focusing horns. G4NuMI simulates the proton beam interactions with the graphite target to produce hadrons, the transport of these hadrons, and the hadron decays that produce the neutrino beam. The flux of neutrinos from this simulation is corrected using constraints from hadron production experiments using PPFX [80, 87]. This package also estimates uncertainties related to the hadron production using a multiverse analysis technique [93], which this analysis uses to characterize the neutrino flux uncertainty, described in Section 3.5.

3.3.1.2 Neutrino Interactions with GENIE

The output of the neutrino beam simulation, including the flux of neutrinos and their kinematics, is provided as input to the neutrino simulation software package GENIE [88, 89] to simulate neutrino interactions in the ND. GENIE uses theoretical models with constraints from neutrino cross section data to simulate neutrino interactions in the detector and the nature of the interaction type. NOvA uses a custom configuration of GENIE v3.0.6 that has been tuned to NOvA neutrino cross section measurements.

The charged-current (CC) quasi-elastic (QE) model was tuned in NOvA’s custom version of GENIE. The CC QE model uses a local Fermi gas model with corrections to include the effects of long-range nuclear correlations known as Random Phase Approximation (RPA) [94–96]. NOvA’s tune to this model is a QE shape suppression. Another adjustment made by NOvA is the inclusion of the meson exchange current (MEC) model, which was previously not included in GENIE. NOvA uses the Nieves MEC model [97] and tunes this model to NOvA ND data using a double-Gaussian re-weight function.

Because neutrino interactions are very rare, the most probable interaction in reality would be none at all. To more efficiently simulate “interesting” neutrino interactions, the cross sections are weighted to increase the likelihood of interaction. The cross section weights are passed along the simulation chain, along with final-state particles from the neutrino interactions, in order to properly normalize the event likelihood at a given POT.

3.3.1.3 Detector Response

After the neutrino interactions have been simulated, the final-state particles from these interactions are passed into Geant4 [90–92] to simulate the detector response to the outgoing particles. Additional secondary particles can be created at this stage based on the detector response. NOvA uses v10.2 of the Geant4 library.

This step simulates the production of scintillation and Cherenkov photons, and the transport of these photons through the WLS fibers to the APDs. The signals into the APDs are simulated as a combination of these photons and a model of thermal noise. Following this, the simulated files are processed with NOvA’s reconstruction algorithm, as would be done with real data.

3.3.1.4 Dedicated Samples of ν - e and MEC Events

To mitigate statistical uncertainties in the backgrounds for this analysis, two samples of specific background processes simulated with higher POT are used. Since ν - e events are the irreducible background for this analysis, it is critical that the uncertainties due to statistical fluctuations in this component be as small as possible. In addition, from studies performed in NOvA’s ν - e cross section analysis group, it was found that “spikes” in the distribution of ν_e -CC events passing their selection cuts were caused by a small number of ν_e -CC MEC events with large weights applied to them. To reduce the statistical uncertainty from these events, a high-statistics sample of ν_e -CC MEC events were generated as well.

To more efficiently produce these samples, the GENIE simulation step was modified to weight the cross section of the interaction types of interest. This further increased the interaction cross sections beyond what is described in Section 3.3.1.2 so that the

interaction type of interest became the dominant process in the simulation. Filters for these interaction types were also used so that only these events of interest were kept. Because the ν - e elastic scattering cross section is very low—about 1/2000 of the total charged-current cross section—a larger weight is necessary to reasonably produce a high-statistics sample of these events.

The 1.35×10^{23} POT enhanced sample of ν - e interactions was produced with the ν - e elastic scattering cross section weighted and filters applied to select an elastic scattering interaction type. In GENIE, the cross section of ν - e elastic scattering is calculated at the tree level. To improve the precision of the simulated ν - e elastic scattering cross section, an additional weight was applied to the ν - e sample to account for radiative corrections. The inclusion of these radiative corrections resulted in a 0.83% increase in the expected number of ν - e events from a tree-level simulation. This effect appears small, but it is significant for analyses such as this, which search for an excess over Standard Model processes.

The 2.0×10^{24} POT enhanced sample of ν_e CC MEC interactions were simulated with filters applied for MEC interaction types, as well as additional filters on the four-momentum transfer q between the neutrino and the nucleus for selected ν_e CC MEC events. These additional filters were $q^2 < 0.2 \text{ GeV}$, $q^0 < 0.4 \text{ GeV}$, and $|\vec{q}| < 0.6 \text{ GeV}$. These particular filters were applied because they are the part of kinematic phase space that produce events that pass selection cuts for this analysis (described in detail in Section 3.4.1).

3.3.2 Dark Matter Simulation

Signal samples are generated in a way similar to the background, but with a slightly different simulation chain:

1. Dark matter production and scattering: dark matter production, propagation, and interaction with electrons in the NOvA ND is done using the BdNMC simulation package [81].
2. Final-state particle propagation and detector response: The scattered electron and detector response are simulated with Geant4 in exactly the same way as for the background simulation.

The scattered single electrons produced by dark matter particles with different masses have different distributions of scattering energy and angle. It is computationally inefficient to produce detector response samples for each dark matter mass; instead of making samples for individual dark matter masses in Step 2 above, a “generic” signal sample was generated and this sample was re-weighted according to a ratio of the energy distribution for specific dark matter masses (the output of Step 1 for a given dark matter mass) and the energy distribution of the generic distribution. This procedure is described in detail in the remainder of this section.

3.3.2.1 Simulation of Production and Scattering

In Step 1, Dark matter particle production and dark matter-electron scattering is simulated using the software package BdNMC [81]. BdNMC is a dark matter simulation tool that simulates the production, propagation, and scattering of dark matter particles on electrons (or nucleons) in fixed-target neutrino detectors. It can be

applied flexibly to various proton beam energies, production targets and detector geometries and materials. It can be configured for a wide array of different fixed-target geometries, and over a dozen production channels at a variety of energies. In addition, it supports four possible scattering interactions between dark matter and the detector material in fixed-target neutrino detectors: elastic electron scattering, elastic neutral current-like nucleon scattering, inelastic neutral current-like nucleon scattering, and deep inelastic scattering. This simulation tool has been used by many other experiments to set limits on sub-GeV dark matter, including COHERENT [68], MiniBooNE [55], MicroBooNE [98], and Coherent CAPTAIN-Mills [99].

This analysis considers dark matter production via π^0 decay, η decay, and proton bremsstrahlung processes. The decays of neutral mesons are the primary source of dark matter production for scenarios with a small $m_{A'}$. As the π^0 and η both have decay lengths of a few nanometers ($c\tau_{\pi^0} = 25.1$ nm and $c\tau_{\eta} \simeq 0.2$ nm), the decay (meson $\rightarrow A'$) can be expected to take place inside of the target. Dark matter particles will then be produced by decays in flight of dark photons.

Due to their very short lifetimes, the production rate and angular distribution of these neutral mesons in NuMI have not been well measured. The π^0 and η primarily decay via radiative channels and do not contribute to the neutrino signal or background of fixed-target neutrino experiments. The underlying π^0 production model used in BdNMC was validated on data from a 400 GeV proton beam on a beryllium target [100], whereas NuMI uses a 120 GeV proton beam on a carbon target. A study was performed to estimate the π^0 flux at NuMI, and this estimation was compared to BdNMC's π^0 production model, allowing corrections to be applied to the expected flux in this analysis.

To estimate the π^0 flux, published data on $\pi^\pm$ flux by the MIPP Collaboration [80]

were analyzed. The π^0 flux was approximated by averaging the flux of charged-pions on the basis of isospin symmetry:

$$N_{\pi^0} \approx \frac{1}{2}(N_{\pi^+} + N_{\pi^-}). \quad (3.11)$$

This approach is reasonable, as the production of pions proceeds through strong processes, and electromagnetic effects related to their charges should lead to minor corrections. The experimental literature supports this method, as production cross sections between the three pion flavors are broadly similar [101], and previous experimental studies have approximated the π^0 momentum distribution by the mean of the charged-pion distributions with some success [102]. This approximation is also consistent with what is seen in simulations of the NuMI target by MicroBooNE [98].

In the simulation of dark matter production, $0.9\pi^0/\text{POT}$ was used, based on what is suggested in BdNMC documentation [81]. From measurements by the MIPP Collaboration, averaging over the π^+ and π^- flux, NuMI has an average of $3.985\pi^0/\text{POT}$ [80]. However, a large contribution to this flux comes from low-momentum pions that would not produce dark matter that propagates to the NOvA ND, and therefore no dark matter-electron scattering signal would be reconstructed from these pions. Instead of using a single normalization factor to correct for pion flux differences, a momentum-dependent correction was made, derived as follows.

First, a distribution in p_z (the momentum component in the beam direction) and p_T (the transverse momentum) of 10,000 π^0 s was generated with BdNMC. The distribution was normalized so that the integral is 0.9, to match the $0.9\pi^0/\text{POT}$ used in the simulation. The π^0 momentum was binned in (p_z, p_T) with the same binning used by the MIPP Collaboration, and then the ratio between the MIPP data and the BdNMC

simulation was taken. Figure 3.3 shows the ratio plot between the MIPP data and the simulated events, along with a contour of the region of momentum space that produces 90% of χ - e scattering events in the NOvA ND.




Figure 3.3: The ratio between the π^0 /POT fluxes in MIPP and BdNMC, with a contour line containing momentum values that produce 90% of single-electron showers in the NOvA Near Detector.

Then, for each simulated signal event producing a single-electron shower in the ND, a histogram was populated with the event's recoil electron energy, weighted by the MIPP/BdNMC ratio for the (p_z, p_T) of the π^0 in the event. The ratio between this weighted distribution and the nominal distribution of recoil electron energies (i.e., with a weight of 1 per event) is taken, shown in Figure 3.4. The weighted-to-nominal ratio of the recoil electron energy was implemented as a weight to the sample re-weighting procedure described in Section 3.3.2.3.

In BdNMC the NOvA near detector is configured as a simple cuboid with 4.1 m width,




Figure 3.4: The recoil electron energy distribution of events from the nominal simulation and from applying MIPP/BdNMC ratio weights to the events.

4.1 m height, and 14.3 m length. The (x, y, z) position of the cuboid center is (13.8 m, -4.19 m, 989.89 m), where the origin of the axes $(0, 0, 0)$ is in the center of the NuMI target, and all production processes (such as meson production and decay, and proton bremsstrahlung) occur at the origin (the z -axis is oriented parallel to the NuMI beam direction). The length of the cuboid is aligned parallel to the z -axis (beam direction), the width is along the x -axis and the height is along the y -axis.

The material composition of this cuboid is BdNMC's liquid scintillator, with an electron number density of $4.14 \times 10^{23}/\text{cc}$ (the atomic number density is set to $5.16 \times 10^{22}/\text{cc}$ and the number of electrons per atom is set to 8). To correct for differences between the detector composition in BdNMC and the real detector composition of the NOvA ND, a correction factor was applied to the expected dark matter flux. Based on the known detector composition of the NOvA ND, shown in Table 3.1, the

average number density of the NOvA ND is $3.71 \times 10^{22}/\text{cc}$ and the electron number density is $3.25 \times 10^{23}/\text{cc}$.

Table 3.1: Composition of NOvA ND.

NOvA ND density: 0.987g/cc		
Element	molar mass (g/mol)	fraction
H	1.01	0.1076
C	12.01	0.6671
O	16	0.0298
Cl	35.5	0.1606
Ti	47.8	0.0322
N	14.01	0.0003
S	32.06	0.0010
Sn	118.71	0.0012
Ca	40.08	0.0003
Na	22.99	0.00003

The signal rate is directly proportional to the electron number density of the detector material; more electrons in the detector volume means more opportunities for the dark matter to scatter off of an electron. Because of this direct proportionality, a single correction factor for the difference in electron number density was applied. A correction factor of $3.25/4.13 \approx 0.79$ is applied in this final steps of the analysis, when converting the normalization of the signal template in the best-fit to the parameter Y . The impact of this correction was a slight decrease in the expected number of signal events from the nominal value in the simulation, which translates to a slight decrease in expected sensitivity.

When a dark matter particle elastically scatters off an electron in the NOvA near detector, it will usually produce a detectable single electron event. In addition to the production of dark matter in the target, BdNMC simulates the production of scattered electrons in a detector from dark matter-electron interactions. The number of scattered electrons is linearly correlated to the number of dark matter particles,

and scaled by the parameter Y . Based on the parameters set in the configuration, BdNMC predicts the total number, the scattering angle, the recoil energy, and the vertex of the scattered electrons in the detector.

Using BdNMC, a simulation of 500 dark matter masses from 1 to 500 MeV were used to predict the expected number of events, the mean energy, and the mean scattering angle of the electrons scattered by dark matter particles, assuming $\alpha_D = 0.5$, $\varepsilon = 1 \times 10^{-3}$, $m_\chi/m_V = 1/3$ (i.e. $Y = 6.173 \times 10^{-9}$) and POT = 1.0×10^{21} . The predicted number of dark matter-electron scattering events as a function of dark matter mass is shown in Figure 3.5.




Figure 3.5: The expected number of dark matter-electron scattering events as a function of dark matter mass. In increasing values of m_χ , the dominant dark matter production channels are π^0 -decays, η -decay, and proton bremsstrahlung process, which becomes resonant near the ρ/ω mass region. This assumes $\alpha_D = 0.5$, $\varepsilon^4 \alpha_D = 5 \times 10^{-13}$, $m_\chi/m_V = 1/3$ and POT = 1.0×10^{21} . Statistical uncertainties are shown with blue error bars.

The electrons scattered by dark matter particles with different masses have different distributions of recoil energy and scattering angles, which means that the distribution

of the recoil electron $E\theta^2$ will vary with the dark matter mass. The number of expected signal events varies by orders of magnitude depending on the mass considered. For higher dark matter masses and the number of POT considered in this analysis, there may be as few as one signal event. To ensure the distribution of events in the ND was accurate, a high-statistics sample of each mass was produced. These high-statistics samples were later re-weighted to the POT considered for this analysis. The generic distributions of the recoil energy (see Section 3.3.2.2) and scattering angle of the electrons were used as the input to the Geant4 simulation step to produce single-electron events and then to simulate the detector response to those events.

3.3.2.2 Generic Single-Electrons Samples

As described above, BdNMC was used to obtain the specific distributions of the recoil energy and the scattering angle of single electrons for a series of dark matter masses. The specific energy distributions of the single electrons were normalized and summed together to produce an envelope over all the distributions to get a “generic” energy distribution. The generic energy distribution, together with five mass-specific distributions, are shown in Figure 3.6 on the left.

The electron’s scattering angle is correlated to its recoil energy, as shown in Figure 3.7 on the right. Due to the off-axis position of the NOvA ND breaking the symmetry between scattering closer to or away from the beam direction, the scattering angles have a double-peak structure.

The generation of scattered electrons in the simulation is the random emission of a single electron within the $4.1 \times 4.1 \times 14.3 \text{ m}^3$ cuboid volume (approximately the size of the NOvA ND). The energy of this electron was randomly generated following the




Figure 3.6: The simulated energy distribution of the single electrons scattered by dark matter particles for various masses. The green curve is the generic energy distribution.




Figure 3.7: The correlation between the scattered electron recoil energy and the scattering angle. The small inset figure shows the scattering angle distribution at 1 GeV recoil energy.

generic energy distribution of the total dark matter distribution packet, shown as the green curve in Figure 3.6, and the scattering angle was randomly assigned following the angular distribution at a given energy. The position of the scattered electrons

were randomly sampled in the cuboid volume following the position distributions of the BdNMC output. The detector response to these scattered electrons was then simulated in a similar way as described in Section 3.3.1.3.

In total, 1.25×10^6 generic signal events were produced. These samples were produced as single-electron showers in the ND without any other particles present. In reality, this does not happen, as the ND sees many showers per event window from the high-luminosity proton beam, seen in Figure 2.7. To address the systematic uncertainty in the lack of an “overlay” procedure for these signal events, a detector pileup systematic was used, which is discussed in Section 3.4.2.

3.3.2.3 Signal Template Generation by Sample Re-weight

The generic signal samples were re-weighted to get specific signal distributions for individual dark matter masses. To do so, the generic samples were weighted by scale factors that depend on the true recoil electron energy. The scale factor is the bin-by-bin ratio between two histograms: the electron energy distribution from the generic sample and the electron energy distribution for a specific dark matter mass. Figure 3.8 shows the ratio histograms for dark matter mass points of 5, 10, 50, 100, and 450 MeV.

After the re-weighting factors were determined for each mass and recoil electron energy, events in the generic distribution were re-weighted according to their true electron recoil energy and the re-weighting function, implemented as a lookup table, to produce a mass-specific sample. Figure 3.9 shows an example of the specific sample production procedure. This procedure produces a signal template for a specific dark matter mass m_χ quite accurately, as seen in Figure 3.10.




Figure 3.8: The bin-by-bin ratio between the generic energy distribution (green curve in Figure 3.6) and the specific energy distributions for specific dark matter masses (curve with corresponding color).

It is in this step that a second re-weighting function was applied to dark matter masses below 40 MeV (those that can undergo dark photon production from pion decay) to account for differences between BdNMC’s pion production model and observed pion flux in NuMI. The calculation of this re-weighting function is described in Section 3.3.2.1.

The re-weighting procedure kept about 10% of the generic samples for each specific signal sample, or 1×10^5 events for a particular mass. A POT equivalence for the specific sample was determined by comparison to the expected number of signal events at 1×10^{21} POT, which is shown in Figure 3.5. Using $Y_{\text{ref}} = 6.17 \times 10^{-9}$ for each simulated dark matter mass, a reference POT can be determined by calculating the




Figure 3.9: The simulation and re-weight procedure of the single electron events. On the left are the recoil electron spectra produced from the BdNMC simulations: the generic energy spectrum, the energy-angle correlation, and the scattering angle spectrum. In the middle are the spectra of the single-electron energies and angles produced from the Geant4 simulation, which was used to produce reconstructed hits. The recoiling energy spectrum was directly sampled from the generic energy distribution, and the scattering angle spectrum was calculated from the energy-angle correlation. After simulation of detector response and event reconstruction, the generic samples were re-weighted according to the ratio histogram, shown at the top of the middle, to produce specific mass samples. On the right, the energy and angle distributions from the re-weighting procedure for dark matter mass 10, 50, 100, and 450 MeV are showed in black. The “true” distributions produced from BdNMC software are plotted on top in red and light blue.




Figure 3.10: A comparison of the $E\theta^2$ spectra between samples reweighted from the generic sample and the “true” distributions generated directly by BdNMC for four different dark matter mass points (5, 10, 20, 100 MeV).

number of POT required to produce 1×10^5 signal events in the ND¹. Table 3.2 shows the number of signal events after re-weighting for eight example dark matter mass, which is the integral of the re-weighted histograms without any selection cuts, and the reference POT_{ref} when $Y_{\text{ref}} = 6.17 \times 10^{-9}$.

In the fit, the free parameter being varied is the normalization scale of the signal template a , not the physical parameter Y . Since the dark matter samples were all simulated with $Y_{\text{ref}} = 6.17 \times 10^{-9}$, the signal template scale a can be converted to Y with

$$Y = \sqrt{a} Y_{\text{ref}}. \quad (3.12)$$

¹The exact number of signal events after re-weighting was used in practice, but for explanatory purposes, the approximate number of 1×10^5 events is stated.

Table 3.2: Summary of the simulated signal samples. The number of signal events N_{events} for a dark matter mass m_χ was calculated by integrating the re-weighted signal histograms without any cuts. The POT_{ref} was calculated with $Y_{\text{ref}} = 6.17 \times 10^{-9}$.

m_χ (MeV)	N_{events}	POT_{ref}
1	10960	9.33×10^{14}
2	32433	1.17×10^{16}
5	53315	1.47×10^{17}
10	78339	1.29×10^{18}
20	88494	1.52×10^{19}
50	134650	3.23×10^{22}
100	113144	8.67×10^{23}
200	131004	6.87×10^{25}

3.4 Event Selection

3.4.1 Single Electron Event Selection

This work aims to select candidate single-electron events from the NOvA ND data. Dark matter particles elastic scatter off electrons, producing a single-electron shower along the beam direction pointing directly back to the target. This signal signature is similar to that of single electrons produced by the background ν - e process. Therefore, the event selection criteria used for this analysis is the same as that used by the mature ν - e analysis [103]. These cuts were optimized using a figure of merit (FOM), defined as

$$S / \sqrt{S + B + \delta B^2}, \quad (3.13)$$

where S is the number of “signal” (single-electron shower) events, B is the number of neutrino-nucleus interaction background events, and δB^2 represents the systematic uncertainty in the background. The cuts optimized for the ν - e analysis are used in this analysis rather than optimized cuts for a dark matter signal for the following reasons. The expected number of χ - e signals is not known, a priori, since neither the

true value for the coupling ε or the dark matter mass m_χ are known. As shown below, the selection efficiency is also mass-dependent, which poses a problem of selecting a single mass value for the optimization of cuts, without knowing the “correct” mass. Because of this, the ν - e selection cuts were used as is.

First, preselection cuts were applied to remove obvious ν_μ -CC interactions produced in the detector, which contain long, straight tracks. These cuts are defined as:

- $L < 800$ cm, where L is the length of the longest prong,
- $N_{\text{plane}} < 120$, where N_{plane} is the number of planes of the longest prong, and
- $N_{\text{cell}} < 600$, where N_{cell} is the total number of cells in a slice.

See Section 2.3 for a reminder of how prongs and slices are defined.

Containment cuts were applied to suppress backgrounds that are induced by neutrino interactions with the rock surrounding the ND and to ensure event interactions are fully contained within the ND. The containment cuts required, for the x , y , and z position of all hits in a prong,

- minimum $x > -190$ cm , maximum $x < 180$ cm,
- minimum $y > -180$ cm , maximum $y < 190$ cm,
- minimum $z > 105$ cm, , maximum $z < 1275$ cm,

where the origin of the coordinate system is at the center of the front face of the ND. Requirements were also made on the position of the reconstructed vertex of the prong:

- $-185 < x_{\text{vtx}} < 175$ cm,

- $-175 < y_{\text{vtx}} < 175 \text{ cm}$,
- $95 < z_{\text{vtx}} < 1095 \text{ cm}$,

where the coordinate system origin is at the center of the front face of the ND.

Single-particle requirement cuts were applied to mitigate backgrounds with multiple final-state particles. These cuts were defined as

- The ratio of the cell energies contained by the primary prong $E_{\text{prong}}/E_{\text{total}}$ should be greater than 0.8, where E_{prong} is the most energetic prong energy and E_{total} is the total energy of the slice.
- The vertex energy should be less than 0.02 GeV, where the vertex energy is the sum of the cell energies above thresholds, excluding the most energetic prong, within ± 8 planes of the event vertex plane.
- The shower gap should be less than 20 cm, where the shower gap is the distance between the start point of the most energetic prong and the event vertex.

An energy cut to select for candidate electrons with specific energies of interest was applied, defined as:

- $0.5 < E_{\text{prong}} < 5 \text{ GeV}$.

Two event classifiers based on convolutional neural networks (CNNs) were trained to identify single-electron events (NuoneID) and another to further reject background with a π^0 in the final state (Epi0ID). The architecture² for these CNNs was the same used for NOvA's Convolutional Visual Network [104]. Both of these classifiers

²Neural networks can be represented as mathematical graphs, and the architecture of the network defines how nodes are interconnected.

were trained on single-electron events as the signal. However, the classifiers were not trained specifically on ν - e events. Instead, they were trained on single electron showers with a broad distribution of energies and recoil angles, which was done to mitigate the model's dependence on the scattering energy and angle. This benefits this analysis, in particular, since the scattering energy and angle distributions for χ - e scattering is different from ν - e events. The discriminator threshold on these classifiers, optimized by the ν - e analysis group, was:

- NuoneID > 0.73,
- Epi0ID > 0.92.

The distributions of the discriminator values for these two neural networks on training samples is shown in Figure 3.11.




Figure 3.11: The discriminator values for the NuoneID (left) and Epi0ID (right) neural networks. The testing samples in these plots consist of ν - e events (“MC Signal”) and ν - N events (“MC Background”). An arrow shows where the discriminator threshold was set and the events that are kept by the networks.

Tables 3.3 and 3.4 show the effect of the event selection cuts on the background and the signal, respectively. The efficiency of each cut, ϵ^{N-1} , is defined as the ratio between the number of events with to that without the cut, with all previous cuts applied.

The cumulative efficiency of each cut, ϵ , is defined as the ratio between the number of events remaining after applying that cut and all previous cuts to the initial number of events. The efficiency of the signal depends on the dark matter mass considered, so the selection efficiency for three different masses (1, 10, and 100 MeV) is shown. The selection efficiency for each cut as a function of $E\theta^2$ is shown in Figure 3.12.

Table 3.3: Cut efficiency for the MC ν - e background and the MC ν - N background events.

Cut	$\epsilon_{\nu-e}^{N-1}$ (%)	$\epsilon_{\nu-e}$ (%)	$\epsilon_{\nu-N}$ (%)
Preselection	99.95	99.95	23.44
Containment	58.38	58.35	4.67
Single Particle Requirement	85.50	49.90	1.10
Primary Shower Energy	66.70	33.29	0.21
Particle ID	55.40	18.47	5.15e-3

Table 3.4: Cut efficiency for the MC χ - e signal events at $m_\chi = 1, 10$, and 100 MeV.

Cut	$\epsilon_{1\text{MeV}}$ (%)	$\epsilon_{1\text{MeV}}^{N-1}$ (%)	$\epsilon_{10\text{MeV}}$ (%)	$\epsilon_{10\text{MeV}}^{N-1}$ (%)	$\epsilon_{100\text{MeV}}$ (%)	$\epsilon_{100\text{MeV}}^{N-1}$ (%)
Preselection	99.9996	99.9996	99.9986	99.9986	99.9919	99.9919
Containment	81.47	81.47	77.00	77.00	64.55	64.55
Single Particle Requirement	68.05	83.53	65.30	84.80	55.64	86.19
Primary Shower Energy	11.13	16.35	33.20	50.84	43.89	78.88
Particle ID	7.78	69.90	23.05	69.43	29.85	68.01

3.4.2 Overlay Effect on Event Reconstruction and Selection

“Overlay” is a procedure to mix simulated single-interaction events on top of a single slice to mimic detector pileup. For the ν - e events, these single-electron events were generated and then overlaid onto a full slice with neutrinos interacting in the detector. These spills don’t have quite the right detector occupancy since there is one more neutrino interaction than there should be, but it is close enough and does not significantly impact the overall beam intensity.




Figure 3.12: Selection efficiency as a function of $E\theta^2$ for ν - e events (Top Left) and dark matter signal events for $m_\chi = 1$ (Top Right), 10 (Bottom Left), and 100 (Bottom Right) MeV.

Due to computational constraints, no overlay samples were produced for χ - e events. Instead, in order to understand the effect of the overlay procedure on the reconstruction and event selection efficiency, the overlay and non-overlay ν - e samples were compared.

The overlaid events have a different event reconstruction efficiency than non-overlaid events due to the presence of other nearby hits influencing the prong finding algorithm. This effect was studied using a 1.26×10^{23} POT overlay sample, which contained 729,152 ν - e events and a 1.48×10^{23} POT non-overlay sample, which contained 858,648 events. From the overlay samples, 472,915 events were successfully reconstructed ($\sim 64.8\%$ efficiency), while 569,475 events were reconstructed ($\sim 66.3\%$ efficiency) from the non-overlay samples. It indicates that the reconstruction efficiency from the overlay samples is $\sim 1.5\%$ less than that from the non-overlay samples.

The overlay procedure also has an impact on the selection of the reconstructed events. The event reconstruction and selection efficiencies after the event selection cuts with respect to the true $E\theta^2$ is shown in Figure 3.13. It shows that for single electron shower events, the event reconstruction and selection efficiency for non-overlay samples is $\sim 4\%$ higher than for overlay samples and that the difference in selection efficiencies between both samples do not depend significantly on the true $E\theta^2$. Therefore, because the difference is a constant $\sim 4\%$, it was not necessary to produce overlay samples for dark matter signal events. In order to account for the 4% reconstruction and selection efficiency difference introduced by ignoring the overlay procedure, a “pileup” systematic uncertainty was introduced to characterize this discrepancy. This systematic uncertainty was implemented as a 4% decrease in the normalization of the non-overlay samples (χ - e and ν - e) at the 1σ level. The ν - N background was simulated with pileup present, so this systematic uncertainty was not applied to this

background.




Figure 3.13: The Monte Carlo single-electron (from ν - e) event reconstruction and selection efficiency on non-overlay (solid line) and overlay samples (dashed line). Each colored line is a step in the event selection described in Section 3.4.1. The ratio of overlay to non-overlay is shown in the lower panel. The ratio is an approximately constant 96% as a function of true $E\theta^2$. Note the plotted $E\theta^2$ range is narrower than the full signal region.

3.5 Systematic Uncertainties

In this analysis the following systematic uncertainties were used:

1. The detector response uncertainties for the three template components ($\mathcal{S}_{\chi-e}$, $\mathcal{B}_{\nu-e}$, and $\mathcal{B}_{\nu-N}$), considering the NOvA ND calibration, light-level, and Cherenkov effects.
2. The flux uncertainties, for both the signal and the background.
3. Pileup uncertainties to address the overall normalization due to the lack of overlay in the signal and ν - e background samples. For this, a 4% uncertainty at 1σ is used and motivated at the end of the previous section.
4. Neutrino interaction cross section modeling uncertainties.

For all systematic effects, the impact of these uncertainties was explicitly determined at $\pm 1, \pm 2, \pm 3\sigma$ shifts. A cubic interpolation function was used to determine the impact between these explicitly determined shifts.

3.5.1 Detector Response Uncertainties

The impact of detector response uncertainties is assessed by creating alternate simulation samples with shifted detector response parameters. Detector response uncertainties included in this analysis are calibration uncertainties, light-level uncertainties, and Cherenkov light uncertainties.

Calibration uncertainties were determined by studying the differences between data and MC in the energy response for various samples (ND and FD cosmic rays, ND

muon tracks from ν_μ interactions, and ND proton tracks from ν_μ interactions). Data-MC comparisons were made for each sample by comparing the dE/dx as a function of position along the particle's track. From these studies, proton samples showed the largest discrepancy between data and MC at 5% [105]. This was taken as a conservative estimate for the calibration uncertainty across all samples.

The photon production model in Geant4 used by NOvA is calculated as

$$N_\gamma = F_{\text{view}}(Y_{\text{scint}}E_{\text{Birks}} + \epsilon_{\text{Cher}}C_{\text{Cher}}), \quad (3.14)$$

where N_γ is the number of photons produced, F_{view} is the scale factor depending on the detector view (x or y direction), Y_{scint} is the number of scintillation photons produced per unit energy deposition, E_{Birks} is the energy deposited according to Birks' law [106], ϵ_{Cher} is the efficiency for Cherenkov photons, and C_{Cher} is the number of Cherenkov photons produced by the charged particle. The Cherenkov uncertainty takes into consideration the 5% discrepancy in proton dE/dx between data and MC. This discrepancy was decreased to less than 1% by adjusting the light response model (adjusting Y_{scint} and ϵ_{Cher}) while keeping the detector response to muons the same. The Cherenkov efficiency was determined to have an uncertainty of 6.2% [105]. Light-level uncertainty was determined by adjusting the detector calibration to keep the muon response the same, and was determined to be 5% [105].

3.5.2 Flux Uncertainties

The flux uncertainty for χ - e events is derived from the uncertainty of pion flux from MIPP data [80]. MIPP reports a $\sim 5\%$ systematic uncertainty and $\sim 1\%$ statistical uncertainty on their measurements in momentum bins that contain 90% of pions pro-

ducing χ - e events in the ND. Based on this, a 6% uncertainty on the dark matter flux was applied to the signal template. This uncertainty is uncorrelated with the background flux uncertainty. Note that this uncertainty should only affect χ - e events that were produced from the decay of a π^0 . This is the dominant production mechanism for masses where NOvA is expected to have competitive sensitivity, so this uncertainty is prioritized over uncertainty from proton bremsstrahlung production.

The flux uncertainty for ν -induced background events arise from two sources: hadron production uncertainties and uncertainties in the beam optics. The first refers to the cascade of particles produced in the proton beam interaction with the NuMI target, which eventually decay into neutrinos. The second refers to the focusing of these charged hadrons that eventually decay into neutrinos. These uncertainties are dependent upon the true neutrino energy. The background component flux uncertainties use a principal component analysis [93]. Using PPFX [87], simulations with randomly generated variations in the hadron production and beam focusing uncertainties determine the variance in the outgoing neutrino beam, along with the correlations between bins of neutrino energy. From analyzing the correlations, a covariance matrix was obtained and diagonalized, and the 7 largest eigenvectors were kept for this analysis. These characterize the flux uncertainty of our background, which induces an $\sim 8\%$ uncertainty.

3.5.3 Cross Section Uncertainties

The largest number of uncertainties introduced to the analysis comes from uncertainties in the neutrino cross section, due to the many and varied interactions neutrinos can have with nuclei. Included in this analysis are uncertainties related to elastic

Table 3.5: Summary of how 1σ systematic uncertainties affect the normalization of ν - N and ν - e backgrounds and χ - e signal events.

Source	ν - N (%)	ν - e (%)	χ - e (%)
Detector model	+10.1/ - 4.4	+1.7/ - 2.6	+3.1/ - 3.0
Flux	+7.2/ - 7.5	+8.2/ - 8.3	+6.0/ - 6.0
Cross sections	+17.8/ - 11.4	+0.0/ - 0.0	+0.0/ - 0.0
Pileup	+0.0/ - 0.0	+4.0/ - 4.0	+4.0/ - 4.0
Total	+24.1/ - 14.6	+9.3/ - 9.5	+8.0/ - 8.0

and quasi-elastic (QE) scattering, MEC, resonance production (RES) and deep inelastic scattering (DIS). These uncertainties are implemented by re-weighting the nominal simulation events. The nominal simulation predicts a certain rate at which these processes can occur. The re-weighting adjusts these rates, thus affecting the overall normalization of event types in the analysis. The uncertainty of the ν - e scattering cross section is negligibly small. However, the cross section uncertainty for ν - N processes induces a $\sim 15\%$ uncertainty on the normalization of this background component.

In summary, the largest source of uncertainty on both ν - e events and χ - e events comes from their respective flux uncertainties, and the largest source of uncertainty on ν - N events comes from cross section modeling. Table 3.5 summarizes the effect of each source of uncertainty on the backgrounds and signal. Importantly, none of the systematic uncertainties induce a shape effect on the neutrino background in $E\theta^2$ that could mimic a false signal.

3.6 Analysis Results

The unblinding procedure for this analysis was done in stages to validate the background models used in the fit. First, data with $0.02 < E\theta^2 < 0.048 \text{ GeV Rad}^2$, which is expected to be dominated by νN events and have no present χe signal based on existing constraints, was compared to the nominal background prediction and then fitted. The nominal prediction agreed with data within cross section uncertainties, the largest source of systematic uncertainty in this region. Next, data in the first bin of the full analysis region ($E\theta^2 < 0.004 \text{ GeV Rad}^2$), which is expected to be dominated by νe events, was compared to the nominal prediction and agreed with data within neutrino flux uncertainties, which is the largest source of systematic uncertainty for the νe background. Finally, the full analysis region was unblinded and fitted.

After selection cuts, 4400.1 background events³ were expected. A breakdown of the production channels contributing to the expected background is shown in Table 3.6. From investigating events in the “Other (νN)” events, they all come from events with a π^0 in the final state, and the two photons from the π^0 decay were reconstructed as separate prongs. Due to a limitation of NOvA simulation software, each prong is assigned a “parent” neutrino that initiates the interaction and these neutrinos are uniquely assigned to one prong. With the parent neutrino already assigned to one photon prong in the π^0 decay, the other photon prong cannot be assigned a parent neutrino. This makes characterizing the nature of the interaction much more difficult, so these events are kept separate as an “Other” category⁴.

A total of 4407 events were selected from data with $0 < E\theta^2 < 0.048 \text{ GeV Rad}^2$.

³Note that the predicted number of events includes fractional events due to event weights applied in the simulation and the re-scaling of the simulated samples’ POT to the data’s POT.

⁴Although they are kept as a separate category, the fact that these events contain a π^0 makes it more likely that these are NC events.

Table 3.6: A summary of background processes and the nominal predicted amount passing selection cuts.

Background	Nominal prediction (2.55×10^{21} POT)
ν - e	1487.6
ν_e CC (ν - N)	1687.8
ν_μ CC (ν - N)	493.9
NC (ν - N)	492.3
Other (ν - N)	238.5
Total	4400.1

The data is consistent with the Standard Model, as a neutrino-only fit had a fit $-2\ln\mathcal{L} = 9.289$ with 12 degrees of freedom. The dark matter best fit has parameters $(m_\chi, Y) = (1 \text{ MeV}, 1.51 \times 10^{-12})$ with a fit $-2\ln\mathcal{L} = 7.132$ for 11 degrees of freedom. The best-fit parameters correspond to 85.2 predicted dark matter events in the data. The best fit slightly reduced the normalization of background events to account for the apparent presence of a signal. However, this helps to smooth out the discrepancies between the data and the background between $0.004 < E\theta^2 < 0.016 \text{ GeV Rad}^2$, where the presence of signal events are expected. Note that the best-fit parameters are already excluded by LSND [51] (not plotted) and NA64 [107] at 90% confidence level. All dark matter masses tested were weakly preferred over the neutrino-only fit. The best fit compared to data and the neutrino-only fit is shown in Figure 3.14.

In order to minimize $-2\ln\mathcal{L}$, some systematic uncertainties were pulled away from their nominal values, although all systematic uncertainties in the dark matter best fit and the neutrino-only fit were pulled away from their nominal value by less than one standard deviation. The largest systematic pulls in the dark matter best fit and neutrino-only fit are summarized in Table 3.7.

In addition to $E\theta^2$, data was compared to the nominal predictions in E and θ separately, shown in Figure 3.15. The data agrees generally well with the prediction




Figure 3.14: Top: The observed $E\theta^2$ spectrum compared to the dark matter best-fit prediction. Middle: The dark matter component of the best-fit prediction. Bottom: The ratio of data to the dark matter best-fit prediction and the neutrino-only best-fit prediction. In the top and bottom panel, the $\pm 1\sigma$ systematic error band is shown in dashed gray.

in these variables, with exception to a low-energy deficit of events in data. Note, however, that this is only a comparison to the nominal prediction, not the best fit, which may have better agreement with data.

Table 3.7: A summary of the largest systematic pulls in the dark matter best fit and neutrino-only fit.

Systematic	Description	$\nu + \chi$	ν -only
Detector	Absolute calibration	0.33σ	0.52σ
Cross section	NC elastic shape and normalization	0.21σ	0.21σ
Cross section	NC elastic shape	0.11σ	0.28σ
Cross section	MEC ($q_0, \vec{q} $) response	0.20σ	0.09σ
Cross section	MEC E_γ shape	0.07σ	-0.22σ
Cross section	MEC initial state neutron-proton pair fraction	0.02σ	0.18σ
Cross section	CC QE normalization	0.16σ	0.30σ
Flux	Principal component 0 of PPFX uncertainty	0.02σ	0.18σ




Figure 3.15: The nominal prediction compared to data in E (left) and θ (right).

The 90% confidence-level upper limit on Y is determined from the best fit at each mass using the profile likelihood ratio as a test statistic and the profile construction method [86]. Figure 3.16 compares the 90% confidence level exclusion results of this analysis with previous searches from BaBar [50], COHERENT [68], NA64 [107], and NOvA's expected sensitivity determined from pseudo-experiments with Poisson-fluctuated distributions of the background. The parameters that yield the relic abundance for scalar dark matter are shown as the dashed black curve [46]. This analysis improves on existing constraints on Y for $10 < m_\chi < 20$ MeV. At $m_\chi = 15$ MeV, where NOvA most strongly improves upon existing constraints, this analysis constrains the number of dark matter events in the data to below 222 at 90% confidence level.




Figure 3.16: Constraint on (m_χ, Y) parameter space for NOvA at 90% confidence level compared to other experimental data.

3.7 Discussion

High-luminosity neutrino experiments can be an effective probe for light dark matter candidates. Using 2.55×10^{21} POT of data recorded by the NOvA ND during NuMI FHC operations, this analysis has improved upon current constraints for dark matter particles produced in accelerators from 11 to 19 MeV.

There are several ways that NOvA's search for sub-GeV dark matter can be improved. The first is to consider proton interactions in the hadron absorber toward the end of the NuMI beam line. These interactions are effectively the same as proton interactions in the target and could produce dark matter. Simulation studies of π^0 flux in the NuMI beam dump show that up to 30 pions are being produced [98]. Whether these pions produce dark matter that has detector acceptance in the NOvA ND is to be determined. If so, these events are already present in the data. The inclusion of this production channel in the signal model would provide a more accurate account of the DM flux and improve exclusion sensitivity. The addition of the beam dump will also not scale up the background in the same way as for target production, since the charged hadrons produced in the beam dump are not focused into a beam, leading to fewer neutrinos reaching the ND per proton in the hadron absorber.

Another way to improve the analysis is to include RHC data. Not only will the inclusion of RHC data lead to a higher statistics measurement, but the slightly different physics processes of ν 's vs $\bar{\nu}$'s will lead to differences in the expected background. Because the dark matter is produced through decays of neutral mesons and proton bremsstrahlung upstream of the focusing horns, the signal is agnostic to the focusing horn current and one expects about the same number of signal events in FHC and RHC at a fixed exposure, whereas the interaction rates in the ND are different for

ν 's vs $\bar{\nu}$'s [108]. A correlated fit on FHC and RHC data would better constrain the presence of a dark matter signal in the two sets of data. NOvA has already collected over 1.2×10^{21} POT of RHC data, and plans to collect more in the future.

This analysis contributes to the effort to search for MeV-scale dark matter. The parameter space for dark matter consistent with thermal relic predictions still remains largely unconstrained. Future accelerator-based searches for light dark matter using detection techniques different from NOvA (and other experiments with world-leading sensitivity) will be able to significantly improve upon current constraints and probe thermal relic targets. One such search with the Light Dark Matter eXperiment will be discussed in the next chapter.

Part II

Visibly Decaying Dark Photons with LDMX

Chapter 4

The Light Dark Matter eXperiment

The Light Dark Matter eXperiment (LDMX) is a planned electron-beam, fixed-target, missing-momentum search for dark matter in the MeV to GeV mass range. LDMX will be mounted at the SLAC National Accelerator Laboratory, where it will parasitically use beam bunches from the Linac Coherent Light Source II (LCLS-II) with an energy of 8 GeV. In one year of operations, LDMX can accumulate approximately 4×10^{14} electrons on target (EOT), which is sufficient to achieve a 100-fold improvement on current sensitivity to sub-GeV dark matter. In 3.5 years of operation and a thicker target¹, LDMX can achieve a 1000-fold improvement on current constraints. This same 1000-fold improvement in sensitivity could also be achieved with a thin target and 10^{16} EOT, but this would take significantly more run time. In what follows, projections for LDMX’s ultimate sensitivity are presented as “ 1×10^{16} EOT.”

The production of dark matter via electron bremsstrahlung will be identified by a significant amount of missing energy and momentum measured in the detector after the primary beam electron’s interaction with a thin target. Processes in the Standard Model—namely electron bremsstrahlung in the target—can mimic the low-momentum recoil electron signature of dark matter production. To identify and veto these background processes, which are many and varied, LDMX will utilize three major subsystems:

¹An increase in target thickness from 0.1 to 0.4 radiation lengths.

1. a tracking system to measure the incident and recoil momentum of the beam electrons, as well as charged particle multiplicity,
2. an electromagnetic calorimeter (ECal) to measure the energy of the recoil products and veto backgrounds based on shower topology in the calorimeter, and
3. a hadronic calorimeter (HCal) to veto neutral hadrons produced in photonuclear reactions in the ECal or target.

A plan view of the LDMX detector is shown in Figure 4.1. This figure shows a view of the detector in the yz plane, and the origin of the coordinate system is at the center of the target.




Figure 4.1: A plan view of the LDMX detector, showing the production of dark matter in the target. The drawing is not to scale. From Ref. [109].

A missing-momentum-style search for dark matter is much more sensitive to the coupling strength between dark matter and the Standard Model than a beam-dump-style search at the same integrated luminosity, as explained in Chapter 1. But the caveat

to this improved sensitivity is that missing-momentum-style experiments require a hermetic detector with precision tracking and the ability to measure the incident particle's momentum. This chapter will describe the experimental design of LDMX and the missing-momentum search for dark matter.

4.1 The Beamline

For LDMX to achieve its physics goals, a low multiplicity, high duty-factor beam is required. A low multiplicity beam is required so that individual electrons can be tagged by the tracker. A high repetition is required, partially to counter the low multiplicity, so that high statistics measurements can be achieved in a reasonable time scale. LDMX will operate in End Station A at SLAC and use the refurbished Linac to End Station A (LESA) beamline. This will deliver a 46 MHz beam of 8 GeV electrons from the LCLS-II beamline to the experiment [110].

The electrons in LCLS-II are sourced from a high-power ultraviolet laser pulse on a Cs₂Te photocathode. The electrons liberated from the photocathode are bunched and accelerated by a very high radio frequency (VHRF) cavity to 750 keV, where they are then sent to the main superconducting linac [111]. In the main linac, the electrons are accelerated to 8 GeV using a series of 37 superconducting cryomodules containing eight niobium cavities [112]. Both the VHRF cavity injector and main linac operate at 1.3 GHz [113].

Low-multiplicity buckets from the beam go unused by the experiments for which the LCLS-II beam is primarily designed. To make use of the low-multiplicity electron buckets from this beam (which would otherwise go to a beam dump), a kicker is added to the LCLS-II beamline to deliver single electrons in 21.5 ns bunches to End

Station A, as shown in Figure 4.2 [110].




Figure 4.2: A diagram showing the Linac to End Station A beamline (labeled “DASEL Beamline”). From Ref. [114].

4.2 The Detector

The LDMX detector consists of three major subsystems to measure missing energy and momentum as a signature of dark matter production in the target. Each subsystem is based on the design of existing detectors.

4.2.1 Tracking System

The key concept behind LDMX’s method to search for dark matter is the missing momentum method, where the recoil electron momentum is measured by a tracker. The LDMX tracking system consists of an tagging tracker upstream of the target to measure the number of incoming beam electrons, their positions, and their four-momenta before interaction with the target and a recoil tracker downstream of the

target to measure the four-momenta of outgoing charged particles and their trajectories. The LDMX trackers consist of silicon strips based on the design and technology of the Silicon Vertex Tracker used in the Heavy Photon Search [115].

Measuring the momentum of charged particles requires a magnetic field. LDMX will use a 18D36 dipole magnet (18" by 36" in size) with a central field of 1 T. The tagging tracker will be located in the uniform field of the magnet, and the recoil tracker will be located in the fringe field of the magnet. The target is a thin tungsten sheet corresponding to 0.1 radiation lengths placed at the start of the magnet's fringe field. A diagram of the magnetic field strength along the beam direction, along with the position of the tagging tracker, target, recoil tracker, and ECal, is shown in Figure 4.3.




Figure 4.3: The magnetic field strength as a function of position along the beamline. Each colored line represents an off-axis distance from the magnet's central field.

The tagging tracker consists of 7 silicon strips, each spaced 10 cm apart (in the beam direction), with the furthest downstream strip placed 1.5 cm from the target. The recoil tracker consists of 4 silicon strips spaced 1.5 cm apart, with the most upstream

strip being 1.5 cm from the target, and two additional strips placed 10 cm from the target and 18 cm from the target, respectively. Both the tagging tracker and recoil tracker have position resolutions of $7\mu\text{m}$ in the x -direction (perpendicular to the magnetic field) and $100\mu\text{m}$ in the y -direction (parallel to the magnetic field) and are capable of reconstructing tracks for charged particles with a momentum of at least $50\text{ MeV}/c$. The momentum resolution for electrons with a few GeV of energy is $\sim 3\text{ MeV}/c$.

4.2.2 Electromagnetic Calorimeter

The LDMX ECal is a high-granularity, high-density silicon-tungsten sampling calorimeter. The ECal design is based on the design of the CMS High Granularity Calorimeter (HGC) for the CMS Phase-2 upgrade [116]. It fulfills LDMX’s requirement for a fast, efficient, radiation hard, high density, and granular calorimeter system.

A diagram of the ECal is shown in Figure 4.4. The ECal is 17 planes deep, with each plane consisting of double-sided layers of tungsten absorbers, silicon sensors, and printed circuit boards (PCBs), with a cooling plate sandwiched between each layer. Each layer contains 7 hexagonal silicon sensors arranged in a “flower” configuration, where each sensor contains 432 readout pads. Each readout pad is 0.56 cm^2 in area, allowing high-granularity shower reconstruction needed to resolve individual electrons and photons with small angular separation. Each plane is 6 mm thick and the entire ECal is 40 radiation lengths deep, providing ample electromagnetic shower containment.

The ECal has ns-level timing resolution, making it capable of resolving individual events from the beam pulses that arrive every 21 ns. The ECal has a Molière radius




Figure 4.4: A diagram of the LDMX ECal.

of ~ 2.5 cm; when combined with the high granularity of the readout pixels, the ECal is capable of resolving individual recoil products from the target interaction. The ECal is also capable of MIP tracking, and the identification of MIPs will be an important veto handle in both the missing momentum analysis (Section 4.4) and searches for long-lived particles (Chapter 5). An example of event displays highlighting the difference between a MIP and an electromagnetic shower is shown in Figure 4.5. The ECal is 40 radiation lengths deep, which allows late-developing showers to be resolved in the detector volume. The tungsten absorbers in the front layers of the ECal are thinner (1–3.5 mm thick) than the absorbers in the final 10 layers (7 mm thick) so that particle shower shapes can be better resolved in the front of the detector while still keeping full shower containment toward the back of the detector.




Figure 4.5: Top: Event display of a MIP traveling through 8 layers of the CMS HGC. Bottom: Event display of an electron traveling through 8 layers of the CMS HGC. Both from Ref. [117]. Note that the LDMX ECal is based on the design of the CMS HGC.

As the beam will not always have an electron multiplicity of one, the number of incident electrons must be known to inform the trigger. A trigger scintillator upstream of the target measures the incident number of electrons, and information from the ECal measures the energy after electrons pass through the target. Information from both the trigger scintillator and ECal are input into the trigger DAQ, and events are kept if at least one electron in the bunch loses $\sim 60\%$ of its energy.

4.2.3 Hadronic Calorimeter

Surrounding the ECal from the top, bottom, sides, and back is the HCal. The portion of the HCal surrounding the top, bottom, and sides of the ECal will be referred to as the Side HCal and the portion behind the ECal will be referred to as the Main HCal. The LDMX Hcal is an iron-scintillator sampling calorimeter, with the active

component based on the scintillation counter design used by the Mu2e cosmic ray veto [118].

The primary function of the HCal is to identify and veto neutral hadrons (mostly K^0 and neutrons) in the energy range from hundreds of MeV up to several GeV. Neutral hadrons that escape the detector and carry away several GeV of energy look like a false signal event. To ensure these neutral hadrons interact in the HCal and leave a detectable amount of energy to veto, the Main HCal is 17 interaction lengths deep.

A diagram of the HCal is shown in Figure 4.6. The Main HCal consists of 96 layers of 25 mm thick steel absorber plates and 20 mm thick doped polystyrene scintillator counters. Each layer of the Main HCal is 2 m tall and 2 m wide. Each scintillator counter is 50 mm wide and is co-extruded with a TiO_2 reflector coating and a single hole into which a 1.8 mm diameter wavelength-shifting fiber is inserted. Light produced in the scintillator is collected by the wavelength-shifting fiber and read out via silicon photomultipliers (SiPMs). The scintillator counters are glued into units of four counters, called a quad-counter, and each layer is comprised of 10 quad-counters. Similar to the NOvA detector described in Chapter 2, the horizontal and vertical orientation of the scintillator bars alternates with each layer. However, currently the position reconstruction along the length of a bar is only based on time-of-arrival of photons, not a clustering algorithm based on hits in neighboring layers.

The Side HCal consists of four stepped modules arranged in a pinwheel-like fashion around the ECal. The primary purpose of the Side HCal is to catch wide-angle particles emitted from the target that fall outside the ECal acceptance. Unlike in the Main HCal, the Side HCal scintillator counters are read out at only one end.




Figure 4.6: Isometric view of the LDMX HCal.

4.3 Simulation Framework

The LDMX simulation framework uses v10.2.3 of the Geant4 [90–92] library, with modifications to the nuclear interaction cross sections and kinematics that are relevant for LDMX’s beam energy. The framework is written in C++, and uses Python wrappers to configure Monte Carlo event simulation, detector response, and event reconstruction.

Events in which the full beam energy is deposited in the ECal are not triggered. As indicated in Figure 4.7, events with an appreciable amount of missing energy in the ECal occur at relative rates below 10^{-4} , so inclusive background samples are an inefficient use of computing resources for the backgrounds of interest to LDMX. To efficiently generate background samples of interest to LDMX, event generator filters and cross-section weights are utilized.

Simulations start with a beam of electrons fired 800 mm upstream of the target, positioned and aimed in such a way that electrons hit the center of the target. The beam is “smeared” in the x and y direction by 20 mm and 80 mm, respectively, to mimic the desired beam profile. Truth-level filters are applied to remove events in which the electron does not reach the target with sufficient energy (>7600 MeV) and does not undergo “hard” bremsstrahlung in the target. The latter is imposed by requiring the electron exit the target with <3000 MeV energy and a bremsstrahlung photon is produced with >5000 MeV energy. Events not satisfying these requirements are vetoed. Other filters can be applied to only keep events in which an interaction takes place in a specific region of the detector; for example, a filter to only keep processes that occur in the ECal.

To enhance the likelihood of particular interaction types (like background photo-nuclear interactions, which are of particular interest for LDMX) and more efficiently produce samples of these particular interactions, a weight to the photon cross section is applied. These weighting factors are propagated through the simulation framework to account for how this weighting affects the equivalent EOT of the sample.

To simulate the process of dark bremsstrahlung that produces signal events, a new tool for Geant4 was developed by LDMX [119]. This Geant4 tool was validated against the ubiquitous dark matter event generator MadGraph [120, 121], showing similar results for the beam energies and target materials/thicknesses relevant to LDMX. This tool was developed to simplify the simulation of signal events, which now only require Geant4 as a simulation library.

4.4 Missing-Momentum Search

The main physics objective of LDMX is to identify the production of dark matter in the target via missing energy and momentum in the detector. The background processes expected in LDMX are many and varied. A breakdown of the expected backgrounds, their rates relative to the beam, and the method used to veto them is shown in Figure 4.7. This section will briefly describe the flagship analysis for LDMX, which will surpass the current experimental sensitivity on MeV-scale dark matter by orders of magnitude. The description of this analysis is primarily to provide context for some of the analysis methods described in Chapter 5.




Figure 4.7: The relative rates of background processes in LDMX, along with the veto handles used to reject them, down to a relative rate of 10^{-16} . The green shading indicates events in which the full beam energy is deposited in the ECal, the blue shading indicates events with low energy depositions in the ECal or particles escaping the ECal, and the red shading indicates irreducible Standard Model backgrounds with real missing energy in the detector. Note that most background sources have more than one veto handle.

For the missing-momentum search, the signal signature is a single, low-momentum electron reconstructed in the recoil tracker and calorimeters. For trigger selection, events are recorded in which <3160 MeV is deposited in the first 20 layers of the ECal. This particular energy requirement was chosen based on trigger studies to keep the signal efficiency of 1 MeV dark matter at 78%. The trigger rate is ~ 250 Hz for a 46 MHz beam. Only the first 20 layers of the ECal are used for triggering instead of the full 34 layers to reduce the load on the data acquisition system. Following the trigger selection, a more stringent requirement on the missing energy in the ECal is required: <3160 MeV deposited in the whole ECal. Additionally, events are required to have exactly one track in the recoil tracker with a momentum less than 2400 MeV/c. The recoil tracker has high momentum resolution, so a more stringent requirement on the momentum is made. In the results shown here, this cut is made on the truth-level momentum of charged particles in the recoil tracker, not a simulated reconstruction of the tracks². Since the true momentum is used for this cut, an additional requirement that the charged particle have at least 50 MeV/c of momentum is made. This is to mimic the reconstruction efficiency of the tracker, which can resolve particles with momentum as small as 50 MeV/c. Although not explicitly used as an analysis handle, the transverse momentum of the recoil electron can be a powerful discriminator between signal and background, as seen in Figure 4.8.

Background events selected by the trigger that also pass the tracker momentum requirement will predominantly be events in which the electron loses most of its energy to a bremsstrahlung photon in the target. This bremsstrahlung photon may not always deposit all of its energy visibly in the ECal. This can occur in photo-nuclear (PN) interactions, which produce neutral or charged hadrons. Charged hadrons can

²At the time of writing, not all simulated samples had reconstructed tracking information. In order to use consistent cuts among the samples, a truth-level cut is used.




Figure 4.8: Transverse momentum (p_T) distribution of the recoil electrons for events passing the ECal missing energy trigger. Photo-nuclear events are background, and distributions from four signal masses are shown. The rightmost bin is an overflow bin. From Ref. [109].

be vetoed based on the presence of additional tracks in the ECal or tracker. Neutral hadrons are more difficult to veto based on ECal activity alone, but it is possible to distinguish these backgrounds from signal, as the interaction creating neutral hadrons may leave an energy deposition signature in the ECal. The high granularity of the ECal is utilized to determine whether the recoil electron shower is accompanied by other, photon-induced energy depositions. To discriminate between signal and photon-induced background events, a boosted decision tree (BDT) trained on ECal shower shape features is used.

The ECal BDT, implemented with XGBoost [122], was trained on 47 input features based on the topology of signal and PN background activity in the ECal. A detailed description of all the features included in the ECal BDT can be found in Ref. [109]. The BDT was trained on four dark matter masses (1, 10, 100, and 1000 MeV) as signal and PN reactions in the ECal as background. Only the PN events were used in training as the background because it is expected to be the most difficult background to veto. The BDT’s discriminator threshold was set to a value of 0.99741, which

was chosen to remove all but two MC events when combined with the HCal veto, described below. This is so events passing all cuts can be studied and to facilitate projections to higher EOT than what has been simulated.

Due to the generally low transverse momentum of recoil electrons in background bremsstrahlung events, as shown in Figure 4.8, photon-induced activity in the ECal may not always be distinguishable from the recoil electron’s shower. This is especially problematic if neutral hadrons carrying away most of the photon’s energy are produced, in which case the activity in the ECal looks indistinguishable from signal. As the HCal is 17 interaction lengths deep, these neutral hadrons produced in the ECal will eventually interact in the HCal and leave a deposition of energy. To discriminate events producing neutral hadrons, events are vetoed if any HCal counter has 8 or more photo-electrons (PEs). A threshold of 8 PEs is sufficient to remove noise hits, most of which are at the 1 PE level but extend up to 6 PEs.

This HCal veto, combined with the ECal BDT, will provide the necessary veto handles for LDMX to be a background-free search while maintaining a high signal efficiency. Figure 4.9 shows the combination of the ECal BDT and HCal veto, with the distribution of the background events’ BDT scores and maximum PEs in the HCal.

A summary of the full set of selection cuts on the background and signal is shown in Table 4.1. Two events in which a PN reaction occurs in the ECal remain after all cuts. These two remaining events are used to characterize background events that look “signal-like,” but can be vetoed with a higher discriminator threshold on the ECal BDT.

Through all of these veto handles, LDMX is able to achieve a background-free search




Figure 4.9: Distribution of BDT discriminator values and the maximum number of photo-electrons in the HCal for 1 GeV signal (left) and PN background (right) events. The BDT discriminator value has been transformed to $-\log(1 - \text{BDT score})$. The color scale shows the number of events in the bin.

Table 4.1: Number of simulated background events remaining at different selection stages of the missing-momentum search. The background classes are separated by where in the detector the reaction took place.

	Photo-nuclear		Muon conversion		Signal efficiency for different $m_{A'}$ (%)			
	Target	ECal	Target	ECal	1 MeV	10 MeV	100 MeV	1000 MeV
EOT equivalent	10^{14}	10^{14}	10^{15}	10^{14}				
Trigger (front ECal energy < 3160 MeV)	1.77×10^7	1.61×10^8	4.56×10^6	1.47×10^7	78.2	89.6	91.3	94.3
Total ECal energy < 3160 MeV	1.22×10^7	3.87×10^7	3.46×10^6	1.02×10^7	78.2	89.6	91.3	94.3
Single track with $p < 2400$ MeV/c	69282	3.50×10^7	< 1	9.86×10^6	77.1	88.4	90.0	93.0
ECal BDT	3919	61145	< 1	< 1	52.9	68.6	73.1	67.8
HCal max PE < 8	< 1	2	< 1	< 1	49.9	64.3	68.3	59.3

for sub-GeV dark matter at the 4×10^{14} EOT level. A zero-background search is also anticipated for the full dataset, 1×10^{16} EOT, although this has not been explicitly verified through simulation. The ECal BDT plot shown in Figure 4.9 does not include explicit MIP tracking cuts. However, if MIP tracking cuts are included (no candidate MIP tracks in the ECal), the background is reduced to <1 event. Since the making of this plot, improvements have been made to include MIP track candidates as features in the ECal BDT, which has improved the BDT's performance. An explicit cut on MIP

tracks is no longer required, and LDMX can achieve a zero-background search. The expected sensitivity of LDMX compared to other experiments is shown in Figure 4.10, which assumes 1×10^{16} EoT collected.




Figure 4.10: The projected sensitivity of LDMX with various background scenarios. Three thermal relic targets are shown as black curves, corresponding to dark matter particles that have (from top to bottom) scalar, Majorana fermionic, or pseudo-Dirac fermionic nature. A 50% signal efficiency is assumed for all m_χ . From Ref. [109].

Chapter 5

Visibly Decaying Dark Photons in LDMX

In addition to the missing momentum search, LDMX offers a diverse physics program due to its unique detector design. LDMX effectively operates as a fully instrumented beam dump with large acceptance and thus is ideal for searching for visible decays of exotic particles that have a fairly long lifetime. Preliminary estimates of such searches have been made for the minimal dark photon model, axion-like particles (ALPs), and Strongly Interacting Massive Particles (SIMPs) [46]. For the minimal dark photon model, the kinematics of A' production are identical to the missing momentum search, in which the A' carries most of the beam energy. Whereas the missing momentum search probes the mass hierarchy scenario $2m_\chi < m_{A'}$, a “visibles” search probes the complimentary scenario $2m_\chi > m_{A'} > 2m_e$. While it is possible for the A' to decay into $\mu^+\mu^-$ if $m_{A'} > 2m_\mu$, LDMX is not able to kinematically produce them for most of the mass range being probed. Also, in the area where they could be kinematically produced, LDMX would not be competitive with other experiments [123, 124]. Thus, this work focuses on the electron decay channel for a visibly decaying A' . The expected signal is an energy deposition near the beam energy in the HCal accompanied by a soft recoil electron in the ECal, as seen in Figure 5.1. The same coordinate system as the missing-momentum search is used: the origin is at the center of the target and the beam travels in the $+z$ direction.

In this work, a reach estimate for a displaced A' visibles search was performed using a

full detector simulation for the expected run schedule of an 8 GeV electron beam with 1×10^{14} electrons on target (EOT). The results from 1×10^{14} EOT are extrapolated to roughly estimate the sensitivity for the anticipated full dataset of 1×10^{16} EOT. First, our trigger and event selection cuts are discussed. Next, the primary backgrounds of concern are discussed—photo-nuclear (PN) and bremsstrahlung conversions—with strategies to mitigate these processes while preserving high signal efficiency. Finally, we calculate our projected sensitivity based on our expected background and signal efficiency after all analysis cuts.




Figure 5.1: A plan view view of the LDMX apparatus with visible search sketch overlaid and subsystems labeled. Instead of energy escaping the detector like in Figure 4.1, a large deposition of energy is displaced from the recoil electron shower.

5.1 Simulated Event Samples

In this section, we discuss the simulation of the visibly decaying A' signal sample and background samples.

5.1.1 Visible A' Simulation

The visibly decaying A' simulation used MadGraph [120, 121] to calculate the kinematics of both the A' production in the target and the decay of the A' to an e^+e^- pair. The A' mass, the allowed decay window of the A' , and the proper lifetime of the A' are configurable in the simulation. Four A' masses were simulated: 5, 10, 50, and 100 MeV. The lower limit of 5 MeV was chosen to avoid the resonance around 1 MeV and because the phase space of these lighter A' s has been thoroughly searched in LDMX's expected sensitivity region. A “uniform” decay was chosen for the A' , meaning it is equally likely to decay at any path length after production. This was done to ensure a high-statistics sample of A' decays was generated over the entire detector volume, so that the efficiency of selection cuts could be accurately determined. As described in Section 5.3.2, a specific mass $m_{A'}$ and coupling ε can be mapped to an expected number of events if the efficiency is known. The output of the MadGraph simulation was then passed into Geant4, which simulates the detector response to the recoil electron and the e^+e^- pair from the A' decay.

5.1.2 Background Simulation

The signal signature is an energy deposition near the beam energy in the HCal. There are several background processes that can mimic this signature and thus must be

mitigated. Though the visibly decaying A' channel is a complimentary search to the standard missing momentum search, the backgrounds are largely the same underlying physical process. Thus, many of the same background samples used to understand the backgrounds in the missing momentum search were used for this analysis.

The main background of concern is a hard bremsstrahlung photon from the target that undergoes a PN interaction in the ECal and produces hadrons. If these hadrons deposit energy in the HCal near the beam energy, they could be mistaken for the signal. PN interactions can also occur in the target. The following samples were used to characterize the PN-induced background:

- 1.00×10^{14} EOT enhanced sample of ECal PN interactions, and
- 1.03×10^{14} EOT enhanced sample of target PN interactions.

In addition, gamma conversions to muon pairs, from either the target or ECal, can occur but can be easily and sufficiently mitigated. The following samples were used to characterize the muon pair background:

- 1.00×10^{14} EOT enhanced sample of ECal di-muon production, and
- 1.00×10^{15} EOT enhanced sample of target di-muon production interactions.

A summary of the weights and thresholds on particle energies in order to weight them for each of these four samples is detailed in Figure 5.1. With the exception of muon production in the target, only up to 1×10^{14} EOT were simulated in these samples, not 4×10^{14} or 1×10^{16} . This was due to computational constraints—both processing time and storage—in generating larger event samples.

Table 5.1: Summary of the filter thresholds and weights used in the background simulations.

	Photo-nuclear		Muon production	
	Target	ECal	Target	ECal
EOT equivalent	10^{14}	10^{14}	10^{15}	10^{14}
Interaction cross section weight	450	550	10^6	10^5
Bremsstrahlung photon energy threshold (MeV)	4640	5000	4840	4840

5.1.3 HCal Reconstruction

Simulated hits in the HCal are digitized in the LDMX software environment to convert truth-level information (energy, position) into digitized data (time over ADC threshold, time of arrival). This process includes summing the energy for all hits within each HCal counter and emulating the readout with the readout board that will be used in the experiment (HGCROC). The total number of photoelectrons (PE) associated with each energy deposit is estimated using standalone measurements of the scintillator response to muons. Specifically, the average energy deposited by muons traversing a scintillator bar is 4.66 MeV, and such a minimum ionizing particle¹ (MIP) generates 68 PEs. Once the number of PEs associated with each energy deposit is computed, the energy is converted to MIPs, and the simulated SiPM pulse is sent to the shaper. Electronic noise hits are added in scintillator counters that do not contain simulated hits. Finally, digitized samples are converted into reconstructed HCal hits by applying the energy calibration to determine the energy response in MeV.

The HGCROC board employs two methods for measuring the amplitude of the signal, depending on the size of the pulse. If the signal is small enough (below 34 MeV), it is sampled by the HGCROC's 10-bit Analog to Digital Converter (ADC). If the signal exceeds the operating range of the ADC, the HGCROC measures the time during

¹A MIP is a charged particle that has a minimum amount of energy loss when traveling through a material. Muons are minimum ionizing at energies of a few hundred MeV.

which the shaper output is above a threshold, known as the time over threshold (TOT) measurement². In addition to measuring the deposited energy, the HGCROC performs a timing measurement of the incoming particles, known as the time of arrival (TOA).

There were two limitations present on the simulated HCal digitization and reconstruction. First, the TOT to energy conversion was not yet implemented in the reconstruction, which results in an upper bound on the energy measurement per hit of 34 MeV (approximately 280 PEs). This affected the energy reconstructed in events that left particularly large amounts of energy deposited in a single bar, such as electromagnetic showers of several GeV like the signal in this analysis. Second, the time of simulated hits did not include photon transport studies through the scintillator and WLS fibers, making TOA measurements less robust. These features will be included in a future version of the simulation.

5.2 Event Selection

This section describes the trigger and event selection cuts used for the analysis, and details how the event selection vetoes specific background processes expected in LDMX. This analysis is optimized for A' decays in the HCal. Since the kinematics of A' production in a visibly decaying scenario are the same as those in an invisibly decaying scenario, the same physics signatures are expected in the recoil tracker and ECal for this analysis as in the missing-momentum analysis, described in detail in Section 4.4. Therefore, this analysis used the same trigger, tracker veto, and ECal event selection cuts as the missing-momentum search—selection cuts for which LDMX was

²Note that this is only true in the simulation, as in reality the TOT is measured on the output of the pre-amplifier, before the ADC shaper.

optimized.

Selection cuts for this analysis diverged from the missing-momentum search in the HCal; instead of requiring no activity in the HCal, a large amount of energy deposited in the HCal was required. Section 5.2.1 describes the HCal-specific event selection cuts, including a boosted decision tree (BDT) to distinguish the electromagnetic showers of $A' \rightarrow e^+e^-$ from hadronic showers originating from PN-induced backgrounds. These cuts were then tested on specific background processes that were expected in the analysis: PN reactions (Section 5.2.2.1) and $\gamma \rightarrow \mu^+\mu^-$ production (Section 5.2.3). The total background after selection cuts is summarized in Section 5.2.4. Finally, possible improvements to the analysis are discussed in Section 5.2.5, along with an extrapolated estimate of the background from 1×10^{14} EOT (the EOT equivalence of the simulated samples) to 1×10^{16} (LDMX's expected full dataset).

5.2.1 HCal Selection Cuts

This analysis differed from the missing-momentum analysis by its requirements on HCal activity. Three cuts based on reconstructed HCal hits were used to select for the visible decay of an A' . First, a visibly decaying A' had to leave an energy deposition in the HCal equal to the amount lost by the incident electron. Since the energy requirement for the recoil electron was less than 3160 MeV deposited in the ECal, greater than 4840 MeV was required to be deposited in the main HCal³. Second, in order to resolve the decay vertex, the shower was required to be fully contained in the HCal, with a requirement of more than 5 PEs deposited in the first layer of the main HCal.

³Note that the requirement was made on the main HCal, not the side HCal, as the A' is very forward and highly unlikely to decay in the side HCal.

5.2.1.1 The Visibles Boosted Decision Tree

Background events that leave very little activity in the ECal and a large amount of energy in the HCal can originate from PN reactions that produce neutral hadrons (such as K_S, K_L, n) and muon pair production. These backgrounds primarily leave MIP tracks or hadronic shower signatures in the HCal, while the signal leaves an electromagnetic shower. To mitigate this remaining background, a BDT was used to distinguish between electromagnetic and hadronic showers in the HCal.

A BDT, hereinafter called the Visibles BDT, was implemented with XGBoost [122] and trained on shower shape and other kinematic features to distinguish signal from background. The Visibles BDT utilized twelve features: the total number of HCal layers hit; the total number of HCal hits; the energy-weighted hit x , y , and r mean position ($r = \sqrt{x^2 + y^2}$); the energy-weighted hit x , y and z standard deviation; the number of isolated hits; the isolated energy; the total energy deposited in the HCal; and the energy-weighted mean hit distance from the projected photon track. Distributions of these features for signal and background are shown in Figures 5.2 and 5.3.

For the HCal, hits were defined to be “isolated” if there were no additional hits in a neighboring counter in the same layer. The energy-weighted mean and standard deviation of the hits in x (as an example) were defined as:

$$\mu_x = \frac{\sum_i x_i E_i}{E_{\text{tot}}}, \quad \sigma_x = \sqrt{\frac{\sum_i E_i (x_i - \mu_x)^2}{E_{\text{tot}}}}. \quad (5.1)$$

Similar definitions follow for other energy-weighted features.

The mean distance of hits from the projected photon track was used as a strong dis-




Figure 5.2: Six of the twelve BDT features showing the distributions for four signal A' masses and the ECal PN background: isolated energy (top left), isolated hits (top right), layers hit (middle left), mean distance from the photon track (middle right), total HCal energy (bottom left), and total HCal hits (bottom right). Distributions are normalized so that their area is one.




Figure 5.3: The remaining six BDT features showing the distributions for four signal A' masses and the ECal PN background: mean hit x (top left), standard deviation in hit x (top right), mean hit y (middle left), standard deviation in hit y (middle right), mean hit r (bottom left), and standard deviation in hit z (bottom right). Distributions are normalized so that their area is one.

criminator between signal and background. For bremsstrahlung—dark or Standard Model—in the target, there are two outgoing particles: the recoil electron and the unobserved dark or Standard Model photon. The recoil tracker measures the momentum of the recoil electron and allows, with the tagger tracker, the momentum of the unobserved particle to be determined. The track of the dark or Standard Model photon was then projected to the HCal, and the distance in the (x,y) plane of the reconstructed hits from the track was calculated. Signal events from $A' \rightarrow e^+e^-$ have an electromagnetic shower centered on this projected axis, so the mean distance of all hits from this projected momentum axis is expected to be small. For an ECal PN event, high particle multiplicities in the final-state could cause these secondary particles to deviate from the photon’s projected momentum. Whether energy is deposited in the HCal from a single particle carrying most of the energy⁴ or from multiple particles, the mean distance of hits from the projected track in background events was larger than for signal events.

The Visibles BDT was trained on ECal PN background events and signal events composed of an approximately equal mix of four A' masses. The number of events in each sample used for training is shown in Table 5.2. All events included for training passed the trigger, the HCal energy requirement, and the tracker veto. Note that the ECal BDT cut was excluded in the training samples. The ECal BDT cut removes a significant fraction of all background events, so this cut was excluded to ensure a high-statistics sample for training. Even if these events had been vetoed before the Visibles BDT cut due to their topology in the ECal, they held useful information for the Visibles BDT on the topology of background events in the HCal.

⁴In the case that a single particle remained on-axis with the photon track, the particle could still leave a hadronic shower, which would have a larger mean distance than a signal event due to hadronic showers generally being wider than electromagnetic showers.

Table 5.2: Number of background and signal events used for training the Visibles BDT.

Sample	Events for training
ECal PN	66409
5 MeV A'	62784
10 MeV A'	69499
50 MeV A'	68909
100 MeV A'	66245

The performance of the Visibles BDT was then tested on a separate sample of ECal PN events and each signal mass individually. ROC curves⁵ and BDT output values from this testing are shown in Figure 5.4.




Figure 5.4: Visibles BDT ROC curve (left) and discriminator value output (right).

The discriminator threshold used for the Visibles BDT was optimized based on a Figure of Merit (FOM). Since the expected number of signal events was not known a priori and depends on both the A' mass and kinetic mixing ε , a traditional FOM like $S/\sqrt{B}$ or $S/\sqrt{S+B}$ could not be used. Instead, a FOM proposed by Punzi [125] was

⁵Receiver Operating Characteristic (ROC) curves, which show the true positive rate vs the false positive rate as a function of the BDT's discriminator threshold.

used:

$$\frac{\epsilon_s}{a/2 + \sqrt{B}}, \quad (5.2)$$

where ϵ_s is the signal efficiency, a is the number of standard deviations corresponding to a one-sided Gaussian test at a chosen significance (in this case, $a = 3$), and B is the number of background events. This FOM did not depend on the number of signal events, and instead depended on the signal efficiency. The FOM for each of the four test masses is shown in Figure 5.5, along with a red line that shows the discriminator value chosen. One can clearly see the discontinuities in the FOM each time there is a change in the number of background events. The point at which the FOM becomes a smooth function again is the region of zero background events, which is where the FOM is highest. A discriminator value at the maximum of the FOM is the optimal value, and at this optimal value the number of background events drops to zero. However, a different sample of background events from the ones simulated for this study may have a slightly different distribution of BDT discriminator values, including higher values than the optimum of the FOM. To insure a zero background search for other samples of background events, a discriminator slightly higher than the optimum point was chosen at the cost of reduced signal efficiency. The chosen discriminator value for the Visible BDT was 0.999965.

5.2.2 Backgrounds

The event selection cuts described above were tested on specific background processes expected in LDMX. A description of these specific processes and the summary of the total background are described here.




Figure 5.5: The FOM as a function of the discriminator value on the Visible BDT for four test masses. The red line indicates the thresholds chosen.

5.2.2.1 Photo-nuclear backgrounds

The selection cuts presented above were tested on a $\sim 1 \times 10^{14}$ EOT equivalent sample of ECal PN events.⁶ A breakdown of the cut flow for this sample is shown in Table 5.3. No ECal PN events remained after all selection cuts.

The selection cuts were also tested on a 1.03×10^{14} EOT equivalent sample of target PN events. Similarly to the ECal PN sample, no target PN events passed all selection cuts.

One type of background from PN reactions that was of particular concern for this analysis was the production of neutral kaons, K_S or K_L , that decay to π^0 s. The

⁶One batch of ECal PN events was reserved for training the Visible BDT, so the EOT equivalence is slightly less than 1×10^{14} . However, this reserved batch is only 2% of all trigger-skimmed ECal PN events, so approximately 1×10^{14} EOT events were tested.

π^0 s promptly decay to $\gamma\gamma$, leaving an electromagnetic shower in the HCal that is hard to distinguish from a signal event. While the selection cuts presented here were able to fully mitigate this background at 1×10^{14} EOT (the extent of the simulated samples), this is not guaranteed for higher exposures. Understanding the rates of these backgrounds at 1×10^{16} EOT is critical for an accurate projection of this analysis's sensitivity to the full expected data set.

5.2.3 Muon Pair Production

Background events due to photons undergoing muon pair production pass the trigger at a rate higher than that of photons undergoing PN conversions, but they are significantly less challenging to veto. These events can originate from either the target or ECal. These events can most easily be vetoed by the straight tracks they leave in the ECal and HCal, which are caught by the BDT. No events from muon pair production in either the target or ECal passed all selection cuts.

5.2.4 Total Background

Table 5.3 shows the full breakdown of the analysis cuts for each background sample and for two signal masses. No background events remain at 1×10^{14} EOT.

5.2.5 Full Dataset Prospects

With the current version of this analysis, a background-free search for visibly decaying dark photons can be achieved at 1×10^{14} EOT. It might be possible to achieve a zero-

Table 5.3: Number of simulated background events remaining at different selection stages. Each type of background is separated by where in the detector the reaction took place.

	Photo-nuclear		Muon conversion		Signal efficiency (%)	
	Target	ECal	Target	ECal	10 MeV A'	100 MeV A'
EOT equivalent	1.03×10^{14}	1.00×10^{14}	1.00×10^{15}	1.00×10^{14}		
Trigger (front ECal energy < 3160 MeV)	1.77×10^7	1.61×10^8	4.56×10^6	1.47×10^7	75.4	69.1
Total ECal energy < 3160 MeV	1.22×10^7	3.87×10^7	3.46×10^6	1.02×10^7	75.2	69.0
Single track with $p < 2400$ MeV/c	69282	3.50×10^7	< 1	9.86×10^6	69.0	65.0
Missing-momentum ECal BDT	3919	61145	< 1	< 1	64.1	62.5
Main HCal energy > 4840 MeV	2190	55283	< 1	< 1	61.7	58.9
Containment	722	30665	< 1	< 1	61.7	58.9
Visibles BDT	< 1	< 1	< 1	< 1	33.4	27.0

background search for visibly decaying A' 's in the HCal at 1×10^{16} EOT, as there are ways to improve this analysis.

The ECal BDT used in the missing momentum analysis can be improved and updated. Some preliminary studies within LDMX on the use of a neural network, like ParticleNet [126], show that this type of machine learning model may be more effective as a veto than a BDT. The ECal's high granularity, with over 100,000 readout pixels, is more effectively leveraged with a neural network than a BDT. Since this analysis is background-free with a BDT, it is likely that a neural network can keep the analysis background-free at the same or larger EOT.

The Visibles BDT can also be improved. The BDT only uses 12 features, so it is a relatively simple BDT. A more sophisticated BDT with more physics-based features may improve background rejection while keeping a high signal efficiency. For example, the Visibles BDT uses global features of the activity in the HCal, but the inclusion of a clustering algorithm and features related to candidate clusters may better distinguish signal and background events. Alternatively, although the HCal is not nearly as granular as the ECal, a neural network may be effective for identifying visible A' decays. Other experiments with scintillator-based detectors that are less granular than the LDMX HCal, such as NOvA, have had success with a convolutional neural

network for classifying events [104]. Although it should be noted that, while the detector cell sizes in the NOvA detectors are larger than in the LDMX HCal, NOvA’s detectors are almost fully active and the HCal is a sampling calorimeter with iron absorbers.

5.3 Signal Rates and Sensitivity Estimates

5.3.1 A' Efficiency

The signal efficiency for the analysis cuts summarized in Table 5.3 varies very slightly with A' mass, as can be seen in Figure 5.6. All A' masses are detected with an efficiency of about 30% after all cuts. An efficiency plot after each cut for 100 MeV A' s, the lowest performing mass, is shown in Figure 5.7. The largest source of inefficiency comes from the Visibles BDT, where this single cut is about 50% efficient to the signal (all other cuts combined are 60% efficient). However, this reduces the background at 1×10^{14} EOT from about 31,000 events to less than one event, and it is an optimal choice according to the FOM.

The Visibles BDT is slightly less efficient to higher mass A' s. This is because the showers from higher mass A' s tend to be longer and wider than those from lower mass A' s. This is because higher mass A' s typically carry away a higher fraction of the beam energy, as seen in Figure 5.8, which needs to be dissipated in an electromagnetic shower. To the BDT, because these showers are longer, wider, and more energetic, they look more hadronic—or more background-like—so they are rejected at a slightly higher rate.




Figure 5.6: Final signal efficiency as a function of the A' decay position after all cuts. The shaded region represents the coverage of the HCal.

5.3.2 A' Rates

Reach estimates for the minimal A' model can be constructed from theoretical calculation of signal production and lifetime coupled with signal efficiency and background estimates from this work. Following the convention from [46], the approximate rate of production of A' s at the target $N_{A'}$ as a function of mass, ε , and EOT is

$$N_{A'} \approx 7 \times \left(\frac{\varepsilon}{10^{-5}} \right)^2 \left(\frac{100 \text{ MeV}}{m_{A'}} \right)^2 \left(\frac{\text{EOT}}{10^{16}} \right). \quad (5.3)$$

The product of the proper decay length $c\tau_{A'}$ and a typical A' boost γ as a function of mass and ε is

$$\gamma c\tau_{A'} \approx 65 \text{ cm} \left(\frac{E_{A'}}{8 \text{ GeV}} \right) \times \left(\frac{10^{-5}}{\varepsilon} \right)^2 \left(\frac{100 \text{ MeV}}{m_{A'}} \right)^2. \quad (5.4)$$




Figure 5.7: Signal efficiency as a function of A' decay distance from the target through each step of the selection cuts for the A' mass with the lowest final efficiency (100 MeV). The shaded region represents the coverage of the HCal.

The actual amount of signal N_{sig} to decay in a certain fiducial region, between the ranges of some $z_{\min}$ and $z_{\max}$, can be found by integrating the exponential decay distribution over that range:

$$N_{\text{sig}} = N_{A'} \times \int_{z_{\min}}^{z_{\max}} \frac{e^{-z/\gamma c \tau_{A'}}}{\gamma c \tau_{A'}} dz. \quad (5.5)$$

The integrand is divided by a factor of $\gamma c \tau_{A'}$ in order to normalize the integral from 0 to infinity, such that the integral can be interpreted as the fraction of signal events that decay between $z_{\min}$ and $z_{\max}$. Finally, to get the expected signal, the distribution must be multiplied by the efficiency $\text{eff}(m_{A'}, z)$ of the signal at every z and for every

mass.

$$N_{\text{sig}}(m_{A'}, \epsilon) = N_{A'}(m_{A'}, \epsilon) \times \int_{z_{\min}}^{z_{\max}} \frac{e^{-z/\gamma c \tau_{A'}(m_{A'}, \epsilon)}}{\gamma c \tau_{A'}(m_{A'}, \epsilon)} \text{eff}(m_{A'}, z) dz \quad (5.6)$$

5.3.3 Sensitivity Estimates

For sensitivity estimates, the following simplifying assumptions were made. First, a conservative approximation of $E_{A'}/E_{\text{beam}} = 0.75$ was made, which was the lowest average fraction of energy carried away by the A' masses considered in this study and supported by Figure 5.8. Higher mass A' 's carry away a higher fraction of the beam energy, on average. Second, the A' 's are very forward in the beam direction, as can be seen in Figure 5.9. Therefore the total distance traveled by the A' was approximated as the distance traveled in the z direction. The beginning of the signal window, $z_{\min}$, was 950 mm; the HCal starts at 870 mm from the target and events with non-zero hits in the first layer were vetoed, so the signal region begins 950 mm from the target. Finally, a constant signal efficiency of 28%—the signal efficiency of the lowest performing mass—was used for every combination of A' mass and z , justified by Figure 5.7.

The signal sensitivity contours are shown for two different scenarios: 0.5 background events at 4×10^{14} EOT (expected from approximately one year of operations) and 0.5 background events at 1×10^{16} EOT (an optimistic estimation of the full dataset). The limits were calculated using the CL_s criterion [127, 128]. With an expected background of 0.5 events, the minimum number of expected signal events required to exclude with 90% confidence is 2.5. The contours are drawn for the two scenarios if Equation 5.6 is greater than these minimum expected signal values, checked at every combination of A' mass and ϵ value.




Figure 5.8: The fraction of beam energy carried away by the A' .




Figure 5.9: The angle of the A' momentum with respect to the beam axis.

The final reach projection can be seen in Figure 5.10. For the 4×10^{14} EOT contours, the number of background events comes from Table 5.3. As mentioned at the end of Section 5.3.1, we take the efficiency of the lowest performing mass point in order to be conservative, which gives us a constant efficiency of 40%.




Figure 5.10: The 90% confidence level expected sensitivity for LDMX with 4×10^{14} EoT and 1×10^{16} EoT. Existing leading limits from colliders [60, 123, 129–132], beam dump [62, 133–138], electron $g-2$ measurements [139–142], and fixed target [61] experiments are shown as the gray region.

5.4 Discussion

A full detector simulation for a long-lived visibly decaying signature for minimal A' s with an 8 GeV beam in LDMX has been done. The results show competitive projections with other beam dump experiments, as shown in Figure 5.10, despite several conservative assumptions in the estimation of sensitivity. No unprobed phase space would be reached with this analysis. As mentioned in Section 5.2.5, improvements can be made to the analysis to better reject background and increase signal efficiency, which is currently a limiting factor in the expected sensitivity of this search.

LDMX's expected sensitivity at 1×10^{16} EOT will cover much of the same parameter space as FASER [132]. This will allow LDMX to verify their findings through a different experimental method. While FASER is situated along the proton beam axis 480 m downstream from the ATLAS interaction point, with about 100 m of solid rock for the dark matter and background particles to pass through, LDMX instead has a fixed target, electron beam, and a fully-instrumented detector system that only spans a few meters. The A' s produced for FASER will also be much higher energy than the ones produced in LDMX, so they will be much more heavily boosted. Another consideration is the primary production mechanism for A' s in each experiment; the primary production channel for the LDMX is through dark bremsstrahlung, whereas in FASER there are a more production modes, such as neutral hadron decay. So, while the area of phase space covered is similar between the two experiments, the method of getting there is quite different.

LDMX can further improve its sensitivity to visibly decaying A' signatures, in particular A' s with shorter lifetimes, by using the ECal as a target. In addition, the production of A' s through the decay of neutral mesons would increase the yield of sig-

nal events. This A' production mechanism is not currently implemented in LDMX's simulation tools, so this would need to be developed in order to estimate the added sensitivity from this production mechanism. A Visibles ECal as target search is currently in development. Another method to improve LDMX's sensitivity to shorter lifetimes is through a recoil tracker search for “dark tridents”, which would require a new trigger and a very different analysis approach to what is presented here.

The full set of selection cuts described here assumes that the signal is a visibly decaying A' , which rely on the kinematics of this underlying model. However, the set of features used to train the Visibles BDT are generic for any long-lived particle that decays to particles with electromagnetic final states. Therefore, the Visibles BDT could be an effective tool for other long-lived particle searches in LDMX, such as ALPs.

Part III

Performance Studies of the Mu2e Cosmic Ray Veto

Chapter 6

The Mu2e Experiment

The following two chapters are tangentially related to the physics objectives of this thesis, although closely related to the hadronic calorimeter used by LDMX. These chapters briefly describe the Mu2e experiment, which aims to detect charged lepton flavor violation, and measurements of key performance metrics of Mu2e’s cosmic ray veto modules. Namely, studies of the efficiency, position resolution, and light yield decline are presented.

The Mu2e experiment intends to search for the conversion of a negatively charged muon to an electron in the presence of an aluminum nucleus: $\mu^- \text{Al} \rightarrow e^- \text{Al}$ [143]. In the extended Standard Model with the inclusion of non-zero neutrino masses, the ratio of conversion to capture by the nucleus occurs at an experimentally unobservable rate of $R(\mu^- \text{Al} \rightarrow e^- \text{Al}) \simeq 10^{-52}$ [144, 145]. Any observation of this process would be unambiguous evidence of new physics. The current experimental limit on the muon-to-electron conversion rate in the presence of a nucleus ($\mu^- \text{N} \rightarrow e^- \text{N}$) relative to the capture rate has been set by the SINDRUM-II Collaboration to be $< 7 \times 10^{-13}$ at 90% confidence level [146]. Mu2e’s Run I has been planned to record 6.0×10^{16} stopped muons and have a sensitivity of $< 6.2 \times 10^{-16}$ at 90% confidence level, a three order of magnitude improvement over the current experimental limit [147]. After an approximately three year downtime to upgrade Fermilab’s accelerator complex, Mu2e will resume data collection to improve the sensitivity by another order of magnitude [143].

6.1 The Mu2e Experimental Apparatus

An isometric cutaway drawing of the Mu2e experimental apparatus (excluding the cosmic ray veto) is shown in Figure 6.1. This section will briefly outline how Mu2e will search for the muon-to-electron conversion process.




Figure 6.1: Isometric drawing of the Mu2e experimental apparatus. The magnetic field strength of each of the solenoids is shown in blue circles. The detectors (tracker and calorimeter) are both located in the detector solenoid. The cosmic ray veto, which surrounds the apparatus, is not shown. From Ref. [147].

6.1.1 Muon beamline

Mu2e will receive 8 GeV protons from the Delivery Ring, which stores proton bunches from the Recycler before they are accelerated to 120 GeV in the Main Injector (see Section 2.1 for a detailed description of Fermilab’s accelerator complex and how proton bunches are produced). The proton beam from the Delivery Ring has a pulsed timing structure, with 250 ns-wide proton pulses appearing every 1695 ns [147]. The beam will enter the production solenoid, a high-field magnet with a graded solenoidal field that varies smoothly from 4.6 T to 2.5 T [143], and interacts with a 1.6 interaction length tungsten target. The field is at its maximum strength downstream of the target in order to reflect back a portion of the charged particles produced in the target. Primarily, pions are produced in the proton interactions with the target, which

then decay to muons.

Charged particles exiting the production solenoid are directed through the S-shaped transport solenoid, which also has a graded magnetic field from 2.5 T to 2 T, toward the detector solenoid. The S-shape of the transport solenoid is designed to mitigate neutral particles, such as neutrons or photons, from entering the detector solenoid and producing signal-like backgrounds. Collimators at the entrance, center, and exit of the transport solenoid are used to filter for muons with momenta with less than 100 MeV/c, which is needed for capture by the Al nucleus. The magnetic field of the transport solenoid causes particles of opposite charge signs to vertically drift in opposite directions, allowing a vertically offset opening in the collimator at the center of the transport solenoid to select the negatively¹ charged particles to be delivered to the detector solenoid [143].

6.1.2 The Detectors

Muons passing through the transport solenoid are delivered to the detector solenoid, which has a graded magnetic field from 2 T to 1 T in the upstream region of the solenoid, and a uniform magnetic field of 1 T in the downstream region. The average muon momentum of muons delivered to the detector solenoid is 50 MeV/c. About one third of these muons are stopped in an aluminum target, which is placed in the graded magnetic field of the detector solenoid. Stopped muons are captured by the Al atoms and rapidly cascade to the 1s orbital state, where they may undergo muon-to-electron conversion. Signal conversion electrons will have an energy close to 105 MeV, which is reconstructed in a straw tube tracker and confirmed in an electromagnetic calorime-

¹The opening is configurable to select for either positively or negatively charged particles, although the positively charged particles are to be used for detector calibration data instead of physics data.

ter.² Accounting for detector resolution of this monoenergetic signal and energy losses in the target and detector material, Mu2e’s expected signal is a reconstructed electron with a measured momentum between 103.6 and 104.9 MeV/c, shown in Figure 6.2.




Figure 6.2: The signal conversion electron and background electron momenta (left) and time after the proton beam pulse (right) distributions. The conversion electron (CE) signal distributions correspond to a conversion rate of $R(\mu^- \text{Al} \rightarrow e^- \text{Al}) = 1.0 \times 10^{-15}$. Note that the cosmic-ray-induced background rate in this plot excludes events vetoed by the cosmic ray veto. From Ref. [147].

Mu2e uses a low-density, low-mass straw tracker downstream of the stopping target to reconstruct the electron’s momentum. The tracker is 3 m long and consists of 18 stations, which are made of rows of 5 mm diameter straw tubes filled with a mixture consisting of 80% Ar and 20% CO₂. It is positioned in the uniform field of the detector solenoid. The expected momentum resolution of the tracker is 300 keV/c [143]. Finally, downstream of the tracker, is an undoped CsI electromagnetic calorimeter, which is used for triggering and particle identification when combined with tracker information.

The largest sources of background expected in Mu2e’s are shown in Table 6.1. For Mu2e to achieve its sensitivity goals, it is vital to keep backgrounds to less than one event over the duration of the experiment. An enormous simulation effort has

²This is a monoenergetic signal calculated as $E_e = m_\mu - m_e - E_b - E_{\text{rec}}$, where $m_\mu - m_e$ is the mass difference between the muon and electron, $E_b \approx Z^2 \alpha^2 m_\mu / 2$ is the binding energy of the muonic atom, and $E_{\text{rec}} \approx m_\mu^2 / (2m_{\text{Al}})$ is the nuclear recoil energy [148].

been made to estimate the various backgrounds, and Mu2e’s Run I is expected to have 0.105 ± 0.032 background events [147], and the total expected background after Run I and Run II is 0.41 ± 0.03 events [149]. The backgrounds come in three types: (1) stopped-muon-induced backgrounds; (2) beam-related backgrounds; and (3) time-dependent backgrounds. The largest source of background originates from cosmic rays. In order to mitigate this background, Mu2e will use a cosmic ray veto surrounding the detector solenoid, described in detail below.

Table 6.1: Expected backgrounds in Mu2e Run I. Only major backgrounds are listed. From Ref. [147].

Background source	Expected events
Cosmic rays	0.046 ± 0.010 (stat) ± 0.009 (syst)
Decay in orbit	0.038 ± 0.002 (stat) $^{+0.025}_{-0.015}$ (syst)
Antiprotons	0.010 ± 0.003 (stat) ± 0.010 (syst)
Radiative pion capture	0.010 ± 0.002 (stat) $^{+0.001}_{-0.003}$ (syst)
Total	0.105 ± 0.032

6.2 The Cosmic Ray Veto

Cosmic rays can produce conversion-like electrons from muon decay-in-flight, the production of delta electrons, neutron interactions, and muons misidentified as conversion electrons. Such events are expected to appear at a rate of roughly one per day and represent the largest background to the experiment. To identify and veto such events, a cosmic ray veto (CRV) detector has been fabricated that will surround much of the Mu2e apparatus, shown in Figure 6.3. The CRV is designed to veto cosmic ray muons with approximately an overall 99.99% efficiency, reducing the cosmic-ray background to less than one event over the full experimental run.




Figure 6.3: The cosmic ray veto (gray) consists of 83 modules that are mounted on the top and sides of concrete shielding blocks (red) that surround the entire Mu2e detector solenoid and half of the transport solenoid.

6.2.1 The Cosmic Ray Veto Detector Design

The CRV consists of four layers of scintillator extrusions, each extrusion with two embedded wavelength-shifting fibers (which together make a “counter”), read out by silicon photomultipliers (SiPMs). A conversion-like electron in the tracker that comes in coincidence with a track stub found in the CRV is vetoed in the offline analysis.

The CRV scintillator counters are grouped in modules, each module (with one exception) consisting of 64 counters, 16 per layer, with aluminum plates of 9.525 mm (3/8”) thickness placed between each layer. The entire CRV consists of 83 total modules. Each scintillator extrusion has a cross section of $51.34 \times 19.80 \text{ mm}^2$, with lengths ranging from 0.985 m to 6.900 m. The extrusions are made of polystyrene that has been doped with 1% PPO and 0.03% POPOP for wavelength-shifting and improved scintillation light yield. The extrusions have a thin TiO_2 reflective coating and two cavities into which the wavelength-shifting (WLS) fibers are inserted [150]. Two di-

ameters of WLS fiber are used in the modules; 1.4 mm fiber is used in 66 modules and 1.8 mm fiber is used in 17 modules placed in the most critical region to detect cosmic-ray muons. A picture of an extrusion is shown in Figure 6.4.




Figure 6.4: A polystyrene scintillator extrusion with two cavities for wavelength-shifting fibers, coated with a TiO_2 -polystyrene mixture.

Light collected by the fibers is read out by $2 \times 2 \text{ mm}^2$ silicon photomultipliers³ (SiPMs) located at each end, except for a select few modules that are only read out on one end. For readout purposes, two counters are glued together side-by-side to form a di-counter, shown in Figure 6.5. An acetal fiber guide bar (FGB) is attached to each end of a di-counter, into which the fibers are glued, after which their ends are polished. Note that the fibers are only glued to the fiber guide bar, not to the extruded cavity of the extrusion. Each SiPM is mounted on a tiny printed circuit board by the manufacturer [151]. The SiPMs are housed in small rectangular wells in a black anodized aluminum fixture called the SiPM mounting block, which is precisely registered to the FGB by a pair of sleeves. A small circuit board, the counter motherboard (CMB), has spring-loaded pogo pins that gently press the SiPMs against the fibers and provide electrical contact to the SiPM carrier boards. The CMB provides the SiPM bias voltage and sends the four SiPM output signals of each di-

³The SiPMs are Hamamatsu S13360-2050VE devices with a $50 \mu\text{m}$ pixel size, the large size being chosen as a compromise between dynamic range and quantum efficiency [151].

counter end to a readout board via an HDMI cable. The CMB also has two flasher LEDs, one for each counter, and a single thermometer⁴ centered on the board. This thermometer’s temperature measurement serves as a proxy for the temperature of the four SiPMs serviced by the CMB.




Figure 6.5: Exploded view of the manifold at the end of a di-counter showing the fiber guide bar, SiPM mounting block, SiPM carrier boards, SiPMs, and counter motherboard. The flasher LEDs, thermometer, and pogo pins are not shown.

The readout manifold, attached to each di-counter end, was designed to be as compact as possible to minimize gaps in the coverage of the CRV that could allow cosmic-ray muons to leak through undetected. Cooling the SiPMs to control their temperature—and hence breakdown voltage—as well as to reduce their dark count rates, was considered when designing the CRV, but rejected due to cost and the strict inter-counter gap requirements to maintain the CRV’s high efficiency. Instead, it was decided to carefully monitor the SiPM temperatures to allow for temperature corrections to be made to accurately determine the photoelectron (PE) yields.

⁴The thermometer is a UMW DS18B20U digital temperature sensor with a range from -55°C to 125°C and an accuracy of $\pm 0.4^{\circ}\text{C}$.

All of the 83 modules that comprise the CRV were fabricated at the University of Virginia and are now at Fermilab in the Wideband detector hall awaiting installation [118]. The tests reported in Chapter 7 come from several modules situated in a cosmic-ray test stand at Wideband.

6.2.2 Readout of the Cosmic Ray Veto

The architecture of the CRV readout is shown in Figure 6.6. Front-end boards (FEBs) provide the bias to the SiPMs and serve to take the signals from the SiPMs, amplify and shape them, digitize them in amplitude and time, zero-suppress them, and finally, store them for later readout. Each FEB serves 64 SiPMs via 16 HDMI cables. Two FEBs are required to read out each end of a CRV module.




Figure 6.6: Architecture of the CRV readout for a “normal” module read out on both ends.

The FEBs receive power through CAT6 cables from a readout controller (ROC), which also supplies the clock timing information. In the normal readout of the CRV, upon receiving a trigger, each ROC accumulates data from up to 24 FEBs and sends

it to the Data Transfer Controllers (DTCs), which serve as the interface between the data acquisition hardware and the online computer farm. For the work described here, the DTCs were not used and the ROCs were only used to provide power, timing information, and triggers. The FEB data were read out via CAT6 cables, through a switch, to a DAQ PC. The PC was used to set the FEB amplifier gain and SiPM bias voltages, as well as to retrieve the CMB temperatures and SiPM waveforms.

Each FEB has eight AFE chips that serve to amplify, shape, and digitize the signals with a 12-bit ADC from eight SiPMs every 12.5 ns [152]. The digitized signals from the eight chips are sent to four FPGAs (each FPGA handling 16 channels). The FPGAs perform a set of functions — such as zero-suppression — that were not used in this work. The data from the FPGAs are sent to a buffer memory where they await a signal to be read out.

Each FEB has a temperature sensor that allows the temperature dependence of the bias it provides to the SiPMs to be determined. The biases are set by an overall bus bias common to all 64 SiPM channels served by one FEB, along with trim settings that allow each SiPM's bias to be set to a precision of 2 mV.

The data acquisition for the results discussed here consisted of a 240 s, two-phase “cycle”: a 220 s live-time period where triggering and data collection occurred, followed by a 20 second period for data transfer to the DAQ PC for storage. Before the onset of each cycle, slow control information was retrieved from the FEBs and written to the raw data file. This information included the settings of the AFEs, bias settings for each SiPM, CMB and FEB temperature, and the AFE gain settings. During the live-time in the cycle, the digitized SiPM-generated data were written to the buffer memory on the FEB every 12.5 ns. Upon receiving a trigger, data before and after the trigger time were recorded. Typically 750 ns of data were recorded before the trigger

(the “pre-signal” region—see Figure 7.3), allowing the pedestal to be determined as well as occasional dark counts from the SiPM noise; the latter was used in the calibration process. For each trigger, a total of 127 ADC samples were recorded for each of the 64 FEB channels. This sequence of data initiated by a single trigger is called an event. For the work described here, the cosmic-ray trigger rate was approximately 3 Hz and the LED flasher trigger rate was 1 kHz.

Chapter 7

Performance studies of the Mu2e CRV

In this chapter, an analysis is presented of three performance metrics from a subset of the CRV modules with cosmic-ray data: the efficiency, the position resolution, and the light-yield decline. First, a description of the experimental setup at the Wideband hall at Fermilab is described. Next, a procedure for calibrating the SiPMs that accounts for their temperature dependence is presented, which is critical for accurate measurement of the light yield decline. Finally, an analysis to determine the efficiency, position resolution, and light-yield decline is presented.

7.1 Cosmic Ray Test Stand

At the Wideband hall at Fermilab, a subset of the modules were stacked for testing, as shown in Figure 7.1. In the studies presented here, three separate stacks of modules were used. For efficiency studies in Section 7.3 and position reconstruction measurements in Section 7.4, a stack of four 6 m-long modules outfitted with 1.8 mm diameter WLS fibers were used, which will henceforth be referred to as CRV-T modules.¹ For aging studies in Section 7.5, a separate stack of four CRV-T modules and a stack of four 4 m-long modules with 1.4 mm diameter WLS fibers (referred to as

¹For efficiency studies, module 169 was tested and modules 171, 172, and 173 were the other modules in the stack. For position resolution studies, module 171 was tested.

CRV-L modules) were used². The modules used for these studies are representative of all CRV production modules. Due to limited availability of production FEBs, these studies were not done in parallel with each other. Instead, FEBs were connected to one of these three module stacks for 1–4 consecutive months for data collection before moving the FEBs to another stack.




Figure 7.1: Top: The cosmic test stand at the Wideband hall at Fermilab. Bottom: Cartoon of the end-view of a module stack, showing the position of the trigger paddles.

²Module 127 was the CRV-L module tested in aging studies. Modules 163, 165, 168 were the CRV-T modules used for aging studies.

Three plastic scintillator trigger paddles read out with photomultiplier tubes (PMTs) were placed 1.0 m from the modules' readout end, perpendicular to the direction of di-counters. To collect cosmic-muon events in the cosmic-ray test stand, a three-fold coincidence among the trigger paddles was required within a 100 ns window. The coincident trigger signals were distributed to the FEBs through a ROC.

7.1.1 Trigger Paddles

A diagram of a trigger paddle, fabricated at the University of Virginia, used in the cosmic-ray test stand is shown in Figure 7.2. The trigger paddles are 100 cm long with a cross sectional area of 10 cm×2 cm and are made of EJ-200 plastic scintillator. A PMMA fishtail light guide connects the paddles to Hamamatsu R329-02 photomultiplier tubes. A 2 mm thick EJ-560 optical interface pad is situated between the light guide and photomultiplier tube. To boost the reflectivity of scintillation light produced, the paddles were wrapped in Tyvek. Then, to keep the paddles light-tight, they were wrapped in Tedlar.




Figure 7.2: A diagram of a trigger paddle used in the cosmic-ray test stand.

PMTs, like the APDs and SiPMs described previously, use the photoelectric effect to amplify single photons into measurable electrical signals. The bias voltage applied to the PMTs was optimized to maintain a high muon detection efficiency ($\sim 99\%$) and a low noise rate (~ 22 Hz). The output of each PMT went to a LeCroy 623B octal discriminator, which converted any < -25 mV electrical pulse from the PMT to

a 60 ns wide square pulse. The three discriminator channels from the trigger paddles went to a logic unit that performed a three-fold coincidence in the paddles within a 100 ns time window.

The coincident trigger signals were distributed to the FEBs through the ROC. Upon receiving a trigger, a total of 127 samples in each FEB channel was read out at 80 MHz both before and after the signal that caused the trigger. An example of the signal traces is shown in Figure 7.3 [153]. Typically, 750 ns of data before the signal pulse (referred to as the pre-signal region) were read out. The pre-signal region contained occasional dark pulses caused by SiPM thermal noise, which were used to calibrate the signal pulse area to the number of PEs. The data processing scheme is discussed in the following section.




Figure 7.3: Example ADC waveform from the two SiPM channels of one end of a CRV counter, 127 ADC samples each. The pre-signal region shows dark-count pulses of 1 PE in both channels at 400 ns and 600 ns. The signal region shows signal pulses with 24 PEs and 30 PEs in both channels. All pulses get fitted with a Gumbel function.

7.2 Data Processing

The signals from the cosmic-ray test stand were calibrated and reconstructed. As the data were taken over several years with different temperatures, and because the electronic response of a SiPM is temperature-dependent, appropriate temperature corrections were applied to ensure that the results were independent of temperature effects. A detailed discussion of the calibration and reconstruction procedure is described below.

7.2.1 Calibration

7.2.1.1 Calibration at a Stable Temperature

In the pre-signal region, the data in each channel underwent a scan for dark pulses. An example dark pulse is shown in Figure 7.3. The pulses underwent pedestal subtraction and were fitted using a Gumbel function [154]:

$$f(t) = A e^{\frac{t-\mu}{\beta}} - e^{\frac{t-\mu}{\beta}}, \quad (7.1)$$

where μ is the time of the pulse peak, A/e is the pulse height, and $A \cdot \beta$ is the pulse area. The Gumbel function was selected empirically based on the quality of pulse fitting.

The fitting process determined the dark count pulse areas, which was used rather than the pulse height to determine PE yields, the distribution of which can be seen in Figure 7.4, showing distinct peaks corresponding to the pulse areas of 1 PE, 2 PEs, 3 PEs. The single photoelectron (SPE) peak in the distribution of Figure 7.4

is used to produce the SPE calibration constant that translates the SiPM waveform area into the reconstructed number of photoelectrons. Due to the low noise and low cross-talk rate—and hence limited statistics for the photoelectron peaks greater than 1 PE—only the SPE peak is utilized to establish the SPE calibration constant.




Figure 7.4: The pedestal-subtracted dark count distribution from a single SiPM displays distinct peaks that correspond to one, two, and three photoelectrons. The location of the first peak (the single photoelectron peak, SPE) is utilized to extract the SPE calibration constant.

The calibration procedure described above works well for stable temperature environments and has been used in test beam experiments to study the CRV performance [155]. However, if the temperature is not stable, such as in measurements made several months apart when seasonal variations affect the ambient temperature, the variations in these calibration constants is too large to ignore. A temperature-dependent calibration procedure for normalizing the CRV response to a reference temperature is described below. This temperature-dependent calibration procedure was used to make measurements of the aging of the CRV counters, described in Section 7.5, but will also be implemented in the operation of the Mu2e experiment.

7.2.1.2 Temperature-Dependent Calibration

The SPE calibration constants and photo-detection efficiencies are impacted by the temperatures of both the SiPMs (as measured by the thermometer on the CMBs) and the FEB (T_{SiPM} and T_{FEB}), the former because the SiPMs have temperature dependent breakdown voltage, and the latter because of the bias provided by the FEBs to the SiPMs is slightly temperature dependent and therefore does not retain its set value. In addition, the gain of the AFE amplifier/digitizers have a temperature dependence. Consequently, temperature corrections are essential to adjust the SPE constants and photo-detection efficiencies to reference temperatures, chosen to be: $T_{\text{SiPM}}^{\text{Ref}} = 20^\circ\text{C}$ for the SiPM and $T_{\text{FEB}}^{\text{Ref}} = 40^\circ\text{C}$ for the FEB (roughly, their normal operating temperatures). This adjustment follows a specific procedure, as outlined in Figure 7.5 and described in detail below.




Figure 7.5: A flow chart of the calibration method to equalize the CRV response as a function of temperature. The calibration, measurement, and rescale steps are described in detail in the text. Note that it takes many dark pulses to derive the SPE calibration constant. Each dark pulse must have its response, taken at temperature, T_{Dark} , corrected to the reference temperature, T_{Ref} . Similarly, A_{Meas} is a signal pulse that can be taken at any temperature, T_{Meas} . Before the SPE calibration constant can be applied it must be corrected from the reference value, C_{Ref} , to the relevant value at the measured temperature, C_{Meas} .

In Step 1, described in Section 7.2.1.3, dark pulses are collected, fitted, and the pulse areas are extracted, A_{Dark} . The SiPMs are not noisy, so the collection of sufficient numbers of dark pulses for a precise measurement of the SPE calibration constant might occur on a time scale in which the temperature is not constant. Hence, the dark pulses must first be corrected to the reference temperature before the SPE calibration constant, C_{Ref} , is extracted from the single photoelectron peak. In Step 2, described in Section 7.2.2, the calibrations are applied to a measured signal pulse area, A_{Meas} , at its measured temperature, T^{Meas} . C_{Ref} must first be corrected to the proper value at the measurement temperature, C_{Meas} . It is then applied to the measured pulse area, A_{Meas} , to estimate the number of photoelectrons. To compare the performance of the CRV over time, Step 3 of Figure 7.5, also described in Section 7.2.2, is applied to rescale the measured PE value to the reference temperature. Using this calibration procedure, CRV performance over time can be compared even if the temperature is not stable, and effects independent of temperature can be determined.

Table 7.1: Temperature correction parameters extracted from the temperature-dependence studies described in Ref. [153]. The correction parameters p_3 , p_4 , and p_5 are valid for a particular gain setting and a nominal over-voltage of 2.5 V.

	Parameter	Value
p_1	Temperature coefficient of the SiPM breakdown voltage	55.4 ± 0.8 mV/K
p_2	Temperature coefficient of the SiPM bias provided by the FEB	4.1 ± 0.7 mV/K
p_3	SPE calibration constant (or pulse area) increase per SiPM bias increase	125 ± 3 ADC·ns/V
p_4	SPE calibration constant (or pulse area) decrease due to AFE temperature dependence	-1.46 ± 0.16 ADC·ns/K
p_5	PE increase due to overvoltage bias increase	$21.8 \pm 2.1\%$ /V

The CRV SiPMs are operated at constant bias voltages, corresponding to about 2.5 V over the SiPM's breakdown voltage at the reference temperature. To determine this, a bias scan procedure was developed to extract the breakdown voltages for each SiPM.

These breakdown voltages depend on the SiPM temperatures (p_1 in Table 7.1). In addition, temperature variations result in relatively small fluctuations in the FEB bias supplied to the SiPMs (p_2 in Table 7.1). Changes in the over-voltage due to temperature fluctuations, in turn, lead to corrections of the SPE constants (p_3 in Table 7.1). The SPE constants are also influenced by a gain change in the AFE (p_4 in Table 7.1)³. Finally, changes to the over-voltage also require corrections to the light yield (p_5 in Table 7.1). A detailed description of the measurement of these parameters can be found in Ref. [153].

7.2.1.3 Obtaining the Reference Calibration

Each pulse collected during the dark-count scan undergoes an adjustment to its measured pulse area A_{Dark} related to its values at the reference temperatures $T_{\text{SiPM}}^{\text{Ref}}$ and $T_{\text{FEB}}^{\text{Ref}}$. For each pulse, the deviation of the over-voltage ΔV_{ov} due to the temperature difference is calculated as

$$\Delta V_{\text{ov}} = -p_1 \cdot (T_{\text{SiPM}}^{\text{Dark}} - T_{\text{SiPM}}^{\text{Ref}}) + p_2 \cdot (T_{\text{FEB}}^{\text{Dark}} - T_{\text{FEB}}^{\text{Ref}}). \quad (7.2)$$

This deviation of the over-voltage is used to calculate what the pulse area at the reference temperatures of the CMB and FEB would have been. The gain change due to the FEB temperature is an additional adjustment:

$$A_{\text{Ref}} = A_{\text{Dark}} - p_3 \cdot \Delta V_{\text{ov}} - p_4 \cdot (T_{\text{FEB}}^{\text{Dark}} - T_{\text{FEB}}^{\text{Ref}}). \quad (7.3)$$

³Parameters p_3 and p_4 depend on the AFE gain settings, and the values in the table are only valid for the specific experimental setup.

These temperature-corrected pulse areas A_{Ref} are then binned and the temperature-corrected SPE constants, C_{Ref} , are determined using the same approach as described in Sec. 7.2.1.1 and shown in Figure 7.4. The parameters in Eq. (7.3) were obtained for temperature corrections of the SPE calibration constant, which is the pulse area of a single photoelectron. Therefore, this equation is only valid for pulse areas of single PEs, and only the single PE peak of the temperature-corrected calibration plot was used to determine the SPE calibration constant.

7.2.2 Reconstructing the Light Yield

SiPM signal pulses were pedestal-subtracted and fitted with a Gumbel function in the same way as the dark counts shown in Eq. (7.1). Their pulse areas, peak times, leading-edge times, and width parameters were recorded.

Similar to the SPE calibration constants C_{Ref} , the measured pulse required temperature corrections. The process began by obtaining the SPE calibration constant for the temperatures at the time the pulse was recorded, C_{Meas} . Since the SPE calibration constant C_{Ref} was initially stored for the reference temperatures $T_{\text{SiPM}}^{\text{Ref}}$ and $T_{\text{FEB}}^{\text{Ref}}$, the temperature correction procedure outlined in the previous section had to be reversed to derive the SPE calibration constant specific to the SiPM and FEB temperatures of the given pulse $T_{\text{SiPM}}^{\text{Meas}}$ and $T_{\text{FEB}}^{\text{Meas}}$. As before, the first step was calculating the deviation of the over-voltage ΔV_{ov} due to the temperatures at the SiPM and FEB (see Eq. (7.2)):

$$\Delta V_{\text{ov}} = -p_1 \cdot (T_{\text{SiPM}}^{\text{Meas}} - T_{\text{SiPM}}^{\text{Ref}}) + p_2 \cdot (T_{\text{FEB}}^{\text{Meas}} - T_{\text{FEB}}^{\text{Ref}}). \quad (7.4)$$

The second step was adjusting the calibration SPE constant from its value at the reference temperatures (C_{Ref}) to a calibration constant at the actual temperatures when

the pulse was recorded (C_{Meas}). This was done using the over-voltage change from Eq. (7.4) and the gain change due to the FEB temperature was an additional adjustment (analogous to Eq. (7.3) above but adjusting the calibration constant instead of the pulse area):

$$C_{\text{Meas}} = C_{\text{Ref}} + p_3 \cdot \Delta V_{\text{ov}} + p_4 \cdot (T_{\text{FEB}}^{\text{Meas}} - T_{\text{FEB}}^{\text{Ref}}). \quad (7.5)$$

Dividing the pulse area A by this calibration constant C_{Meas} resulted in the number of photoelectrons of this pulse:

$$\text{PE}_{\text{Meas}} = A_{\text{Meas}} / C_{\text{Meas}}. \quad (7.6)$$

The change in over-voltage affected the photon detection efficiency. To obtain the PE yield at the reference temperature, PE_{Ref} , the effect of the over-voltage change had to be removed in a second step involving p_5 from Table 7.1 and ΔV_{ov} from Eq. (7.4):

$$\text{PE}_{\text{Ref}} = \frac{\text{PE}_{\text{Meas}}}{1 + p_5 \cdot \Delta V_{\text{ov}}}. \quad (7.7)$$

The temperature-corrected number of photoelectrons, PE_{Ref} , were accumulated in a histogram (an example is shown in Figure 7.6). The resultant PE distribution was fitted with a convoluted Gauss-Landau function for cosmic-ray data and a Poisson function for LED data. These fits determines the most probable value of the PE distribution and were a useful figure of merit to characterize the CRV performance.

7.3 Efficiency Studies

In most areas covered by the CRV, a cosmic-ray muon will be identified by a coincidence in space and time of three of the four layers of counters. To keep the background




Figure 7.6: Typical SiPM photoelectron distribution example from cosmic-ray muons, fitted with a convoluted Gauss-Landau function.

levels in Mu2e below one event during the experiment’s run time, the required detection efficiency for these muons must be as high as 99.99% in some critical regions of the CRV. In this section, the measurement of the efficiency of CRV modules is described.

7.3.1 Event Selection

To measure the efficiency to the required level, a very pure sample of cosmic-ray muons is needed. In addition to the trigger requirement, three of the four modules (nicknamed “defining modules” for this analysis) in the stack were used in the offline analysis to define a good cosmic-ray muon. The efficiency of the fourth module (nicknamed the “test module”) in the stack was measured. Three selection cuts were applied to data in the defining modules to isolate likely muon events:

1. At least 10 PEs had to be deposited in each of the twelve layers of the defining modules,
2. a linear fit to the hits in the defining modules had to have a $\chi^2/\text{n.d.f.} < 2$, and
3. there had to be less than 3000 PEs in all modules.

In this analysis, a “hit” was defined as an energy deposition of >5 PEs read out by a SiPM.⁴ These criteria were based on the fact that the muons are minimum ionizing particles, and hence they leave straight tracks through the stack with an energy deposition proportional to their path length through the modules. Events originating from electromagnetic or hadronic showers were more likely to leave many more hits that do not present as straight tracks.

Using a Monte Carlo simulation of the configuration at Wideband, described in more detail in Section 7.3.3, the efficiency of these cuts for selecting muons was verified. The initial sample of triggered Monte Carlo events, generated with the CRY cosmic ray event generator [156], consisted of 86.7% muons. After the event selection cuts described above, this muon sample purity was improved to 99.1%. Figure 7.7 shows so-called $N - 1$ plots⁵ for the three cuts used to select muons, and how those cuts affect muons and non-muons in the Monte Carlo sample. The improved purity of muons after selection cuts is critical for understanding the modules’ veto efficiency to the particles that create a challenging background for Mu2e’s experimental run.

⁴5 PEs is the minimum requirement for a hit because pulses below 5 PEs are suppressed to zero.

⁵These are plots of the parameters used for event selection with all other cuts applied except the one being plotted.




Figure 7.7: Distributions of simulated muon and non-muon events for the three parameters used in event selection cuts. These are $N - 1$ cuts, meaning all other cuts are applied in the panel besides the one being plotted. The red vertical line shows where the cut is applied and the gray region is excluded by the selection cut.

7.3.2 Measurement of the Track Stub Efficiency

From approximately two months of data collection, 3.4×10^6 events passed the selection cuts described above. After event selection, the efficiency of a module was measured as a function of the PE threshold. The 3/4 efficiency of a module was determined in two ways: (1) measuring the single-layer efficiency of the module and extrapolating this efficiency to a 3/4 layer efficiency using combinatorics, and (2) a direct measurement of the 3/4 layer coincidence efficiency.

The single-layer inefficiency as a function of the PE threshold on the sum of all PEs

measured by a single end in the layer is shown in Figure 7.8. The efficiencies of each layer are quite similar. Each layer is functionally identical, so this is expected. Since the efficiencies of each layer are more-or-less the same, the average value was used when extrapolating to a 3/4 efficiency instead of the efficiency of each individual layer. Assuming each single-layer efficiency ϵ_{SL} is independent, the efficiency $\epsilon_{3/4}$ of at least 3 layers being hit is

$$\epsilon_{3/4} = 4\epsilon_{\text{SL}}^3 - 3\epsilon_{\text{SL}}^4. \quad (7.8)$$




Figure 7.8: The single layer inefficiency of a CRV-T module as a function of the PE threshold. Note how each layer has a very similar efficiency.

This extrapolation of the single-layer efficiency is compared to the measured 3/4 layer coincidence efficiency in Figure 7.9. The two methods of measuring the efficiency agree for high PE thresholds, but diverge significantly for lower PE thresholds, with the extrapolation underestimating the efficiency.

The reason for the low PE divergence is as follows. The efficiency of a counter




Figure 7.9: Inefficiency of hit coincidence in 3/4 layers of a module from direct measurements and the extrapolated single layer efficiency using Eq. (7.8).

producing a hit is driven by the number of PEs produced in the counter, which is proportional to the path length of the muon through the counter. At high thresholds, for a hit to be registered the muon had to have gone through the counter and not a gap between counters. This inefficiency is layer-independent. However, as the PE threshold is lowered, it is much more likely that enough PEs will be produced to meet the threshold— even for minimal path lengths—and the inefficiency is driven by the small gaps between di-counters. Because of staggering in the module layers, a muon is much less likely to pass through a second gap if it has already passed through one, as illustrated in Figure 7.10. The staggering of the module layers (42 mm) was determined by optimizing the 3/4 efficiency from extensive Monte Carlo studies of the response of top and side modules to cosmic-ray muons. Hence a low efficiency in one layer is typically correlated with a high efficiency in the other layers. Equa-

tion (7.8) assumes that each layer's efficiency is independent of the others, which is not true. Therefore, Eq. (7.8) does not accurately reflect the $3/4$ efficiency for a low PE threshold.




Figure 7.10: A cartoon illustrating two muon tracks passing through a module. The vertical red muon track goes through the external gap (exaggerated in width) between two di-counters in the second layer. However, due to the staggering of the layers, it is less likely to pass through another gap in a different layer. The orange muon track must be well inclined to pass through the external gap in all four layers, and hence goes through a significant amount of scintillator on the counters on either side of the gap. The staggered position of gaps is drawn to scale, but the size of the gaps is exaggerated for effect.

The measured $3/4$ efficiency and the extrapolation start to diverge at a threshold of ~ 75 PEs, with the measured $3/4$ efficiency being more efficient. Below a 50 PE threshold, the measured inefficiency plateaus. This plateau is likely due to non-muon events that pass selection cuts, which is supported by Monte Carlo studies, described as follows.

7.3.3 Monte Carlo Studies

A Monte Carlo simulation was used to study the module efficiency. In particular, it was used to understand whether the plateau in inefficiency at low PE thresholds observed in Figure 7.9 is due to muons or some other incident particle.

The generation of cosmic rays was performed using the CRY event generator [156]. The simulation of the CRV modules was based on the Geant4 library [90–92], which handles the particle transport through the modules and all particle interactions, including the production of photons due to the scintillation process and Cherenkov effect in the scintillator counters. It also tracks these photons through the scintillator and the embedded wavelength-shifting fiber all the way to the SiPMs and simulates all processes such as wavelength shifting, refraction, reflection, and absorption. However, tracking each individual photon through all counters requires an unreasonably high computation time, which is not feasible for simulations with a large number of events. Therefore, lookup tables were produced (using the Geant4) that store arrival probabilities and time delay distributions for photons arriving at the SiPMs based on the location where the original particle went through the counter and the amount of energy that was deposited there [157]. Based on this information, a SiPM response with an appropriate pulse shape was simulated. This step takes the wavelength-depending photon-detection efficiency of the SiPM pixels into consideration as well as the typical geometrical distribution of the photons at the end of the fiber. The simulation of the SiPM response models the pixel discharge/recharge behavior, saturation effects, cross talk, and dark counts. The final SiPM pulse shape is the sum of the individual pulses of all SiPM pixels, where the single-pixel pulse shape is determined by an independent simulation based on the properties of the SiPM. The last step of the simulation uses the responses of the SiPMs of all counters and generates a digitized waveform similar to the waveforms produced by the experiment. The simulation was tuned to match the results of test-beam studies and studies with cosmic rays [155].

The summed light yield in each module layer was used to validate the model. This

serves as a “figure of merit” for the accuracy of the model to the real configuration. Figure 7.11 shows the light yield in the model compared to data. The simulation agrees very well with data, except in instances of low PE depositions (<80 PEs). For the purpose of this analysis it is not critical that the light yield perfectly match.




Figure 7.11: The summed light yield from all SiPMs in a module layer (in PEs) for Monte Carlo and data, with the ratio data/MC shown in the bottom panel. The distributions have been normalized to have an area of one.

A total of 5.6×10^7 simulated events were generated and passed through the selection cuts described in Section 7.3.1. The measured 3/4 inefficiency and extrapolation from the single-layer efficiency of the MC sample, with a comparison to data, is shown in Figure 7.12. The simulation reproduces the module efficiency observed in the data, including the plateau at low PE thresholds.

Although the Monte Carlo simulation is able to reproduce the plateau in efficiency observed in data, there are still large discrepancies between the inefficiencies measured in the simulation and data. A source of discrepancy most likely comes from how




Figure 7.12: The directly measured 3/4 inefficiency and the extrapolated 3/4 inefficiency from the single layer efficiency for MC and data.

the detector hall was modeled in simulation, which was found to have non-negligible effects on the MC results. The simulated building mimics Wideband, but is not an exact match. Another potential source of error in the modeling is discrepancies between the distribution of cosmic-ray events in simulation and reality. It is not guaranteed that CRY accurately reproduces cosmic showers in the exact environment at Fermilab. Detailed studies by MicroBooNE comparing CRY to other cosmic-ray generators like CORSIKA show that the expected event rates are different between generators, particularly for neutrons and protons [158]. Nevertheless, inefficiency and light yield discrepancies between this Monte Carlo simulation and data are $\mathcal{O}(10^{-3})$ and below. As stated above, the selection cuts are expected to only produce a 99% purity sample of muons, so the impurities of other cosmic-ray particles are on the level of 10^{-2} , within the discrepancy between data and Monte Carlo. For this analysis, it is not critical to perfectly model the module efficiency in simulation; instead, the reproduc-

tion of the plateau in efficiency at low PE thresholds is the essential component of these simulation studies.

Upon inspection of events in the plateau of the measured MC efficiency, it was found that all of the inefficient events originate from incident particles that are not muons. Figure 7.13 shows three examples of inefficient simulated non-muon events at low PE thresholds with hit patterns similar to inefficient events observed in data. Inefficient events such as these occur in the test stand at Wideband, but they are unlikely to be a source of inefficiency for the CRV during Mu2e operations due to concrete shielding blocks that will be above the experimental apparatus, which greatly suppress non-muon events. The suppression of events in the plateau from shielding blocks is supported by further simulation studies with concrete slabs situated above the modules, not reported here.

7.3.4 Efficiency without Defining Modules

When running Mu2e, it will not be possible to perform the cleanup cuts used to define good candidate tracks described in Section 7.3.1 because defining modules cannot be used. Instead, the module efficiency will need to be measured in-situ, using only information from a single module. To do this, candidate events can be “tagged” by requiring hits in the top and bottom layers of a module and measuring the single-layer efficiency of the middle two layers. However, without cleanup cuts using defining modules, the candidate events will have a lower muon purity, which will affect the measured efficiency. Using the events triggered from the external scintillator paddles and by requiring at least 10 PEs deposited in the top and bottom layers of the testing module, the single-layer efficiency of the middle two layers was





Figure 7.13: Event displays for simulated non-muon events (top) with similar hit patterns to data events (bottom). These events do not pass the 3/4 layer veto requirement for low PE thresholds.

measured and compared to the single-layer efficiency shown in Figure 7.8. As seen in Figure 7.14, the single-layer efficiency is lower without cleanup cuts to improve the muon purity of the selected events, particularly for low thresholds. This makes the modules seem less efficient, but many of these inefficient events are not of importance to the CRV’s veto capabilities, as they come from particles other than muons. Also note that this in-situ method of measuring the efficiency can only measure the single-layer efficiency, which would then be extrapolated to a $3/4$ efficiency. As shown in Figure 7.9, the extrapolation of the single-layer efficiency underestimates the true $3/4$ efficiency. Therefore, only a lower bound on the module efficiency can be measured in-situ. Although defining modules cannot be used when running Mu2e, plans have been made to place external di-counters above and below CRV modules to more accurately “tag” candidate muon events, allowing a more robust measurement of the module efficiencies.

7.4 Position Resolution Study

The position of the energy deposition along the length of the scintillator counters can be determined by time of arrival (TOA) of the scintillation light to the SiPMs on each end of the detector. The cosmic ray veto requires a coincidence of layers hit in both space and time to avoid a false veto caused by noise, so determining the position resolution of the CRV modules is important to reduce the number of false coincidences.

To study the reconstruction of a hit’s position from the TOA information, a modified version of the cosmic-ray test stand described in Section 7.3 was used. Eight FEBs were used in this study, but instead of reading out four modules on one side, two




Figure 7.14: The single-layer inefficiency of the middle two layers of the testing module with the event selection described in Section 7.3.1 (purple triangles) and by only requiring hits in the top and bottom layers of the testing module (orange circles).

modules were read out on both sides⁶. Trigger paddles placed 1.0 m from one end of the modules were used for event triggering, as was done in the previous section⁷. As the trigger paddles are 10 cm wide, the expected position of cosmic rays from one end was 105 ± 5 cm.

First, air-shower events were analyzed to ensure the time was calibrated properly between FEBs connected to each end of the module. An air shower is a large cascade of particles originating from cosmic ray interactions in the atmosphere or the Wideband building material. These events leave particle tracks throughout the whole module simultaneously, which allows the time offset between individual SiPMs to be measured.

Air shower events were selected from the triggered data by requiring at least 20 PEs in every SiPM at both ends of the module. This ensured that there was activity in the entire detector, and increased the likelihood that simultaneous hits occurred on both ends of the module. The time difference between the SiPM pulse maximum on the far side and the near side ($\Delta t = t_{\text{far}} - t_{\text{near}}$) was measured for each of the 8 AFE chips in the two FEBs. This was done at the level of the AFE chips, instead of by individual channels or the whole FEB, because the clock settings of the FEB are set by AFE chip. Figure 7.15 shows the measured time difference between the AFE chips for the two FEB pairs used in this study, which comprised the full double-ended readout of a single CRV-T module. The time values given in each panel are the average time differences from the distributions, which were the values used to correct the time offsets between channel pairs on opposite ends of the module. Note that while the >20 PEs in all channels requirement increased the likelihood of simultaneous

⁶Note that this configuration of the electronics was due to the fact that only eight FEBs were available.

⁷Due to a limited number of trigger paddles, and to not disturb the configuration for aging measurements (Section 7.5), the same trigger paddle configuration as the efficiency study was used.

hits everywhere in the module, it did not guarantee that this occurred. Because events were triggered by a coincidence in scintillator paddles located 1 m from the readout on one end, there was a risk of biasing the location of the hits to where these paddles were located. This can be observed in the distributions of Figure 7.15 as the “shoulder” on the right side of the distributions, at around 20 ns. Despite this bias, the time calibration described here was an in-situ calibration, using the same collection of data and data-taking configuration as the measurement of the position. This configuration was used instead of placing paddles at the center of the modules, where the TOA should be the same, due to a lack of spare trigger paddles.




Figure 7.15: The time offset between AFE chips in FEBs on opposite ends of the module. The left plot grouping is for the top two layers of the module, and the right plot grouping is for the bottom two layers of the module.

After applying these small time offset corrections in offline analysis, the position measurement was made. To ensure events selected for this measurement were likely a

single particle passing through the trigger paddles, a maximum of two SiPMs in each layer (the two neighboring SiPMs in a counter) with more than 20 PEs registered was required. As seen in Figure 7.11, the most probable value of light yield in a layer was about 102 PEs, or 51 PEs per SiPM, when particles passed 1 m from the module end. The light yield was expected to be attenuated $\sim 50\%$ at 5 m away compared to 1 m away [159]. To be conservative and capture as many muon events as possible, a threshold of 20 PEs was set.

The position x (in cm) of a particle from the near side passing through a counter of length L (in this case, $L = 6$ m) was determined from the light arrival time difference between the far side and the near side Δt and the speed of light in the WLS fiber c_f as

$$x = \frac{1}{2}(L - c_f \Delta t). \quad (7.9)$$

The speed of light in the WLS fiber was determined to be 17.3 cm/ns from test beam studies on counters with 1.4 mm WLS fiber [155]. As each end of a counter is read out by two SiPMs, their measured time difference was averaged. Figure 7.16 shows the reconstructed position based on the light arrival time difference. A gaussian was fitted to the distribution to determine the mean value and standard deviation. The mean of 101.0 cm is consistent with the position of the trigger paddles at 105 ± 5 cm from the closer end of the module. The 1σ uncertainty of the measurement is 30.4 cm, which is about twice as large as the uncertainty from similar measurements using a test beam [155]. The 10 cm width of the trigger paddles may contribute to the larger uncertainty than what was observed in test beam studies, as well as the shorter (3 m) length of the counters used in the test beam.




Figure 7.16: The reconstructed position of cosmic-ray events from the end closest to the trigger paddles using the light arrival time differences from the two ends of a module and assuming a speed of light in the WLS fiber of 17.3cm/ns. Events are triggered from a coincidence in two trigger paddles, represented by the shaded gray region.

7.5 Aging of CRV Modules

It has long been known that the light yield from plastic scintillators decreases over time [160]. This phenomenon is referred to as the “aging” of plastic scintillators.

As the light yield of the CRV counters directly influences the veto efficiency, the aging rate of the CRV needs to be understood to ensure the CRV can maintain its required veto efficiency throughout the experiment’s lifetime. When designed, an aging rate of about 3% per year was anticipated, similar to what was observed in the MINOS experiment [161], which used very similar scintillator extrusions. It should be pointed out that plastic scintillators age non-uniformly over time, as observed, for example, by the T2K experiment [162]. The scintillator experiences a more rapid aging in the first

1–2 years after production, and then the aging decreases to a constant long-term rate. This section describes the measurement the long-term aging rate of the CRV counters. A series of cosmic-ray test stands using the CRV modules at the Fermilab Wideband facility were set up to measure the aging. The modules were manufactured using scintillator extruded in early 2018, and the first measurements from this cosmic-ray test stand were done in 2021, so none of the early, rapid aging behavior was captured. The aging of the WLS fiber was not expected to greatly affect the light yield decline measured here [163].

The configuration of the cosmic-ray test stand for aging measurements is described in Section 7.1. Two similar test stands were set up for aging measurements. The first consisted of four CRV-L modules containing WLS fibers of 1.4 mm diameter. Measurements on this module stack began in July 2021 with only 2 pre-production FEBs reading out the counters on a single end of one module, which were the only FEBs available at the time. The light yield on these counters has been tracked over nearly 3 years. Four more FEBs were added to this test stand in October 2022. The next test stand used for aging was set up in August 2022 and was instrumented with 6 pre-production FEBs. This test stand used CRV-T modules with 1.8 mm diameter WLS fibers.

The data collected for aging were calibrated using the temperature-dependent procedure described in Section 7.2.1. This calibration procedure was critical for an accurate measurement of the aging, since the measurements occurred several months apart with seasonal variations affecting the temperature of the electronics. Data for the aging measurements was not continuously read out, but was instead read out for about a month at a time before moving the electronics to another configuration for other tests. It was observed that in some SiPMs the single photo-electrons (SPEs)

experienced sudden jumps when the configuration was unplugged and plugged back in (see Figure 7.17). This was due to HDMI cable changes between measurements, and standalone measurements have shown that not all of the HDMI cables of the same length have the same electrical properties.




Figure 7.17: An example of the SPE change for a “bad” channel over time. The vertical lines indicate when the electronics were unplugged and moved to another study.

The stability of the SPEs was critical for the aging analysis, as the SPEs were used to calibrate the light yield. A deviation in the SPEs as large as $\sim 4\%$, seen in Figure 7.17, shifts the signal pulses by roughly the same amount, which is significantly more than the uncertainty of the light yield measurement ($\lesssim 1\%$). To ensure stability in the SPEs—and therefore the light yield measurement—quality cuts were applied to remove sub-runs in which the SPEs deviated by more than 0.5% from their median values. This requirement was to ensure the SPEs were stable within the uncertainty of the light yield measurement. If the time between the first remaining sub-runs and last remaining sub-run was less than one year, the whole SiPM was excluded. This is to ensure a sufficient baseline was recorded for the detector aging.

The light yield central value and uncertainty for each data point were extracted from the light yield distribution of each sub-run. As described in detail in Section 7.2.2, the light yield distribution is fitted with a convoluted Gauss-Landau function to determine the most probable value of the distribution. The light yield central value and uncertainty for this aging measurement were taken to be the fitted most probable value and the uncertainty on the fit. Each channel had its aging measured by fitting sub-runs passing the SPE stability requirement to the function

$$A_0(1 - r/100)^t, \quad (7.10)$$

where A_0 is the light yield at $t = 0$ in PEs, r is the aging rate in %/year, and t is the time in years since the initial measurement. The variable of interest to extract from this fit is r . An example aging measurement on one channel is shown in Figure 7.18. Since light produced in each scintillator counter is read out by two SiPMs, the fitted aging rates for the two SiPMs in a counter were averaged to determine the aging rate of the counter.

An average aging rate of 2.50 ± 0.35 %/year over approximately 2 years and 6 months was observed from 144 counters in CRV-T modules, as shown in Figure 7.19. A measurement on a separate set of 31 counters from a CRV-L module showed an aging rate of 2.49 ± 0.93 %/yr over approximately 2 years and 11 months, which is consistent with the aging observed in CRV-T modules.

Since the aging measured in SiPMs in the same counter were averaged to determine the aging of the counter, the difference in the measured aging between neighboring SiPMs was calculated, shown in Figure 7.20. The measured difference in aging rate of most SiPM pairs falls within the 1σ uncertainty of the aging measurement, 0.35%/yr.




Figure 7.18: Light yield over time of an example SiPM. Each point corresponds to a sub-run.




Figure 7.19: Distribution of measured aging in all counters with 1.8 mm WLS fibers. The mean aging is 2.50 %/yr with a standard deviation of 0.35 %/yr. A normal distribution with these values is drawn as the red dashed curve.




Figure 7.20: The difference in the measured aging rate of two SiPMs in the same counter.

7.6 Extrapolation to Overall CRV Efficiency

The expected background for Mu2e Run I from cosmic-ray-induced backgrounds was determined to be 0.046 ± 0.010 , which assumed an aging rate of 8.7%/yr [147]. Note that this assumed aging rate came from preliminary measurements, done on test counters before temperature-corrected calibrations, described in Section 7.2.1, were implemented. Keeping the expected background for the two Mu2e runs below one event relies on the CRV being 99.99% efficient. This study demonstrates that the efficiency of CRV modules exceeds 99.99% for reasonably low PE thresholds, although high enough to avoid false coincidences due to noise, and the aging rate of counters is approximately 2.5%/yr. Based on the module efficiency exceeding its 99.99% requirement and the aging being significantly lower than that used in previous studies, the CRV is expected to maintain an efficiency that keeps the anticipated cosmic-ray-induced background below one event over the entire Mu2e run time.

7.7 Discussion

The Mu2e experiment will be a sensitive probe of charged lepton flavor violation, vastly improving upon current experimental constraints on the $\mu^- N \rightarrow e^- N$ process. Mu2e's unprecedented sensitivity will only be possible with sufficient background vetoing, and the largest source of background comes from cosmic-ray-induced backgrounds. For Mu2e to meet its sensitivity goals, the CRV must maintain an efficiency of 99.99% over the experiment's run time.

Measurements of the cosmic-ray muon detection efficiency, reconstructed position resolution, and aging rate for CRV counters have been reported here. Using test stands consisting of CRV modules and scintillator paddles for external triggering, this study confirms that, at appropriate PE threshold levels, the efficiency of CRV modules can exceed the designed muon detection efficiency of 99.99%. Simulation studies show good agreement with the measurements regarding the relationship between module detection efficiency and the PE threshold, explaining features in this relationship caused by module geometry and non-muonic components in the cosmic rays. The veto is determined by a coincidence in space and time of energy deposition in at least three out of four layers of counters. Based on a light arrival time study, the reconstructed position of hits was consistent with the position of the triggering paddles with an uncertainty of 30.4 cm, or 5% of the module's length. The long-term aging rates of the production modules, with counters embedded with wavelength-shifting (WLS) fibers of two different diameters, were measured over multiple years. A similar aging rate of approximately 2.5% per year was reported for both types of counters. With the measured detection efficiency and an aging rate much lower than the assumed value in previous studies, it is expected that the CRV will remain over 99.99% efficient and will be capable of suppressing the cosmic-ray-induced background for the Mu2e

experiment to below one event over Mu2e's run time.

Bibliography

- [1] S. Chatrchyan et al. (CMS Collaboration), “Observation of a New Boson at a Mass of 125 GeV with the CMS Experiment at the LHC”, [Phys. Lett. B 716, 30–61 \(2012\)](#).
- [2] G. Aad et al. (ATLAS Collaboration), “Observation of a new particle in the search for the Standard Model Higgs boson with the ATLAS detector at the LHC”, [Phys. Lett. B 716, 1–29 \(2012\)](#).
- [3] J. M. Cline, “Baryogenesis”, in Les Houches Summer School - Session 86: Particle Physics and Cosmology: The Fabric of Spacetime (Sept. 2006).
- [4] Y. Fukuda et al. (Super-Kamiokande Collaboration), “Evidence for oscillation of atmospheric neutrinos”, [Phys. Rev. Lett. 81, 1562–1567 \(1998\)](#).
- [5] Q. R. Ahmad et al. (SNO Collaboration), “Measurement of the rate of $\nu_e + d \rightarrow p + p + e^-$ interactions produced by ^8B solar neutrinos at the sudbury neutrino observatory”, [Phys. Rev. Lett. 87, 071301 \(2001\)](#).
- [6] F. Zwicky, “Die rotverschiebung von extragalaktischen nebeln”, *Helvetica Physica Acta* 6, 110–127 (1933).
- [7] V. C. Rubin, W. K. Ford, and N. Thonnard, “Rotational properties of 21 sc galaxies with a large range of luminosities and radii, from NGC 4605 ($R = 4$ kpc) to UGC 2885 ($R = 122$ kpc)”, [Astrophysical Journal 238, 471–487 \(1980\)](#).
- [8] K. Freese, “Status of dark matter in the universe”, [International Journal of Modern Physics D 26, 1730012 \(2017\)](#).

- [9] N. Aghanim et al. (Planck Collaboration), “Planck 2018 results - vi. cosmological parameters”, [A&A 641, A6 \(2020\)](#).
- [10] D. N. Schramm and M. S. Turner, “Big-bang nucleosynthesis enters the precision era”, [Rev. Mod. Phys. 70, 303–318 \(1998\)](#).
- [11] V. Mossa, G. Mangano, O. Pisanti, A. Coc, H. Davoudiasl, J. Heeck, G. Hütsi, and J.-P. Uzan, “The baryon density of the universe from an improved rate of deuterium burning”, [Nature 587, 210–213 \(2020\)](#).
- [12] R. Massey et al., “Dark matter maps reveal cosmic scaffolding”, [Nature 445, 286–90 \(2007\)](#).
- [13] F. Leclercq, J. Jasche, P. Sutter, N. Hamaus, and B. Wandelt, “Dark matter voids in the sdss galaxy survey”, [Journal of Cosmology and Astroparticle Physics 2015, 047 \(2015\)](#).
- [14] M. Milgrom, “A modification of the newtonian dynamics as a possible alternative to the hidden mass hypothesis”, [Astrophys. J. 270, 365–370 \(1983\)](#).
- [15] D. Clowe, M. Bradač, A. H. Gonzalez, M. Markevitch, S. W. Randall, C. Jones, and D. Zaritsky, “A direct empirical proof of the existence of dark matter”, [The Astrophysical Journal 648, L109 \(2006\)](#).
- [16] Paraficz, D., Kneib, J.-P., Richard, J., Morandi, A., Limousin, M., Jullo, E., and Martinez, J., “The bullet cluster at its best: weighing stars, gas, and dark matter”, [A&A 594, A121 \(2016\)](#).
- [17] M. Markevitch, A. H. Gonzalez, D. Clowe, A. Vikhlinin, W. Forman, C. Jones, S. Murray, and W. Tucker, “Direct Constraints on the Dark Matter Self-Interaction Cross Section from the Merging Galaxy Cluster 1E 0657-56”, [The Astrophysical Journal 606, 819–824 \(2004\)](#).

- [18] W. J. G. de Blok, “The Core-Cusp Problem”, [Adv. Astron. 2010, 789293 \(2010\)](#).
- [19] S. Garrison-Kimmel, M. Boylan-Kolchin, J. S. Bullock, and E. N. Kirby, “Too big to fail in the local group”, [Monthly Notices of the Royal Astronomical Society 444, 222–236 \(2014\)](#).
- [20] I. de Martino, S. S. Chakrabarty, V. Cesare, A. Gallo, L. Ostorero, and A. Diaferio, “Dark matters on the scale of galaxies”, [Universe 6, 10.3390/universe6080107 \(2020\)](#).
- [21] L. Angel, G. Gambini, L. Guedes, F. S. Queiroz, and V. de Souza, “Constraining gamma-ray lines from dark matter annihilation using fermi-lat and h.e.s.s. data”, [Journal of Cosmology and Astroparticle Physics 2024, 028 \(2024\)](#).
- [22] J. D. Bekenstein, “Relativistic gravitation theory for the modified newtonian dynamics paradigm”, [Phys. Rev. D 70, 083509 \(2004\)](#).
- [23] S. Boran, S. Desai, E. O. Kahya, and R. P. Woodard, “Gw170817 falsifies dark matter emulators”, [Phys. Rev. D 97, 041501 \(2018\)](#).
- [24] B. Carr, F. Kühnel, and L. Visinelli, “Constraints on stupendously large black holes”, [Monthly Notices of the Royal Astronomical Society 501, 2029–2043 \(2020\)](#).
- [25] R. Laha, J. B. Muñoz, and T. R. Slatyer, “*INTEGRAL* Constraints on primordial black holes and particle dark matter”, [Phys. Rev. D 101, 123514 \(2020\)](#).
- [26] G. Steigman, D. N. Schramm, and J. E. Gunn, “Cosmological limits to the number of massive leptons”, [Physics Letters B 66, 202–204 \(1977\)](#).
- [27] E. W. Kolb and M. S. Turner, *The Early Universe*, Vol. 69 (Taylor and Francis, May 2019).

- [28] G. Jungman, M. Kamionkowski, and K. Griest, “Supersymmetric dark matter”, [Physics Reports 267, 195–373 \(1996\)](#).
- [29] K. Griest and M. Kamionkowski, “Unitarity limits on the mass and radius of dark-matter particles”, [Phys. Rev. Lett. 64, 615–618 \(1990\)](#).
- [30] R. J. Scherrer and M. S. Turner, “On the relic, cosmic abundance of stable, weakly interacting massive particles”, [Phys. Rev. D 33, 1585–1589 \(1986\)](#).
- [31] J. R. Ellis, J. S. Hagelin, D. V. Nanopoulos, K. A. Olive, and M. Srednicki, “Supersymmetric Relics from the Big Bang”, [Nucl. Phys. B 238, edited by M. A. Srednicki, 453–476 \(1984\)](#).
- [32] S. Dimopoulos and H. Georgi, “Softly broken supersymmetry and $su(5)$ ”, [Nuclear Physics B 193, 150–162 \(1981\)](#).
- [33] V. A. Mitsou, “Searches for Supersymmetry with the ATLAS and CMS detectors”, in [Proceedings of corfu summer institute 2022 “school and workshops on elementary particle physics and gravity”](#), Vol. 436 (2023), p. 074.
- [34] P. Agnes et al. (DarkSide-50 Collaboration), “Search for low-mass dark matter wimps with 12 ton-day exposure of darkside-50”, [Phys. Rev. D 107, 063001 \(2023\)](#).
- [35] D. Huang et al. (PandaX Collaboration), “Search for dark-matter–nucleon interactions with a dark mediator in pandax-4t”, [Phys. Rev. Lett. 131, 191002 \(2023\)](#).
- [36] J. Aalbers et al. (LUX-ZEPLIN Collaboration), “First dark matter search results from the lux-zeplin (lz) experiment”, [Phys. Rev. Lett. 131, 041002 \(2023\)](#).
- [37] E. Aprile et al. (XENON Collaboration), “First search for light dark matter in the neutrino fog with xenonnt”, [Phys. Rev. Lett. 134, 111802 \(2025\)](#).

- [38] C. A. J. O’Hare, “New definition of the neutrino floor for direct dark matter searches”, [Phys. Rev. Lett. 127, 251802 \(2021\)](#).
- [39] J. F. Beacom, “The diffuse supernova neutrino background”, [Annual Review of Nuclear and Particle Science 60, 439–462 \(2010\)](#).
- [40] D. S. Akerib et al., Snowmass2021 cosmic frontier dark matter direct detection to the neutrino fog, 2022.
- [41] M. Pospelov, A. Ritz, and M. Voloshin, “Secluded wimp dark matter”, *Physics Letters B* 662, 53–61 (2008).
- [42] B. W. Lee and S. Weinberg, “Cosmological lower bound on heavy-neutrino masses”, [Phys. Rev. Lett. 39, 165–168 \(1977\)](#).
- [43] J. D. Bjorken, R. Essig, P. Schuster, and N. Toro, “New fixed-target experiments to search for dark gauge forces”, [Phys. Rev. D 80, 075018 \(2009\)](#).
- [44] B. Holdom, “Two $u(1)$ ’s and epsilon charge shifts”, [Phys. Lett. B 166, 196–198 \(1986\)](#).
- [45] N. Arkani-Hamed and N. Weiner, “Lhc signals for a superunified theory of dark matter”, [Journal of High Energy Physics 2008, 104 \(2008\)](#).
- [46] A. Berlin, N. Blinov, G. Krnjaic, P. Schuster, and N. Toro, “Dark matter, millicharges, axion and scalar particles, gauge bosons, and other new physics with ldmx”, [Phys. Rev. D 99, 075001 \(2019\)](#).
- [47] H. Davoudiasl and W. J. Marciano, “Running of the $u(1)$ coupling in the dark sector”, *Physical Review D* 92, 035008 (2015).
- [48] E. Izaguirre, G. Krnjaic, P. Schuster, and N. Toro, “Analyzing the discovery potential for light dark matter”, *Physical Review Letters* 115, 251301 (2015).

- [49] J. L. Feng and J. Smolinsky, “Impact of a resonance on thermal targets for invisible dark photon searches”, [Phys. Rev. D 96, 095022 \(2017\)](#).
- [50] J. Lees et al. (BaBar Collaboration), “Search for invisible decays of a dark photon produced in $e^+ e^-$ collisions at babar”, *Phys. Rev. Lett.* 119, 131804 (2017).
- [51] L. Auerbach et al., “Measurement of electron-neutrino electron elastic scattering”, *Physical Review D* 63, 112001 (2001).
- [52] P. DeNiverville, M. Pospelov, and A. Ritz, “Observing a light dark matter beam with neutrino experiments”, *Physical Review D* 84, 075020 (2011).
- [53] B. Batell, R. Essig, and Z. Surujon, “Strong constraints on sub-gev dark sectors from slac beam dump e137”, *Physical Review Letters* 113, 171802 (2014).
- [54] D. Banerjee, V. Burtsev, A. Chumakov, D. Cooke, P. Crivelli, E. Depero, A. Dermenev, S. Donskov, R. Dusaev, T. Enik, et al., “Dark matter search in missing energy events with na64”, *Physical review letters* 123, 121801 (2019).
- [55] A. Aguilar-Arevalo et al. (MiniBooNE-DM Collaboration), “Dark matter search in a proton beam dump with miniboone”, *Phys. Rev. Lett.* 118, 221803 (2017).
- [56] J. Beacom, N. Bell, and G. Bertone, “Gamma-ray constraint on galactic positron production by mev dark matter”, *Phys. Rev. Lett.* 94 (2005).
- [57] G. W. Bennett et al. (Muon $g-2$ Collaboration), “Final report of the e821 muon anomalous magnetic moment measurement at bnl”, [Phys. Rev. D 73, 072003 \(2006\)](#).
- [58] D. P. Aguillard et al. (Muon $g-2$ Collaboration), “Measurement of the positive muon anomalous magnetic moment to 0.20 ppm”, [Phys. Rev. Lett. 131, 161802 \(2023\)](#).

- [59] M. Pospelov, “Secluded $u(1)$ below the weak scale”, [Physical Review D 80, 10.1103/physrevd.80.095002 \(2009\)](#).
- [60] B. Aubert et al. (BABAR Collaboration), “Search for dimuon decays of a light scalar boson in radiative transitions $\Upsilon \rightarrow \gamma A^0$ ”, [Phys. Rev. Lett. 103, 081803 \(2009\)](#).
- [61] J. R. Batley et al. (NA48/2 Collaboration), Search for the dark photon in π^0 decays, 2015.
- [62] D. Banerjee et al. (NA64 Collaboration), “Improved limits on a hypothetical $X(16.7)$ boson and a dark photon decaying into e^+e^- pairs”, [Physical Review D 101, 10.1103/physrevd.101.071101 \(2020\)](#).
- [63] A. Bazavov et al., “Hadronic vacuum polarization for the muon $g-2$ from lattice QCD: Long-distance and full light-quark connected contribution”, (2024).
- [64] A. J. Krasznahorkay et al., “Observation of anomalous internal pair creation in ^8Be : a possible indication of a light, neutral boson”, [Phys. Rev. Lett. 116, 042501 \(2016\)](#).
- [65] M. F. Albakry et al. (SuperCDMS Collaboration), “Search for low-mass dark matter via bremsstrahlung radiation and the migdal effect in supercdms”, [Phys. Rev. D 107, 112013 \(2023\)](#).
- [66] P. Adari et al. (SENSEI Collaboration), “First direct-detection results on sub-gev dark matter using the sensei detector at snolab”, [Phys. Rev. Lett. 134, 011804 \(2025\)](#).
- [67] G. Krnjaic et al., A snowmass whitepaper: dark matter production at intensity-frontier experiments, 2022.

- [68] D. Akimov et al. (COHERENT Collaboration), “First probe of sub-gev dark matter beyond the cosmological expectation with the coherent csi detector at the sns”, *Physical Review Letters* 130, 051803 (2023).
- [69] P. Adamson et al., “The numi neutrino beam”, *Nucl. Instrum. Methods Phys. Res. Sect. A* 806, 279–306 (2016).
- [70] Fermilab, Concepts rookie book, <https://rookiebooks.fnal.gov/concepts/>, Fermi National Accelerator Laboratory.
- [71] S. Mufson et al., “Liquid scintillator production for the NOvA experiment”, *Nucl. Instrum. Meth. A* 799, 1–9 (2015).
- [72] A. Rodriguez and A. Laio, “Clustering by fast search and find of density peaks”, *Science* 344, 1492–1496 (2014).
- [73] R. C. Prim, “Shortest connection networks and some generalizations”, *Bell System Technical Journal* 36, 1389–1401 (1957).
- [74] L. A. Fernandes and M. M. Oliveira, “Real-time line detection through an improved hough transform voting scheme”, *Pattern Recognition* 41, 299–314 (2008).
- [75] M. Ohlsson, C. Peterson, and A. L. Yuille, “Track finding with deformable templates — the elastic arms approach”, *Computer Physics Communications* 71, 77–98 (1992).
- [76] M. Baird et al., “Event reconstruction techniques in nova”, *J. Phys. Conf. Ser.* 664, 072035 (2015).
- [77] M.-S. Yang and K.-L. Wu, “Unsupervised possibilistic clustering”, *Pattern Recognition* 39, 5–21 (2006).

- [78] J. C. Dunn, “A fuzzy relative of the isodata process and its use in detecting compact well-separated clusters”, [Journal of Cybernetics 3, 32–57 \(1973\)](#).
- [79] P. deNiverville and C. Frugiuele, “Hunting sub-gev dark matter with the nova near detector”, *Physical Review D* 99, 051701 (2019).
- [80] J. Paley et al. (MIPP Collaboration), “Measurement of charged pion production yields off the numi target”, *Physical Review D* 90, 032001 (2014).
- [81] P. Deniverville, C.-Y. Chen, M. Pospelov, and A. Ritz, “Light dark matter in neutrino beams: production modeling and scattering signatures at miniboone, t2k, and ship”, *Physical Review D* 95, 035006 (2017).
- [82] G. Agakishiev et al. (HADES Collaboration), “Inclusive pion and η production in p +nb collisions at 3.5 gev beam energy”, [Phys. Rev. C 88, 024904 \(2013\)](#).
- [83] T. Åkesson et al., “Inclusive π production at low transverse momentum in 63 gev ppcollisions at the cern intersecting storage rings”, [Phys. Lett. B 178, 447–451 \(1986\)](#).
- [84] J. A. Formaggio and G. P. Zeller, “From ev to eev: neutrino cross sections across energy scales”, [Rev. Mod. Phys. 84, 1307–1341 \(2012\)](#).
- [85] L. Lyons, “Raster scan or 2-d approach?”, arXiv preprint [arXiv:1404.7395](#) (2014).
- [86] M. A. Acero et al. (NOvA Collaboration), “Monte Carlo method for constructing confidence intervals with unconstrained and constrained nuisance parameters in the NOvA experiment”, [J. Instrum. 20, T02001 \(2025\)](#).
- [87] L. Aliaga et al. (MINERvA Collaboration), “Neutrino flux predictions for the numi beam”, [Phys. Rev. D 94, 092005 \(2016\)](#).

- [88] C. Andreopoulos et al., “The GENIE neutrino Monte Carlo generator”, [Nuclear Instruments and Methods in Physics Research A 614, 87–104 \(2010\)](#).
- [89] C. Andreopoulos, C. Barry, S. Dytman, H. Gallagher, T. Golan, R. Hatcher, G. Perdue, and J. Yarba, The genie neutrino monte carlo generator: physics and user manual, 2015.
- [90] S. Agostinelli et al., “Geant4—a simulation toolkit”, [Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment 506, 250–303 \(2003\)](#).
- [91] J. Allison et al., “Geant4 developments and applications”, *IEEE Transactions on Nuclear Science* 53, 270–278 (2006).
- [92] J. Allison et al., “Recent developments in geant4”, [Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment 835, 186–225 \(2016\)](#).
- [93] L. A. Soplin, “Neutrino flux prediction for the numi beamline”, Ph.D. dissertation (College of William and Mary, Williamsburg, Virginia, Fall 2016).
- [94] D. Bohm and D. Pines, “A collective description of electron interactions. i. magnetic interactions”, [Phys. Rev. 82, 625–634 \(1951\)](#).
- [95] D. Pines and D. Bohm, “A collective description of electron interactions: ii. collective vs individual particle aspects of the interactions”, [Phys. Rev. 85, 338–353 \(1952\)](#).
- [96] D. Bohm and D. Pines, “A collective description of electron interactions: iii. coulomb interactions in a degenerate electron gas”, [Phys. Rev. 92, 609–625 \(1953\)](#).

- [97] J. Nieves, J. E. Amaro, and M. Valverde, “Inclusive quasielastic charged-current neutrino-nucleus reactions”, [Phys. Rev. C 70, 055503 \(2004\)](#).
- [98] P. Abratenko et al. (MicroBooNE Collaboration), “First search for dark-trident processes using the microboone detector”, [Phys. Rev. Lett. 132, 241801 \(2024\)](#).
- [99] A. A. Aguilar-Arevalo et al. (CCM Collaboration), “First dark matter search results from Coherent CAPTAIN-Mills”, [Phys. Rev. D 106, 012001 \(2022\)](#).
- [100] M. Bonesini, A. Marchionni, F. Pietropaolo, and T. Tabarelli de Fatis, “On particle production for high energy neutrino beams”, *Eur. Phys. J. C* 20 (2001).
- [101] S. Teis, W. Cassing, M. Effenberger, A. Hombach, U. Mosel, and G. Wolf, “Pion-production in heavy-ion collisions at sis energies”, *Zeitschrift für Physik A Hadrons and Nuclei* 356, 421–435 (1996).
- [102] K. Jaeger, J. Campbell, G. Charlton, D. Swanson, C. Fu, H. Rubin, R. Glasser, D. Koetke, and J. Whitmore, “Inclusive γ , π^0 , K^0 , and Λ production in 12.4-gev/c p p interactions”, *Physical Review D* 11, 1756 (1975).
- [103] W. Wu and Y. Xiao, “Constraint of the integrated neutrino flux from neutrino-electron elastic scattering in the nova near detector”, *NOvA DocDB-56383* (2025).
- [104] A. Aurisano, A. Radovic, D. Rocco, A. Himmel, M. Messier, E. Niner, G. Pawloski, F. Psihas, A. Sousa, and P. Vahle, “A convolutional neural network neutrino event classifier”, *J. Instrum.* 11, P09001 (2016).
- [105] A. T. C. Sutton, “Domain Generalization with Machine Learning in the NOvA Experiment”, PhD thesis (University of Virginia, 2023).

- [106] J. B. Birks, “Scintillations from organic crystals: specific fluorescence and relative response to different radiations”, [Proceedings of the Physical Society. Section A 64](#), 874 (1951).
- [107] Y. M. Andreev et al. (NA64 Collaboration), “Search for light dark matter with na64 at cern”, [Phys. Rev. Lett. 131](#), 161801 (2023).
- [108] L. Ren et al. (MINERvA Collaboration), “Measurement of the antineutrino to neutrino charged-current interaction cross section ratio in minerva”, [Phys. Rev. D 95](#), 072009 (2017).
- [109] T. Åkesson et al. (LDMX Collaboration), “Photon-rejection power of the Light Dark Matter eXperiment in an 8 GeV beam”, [JHEP 12](#), 092 (2023).
- [110] T. Markiewicz, T. Raubenheimer, and N. Toro (LESA construction team), “The SLAC Linac to ESA (LESA) Beamline for Dark Sector Searches and Test Beams”, in Snowmass 2021 (May 2022).
- [111] J. F. Schmerge et al., “The LCLS-II Injector Design”, in Cosmology after Planck (Feb. 2015).
- [112] T. Raubenheimer, “The LCLS-II-HE, A High Energy Upgrade of the LCLS-II”, in [60th ICFA Advanced Beam Dynamics Workshop on Future Light Sources](#) (2018), MOP1WA02.
- [113] A. Brachmann, B. Dunham, and J. Schmerge, “LCLS-II - Status and Upgrades”, in [39th International Free Electron Laser Conference](#) (2019), FRA02.
- [114] T. Raubenheimer, A. Beukers, A. Fry, C. Hast, T. Markiewicz, Y. Nosochkov, N. Phinney, P. Schuster, and N. Toro, “DASEL: Dark Sector Experiments at LCLS-II”, (2018).

- [115] P. H. Adrian, “The silicon vertex tracker for the heavy photon search experiment”, in [2015 IEEE Nuclear Science Symposium and Medical Imaging Conference \(NSS/MIC\)](#) (2015), pp. 1–6.
- [116] “The Phase-2 Upgrade of the CMS Endcap Calorimeter”, [10.17181/CERN.IV8M.1JY2](#) (2017).
- [117] A. Martelli, The cms hgcaldetector for hl-lhc upgrade, 2017.
- [118] S. Boi, “Design and Fabrication of a Novel Large-area, High-efficiency Cosmic Ray Veto Detector for the Mu2e Experiment”, PhD thesis (Virginia U., 2021).
- [119] T. Eichlersmith, J. Mans, O. Moreno, J. Muse, M. Revering, and N. Toro, “Simulation of dark bremsstrahlung in Geant4”, [Comput. Phys. Commun.](#) **287**, 108690 (2023).
- [120] J. Alwall, R. Frederix, S. Frixione, V. Hirschi, F. Maltoni, O. Mattelaer, H. -. Shao, T. Stelzer, P. Torrielli, and M. Zaro, “The automated computation of tree-level and next-to-leading order differential cross sections, and their matching to parton shower simulations”, [JHEP](#) **07**, 079 (2014).
- [121] F. Ambrogio, C. Arina, M. Backovic, J. Heisig, F. Maltoni, L. Mantani, O. Mattelaer, and G. Mohlabeng, “MadDM v.3.0: a Comprehensive Tool for Dark Matter Studies”, [Phys. Dark Univ.](#) **24**, 100249 (2019).
- [122] T. Chen and C. Guestrin, “Xgboost: a scalable tree boosting system”, in [Proceedings of the 22nd ACM SIGKDD international conference on knowledge discovery and data mining](#), KDD ’16 (2016), pp. 785–794.
- [123] R. Aaij et al. (LHCb Collaboration), “Search for $A' \rightarrow \mu^+\mu^-$ decays”, [Phys. Rev. Lett.](#) **124**, 041801 (2020).

- [124] E. Cortina Gil et al. (NA62), “Search for dark photon decays to $\mu^+\mu$ at NA62”, [JHEP 09, 035 \(2023\)](#).
- [125] G. Punzi, “Sensitivity of searches for new signals and its optimization”, eConf C030908, edited by L. Lyons, R. P. Mount, and R. Reitmeyer, MODT002 (2003).
- [126] H. Qu and L. Gouskos, “Jet tagging via particle clouds”, [Phys. Rev. D 101, 056019 \(2020\)](#).
- [127] T. Junk, “Confidence level computation for combining searches with small statistics”, [Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment 434, 435–443 \(1999\)](#).
- [128] A. L. Read, “Presentation of search results: The CL_s technique”, [J. Phys. G 28, edited by M. R. Whalley and L. Lyons, 2693–2704 \(2002\)](#).
- [129] F. Archilli et al. (KLOE-2 Collaboration), “Search for a vector gauge boson in ϕ meson decays with the kloe detector”, [Physics Letters B 706, 251–255 \(2012\)](#).
- [130] D. Babusci et al. (KLOE-2 Collaboration), “Limit on the production of a light vector gauge boson in ϕ meson decays with the kloe detector”, [Physics Letters B 720, 111–115 \(2013\)](#).
- [131] R. Aaij et al. (LHCb Collaboration), “Search for dark photons produced in 13 tev pp collisions”, [Phys. Rev. Lett. 120, 061801 \(2018\)](#).
- [132] B. Petersen (FASER Collaboration), First Physics Results from the FASER Experiment, tech. rep., Contribution to the 2023 Electroweak session of the 57th Rencontres de Moriond. 6 pages, 4 figures (CERN, Geneva, 2023).

- [133] A. Konaka et al. (KEK Collaboration), “Search for neutral particles in electron-beam-dump experiment”, [Phys. Rev. Lett. 57, 659–662 \(1986\)](#).
- [134] M. Davier and H. Nguyen Ngoc, “An Unambiguous Search for a Light Higgs Boson”, [Phys. Lett. B 229, 150–155 \(1989\)](#).
- [135] S. Andreas, C. Niebuhr, and A. Ringwald, “New limits on hidden photons from past electron beam dumps”, [Phys. Rev. D 86, 095019 \(2012\)](#).
- [136] J. Blumlein et al., “Limits on neutral light scalar and pseudoscalar particles in a proton beam dump experiment”, [Z. Phys. C 51, 341–350 \(1991\)](#).
- [137] J. Blumlein et al., “Limits on the mass of light (pseudo)scalar particles from Bethe-Heitler $e^+ e^-$ and $\mu^+ \mu^-$ pair production in a proton - iron beam dump experiment”, [Int. J. Mod. Phys. A 7, 3835–3850 \(1992\)](#).
- [138] D. Banerjee et al. (NA64 Collaboration), “Dark matter search in missing energy events with na64”, [Physical Review Letters 123, 10.1103/physrevlett.123.121801 \(2019\)](#).
- [139] R. Bouchendira, P. Cladé, S. Guellati-Khélifa, F. Nez, and F. Biraben, “New determination of the fine structure constant and test of the quantum electrodynamics”, [Phys. Rev. Lett. 106, 080801 \(2011\)](#).
- [140] T. Aoyama, M. Hayakawa, T. Kinoshita, and M. Nio, “Tenth-order qed contribution to the electron $g - 2$ and an improved value of the fine structure constant”, [Phys. Rev. Lett. 109, 10.1103/physrevlett.109.111807 \(2012\)](#).
- [141] D. Hanneke, S. Fogwell, and G. Gabrielse, “New measurement of the electron magnetic moment and the fine structure constant”, [Phys. Rev. Lett. 100, 120801 \(2008\)](#).

- [142] H. Davoudiasl, H.-S. Lee, and W. J. Marciano, “Dark side of higgs diphoton decays and muon $g - 2$ ”, [Physical Review D 86, 10.1103/physrevd.86.095009 \(2012\)](#).
- [143] L. Bartoszek et al. (Mu2e Collaboration), “Mu2e Technical Design Report”, [10.2172/1172555 \(2014\)](#).
- [144] W. J. Marciano, T. Mori, and J. M. Roney, “Charged Lepton Flavor Violation Experiments”, [Ann. Rev. Nucl. Part. Sci. 58, 315–341 \(2008\)](#).
- [145] S. Davidson, B. Echenard, R. H. Bernstein, J. Heeck, and D. G. Hitlin, Charged lepton flavor violation, 2022.
- [146] W. Bertl et al., “A search for $\mu \rightarrow e$ conversion in muonic gold”, [The European Physical Journal C 47, 337–346 \(2006\)](#).
- [147] F. Abdi et al. (Mu2e Collaboration), “Mu2e Run I Sensitivity Projections for the Neutrinoless $\mu^- \rightarrow e^-$ Conversion Search in Aluminum”, [Universe 9, 54 \(2023\)](#).
- [148] A. Czarnecki, X. Garcia i Tormo, and W. J. Marciano, “Muon decay in orbit: spectrum of high-energy electrons”, [Phys. Rev. D 84, 013006 \(2011\)](#).
- [149] R. H. Bernstein, “The Mu2e Experiment”, [Frontiers in Physics 7, 1, 1 \(2019\)](#).
- [150] Kuraray America, Inc., 200 Park Ave., N,Y 10166 USA; 3-1-6, NIHONBASHI, CHUOKU, TOKYO 103-8254, JAPAN. ().
- [151] Hamamatusu Corp. (360 Foothill Rd., PO Box 6910, Bridgewater, NJ 08807-0919, USA).
- [152] AFE5807, Texas Instruments, Inc., 2500 TI Blvd, Suite T, Dallas, TX ().

- [153] L. Curtis et al., Temperature-dependent calibration procedures for the silicon photomultiplier readout of the cosmic ray veto detector for the mu2e experiment, 2025.
- [154] E. J. Gumbel, “Les valeurs extrêmes des distributions statistiques”, [Annales de l’institut Henri Poincaré](#) 5, 115–158 (1935).
- [155] A. Artikov et al. (Mu2e Collaboration), “Photoelectron Yields of Scintillation Counters with Embedded Wavelength-Shifting Fibers Read Out With Silicon Photomultipliers”, [Nucl. Instrum. Meth. A](#) 890, 84–95 (2018).
- [156] C. Hagmann, D. Lange, and D. Wright, “Cosmic-ray shower generator (cry) for monte carlo transport codes”, in [2007 ieee nuclear science symposium conference record](#), Vol. 2 (2007), pp. 1143–1146.
- [157] E. C. Dukes, R. Ehrlich, C. Group, S. Hansen, Y. Oksuzian, A. Pla-Dalmau, P. Rubinov, E. Song, and Y. Wu, Simulation of scintillator counters with embedded wavelength shifting fibers, Internal document. Mu2e DocDb 23929, 2019.
- [158] “Cosmic Shielding Studies at MicroBooNE”, [10.2172/1573042](#) (2016).
- [159] G. DeZoort, E. C. Dukes, R. C. Group, H. Kessenich, Y. Oksuzian, T. Rase, and D. Shooltz, “Performance of Wavelength-Shifting Fibers for the Mu2e Cosmic Ray Veto Detector”, (2015).
- [160] C. F. Barnaby and J. C. Barton, “Ageing of plastic scintillators”, [Journal of Scientific Instruments](#) 39, 176 (1962).
- [161] D. Michael et al., “The magnetized steel and scintillator calorimeters of the minos experiment”, [Nuclear Instruments and Methods in Physics Research](#)

Section A: Accelerators, Spectrometers, Detectors and Associated Equipment

596, 190–228 (2008).

- [162] K. Abe et al., “Scintillator ageing of the t2k near detectors from 2010 to 2021”, [Journal of Instrumentation](#) **17**, P10028 (2022).
- [163] M. Solt, D. Coveyou, E. C. Dukes, R. C. Group, Y. Oksuzian, and S. Roberts, “Performance of the wavelength-shifting fiber upgrade for the Mu2e cosmic-ray veto detector”, [JINST](#) **18**, T05004 (2023).

Appendices