

DISSERTATION

EXTENDING ICARUS'S SENSITIVITY TO HIDDEN SECTOR SCALARS WITH A
SEARCH FOR UNCONTAINED DIMUON FINAL STATES

Submitted by

Jamie Christine Dyer

Department of Physics

In partial fulfillment of the requirements

For the Degree of Doctor of Philosophy

Colorado State University

Fort Collins, Colorado

Summer 2026

Doctoral Committee:

Advisor: Joshua Berger

Robert Wilson

Michael Mooney

Asa Ben-Hur

Copyright by Jamie Dyer 2026

All Rights Reserved

ABSTRACT

EXTENDING ICARUS'S SENSITIVITY TO HIDDEN SECTOR SCALARS WITH A SEARCH FOR UNCONTAINED DIMUON FINAL STATES

The ICARUS detector is a liquid argon time projection chamber (LArTPC) acting as the far detector in the Short Baseline Neutrino program at Fermilab. Complementing its ability to better constrain neutrino physics, ICARUS is also able to probe any dark sector state present in the exposed flux, that results in a final state of charged Standard Model particles within the active volume of the detector. This work exploits such sensitivity to the Higgs portal scalar (HPS) and heavy QCD axion (ALP, for "axion-like particles") models with a selection for final states with exactly two muons, where at least one of the muons is not contained in the fiducial volume of the detector. The ALPs and HPSs are generated from meson-decay or -mixing in the Neutrinos at the Main Injector (NuMI) beam at Fermilab, propagate to ICARUS, and then decay into the dimuon pair. The work expands on the sensitivity probed by an earlier ICARUS result that selected final states of two fully-contained muons, extending ICARUS's reach to a higher phase space of dark sector models. The ICARUS detector, HPS and ALP models, simulation pipelines, uncontained dimuon event selection, systematic uncertainties, and strategy for extracting results are discussed.

ACKNOWLEDGEMENTS

My ability to complete this dissertation is owed entirely to the vast network of support that has carried me through the last several years.

First, I would like to thank my PhD advisor, Josh Berger, for the continued mentorship he offered through the entirety of my time as a researcher at CSU. Josh was always available to answer questions and help where needed, and I appreciate the significant time and effort he invested in me and the projects we have worked on together.

I would also like to thank my committee members Bob Wilson and Mike Mooney, who, in addition to Josh, provided the resources, opportunity, and encouragement to gain first-hand research experience at Fermilab, the home of the ICARUS experiment, and home to me for two years of my PhD experience. My time at Fermilab profoundly enriched my research contributions and experience. Furthermore, Bob and Mike lent ample expertise on experimental and collaboration matters, which was very useful as I navigated the later stages of the work presented in this dissertation.

I also owe significant thanks to Gray Putnam, who has provided me with invaluable advice and assistance with the analysis presented here, including helping to set me up with the tools used for analysis, managing data sets, and talking through various analysis strategies with me. In addition to their mentorship, I also benefit from having a scientific role model and a friend in Gray.

PhD's are long, and even this 100+ page thesis does not encapsulate all the scientific tinkering I engaged with throughout the course of this degree. Throughout my somewhat meandering research journey, I have had the honor of collaborating with and learning from many physicists with diverse areas of expertise. I am particularly grateful to the numerous ICARUS Collaborators I got to work with and learn from along the way. Though there are far too many to list, here are a handful meriting special acknowledgment:

Christian Farnese provided useful guidance with the hand-scanning of ICARUS data in my early days as an ICARUS collaborator, laying a foundation on which I built further understanding of TPC signals and that I leveraged heavily in designing the event selection for the analysis

presented here. Gianluca Petrillo helped to incorporate work completed prior to officially joining the ICARUS collaboration into the collaboration's code repository after I did join. Tyler Boone and Anna Heggstuen helped me navigate the ICARUS computing system and code bases when I first joined the collaboration. Vito Di Benedetto and Promita Roy offered significant support in the creation and management of the many datasets used in the analysis presented here. Francois Drielsma, Tracy Usher, and Justin Mueller offered valuable guidance during a period of time when I dabbled with using a machine learning based reconstruction framework for my analysis. Bruce Howard helped familiarize me with the Pandora reconstruction software and outputs, and offered mentorship on studies I did to gauge Pandora's shower reconstruction performance. Geoff Savage taught me a great deal about the software systems that facilitate ICARUS operations, and offered invaluable guidance for the projects I did contributing to the detector control system (DCS) for ICARUS. Antoni Giosa, Niccolo Mogi, Donatella Torrato, and Polina Abratenko also supported me in various projects related to DCS work. Matteo Vicenzi and Gian Luca Raseli trained me as a PMT hardware expert, and offered further expertise when I did not know enough to address issues that arose with the hardware. Filippo Varanini offered useful suggestions for the Multiple Coulomb Scattering momentum reconstruction technique used in my analysis. Finally, Vittorio Paolone, Animesh Chatterjee, Dan Cherdack, and Minerba Betancourt, along with other members of the NuMI BSM and Neutrino Cross Section working groups, offered valuable input at every stage of my analysis, which helped to make it as robust as possible and shape it into the form presented here.

Furthermore, I would like to extend thanks to my CSU cohort mates, especially Anna Heggstuen, Adam Whitney, Ryan Baker, and Camden Kasik. The intellectual exchanges and comradery I shared with all of them helped me get through the challenging first two years of graduate coursework and laid the foundation for what are now some of my dearest friendships. I am especially grateful to have found a friend and sister in Anna, with whom I shared many of the key experiences of this program - coursework at CSU, collaboration on ICARUS, life at Fermilab, and navigating it all as a woman in STEM.

I appreciate, also, the invested time and effort of my professors, for instilling in me the foundational knowledge required for a fruitful career as a physicist.

I would be remiss not to also acknowledge the friends that have enriched my personal life, and graduate school experience in turn, over the last several years.

I am grateful for the dear friends I have made in Fort Collins. In addition to cohort mates listed above, this includes Jacqueline Whitney, Jessica Bowse, Emily Pinckert, Biswajit Padhi, Mark Lide, and Mia Sawaya. Thanks for all the good conversations, family dinners, movie nights, volleyball games, and mountain adventures.

I also found a strong sense of community during my time at Fermilab, which, more than chance, I believe to be the product of the concerted effort of the people there. It was a joy to meet, work, and live alongside researchers from all over the world, united in pursuit and celebration of science. In particular, thanks to Nate Otto, Suyash Tripathi, Gianfranco Ingratta, Joshua Barrow, Guadalupe Moreno, Maria Casalas, Everardo Granados, Barbara Yaeggy, and Luis Ruiz, among many more people - too many to list - for all the good times. Tania Brown is another wonderful friend I made at the lab, and I would like to further thank her for cultivating a space at Fermilab (the User's Center) where everyone, self included, feels so welcome and comfortable to be themselves – many positive memories flourished there.

I would further like to express heartfelt thanks to my wonderful boyfriend, Hendrik Pot, who has provided a well of comfort, support, and happy memories over the last two years of this graduate program.

Finally, my deepest gratitude extends to my family members, Holly Roemeling, Ben Roemeling, Bode Roemeling, Julie Ansbaugh, Nathan Ansbaugh, Eleanor Ansbaugh, Samuel Ansbaugh, Nicholas Dyer, Mitch Usher, Mom, and Dad. Even though some of you don't even know what graduate school is (you can't help being born in the 2020's!), you brought so many smiles over the last few years. For the rest of you, your love and support have carried me not just through graduate school, but through every endeavor of my life to date. Thanks for everything.

This work has been funded by the United States Department of Energy, Award Number DE-SC0017740.

DEDICATION

To my parents:

Shelley Dyer and Dr. James Steven Dyer

TABLE OF CONTENTS

ABSTRACT	ii
ACKNOWLEDGEMENTS	iii
DEDICATION	vii
LIST OF TABLES	x
LIST OF FIGURES	xi
Chapter 1 Introduction	1
1.1 The Standard Model	1
1.2 The SBN Program and LArTPC Technology	5
1.3 Problem Statement and Roadmap for this Work	10
Chapter 2 The ICARUS Detector	14
2.1 History	14
2.2 Hardware	15
2.3 Operations and Data Collection	18
Chapter 3 Theory	22
3.1 Higgs Portal Scalars	22
3.1.1 Production	22
3.1.2 Decay	23
3.2 Heavy QCD Axions	24
3.2.1 Production	26
3.2.2 Decay	27
Chapter 4 Simulated Datasets	33
4.1 Simulated Neutrino Interaction Samples	33
4.2 MeVPrt1: Event Generator for BSM Processes within SBN	36
4.2.1 IMesonGen: Meson Extraction	38
4.2.2 IMeVPrt1Flux: Portal Particle Flux Determination	39
4.2.3 IRayTrace: Ray Tracing	40
4.2.4 IMeVPrt1Decay: Portal Particle Decay	43
4.2.5 Event Weight Validation	43
4.3 Simulated HPS and ALP Samples	47
4.4 Downstream Simulation	53
Chapter 5 Event Selection	55
5.1 Cut Flow	55
5.1.1 Preselection	56
5.1.2 Object Cuts	57
5.1.3 Particle Identification	58
5.1.4 Kinematic Cuts	59

5.1.5	Track Length Cut	63
5.2	Exercises to Optimize the Event Selection	63
5.2.1	Event Selection Performance	64
Chapter 6	Systematic Uncertainties	74
6.1	Calculation of Systematic Uncertainties	74
6.2	Sources of Systematic Uncertainty	77
Chapter 7	Data/ Monte Carlo Comparisons in a Far Sideband	85
Chapter 8	Projected Sensitivity	89
Chapter 9	Unblinding Procedure	93
Chapter 10	Conclusions	100
Bibliography		104
Appendix A	ICARUS Detector Controls System	116
A.1	GUIs to Monitor the White Rabbit Switch	118
A.2	Overhaul to Presentation of all ICARUS Slow Controls GUIs	120
Appendix B	Mock Data Tests for Unblinding the Near Sideband and Signal Box	125
B.1	Mock Data Test: Background Only, Throw 1	125
B.2	Mock Data Test: Background Only, Throw 2	126
B.3	Mock Data Test: Background Only, Throw 3	127
B.4	Mock Data Test: Background + ALP Signal with $m_a = 400$ MeV, $f_a =$ 25000 GeV, Throw 1	128
B.5	Mock Data Test: Background + ALP Signal with $m_a = 400$, $f_a = 25000$ GeV, Throw 2	129
B.6	Mock Data Test: Background + ALP Signal with $m_a = 400$, $f_a = 5000$ GeV, Throw 1	130

LIST OF TABLES

2.1	Dates, exposure, and horn current configuration for ICARUS's completed and ongoing physics data runs.	20
4.1	The number of generated and POT-normalized HPS events occurring in ICARUS according to <code>MeVPrtl</code> and the validation approach for three different benchmark scalar masses. The events generated are normalized to 6×10^{20} protons on target. The p-value for χ^2 given n for each model benchmark is also shown here.	46
5.1	Selection efficiency of each type of event, for every consecutive cut in the event selection. Fractions are relative to the preselected samples of events. The same set of cuts is shown in all tables – a different naming convention is used for the ALP and background events for formatting purposes.	71
5.2	Number of Monte Carlo events surviving each cut in the event selection, for each type of event.	72
5.3	Number of background events surviving each cut in the event selection, scaled to an exposure of 2.41×10^{20} protons on target.	73
6.1	Systematics budget in the signal box ($\% \delta = \text{absolute uncertainty} / \text{number of POT-normalized events expected in the signal box}$) for the present analysis, by simulated interaction type. The first column shows the number of Monte Carlo events surviving the event selection, while the subsequent columns show the percent uncertainty on the surviving events, by uncertainty type. The percent uncertainty on the background due to detector modeling effects, 54%, is presently applied to the signal events as well as a temporary placeholder. Similarly, the 10% uncertainty associated with the reconstruction performance of cathode-crossing tracks applied to all events is a temporary placeholder value, as described in the text. These values will be updated by analysis completion.	75
6.2	Specific neutrino interaction variations from the GENIE <code>Reweight</code> package applied in order to determine the overall uncertainty due to neutrino interaction cross sections. The technical name, type of reweighting scheme, and a brief description are provided for each variation. This table is a copy of Table 8.2 in [1].	80
6.3	Percent uncertainty on the background estimation in the signal box, due to specific variations of the detector modeling. The percent uncertainties shown here combine to give the total 54% uncertainty due to detector effects from coh-like events, that is shown in Table 6.1. The first column shows the number of Monte Carlo events surviving the event selection, for each simulated detector variation sample.	83

LIST OF FIGURES

1.1	Schematic depiction of the particles comprising the Standard Model of particle physics, showing the three generations of fermions and the force carrying particles, or bosons. Image is from Wikipedia Commons [2].	2
1.2	Artistic rendering of the detectors of the SBN program (SBND, MicroBooNE, and ICARUS), and the beams they are exposed to (the BNB and NuMI beam), on the Fermilab campus. Image is adapted from [3].	6
1.3	Schematic of the ICARUS detector, highlighting the scintillation light and ionization charge emitted when an incident charged particle interacts with the liquid argon. At the anode plane, the ionization light is measured with an array of photomultiplier tubes. The ionization charge, drifted to the anode by an externally applied electric field, is measured by three planes of parallel wires. The charge induces signal on the first two planes, and is collected on the third plane. Image created by ICARUS collaborator Guadalupe Moreno.	8
1.4	dE/dx , vs. residual range for different types of charged particles traversing LAr, as predicted by the Bethe-Bloch formalism. Residual range is defined as the distance along the track, measured from the particle's ultimate stopping point.	10
1.5	Calibrated dE/dx measurements as a function of residual range, for a sample of stopping muons seen by ICARUS. Figure is from [4].	11
1.6	Fraction of simulated BSM events with dimuon final states that are fully contained in the ICARUS fiducial volume (defined for the present work in Section 5.1.1), as a function of binned true BSM particle total energy. The BSM events are generated according to the Higgs portal scalar and heavy QCD axion models, as described in Sections 3 and 4.	12
2.1	Photograph of the inside of one of ICARUS's cryogenic modules. The vertical metal sheet separating the two time projection chambers is the cathode plane. The longitudinal bars across the top and bottom of the TPC are the field cage. Arrays of photomultiplier tubes are shown at the anode plane of both TPCs. Light flares from the wire planes are also visible at both anode planes. Image is from [5].	17
2.2	Profile of BNB (left) and NuMI (right) ICARUS on- and off-beam triggers as a function of time relative to the start of each beam window, for ICARUS Run 2 of data taking. Image is from [6].	19
3.1	Left: Higgs portal scalar (S) production from kaons. Right: Higgs portal scalar decay to two leptons, in this case muons. The filled dots represent mixing between S and the Higgs boson (h).	23
3.2	Higgs portal scalar branching ratios.	24

3.3	Feynman diagrams for ALP production via kaon decay (left) and mixing with pseudoscalars π^0 , η , or η' (right). The crossed dot is a weak vertex in the chiral effective theory. The solid dot on the left is a shorthand representation of the axion mixing with the pseudoscalars, i.e. the process shown in detail on the right. The mixing is mediated by quarks (solid lines) and gluons (looped lines).	26
3.4	Feynman diagrams for the following ALP decays: $a \rightarrow \gamma\gamma$, $a \rightarrow 3\pi$, $a \rightarrow 2\pi\gamma$, and $a \rightarrow \ell\ell$. The solid dots represent ALP/meson or ρ/γ mixing.	28
3.5	ALP branching fractions for $c_\ell = c_\mu = 0.01$ and $c_1 = c_2 = c_3 = 1$	29
3.6	Values of the axion-muon coupling parameter, c_μ , induced from the other couplings at loop level for the codominance scenario, as modeled by [7].	32
4.1	Diagram of how the MeVPrtl generator fits into SBN Monte Carlo production. Image is from a forthcoming paper describing the MeVPrtl event generator.	37
4.2	Schematic representation of the MeVPrtl framework. Each of the four main stages of the generator are shown in consecutive order at left in black. Specific example implementations that could be run for each stage are shown in colored blocks to the right of the corresponding stage. There is exactly one implementation for the RayTrace stage, WeightedRayTrace, as that physics is common to any dark sector particle originating in the BNB or NuMI beam and traveling to the SBN detectors. Image is from a forthcoming paper describing the MeVPrtl event generator.	38
4.3	Example solutions to the lab frame momentum (4.5) and the ray weight (4.7). The solutions are specified by the parent velocity in the lab frame (β) and the child mass (m) and momentum in the parent rest frame (p'). Plots are from a forthcoming paper describing the MeVPrtl event generator.	42
4.4	Demonstrated agreement between the MeVPrtl generator and the validation approach for three different HPS model benchmarks. The spectra generated by MeVPrtl show consistency with those generated with the validation approach for all three model benchmarks, with p-values for χ^2 given n equal to 0.69, 0.12, 0.12 for the $m_S = 100$ MeV with $\theta_S = 2 \times 10^{-4}$, $m_S = 240$ MeV with $\theta_S = 10^{-5}$, and $m_S = 340$ MeV with $\theta_S = 10^{-5}$ HPS benchmarks respectively.	47
4.5	Energy spectra for some simulated Higgs portal scalar (left) and heavy QCD axion (right) benchmarks. Generally, more massive portal particles that arrive at and decay in ICARUS are more energetic. The energy spectrum for the 540 MeV heavy QCD axions is far more energetic than the other spectra due to the ALP's proximity to the η resonance there, as described in the text.	50
4.6	ALP rates predicted from mesons generated with PYTHIA8 (blue) and Geant4 with (green) and without (orange) the inclusion of secondary proton interactions in the NuMI target.	54
5.1	First four panels: dE/dx in the collection plane, calculated along various regions near the start of the track, for all simulated events that pass the preselection. Bottom Panel: Distribution in lengths for the longest reconstructed shower in each event. Events that do not have any reconstructed showers are assigned a flagged value of -10, attributing to the most populated bin.	66

5.2	Top left: Muon-PID scores for contained candidate muon tracks, for all events that pass the preselection. Top right: θ_{NuMI} determine using range-based momentum reconstruction for contained tracks and the corrected MCS momentum reconstruction for uncontained tracks. Bottom left: Opening angle $\theta_{\mu\mu}$ between the candidate muon tracks. Bottom right: length of the shorter of the branch or trunk tracks for each event. The secondary peak at 1 m is a feature of the preselection requirement that any uncontained track is at least 1 m long. All: Plots are for preselected events from each simulated sample.	67
5.3	θ_{NuMI} is the reconstructed angle between the candidate muons' (solid orange) parent particle momentum (dashed orange) and the vector connecting the NuMI target to the origin of the ICARUS coordinate system (cyan). The ICARUS coordinate system origin is 23cm above the geometric center of the detector. $\theta_{\mu\mu}$ is the opening angle between the reconstructed candidate muon tracks.	67
5.4	$(p_{\text{MCS}} - p_{\text{true}})/p_{\text{true}}$ distributions for contained (left) and exiting (right) muon-matched tracks from the simulated neutrino and BSM samples. The blue curve is made using the default MCS-reconstructed momenta. The orange curve shows the distribution after the correction is applied. Muon tracks with an unphysical MCS momenta (before the correction is applied) are excluded from the distributions. Unphysical values of p_{MCS} result from computational underflows, overflows, and failures of the MCS momentum reconstruction algorithm, and occur less than 3% of the time.	68
5.5	MCS-reconstructed momentum vs. true momentum for candidate muon tracks from the simulated BSM and ν samples. The left panel shows the nominal MCS-reconstructed momentum, while the right panel shows the MCS-reconstructed momentum after the correction from Equation 5.1.4 is applied. The solid black line in both panels is the line $y=x$, which should fit the data if there is no bias in the MCS momentum reconstruction. The dotted black line in the left panel is the best fit to the data using the hottest p_{MCS} bin per p_{true} bin, as described in the text. That best fit line has a slope of 1.18 and a y-intercept of -0.04.	68
5.6	The distribution in θ_{NuMI} for simulated BSM and neutrino interactions in the analysis, according to different methods for calculating θ_{NuMI} . Solid lines are used for the BSM events while dashed lines are used for neutrino events. Black is used to show θ_{NuMI} as calculated from the true momenta corresponding to the trunk and branch candidate muon tracks. Green shows θ_{NuMI} as calculated using range-based momentum reconstruction for contained tracks and the default MCS momentum reconstruction for uncontained tracks. Pink is the same as green, except that the MCS-momentum has been corrected in-post, as described above. Purple shows θ_{NuMI} calculated according to the "weight by length" method.	69
5.7	The distribution in dimuon opening angle, $\theta_{\mu\mu}$, for different simulated BSM model benchmarks. Generally, the opening angle is larger for muon pairs produced in the decay of higher-mass portal particles, as explained in the text. An exception to this trend is seen for muon pairs coming from heavy QCD axions with mass $m_a = 540$ MeV, which are close to the η resonance. This is because events near the eta resonance tend to be much more boosted, as shown in Figure 4.5 and described in Section 4.3.	70

7.1	Data/Monte Carlo comparisons in the far sideband for the differential energy deposition with respect to position, at various points along the reconstructed tracks. Collectively, cuts on these parameters comprise the “proton stub cuts.” The simulated background events are presented as stacked histograms, with hashed gray bands showing the total systematic and statistical uncertainty on the Monte Carlo for each bin. The black points correspond to data, with error bars equal to their statistical uncertainty.	87
7.2	Same as Figure 7.1, for the other cut parameters used in the event selection.	88
8.1	Projected sensitivity to the Higgs portal scalar (left) and heavy QCD axion (right) model spaces. The ALP sensitivity is drawn for the the codominance scenario for the Gauge couplings, and the minimalist scenario of the muon coupling where c_μ is induced by (or “runs with”) the axion decay constant f_a	90
9.1	Top panel: Background events are presented as stacked histograms, categorized by event type. The gray hashed bands show the total uncertainty on the background expectation. The expected signature from one heavy QCD axion model benchmark (mass given in units of MeV; f_a in units of GeV), near the expected edge of sensitivity, is shown in yellow, with uncertainty bands. Bottom panel: A breakdown of the fractional uncertainty on the background, per bin. The legend shows the fractional contribution of each type of uncertainty in the signal box.	94
9.2	Same as top panel of Figure 9.1, but with the curve depicting expected signal also including the underlying background expectation. The plot on the right includes all the same inputs as that on the left, with an additional curve for an ALP model benchmark capable of producing about as many events as possible at ICARUS. Such a signature would entirely swamp the underlying background expectation, even in the near sideband region. ALP masses are presented in units of MeV; f_a are presented in units of GeV.	95
9.3	Rate of occurrence of a particular toy value for the number of background-only data events in the signal box, for 1 million toy throws.	98
A.1	Info-graphic illustrating the key components (hardware, files, applications, and data flows) required to monitor the White Rabbit Switch, with accompanying informative text.	121
A.2	Screenshot of the CSS-displayed WRS expert GUI used by ICARUS White Rabbit switch experts, as described in the text.	122
A.3	CSS GUI for use by ICARUS shifters, featuring a navigation tab to access monitoring GUIs required for performing checklist tasks.	124
A.4	CSS GUI for use by ICARUS hardware experts, featuring a navigation tab to access any and all ICARUS slow controls GUIs.	124
B.1	Number of events per θ_{NuMI} bin for simulated background (blue) and fake data sampled from the simulated background (black), Throw 1. The signal box is the left-most bin, and the sideband region is comprised of the second and third bins.	126
B.2	The profile likelihood probability density from Equation 9.2 for the background-only mock data, Throw 1, with respect to $m_a = 300, 350, 400$, and 450 MeV ALP benchmarks for s. (The values of f_a are given in GeV.)	127

B.3	Exclusion contours for the Higgs portal scalar heavy QCD axion parameter spaces. The blue bands are from the expected sensitivity of the search, as described in section 8. The black curves are drawn in accordance with the background-only, Throw 1 mock data, i.e. an observation of 0 events observed in the signal box.	128
B.4	Same as Figure B.1, for Mock Data Test with Background Only, Throw 2.	129
B.5	Same as Figure B.2, for Mock Data Test with Background Only, Throw 2.	130
B.6	Same as Figure B.3, for Mock Data Test with Background Only, Throw 2.	131
B.7	Same as Figure B.1, for Mock Data Test with Background Only, Throw 3.	132
B.8	Same as Figure B.2, for Mock Data Test with Background Only, Throw 3.	132
B.9	Same as Figure B.3, for Mock Data Test with Background Only, Throw 3.	132
B.10	Same as Figure B.1, but for mock data corresponding to background + ALP Signal with $m_a = 400$, $f_a = 25000$ GeV, Throw 1.	133
B.11	Same as Figure B.2, but for mock data corresponding to background + ALP Signal with $m_a = 400$, $f_a = 25000$ GeV, Throw 1.	133
B.12	Same as Figure B.3, but for mock data corresponding to background + ALP Signal with $m_a = 400$, $f_a = 25000$ GeV, Throw 2.	134
B.13	Same as Figure B.2, but for mock data corresponding to background + ALP Signal with $m_a = 400$, $f_a = 25000$ GeV, Throw 2.	134
B.14	Same as Figure B.1, but for mock data corresponding to background + ALP Signal with $m_a = 400$, $f_a = 5000$ GeV, Throw 1. The range of predicted signal box values from the Likelihood fit to near sideband data (see Figure B.15) is shown in this θ_{NuMI} plot, since it does not overlap with the background-only expectation predicted by the Monte Carlo. Notably, the mock data value in the signal box is consistent with the range of values predicted by the fit to the near sideband data.	135
B.15	Same as Figure B.2, but for mock data corresponding to background + ALP Signal with $m_a = 400$, $f_a = 5000$ GeV, Throw 1.	135

Chapter 1

Introduction

1.1 The Standard Model

Our understanding of the fundamental building blocks of nature is encapsulated in The Standard Model of particle physics, a rich theoretical description describing the interactions of forces and elementary particles. The Standard Model is a quantum field theory in which individual forces and matter particles are described as fields with quantum numbers that dictate how a given particle responds to various underlying mathematical symmetries. Charge, for example, is a quantum number that corresponds to a $U(1)$ symmetry, while spin is a quantum number that determines particles' behavior with respect to $SU(2)$ symmetries.

In the Standard Model, spin-1/2 particles called fermions comprise matter, while integer-spin particles called bosons are force-carriers. There are four force-carrying particles with spin equal to one, called vector or gauge bosons. The gluon mediates the strong force, which is responsible for holding the nuclei of atoms together. As its name suggests, it is the strongest force within the SM. Photons are massless gauge bosons that carry the electromagnetic force. The W and B gauge bosons, with charges ± 1 and 0 respectively, mediate the weak force. All of the fermions interact via the weak force, but as the name “weak” suggests, these interactions are very feeble and subsequently difficult to observe and study. The Higgs boson is the only spin zero (aka “scalar”) particle presently included in the SM, and is responsible for generating the masses of the other Standard Model particles.

Fermions are sub-divided on the basis of whether or not they interact via the strong force. Those that do are called quarks, and there are six of them: the up (u), down (d), charm (c), strange (s), top (t), and bottom (b). The up, down, and charm quarks have charge equal to $2/3$ (in units of the absolute value of electron charge), while the others have charge equal to $-1/3$. Fermions that do not interact via the strong force are called leptons. The electron (e), muon (μ), and tau (τ) leptons

Standard Model of Elementary Particles




Figure 1.1: Schematic depiction of the particles comprising the Standard Model of particle physics, showing the three generations of fermions and the force carrying particles, or bosons. Image is from Wikipedia Commons [2].

have charge equal to -1. Neutrinos, on the other hand, are charge-less and thus do not interact electromagnetically. They interact with the other leptons via the weak force, and there is one type or “flavor” of neutrino per lepton. Electron flavored neutrinos interact only with electrons, and so forth for the other two fermion generations. The Standard Model is depicted schematically in Figure 1.2.

The SM consists of just 17 fundamental particles and forces, and its Lagrangian is expressed with only 19 free parameters (including the elementary particle masses and coupling strengths, for example) to be determined empirically. Despite its brevity, the Standard Model is consistent with a remarkable breadth of observed physics, spanning subatomic to cosmological scales. As

a few examples: the Higgs boson was discovered at the Large Hadron Collider in 2012, exactly as postulated by the SM [8, 9]; thus far the Muon $g-2$ experiment has measured the magnetic momentum of the muon to 127 parts per billion precision, yielding a result still consistent with the SM expectation [10]; and the observed abundance of light elements in the Universe matches the Standard Model prediction as applied to a cooling, expanding Universe a few minutes after the Big Bang (in a period of cosmological evolution referred to as Big Bang Nucleosynthesis) [11–14].

This dissertation would not exist if some shortcomings of the Standard Model did not accompany its many successes. One such example is that the Standard Model describes only baryonic matter, i.e. matter that interacts via strong, electromagnetic, and weak forces. Astrophysical and cosmological evidence suggests, however, that baryonic matter comprises only 20% of the matter in the Universe. The dominating 80% is comprised of matter that interacts only gravitationally, or via the weak force at scales presently undetectable with current technology. To date, model-independent bounds on the mass of this so called “dark matter” range from 10^{-22} eV to $5 M_{\odot}$. ([15] and references therein provide a wholistic review of dark matter.) Further, it is yet unknown whether dark matter is comprised of a single constituent of matter, or several constituents, possibly subject to as much or more rich physics as is contained in our own Standard Model. The former case inspires direct and indirect searches for dark matter particles. The latter case, however, opens an even wider scope of possible research, as there could exist “portal” particles bridging the gap between the SM and dark sector that are not subject to the constraints of dark matter itself. Such extensions to the Standard Model are parameterized into scalar, psuedoscalar, vector, and neutrino portal models. As will be expounded upon later, the Higgs portal scalar extension to the SM is one such possibility of New Physics. Evidence for Higgs portal scalar particles is explicitly investigated by this dissertation.

The contention between the prediction of an electric dipole moment for the neutron within the SM Lagrangian and the non-observance of its dipole moment experimentally [16–20], referred to as the Strong CP problem, posits an additional puzzle thus far unaddressed by the SM. The introduction of a new type of particle called an axion, with gluonic couplings to the Standard

Model, could resolve the Strong CP problem. This dissertation further probes this axionic parameter space, searching for evidence of these axions, which could alleviate this shortcoming of the Standard Model.

Another pocket of mystery within the Standard Model, and the one entirely responsible for the existence of the indispensable tools and technology used for the present work, is on the subject of neutrinos. Neutrinos are the most abundant particles in the Universe, with an estimated 1 trillion passing through any person's human body every second. Despite their abundance, neutrinos are extremely difficult to detect because they are so weakly interacting: even with such a high number of neutrinos passing through your body at any given time, only one neutrino is expected to interact with the matter in the body in a 100 year period [21]. Mapping onto the three-generation paradigm of the other fermions, neutrinos come in three flavors, ν_e , ν_μ , and ν_τ . The weak force only mediates interactions between neutrinos and their respective lepton "partner" of the same generation. Unlike the other fermions, however, neutrinos also exhibit a unique property of "oscillating" between generations, or flavor states, as they propagate through space. This behavior was first observed by the Super-Kamiokande experiment [22] and the Sudbury Neutrino Observatory [23, 24], which led to a Nobel Prize in physics awarded to the corresponding principal investigators in 2015. Oscillation is a quantum mechanical phenomenon that can occur between two non-orthogonal sets of states of a particle. In particular, neutrino oscillation is mediated between the flavor states ν_e , ν_μ , and ν_τ and mass states ν_1 , ν_2 , and ν_3 , which are eigen values of the Hamiltonian. When a particle is in a given flavor state it is in a superposition of mass states, and vice versa. This relationship merits that at least two of the three known neutrinos have non-zero mass (so that distinct mass states are possible), which differs from the baseline expectation from the Standard Model that neutrinos are massless. The origin of neutrino mass remains an open question in theoretical particle physics, and experiments are working to improve estimates of what the neutrino masses are.

The flavor and mass eigen vectors for neutrinos are related via the Pontecorvo-Maki-Nakagawa-Sakata (PMNS) mixing matrix, whose elements are functions of a CP-violating phase δ_{CP} and mixing angles between pairs of neutrinos θ_{12} , θ_{13} , and θ_{23} [25]. Each of these are free parameters

in the Standard Model that must be determined empirically, and there is much ongoing research seeking to do so. δ_{CP} is of particular interest, as charge-parity violation is one of the conditions that must be met in order to explain the matter/antimatter asymmetry in the Universe. δ_{CP} has thus far not been observed in the lepton sector, but a value not equal to zero could help to explain the origin of this imbalance. The matter/antimatter asymmetry is of great anthropic interest since we, as humans made of matter, would simply not exist without it. Increasingly precise measurements of neutrino oscillation place greater constraints on δ_{CP} in the lepton sector.

Experiments targeting solar, atmospheric, reactor, and long baseline accelerator-sourced neutrinos have turned results consistent with the three-neutrino oscillation model described above, and have placed limits on neutrino mass differences, the mixing angles, and δ_{CP} (see, e.g., [26] for a global analysis of such results). Interestingly, however, short-baseline accelerator experiments have raised even more questions in neutrino physics. The LSND and MiniBooNE experiments have both reported anomalous oscillation measurements with excesses of ν_e events [27, 28], that are in contention with the established 3-neutrino paradigm. It has been postulated that an extension of the present Standard Model description of neutrinos to include one or more neutrinos that participate in oscillation but do not interact via the weak force could explain these short baseline anomalies. The potential existence of such “sterile” neutrinos further brings into question whether or not the Standard Model provides a complete picture of the lepton sector. Much research continues today to more precisely quantify the parameters of the PMNS mixing matrix, determine the neutrino masses, and settle whether or not there are more than three neutrinos. Some of this research is being carried out within the Short Baseline Neutrino Program (SBN), located at Fermi National Accelerator (“Fermilab”). The SBN program hosts the ICARUS detector, which is the detector used to collect data for the analysis presented in this dissertation.

1.2 The SBN Program and LArTPC Technology

The Short Baseline Neutrino Program, or SBN, was proposed to investigate possible explanations for the observed LSND and MiniBooNE short baseline anomalies, with neutrino oscillation




Figure 1.2: Artistic rendering of the detectors of the SBN program (SBND, MicroBooNE, and ICARUS), and the beams they are exposed to (the BNB and NuMI beam), on the Fermilab campus. Image is adapted from [3].

measurements of percent-level precision [29]. The program leverages multiple detectors at varying short-baseline positions along Fermilab’s Booster Neutrino Beam (BNB) to measure oscillations. The detectors included in the program are the Short Baseline Neutrino Detector (SBND), located 110 m from the BNB target (source of neutrinos), MicroBooNE, located 470 m from the target, and ICARUS, 600 m downstream of the target. The BNB is created from 8 GeV protons incident on a beryllium target in Fermilab’s accelerator complex. ICARUS and MicroBooNE are exposed to additional neutrinos from Fermilab’s Neutrinos at the Main Injector (NuMI) beam, owing to their small off-axis positions relative to NuMI (5.74° and $^\circ$, respectively). NuMI is higher energy than BNB, resulting from 120 GeV protons incident on a graphite target [30].

SBND, ICARUS, and MicroBooNE are liquid argon time projection chambers (LArTPCs), which are state-of-the-art detectors with an ability to reconstruct particle interactions that occur in the detector with high spatial, temporal, and calorimetric precision. These features of LArTPC technology further enable high-confidence particle identification. Flux systematics and some correlated detector uncertainties (associated with any shared features of various detectors) will be highly constrained in joint oscillation analyses performed with multiple SBN detectors’ data, while other sources of uncertainty related to interaction models and uncorrelated detector effects continue to diminish with ongoing measurement and calibration. These highly constrained systematic uncertainties, combined with the high statistics of neutrinos delivered and to-be delivered by the BNB during the Program’s lifetime, will enable short baseline oscillation and neutrino-argon cross sec-

tion measurements with expected percent-level uncertainties in coming years. Furthermore, SBN also serves as an extensive research and development platform for Fermilab’s flagship experiment, DUNE, whose unprecedented scale will trailblaze further discovery in neutrino physics in coming years [31].

LArTPCs were first developed as a high-precision particle detector at Harvard in 1977 [32]. Fundamentally, a LArTPC consists of a cryostat full of liquid argon, in which particle interactions occur. When charged particles interact with liquid argon, the Ar atoms scintillate and ionize. The cryostat of the LArTPC is immersed in a uniform electric field, so that free, ionized electrons “drift” to one side of the detector, at which there is an array of electrical instrumentation to read-out charge signal associated with 2-dimensional spatial coordinates (lending to the name “time projection chamber”). This instrumentation can consist of wire planes or CCDs, for example. LArTPCs further leverage a photon detection system, which is used to $\sim$ instantaneously collect the scintillation light emitted at the time of the particle interaction. Using the electron drift velocity in the argon (which depends on the argon purity, and can be measured) together with the time of the scintillation and charge readout signals, the third spatial dimension (i.e. along the drift direction) of the particle interaction can be reconstructed.

Figure 1.3 shows a schematic of the ICARUS detector, highlighting these basic working principles of a generic LArTPC. An electric field, equal to 500 V/cm in this example, drifts the ionization charge to the anode plane. In this case, charge signals are collected on three arrays of parallel wires labeled Induction 1, Induction 2, and Collection, and PMTs are implemented for the photon detection system. The details of this figure that are specific to the ICARUS detector are elaborated upon in Section 2.2.

These light and charge signals are emitted all along the incident particle’s trajectory through the LAr, thus enabling its reconstruction in 3 dimensions. The spatial resolution of the reconstructed particle trajectory is set by the resolution of the electrical TPC readout. Furthermore, the amount of charge collected at each point along the trajectory provides a profile of the incident particle’s energy loss as it traverses the Argon. If a particle comes to rest within the active volume of the




Figure 1.3: Schematic of the ICARUS detector, highlighting the scintillation light and ionization charge emitted when an incident charged particle interacts with the liquid argon. At the anode plane, the ionization light is measured with an array of photomultiplier tubes. The ionization charge, drifted to the anode by an externally applied electric field, is measured by three planes of parallel wires. The charge induces signal on the first two planes, and is collected on the third plane. Image created by ICARUS collaborator Guadalupe Moreno.

detector, the charge collected along its entire trajectory can be integrated to determine the particle's kinetic energy at incidence.

Finally, the topological signature of the deposited charge can inform about the particle's identity. Reconstructed topological signatures of ionization charge in a LArTPC generally fall into one of two categories: "tracks" and "showers." Tracks manifest as a continuous line segment of charge that "paints" the incident particle's trajectory through the LAr. All charged particles other than electrons or positrons produce tracks (given they are above the energy threshold required for reconstruction – most LArTPC experiments to date report this threshold to be around 50 MeV). By contrast, showers are more complex topological signatures, wherein the signal resembles a "spray" of charge rather than a discrete, continuous line segment. Showers result from photons, electrons,

and positrons interacting with the argon. These light, electromagnetic particles radiate additional photons which then pair produce to create e^+ and e^- , which radiate more photons, and so forth. The cycle repeats itself until the expendable energy from the incident particle disseminates, resulting in a disperse, conical signature of ionization charge reconstructed as a shower.

A few more tricks are employed to distinguish particle type in LArTPCs, even beyond the classification of reconstructed objects as tracks or showers. The Bethe-Bloch formalism (see Section 34 of [15], and references therein) provides a model for the differential energy loss due to ionization and scintillation processes with respect to distance traversed, dE/dx , of different types of charged particles traversing various materials. Measures of the differential energy loss along reconstructed tracks can be compared to the signature expected from the Bethe-Bloch expectation in argon, for given particle types, in order to determine the most likely specimen. Muons and pions tend to be difficult to distinguish from one another with this method, due to the close proximity in their signatures of differential energy loss, which can become indistinguishable at the level of LArTPCs' reconstruction ability. Protons, however, have a significantly discrepant signature, making them easier to distinguish from muons and pions. Figure 1.4 shows the dE/dx signatures for each of these particles, as predicted from the Bethe-Bloch formalism. Figure 1.5 shows the calibrated dE/dx for a sample of stopping muons in the ICARUS detector, with the theoretical expectation for muons in argon overlaid, demonstrating the LArTPC's ability to use the differential energy loss as a means for particle identification.

For showers, photons can be distinguished from $e^\pm$ based on the presence or absence of what is referred to as a "conversion gap." Interaction vertices can be determined from TPC readout by identifying the intersection of multiple particle's starting point. Electrons and positrons originating from this interaction begin ionizing argon atoms immediately, leading to electromagnetic showers beginning right at the interaction vertex. Photons, on the other hand, do not directly ionize the argon atoms themselves. Rather, the $e^\pm$ pair-produced from the photon are what induce the ionization. Since ionization from the radiation chain is not visible until the initial $e^\pm$ are produced from the photon, the visible shower's starting point is typically displaced from the interaction vertex.




Figure 1.4: dE/dx , vs. residual range for different types of charged particles traversing LAr, as predicted by the Bethe-Bloch formalism. Residual range is defined as the distance along the track, measured from the particle’s ultimate stopping point.

This displacement is referred to as a “conversion gap.” Typically, photon conversion gaps in argon are about 10 cm long. The ability to confidently distinguish $e^\pm$ and photons is a novel ability of LArTPCs – other particle detectors used in neutrino physics before the advent of LArTPCs, such as Kamiokande which is a Cherenkov detector, are much more limited in their ability to discriminate between electrons (or positrons) and photons.

1.3 Problem Statement and Roadmap for this Work

Though the SBN Program and its participating experiments were designed for the primary purpose of advancing neutrino physics, the same characteristics that enable detailed observation and characterization of the neutrino interactions in the LArTPCs also lend themselves to detecting interactions of *any* exotic particle interaction occurring in the detector, that results in visible (i.e. charged) SM particles in the final state. Various phenomenological works have demonstrated that the SBN experiments could be sensitive to a wide spectrum of candidate BSM particles, including Higgs portal scalars [33], heavy QCD axions [34], heavy neutral leptons [35], millicharged particles [36], and inelastic dark matter [37], for example. The anticipated signatures of these particles




Figure 1.5: Calibrated dE/dx measurements as a function of residual range, for a sample of stopping muons seen by ICARUS. Figure is from [4].

are diverse owing to their wide range in possible decay channels and kinematics. They therefore pose a wide breadth of opportunity to discover New Physics. As such, robust programs to search for (non-neutrino) beyond-the-Standard-Model (BSM) physics exist within each of the three SBN experiments, complementing the experiments' primary research objectives in neutrino physics.

This thesis presents work towards a search for BSM particles that could arise in the NuMI beam at Fermilab, and decay into a pair of muons inside the ICARUS detector. Higgs portal scalars and heavy QCD axions (sometimes abbreviated as “ALPs,” for “axion-like particles”) are two such new particles that could produce this signature. In particular, the analysis selects for dimuon candidate events, where one or both of the muons exits the active volume of the detector before coming to rest. The analysis presented here compliments ICARUS's first published physics results [38], which selected for dimuon pairs that are fully contained in the detector, by probing a higher energy phase space of the BSM models of interest: when the BSM particle has more energy, its decay products (i.e. the muons) are more boosted and thus less likely to come to rest before exiting the detector. This trend is exemplified in Figure 1.6, which shows the fraction of dimuon events fully contained in ICARUS as a function of true portal particle energy, for simulated HPSs and ALPs at ICARUS.




Figure 1.6: Fraction of simulated BSM events with dimuon final states that are fully contained in the ICARUS fiducial volume (defined for the present work in Section 5.1.1), as a function of binned true BSM particle total energy. The BSM events are generated according to the Higgs portal scalar and heavy QCD axion models, as described in Sections 3 and 4.

Chapter 2 describes in more detail the legacy, hardware, and data collection of the ICARUS experiment, laying the foundation for how data is collected and leveraged in this analysis. Chapter 3 describes the pertinent theory underlying the Higgs portal scalar and heavy QCD axion models. These portal particles’ theoretical characteristics are what dictate their rate of production and decay in NuMI and the ICARUS detector, and thus affect the degree to which ICARUS is sensitive to the models. In the field of high energy particle physics, it is customary to simulate signatures of signal and background that are expected in data, to 1) verify a working understanding of the detector response systems and data reconstruction tools, 2) develop and optimize your analysis procedures on simulated data, without spoiling the actual results in data before the analysis is mature, and 3) determine limits on specific models set by the data, or to verify that the data is consistent with specific models. This custom calls for detailed modeling of each of the following for both the signal (HPS and ALP events, in this case) and sources of background (cosmic rays and neutrinos): particle production in the neutrino beam or atmosphere, particle propagation to ICARUS, particle decay or interaction with the liquid argon nuclei, and the detector response to interactions taking

place in the LAr. A new event generator used to simulate BSM particles within the SBN Program, called `MeVPrtl` [39]¹, was used to simulate the signal events. The simulation approaches and tools (including `MeVPrtl`) used to generate Monte Carlo samples for the analysis are described in Chapter 4. Chapter 5 describes the event selection used to discriminate between signal and background in the search for uncontained dimuon final states indicative of BSM physics, and Chapter 6 describes the sources of systematic uncertainty inherent to the analysis. Comparisons between Monte Carlo events and data in a sideband region far from the signal box are shown in Chapter 7, demonstrating adequate modeling of neutrino fluxes, neutrino interactions, and the detector response. The projected sensitivity of this search to the Higgs portal and ALP models spaces is presented in Chapter 8, taking into account the event selection performance on simulated data sets and the estimated uncertainties. Chapter 9 presents a plan for unblinding a near sideband and the signal box, as well as interpreting the observed data in those regions in the context of a discovery of evidence or as an exclusion limit, for the HPS and heavy QCD axion particles.

This dissertation presents most of the work required for this analysis. That said, there are a small number of remaining tasks to be completed after the time of writing, that are required to bring the analysis into a fully mature state. These outstanding tasks are noted throughout the text, and summarized in Chapter 10, along with other conclusions.

¹P. Abratenko et al. (ICARUS and SBND Collaborations for the SBN Program) `MeVPrtl`: An Event Generator for Dark Sector Particles in the Short- Baseline Neutrino Program. 5 2026.

Chapter 2

The ICARUS Detector

2.1 History

The ICARUS detector of the SBN experiment is largest operating LArTPC detector ever built, and the culmination of roughly two decades of research and development across various scientific institutions. The ICARUS detector, also referred to as the T600 detector, was first built to observe neutrinos from the CERN to Gran Sasso (CNGS) beam, housed under Grans Sasso National Park at Gran Sasso National Laboratory (LNGS) in Italy. The Italian Institute for Nuclear Physics (INFN) oversaw the commissioning and operation of ICARUS during its “first life” in Europe. Assembly of T600 began in 1999 in Pavia, leading to a successful 3 month test run to observe cosmic rays in 2001 [40]. In 2004, ICARUS was transported to the underground laboratory of LNGS, where it underwent 6 more years of commissioning and was used for its physics analysis [41]. The LNGS laboratory is strategically located underground beneath about 1400 meters of rock, which significantly shields the experiments housed there from exposure to cosmic rays. A 3-year run of data taking in the CNGS beam began in 2010, leading to results that placed limits on ν_e disappearance [42] and refuted the OPERA experiment’s earlier claim of super-luminal neutrinos [43,44]. The successful run also demonstrated the reliability and promise of LArTPC technology for producing high confidence physics results.

From 2013 to 2014, ICARUS was decommissioned and transported to CERN, where upgrades were made to ICARUS for its “second” life as the far detector in the SBN Program at Fermilab [45,46]. ICARUS arrived at FNAL in 2017 and was installed in the designated Far Detector (FD) building along the Booster Neutrino Beamline in 2018. Warm commissioning ensued for two years, followed by cold commissioning and the filling of liquid argon in 2020. ICARUS recorded its first cosmic activity at Fermilab with the cathode and anode ramped to nominal voltage in August of 2020, and the first of its physics-data-taking runs commenced in June 2022 [4].

2.2 Hardware

The ICARUS detector is comprised of two cryogenic modules containing LAr, each 3.6 m wide by 3.9 m tall by 19.6 m long (along the beam direction). The detector coordinate system is a right-handed system defined with $\hat{x}$ going from side-to-side across the detector, $\hat{y}$ pointing vertically up, and $\hat{z}$ along the neutrino beamline. The cryostats are filled with a cumulative total of 760 tons of ultra-pure liquid argon [47]. There are three distinct subsystems of hardware responsible for ICARUS's data collection and readout: the time projection chambers (TPCs) [45, 48, 49], the photon detection system (PDS) [46], and the cosmic ray tagger (CRT) [50].

The TPC system refers to the electronic readout instrumentation used to measure ionization charge from particle interactions in the active volume of the cryostats. Each of ICARUS's two cryostats is further subdivided into two TPCs that share a cathode at the center and have respective anode planes at the outer sides. A uniform electric field of 500 V/cm is applied to each TPC by a surrounding field cage, so that the drift velocity of ionization electrons from the cathode to the anode is uniform with respect to position along the drift coordinate ($\hat{x}$). The drift speed of ionized electrons in ICARUS is 1.6 mm/ μ s for the electric field of 500 V/cm. Given the high average electron lifetime of 15 ms in ICARUS, attributed to by the high purity of the argon in the detector [51], this drift speed is sufficient for detecting ionization charge originating even at the maximum drift distance from the anode plane.

The anode plane for each of ICARUS's four TPCs is comprised of three parallel planes of wires normal to $\hat{x}$. The wires in the first of these planes are parallel to $\hat{z}$, while the wires in the second and third planes are oriented at 60° and -60° relative to $\hat{z}$. The first two planes do not interfere with the ionization charge incident onto the cathode. Instead, the ionization charge merely induces waveforms of current on these wires. As such, the first and second wire planes are referred to as Induction 1 and Induction 2. The third wire plane actually collects the ionized charge and is thus referred to as the Collection Plane.

To understand how the signal from each wire plane comprising the anode is used to reconstruct a 3D image of particle trajectories using the ionization charge, first consider just the collection

plane. The collection plane is oriented parallel to $\hat{z}$ so that each wire is perfectly horizontal and resides at a discrete vertical position along $\hat{y}$. For this plane, charge collected on a wire located at position y corresponds to the origination of charge, or interaction between the incident particle and the LAr, at the same coordinate y . Furthermore, scintillation light collected by PMTs surrounding each TPC (discussed further below) are able to provide a timestamp for particle interactions with the LAr, marking the production of the ionization charge. Together with the known drift speed of the ionized electrons through the argon, this time stamp can be used to determine the incident particle's interaction position along the drift coordinate $\hat{x}$. Thus, the x and y coordinates of the incident particle's trajectory can be reconstructed using just the Collection plane of wires. With the inclusion of an additional wire plane whose pitch is offset relative to the Collection plane, the incident particle's trajectory can be reconstructed in x and a direction with both $\hat{y}$ and $\hat{z}$ components. Geometric computations can be performed to determine the incident particle's trajectory in xyz coordinates using the two sets of coordinates obtained from each wire plane. The inclusion of a third plane of wires is not necessary for this computation, but offers redundancy that serves as a useful cross validation tool.

Each of the anode wire planes consists of wires separated by 3 mm, allowing for an impressive degree of spatial resolution across the entire active volume of the detector.

The ICARUS photon detection system is comprised of 360 Hamamatsu R5912-MOD photo multiplier tubes (PMTs) that are used to detect scintillation light from particle interactions in the LAr. The PMTs are arranged in arrays behind each of the sets of anode wire planes, with an array of 90 PMTs spanning each TPC. The arrangement of the PMTs is such that PMTs are spatially clustered in pairs, as seen in Figure 2.1. The PMTs are sensitive to scintillation signals below a threshold of 100 MeV, with a timing resolution on the order of one nanosecond. Their high light yield and temporal precision enable their use in determining the incident time of particle interactions in the LAr from the $\sim$ instantaneous scintillation flash produced upon interaction. This scintillation signal is used to reconstruct the interaction's drift coordinate, as described above. Furthermore, light collected by the PMTs is used to identify activity in the TPCs that is consistent



Figure 2.1: Photograph of the inside of one of ICARUS’s cryogenic modules. The vertical metal sheet separating the two time projection chambers is the cathode plane. The longitudinal bars across the top and bottom of the TPC are the field cage. Arrays of photomultiplier tubes are shown at the anode plane of both TPCs. Light flares from the wire planes are also visible at both anode planes. Image is from [5].

with the expected arrival times of particles from the BNB or NuMI beams. This signal is used to “trigger” data readout of all detector subsystems, as described in Section 2.3.

Figure 2.1 shows the inside of one of ICARUS’s cryogenic modules during installation, before getting filled with the liquid argon. The field cage, cathode plane separating two TPCs, and arrays of photomultiplier tubes at the anodes are shown. There are also light glares in the photograph on both of the anode planes, which are caused by the TPC wire planes. The individual wires themselves are too fine to depict in the image.

The cosmic ray tagger is ICARUS’s final detector subsystem, and plays the role of identifying activity in the active volume that originated anywhere other than from the neutrino beams. At Fermilab, the ICARUS detector is located at surface level, making the detector significantly more exposed to cosmic rays than it was at LNGS. A 3 m-thick concrete overburden was installed above the detector at Fermilab, which helped to reduce the cosmic flux by about a factor of 2. Still, the rate of cosmics entering the detector even with the overburden in place is 11 kHz, posing a substantial

source of background for searches of physics originating in the neutrino beams. The cosmic ray tagger exists to mitigate this background by identifying the time and location of particles entering and exiting the detector, where these “tags” are used to discriminate between activity originating in the neutrino beams or elsewhere. The cosmic ray tagger system as a whole is composed of the top CRT, side CRT, and bottom CRT, for nearly 4π coverage around ICARUS. Each of these CRT subsystems have slightly different designs, as described in [50], but they all utilize two layers of scintillating plastic embedded with optical fibers. The light signals in these fibers is read out either by silicon photomultipliers (top and side CRTs) or PMTs (bottom CRT).

For the BSM analysis presented in this thesis, all cosmic background interactions are rejected using information from just the TPCs and PMTs, and the primary background consists of neutrino interactions originating in the NuMI beam. Since the primary background has the same CRT signatures (i.e. detector entry points) as the signal events, CRT data is not used in the analysis. On the other hand, cosmics are the primary background when selecting for neutrino interactions, making the CRT an indispensable tool for ICARUS in the context of its primary physics objectives dealing with neutrinos, and especially given its high susceptibility to cosmic rays at the surface level.

2.3 Operations and Data Collection

The BNB and NuMI beams each deliver protons to their respective targets in bursts colloquially referred to as beam “windows” or “spills” that are $1.6 \mu\text{s}$ and $9.6 \mu\text{s}$ long for BNB and NuMI, respectively. For each beam, an “early warning” signal is sent from the Fermilab accelerator complex to the various experiments a fixed time before protons hit the target, providing a time reference for the start of each beam spill [6].

In conjunction with beam early warning signals, PMT light signals are used to trigger the readout of activity in the active LAr that is “in-time” with the arrival of neutrinos from the BNB and NuMI beams. The logic employed for this triggering, referred to as the “majority trigger,” requires the majority of PMT pairs in the array to see a light flash above some programmable




Figure 2.2: Profile of BNB (left) and NuMI (right) ICARUS on- and off-beam triggers as a function of time relative to the start of each beam window, for ICARUS Run 2 of data taking. Image is from [6].

threshold. Majority trigger samples collected in time with the neutrino beam are referred to as “on-beam” data sets. Due to the constant onslaught of cosmic rays, a portion of all on-beam data is from cosmic events that satisfy the majority trigger and just happen to occur in ICARUS during the expected time window for neutrinos. To account for this contamination, the majority trigger is also used to collect samples of events that are out-of-time with neutrinos from the beams, referred to as “off-beam” samples. The difference between on- and off-beam data corresponds to the number of majority triggered events originating in BNB or NuMI, and can be used in physics analyses to study neutrinos and other exotic particles from the beam that would arrive at ICARUS during the beam windows.

Figure 2.2 shows the number of triggers for the BNB and NuMI beams as a function of trigger time (defined with respect to the early warning signal), for on- and off-beam data sets in a given beam window. Six $1.6\mu\text{s}$ “bunches” of protons are delivered to the NuMI target during a single beam spill, and this substructure is discernible in the figure due to the resolution of the triggering system.

A second triggering mechanism, referred to as the “minimum bias trigger,” triggers data collection every N number of beam spills, with no requirement on the amount of scintillation light observed by the PMTs. The minimum bias trigger is used for to evaluate the efficiency of the majority trigger, and could be used to study processes that are out of-time-time with neutrinos from

the beam, such as cosmics themselves or any heavy particles originating in the beam that would travel more slowly than the near-light-speed neutrinos.

Once an event is triggered, TPC data is recorded for 1.5 ms, a time sufficient for even the longest drift distance for ionization charge, PMT data is recorded for a 2 ms time window around the beam spill in order to collect light signals for all cosmics that enter the LAr during the ionization charge drift time [4], and CRT data is recorded for 6 ms. Data collected across the TPC, PDS, and CRT subsystems is time-synchronized with proton beam spill signals from the accelerator complex to within a few ns using a White Rabbit network [52]. The data is then sent to centralized processors, the “EventBuilders,” where it is arranged into “events,” or collections of data spanning all detector subsystems but common to a particular trigger.

ICARUS has completed four physics quality data runs since 2022, with the fifth ongoing. Runs 1-3 collected data from both the BNB and NuMI beams, while Runs 4 onward have only collected BNB data due to technical issues with the infrastructure needed to run NuMI at Fermilab. The first two NuMI runs were collected with the forward horn current (FHC) configuration, referring to the direction of current in the magnetic focusing horns used to focus charged particles in the beam. The third run operated with reverse horn current (RHC). The exposure of each run is quantified by the number of incident protons on the target (POT) throughout the run. Table 2.3 summarizes ICARUS’s physics data runs. The physics analysis presented in this thesis utilizes NuMI data from Runs 1 and 2.

Run #	Dates	BNB (FHC) POT	NuMI (FHC) POT	NuMI (RHC) POT
Run 1	June - July 2022	0.42×10^{20}	0.69×10^{20}	-
Run 2	Dec. 2022 - July 2023	2.13×10^{20}	2.84×10^{20}	-
Run 3	March - July 2024	1.36×10^{20}	-	2.82×10^{20}
Run 4	Dec. 2024 - July 2025	3.63×10^{20}	-	-
Run 5	Oct. 2025 - present	0.80×10^{20}	-	-

Table 2.1: Dates, exposure, and horn current configuration for ICARUS’s completed and ongoing physics data runs.

It is important for ICARUS to collect as much physics-quality data as possible whenever the BNB and NuMI beams are running at Fermilab. This priority drives a policy of continuous monitoring of detector operations, where members of the collaboration are assigned shifts to oversee successful running of all detector and data acquisition subsystems. The designated team of “detector slow controls” (DCS) experts, which I am one of, is responsible for creating and maintaining the software systems used by shifters and hardware experts to monitor operations, and intervene when needed (intervention could entail turning a subsystem off or on, or adjusting the voltage delivered to a PMT, as a couple of examples). Throughout the course of my PhD work, I have made many contributions to ICARUS DCS efforts, some of which are described in Appendix A.2.

Chapter 3

Theory

3.1 Higgs Portal Scalars

Solutions to the Standard Model's shortcomings may lie in a hidden sector weakly coupled to the Standard Model. Models in which these couplings are mediated through renormalizable extensions to the SM are preferred for their analytic simplicity. The Higgs portal is one such model, and can be used to probe light dark matter (eg [53]) and describe reheating in the Early Universe (eg [54]), among other anomalous phenomena.

In the Higgs portal scalar model, the mediator between the SM and dark sector is a neutral scalar particle (S) that mixes with the Higgs boson. The model is described in the mass basis by the interaction Lagrangian

$$\mathcal{L}_{\text{int}} = \sin \theta S \left(\frac{2m_W^2}{v} W^{+\mu} W_{\mu}^{-} + \frac{2m_Z^2}{v} Z^{\mu} Z_{\mu} - \sum_f \frac{m_f}{v} \bar{f} f \right), \quad (3.1)$$

where θ is the mixing angle of the scalar S with the Higgs, v is the Higgs vacuum expectation value, and f are the Standard Model fermions. The parameter space is in this way characterized by the mass of the scalar m_S and the mixing angle θ . Decay products of mediators with masses below the K - π mass difference and mixing angles down to order 10^{-4} can be detected by the SBN LArTPCs. In particular, scalars produced from NuMI kaons and decaying to e^+e^- , $\mu^+\mu^-$, or a pair of pions in ICARUS have the greatest projected sensitivity in the range $m_S \approx 40 - 360$ MeV [33]. Searches at MicroBooNE have demonstrated sensitivity to the Higgs portal scalar parameter space [55].

3.1.1 Production

Kaon decay $K \rightarrow \pi S$ is the dominant production channel for Higgs portal scalars at the SBN Program [33], where K can be either $K^{\pm}$ or K_L^0 . ($K_s^0 \rightarrow \pi^0 S$ is significantly suppressed by the




Figure 3.1: Left: Higgs portal scalar (S) production from kaons. **Right:** Higgs portal scalar decay to two leptons, in this case muons. The filled dots represent mixing between S and the Higgs boson (h).

K_s^0 's short lifetime and CP violation) The leading diagram for this process is a W boson-top quark penguin diagram with partial width

$$\Gamma(K \rightarrow \pi S) = \frac{\theta^2 \tau_K}{16\pi m_K} \left| \frac{3V_{CKM} m_t^2 m_K^2}{32\pi^2 v^3} \right|^2 \frac{2p_S}{m_K}, \quad (3.2)$$

where v is the Higgs vacuum expectation value and V_{CKM} are elements from the CKM Quark-Mixing Matrix [33]. $V_{CKM} = V_{td} \times V_{ts}$ for $K^\pm$, and $\text{Re}(V_{td}) \times V_{ts}$ for K_L^0 ; the Wolfenstein parameterization for these constants [15] is used for the present analysis. The Feynman diagram for HPS production via kaon decay is shown in the left panel of Figure 3.1.

Scalars may also be produced via B meson decays and proton bremsstrahlung processes [54]. However, very few B mesons are produced at the SBN detectors given the program's relatively low energy flux, e.g. relative to collider experiments. Furthermore, [33] finds that proton bremsstrahlung processes do not improve SBN's sensitivity beyond kaon-decay production channels. This analysis thus neglects both of these production modes.

3.1.2 Decay

For the scalar mass ranges that the SBN Program is sensitive to, Higgs portal scalars may decay into lepton or pion pairs. The partial decay width into leptons is given by

$$\Gamma(S \rightarrow l^+ l^-) = \theta^2 \frac{m_l^2 m_S}{8\pi v^2} \left(1 - \frac{4m_l^2}{m_S^2} \right)^{3/2}, \quad (3.3)$$




Figure 3.2: Higgs portal scalar branching ratios.

and is represented in the right panel of Figure 3.1. Scalar decay to pions is calculated with chiral perturbation theory, resulting in

$$\Gamma(S \rightarrow \pi^0 \pi^0) = \theta^2 \frac{|G_\pi(m_S^2)|^2}{32\pi v^2 m_S} \left(1 - \frac{4m_\pi^2}{m_S^2}\right)^{1/2}. \quad (3.4)$$

The hadronic form factor $G_\pi(x)$ can be calculated numerically from [56–59]. $\Gamma(S \rightarrow \pi^+ \pi^-) = 2 \times \Gamma(S \rightarrow \pi^0 \pi^0)$, since the charged pions comprise two components of the pions’ isospin triplet. The branching ratios for the Higgs portal scalars are shown in Figure 3.2.

3.2 Heavy QCD Axions

Heavy QCD axions are motivated as a solution to the Strong CP problem in the Standard Model, which is a tension between an electron dipole moment for the neutron predicted by the Standard Model term $\mathcal{L} \supset \theta \frac{g^2}{32\pi^2} G^{\mu\nu} \tilde{G}_{\mu\nu}$ and non-observation of the neutron dipole moment experimentally [16–20]. The QCD axion field a , defined by a gluonic coupling, resolves the Strong CP problem by shifting the effective potential of θ to zero: $\theta \rightarrow \theta_{\text{eff}} = (\theta + a/f_a)$. f_a is a necessary model parameter referred to as the axion decay constant, with units of energy (units of GeV should be assumed anywhere units are omitted in this text).

The Peccei-Quinn (PQ) mechanism [60, 61] achieves this by coupling the QCD axions to gluons. In this scenario, upon confinement, the QCD axion generates a potential with a CP-conserving minimum so that the solution is achieved dynamically. The Peccei-Quinn mechanism leads to “light” QCD axions with mass $m_a \lesssim 5.7$ MeV. A shortcoming of the PQ solution to the CP problem, however, is that the $U(1)_{PQ}$ PQ symmetry is subject to breaking by anomalous gravitational effects. This symmetry breaking in turn reintroduces otherwise Planck-suppressed terms which drive the potential back away from the CP-conserving minimum needed to address the Strong CP problem. This predicament is coined the “quality problem” [62–66].

Heavy ($m_a \sim \text{GeV}$) QCD axions provide a solution to the Strong CP problem while avoiding the quality problem that arises in the PQ mechanism (see [34] and references therein). When $m_a \gtrsim 100$ MeV, the axions are too heavy to be dictated by strong dynamics, leaving the minimum potential impervious to gravitationally-induced CP violating terms. Notably, heavy QCD axions are too unstable to be considered a viable dark matter candidate.

While gluonic couplings are the definitional component of QCD axions, couplings to the other SM gauge bosons and to leptons are also allowed. The complete Lagrangian for these axions is as follows:

$$\mathcal{L}_{\text{axion}} = \frac{1}{2}(\partial_\mu a)^2 - \frac{1}{2}m_a^2 a^2 + \mathcal{L}_{\text{gauge}} + \mathcal{L}_{\text{lepton}} \quad (3.5)$$

with

$$\mathcal{L}_{\text{gauge}} = c_3 \frac{\alpha_3}{8\pi f_a} a G \tilde{G} + c_2 \frac{\alpha_2}{8\pi f_a} a W \tilde{W} + c_1 \frac{\alpha_1}{8\pi f_a} a B \tilde{B} \quad (3.6)$$

and

$$\mathcal{L}_{\text{lepton}} = \sum_{\ell=e,\mu,\tau} c_\ell \frac{\partial_\mu a}{2f_a} \bar{\ell} \gamma^\mu \gamma_5 \ell. \quad (3.7)$$

For simplicity in the analysis, I consider a “codominance” scenario in which all the gauge couplings are equal ($c_1 = c_2 = c_3$), and assume c_μ is the only non-zero leptonic coupling. These assumptions are commonly employed for other QCD axion searches as well (e.g. [34, 38]).




Figure 3.3: Feynman diagrams for ALP production via kaon decay (left) and mixing with pseudoscalars π^0 , η , or η' (right). The crossed dot is a weak vertex in the chiral effective theory. The solid dot on the left is a shorthand representation of the axion mixing with the pseudoscalars, i.e. the process shown in detail on the right. The mixing is mediated by quarks (solid lines) and gluons (looped lines).

3.2.1 Production

Axions can be produced at the SBN Program from kaon decay or through mixing with pseudoscalar mesons π^0 , η , and η' , as depicted in 3.3. Production via kaon decay dominates below the $m_K - m_\pi$ axion mass threshold [7]. At higher masses, ALPs are primarily produced via mixing with pseudoscalar mesons π^0 , η , and η' . Reference [67] points out that the Primakoff process $\gamma + A \rightarrow a + A$ in which a photon scatters off of a nucleus to produce an ALP would be the dominant production mechanism if the ALP's gluonic coupling were subdominant to its photonic coupling, but this scenario is not considered in the present analysis.

Production of ALPs in kaon decays is described as follows. CP violation suppresses the production of axions (which are CP odd) from K_L^0 , leaving just the following relevant branching fractions for axion production via kaon decay [68, 69]

$$\text{Br}(K^\pm \rightarrow a\pi^\pm) = \frac{\tau_{K^\pm}}{\tau_{K_S^0}} \frac{2f_\pi^2 c_3^2}{f_a^2} \left(\frac{m_K^2 - m_a^2}{4m_K^2 - 3m_a^2 - m_\pi^2} \right)^2 \sqrt{\frac{\lambda(1, m_\pi^2/m_K^2, m_a^2/m_K^2)}{1 - 4m_\pi^2/m_K^2}} \times \text{Br}(K_S^0 \rightarrow \pi^+\pi^-) \quad (3.8)$$

and

$$\text{Br}(K_S^0 \rightarrow \pi^0 a) \approx \frac{\tau_{K_S^0}}{\tau_{K^\pm}} \text{Br}(K^\pm \rightarrow \pi^\pm a), \quad (3.9)$$

where τ_K are the respective kaon lifetimes and $f_\pi \approx 130$ MeV is the pion decay constant. Production from K_S^0 is highly suppressed by the ratio $\tau_{K_S}/\tau_{K^\pm}$ in 3.9, so to good approximation, only charged kaons can significantly produce the axions.

For heavy QCD axion production via mixing, the rate of mixing with a meson m is equal to the magnitude-squared of the relevant mixing angle, $|\theta_{am}|^2$. The relevant mixing angles are given by [34, 70]

$$\begin{aligned}\theta_{a\pi} &= \frac{1}{6\sqrt{2}} \frac{f_\pi}{f_a} \frac{m_a^2}{m_a^2 - m_\pi^2}, \\ \theta_{a\eta} &= \frac{1}{2\sqrt{3}} \frac{f_\pi}{f_a} \frac{m_a^2 - 4m_\pi^2/9}{m_a^2 - m_\eta^2}, \\ \theta_{a\eta'} &= \frac{1}{2\sqrt{6}} \frac{f_\pi}{f_a} \frac{m_a^2 - 16m_\pi^2/9}{m_a^2 - m_{\eta'}^2}.\end{aligned}\tag{3.10}$$

3.2.2 Decay

Partial widths for heavy QCD axion decay channels that are relevant to ICARUS are enumerated below. Feynman diagrams for some of the processes are shown in 3.4. Fig. 3.5 shows the branching ratios for ALPs with model parameters $c_\ell = c_\mu = 0.01$ and $c_1 = c_2 = c_3 = 1$.

1. Photons:

The width for axion decay into two photons is given by [34, 70, 71]

$$\Gamma_{a \rightarrow \gamma\gamma} = \frac{\alpha^2 |c_\gamma|^2 m_a^3}{256\pi^3 f_a^2},\tag{3.11}$$

where α is the fine structure constant and c_γ is the axion coupling to photons. c_γ is enhanced by all other axion couplings at loop level [34, 72]:

$$\begin{aligned}c_\gamma = c_3 \left(-1.93 + \frac{1}{3} \frac{m_a^2}{m_a^2 - m_\pi^2} + \frac{8}{9} \frac{m_a^2 - 4m_\pi^2/9}{m_a^2 - m_\eta^2} + \frac{7}{9} \frac{m_a^2 - 16m_\pi^2/9}{m_a^2 - m_{\eta'}^2} \right) \\ + \frac{5}{3} c_1 + c_2 + 2 \sum_{\ell=e,\mu,\tau} c_\ell B(4m_\ell^2/m_a^2),\end{aligned}\tag{3.12}$$




Figure 3.4: Feynman diagrams for the following ALP decays: $a \rightarrow \gamma\gamma$, $a \rightarrow 3\pi$, $a \rightarrow 2\pi\gamma$, and $a \rightarrow \ell\ell$. The solid dots represent ALP/meson or ρ/γ mixing.

where $B(x) = 1 - xf(x)^2$ and

$$f(x) = \begin{cases} \sin^{-1}\left(\frac{1}{\sqrt{x}}\right) & \text{if } x \geq 1 \\ \frac{\pi}{2} + \frac{i}{2} \log\left(\frac{1+\sqrt{1-x}}{1-\sqrt{1-x}}\right) & \text{if } x < 1. \end{cases} \quad (3.13)$$

The factor or -1.93 is an empirical value derived from a fit of chiral perturbation theory and lattice QCD to decay and scattering data [72].

2. Hadronic decay: 3π

Analytic results from [73] are used to model hadronic axion decays in this work, yielding amplitudes for the two possible $a \rightarrow 3\pi$ channels equal to




Figure 3.5: ALP branching fractions for $c_\ell = c_\mu = 0.01$ and $c_1 = c_2 = c_3 = 1$.

$$\begin{aligned}
A(a \rightarrow 3\pi^0) &= \frac{m_\pi^2}{f_a f_\pi} \left[\langle a\pi \rangle - \delta_I \left(\frac{1}{\sqrt{3}} + \sqrt{2}S_{\eta\pi} + S_{\eta'\pi} \right) \left(\sqrt{2}\langle a\eta \rangle + \langle a\eta' \rangle \right) \right. \\
&\quad \left. + \sqrt{3}\delta_I \frac{m_\pi^2 - 2m_\eta^2}{m_\pi^2 - 4m_\eta^2} \left(\sqrt{2}S_{\eta\pi} + S_{\eta'\pi} \right) \right], \\
A(a \rightarrow \pi^+\pi^-\pi^0) &= \frac{1}{3f_a f_\pi} \left[(3m_{\pi^+\pi^-}^2 - m_a^2 - 2m_\pi^2) \langle a\pi \rangle \right. \\
&\quad - \delta_I m_\pi^2 \left(\frac{1}{\sqrt{3}} + \sqrt{2}S_{\eta\pi} + S_{\eta'\pi} \right) \left(\sqrt{2}\langle a\eta \rangle + \langle a\eta' \rangle \right) \\
&\quad + \delta_I \frac{m_\pi^2 - 2m_\eta^2}{m_\pi^2 - 4m_\eta^2} \times \\
&\quad \left. \left(\sqrt{3}m_\pi^2 \left(\sqrt{2}S_{\eta\pi} + S_{\eta'\pi} \right) - 3m_{\pi^+\pi^-}^2 + m_a^2 + 3m_\pi^2 \right) \right].
\end{aligned} \tag{3.14}$$

$\delta_I = (m_u - m_d)/(m_u + m_d) \approx 1/3$ is the isospin symmetry breaking factor, computed from the up and down quark masses. $S_{P_1 P_2} = M_{P_1 P_2}/(m_{P_1}^2 - m_{P_2}^2)$, where m_{P_i} is the mass of meson P_i (π , η , or η'), $M_{\pi\eta}^2 = -m_\pi^2 \sqrt{2/3}$, and $M_{\pi\eta'}^2 = -m_\pi^2/\sqrt{3}$. The brackets $\langle \rangle$ are used to denote expectation values, which are given by

$$\begin{aligned}
\langle a\pi \rangle &\approx \frac{\delta_I}{2} \frac{m_a^2}{m_a^2 - m_\pi^2}, \\
\langle a\eta \rangle &\approx \left[\frac{m_a^2}{\sqrt{6}} - \frac{m_\pi^2}{2\sqrt{6}} \right] \frac{1}{m_a^2 - m_\eta^2}, \\
\langle a\eta' \rangle &\approx \left[\frac{m_a^2}{2\sqrt{3}} - \frac{m_\pi^2}{\sqrt{3}} \right] \frac{1}{m_a^2 - m_{\eta'}^2}.
\end{aligned} \tag{3.15}$$

The decay widths are derived from these amplitudes as

$$\Gamma_{a \rightarrow 3\pi} = \frac{k}{2Sm_a} \int |A(a \rightarrow 3\pi)|^2 d\Phi_3. \tag{3.16}$$

Here, $k = 2.7$ is an empirical factor introduced to bring the equivalent analytic result for the $\eta' \rightarrow 3\pi$ decay width into agreement with experimental data, and S is a symmetry factor equal to one for the $a \rightarrow \pi^+\pi^-\pi^0$ channel, and equal to six for the $a \rightarrow \pi^0\pi^0\pi^0$ channel. For three body decays of the axion with mass m_a into particles with masses m_1, m_2 , and m_3 , the phase space differential is

$$d\Phi_3 = \frac{dm_{12}^2 dm_{23}^2}{128\pi^3 m_a^2}. \tag{3.17}$$

3. Hadronic decay: $\pi^+\pi^-\gamma$

Axion decays to this final state undergo an intermediary ρ vector resonance, so that the overall decay width is given by

$$\Gamma_{a \rightarrow \gamma(\rho \rightarrow \pi^+\pi^-)} = \frac{3\alpha m_a^3}{2^{11}\pi^6 f_a^2} \int dm_{\pi\pi}^2 |12\pi \text{BW}_\rho(m_{\pi\pi}^2) \langle a\rho\rho \rangle|^2 m_{\pi\pi}^2 \left[1 - \frac{m_{\pi\pi}^2}{m_a^2}\right]^3 \left[1 - \frac{4m_\pi^2}{m_{\pi\pi}^2}\right]^{3/2}, \tag{3.18}$$

where BW_ρ is the Breit-Wigner function for the ρ resonance. For axion masses above 1.4 GeV (which are not relevant for ICARUS), the amplitude in the integrand requires an addi-

tional form factor. The expectation value is evaluated as

$$\langle a\rho\rho\rangle = \frac{\langle a\eta\rangle}{\sqrt{6}} + \frac{\langle a\eta'\rangle}{\sqrt{12}}, \quad (3.19)$$

and [74] provides the Breitt-Wigner function as

$$\begin{aligned} \text{BW}_\rho(s) &= \text{BW}^{\text{GS}}(s, m_\rho, \Gamma_\rho), \\ \text{BW}^{\text{GS}}(s, m_\rho, \Gamma_\rho) &= \frac{1 + d(m_\rho)\Gamma_\rho/m_\rho}{m_\rho^2 - s + f(s, m_\rho, \Gamma_\rho) - im_\rho\Gamma(s, m_\rho, \Gamma_\rho)}, \\ \Gamma(s, m_\rho, \Gamma_\rho) &= \Gamma_\rho \frac{s}{m_\rho^2} \left(\frac{\beta_\pi(s)}{\beta_\pi(m_\rho^2)} \right)^3, \\ \beta_\pi(s) &= \sqrt{1 - 4m_\pi^2/s}, \\ d(m_\rho) &= \frac{3}{\pi} \frac{m_\pi^2}{k^2(m_\rho^2)} \ln \left(\frac{m_\rho + 2k(m_\rho^2)}{2m_\pi} \right) + \frac{m_\rho}{2\pi k(m_\rho^2)} - \frac{m_\pi^2 m_\rho}{\pi k^3(m_\rho^2)}, \\ f(s, m_\rho, \Gamma_\rho) &= \frac{\Gamma_\rho m_\rho^2}{k^3(m_\rho^2)} [k^2(s)(h(s) - h(m_\rho^2)) + (m_\rho^2 - s)k^2(m_\rho^2)h'(m_\rho^2)], \\ k(s) &= \frac{1}{2} \sqrt{s} \beta_\pi(s), \\ h(s) &= \frac{2}{\pi} \frac{k(s)}{\sqrt{s}} \ln \left(\frac{\sqrt{s} + 2k(s)}{2m_\pi} \right), \\ h'(s) &= \frac{dh(s)}{ds}. \end{aligned} \quad (3.20)$$

4. Leptonic decays: $a \rightarrow \ell\ell$

Tree level couplings between heavy QCD axions and leptons result in the decay width

$$\Gamma_{a \rightarrow \ell\ell} = \frac{c_\ell^2 m_a m_\ell^2}{8\pi f_a^2} \sqrt{1 - \frac{4m_\ell^2}{m_a^2}}. \quad (3.21)$$

Fortuitously, leptonic decays can also arise from the other gauge couplings at loop-level, qualifying leptonic decays as a viable search channel even in minimalist model scenarios that do not explicitly include the leptonic Lagrangian term (see [7] and references therein).




Figure 3.6: Values of the axion-muon coupling parameter, c_μ , induced from the other couplings at loop level for the codominance scenario, as modeled by [7].

Under such a scenario, which will sometimes be referred to as the “running” (leptonic or muonic) coupling, the induced leptonic decays depend on the relative strengths of the gauge couplings c_1 , c_2 , and c_3 . [7] describe this scenario assuming values of c_1 , c_2 , c_3 defined at an energy scale of $4\pi f_a$. Starting with the initial condition that $c_\ell = 0$ at this scale, they then solve the one-loop renormalization group equations to run the couplings down to the scale of the top quark mass $m_t \approx 172.69 \text{ GeV}$ [15]. A non-zero c_ℓ is generated by this evolution. Below the top mass, the running of c_ℓ stops at one loop, so it is considered to be the low-scale effective value of this coupling. Figure 3.6 depicts the running coupling c_μ associated with different values of f_a .

Chapter 4

Simulated Datasets

This analysis leverages simulated datasets of HPS, heavy QCD axion, neutrino, and cosmic activity in the active volume of ICARUS’s TPCs. Specific Monte Carlo samples are generated for the BSM signal events and for neutrino interactions arising in the NuMI neutrino beam, and cosmic ray signatures generated with `CORSIKA` [75] are added to events in each of these datasets to represent the cosmics that would enter the ICARUS TPCs during the TPC electronics’ readout time of a neutrino- or portal particle-triggered event.

4.1 Simulated Neutrino Interaction Samples

The simulation of neutrino interactions in ICARUS relies on the accurate prediction of off-axis flux (quantity, flavor, and kinematics) of neutrinos from the NuMI beam, and on how neutrinos interact with argon nuclei in the ICARUS TPCs.

To produce the NuMI beam, hydrogen atoms are stripped of their electrons, accelerated to 120 GeV, and directed at a thin graphite target approximately 1 meter long [30]. The high energy interactions between the accelerated protons and graphite produce hadrons, which are forward-focused by magnetic focusing horns located downstream of the target. The direction of current in the horns determines whether negatively or positively charged hadrons are focused, where the FHC (RHC) configuration focuses positively (negatively) charged hadrons. Kaons and pions produced in the target subsequently decay into neutrinos and other particles. Most resulting neutrinos tend to be collimated due to the kinematic boost they inherit from their focused parent hadrons, but since they are charge-less and thus impervious to the magnetic focusing horns, they do tend to disperse to off-axis angles. The neutrino beam is also dispersed due to some of the hadrons decaying to neutrinos before reaching the horns. This dispersion is what enables ICARUS to see a fraction of NuMI-beam-generated neutrinos and other neutral particles.

The MINOS experiment, which was the first experiment designed to study neutrinos from NuMI, developed the first simulation for the NuMI neutrino flux, which is based on the *Geometry and Tracking* (GEANT4) simulation toolkit [76–78]. Hadrons are transported and interacted with material in the NuMI target and decay pipe in this simulation. The output of the simulation is stored in the DK2NU [79] flux file format, which preserves kinematic information for every particle produced throughout the chain from protons interacting with the graphite to the eventual neutrinos. An additional software toolkit called the *Package to Predict the Flux* (PPFX) [80,81], also developed by MINOS, compares the flux output by GEANT4 with data and applies corrections and estimates of systematic uncertainties associated with hadron interaction modeling in order to bring the simulation and data into agreement. These corrections and uncertainties are determined using experimental cross section data, and can be propagated downstream for analyses that utilize the NuMI beam (such as the one presented in this thesis).

The simulation approach was later upgraded by MINERvA for use in the MINERvA and NO ν A experiments (both of which also lie on the NuMI beamline axis). More recently, ICARUS has dedicated effort towards characterizing the *off-axis* NuMI flux it sees. For ICARUS’s off-axis angle of 5.75° relative to the NuMI beam, much of the beamline geometry downstream of the target is negligible. On the other hand, there are other materials, such as the presence of large concrete slabs positioned along the vector from the NuMI target to ICARUS that must be accounted for, as this dense material attenuates hadrons and neutrinos alike. Furthermore, some of the operating conditions of the NuMI beam have changed since its inception and use in older experiments, such as the beam spot size. These changes are also accounted for in ICARUS’s dedicated study and characterization of the off-axis flux. [82] describes in detail all that goes into the off-axis NuMI neutrino flux simulation for ICARUS, and the associated systematic uncertainties.

With a simulation of the NuMI neutrino flux as seen by ICARUS in place, charged current (CC) and neutral current (NC) interactions between the neutrinos and the LAr in ICARUS are then modeled. Neutral current interactions are mediated by the Z boson, and in these interactions the incident neutrino persists to the final state. On the contrary, charged current interactions are medi-

ated by the Z^+ or Z^- boson, and the incident neutrino is replaced by its lepton partner in the final state. CC interactions, which dominate over NC interactions as background for the present analysis since they are more easily mis-reconstructed as dimuon final states, can be further classified into the following types:

- **Charged current quasi-elastic and meson exchange current (CC QE + MEC):** no pions or other hadrons are produced in these interactions. For QE interactions, the ν interacts with a single neutron, resulting in a lepton and one nucleon in the final state. In MEC interactions, the neutrino interacts with two nucleons, resulting in a lepton and two nucleons in the final state.
- **Charged current resonant scattering (CC RES):** a neutrino and nucleon interact to produce a lepton and a Δ resonance, which is an excited state of the nucleon. The Δ subsequently decays back into a nucleon and typically one or more pions.
- **Charged current deep inelastic scattering (CC DIS):** typically occurring at higher energies ($E_\nu \sim 10$ GeV), the neutrino interacts with an individual quark in the nucleon, producing a shower of other hadrons (pions and nucleons) in the final state that accompany the final state lepton.
- **Charged current coherent pion production (CC COH or CC Coh π):** a small energy transfer from the neutrino to the nucleon coherently produces a pion. The final state is one lepton and one pion.

The ICARUS collaboration uses the software package `GENIE` to generate ν -Ar interactions, though there are a suite of alternate event generators that can also be used (e.g. `GiBUU` [83], `NEUT` [84], `NuWro` [85], and `ACHILLES` [86]). Each of these event generators have their own implementation of neutrino-Ar cross sections, resulting in slightly different outputs for the distribution of observable final state particles, given the same input of neutrino flux. These differences between generators, as well as comparisons to experimental data, are used to determine systematic uncertainties associated with the relevant cross section modeling.

For the present analysis, I use a generic sample of NuMI neutrinos generated for ICARUS in the production campaign MCP2023A. As will be presented in Chapter 5, the leading background in the search for dimuon final states comes from CC Coh π interactions, which have one pion and one muon in the final state. Despite being the primary background for the present analysis, these interactions are very rare relative to other neutrinos, comprising less than 1% of the neutrino interactions in the NuMI beam. To bolster the statistics of the leading background, I use an additional sample of “coh-like” events, which includes CC Coh π interactions and other neutrino interactions with topologically similar final states. Generated neutrino events must satisfy the following to be included in the coh-like sample: has ≥ 2 generated particles of μ , $\pi^\pm$, or $\Sigma^\pm$, and does not have a proton generated with kinetic energy > 50 MeV. The sample of coh-like events has $> 50\times$ the statistics of coh-like events expected for data, compensating for the rarity of these interactions relative to others in NuMI which otherwise leads to insufficient statistics for studying background in the signal region.

4.2 MeVPrt1: Event Generator for BSM Processes within SBN

Members across the Beyond Standard Model working groups of ICARUS and SBND, self included, have collaborated to develop an in-house event generator for simulating the production and decay or interaction of 10s-100s MeV mass scale portal particles in the SBN detectors. This event generator is called MeVPrt1, and is described in Ref. [39]. The MeVPrt1 generator is used to make the simulated samples of Higgs portal scalars and heavy QCD axions produced via mixing used in this analysis. It simulates the production of BSM particles via meson decay or mixing in BNB or NuMI, the BSM particles’ decay or interaction inside ICARUS or SBND, and the associated weight of each simulated event. At present the MeVPrt1 event generator includes implementations for Higgs portal scalars, heavy QCD axions, and heavy neutral leptons. Event weights are needed to recover the correct probability of the event occurring in actual data. The positioning of the MeVPrt1 generator relative to the other simulation and reconstruction codes used within SBN is shown in Figure 4.1.




Figure 4.1: Diagram of how the MeVPrtl generator fits into SBN Monte Carlo production. Image is from a forthcoming paper describing the MeVPrtl event generator.

Event generation is split into four stages in MeVPrtl: extraction of mesons from provided flux files (such as those containing simulated neutrinos, as discussed in the previous section), determination of the portal particle flux, tracing portal particle trajectories from their production point to the detector, and determination of the decay in the detector. Each of the four stages has its own interfacing C++ module: IMesonGen, IMeVPrtlFlux, IRayTrace, and IMeVPrtlDecay, respectively. The modular organization of this event generator eliminates redundancy in simulating different BSM particles, where there is shared overlap in the pertinent phenomenology. A diagram depicting the four stages of event generation, along with example implementations for each stage, is presented in Figure 4.2.

The total weight corresponding to each event produced by MeVPrtl is a product of contributions from the various stages of generation, normalized by the number of protons on target (POT) generated versus expected:

$$w_{\text{total}} = w_{\text{flux}} \times w_{\text{ray}} \times w_{\text{decay}} \times \frac{\text{POT}_{\text{expected}}}{\text{POT}_{\text{generated}}}. \quad (4.1)$$

w_{flux} , w_{ray} , and w_{decay} are discussed in the following subsections, in the context of the module in which each term is calculated. $\text{POT}_{\text{generated}}$ is the number of simulated protons on target and $\text{POT}_{\text{expected}}$ is the number of protons on target expected for the experiment being studied.

The analysis presented here uses signal events that have been de-weighted by the event generator. MeVPrtl de-weights events according to an accept-reject algorithm, where it calculates the maximum possible weight m , and then rejects simulated events with weight w according to the probability $1 - w/m$.




Figure 4.2: Schematic representation of the MeVPrt1 framework. Each of the four main stages of the generator are shown in consecutive order at left in black. Specific example implementations that could be run for each stage are shown in colored blocks to the right of the corresponding stage. There is exactly one implementation for the RayTrace stage, WeightedRayTrace, as that physics is common to any dark sector particle originating in the BNB or NuMI beam and traveling to the SBN detectors. Image is from a forthcoming paper describing the MeVPrt1 event generator.

MeVPrt1 is fully integrated with the SBN Program’s LArSoft-based [87] simulation and reconstruction pipeline. The generator is hosted by the SBN Program on its GitHub repository for common analysis software, SBNSoftware/sbncode. The generator’s code can be found in sbncode/EventGenerator/MeVPrt1. It is accessible on GitHub [88].

The physics for which each stage of the MeVPrt1 generator is responsible is discussed in Subsections 4.2.1, 4.2.2, 4.2.3, and 4.2.4. Subsection 4.2.5 presents work that I completed to validate the event weighting procedure implemented in the MeVPrt1.

4.2.1 IMesonGen: Meson Extraction

The IMesonGen class sources the parent particles, usually mesons, that will decay to or mix with the BSM particles, and returns their information in the same format as the neutrino flux inputs currently utilized in the SBN Program’s LArSoft [87] analysis pipeline. The IMesonGen class also records the number of protons on target simulated for all of the provided parent particles.

`MeVPrt1` encodes all needed meson information in `LArSoft`'s `simb::MCFlux` [87, 89, 90] object, which is used to handle neutrino fluxes, for easier incorporation into the standard neutrino simulation code. The `simb::MCFlux` objects also have fields reserved for neutrino information; these fields are ignored by `MeVPrt1`.

Information is extracted differently for long- and short-lived parent particles. To simulate long-lived mesons K and $\pi^\pm$, the generator relies on the neutrino flux simulation in NuMI [30] and BNB [91]. Kaons are extracted from `g4numi` [82] and `g4bnb` [92–94] `Geant4` simulation outputs in `Dk2Nu` [79] and `BooNE` format, respectively.

The short-lived π^0 , η , and η' parent particles do not produce significant numbers of neutrinos, and so they are not kept in standard neutrino flux files, and therefor cannot be sourced from `g4numi` or `g4bnb` flux files as the kaons can. The `MeVPrt1` generator instead sources the π^0 , η , and η' particles from a sample simulated with `PYTHIA8` [95], where primary proton-nucleon interactions in the NuMI target are simulated at the center-of-mass energy $E_{\text{cm}} = \sqrt{m_p^2 + m_N^2 + 2m_N E_p}$. Here, m_N is the mass of the target nucleon and m_p and E_p are the mass and energy of the incident proton, respectively. The resulting E_{cm} is 15.123 GeV for proton-proton collisions and 15.133 GeV for proton-neutron collisions in the target. Nuclear motion, nuclear final state interactions, and secondary interactions are neglected in this approximation of E_{cm} . The protons and nucleons are collided with the `SoftQCD` process, which simulates the low energy QCD interactions primarily responsible for light meson production. The four-momenta of the resulting π^0 , η , and η' particles is recorded for use in generating the portal particle fluxes in the next module. Notably, the full NuMI beam geometry downstream of the target is not included in the `PYTHIA8` simulation used to source these mesons, but it is reasonable to neglect such a detailed simulation since the π^0 , η , and η' are unlikely to propagate far after production.

4.2.2 `IMeVPrt1Flux`: Portal Particle Flux Determination

The `IMeVPrt1Flux` class handles the production of the BSM particles via meson decay or mixing, and sets the multiplicative weight associated with the production, w_{flux} .

For Higgs portal scalars, the decay width of kaons into scalars presented in Equation 3.2 is used. For simulated heavy QCD axions, the mixing angles in Equation 3.10 are used.

For portal particles originating from long-lived mesons (M), the flux weight takes the form

$$w_{\text{flux}} = w_{\text{imp}} \times \frac{\text{Br}(M \rightarrow S + X)}{\text{Br}(M \rightarrow \nu + X)}. \quad (4.2)$$

The `g4bnb` and `g4numi` codes force mesons to decay into neutrinos for increased statistics of neutrino interactions. This unnatural bias is accounted for by dividing by the SM branching fraction of meson to neutrinos. $\text{Br}(M \rightarrow S + X)$ is then included to account for the probability that the meson will decay to the portal particle of interest. The importance weight [92] w_{imp} is a statistical weight output by `g4bnb` or `g4numi`, reflecting the modification of used cross sections (`g4bnb`) or encoding the survival probability of a parent particle (`g4numi`). For portal particles that come from short-lived mesons, $w_{\text{flux}} = \text{Br}(M \rightarrow S + X)$. If it is impossible for a given meson to produce a portal particle, the meson is skipped.

4.2.3 IRayTrace: Ray Tracing

The `IRayTrace` class takes as input the portal particle flux, and outputs a direction such that the particle impinges ICARUS or SBND, along with the associated weight w_{ray} .

The geometry of the detector is modeled as a box containing the active volume(s) of the detector. The dimensions are specified in a coordinate system where z is the BNB axis and y is vertical. In the case of a NuMI-generated flux, the detector is then rotated to a coordinate system where z is the NuMI beam axis. The coordinates of the center of the detector are specified in the relevant beam frame for the given simulation, with θ being the polar angle with respect to the beam axis and ϕ being the azimuthal angle around the beam axis. This reference frame is referred to as the lab frame below.

To ensure efficient event generation, simulated particles are forced to follow a trajectory that impinges a detector volume. This is done by sampling the portal particles' direction from all those that would intersect the detector, and accounting for the probability that the portal particle would

actually take that path in the ray weight, w_{ray} . For portal particles that are produced isotropically in the decay of a parent particle, this probability is equal to

$$w_{\text{ray}} = \frac{1}{4\pi} d\Omega', \quad (4.3)$$

where $\Omega' = (\theta', \phi')$ is the solid angle spanned by the detector, in the rest frame of the parent particle. A change of variables from the parent particle rest frame to the lab frame is applied, so that the solid angle of the detector can be expressed in lab frame coordinates $\Omega = (\theta, \phi)$:

$$w_{\text{ray}} = \frac{1}{4\pi} \left| \frac{d\Omega'}{d\Omega} \right| d\Omega. \quad (4.4)$$

The `IRayTrace` module performs this integration by sampling particle trajectories from Ω , and applying the Jacobian $|d\Omega'/d\Omega|$ as weight on the event. This weight is known for neutrinos ($d\Omega'/d\Omega = E^2/E'^2$), where the mass of the child particle is approximately zero [96]. The Jacobian for massive child particles (i.e. the portal particles `MeVPrtl` simulates) is computed as follows.

The momentum and energy of the child particle in the parent rest frame (p' and E') are determined at the time of production via two-body decay. Furthermore, the child particle direction in the lab frame (θ, ϕ) is known, since it is sampled from the solid angle spanned by the detector. The child particle momentum in the lab frame (p) and the polar angle in the parent particle rest frame (θ') are expressed in terms of the known quantities as

$$p_{\pm} = \frac{E' \beta \cos \theta / \gamma \pm \sqrt{E'^2 / \gamma^2 - m^2 + m^2 \beta^2 \cos^2 \theta}}{1 - \beta^2 \cos^2 \theta} \quad (4.5)$$

and

$$\cos \theta'_{\pm} = \frac{-E' \beta \gamma \sin^2 \theta \pm \cos \theta \sqrt{E'^2 / \gamma^2 - m^2 + m^2 \beta^2 \cos^2 \theta}}{p' \gamma (1 - \beta^2 \cos^2 \theta)}, \quad (4.6)$$

where β and γ are the usual quantities from special relativity given by the velocity and boost of the parent particle in the lab frame, and m is the child particle mass. Unlike the massless (neutrino)




Figure 4.3: Example solutions to the lab frame momentum (4.5) and the ray weight (4.7). The solutions are specified by the parent velocity in the lab frame (β) and the child mass (m) and momentum in the parent rest frame (p'). Plots are from a forthcoming paper describing the `MeVPrt1` event generator.

case where there is exactly one solution to the Jacobian, there are two analytic solutions to the Jacobian for massive child particles

$$\begin{aligned} \left(\frac{d\Omega'}{d\Omega} \right)_{\pm} &= \frac{E'}{E_{\pm}} \frac{p_{\pm}^2}{p'^2} \frac{dp}{dp'_{\pm}}, \\ \left(\frac{dp}{dp'} \right)_{\pm} &= \frac{p'}{E'\gamma} \frac{\beta \cos \theta \pm E' / \sqrt{E'^2 - \gamma^2 m^2 (1 - \beta^2 \cos^2 \theta)}}{1 - \beta^2 \cos^2 \theta}. \end{aligned} \quad (4.7)$$

One, both, or neither of these solutions are allowed for a given event, depending on the relative kinematics of the parent and child particles. In the case that $v'_{\text{child}} > v_{\text{parent}}$, only $(d\Omega'/d\Omega)_{-}$ is physically possible. In the case that $v'_{\text{child}} < v_{\text{parent}}$, whether both or neither of $(d\Omega'/d\Omega)_{\pm}$ are allowed is dictated by the relative directions of the particles. Both solutions are allowed (of which the generator randomly applies one) below a critical angle defined by $\cos \theta_c = \sqrt{m^2 - E'^2/\gamma^2}/(m\beta)$, and neither solution is allowed for polar angles greater than θ_c . The critical angle is the maximum polar angle the particle can have and still arrive at the detector, when $v'_{\text{child}} < v_{\text{parent}}$. Example numerical solutions for the ray weight (i.e. Jacobian) and child particle lab frame momentum under this scenario are shown in Figure 4.3.

As is evident from the figure, the ray weight corresponding to the critical angle is infinite, meaning that de-weighting with an accept-reject algorithm based on the maximum possible weight cannot be performed. To overcome this issue, `MeVPrt1` provides an alternative computation that only forces the ϕ angular variable to hit the detector, while still picking θ' randomly. In this case, the Jacobian contributing to the ray weight is trivial ($d\phi'/d\phi = 1$), but the weight is scaled by the fraction of successful throws of θ' that impinge the detector, from a (configurable) set number of throws for a single event. The event generation is still reasonably efficient with only one angular variable forced to intersect the detector volume ($\gtrsim 0.1\%$ of throws are successful). Events that cannot possibly impinge the detector are dropped during the ray tracing stage of generation.

4.2.4 `IMeVPrt1Decay`: Portal Particle Decay

The `IMeVPrt1Decay` class also takes the portal particle flux as input, and outputs the decay information of the particle. This information includes the portal particle's decay position and time, the list of particles it decays to, and the kinematics of the portal particle and its decay products. The associated probability for the particle decaying to the final state of interest in a detector, w_{decay} , is given by

$$w_{\text{decay}} = (e^{-L_{\text{enter}}/\gamma\beta\tau} - e^{-L_{\text{exit}}/\gamma\beta\tau}) \times \text{Br}(S \rightarrow \text{final state}) \quad (4.8)$$

where γ and β are the Lorentz boost and velocity for the portal particle (S) with respect to the lab frame and τ is its proper lifetime. L_{enter} and L_{exit} correspond to the points along the particle's trajectory where, if the particle is stable enough, it would enter and exit the detector. "Final state" refers to the decay products being simulated (e.g. $\mu^+\mu^-$ for an HPS decaying into the dimuon channel).

4.2.5 Event Weight Validation

`MeVPrt1`'s determination of event weights has been validated in the context of HPS generation at ICARUS by an independent approach described by the following formalism.

For a given event, the differential probability of HPS decay in terms of its decay position relative to the parent kaon is given by

$$dP = \frac{1}{4\pi} \frac{1}{\gamma v \tau} e^{-r/\gamma v \tau} dr d\cos\theta' d\phi', \quad (4.9)$$

where r is now the length of the HPS trajectory, v is the HPS speed, and τ is the HPS lifetime. In terms of ICARUS detector coordinates x_d, y_d, z_d this becomes

$$dP = \frac{1}{4\pi} \frac{1}{\gamma v \tau} e^{-r(x_d, y_d, z_d)/\gamma v \tau} \left| \frac{\partial(\cos\theta', \phi')}{\partial(\cos\theta, \phi)} \right| \left| \frac{\partial(r, \cos\theta, \phi)}{\partial(x_d, y_d, z_d)} \right| dx_d dy_d dz_d. \quad (4.10)$$

The first Jacobian is the same as $|d\Omega'/d\Omega|$ in 4.7. The second takes the form $1/r^2$, following the usual transformation from a spherical to Cartesian system. For any given event, the scalar's trajectory $r = r(x_d, y_d, z_d)$ is well defined by the scalar and kaon decay positions.

Event weights are defined such that a sum over the event weights in a sample of weighted events, to be normalized by the POT as explained in 4.1, gives the number of expected interactions in the detector (N'). This statement assumes the POT is exactly known. Monte Carlo generation performs this sum via implicit integration of the likelihood of each event, over the detector volume. Leaving the POT-normalization to be accounted for downstream, the kaon importance weight, relevant branching fractions, and differential probability of decay from 4.10 contribute to the likelihood of each event:

$$N' = \int_{\text{detector volume}} w_{\text{imp}} \frac{\Gamma(K^\pm \rightarrow \pi^\pm S)}{\Gamma(K^\pm \rightarrow \nu X)} \frac{\Gamma(S \rightarrow \text{final state})}{\Gamma(S \rightarrow \text{any})} \frac{dP}{dx_d dy_d dz_d} dx_d dy_d dz_d. \quad (4.11)$$

In accordance with the mean value theorem, the integral can be approximated as a product of the detector volume and a sum over the form of the integrand for each generated event, normalized by the number of generated events in the sample (N):

$$N' = \frac{V_{\text{detector}}}{N} \sum_{i=1}^N w_{\text{imp}i} \frac{\Gamma(K^\pm \rightarrow \pi^\pm S)}{\Gamma(K^\pm \rightarrow \nu X)} \frac{\Gamma(S \rightarrow \text{final state})}{\Gamma(S \rightarrow \text{any})} \left(\frac{dP}{dx_d dy_d dz_d} \right)_i. \quad (4.12)$$

Further leaving N to be re-accounted for downstream, and distributing the factor of the detector volume into the form of the integrand for each event, the total weight of a given event is given by

$$w_{\text{total}} = w_{\text{imp}} \frac{\Gamma(K^\pm \rightarrow \pi^\pm S)}{\Gamma(K^\pm \rightarrow \nu X)} \frac{\Gamma(S \rightarrow \text{final state})}{\Gamma(S \rightarrow \text{any})} V_{\text{detector}} \frac{dP}{dx_d dy_d dz_d}. \quad (4.13)$$

This formalism lends itself to sampling HPS decay positions from the ICARUS detector volume. Indeed, this is done in practice for the validation approach. As long as the scalar's velocity in the kaon rest frame, v'_S , is greater than the kaon's lab-frame velocity v_K , sampling scalar decay positions in the detector volume will result in a finite maximum weight. With the scalar decay position fixed entirely by sampling, though, it remains possible that the singularity in $\left| \frac{d\Omega'}{d\Omega} \right|$ is encountered when $v_K > v'_S$ (as discussed in 4.2.3). In the validation approach, event de-weighting is thus performed with an accept/reject algorithm that uses the maximum *generated* weight in the sample, in contrast to the maximum possible weight as is done by `MeVPrt1`. While one must be extraordinarily unlucky to sample and generate an event with kinematics landing exactly on the singularity, at high statistics, generating at least one event near the singularity (i.e. on the asymptote to the infinite weight) is not impossible. Since encountering a maximum weight near the singularity any time the generator is run is left to chance, leading to unpredictable and possibly very poor de-weighting efficiency, this method is reserved for validation purposes only.

To validate that `MeVPrt1` is correctly weighting HPS events, the expected event rates and energy spectra of HPS generated by `MeVPrt1` have been compared to samples generated with the validation approach. The validation was performed for three different model benchmarks: $m_S = 100$ MeV with $\theta_S = 2 \times 10^{-4}$, $m_S = 240$ MeV with $\theta_S = 10^{-5}$, and $m_S = 340$ MeV with $\theta_S = 10^{-5}$. For the 240 MeV and 340 MeV mass benchmarks, samples consisting only of scalars that decay to muons are compared. This comprises the vast majority of 240 MeV scalars,

since pion final states are not kinematically allowed there and the lepton decay width scales with the lepton mass squared. Decay to pions dominates for 340 MeV scalars, but a comparison for the fraction of scalars that decay to two muons verifies correct handling of relative branching fractions in the generators. 4.1 shows the number of de-weighted Monte Carlo simulation events for each sample used for the comparison, along with the corresponding projected number of events to occur in ICARUS for 6×10^{20} protons on target from the NuMI beam (roughly equal to ICARUS's exposure to NuMI during its first three runs of data collection). The number of events expected according to each generation method agrees within the statistical uncertainty, thus validating the approach employed in `MeVPrtl`. The POT-normalized spectra are shown in 4.4. The statistics and event rates in the legends in 4.4 reflect only the portions of each sample included in the plotted range (0–4 GeV). A χ^2 comparison is computed for each pair of spectra over the plotted range, so that the number of degrees of freedom n for each fit is equal to the number of non-empty bins. The spectra generated by `MeVPrtl` show consistency with those generated with the validation approach for all three model benchmarks, with p-values for χ^2 given n shown in 4.1.

HPS model benchmark (final state)	Events (<code>MeVPrtl</code>)	Events (Validation)	p-value
$m_S = 100 \text{ MeV}, \theta_S = 2 \times 10^{-4} (ee)$	21.41	22.09	0.69
$m_S = 240 \text{ MeV}, \theta_S = 10^{-5} (\mu\mu)$	3.19	3.19	0.12
$m_S = 340 \text{ MeV}, \theta_S = 10^{-5} (\mu\mu)$	5.05	5.32	0.12

Table 4.1: The number of generated and POT-normalized HPS events occurring in ICARUS according to `MeVPrtl` and the validation approach for three different benchmark scalar masses. The events generated are normalized to 6×10^{20} protons on target. The p-value for χ^2 given n for each model benchmark is also shown here.




Figure 4.4: Demonstrated agreement between the MeVPrt1 generator and the validation approach for three different HPS model benchmarks. The spectra generated by MeVPrt1 show consistency with those generated with the validation approach for all three model benchmarks, with p-values for χ^2 given n equal to 0.69, 0.12, 0.12 for the $m_S = 100 \text{ MeV}$ with $\theta_S = 2 \times 10^{-4}$, $m_S = 240 \text{ MeV}$ with $\theta_S = 10^{-5}$, and $m_S = 340 \text{ MeV}$ with $\theta_S = 10^{-5}$ HPS benchmarks respectively.

4.3 Simulated HPS and ALP Samples

For the present analysis, simulated samples of HPS and mixing-produced heavy QCD axions are generated with MeVPrt1. Because the kinematics and thus final state observables of hypothetical portal particles depend on the portal particle mass, it is important to simulate samples across various mass benchmarks for each BSM model being probed. From there, predicted event rates for a given benchmark mass can be reweighted in post to reflect a change in the other (non-mass)

model parameters. This analysis utilizes simulated samples for the following HPS and heavy QCD axion model benchmarks:

Higgs portal scalars:

- $m_S = 220 \text{ MeV}, \theta_S = 10^{-5}$
- $m_S = 240 \text{ MeV}, \theta_S = 10^{-5}$
- $m_S = 260 \text{ MeV}, \theta_S = 10^{-5}$
- $m_S = 280 \text{ MeV}, \theta_S = 10^{-5}$
- $m_S = 300 \text{ MeV}, \theta_S = 10^{-5}$
- $m_S = 320 \text{ MeV}, \theta_S = 10^{-5}$
- $m_S = 340 \text{ MeV}, \theta_S = 10^{-5}$

Mixing-produced ALPs:

- $m_a = 300 \text{ MeV}, f_a = 10^6 \text{ GeV}, c_\mu = 0.01$
- $m_a = 350 \text{ MeV}, f_a = 10^6 \text{ GeV}, c_\mu = 0.01$
- $m_a = 400 \text{ MeV}, f_a = 10^6 \text{ GeV}, c_\mu = 0.01$
- $m_a = 450 \text{ MeV}, f_a = 10^6 \text{ GeV}, c_\mu = 0.01$
- $m_a = 500 \text{ MeV}, f_a = 10^6 \text{ GeV}, c_\mu = 0.01$
- $m_a = 540 \text{ MeV}, f_a = 10^6 \text{ GeV}, c_\mu = 0.01$
- $m_a = 600 \text{ MeV}, f_a = 10^6 \text{ GeV}, c_\mu = 0.01$
- $m_a = 650 \text{ MeV}, f_a = 10^6 \text{ GeV}, c_\mu = 0.01$
- $m_a = 700 \text{ MeV}, f_a = 10^6 \text{ GeV}, c_\mu = 0.01$

For the generated ALPs, the other lepton couplings are assumed to be zero, and the codominance scenario is assumed for the gauge boson couplings, with $c_1 = c_2 = c_3 = 1$. The even 50 MeV spacing in ALP mass benchmarks is interrupted where 540 MeV ALPs are generated in lieu of 550 MeV ALPs. This was done intentionally to avoid the η resonance encountered in hadronic decay modes of the axion, which blow up the axion decay width (see Equations 3.14 and 3.15).

Figure 4.5 shows the BSM particle energy spectra for a subset of the simulated signal benchmarks, demonstrating the general trend that more massive portal particles arriving at ICARUS tend to have a greater total energy. The energy spectrum shown for 540 MeV heavy QCD axions is much more boosted than the other benchmarks due to its proximity to the η resonance encountered in hadronic decay modes of the axion. This and other resonances result in infinite spikes in the total decay width of the axion, thus infinitely shortening its lifetime and mean free path. As such, an axion with a mass near a resonance will need to be much more boosted in order to arrive at ICARUS before decaying, lending to the shifted energy distribution seen for 540 MeV ALPs.

For Higgs portal scalar events, a reweighting from a particular mixing angle θ to a new mixing angle θ' is given by

$$w_{\theta \rightarrow \theta'} = \frac{\theta'^2}{\theta^2} \times \frac{e^{-\theta'^2 L_{\text{enter}}/\theta^2 L_S} - e^{-\theta'^2 L_{\text{exit}}/\theta^2 L_S}}{e^{-L_{\text{enter}}/L_S} - e^{-L_{\text{exit}}/L_S}}, \quad (4.14)$$

where L_S is the mean lab frame decay length of the scalar with mixing angle θ . L_{enter} and L_{exit} are the points along the scalar's trajectory where, if the particle is stable enough, it would enter and exit ICARUS.

A reweighting for ALPs from parameters f_a and c_μ to new parameters f'_a and c'_μ is given by

$$w_{f_a, c_\mu \rightarrow f'_a, c'_\mu} = \frac{f_a^2}{f_a'^2} \frac{e^{-L_{\text{enter}}/L'_a} - e^{-L_{\text{exit}}/L'_a} \text{Br}(a \rightarrow \mu\mu)'}{e^{-L_{\text{enter}}/L_a} - e^{-L_{\text{exit}}/L_a} \text{Br}(a \rightarrow \mu\mu)}, \quad (4.15)$$

where L_a and $\text{Br}(a \rightarrow \mu\mu)$ are the axion mean lab frame decay length and branching fraction for decay to two muons, respectively, given f_a and c_μ . L'_a and $\text{Br}(a \rightarrow \mu\mu)'$ are the same parameters, but given f'_a and c'_μ . Written in terms of their un-primed counterparts, these are




Figure 4.5: Energy spectra for some simulated Higgs portal scalar (left) and heavy QCD axion (right) benchmarks. Generally, more massive portal particles that arrive at and decay in ICARUS are more energetic. The energy spectrum for the 540 MeV heavy QCD axions is far more energetic than the other spectra due to the ALP’s proximity to the η resonance there, as described in the text.

$$L'_a = \frac{f_a^2}{f_a'^2} \frac{L_a}{(1 - \text{Br}(a \rightarrow \mu\mu) + c_\mu'^2 \text{Br}(a \rightarrow \mu\mu)/c_\mu^2)} \quad (4.16)$$

and

$$\text{Br}(a \rightarrow \mu\mu)' = \frac{c_\mu'^2}{c_\mu^2} \frac{\text{Br}(a \rightarrow \mu\mu)}{(1 - \text{Br}(a \rightarrow \mu\mu) + c_\mu'^2 \text{Br}(a \rightarrow \mu\mu)/c_\mu^2)}. \quad (4.17)$$

The HPS samples produced by `MeVPrt1` can be used as-is. The ALPs, on the other hand, require further processing or consideration.

First, `MeVPrt1` only produces ALPs arising from mixing with neutral mesons, so ALPs originating in kaon decay need to be derived from the (kaon-produced) HPS samples. This is performed

with a reweighting, where Higgs portal scalar events are rescaled in accordance with the heavy QCD axion model according to

$$w_{\theta \rightarrow f_a, c_\mu, c_3} = \frac{\text{Br}(K \rightarrow a) \text{Br}(a \rightarrow \mu\mu)}{\text{Br}(K \rightarrow S) \text{Br}(S \rightarrow \mu\mu)} \frac{(e^{-L_{\text{enter}}/L_a} - e^{-L_{\text{exit}}/L_a})}{(e^{-L_{\text{enter}}/L_S} - e^{-L_{\text{exit}}/L_S})}. \quad (4.18)$$

L_a can be written in terms of L_S as

$$L_a = \frac{\Gamma(S \rightarrow \text{any})}{\Gamma(a \rightarrow \text{any})} L_S. \quad (4.19)$$

The $\text{HPS} \rightarrow \text{ALP}$ reweighting is only applied to HPSs with charged kaon parents, since axions from K_L^0 are CP-violating-suppressed. This model reweighting is valid since the decay kinematics are identical for the processes $K \rightarrow \pi S$ and $K \rightarrow \pi a$, whenever $m_a = m_S$.

All parameters needed for the aforementioned reweights are output by the `MeVPrt1` event generator.

Furthermore, the modeling for mixing-produced ALPs incorporated into `MeVPrt1` does not consider all the relevant physics. `MeVPrt1` generates ALPs produced via mixing with neutral mesons (π^0 , η , and η') by applying a corresponding mixing angle weight to meson events that are generated with `PYTHIA8`. The `PYTHIA8` simulation is not a full simulation of protons interacting within the target, but rather a simulation of individual primary protons interacting with protons and neutrons at rest. The output therefor only includes neutral mesons whose parent particle was a primary proton, excluding the extras that would come from secondary interactions within the target or NuMI decay pipe. Such a detailed simulation of these neutral mesons is not available in the official Monte Carlo fluxes for the NuMI beam as they are for charged mesons, which are of primary interest for neutrino physics. This shortcoming of the `PYTHIA`-based `MeVPrt1` generation is addressed by performing a partial beam simulation for the neutral mesons using `Geant4`. This simulation was performed by Joshua Berger and described below.

Events are simulated in `Geant4`, where a proton smashes into a box of graphite with dimensions $940 \times 15 \times 4$ mm, approximately equal to the dimensions of the NuMI target. The density of the graphite in the simulation is adjusted to that of the NuMI target, 1.78 g/cm^3 .

From there, π^0 , η , and η' mesons produced for each proton hitting this target were identified. Two sets of the generated mesons are then considered. The first is a subset that includes only mesons classified as coming from primary interactions, where the meson's parent particle is the incident proton. These interactions are consistent with those generated in the `PYTHIA8` simulation. The second set of events includes *all* mesons produced as the particles are propagated through the target. This set includes mesons produced in secondary interactions within the target, in addition to those produced in primary interactions, and is a more realistic prediction of the actual meson flux that could contribute to ALPs at ICARUS.

To ensure the kinematics of the mesons closely match those that would be observed by ICARUS, only mesons traveling at an angle of between 5.557° and 5.956° with respect to the beam axis, replicating the solid angle spanned by the ICARUS detector in the standard beamline geometry, were used in each set of events. The `Geant4` simulation was generated to correspond to 10^5 POT, and the energy spectrum of each meson, π^0 , η , and η' , was recorded. For comparison with the baseline simulation, meson spectra for 10^6 proton interactions using `PYTHIA8` with the same settings as in `MeVPrtl` were also generated.

Next, the total event rate for ALPs as a function of energy was computed, for the three scenarios: `PYTHIA8` (without secondaries), `Geant4` without secondaries, and `Geant4` with secondaries. The total event rate for the ALPs was calculated from the meson spectra in each case by

$$N(E \text{ bin}) = \theta_\pi^2 N_\pi(E \text{ bin}) + \theta_\eta^2 N_\eta(E \text{ bin}) + \theta_{\eta'}^2 N_{\eta'}(E \text{ bin}), \quad (4.20)$$

where “ $E \text{ bin}$ ” denotes a given bin of meson energy and the mixing angles θ are as given in [?]. These rates are presented per POT, per GeV in Figure 4.6, spanning 20 total energy bins from 135 MeV to 22.143 GeV. The error bars for the vertical axis come from the statistical uncertainty in

each bin, according to

$$\sigma_N^2(E \text{ bin}) = \theta_\pi^4 N_\pi(E \text{ bin}) + \theta_\eta^4 N_\eta(E \text{ bin}) + \theta_{\eta'}^4 N_{\eta'}(E \text{ bin}). \quad (4.21)$$

Central value reweights to account for `MeVPrtl`’s failure to include neutral mesons produced in secondary interactions in the target are derived from the ratio of the ALP rates as predicted by the `Geant4` simulations with and without secondaries. These are shown by the green and orange data in Figure 4.6, respectively. Logarithmic binning is used for this derivation to mitigate the effect of prohibitively low statistics in the higher energy bins. In rare cases where the ALP energy is greater than the upper bound of the highest energy bin, the reweight derived for the highest energy bin is applied to that event.

As is discussed in Section 6, a comparison between the *primary* spectra from `PYTHIA8` and `Geant4` is used to determine a one-sided flux uncertainty for ALP production.

4.4 Downstream Simulation

The above subsections describe how the signal BSM and background neutrino events are generated in ICARUS, but further simulation is needed to propagate those interactions into observable signals ready for analyzing. This is performed with two more steps in the ICARUS simulation pipeline: the “g4” simulation and the “detsim” simulation. The g4 simulation uses the `Geant4` software in order to determine how the final state particles interact with the liquid argon. The Bethe-Bloch formulation, for example, is built into the g4 simulation stage in order to predict the energy deposited by the final state particles as they traverse through the LAr. (Generally speaking, `Geant4` can be used to model the propagation of any particle through any medium, which is why it is also used to generate the neutrino flux, e.g. starting with protons interacting with the graphite target.) The output of ICARUS’s g4 simulation is used as input to its detsim simulation, which models the detector response to the simulated charge and light signals, giving an expectation for what the actual collected data will look like.




Figure 4.6: ALP rates predicted from mesons generated with PYTHIA8 (blue) and Geant4 with (green) and without (orange) the inclusion of secondary proton interactions in the NuMI target.

Chapter 5

Event Selection

Physical parameters such as interaction vertices, primary and secondary particles, particle types, and track and shower start point, length, direction, and energy are reconstructed from the simulated and collected data using the reconstruction software Pandora [97]. The present analysis employs a cut-and-count approach to signal identification, where cuts on these reconstructed quantities are used to distinguish signal and background. The chosen event selection has been designed to suppress the background as much as possible, so that a BSM signal can be identified as any significant excess. The selection for uncontained dimuon tracks results in ~ 0 cosmic background events and ~ 2 neutrino background events expected for an exposure of 2.41×10^{20} POT. The event selection was designed, tested, and optimized using simulated data sets, and will be applied to collected data to probe the HPS and ALP models.

The selection criteria employed in this analysis can be classified into categories of Preselection, Object Cuts, Particle Identification, Kinematic Cuts, and Track Length. The event selection, along with descriptions of the parameters used to define cuts, is presented in the following subsection. Efforts to optimize the event selection are described in Section 5.2, followed by a summary of the event selection's performance in Section 5.2.1.

5.1 Cut Flow

This subsection describes the reconstructed quantities and corresponding cuts used to discriminate between signal and background events. Figures 5.1 and 5.2 show the distributions of these quantities by simulated interaction type: 'BSM,' ' ν ,' or 'cosmic.' These interaction classifications are based on the type of particle that deposits the most energy in the reconstructed slice. The 'BSM' category lumps together interactions from all the simulated HPS and ALP model benchmarks (that do not get classified as 'cosmic'). The ' ν ' category includes simulated events from the nominal and coh-like neutrino samples, also excluding those that do not get classified as 'cosmic.'

5.1.1 Preselection

The preselection for this analysis requires a reconstructed interaction vertex in the fiducial volume, two reconstructed muon-like tracks starting within 10 cm of the reconstructed interaction vertex, and that at least one of the muon-like tracks has a reconstructed end point not in the fiducial volume.

The fiducial volume is defined to exclude the outermost 10 cm of the active volume boundaries in the x and y coordinates and a 15 cm buffer from the upstream detector boundary in z . For the downstream buffer in z , interaction vertices must be 1 m from the detector edge to be considered fiducial, while track depositions must be a minimum of 15 cm from the detector edge to be considered fiducial. The fiducial volume definition excludes two additional regions of the detector in order to avoid complications with modeling the detector response around a hanging cable in the western-most TPC and a field cage distortion in the eastern-most TPC. The non-fiducial region around the hanging cable is defined as $210.29 < x < 358.49$ cm, $y > 70$ cm, and $z > 0$ cm. The non-fiducial region around the field cage distortion is defined as $-358.49 < x < -210.29$ and either $y > 115$ cm or $y < -161.86$ cm. Finally, preselected events must not be marked as obviously corresponding to a cosmic event, as identified with the `clear_cosmic` tag by Pandora.

Calorimetry and track length are considered for “muon-likeness.” Using calorimetry, a χ^2 fit of the reconstructed dE/dx vs. residual range for the track is computed with respect to the expected signature for muons and protons in liquid argon [98, 99]. Henceforth, these χ^2 values may sometimes be referred to as the “PID scores” of the track. To be considered muon-like by the preselection, contained tracks must be at least 10 cm long and be loosely consistent with a muon and inconsistent with a proton, with $\chi_\mu^2 < 30$ and $\chi_p^2 > 80$. For uncontained tracks, the preselection requires that the track is at least one meter long. There is no calorimetric requirement imposed on exiting tracks in the preselection. For some interactions, there may be more than 2 tracks in the reconstructed interaction satisfying these criteria. In cases where there are more than two such reconstructed tracks in an event, the two muon candidates are taken to correspond to the tracks with start points that are closest to the interaction vertex. The track starting closest to the

interaction vertex is called the “trunk” track, and the track with the next-closest start point is called the “branch” track. Up to the requirement of track (un)containment, this preselection is identical to that used in the contained dimuon analysis.

5.1.2 Object Cuts

Object cuts are applied to identify and eliminate interactions with more primary particles than the two candidate muon events identified in the preselection. Low energy ($\lesssim 50$ MeV kinetic energy) primary protons are targeted first. Such low energy protons are unlikely to traverse through the liquid argon far enough to leave a discernible track of ionization charge that Pandora can reconstruct. On the other hand, ionization charge from the Bragg peak of these particles results in high charge deposition near the interaction vertex, which Pandora has a tendency to cluster into other particles’ reconstructed tracks that originate at the same vertex. These signatures, colloquially referred to as proton “stubs,” are identified using an algorithm that checks the collection plane dE/dx of the trunk and branch tracks in regions close to the track start point. This is done in four regions along the track: along the first half-centimeter, between 0.5-1 cm, between 1-2 cm, and between 2-3 cm from the start of the track. Distributions of the dE/dx in each in each of these regions along the track are shown for BSM, ν , and cosmic events in Figure 5.1. Discrimination between events with and without low energy protons in the final state is found with the application of acceptable maximum values of dE/dx in each of these regions:

- dE/dx along first half-centimeter of track ≤ 50 MeV/cm (“Stub Cut 1”)
- dE/dx from 0.5 cm to 1 cm along track ≤ 35 MeV/cm (“Stub Cut 2”)
- dE/dx from 1 cm to 2 cm along track ≤ 15 MeV/cm (“Stub Cut 3”)
- dE/dx from 2 cm to 3 cm along track ≤ 10 MeV/cm (“Stub Cut 4”).

Henceforth, these cuts will be collectively referred to as the “Stub ID” cuts.

Next, events with primary electrons or photons are eliminated from the sample via the identification of sufficiently long reconstructed showers. Though the final state of interest contains exactly

zero electrons or photons, Delta rays or misclassified charge depositions from a minimum ionizing particle can still result in reconstructed low-energy showers in signal events. The number of reconstructed showers above some energy threshold, then, could be a powerful variable to cut on instead in order to eliminate background events containing a primary electron or photon. However, present challenges with the current shower energy reconstruction in Pandora at the time this event selection was designed deter use of reconstructed shower energy in the analysis. Shower length, which has a positive correlation with energy, is therefore chosen instead as the variable used to identify and eliminate interactions with primary electrons or photons in the final state. An upper limit on the length of the longest reconstructed shower in selected events is enforced. Interactions with no reconstructed showers are assigned an unphysical shower length quantity equal to -10, so that these events are accepted by cuts that also accept short reconstructed showers. A threshold set at zero cm would thus only select events with no reconstructed shower. The maximum reconstructed shower length for all reconstructed BSM, ν , and cosmic interactions that pass the preselection is shown in the top left panel of Figure 5.2. A cut at a **maximum shower length** < 0 cm, therefor accepting only events with no reconstructed shower, has been adopted for this analysis, as informed by the exercises described in Section 5.2.

5.1.3 Particle Identification

As was done in the contained dimuon analysis, I reconsider tighter cuts on contained tracks' PID scores beyond what was imposed in the preselection. PID scores are calculated by comparing the reconstructed dE/dx vs. residual range of a track to theoretical expectations from the Bethe-Bloch formulation. In the Bethe-Bloch formulation, residual range is defined with respect to where particles come to rest. In the reconstruction software, residual range is defined with respect to the reconstructed endpoint of the track, which does not correspond to the point at which the particle traversing the detector comes to rest if the particle exits the detector. For exiting tracks, then, use of reconstructed track end point as the reference point for residual range folds a non-physical systematic effect into the determination of residual range, and thus also into the calculated PID

scores. With this in mind, I only apply a further PID score cut to contained muon candidate tracks. 94% of all preselected background events have a muon candidate track that is contained, while only 59% of preselected signal events have a muon candidate track that is contained. The muon PID scores for all contained trunk and branch tracks from preselected events are shown in Figure 5.2. The cut is placed at $\chi_\mu^2 < 8$ for contained muon candidate tracks.

5.1.4 Kinematic Cuts

Two kinematic variables that provide substantial discrimination power between signal and background events have been identified: θ_{NuMI} and $\theta_{\mu\mu}$. These are depicted schematically in Figure 5.3, and their distributions for preselected BSM, ν , and cosmic interactions are shown in Figure 5.2.

θ_{NuMI} is the reconstructed angle between the incoming parent particle, and the vector that connects the NuMI target to the origin of the ICARUS detector coordinate system. For signal events, the parent particle is the dark sector scalar, while for neutrino background events it is the incoming neutrino. The parent particle direction is determined from its reconstructed momentum, $\mathbf{p}_{\text{parent}} = \mathbf{p}_{\text{trunk}} + \mathbf{p}_{\text{branch}}$, assuming that the candidate muons originate from a two body decay. This assumption is true for BSM particles producing dimuon final states. Since most kaons and psuedo-scalar mesons would decay into the BSM particle near the NuMI target shortly after being produced, the distribution in θ_{NuMI} is then highly peaked at zero for signal. On the other hand, despite most neutrinos also originating near the target [82], the distribution of (reconstructed) θ_{NuMI} is less peaked towards zero for neutrino events since typically not all of the incoming neutrino's momentum is shared between the two candidate muon tracks. This is attributed to by various effects including “missing” reconstructed energy of scattered neutrinos in the final state, momentum-smearing from the Fermi-momentum of nucleons in ν -Argon interactions, and neutrino interaction final states comprised of more than two visible primary particles and subsequently not fully accounted for by consideration of just the selected trunk and branch muon candidate tracks. θ_{NuMI} is peaked far from zero for cosmic activity, which primarily comes from above the detector.

By default, Pandora uses the dE/dx signature of a track in conjunction with the continuous slowing down approximation (CSDA) to determine a particle's incident momentum from the length of its ionizing track:

$$\text{Track Length} = \text{Range}_{\text{CSDA}} = \int_{E_{\text{incident}}}^0 \frac{1}{(dE/dx)} dE. \quad (5.1)$$

The range (length of track) and dE/dx are measured from the data, then the predicted particle type is used to solve for the incident energy or momentum. Similarly to the case of PID scores, the range based momentum reconstruction biases the reconstructed momentum for uncontained tracks, and thus biases the calculated θ_{NuMI} for uncontained muon tracks.

Particle momenta can alternatively be determined from principles of Multiple Coulomb Scattering (MCS), instead of Pandora's default range-based momentum reconstruction. Multiple Coulomb Scattering refers to a series of electromagnetic scattering events of an incident charged particle off atomic nuclei. MCS happens as charged particles traverse the liquid argon in ICARUS, so that on short scales, reconstructed particle tracks resemble a series of splines as opposed to one long straight line. Given an initial momentum p of the incident particle, the scattering angles of these coulomb interactions follow a Gaussian distribution centered at zero, with an RMS prescribed analytically by the Highland Formula [100]:

$$\sigma_{\text{HL}} = \frac{s_2}{p\beta c} z \sqrt{\frac{l}{X_0} \left(1 + \epsilon \ln \frac{l}{X_0} \right)}, \quad (5.2)$$

where, z is the charge of the incoming particle, X_0 is the radiation length of the target material, l is the track segment or spline length, and s_2 and ϵ are fit parameters. By measuring the RMS of the MCS scattering angles, one can thus determine the incident particle's momentum.

To determine θ_{NuMI} , range-based momentum reconstruction is used for contained tracks, and the momentum estimated from Multiple Coulomb Scattering (MCS) is used for uncontained tracks. The algorithm ICARUS uses to determine particle momentum from MCS was developed by the MicroBoonNE Collaboration, and ported into sbncode by SBN collaborator Mun Jung. Jung val-

idated use of this algorithm in ICARUS using a sample of stopping, cathode-crossing cosmic tracks ≥ 20 cm long and with a median $dQ/dx > 1000$ ADC/cm. Jung’s studies (unpublished), demonstrated no bias in the MCS-reconstructed momentum with respect to range-based momentum reconstruction for simulated events, and a 3% disagreement between p_{MCS} and p_{range} in data.

I have repeated studies on the efficacy of the MCS momentum reconstruction for the present analysis, comparing p_{MCS} and p_{true} for candidate muon tracks that pass the preselection, and that are truth-matched to a muon, from the simulated BSM and ν interactions. This comparison revealed a significant bias in p_{MCS} not captured by the previous studies. The bias can be seen in the blue curves of Figure 5.4 and in the left panel of Figure 5.5. As seen in the right panel of Figure 5.4, the bias present in uncontained tracks is confounded by the algorithm’s (expected) increased difficulty resolving the momentum when the track is not fully contained. Reconstructed momentum of contained tracks thus gives better insight to the systematic bias present in the MCS momentum reconstruction algorithm.

For the present analysis, I apply a correction to the reconstructed p_{MCS} that re-centers the percent difference between p_{MCS} and p_{true} at zero. To derive this correction, I first plotted the distribution p_{MCS} vs. p_{true} for contained and exiting tracks from the simulated BSM and ν samples, as shown in the left panel of Figure 5.5. For every p_{true} bin, I identified the hottest p_{MCS} bin, and fit a line to those points, using the bin centers as coordinates. Each point used for the fit was weighted by the number of events in that bin. Using this best fit, I then identified the transformation one would need to apply to reconstructed p_{MCS} values in order to bring the best fit line back to $y = x$ in the p_{MCS} vs. p_{true} space. The needed correction is

$$p'_{\text{MCS}} = 0.846 \times p_{\text{MCS}} + 0.036, \quad (5.3)$$

where p'_{MCS} is the corrected value and p_{MCS} is the value output by the default algorithm. The orange curves in Figure 5.4 show $(p'_{\text{MCS}} - p_{\text{true}})/p_{\text{true}}$; the peaks for both the exiting and contained tracks are brought closer to zero after the correction. The right panel of Figure 5.5 shows p'_{MCS} vs.

p_{true} , exemplifying that the distribution in this space lies along the line $y = x$ after the correction is applied.

I would like to note that other methods of calculating θ_{NuMI} that avoid the dependency of the variable on the MCS momentum reconstruction were also considered. One alternative method, referred to as the “weight by length” method, avoids reconstructing the candidate muons’ momenta altogether. With that method, the direction of the incoming parent particle is instead computed by adding the candidate muons’ tracks as vectors with directions equal to the reconstructed track directions and magnitudes equal to their lengths. The distribution of the θ_{NuMI} variable as reconstructed with the “weight by length” method is more discrepant with the distribution made from truth-level momentum than are the momentum-informed θ_{NuMI} distributions, resulting in less discrimination power with the $\theta_{\text{NuMI}_{\text{wgt_by_len}}}$ method. Furthermore, the bias in the MCS momentum reconstruction has very little impact on the distribution of θ_{NuMI} . These statements are evidenced in Figure 5.6. For these two reasons, I opt to cut on the momentum-informed version of θ_{NuMI} in the analysis, as described above.

An exception to the typical profile of θ_{NuMI} being displaced from zero for neutrino interactions lies CC Coh π interactions:

$$\nu_{\mu} + A \rightarrow \mu + \pi + A, \quad (5.4)$$

where A represents an argon nucleus. The energy transfer from the neutrino to the nucleus in these interactions is small, so that most of the final state visible momentum $\vec{p}_{\mu} + \vec{p}_{\pi}$ is oriented along $\vec{p}_{\nu}$. An aggressive cut at $\theta_{\text{NuMI}} 3^{\circ}$ is adopted for the present analysis in order to eliminate as much CC Coh π background as possible.

The other kinematic cut employed in the event selection is on the reconstructed dimuon opening angle, $\theta_{\mu\mu}$. $\theta_{\mu\mu}$ is calculated as the angle between the reconstructed track directions of the trunk and branch candidate muon tracks. Unlike the other parameters used in the event selection, the distributions of $\theta_{\mu\mu}$ vary significantly across the BSM model benchmark masses. Samples with a higher portal particle mass tend to produce muon pairs with a wider opening angle. This is accounted for by the available phase space, which is fixed by the NuMI meson flux, being consumed

by the downstream particles' rest energy in addition to their kinetic energy. When the portal particle mass is higher, there is less available energy left to boost its decay products, resulting in a less columnated, i.e. wider, opening angle. This trend is demonstrated in Figure 5.7, and taken into account when setting the upper limit on acceptable opening angles in the event selection at $\theta_{\mu\mu} < 25^\circ$.

5.1.5 Track Length Cut

CC Coh π events, leading to a final state of $1\mu 1\pi$, are the toughest background to eliminate in the present analysis. On average, muons are expected to travel farther without scattering or reinteracting in LAr than are pions. A requirement on aggressively long tracks is adopted in an effort to squander persistent $1\mu 1\pi$ events surviving the other cuts. The distribution of the minimum track length for trunk and branch tracks in preselected events is shown in Figure 5.2. The secondary bump at 1 m shown there is a feature of the preselection, where any exiting track was required to be at least one meter long. The final track length requirement accepts only events where **both the dimuon candidate tracks (contained or uncontained) are at least one meter long**.

5.2 Exercises to Optimize the Event Selection

Initial cut thresholds were first “eyeballed” from the distributions for the BSM signal, background, and cosmics shown in Figures 5.1 and 5.2. From there, two exercises were performed to better optimize the placement of each cut threshold.

The first optimization exercise optimized for statistical significance $S = n_{\text{sig}} / \sqrt{n_{\text{sig}} + n_{\text{bg}}}$. To optimize the cut threshold for a given cut, various thresholds were scanned over with the other cut thresholds fixed to their initial values. At each scan, the statistical significance was computed. This process was repeated for all cuts except for the final requirement on track length. The reason the track length cut was excluded from this exercise is because that cut is intentionally designed to eliminate all remaining background (or as much as possible) - an optimization for statistical significance distracts from the simpler minimization of background events, which this cut is intended

for. To stay model-agnostic here, I then chose updated cut thresholds that appeared to average the optimal choices across the various model benchmarks.

A second exercise to further optimize, and ultimately finalize, the event selection was performed by comparing the projected sensitivities to the Higgs portal scalar and heavy QCD axion models for event selections slightly varying from the ones benchmarked after optimizing for statistical significance. (The procedure for determining these sensitivities is described in Section 8.) Variations among the tested event selections included tightening and loosening the cuts already in the selection, using different versions of the θ_{NuMI} variable, and excluding certain cuts altogether.

Finally, introductions of a few additional cuts were considered, inspired by a visual scan of survived background events. One additional cut checked for reconstructed primary particles in excess of the muon candidates. Another cut rejected any events with a contained candidate muon track with reconstructed daughter particles. Muons, including those comprising the signal, can have delta rays or Michel electrons at the ends of the tracks, but these are unlikely to be reconstructed as daughter particles. Higher energy minimum ionizing daughter particles resulting from hard pion interactions, by contrast, are more likely to be reconstructed. I also tested additional cuts designed to identify substantial “kinks” in the candidate muon tracks, indicative of pions undergoing a hard scatter that a muon would not (e.g., not just multiple coulomb scattering). The distribution of scattering angles for track segments identified by the MCS momentum reconstruction algorithm was probed in an attempt to identify hard scattering kinks. While some of these cuts excluded single surviving background events from the benchmarked event selection, none improved the overall event selection in a statistically significant way. Therefore, none of these additional cuts was adopted into the cut flow.

The final set of cuts used in the analysis is presented in Tables 5.2.1, 5.2.1, and 5.2.1.

5.2.1 Event Selection Performance

For an exposure of 2.41×10^{20} protons on target, the event selection results in 0 cosmic events, and 1.78 selected neutrino background events. These values correspond to 0 and 55 simulated

Monte Carlo events, respectively. 49 (89%) of the remaining background events are CC Coh π interactions. The average selection efficiency across all simulated signal model benchmarks is about 15%. Tables 5.2.1, 5.2.1, and 5.2.1 show the event selection performance on each of the simulated signal model benchmarks and each type of background interaction.




Figure 5.1: First four panels: dE/dx in the collection plane, calculated along various regions near the start of the track, for all simulated events that pass the preselection. **Bottom Panel:** Distribution in lengths for the longest reconstructed shower in each event. Events that do not have any reconstructed showers are assigned a flagged value of -10, attributing to the most populated bin.




Figure 5.2: **Top left:** Muon-PID scores for contained candidate muon tracks, for all events that pass the preselection. **Top right:** θ_{NuMI} determine using range-based momentum reconstruction for contained tracks and the corrected MCS momentum reconstruction for uncontained tracks. **Bottom left:** Opening angle $\theta_{\mu\mu}$ between the candidate muon tracks. **Bottom right:** length of the shorter of the branch or trunk tracks for each event. The secondary peak at 1 m is a feature of the preselection requirement that any uncontained track is at least 1 m long. **All:** Plots are for preselected events from each simulated sample.




Figure 5.3: θ_{NuMI} is the reconstructed angle between the candidate muons' (solid orange) parent particle momentum (dashed orange) and the vector connecting the NuMI target to the origin of the ICARUS coordinate system (cyan). The ICARUS coordinate system origin is 23cm above the geometric center of the detector. $\theta_{\mu\mu}$ is the opening angle between the reconstructed candidate muon tracks.




Figure 5.4: $(p_{\text{MCS}} - p_{\text{true}})/p_{\text{true}}$ distributions for contained (left) and exiting (right) muon-matched tracks from the simulated neutrino and BSM samples. The blue curve is made using the default MCS-reconstructed momenta. The orange curve shows the distribution after the correction is applied. Muon tracks with an unphysical MCS momenta (before the correction is applied) are excluded from the distributions. Unphysical values of p_{MCS} result from computational underflows, overflows, and failures of the MCS momentum reconstruction algorithm, and occur less than 3% of the time.




Figure 5.5: MCS-reconstructed momentum vs. true momentum for candidate muon tracks from the simulated BSM and ν samples. The left panel shows the nominal MCS-reconstructed momentum, while the right panel shows the MCS-reconstructed momentum after the correction from Equation 5.1.4 is applied. The solid black line in both panels is the line $y=x$, which should fit the data if there is no bias in the MCS momentum reconstruction. The dotted black line in the left panel is the best fit to the data using the hottest p_{MCS} bin per p_{true} bin, as described in the text. That best fit line has a slope of 1.18 and a y-intercept of -0.04.




Figure 5.6: The distribution in θ_{NuMI} for simulated BSM and neutrino interactions in the analysis, according to different methods for calculating θ_{NuMI} . Solid lines are used for the BSM events while dashed lines are used for neutrino events. Black is used to show θ_{NuMI} as calculated from the true momenta corresponding to the trunk and branch candidate muon tracks. Green shows θ_{NuMI} as calculated using range-based momentum reconstruction for contained tracks and the default MCS momentum reconstruction for uncontained tracks. Pink is the same as green, except that the MCS-momentum has been corrected in-post, as described above. Purple shows θ_{NuMI} calculated according to the “weight by length” method.




Figure 5.7: The distribution in dimuon opening angle, $\theta_{\mu\mu}$, for different simulated BSM model benchmarks. Generally, the opening angle is larger for muon pairs produced in the decay of higher-mass portal particles, as explained in the text. An exception to this trend is seen for muon pairs coming from heavy QCD axions with mass $m_a = 540$ MeV, which are close to the η resonance. This is because events near the eta resonance tend to be much more boosted, as shown in Figure 4.5 and described in Section 4.3.

Events w/ $m_S = X$ (MeV)	220	240	260	280	300	320	340
Preselection	1.0	1.0	1.0	1.0	1.0	1.0	1.0
$dE/dx \leq 50$ MeV/cm up to 0.5 cm	0.921	0.924	0.908	0.915	0.911	0.908	0.913
$dE/dx \leq 50$ MeV/cm 0.5-1 cm	0.914	0.921	0.906	0.912	0.907	0.905	0.909
$dE/dx \leq 50$ MeV/cm 1-2 cm	0.887	0.895	0.878	0.888	0.880	0.878	0.885
$dE/dx \leq 50$ MeV/cm 2-3 cm	0.849	0.862	0.843	0.856	0.850	0.848	0.857
Max shower length < 0 cm	0.493	0.539	0.526	0.531	0.530	0.522	0.520
Contained track $\chi_\mu^2 < 8$	0.360	0.427	0.415	0.424	0.412	0.404	0.397
$\theta_{\text{NuMI}} < 3^\circ$	0.177	0.242	0.252	0.246	0.243	0.240	0.229
$\theta_{\mu\mu} < 25^\circ$	0.174	0.238	0.240	0.218	0.201	0.182	0.161
Minimum track length > 1m	0.158	0.199	0.188	0.167	0.158	0.145	0.131

Events w/ $m_a = X$ (MeV)	300	350	400	450	500	550	600	650	700
Preselection	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0
Stub Cut 1	0.927	0.929	0.929	0.934	0.928	0.932	0.925	0.929	0.932
Stub Cut 1	0.924	0.925	0.926	0.931	0.925	0.929	0.921	0.925	0.928
Stub Cut 1	0.897	0.898	0.899	0.904	0.897	0.898	0.893	0.896	0.900
Stub Cut 1	0.859	0.866	0.866	0.873	0.867	0.860	0.864	0.867	0.875
Shower Cut	0.485	0.489	0.476	0.460	0.455	0.423	0.445	0.447	0.456
χ_μ^2 Cut	0.389	0.396	0.379	0.371	0.364	0.353	0.355	0.345	0.355
$\theta_{\text{NuMI}} < 3^\circ$	0.255	0.261	0.249	0.234	0.226	0.208	0.209	0.199	0.196
$\theta_{\mu\mu} < 25^\circ$	0.235	0.218	0.190	0.166	0.152	0.181	0.117	0.098	0.077
Trk Length Cut	0.208	0.195	0.172	0.156	0.143	0.174	0.113	0.096	0.075

	Background	ν NC	ν_μ CC QE+MEC	ν_μ CC RES	ν_μ CC DIS	ν_μ CC COH	ν Other	Cosmic
Preselection	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0
Stub Cut 1	0.874860	0.899001	0.863089	0.880692	0.917506	0.876506	0.918908	
Stub Cut 2	0.823091	0.866468	0.792384	0.829567	0.909894	0.817771	0.886380	
Stub Cut 3	0.712104	0.788428	0.669118	0.714344	0.874540	0.664157	0.810143	
Stub Cut 4	0.642458	0.749923	0.611251	0.641603	0.844341	0.588855	0.768607	
Shower Cut	0.149348	0.391640	0.211486	0.081891	0.328259	0.094880	0.493888	
χ_μ^2 Cut	0.051769	0.166684	0.082204	0.027940	0.129143	0.028614	0.156345	
$\theta_{\text{NuMI}} < 3^\circ$	0.000372	0.003809	0.003290	0.000752	0.044930	0.000000	0.000011	
$\theta_{\mu\mu} < 25^\circ$	0.000000	0.000206	0.000113	0.000065	0.018905	0.000000	0.000000	
Trk Length Cut	0.000000	0.000000	0.000032	0.000065	0.012030	0.000000	0.000000	

Table 5.1: Selection efficiency of each type of event, for every consecutive cut in the event selection. Fractions are relative to the preselected samples of events. The same set of cuts is shown in all tables – a different naming convention is used for the ALP and background events for formatting purposes.

Events w/ $m_S = X$ (MeV)	220	240	260	280	300	320	340
Preselection	2710	6052	7850	5227	8962	10525	12930
$dE/dx \leq 50$ MeV/cm up to 0.5 cm	2495	5592	7130	4781	8163	9556	11809
$dE/dx \leq 50$ MeV/cm 0.5-1 cm	2478	5575	7110	4769	8131	9521	11751
$dE/dx \leq 50$ MeV/cm 1-2 cm	2400	5418	6890	4642	7889	9246	11438
$dE/dx \leq 50$ MeV/cm 2-3 cm	2302	5218	6621	4473	7620	8922	11075
Max shower length < 0 cm	1336	3265	4133	2775	4752	5496	6729
Contained track $\chi_\mu^2 < 8$	975	2586	3261	2215	3693	4256	5136
$\theta_{\text{NuMI}} < 3^\circ$	479	1464	1980	1286	2178	2527	2967
$\theta_{\mu\mu} < 25^\circ$	471	1439	1886	1139	1797	1918	2081
Minimum track length > 1 m	429	1205	1475	873	1412	1528	1697

Events w/ $m_a = X$ (MeV)	300	350	400	450	500	550	600	650	700
Preselection	14466	17668	19617	22104	24056	25684	24829	24667	24409
Stub Cut 1	13412	16412	18230	20646	22320	23953	22955	22909	22739
Stub Cut 1	13367	16351	18164	20574	22247	23863	22867	22809	22652
Stub Cut 1	12977	15860	17630	19980	21576	23068	22175	22097	21968
Stub Cut 1	12423	15293	16981	19287	20867	22078	21449	21379	21350
Shower Cut	7013	8644	9334	10174	10948	10869	11060	11022	11124
χ_μ^2 Cut	5627	6999	7443	8196	8764	9079	8818	8501	8656
$\theta_{\text{NuMI}} < 3^\circ$	3696	4610	4891	5181	5426	5340	5199	4905	4779
$\theta_{\mu\mu} < 25^\circ$	3405	3843	3720	3674	3647	4654	2902	2426	1868
Trk Length Cut	3005	3439	3375	3440	3438	4458	2802	2360	1819

	Background	ν NC	ν_μ CC QE+MEC	ν_μ CC RES	ν_μ CC DIS	ν_μ CC COH	ν Other	Cosmic
Preselection		2685	9713	62004	61203	4073	664	88356
$dE/dx \leq 50$ MeV/cm up to 0.5 cm		2349	8732	53515	53901	3737	582	81191
$dE/dx \leq 50$ MeV/cm 0.5-1 cm		2210	8416	49131	50772	3706	543	78317
$dE/dx \leq 50$ MeV/cm 1-2 cm		1912	7658	41488	43720	3562	441	71581
$dE/dx \leq 50$ MeV/cm 2-3 cm		1725	7284	37900	39268	3439	391	67911
Max shower length < 0 cm		401	3804	13113	5012	1337	63	43638
Contained track $\chi_\mu^2 < 8$		139	1619	5097	1710	526	19	13814
$\theta_{\text{NuMI}} < 3^\circ$		1	37	204	46	183	0	1
$\theta_{\mu\mu} < 25^\circ$		0	2	7	4	77	0	0
Minimum track length > 1 m		0	0	2	4	49	0	0

Table 5.2: Number of Monte Carlo events surviving each cut in the event selection, for each type of event.

Background	ν NC	ν_μ CC QE+MEC	ν_μ CC RES	ν_μ CC DIS	ν_μ CC COH	ν Other	Cosmic
Preselection	394.52	2837.78	6652.41	4682.57	154.31	147.93	48567.53
Stub Cut 1	344.21	2558.59	5653.74	4025.13	140.90	128.39	44629.07
Stub Cut 2	325.76	2459.64	5122.40	3698.97	139.80	119.31	43049.29
Stub Cut 3	278.98	2256.49	4184.20	3008.02	135.02	100.65	39346.65
Stub Cut 4	253.45	2159.46	3666.96	2568.13	130.10	91.00	37329.32
Shower Cut < 0 cm	70.38	1142.90	948.69	317.91	54.42	19.42	23986.94
χ_μ^2 Cut	24.07	483.71	340.82	114.25	22.37	6.42	7593.28
$\theta_{\text{NuMI}} < 3^\circ$	0.02	6.13	7.12	2.39	9.07	0.00	0.55
$\theta_{\mu\mu} < 25^\circ$	0.00	0.03	0.15	0.31	3.84	0.00	0.00
Trk Length Cut	0.00	0.00	0.06	0.31	1.42	0.00	0.00

Table 5.3: Number of background events surviving each cut in the event selection, scaled to an exposure of 2.41×10^{20} protons on target.

Chapter 6

Systematic Uncertainties

In addition to the statistical uncertainty associated with any dataset, this analysis is further subject to systematic uncertainties associated with the modeling of signal and background in the simulated data sets. Sources of these systematic uncertainties lie in the modeling of the portal particles' flux, the neutrino flux, neutrino interaction cross sections, effects related to particle propagation through the LAr, the detector response, the reconstruction performance for cathode-crossing tracks, and the statistical uncertainty on Monte Carlo events. At present, the analysis has all but the cathode crossing and detector variation effects for signal events accounted for. The various ways systematic uncertainties are computed are discussed below, followed by descriptions of each of the sources of uncertainty presently incorporated into the analysis in Section 6.2. Table 6.1 shows the scale of each of these sources of uncertainty on the simulated signal and background events that pass the event selection.

6.1 Calculation of Systematic Uncertainties

The systematic uncertainties on the flux, neutrino cross sections, and particle propagation through LAr are quantified with 100 alternate “systematic Universes.” For each of these systematic uncertainties, each simulated event is assigned a reweighting factor that is sampled in accordance with the modeling of that source of uncertainty. This is repeated 100 times, lending to unique reweighting factors corresponding to a single source of uncertainty across the 100 universes. For a given event in a given systematic universe i , the reweighting factors from every source of uncertainty j are multiplied in order to arrive at a single systematic universe weight for that event,

$$w_i = \prod_j w_{ij}. \quad (6.1)$$

Simulated EvtS (# MC EvtS)	flux $\% \delta$	σ_{incoh} $\% \delta$	$\sigma_{\text{CC Coh } \pi}$ $\% \delta$	g_A $\% \delta$	detector $\% \delta$	cathode crossing $\% \delta$	MC stats $\% \delta$	Total $\% \delta$
$m_S = 220$ MeV (1205)	6	-	-	-	54	10	5	56
$m_S = 240$ MeV (1205)	6	-	-	-	54	10	3	56
$m_S = 260$ MeV (1475)	6	-	-	-	54	10	3	56
$m_S = 280$ MeV (873)	5	-	-	-	54	10	3	56
$m_S = 300$ MeV (1412)	5	-	-	-	54	10	3	55
$m_S = 320$ MeV (1528)	5	-	-	-	54	10	3	56
$m_S = 340$ MeV (1697)	10	-	-	-	54	10	3	56
$m_a = 300$ MeV (3005)	65	-	-	-	54	10	2	85
$m_a = 350$ MeV (3439)	55	-	-	-	54	10	2	78
$m_a = 400$ MeV (3375)	34	-	-	-	54	10	2	65
$m_a = 450$ MeV (3440)	18	-	-	-	54	10	2	58
$m_a = 500$ MeV (3438)	6	-	-	-	54	10	2	56
$m_a = 540$ MeV (4458)	4	-	-	-	54	10	2	55
$m_a = 600$ MeV (2802)	2	-	-	-	54	10	2	55
$m_a = 650$ MeV (2360)	1	-	-	-	54	10	2	55
$m_a = 700$ MeV (1819)	3	-	-	-	54	10	2	55
ν NC (0)	-	-	-	-	-	-	-	-
ν_μ CC QE+MEC (0)	-	-	-	-	-	-	-	-
ν_μ RES (2)	5	43	0	33	54	10	74	107
ν_μ DIS (4)	17	9	0	73	54	10	81	124
ν_μ CC COH (49)	5	0	48	11	54	10	16	76
ν Other (0)	-	-	-	-	-	-	-	-
Cosmic (0)	-	-	-	-	-	-	-	-
Total BG (55)	7	2	38	14	54	10	19	71

Table 6.1: Systematics budget in the signal box ($\% \delta = \text{absolute uncertainty} / \text{number of POT-normalized events expected in the signal box}$) for the present analysis, by simulated interaction type. The first column shows the number of Monte Carlo events surviving the event selection, while the subsequent columns show the percent uncertainty on the surviving events, by uncertainty type. The percent uncertainty on the background due to detector modeling effects, 54%, is presently applied to the signal events as well as a temporary placeholder. Similarly, the 10% uncertainty associated with the reconstruction performance of cathode-crossing tracks applied to all events is a temporary placeholder value, as described in the text. These values will be updated by analysis completion.

The scheme to determine the reweighting factors w_j associated with a particular source of uncertainty differs depending on the uncertainty, but falls into one of these categories:

1. $\pm 1\sigma$: The 1σ relative uncertainty in the weight for a given neutrino is determined. A Gaussian random number is sampled centered at 1 and with a spread given by the relative uncertainty. This generated number is used for the reweighting due to the uncertainty in the parameter. In equations, this amounts to

$$w_{ij} = 1 + \mathcal{N}(0, \sigma_j), \quad (6.2)$$

where i labels the universe, j labels the parameter, and $\mathcal{N}(\mu, \sigma)$ is a random normal variable with mean μ and standard deviation σ .

2. Morph: This category is similar to $\pm 1\sigma$, but instead of a Gaussian reweighting in either direction, the reweighting is unidirectional. Concretely, in equations,

$$w_{ij} = 1 + |\mathcal{N}(0, \sigma_j)|, \quad (6.3)$$

using the same notation as for (6.2).

3. Unisim: In this category, the reweighting factors are predetermined using the packages GENIE Reweight and Geant4Reweight for the cross-section and particle propagation uncertainties respectively. For each universe, the simulation is rerun entirely with varied parameters according to those packages to determine an alternate weight.

In each universe, for every source of systematic uncertainty, a weighted histogram of a variable is constructed using the reweighting factors w_{ij} . The covariance between bins α and β for each source of uncertainty j , $\epsilon_{\alpha\beta,j}$, is then constructed for the histogram over all 100 universes ($N = 100$) as

$$\epsilon_{\alpha\beta,j} = \frac{1}{N} \sum_{i=1}^N (n_{\text{CV}}^\alpha - n_{i,j}^\alpha)(n_{\text{CV}}^\beta - n_{i,j}^\beta), \quad (6.4)$$

where n_{CV} is the number of events in a given bin for the central value sample (i.e. baseline simulation), and n_i is the number of events in the same bin but for systematic universe i .

To determine the systematic uncertainties associated with detector modeling, alternate coh-like neutrino samples are generated with tweaked parameters to account for plausible variations of the detector operations and performance, which can affect track and/or energy scale reconstruction. Alternate histograms are constructed for each of these variations, and scaled by the inverse of the fraction of coh-like events in the CV sample in order to recover a detector variation histogram representative of all NuMI neutrinos. A covariance matrix with respect to the baseline simulation is then computed.

For every source of uncertainty, the square of the diagonal of the corresponding covariance matrix is taken to be the resulting per-bin absolute uncertainty. The systematic uncertainty associated with the Monte Carlo events in each bin is taken as the square root of the number of events in the bin, in the CV sample. All the individual systematic uncertainties are then summed in quadrature to get a total systematic uncertainty per bin.

6.2 Sources of Systematic Uncertainty

Neutrino flux uncertainties are determined with the `PPFX` (as discussed in Section 4.1), where bidirectional variations to the beamline geometry and hadronic cross sections, based on a combination of hadron production data where available and `Geant4` templates where no data is available, are applied. The bidirectional variations applied are:

- shift in beam x position by ± 1 mm
- shift in beam y position by ± 1 mm
- beam spot size variation by ± 2 mm (relative to the CV of 5mm, which has been updated from 3mm since the contained muons analysis)
- shift in Horn 1 x position by ± 3 mm

- shift in Horn 1 y position by ± 3 mm
- shift in Horn 2 x position by ± 3 mm
- shift in Horn 2 y position by ± 3 mm
- shift in horn current by ± 2 kA
- increase/decrease in horn water level by ± 7 mm.

For some of these parameters, variations from the central value applied in either direction have the same impact on the resulting flux. This is true for the beam shift in the x direction, for example, since the ICARUS detector lies at $x=0$ in the NuMI reference frame, so an off-centering of the beam in either direction would lead to a reduction in the exposed flux. On the other hand, since ICARUS lies at $y > 0$ in the NuMI reference frame, an off-centering of the beam upwards in y would result in increased flux at ICARUS while an off-centering downwards would result in decreased flux at ICARUS. For each systematic universe, for every neutrino event, I randomly choose to apply either the positive or negative systematic beam variation as a “morph” type uncertainty. This approach preserves the overall trending effect of each type of variation, regardless of whether the two directions of the variation have the same or opposing effects on the resulting flux. The resulting systematic uncertainty on the flux, derived from variations within PPFX, are applied to the simulated neutrino interactions used in this analysis based on the energy of each simulated neutrino. Similarly, the flux uncertainty is applied to kaon-produced signal events based on the energy of kaon parents of neutrinos that come from PPFX.

As mentioned at the end of Section 4.3, the flux uncertainty on simulated mixing-produced ALPs is derived from a comparison of the *primary* spectra from PYTHIA8 and Geant4 presented in Figure 4.6. The relative one-sided (morph) uncertainty is defined as the difference

$$\frac{N_{\text{prim,Geant4}}(E \text{ bin})}{N_{\text{PYTHIA8}}(E \text{ bin})} - 1. \quad (6.5)$$

The `Geant4` simulation almost always gives a higher number of mesons than the `PYTHIA8` simulation, so this uncertainty pushes the baseline simulation in the direction of increased flux. Though the resulting flux uncertainty for mixing-produced ALPs is large (particularly for the lower-mass ALPs²; see Table 6.1), the uncertainty is irreducible without further improvements to the theory underlying the production of neutral mesons and mixing-produced ALPs. Furthermore, the estimation of the ALP flux uncertainty that is presented here is a more thorough and realistic account of this source of uncertainty than has been employed by other searches for mixing-produced ALPs to date (e.g. [101]). It is likely that other searches have underestimated the uncertainty associated with the heavy QCD axion flux.

The determination of all other sources of uncertainty in the present analysis is the same as was done for ICARUS’s preceding contained muons search [38], which is described in more detail in Ref. [1]. The content below paraphrases the key characteristics of these sources of uncertainty.

Neutrino interactions are varied using the `GENIE Reweight` package [102], accounting for effects of neutrino cross section uncertainties. The specific `GENIE Reweight` variations performed are shown in Table 6.2, along with the scheme that is used to determine the reweight and a brief description of the variation.

The `GENIE Reweight` package uses the classical Rein-Sehgal model to simulate CC Coh π interactions, and assigns a conservative flat uncertainty of 100% to them. The later-developed Berger-Sehgal model avoids a classical ansatz present in the Rein-Sehgal model, and agrees better with CC Coh π cross section data from K2K and SciBooNE. In a study by Mun Jung and Gray Putnam, the Berger-Sehgal model was tuned to fit more recent MINERvA cross section data on iron and hydrocarbon [103] and interpolated to argon. They embedded the results of this tune into the central value simulation for CC Coh π events at SBN, improving upon the baseline simulation

²The relative production modes for ALPs produced in NuMI via mixing with π^0 , η , and η' depends on 1) the rate of each of the parent mesons in NuMI, and 2) the mixing angles for each of the mesons with ALPs (see Equation 3.10). At the lower mass benchmarks, the ALPs are primarily produced from π^0 ’s. ALP production from (unprimed) η ’s dominates at the higher simulated mass benchmarks. The difference in the meson flux as predicted by `PYTHIA8` and `Geant4` is larger for π^0 ’s than it is for η ’s, leading to the inflated flux uncertainties for the lower mass ALPs. (The `PYTHIA8`- and `Geant4`-predicted rates for η' ’s are most discrepant, but this production mode for the ALPs is very sub-dominant for the mass benchmarks we simulate.)

Name	Weighting Scheme	Description
VecFFCCQEshape	unisim	Shape of vector form factor in charged- current quasi-elastic (CCQE) scattering
RPA CCQE	unisim	Variation on random-phase-approximation (RPA) in CCQE scattering
Coulomb CCQE	unisim	Varies the strength of Coulomb corrections
ZExpVariationResponse	unisim	Variations in z-expansion of axial form factor in CCQE scattering (to 4th order)
Norm CCMEC	unisim	Normalization variation in charged-current meson exchange current (MEC) scattering
NORM NCMEC	unisim	Normalization variation in neutral-current MEC scattering
Decay Angle MEC	morph	Variation of model of decay angle in MEC scattering
CCRESVariationResponse	unisim	Normalization uncertainty on charged-current resonant scattering
NCRESVariationResponse	unisim	Normalization uncertainty on neutral-current resonant scattering
RDecBR1gamma	unisim	Scale factor for branching fraction of $X + \gamma$
RDecBR1eta	unisim	Scale factor for branching fraction of $X + \eta$
ThetaDelta2Npi	unisim	Reweight from isotropic to Rein-Segall prediction for the π angle dist. in Δ decay
ThetaDelta2NRad	unisim	Reweight from isotropic to $\propto \cos^2 \theta$ for the γ angular distribution in Δ decay
NonRESBGvp[n]CC(NC)12pi	unisim	Non resonant neutrino-proton charged (neutral) current scattering with 12π , one proton [neutron] ($W < 2$ GeV)
NonRESBGvbarp[n]CC(NC)12pi	unisim	Non resonant antineutrino-proton charged (neutral) current scattering with 12π , one proton [neutron] ($W < 2$ GeV)
DISBYVariationResponse	unisim	Variations in the response of the Bodek-Yang DIS cross section model
FSI pi VariationResponse	unisim	Variations in pion final state interactions (FSI)
FSI N VariationResponse	unisim	Variations in FSI of nucleons
NCELVariationResponse	unisim	Variations in response of neutral-current elastic scattering

Table 6.2: Specific neutrino interaction variations from the GENIE Reweight package applied in order to determine the overall uncertainty due to neutrino interaction cross sections. The technical name, type of reweighting scheme, and a brief description are provided for each variation. This table is a copy of Table 8.2 in [1].

provided by GENIE. Systematic uncertainties on the CC Coh π cross section were derived from the uncertainty of this fit and the first principle component of the covariance matrix from the comparison of the GENIE and MINERvA-tuned fit. The uncertainty is driven primarily by the cross section’s dependence on atomic mass, and how this is scaled from Fe and CH to Ar in the interpolation. The Appendix in [38] provides more detail on the fit and uncertainties associated with the CC Coh π cross section.

Systematic uncertainties related to particle propagation through the liquid argon (i.e. energy dissipation, ionization, and scintillation processes), are determined from variations within the `Geant4Reweight` package [104]. In particular, variations are applied to p -Ar, π^+ -Ar, and π^- -Ar, and calculated as unisim uncertainties.

Detector systematics are obtained for the coh-like sample of events by running alternate ICARUS detector simulations reflective of possible deviations of the detector performance from the baseline simulation. Seven dominant variations of the detector modeling are simulated:

- Middle induction plane transparency $+1\sigma$ and -1σ . While the baseline simulation assumes a uniform transparency for the middle induction plane, studies using through-going cosmic muons have shown that some regions of the detector are more opaque and some are more transparent. Variations of the simulation are run in which the transparency is uniformly increased and decreased to the maximum and minimum observed transparency respectively. These are interpreted as 1σ upper and lower values of the transparency. Their variances with respect to the baseline Monte Carlo are averaged, and applied towards the total systematics budget as a single source of uncertainty related to the middle induction plane transparency.
- Front induction plane gain -15% (“low”) and $+15\%$ (“high”). These two variations are implemented as separate simulations. Their variances with respect to the baseline Monte Carlo are averaged, and applied towards the total systematics budget as a single source of uncertainty related to the middle induction plane transparency.

- Increase in noise by 20%. A detector simulation with increased noise is run and its variance with respect to the baseline Monte Carlo is calculated.
- Increased space charge effect by a factor of 2. In LArTPCs, “space charge effect” refers to the effect of distortions of the internal electric field caused by a buildup of slow-moving ions in the cryostat. There are two known issues with the space charge simulation used to generate the Monte Carlo samples. First, there is a bug that reverses the direction of the effect in the east TPC of both cryostats. Second, it is known to disagree in magnitude by roughly 1.5. To account for this, an alternate simulation is run with the space charge effect doubled. The variance of this change with respect to baseline Monte Carlo is calculated.
- Use wire-cell to model signal shape on front induction plane. The baseline Monte Carlo uses a tuned signal shape to model the signal in this plane, but it has been found that wire-cell better describes low-level data. A simulation is run with wire-cell signal shape and the variance with respect to baseline Monte Carlo is calculated.

The percent uncertainty derived from detector modeling variations for the coh-like sample is applied to all the background. It is safe to do this, and save on computing and storage resources that would otherwise be required to simulate these seven samples for the nominal NuMI neutrinos as well, since the background in the signal box for this analysis is almost entirely from CC Coh π events.

As seen in Table 6.1, the uncertainty due to detector modeling is the dominant source of uncertainty for background, at a value of 54%. This number is itself somewhat uncertain, as it is subject to statistical uncertainty associated with the number of events from each detector variation sample that pass the event selection. Table 6.3 shows the percent uncertainty associated with each detector variation, that contributes to the overall detector uncertainty of 54%, accompanied by the number of events from each detector variation sample that passed the event selection. The greatest contribution to the detector uncertainty is from the modeling of the front induction plane gain, but there is a wide spread in the uncertainties across all variations overall, from a few to $\sim 50\%$. It is

Detector Modeling Variation (# MC Evts)	% δ
Middle induction plane transparency (high and low averaged) (n/a)	25
Middle induction plane transparency -1σ (57)	36
Middle induction plane transparency $+1\sigma$ (26)	4
Front induction plane gain (high and low averaged) (n/a)	45
Front induction plane gain low (40)	38
Front induction plane gain high (44)	51
Increased noise (44)	9
Increased space charge effect (55)	12
Front induction plane signal shape modeling (55)	9

Table 6.3: Percent uncertainty on the background estimation in the signal box, due to specific variations of the detector modeling. The percent uncertainties shown here combine to give the total 54% uncertainty due to detector effects from coh-like events, that is shown in Table 6.1. The first column shows the number of Monte Carlo events surviving the event selection, for each simulated detector variation sample.

likely that this wide spread in the uncertainties, including the particularly high values associated with the front induction gain and low transparency samples, is due to significant statistical fluctuations given the modest numbers of events that pass the event selection for each detector variation sample. The detector uncertainties for this analysis could be made more precise, which could in turn cause them to be smaller, by increasing the statistics of the detector variation samples used to compute them.

It should be noted, also, that much work within the ICARUS Collaboration has been and continues to be done in order to better understand the detector response systems and improve the detector modeling. Such improvements have been made since the central value and detector variation simulated samples used for this analysis were generated. Subsequently, the uncertainty due to detector modeling for this analysis would be smaller if the analysis were to be repeated with more recently generated Monte Carlo samples.

At the time of writing, the uncertainty on *signal events* due to detector modeling has not been quantified, but there are plans in place to do this in the same way as for the coh-like neutrinos, using alternate detector variations for some simulated signal benchmarks. For now, temporary placeholder values for the detector modeling uncertainty on signal events are assumed to equal that of the coh-like neutrinos.

Furthermore, the preceding contained muons analysis [38] performed a data-driven study to derive a signal and background systematic uncertainty associated with the reconstruction performance for cathode-crossing tracks. The study revealed a 5% uncertainty on signal events and 6.4% uncertainty on neutrino background events in that analysis's signal box. A repeat of this exercise has not yet been performed for the present analysis, so for now a temporary placeholder value of 10% is assigned to this uncertainty for both signal and background events. A value larger than that observed for the contained muons analysis is adopted since the sample of uncontained (typically more energetic) muon candidate tracks this analysis targets tend to be longer than the fully contained muons in the other analysis, and so more frequently cathode-crossing.

Finally, statistical uncertainty for every bin is computed as the square root of the number of Monte Carlo simulated events contributing to the bin.

Chapter 7

Data/ Monte Carlo Comparisons in a Far Sideband

Data / Monte Carlo comparisons in a far sideband region have been performed to verify adequate modeling of neutrino and detector physics in the Monte Carlo simulation.

The far sideband is defined as events that pass the preselection and satisfy $15^\circ < \theta_{\text{NuMI}} < 50^\circ$.

Less than 10% of the data from ICARUS Runs 1 and 2 pass this far sideband definition, deeming this sample too statistically limited to risk observation of neutrino oscillation, which could spoil other ICARUS or SBN analyses. The lower bound on θ_{NuMI} in the far sideband definition is set to be a conservative distance from the signal box, which is at $\theta_{\text{NuMI}} < 3^\circ$. The upper bound of $\theta_{\text{NuMI}} < 50^\circ$ is imposed to reduce the number of cosmic events in the far sideband (see upper right panel of Figure 5.2), zero of which are expected in the signal box.

To compare the simulated Monte Carlo expectations for neutrino interactions to data, all samples are normalized to the statistics of the onbeam data. For the offbeam data, this is done using trigger livetime. For simulation, the number of protons on target is used.

After normalization, offbeam data is subtracted from onbeam data to statistically remove cosmic-triggered events. The simulated data is representative of this off-beam-subtracted difference, since for all of the simulated data the neutrino (or BSM) event is assumed to have triggered the readout, and the overlaid cosmics are out-of-time events that are reconstructed during event readout.

In all comparison plots that follow, simulated events are sorted by interaction type and displayed as stacked histograms. The data is plotted with black points, with error bars corresponding to statistical uncertainty.

Data/Monte Carlo comparisons are shown for the parameters used in the proton stub cuts in Figure 7.1, and for all the other cuts' parameters in Figure 7.2. The level of agreement between the

data and simulated event histograms is quantified with a reduced chi-squared value, χ^2/n , where

$$\chi^2 = \sum_{i=1} \frac{(x_{\text{data},i} - x_{\text{MC},i})^2}{\delta_{\text{data},i}^2 + \delta_{\text{MC},i}^2} \quad (7.1)$$

where i indexes the bins and x and δ are the bin values and uncertainties for Data and MC. n is the number of degrees of freedom in the histogram, which is the same as the number of bins. Only non-empty bins that are shown in the plotted ranges in Figures 7.1 and 7.2 are considered when computing χ^2/n . Each comparison plot in those figures is also accompanied by a plot showing the ratio of data to MC events in each bin, with the systematic uncertainty from the simulation shown as a gray band around 1, and statistical error bars on the data.

The data and Monte Carlo agree reasonably well for all parameters, as demonstrated by the fact that 1 is within the error of the ratio of data / MC for most bins. I will note that the reduced chi-squared for all of these comparisons is almost always <1 , which may indicate overly conservative uncertainties. This could be attributed to by setting conservative placeholder values to the uncertainty on signal due to detector modeling effects and/or the estimated uncertainty associated with track splitting when tracks cross the cathode. Another interesting finding from this exercise, from the comparison for θ_{NuMI} specifically, is what appears to be a tendency for the modeling to over-predict the presence of cosmics and underestimate how many neutrino interactions there are. If this is true, these effects could cancel each other out in the comparisons of MC and data for the other cut parameters, which have more constant ratios of ν 's to cosmics across their plotted distributions. At this stage, the potential underestimation of neutrino events hinted at by the Data/MC comparison for the θ_{NuMI} variable is not too concerning, since CC Coh π interactions, which are the primary source of uncertainty in the signal box, are hardly represented in the far sideband at all. Nonetheless, when advancing to looking at data in the near sideband and signal box, it is important to have a plan in place for differentiating an underestimation of the underlying background and an excess indicative of the presence of signal events. A procedure to interpret data in the near sideband and signal box in a way that does this is presented in Chapter 9.




Figure 7.1: Data/Monte Carlo comparisons in the far sideband for the differential energy deposition with respect to position, at various points along the reconstructed tracks. Collectively, cuts on these parameters comprise the “proton stub cuts.” The simulated background events are presented as stacked histograms, with hashed gray bands showing the total systematic and statistical uncertainty on the Monte Carlo for each bin. The black points correspond to data, with error bars equal to their statistical uncertainty.




Figure 7.2: Same as Figure 7.1, for the other cut parameters used in the event selection.

Chapter 8

Projected Sensitivity

The projected sensitivities of this analysis to the Higgs portal scalar and heavy QCD axion models, given the event selection performance and estimated uncertainty on the modeling, are now presented. Here, “sensitivity” corresponds to regions of the model spaces that could be ruled out by the search, up to some confidence level, in the case of data that agrees with the background-only expectation. At the time of writing, the CLs method (see, e.g., Chapter 40 of [15] and [105]) as implemented in the `pyhf` python library [106, 107] is used to find what number of selected signal events corresponds to these limits, and the determined number of signal events for the limits are used to draw sensitivity contours in each models’ parameter space³.

To obtain an *expected* sensitivity, that is, one that assumes agreement with the Standard Model null hypothesis, a hypothetical data observation equal to the central value expectation for background from simulations is assumed. Applying a flat uncertainty to every simulated signal benchmark of 60% in this technique⁴, and the true background uncertainty of 71% in the signal box, the median 90% confidence limit is obtained at 4.78 signal events. The contours of 4.78 selected events for HPSs and ALPs are shown with blue lines in Figure 8.1. For the ALPs, the dash-dot lines correspond to axions produced in kaon decay, while the dashed lines correspond to those produced via mixing. Bands encompassing the $\pm 1, 2\sigma$ 90% CL expected limits, at 1.64, 2.59, 11.28, and 35.19 selected events, are also shown as shaded blue bands around the central value expectation. The interpretation of the resulting plots is the following: every HPS and ALP model benchmark enclosed within these contours can be expected to be ruled out by the search to within a 90%

³It has recently been realized that the implementation of the CLs method performed by the `pyhf` python library is not completely robust for this analysis. A corrected version of the CLs limit-setting procedure will be implemented into the analysis before final publication. Nevertheless, the projected sensitivities produced via `pyhf` still represent results in the ballpark of what can be expected from the improved implementation.

⁴As seen in Table 6.1, each of the simulated signal benchmarks has its own uncertainty estimate. The limit-setting procedure used for this analysis will be updated to take into account the different levels of uncertainty for different signal masses before final publication of the analysis.

confidence level, assuming the observed signal box data is in agreement with the background-only expectation.




Figure 8.1: Projected sensitivity to the Higgs portal scalar (left) and heavy QCD axion (right) model spaces. The ALP sensitivity is drawn for the the codominance scenario for the Gauge couplings, and the minimalist scenario of the muon coupling where c_μ is induced by (or “runs with”) the axion decay constant f_a .

Figure 8.1 includes the sensitivities to the HPS and heavy QCD axion model spaces of other experiments, for comparison to the projected sensitivity for this analysis.

The projected sensitivity of this search is slightly offset towards larger mixing angles, relative to ICARUS’s search using contained muons. This effect is attributed to by the fact that Higgs scalars with larger mixing angles are likely to have a shorter mean free path before decaying, and so they must be more boosted to arrive at ICARUS before decay. The uncontained muons search, which tends to target higher energy portal particles than the contained muons search, is thus more sensitive to the scalars with larger mixing angles. The median 90% C.L. projected sensitivity of the present analysis to the HPS parameter space exceeds other searches in a modestly sized region around $m_S \approx 260\text{--}265$ MeV and $\theta_S^2 \approx 2 \times 10^{-8}$. The greater contribution of the present analysis is to the heavy QCD axions, whose sensitivity is discussed now.

This selection of uncontained dimuon pairs is not competitive with ICARUS's contained muons search in the region where ALPs dominantly originate in kaon decay. This is because ALPs in this region of the parameter space tend to have relatively lower energy, and thus more frequently result in dimuon signatures that come to rest in the LAr before escaping ICARUS's fiducial volume. On the other hand, this analysis's sensitivity to mixing-produced heavy QCD axions, which probe a higher-energy phase space of the ALPs, provides complementarity to the contained muons search in the mass range of $m_a = 350 - \sim 450$ MeV, extending ICARUS's reach. Henceforth I will focus the search and discussion on the heavy QCD axions model mostly, since this is the model space where the present analysis makes the greatest impact.

There are well understood physical explanations for why there is not sensitivity to axion masses surpassing $m_a = 500$ MeV. First, the overall flux of high mass ALPs in the NuMI beam is heavily constrained by the fixed energy distributions of their parent π_0 , η , and η' particles in the beam. The heavier the ALP, the more energetic its parent particle must be to produce it, and thus the less likely it is originate in the NuMI beam at all. Second, sensitivity to heavy QCD axions is exponentially suppressed at alp masses equal to m_{η} and $m_{\eta'}$ due to a resonance in the ALP decay width there. This feature of the model attributes to an unavoidable “pinching” of the sensitivity on either side of $m_a \sim 548$ MeV. Finally, the steep falloff from sensitivity over a wide range in $1/f_a$ at $m_a = 400$ MeV to a narrow range in $1/f_a$ close to $m_a = 450$ MeV is attributed to kinematic availability to hadronic decay modes of the axion that competes with the axion branching fraction to muons starting at $m_a = 405$ MeV. It should be pointed out that each of these characteristics of the sensitivity is entirely due to nuances in the theoretical description of the heavy QCD axion. The search may still be sensitive to high mass scalars from alternate new physics models that avoid limitations so specific to the axion modeling.

Finally, I draw special attention to the reinterpreted CHARM limits of exclusion for the heavy QCD axion model, that are shown by the dotted gray line in the right panel of Figure 8.1. These limits technically encompass the reach of sensitivity of the present analysis, but they should be considered with skepticism. The CHARM experiment performed a search for QCD axions decay-

ing to $\gamma\gamma$ and $\mu\mu$ pairs in their detector, turning exclusionary, null results [101]. Their limits were cast in terms of the ALP mass and an axion decay constant F_x parameterizing a generic QCD axion model. In their work, F_x was used to set the scale of both the axion production and the axion decay width to muons. The extension of the single parameter F_x to both of these use cases, which rely on independent physics, is not theoretically motivated and prohibits an analytic reinterpretation of their exclusion limits in the context of the more detailed model scenarios this analysis investigates. In general, their description of the heavy QCD axion model is very limited, so it is unclear exactly what they probed, or if their treatment of the axions was correct.

Berger and Putnam recently recalculated the expected sensitivity of the CHARM experiment to the co-dominance, running c_μ scenario of the heavy QCD axion model, by simulating the axion production and decay at CHARM with `Geant4` [7]. Their reinterpretation of the CHARM results improve upon the earlier reinterpretation of Winkler [108], as discussed in their paper. Berger and Putnam adopted the reconstruction and selection efficiency of 85% for $a \rightarrow \mu\mu$ events quoted by the original CHARM result, producing the reinterpreted results that are shown in Figure 8.1. While their reinterpretation of the CHARM sensitivity is now physically motivated with respect to improved modeling of the axion, one should still be skeptical of these results due to probable oversimplified simulation of the experimental configuration and systematic uncertainties, and adoption of what is likely an unrealistic selection efficiency from CHARM. The CHARM results, published in 1985, do not discuss any sources of uncertainty and do little to demonstrate the high reconstruction efficiency they quote. This would not stand up to the level of rigor required of modern experiments, that are held to increasingly high standards of systematics accounting. As such, it is likely that many sources of uncertainty were underestimated or left out entirely. As one such example, neither the CHARM nor reinterpreted CHARM results quantify an uncertainty on the axion flux, which is the leading source of uncertainty for the axions in the present analysis.

In the case of no significant excess above the Standard Model background expectation in the signal box, this analysis is projected to set leading bounds on high mass heavy QCD axions, among modern, dedicated searches.

Chapter 9

Unblinding Procedure

With a finalized event selection, (mostly) characterized uncertainty budget, and validation of good modeling with adequate agreement between data and Monte Carlo in a far sideband, this chapter now presents a procedure by which the analysis can check for BSM signatures of interest in the signal box, and report a final result. The high-level approach follows a simple cut-and-count strategy: check how many data events are observed in the signal box, and report the level of consistency of this number with the background-only hypothesis, given the estimated uncertainty for the signal box background expectation. A significant excess will point to a discovery, a lack thereof will be used to set limits of exclusion in the Higgs portal and heavy QCD axion parameter spaces following the CLs procedure.

The unblinding procedure presented here proceeds through a two-step process, first examining data in a near sideband close to the signal box, and then checking the signal box itself. The near sideband is defined as events satisfying $3^\circ \leq \theta_{\text{NuMI}} < 9^\circ$, and surviving all other cuts included in the event selection. This sideband is chosen in order to capture as much of the rare CC Coh π events as possible, while excluding types of neutrino interactions not expected to be significant signal box.

Figure 9.1 shows distributions of the expected number of events per θ_{NuMI} bin, with all the other event selection cuts applied, encapsulating both the signal box and proposed near sideband region. The signal box lies to the left of the dotted red line, and the near sideband is comprised of the two bins immediately to the right of the dashed red line. As seen in the figure, there is an approximately flat background expectation spanning from the near sideband region into the signal box, at about 2-3 events. The signal on the other hand, is expected to rise from right to left (this is by-design with the event selection), with an overall amplitude dependent on the exact model parameters for the HPSs or ALPs. The signal benchmark shown in Figure 9.1, $m_a = 400$ MeV, $c_\mu = 6.5 \times 10^{-4}$, $f_a = 2.5 \times 10^{-4}$ GeV, corresponds to a set of model parameters expected to be

near the edge of sensitivity of this search - in other words, the signal strength shown for this model is about as much as could be ruled out within the expected 90% CL, given no significant excess observed above the background expectation in the signal box. However, other sets of ALP model parameters could lead to signal strengths substantially above the background expectation, even in the proposed near sideband region. Figure 9.2 demonstrates this statement with a side-by-side comparison of two sets of ALP model benchmarks across the signal box and near sideband region.




Figure 9.1: Top panel: Background events are presented as stacked histograms, categorized by event type. The gray hashed bands show the total uncertainty on the background expectation. The expected signature from one heavy QCD axion model benchmark (mass given in units of MeV; f_a in units of GeV), near the expected edge of sensitivity, is shown in yellow, with uncertainty bands. **Bottom panel:** A breakdown of the fractional uncertainty on the background, per bin. The legend shows the fractional contribution of each type of uncertainty in the signal box.

While a signal as strong as the ALP model benchmark shown in the right panel of Figure 9.2 is unlikely to have been missed by the CHARM experiment [101], for example, their results must be taken with a grain of salt and so do not explicitly rule out the possibility. In an abundance




Figure 9.2: Same as top panel of Figure 9.1, but with the curve depicting expected signal also including the underlying background expectation. The plot on the right includes all the same inputs as that on the left, with an additional curve for an ALP model benchmark capable of producing about as many events as possible at ICARUS. Such a signature would entirely swamp the underlying background expectation, even in the near sideband region. ALP masses are presented in units of MeV; f_a are presented in units of GeV.

of caution, a fitting procedure has been developed for application to the near sideband, which can be used to identify a scenario of encountering such a strong signal there, and offer further understanding of both the sideband and signal box data in the context of the heavy QCD axion model in this case. (This fitting procedure is not relevant for Higgs portal scalars, since any region of that parameter space that could produce a large signal in the near sideband is already definitively ruled out by other experiments.)

To perform this fit, the sideband data $\mathbf{n}$ (which is a vector of length two, corresponding to the two sideband bins) is modeled as a linear combination of the underlying background expectation $\mathbf{b}_0$ and some signal benchmark $\mathbf{s}$ of a particular portal particle mass:

$$\mathbf{n} = A \times \mathbf{b} + B \times \mathbf{s}. \quad (9.1)$$

(For this step in the process, it is useful to use a model benchmark expected to lie near the edge of your expected sensitivity for $\mathbf{s}$.) In this model, the best fit normalization for the background, A , is expected to be $1 \pm$ the average fractional uncertainty on the background expectation across the bins. The best fit normalization for the signal, B , is likely to be zero in the case of sideband

data consistent with the background-only expectation. Alternately, in the presence of signal, B should capture scaling of the signal event rate with respect to non-mass model parameters, such as the overall scale factor accounting for the differences in the two model benchmarks shown in Figure 9.2, which have the same general shape. Note that portal particle mass affects the kinematics of such particles that make it to ICARUS, and thus can also affect the shape of the distribution in θ_{NuMI} . As such, the fit presented in Equation 9.1 is performed for multiple s with various benchmark masses simultaneously.

The best fit A and B from Equation 9.1 in the sideband bins are determined according to the likelihood⁵

$$\mathcal{L}(A, B|\mathbf{n}) = \mathcal{P}(\mathbf{n}|A \times \mathbf{b} + B \times \mathbf{s}) \times \mathcal{G}(A \times \mathbf{b}|\mathbf{b}, \Sigma_b), \quad (9.2)$$

where $\mathcal{P}(\mathbf{n}|\boldsymbol{\lambda})$ is the Poisson probability distribution of observing $\mathbf{n}$ data given an expectation of $\boldsymbol{\lambda}$,

$$\mathcal{P}(\mathbf{n}|\boldsymbol{\lambda}) = \prod_i \frac{\lambda_i^{n_i} e^{-\lambda_i}}{(n_i)!}, \quad (9.3)$$

and $\mathcal{G}(\mathbf{x}|\boldsymbol{\mu}, \Sigma)$ is the multivariate normal probability distribution with mean values $\boldsymbol{\mu}$ and covariance matrix Σ :

$$\mathcal{G}(\mathbf{x}|\boldsymbol{\mu}, \Sigma) = \frac{1}{\sqrt{2\pi \det \Sigma}} e^{-(\mathbf{x}-\boldsymbol{\mu})\Sigma^{-1}(\mathbf{x}-\boldsymbol{\mu})/2}. \quad (9.4)$$

The associated profile likelihood $\mathcal{L}(B|\mathbf{n})$ is obtained by adopting the A that maximizes $\mathcal{L}(A, B|\mathbf{n})$ for each B . The best fit values of A and B are those that maximize the probability density, as calculated from the profile likelihood, across all simulated portal particle mass benchmarks. Whichever model benchmark s led to the best fit values should be used in Equation 9.1 to derive an updated, data driven and model-informed expectation for data in the signal box, according to the fit. In cases where the best fit values of A and B predict a negative number of data in the signal box, a prediction of zero events from the best fit is reported.

⁵It is known that this likelihood function needs to be edited to include another multivariate normal probability distribution to account for the uncertainty on the simulated signal events - this task is outstanding. Once complete, the mock data tests will be rerun as well.

To obtain upper and lower limits on the data-driven signal box expectation within confidence level α , the set of integration bounds $B_{s,low}$, $B_{s,high}$ such that the sum of the integrated probability density across all s is equal to α , are used. All the $B_{s,low}$, $B_{s,high}$ (from the various s model benchmarks) must correspond to the same singular probability density value. Once these bounds of integration are found, the α confidence level upper bound expectation for the signal box is determined from which ever $B_{s,high}$ and corresponding value of A from the profile likelihood lead to the greatest number of events. The same is true for α confidence level lower bound expectation for the signal box, using the found values of $B_{s,low}$, unless this returns a negative number of expected events in the signal box. In this scenario, the lower bound expectation from the fit is taken to be zero, as is done for the best fit value.

In the cases that there is not a strong excess of data above the background-only expectation in the sideband, the $\alpha = 95\%$ confidence level from the fit is likely to encapsulate the background-only expectation of 1.78 events, or even 0 events, for the signal box. In these cases, the value of performing the fit to the near sideband data is somewhat obsolete. Rather, the value in performing the fit is to check for consistency between the data and the targeted BSM signature, namely heavy QCD axions, in the case that a significant excess *is* observed in either the sideband region or signal box. In the presence of a very strong ALP signal, the 95% CL signal box expectation from the fit would not overlap with the background-only expectation.

After this near sideband fit is performed and reported, the signal box data will be unblinded. The discovery sensitivity significance of the signal box observation will be determined from the outcome of 1,000,000 toy throws of hypothetical signal box data containing no new physics. Each toy throw is sampled from a Poisson distribution centered on λ , with λ sampled from a Gaussian distribution with $\mu = b_0$ and a spread equal to the uncertainty on the background in the signal box as determined from the Monte Carlo. The distribution of these toy throws is shown in Figure 9.3. A discovery sensitivity p-value for a given signal box observation n_0 will be computed from this distribution as the fraction of toy throws greater than the observed signal box data. If the p-value is significant, the signal box data will be compared with the expected range of data as determined

from the sideband fit, in order to investigate consistency of the discovery with the BSM signatures we have modeled and targeted in this analysis.




Figure 9.3: Rate of occurrence of a particular toy value for the number of background-only data events in the signal box, for 1 million toy throws.

Appendix B.6 demonstrates the application of the presented plan for unblinding on several mock data sets that are sampled from the background-only expectation, as well as various background + signal expectations. These mock data tests demonstrate potential outcomes of the analysis, depending on types of data that could be lurking in the presently-blinded regions.

This plan for unblinding the near sideband and signal box data was approved internally by the ICARUS Collaboration shortly before the writing of this dissertation. The actual unblinding of the near sideband region, step one of the proposed process, has been delayed to accommodate the completion and defense of this thesis. During that analysis “downtime,” it has since been suggested that the procedure for extracting final results be consolidated into a single step of observing data in all three θ_{NuMI} bins $< 9^\circ$ at once. The philosophy behind this suggestion is that, considering the potential (albeit unlikely) contamination of signal events into the near sideband region for possible ALP model parameters, the proposed sideband region of $3^\circ \leq \theta_{\text{NuMI}} < 9^\circ$ cannot actually be considered a sideband at all, and ought to instead be considered part of the signal box. (There is no

alternative near sideband region that captures the dominant background of CC Coh π interactions, while avoiding all possible contamination of ALP events.) It has been agreed upon by myself and the committee members that this suggestion should be adopted for the analysis. It is not clear at this time whether data in the updated signal box region at $\theta_{\text{NuMI}} < 9^\circ$ will be analyzed as a single-bin or three-bin measurement, but this decision will be investigated and finalized after the completion of this dissertation. While the final extraction of results for the analysis will likely vary from the procedure presented here and demonstrated in Appendix B.6, the procedure shown here still demonstrates the ability of this analysis to draw statistically rigorous conclusions about the presence or absence of dark sector states resulting in dimuon final states, namely, Higgs portal scalars and heavy QCD axions.

Chapter 10

Conclusions

The ICARUS detector at Fermilab, commissioned with the primary goal of probing neutrino physics for the SBN Program, is capable of probing *any* physics process that results in charged Standard Model particles in the final state. The analysis presented in this dissertation probes for BSM particles that could produce exactly two muons in the final state, where one or both of the muons is uncontained in the detector’s fiducial volume. In particular, this signature will be used as evidence for, or to set limits on, the Higgs portal scalar and heavy QCD axion models. It complements ICARUS’s sensitivity to these models that has already been demonstrated by the (Fermilab-headquartered) Collaboration’s first physics analysis, which selected for fully contained dimuon final states. In particular, this analysis is projected to have leading sensitivity to high-mass heavy QCD axions, among dedicated searches.

The underlying theory for the HPS and ALP models, ICARUS detector design and data acquisition, signal and background simulation, event selection, systematic uncertainties, data/MC comparisons in a far sideband, projected sensitivities, and a procedure for unblinding the signal box are all presented. These contents comprise the majority of work required for the completion of the analysis. Outlying tasks required for the analysis to reach full maturity include the determination of the size of detector modeling uncertainties for simulated signal events, the determination and implementation of the uncertainty associated with the reconstruction of cathode-crossing tracks, an update to the implementation of the CLs method that is used to set limits on the HPSs and ALPs, incorporation of different levels of uncertainty for the different simulated signal benchmarks for drawing exclusion curves, and, optionally but encouraged, a re-visitation for how the signal box is defined and unblinded, as discussed at the end of Chapter 9. Of course, the last step is for the signal box data to *actually* be unblinded, so that the final result of the analysis can be determined and shared with the broader physics community. There are plans in place for these outstanding

tasks to be completed by other ICARUS members after the completion of this dissertation, with a targeted analysis completion date before the end of the calendar year.

In addition to the aforementioned tasks required to *complete* the present analysis, there are a number of tasks that could be undertaken to *improve* it. As discussed in Section 6.2, the dominant source of uncertainty on simulated background (and the projected dominant reduce-able source of uncertainty on simulated signal) is from detector modeling, at 54%. This source of uncertainty in turn has statistical uncertainty associated with the number of events from each alternate detector variation simulation used to calculate it. This detector modeling uncertainty could be improved (i.e. made more certain), and potentially reduced, by increasing the statistics of the detector variation samples used to compute it. Better yet, the uncertainty due to detector modeling could certainly be reduced if one were to repeat the analysis with Monte Carlo samples generated with a more recent ICARUS code-base that accounts for improved detector calibration and modeling achieved by the ICARUS Collaboration since the samples presently used were generated. This improvement would require a dedicated production campaign within ICARUS: the analysis could re-purpose already-generated nominal NuMI neutrino samples, but would require all-new generation for the coh-like and signal samples (including central value and detector variation simulations).

This analysis could also benefit from increasing the dataset used for the result to include NuMI Run 3 data, taken March-July of 2024. Inclusion of Run 3 data would nearly double the total amount of data used (increasing the rates of both signal and background), therefor reducing its fractional statistical uncertainty. NuMI Run 3 data was acquired with the NuMI focusing horns running in the RHC configuration, contrasting the FHC configuration used for Runs 1 and 2. This difference changes the abundance of charged mesons produced from interactions in the targets, and which charged particles are focused by the horns, therefor affecting the flux seen by ICARUS. As such, one would need to generate new simulated samples produced from an RHC configuration that can be studied in order to include Run 3 data in the analysis.

ICARUS's sensitivity to Higgs portal scalars and heavy QCD axions (among other BSM models that can be probed at ICARUS) could also further be enhanced with complimentary searches for

signals of the other decay modes of the BSM particles, or signals of any of the decay modes of the BSM particles where the BSM particle originates from meson decay at rest in the NuMI absorber. Approximately 13% of the protons delivered by the NuMI beam travel through the NuMI target without interacting, and instead travel roughly 700 meters downstream to the NuMI absorber [109]. Protons that interact in the absorber produce low-energy mesons, which in turn typically come to rest in the absorber before decaying (lending to the term “decay at rest”). BSM particles originating from meson decay at rest in the absorber and arriving at ICARUS would have a distinctly different incoming direction contrasting the majority of beam-originating events that come from near the NuMI target. Furthermore, events from a given meson parent would be mono-energetic, as the stopped mesons’ decay products would only inherit energy from the (non-boosted) meson mass. These unique characteristics can be reconstructed from the observable final states, and used to distinguish meson decay at rest BSM signatures from background, complimenting searches that target the majority of beam-originating events which come from near the target.

Finally, there are a couple of additional tasks that could be performed to cast more impactful, or at least more general, results for this analysis, without changing anything about the event selection or systematics already presented. The first is to determine the analysis’s sensitivity to tau-coupled axions, as motivated by [110]. This third production mode of the axions, in excess of the kaon-decay and meson-mixing modes already accounted for here, would lead to an increased rate of axions at ICARUS, possibly extending the analysis’s sensitivity to heavy QCD axions to higher masses than what has already been shown. The tau flux from NuMI would need to be known in order to probe these types of ALPS. From there, a new sample of tau-coupled axions could be generated, or reweights could be derived and applied to make the axion samples already simulated representative of having originated from mixing with taus in NuMI. Second, it would be good for the final sensitivity of this analysis to be recast in a model-independent way before publication, as was done within the contained muons analysis as well. The shape and limits of the sensitivity this analysis has to heavy QCD axions in particular is highly restricted by nuanced details of the ALP modeling (i.e. resonances in production and decay widths), as discussed in Section 8. A recast of

the results of this analysis in the more generalized 3D space spanned by the BSM particle mass, product of production and decay branching fractions, and average lifetime, as opposed to specific model parameters, would make it easier to reinterpret the results to any other BSM models that could lead to the same uncontained dimuon final state in the future.

Bibliography

- [1] Gray Putnam. *Search for New Physics through a Long-Lived Di-Muon Resonance in the NuMI Beam with the ICARUS Detector*. PhD dissertation, University of Chicago, Chicago, IL, June 2024.
- [2] Cush. File:standard model of elementary particles.svg. https://commons.wikimedia.org/wiki/File:Standard_Model_of_Elementary_Particles.svg.
- [3] Fermi National Accelerator Laboratory SBN Program. Short-baseline neutrino program. <https://sbn.fnal.gov>. Accessed: 2026-04-02.
- [4] ICARUS Collaboration. ICARUS at the Fermilab Short-Baseline Neutrino program: initial operation. *Eur. Phys. J. C*, 83(6):467, 2023.
- [5] Fermi National Accelerator Laboratory ICARUS Collaboration. Icarus: liquid-argon neutrino hunter. <https://icarus.fnal.gov>. Accessed: 2026-03-31.
- [6] ICARUS Collaboration. Operation of the trigger system for the icarus detector at fermilab. *Journal of Instrumentation*, 20(10):P10044, oct 2025.
- [7] Joshua Berger and Gray Putnam. Sensitivity to kaon decays to alps at fixed target experiments. *Physical Review D*, 110(5), September 2024.
- [8] ATLAS Collaboration. Observation of a new particle in the search for the standard model higgs boson with the atlas detector at the lhc. *Physics Letters B*, 716(1):1–29, 2012.
- [9] CMS Collaboration. Observation of a new boson at a mass of 125 gev with the cms experiment at the lhc. *Physics Letters B*, 716(1):30–61, 2012.
- [10] Muon $g - 2$ Collaboration. Measurement of the positive muon anomalous magnetic moment to 127 ppb. *Phys. Rev. Lett.*, 135:101802, Sep 2025.

- [11] David N. Schramm and Michael S. Turner. Big-bang nucleosynthesis enters the precision era. *Rev. Mod. Phys.*, 70:303–318, Jan 1998.
- [12] Gary Steigman. Primordial nucleosynthesis in the precision cosmology era. *Annual Review of Nuclear and Particle Science*, 57(Volume 57, 2007):463–491, 2007.
- [13] Fabio Iocco, Gianpiero Mangano, Gennaro Miele, Ofelia Pisanti, and Pasquale D. Serpico. Primordial nucleosynthesis: From precision cosmology to fundamental physics. *Physics Reports*, 472(1):1–76, 2009.
- [14] Richard H. Cyburt, Brian D. Fields, Keith A. Olive, and Tsung-Han Yeh. Big bang nucleosynthesis: Present status. *Rev. Mod. Phys.*, 88:015004, Feb 2016.
- [15] Particle Data Group. Review of Particle Physics. *Progress of Theoretical and Experimental Physics*, 2020(8), 08 2020. 083C01.
- [16] A. A. Belavin, Alexander M. Polyakov, A. S. Schwartz, and Yu. S. Tyupkin. Pseudoparticle Solutions of the Yang-Mills Equations. *Phys. Lett. B*, 59:85–87, 1975.
- [17] Gerard 't Hooft. Symmetry Breaking Through Bell-Jackiw Anomalies. *Phys. Rev. Lett.*, 37:8–11, 1976.
- [18] R. Jackiw and C. Rebbi. Vacuum Periodicity in a Yang-Mills Quantum Theory. *Phys. Rev. Lett.*, 37:172–175, 1976.
- [19] Curtis G. Callan, Jr., R. F. Dashen, and David J. Gross. The Structure of the Gauge Theory Vacuum. *Phys. Lett. B*, 63:334–340, 1976.
- [20] R. J. Crewther, P. Di Vecchia, G. Veneziano, and Edward Witten. Chiral Estimate of the Electric Dipole Moment of the Neutron in Quantum Chromodynamics. *Phys. Lett. B*, 88:123, 1979. [Erratum: *Phys.Lett.B* 91, 487 (1980)].
- [21] IceCube Collaboration. Icecube quick facts. <https://icecube.wisc.edu/about-us/facts/>.

- [22] Super-Kamiokande Collaboration. Evidence for oscillation of atmospheric neutrinos. *Phys. Rev. Lett.*, 81:1562–1567, Aug 1998.
- [23] 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.
- [24] SNO Collaboration. Direct evidence for neutrino flavor transformation from neutral current interactions in the Sudbury Neutrino Observatory. *Phys. Rev. Lett.*, 89:011301, 2002.
- [25] Ziro Maki, Masami Nakagawa, and Shoichi Sakata. Remarks on the unified model of elementary particles. *Prog. Theor. Phys.*, 28:870–880, 1962.
- [26] Ivan Esteban, M.C. Gonzalez-Garcia, Michele Maltoni, Thomas Schwetz, and Albert Zhou. The fate of hints: updated global analysis of three-flavor neutrino oscillations. *Journal of High Energy Physics*, 2020(9), 2020.
- [27] LSND Collaboration. Evidence for neutrino oscillations from the observation of $\bar{\nu}_e$ appearance in a $\bar{\nu}_\mu$ beam. *Phys. Rev. D*, 64:112007, Nov 2001.
- [28] MiniBooNE Collaboration. Updated miniboone neutrino oscillation results with increased data and new background studies. *Phys. Rev. D*, 103:052002, Mar 2021.
- [29] MicroBooNE Collaboration, ICARUS Collaboration, and SBND Collaboration. A Proposal for a Three Detector Short-Baseline Neutrino Oscillation Program in the Fermilab Booster Neutrino Beam. 3 2015.
- [30] P. Adamson et al. The NuMI Neutrino Beam. *Nucl. Instrum. Meth. A*, 806:279–306, 2016.
- [31] Fermi National Accelerator Laboratory DUNE Collaboration. Dune at lbnf: An international flagship neutrino experiment hosted in a world class facility. <https://lbnf-dune.fnal.gov>.

- [32] C. Rubbia. The Liquid Argon Time Projection Chamber: A New Concept for Neutrino Detectors. 5 1977.
- [33] Brian Batell, Joshua Berger, and Ahmed Ismail. Probing the higgs portal at the fermilab short-baseline neutrino experiments. *Phys. Rev. D*, 100:115039, Dec 2019.
- [34] Raymond T. Co, Soubhik Kumar, and Zhen Liu. Enhancing searches for heavy QCD axions via dimuon final states. *JHEP*, 02:111, 2023.
- [35] A. Chatterjee, J. Hernandez-Garcia, and A. De Roeck. Heavy Neutral Lepton Searches at an ICARUS-like detector using NuMI Beam. *The European Physical Journal C*, 85:195, 2025.
- [36] Gabriel Magill, Ryan Plestid, Maxim Pospelov, and Yu-Dai Tsai. Millicharged particles in neutrino experiments. *Phys. Rev. Lett.*, 122:071801, Feb 2019.
- [37] Brian Batell, Joshua Berger, Luc Darmé, and Claudia Frugiuele. Inelastic dark matter at the fermilab short baseline neutrino program. *Physical Review D*, 104(7), October 2021.
- [38] ICARUS Collaboration. Search for a hidden sector scalar from kaon decay in the dimuon final state at icarus. *Phys. Rev. Lett.*, 134:151801, Apr 2025.
- [39] ICARUS Collaboration and SBND Collaboration. MeVPrtl: An Event Generator for Dark Sector Particles in the Short-Baseline Neutrino Program. 5 2026.
- [40] ICARUS Collaboration. Design, construction and tests of the ICARUS T600 detector. *Nucl. Instrum. Meth. A*, 527:329–410, 2004.
- [41] N. Canci. The icarus t600 detector at lngs underground laboratory. *Physics Procedia*, 37:1257–1265, 2012. Proceedings of the 2nd International Conference on Technology and Instrumentation in Particle Physics (TIPP 2011).
- [42] M Antonello et al. Experimental search for the “LSND anomaly” with the ICARUS detector in the CNGS neutrino beam. *Eur. Phys. J. C*, 73(3):2345, 2013.

- [43] OPERA Collaboration. Measurement of the neutrino velocity with the opera detector in the cngs beam. *Journal of High Energy Physics*, 2012(10), October 2012.
- [44] M. Antonello. The icarus experiment and the neutrino velocity measurement. *Nuclear Physics B - Proceedings Supplements*, 237-238:190–192, 2013. Proceedings of the Neutrino Oscillation Workshop.
- [45] L. Bagby, B. Baibussinov, V. Bellini, M. Bonesini, A. Braggiotti, L. Castellani, S. Centro, T. Cervi, A.G. Cocco, F. Fabris, A. Falcone, C. Farnese, A. Fava, F. Fichera, D. Franciotti, G. Galet, D. Gibin, A. Guglielmi, R. Guida, W. Ketchum, S. Marchini, A. Menegolli, G. Meng, G. Menon, C. Montanari, M. Nessi, M. Nicoletto, R. Pedrotta, P. Picchi, F. Pietropaolo, G. Rampazzo, A. Rappoldi, G.L. Raselli, M. Rossella, C. Rubbia, A. Scaramelli, F. Sergiampietri, M. Spanu, M. Torti, F. Tortorici, F. Varanini, S. Ventura, C. Vignoli, A. Zani, and P.G. Zatti. New read-out electronics for icarus-t600 liquid argon tpc. description, simulation and tests of the new front-end and adc system. *Journal of Instrumentation*, 13(12):P12007, dec 2018.
- [46] B. Ali-Mohammadzadeh, M. Babicz, W. Badgett, L. Bagby, V. Bellini, R. Benocci, M. Bonesini, A. Braggiotti, S. Centro, A. Chatterjee, A.G. Cocco, M. Diwan, A. Falcone, C. Farnese, A. Fava, D. Gibin, A. Guglielmi, W. Ketchum, U. Kose, A. Menegolli, G. Meng, C. Montanari, M. Nessi, F. Pietropaolo, A. Rappoldi, G.L. Raselli, M. Rossella, C. Rubbia, P. Sala, A. Scaramelli, F. Sergiampietri, M. Spanu, D. Torretta, M. Torti, F. Tortorici, F. Varanini, S. Ventura, C. Vignoli, A. Zhang, and A. Zani. Design and implementation of the new scintillation light detection system of icarus t600. *Journal of Instrumentation*, 15(10):T10007, oct 2020.
- [47] F. Abd Alrahman et al. Cryogenics and purification systems of the icarus t600 detector installation at fermilab, 2026.

- [48] ICARUS Collaboration. Calibration and simulation of ionization signal and electronics noise in the icarus liquid argon time projection chamber. *Journal of Instrumentation*, 20(01):P01032, jan 2025.
- [49] ICARUS Collaboration. Angular dependent measurement of electron-ion recombination in liquid argon for ionization calorimetry in the icarus liquid argon time projection chamber. *Journal of Instrumentation*, 20(01):P01033, jan 2025.
- [50] A. Aduszkiewicz, L. Bagby, B. Behera, P. Bernardini, S. Bertolucci, M. Betancourt, H. Budd, T. Boone, A. Campos, D. Casazza, V. Cicero, D. Cherdack, T. E. Coan, L. Degli Esposti, D. Di Ferdinando, L. Di Noto, C. Guandalini, M. Guerzoni, A. Heggestuen, C. Hilgenberg, R. Howell, M. Iliescu, G. Ingratta, T. Kim, U. Kose, G. Laurenti, C. Mariani, N. Mauri, K. S. McFarland, E. Montagna, A. Montanari, N. Moggi, M. Murphy, L. Pasqualini, V. Paolone, L. Patrizii, G. Pellegrini, V. Pia, F. Poppi, M. Pozzato, A. Ruggeri, P. Sala, P. Sapienza, A. Schukraft, G. Sirri, L. Stanco, A. Surdo, M. Tenti, V. Togo, N. Tosi, M. Vicenzi, R. J. Wilson, and S. Zucchelli. Design and implementation of the cosmic ray tagger system for the icarus detector at fnal, 2025.
- [51] ICARUS Collaboration. Experimental observation of an extremely high electron lifetime with the icarus-t600 lar-tpc. *Journal of Instrumentation*, 9(12):P12006, dec 2014.
- [52] J Serrano, P Alvarez, M Cattin, E Garcia Cota, J Lewis, P Moreira, T Wlostowski, G Gaderer, P Loschmidt, J Dedic, R Bär, T Fleck, M Kreider, C Prados, and S Rauch. The White-Rabbit Project. In *12th International Conference On Accelerator And Large Experimental Physics Control Systems, Kobe, Japan, 12 - 16 Oct 2009*, 11 2009.
- [53] Chris Bird, Paul Jackson, Robert V. Kowalewski, and Maxim Pospelov. Search for dark matter in $b \rightarrow s$ transitions with missing energy. *Phys. Rev. Lett.*, 93:201803, 2004.
- [54] F. Bezrukov and D. Gorbunov. Light inflaton Hunter’s Guide. *JHEP*, 05:010, 2010.

- [55] MicroBooNE Collaboration. Search for long-lived heavy neutral leptons and Higgs portal scalars decaying in the MicroBooNE detector. *Phys. Rev. D*, 106(9):092006, 2022.
- [56] John F. Donoghue, J. Gasser, and H. Leutwyler. The Decay of a Light Higgs Boson. *Nucl. Phys. B*, 343:341–368, 1990.
- [57] F. Bezrukov and D. Gorbunov. Light inflaton Hunter’s Guide. *JHEP*, 05:010, 2010.
- [58] F. Bezrukov and D. Gorbunov. Light inflaton after LHC8 and WMAP9 results. *JHEP*, 07:140, 2013.
- [59] Jonathan L. Feng, Iftah Galon, Felix Kling, and Sebastian Trojanowski. Dark Higgs bosons at the ForwArd Search ExpeRiment. *Phys. Rev. D*, 97(5):055034, 2018.
- [60] R. D. Peccei and Helen R. Quinn. CP conservation in the presence of pseudoparticles. *Phys. Rev. Lett.*, 38:1440–1443, Jun 1977.
- [61] R. D. Peccei and Helen R. Quinn. Constraints imposed by CP conservation in the presence of pseudoparticles. *Phys. Rev. D*, 16:1791–1797, Sep 1977.
- [62] Richard Holman, Stephen D. H. Hsu, Thomas W. Kephart, Edward W. Kolb, Richard Watkins, and Lawrence M. Widrow. Solutions to the strong CP problem in a world with gravity. *Phys. Lett. B*, 282:132–136, 1992.
- [63] Stephen M. Barr and D. Seckel. Planck scale corrections to axion models. *Phys. Rev. D*, 46:539–549, 1992.
- [64] Marc Kamionkowski and John March-Russell. Planck scale physics and the Peccei-Quinn mechanism. *Phys. Lett. B*, 282:137–141, 1992.
- [65] Michael Dine. Problems of naturalness: Some lessons from string theory. In *Conference on Topics in Quantum Gravity*, 7 1992.

- [66] S. Ghigna, Maurizio Lusignoli, and M. Roncadelli. Instability of the invisible axion. *Phys. Lett. B*, 283:278–281, 1992.
- [67] Vedran Brdar, Bhaskar Dutta, Wooyoung Jang, Doojin Kim, Ian M. Shoemaker, Zahra Tabrizi, Adrian Thompson, and Jaehoon Yu. Axionlike Particles at Future Neutrino Experiments: Closing the Cosmological Triangle. *Phys. Rev. Lett.*, 126(20):201801, 2021.
- [68] Martin Bauer, Matthias Neubert, Sophie Renner, Marvin Schnubel, and Andrea Thamm. Consistent treatment of axions in the weak chiral lagrangian. *Physical Review Letters*, 127(8), August 2021.
- [69] Martin Bauer, Matthias Neubert, Sophie Renner, Marvin Schnubel, and Andrea Thamm. Flavor probes of axion-like particles. *Journal of High Energy Physics*, 2022(9), September 2022.
- [70] Kevin J. Kelly, Soubhik Kumar, and Zhen Liu. Heavy axion opportunities at the DUNE near detector. *Phys. Rev. D*, 103(9):095002, 2021.
- [71] ArgoNeuT Collaboration. First Constraints on Heavy QCD Axions with a Liquid Argon Time Projection Chamber using the ArgoNeuT Experiment. 7 2022.
- [72] Giovanni Grilli di Cortona, Edward Hardy, Javier Pardo Vega, and Giovanni Villadoro. The QCD axion, precisely. *JHEP*, 01:034, 2016.
- [73] Daniel Aloni, Yotam Soreq, and Mike Williams. Coupling QCD-Scale Axionlike Particles to Gluons. *Phys. Rev. Lett.*, 123(3):031803, 2019.
- [74] BaBar Collaboration. Precision measurement of the $e^+e^- \rightarrow K^+K^-\gamma$ cross section with the initial-state radiation method at BABAR. *Phys. Rev. D*, 88(3):032013, 2013.
- [75] D. Heck, J. Knapp, J. N. Capdevielle, G. Schatz, and T. Thouw. CORSIKA: A Monte Carlo code to simulate extensive air showers. 2 1998.

- [76] J. Allison et al. Recent developments in Geant4. *Nucl. Instrum. Meth. A*, 835:186–225, 2016.
- [77] GEANT4 Collaboration. GEANT4—a simulation toolkit. *Nucl. Instrum. Meth. A*, 506:250–303, 2003.
- [78] John Allison et al. Geant4 developments and applications. *IEEE Trans. Nucl. Sci.*, 53:270, 2006.
- [79] R. Hatcher. Proposal for a unified “flux” n-tuple format. Technical Report MINOS-doc-9070-v4, 2012.
- [80] MINERvA Collaboration. Neutrino Flux Predictions for the NuMI Beam. *Phys. Rev. D*, 94(9):092005, 2016. [Addendum: *Phys.Rev.D* 95, 039903 (2017)].
- [81] L. Aliaga and M. Kordosky. PPFx repository. <https://github.com/kordosky/ppfx>.
- [82] Anthony P. Wood. *The NuMI Neutrino Flux Prediction at ICARUS*. PhD thesis, Houston U., Houston U., 12 2024.
- [83] Ulrich Mosel et al. The giessen boltzmann-uehling-uhlenbeck project. <https://gibuu.hepforge.org/trac/wiki>.
- [84] Yoshinari Hayato and Luke Pickering. The NEUT neutrino interaction simulation program library. *Eur. Phys. J. ST*, 230(24):4469–4481, 2021.
- [85] T. Golan, J. T. Sobczyk, and J. Zmuda. NuWro: the Wroclaw Monte Carlo Generator of Neutrino Interactions. *Nucl. Phys. B Proc. Suppl.*, 229-232:499–499, 2012.
- [86] Joshua Isaacson, William I. Jay, Alessandro Lovato, Pedro A.N. Machado, and Noemi Rocco. Introducing a novel event generator for electron-nucleus and neutrino-nucleus scattering. *Physical Review D*, 107(3), February 2023.

- [87] E. L. Snider and G. Petrillo. LArSoft: Toolkit for Simulation, Reconstruction and Analysis of Liquid Argon TPC Neutrino Detectors. *J. Phys. Conf. Ser.*, 898(4):042057, 2017.
- [88] MeVPrtl GitHub Repository.
- [89] Eric D. Church. LArSoft: A Software Package for Liquid Argon Time Projection Drift Chambers. 11 2013.
- [90] Ruth Pordes and Erica Snider. The Liquid Argon Software Toolkit (LArSoft): Goals, Status and Plan. *PoS, ICHEP2016*:182, 2016.
- [91] I. Stancu et al. Technical Design Report for the 8 GeV Beam. 5 2001.
- [92] MiniBooNE Collaboration. The Neutrino Flux Prediction at MiniBooNE. *Phys. Rev. D*, 79:072002, 2009.
- [93] MicroBooNE Collaboration. Booster Neutrino Flux Prediction at MicroBooNE. 7 2018.
- [94] Josephine L. Paton. Towards An Updated Simulation of the Booster Neutrino Beam. In *25th International Workshop on Neutrinos from Accelerators*, 1 2025.
- [95] Christian Bierlich et al. A comprehensive guide to the physics and usage of PYTHIA 8.3. *SciPost Phys. Codeb.*, 2022:8, 2022.
- [96] R. H. Milburn. Neutrino Beam Simulation using PAW with Weighted Monte Carlo. *Fermilab Internal Documents*, 1995.
- [97] MicroBooNE Collaboration. The pandora multi-algorithm approach to automated pattern recognition of cosmic-ray muon and neutrino events in the microboone detector, 2017.
- [98] C. Anderson et al. The ArgoNeuT Detector in the NuMI Low-Energy beam line at Fermilab. *JINST*, 7:P10019, 2012.

- [99] M. Miyajima, T. Takahashi, S. Konno, T. Hamada, S. Kubota, H. Shibamura, and T. Doke. Average energy expended per ion pair in liquid argon. *Phys. Rev. A*, 9:1438–1443, Mar 1974.
- [100] Virgil L. Highland. Some practical remarks on multiple scattering. *Nuclear Instruments and Methods*, 129(2):497–499, 1975.
- [101] F. Bergsma, J. Dorenbosch, J.V. Allaby, U. Amaldi, G. Barbiellini, C. Berger, W. Flegel, L. Lanceri, M. Metcalf, C. Nieuwenhuis, J. Panman, C. Santoni, K. Winter, I. Abt, J. Aspiazu, F.W. Büsser, H. Daumann, P.D. Gall, T. Hebbeker, F. Niebergall, P. Schütt, P. Stähelein, P. Gorbunov, E. Grigoriev, V. Kaftanov, V. Khovansky, A. Rosanov, A. Baroncelli, L. Barone, B. Borgia, C. Bosio, A. Capone, M. Diemoz, U. Dore, F. Ferroni, L. Luminari, P. Monacelli, F. De Notaristefani, P. Pistilli, R. Santacesaria, L. Tortora, and V. Valente. Search for axion-like particle production in 400 gev proton-copper interactions. *Physics Letters B*, 157(5):458–462, 1985.
- [102] Costas Andreopoulos, Christopher Barry, Steve Dytman, Hugh Gallagher, Tomasz Golan, Robert Hatcher, Gabriel Perdue, and Julia Yarba. The GENIE Neutrino Monte Carlo Generator: Physics and User Manual. 2015.
- [103] MINERvA Collaboration. Neutrino-Induced Coherent π^+ Production in C, CH, Fe, and Pb at $\langle E_\nu \rangle \sim 6$ GeV. *Phys. Rev. Lett.*, 131(5):051801, 2023.
- [104] J. Calcutt, C. Thorpe, K. Mahn, and Laura Fields. Geant4Reweight: a framework for evaluating and propagating hadronic interaction uncertainties in Geant4. *JINST*, 16(08):P08042, 2021.
- [105] A L Read. Presentation of search results: the cls technique. *Journal of Physics G: Nuclear and Particle Physics*, 28(10):2693, sep 2002.
- [106] Lukas Heinrich, Matthew Feickert, and Giordon Stark. pyhf: v0.7.6. <https://github.com/scikit-hep/pyhf/releases/tag/v0.7.6>.

- [107] Lukas Heinrich, Matthew Feickert, Giordon Stark, and Kyle Cranmer. pyhf: pure-python implementation of histfactory statistical models. *Journal of Open Source Software*, 6(58):2823, 2021.
- [108] Martin Wolfgang Winkler. Decay and detection of a light scalar boson mixing with the Higgs boson. *Phys. Rev. D*, 99(1):015018, 2019.
- [109] P. Adamson et al. The NuMI neutrino beam. *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment*, 806:279–306, jan 2016.
- [110] Enrico Bertuzzo, Ana Luisa Foguel, Gabriel M. Salla, and Renata Zukanovich Funchal. New Limits on Leptophilic Axionlike Particles and Majorons from ArgoNeuT. *Phys. Rev. Lett.*, 130(17):171801, 2023.
- [111] Phoebus Developers. Cs-studio/phoebus documentation. <https://control-system-studio.readthedocs.io/en/latest/>.
- [112] EPICS Collaboration. Epics application developer’s guide. <https://docs.epics-controls.org/en/latest/appdevguide/gettingStarted.html>.
- [113] EPICS Collaboration. Epics documentation, chapter 6.12. <https://epics.anl.gov/base/R3-15/6-docs/AppDevGuide/DatabaseDefinition.html#x7-2670006.12>.

Appendix A

ICARUS Detector Controls System

The ICARUS Collaboration leverages Control System Studio (CSS) [111] for the design of graphical user interfaces (GUIs) used to monitor and control hardware systems. CSS GUIs enable the developer to build displays with a wide variety of widgets including but not limited to text labels (liken to a “text box” in PowerPoint, with a set display string), various graphics (e.g. that can be used to make illustrative diagrams or schematics of detector systems), LEDs (filled in circles that take on a particular color in correspondence with acceptable or unacceptable values of data), plots, and text updates (like a text label, but with text that displays data or metadata in real-time and so is not set to a static display string). Each type of widget has a suite of programmable properties that determine its name, what it displays (e.g. if showing a particular datum), how it displays content, its position in the GUI, and any “control features” of the widget, among other characteristics. Examples of a control feature of a widget could be that it opens a webpage or another GUI if clicked on, or that it sends a signal to change the amount of power sent to a particular piece of hardware on the network.

CSS relies on the Experimental Physics and Industrial Control System software (EPICS) to display data in the GUIs in live time. In particular, widgets that relay data have a property called Process Variable (PV) Name, which needs to be set to the name of an EPICS data product, also referred to as a “record.” For hardware devices that are not manufactured to collect and transmit data according to the EPICS protocol already, the DCS experts write parser scripts to translate that collected data into the EPICS format. An EPICS application called an Input/Output controller (IOC) is used to “activate” the EPICS datasets, enabling the live time transmission of the hardware data to the CSS GUI in the EPICS protocol, and enabling live time control of the hardware systems via CSS control widgets or EPICS control commands. In other words, the IOC is what establishes two-way transmission of EPICS data between the GUIs and hardware devices.

The EPICS developers guide [112] describes the different types of EPICS records that can contribute to an EPICS database. A basic type of record frequently used by ICARUS's DCS GUIs is an Analogue Input (AI) record. AI records take an analog value from hardware as input, in the "INP" field, and convert it to a value with engineering units assigned to it. A more complex and arguably more useful type of EPICS record, that is also frequently utilized in the ICARUS slow controls GUIs, is called an MBBI record. MBBI records also have a field for the name of the hardware data it corresponds to, but it has additional programmable properties that can be used to convey pertinent information about possible values of that data. For example, if the analog input could take on values of 0 or 1, corresponding to "ok" or "error" states, these states can be programmed into the MBBI record so that when the record is displayed in the GUI it shows as the more helpful information of "ok" or "error," as opposed to a 0 or 1, which most users would not know how to interpret. Alarms of varying severity can be programmed within EPICS records, such that when a CSS widget points to a record, it will display the data with a colored border indicating a warning or alarm if the data is outside the (programmed) acceptable range. By creating a visual queue for alarm states in the GUIs, this feature enables issues with detector operations to be spotted more efficiently by shifters, and thus facilitating quicker intervention to get all systems back to spec for optimal data taking.

I have contributed to the creation and/or improvement of many ICARUS DCS GUIs and databases used to monitor and control detector hardware. When making new GUIs, special care was taken to design them in ways aligned with optimal user experience. This included streamlining displays for general ICARUS shifters (e.g. not including more information than is necessary for the shifter to check, making alarm states clearly visible, taking into account common personal laptop display sizes when establishing the layout of the GUI, and removing control widgets that could inadvertently be used to change hardware configurations by mistake), who are required only to check high-level operations, and identify and relay any alarm state to subsystem experts for further intervention. When designing GUIs for use by hardware experts, I worked closely with the experts (including leveraging my own experience as a PMT subsystem expert for work on any

PMT-related GUIs) in order to understand and cater to their needs with the GUI. The following two subsections describe two of the more involved slow controls projects I completed during my service as a DCS expert.

A.1 GUIs to Monitor the White Rabbit Switch

As mentioned in Section 2.3, a White Rabbit network [52] is used to time-synchronize the various ICARUS subsystems with beam signals from Fermilab’s accelerator complex. The device used for this synchronization is referred to as the White Rabbit Switch (WRS). The WRS is crucial to the ICARUS triggering system, and therefor must be continually tested, monitored, and backed up to ensure high integrity data acquisition. The WRS stores and transmits data according to the snmp protocol, and so must be parsed into the EPICS format. A description of every snmp data object available to and for the switch is contained in the firmware version-specific `WR-SWITCH-MIB.txt` file provided by the White Rabbit network developers [52]. ICARUS has two WRS devices - one used for actual ICARUS triggering, referred to as the “production” switch, and another switch used for offline testing and as a backup in case of failures of the production switch, the “spare” switch.

At the time I became an ICARUS slow controls expert, the only slow controls monitoring tools available for the WRS was a simple script that parses the production WRS snmp objects into EPICS AI type records, and a GUI that displayed the name and value of each of these records as a list. I improved upon these monitoring tools by creating sets of more user friendly GUIs for both the production and spare switches. For each switch, there is a “shifter” GUI designed to display only the high-level specifications and status of the switch, an “expert” GUI designed to offer more insight than the shifter GUI, and an “All records” GUI that has all available data products corresponding to the switch (like the original GUI). In the shifter and expert GUIs, records are displayed as a dependency tree, so that a high level alarm can intuitively be traced back to lower level causes. I collaborated closely with ICARUS’s White Rabbit Switch experts, to come up with and implement this design.

Furthermore, I made significant improvements to the parser script use to translate the WRS data from the snmp to the EPICS protocol. In particular, the new parser script creates MBBI type records for data products relaying information about the switch status, and these MBBI records are thus incorporated into the GUIs so that in compliant values can easily be spotted with warning or alarm displays.

Finally, my work introduced the use of EPICS templates (see [113], for example) into the WRS monitoring systems, in order to eliminate redundant work around creating similar systems for both the production and spare systems. The use of templates enables flexibility for configuring databases corresponding to whichever White Rabbit firmware version a switch is operating with. This feature proved extremely useful, as ICARUS used the spare switch to test upgraded firmware versions before porting the updated versions to the production switch: with the use of templates, new databases can be configured for either switch, for any firmware version with a few clicks of the button.

This project was worked on iteratively over several months, and reached its final state for use by ICARUS in the spring of 2023. About 6 months ahead of the SBND experiment coming online to take neutrino data, I was consulted by their DCS experts for guidance in setting up their own WRS monitoring systems. Using the systems I built and documentation I wrote, SBND was able to implement their complete WRS slow controls monitoring in about a day.

Figure A.1 is a graphic created upon completion of this project, which illustrates how the relevant hardware, data, and applications interface to produce usable GUIs for monitoring and troubleshooting a White Rabbit switch. Figure A.2 shows a screenshot of the WRS Expert GUI designed by me, featuring a section of header-type information at the top of the display showing basic information about the switch, and a dependency tree for various status-related snmp objects, converted into MBBI type EPICS records by the parser script. The screenshot shows an “error” overall status for the script, which is programmed as an alarm state and so displayed with an eye-catching red border around the widget. The error can easily be traced through the dependency tree to the “Soft PLL Status” error, facilitating streamlined troubleshooting by WRS hardware experts.

Widgets at the bottom of the display can be used to navigate to the “All Records” GUI and to the ICARUS-managed WRS information wiki, for further investigation by experts.

A.2 Overhaul to Presentation of all ICARUS Slow Controls GUIs

ICARUS collaborators access the slow controls GUIs by connecting to the icarus-dcs01 server via a VNC secure connection. CSS is left running on that remote desktop, and has tabs across the top of the window for the various GUIs that shifters and experts alike need to access in order to check on and interact with the hardware. Internal ICARUS wikis for shifters and hardware experts orient users with the contents of the GUIs, and provide instruction as to which widgets or data products are important to check. Since the detector is continuously monitored 24/7, the CSS window and various GUI tabs are left open and running on the server. Often times, when experts access GUIs not needed for the general shifters, those GUIs are left open in the CSS window as well. Over time, the CSS window can become cluttered with many tabs, causing confusion for collaborators not well oriented with all the slow controls tools, and trying to find a particular display or data value. Furthermore, such as the case with the newly designed White Rabbit Switch GUIs, many displays have two versions – one for simply monitoring the hardware, which is accessed by shifters, and one with system “control” capabilities that is used by experts to intervene with detector operations, such as changing the amount of power supplied to a device. These pairs of GUIs tend to look very similar, so it is possible for shifters to unknowingly interact with the control GUIs instead of monitoring GUIs, increasing the risk of inadvertent and/or ill informed changes getting made to actual detector operations, putting the hardware and data streams at risk. Finally, the slow controls team has designed many GUIs to assist with operations over the years, and few collaborators are aware of all of these tools that are available to them for investigating hardware systems. Without a complete knowledge of the full set of GUIs available, collaborators may work harder than necessary to investigate and address various issues.




Figure A.1: Info-graphic illustrating the key components (hardware, files, applications, and data flows) required to monitor the White Rabbit Switch, with accompanying informative text.




Figure A.2: Screenshot of the CSS-displayed WRS expert GUI used by ICARUS White Rabbit switch experts, as described in the text.

To address all of these issues, I took on a project of reorganizing the presentation of *all* the ICARUS slow controls GUIs into a more streamlined, intuitive, and user-friendly arrangement. This was done with the creation of new “top level” GUIs featuring CSS navigation tabs, where each tab redirects to an existing GUI for a piece of ICARUS hardware. The navigation tabs are oriented in a navigation tree that always remains visible on the screen, helping users to orient themselves with what GUI is actively being viewed, as well as what other GUIs are available for use. I made one organization tree intended for use by general shifters, which has just the (monitoring versions of) GUIs they are instructed to reference for their shifter checklists. Another organization tree was made to display *all* ICARUS GUIs available, for use by hardware experts. The latter organization tree includes both monitoring and control GUIs, and the display of the organization tree makes it clear to the user which version of the GUI is being accessed at a given

time. During this effort, I also performed the exercise of carefully comparing all GUIs to their respective descriptions in the ICARUS wikis. Where inconsistencies existed, I worked with experts to settle on which actions should be taken to bring them into agreement - whether that be updating the GUIs or the wikis, or both. On the software side, this intervention included changes to both GUI design as well as some EPICS databases, for example, to update alarm thresholds for various records. Furthermore, all existing ICARUS GUIs were quality checked for broken links or widgets, outdated design, proper security measures (i.e. password protections where needed or desired, such as for accessing the CRT control GUI, which was a request made by CRT experts), and any data streams compromised by the reorganization. Finally, shortcuts were programmed within CSS to easily reset the GUIs displayed in the window to just the two containing the organizational trees that can be used to access everything, providing a quick and easy way for users to “clean up” the window in the case of accumulated clutter of unnecessary GUIs. This overhaul of all of ICARUS’s slow controls GUIs was deployed in the spring of 2025, complemented by thorough efforts to educate all collaborators about the changes, and has since improved the efficiency and efficacy of monitoring and intervening with detector operations across the collaboration. Figures A.3 and A.4 show screenshots of the newly created CSS displays featuring navigation trees for the GUIs, for shifter use and expert use respectively. The navigation tabs are shown on the left, with the top-most level of tabs organized by detector hardware type. Tabs highlighted in blue indicate the path followed in order to display which ever GUI is shown to the right. As seen in Figure A.4, the users can reference the navigation tabs there to easily switch between monitoring and control GUIs, without losing track of which is being accessed.




Figure A.3: CSS GUI for use by ICARUS shifters, featuring a navigation tab to access monitoring GUIs required for performing checklist tasks.




Figure A.4: CSS GUI for use by ICARUS hardware experts, featuring a navigation tab to access any and all ICARUS slow controls GUIs.

Appendix B

Mock Data Tests for Unblinding the Near Sideband and Signal Box

Here, I demonstrate the sideband and signal box unblinding procedure outlined in Chapter 9 on mock data sets representative of no BSM signal, the presence of a weak BSM signal, and the presence of a strong BSM signal.

All mock data tests below include a figure that shows the number of events from the simulated background expectation and from the mock data, per θ_{NuMI} bin, with all other event selection cuts applied. These plots encapsulate both the near sideband region and signal box. The number of events in those Figures are represented as scatter plot points with x coordinates equal to the θ_{NuMI} bin center.

Each mock data test also includes a plot showing the profile likelihood probability density from Equation 9.2, for simulated ALP benchmarks with $m_a = 300, 350, 400, \text{ and } 450 \text{ MeV}$. In each mock data test, this probability density was used to determine the best fit A and B values for the data, according to Equation 9.1, and in turn set the sideband-data-informed 95% CL range of expected signal box values from the fit. In cases where the 95% CL range of expectation values for the signal box does not overlap with the background-only expectation from simulation, the CV and 95% upper and lower bounds for the signal box are also shown in the plot of θ_{NuMI} .

For mock data tests that do not have a significant excess of events in the signal box, “observed” exclusion curves according to the mock data are shown for the ALP model space.

B.1 Mock Data Test: Background Only, Throw 1

This mock data was sampled from the background-only expectation. The best fit values to Equation 9.1 with the near sideband data were $A = 1.13$ and $B = 0.05$, leading to a most likely signal box expectation of 2 data events. The 95% CL upper and lower bounds to the signal box

expectation from the sideband data fit were 21 and 0, respectively. In this case, these fit values are not shown in the θ_{NuMI} plot for this mock data test, since the range in expected signal box values overlaps with the background-only expectation for the signal box from Monte Carlo.

The signal box was opened and revealed zero data events, consistent with no new physics. The corresponding discovery sensitivity p-value is equal to 0.74. An “observed” exclusion curve in the ALP model space was computed for the observation of zero events in the signal box.




Figure B.1: Number of events per θ_{NuMI} bin for simulated background (blue) and fake data sampled from the simulated background (black), Throw 1. The signal box is the left-most bin, and the sideband region is comprised of the second and third bins.

B.2 Mock Data Test: Background Only, Throw 2

This mock data was sampled from the background-only expectation. The best fit values to Equation 9.1 with the near sideband data were $A = 0.84$ and $B = -0.50$, leading to a most likely signal box expectation of 0 data events. The 95% CL upper and lower bounds to the signal box expectation from the sideband data fit were 0 and 10, respectively. In this case, these fit values are not shown in the θ_{NuMI} plot for this mock data test since the range in expected signal box values overlaps with the background-only expectation for the signal box from Monte Carlo.




Figure B.2: The profile likelihood probability density from Equation 9.2 for the background-only mock data, Throw 1, with respect to $m_a = 300, 350, 400,$ and 450 MeV ALP benchmarks for s . (The values of f_a are given in GeV.)

The signal box was opened and revealed 2 data events, consistent with no new physics. The corresponding discovery sensitivity p-value is equal to 0.29. An “observed” exclusion curve in the ALP model space was computed for the observation of 2 events in the signal box.

B.3 Mock Data Test: Background Only, Throw 3

This mock data was sampled from the background-only expectation. The best fit values to Equation 9.1 with the near sideband data were $A = 1.33$ and $B = 0.70$, leading to a most likely signal box expectation of 9 data events. The 95% CL upper and lower bounds to the signal box expectation from the sideband data fit were 33 and 0, respectively. In this case, these fit values are not shown in the θ_{NuMI} plot for this mock data test since the range in expected signal box values overlaps with the background-only expectation for the signal box from Monte Carlo.

The signal box was opened and revealed 1 data event, consistent with no new physics. The corresponding discovery sensitivity p-value is equal to 0.49. An “observed” exclusion curve in the ALP model space was computed for the observation of 1 event in the signal box.




Figure B.3: Exclusion contours for the Higgs portal scalar heavy QCD axion parameter spaces. The blue bands are from the expected sensitivity of the search, as described in section 8. The black curves are drawn in accordance with the background-only, Throw 1 mock data, i.e. an observation of 0 events observed in the signal box.

B.4 Mock Data Test: Background + ALP Signal with $m_a = 400$ MeV, $f_a = 25000$ GeV, Throw 1

This mock data set and the mock data sets used in the following test were both sampled from the expectation for background + alps with model parameters $m_a = 400$, $f_a = 25000$ GeV (and the associated running coupling c_μ). This is a heavy QCD axion model benchmark that lies near the expected reach of sensitivity, meaning that if this model exists, one can expect a small number of signal events in the sideband, and just enough to see an excess above the background-only expectation in the signal box.

The best fit values to Equation 9.1 with the near sideband data were $A = 0.52$ and $B = 1.06$, leading to a most likely signal box expectation of 4 data events. The 95% CL upper and lower bounds to the signal box expectation from the sideband data fit were 20 and 0, respectively. In this case, these fit values are not shown in the θ_{NuMI} plot for this mock data test since the range in




Figure B.4: Same as Figure B.1, for Mock Data Test with Background Only, Throw 2.

expected signal box values overlaps with the background-only expectation for the signal box from Monte Carlo.

The signal box was opened and revealed 9 data events, an observation above the background-only expectation with a discovery sensitivity p-value of 0.0008. This p-value would correspond to a z-value of 3.16 for a one-sided Gaussian distribution centered at 0 with a width of 1. This result thus constitutes evidence for new physics with $\sigma = 3.16$.

B.5 Mock Data Test: Background + ALP Signal with $m_a = 400$,

$f_a = 25000$ GeV, Throw 2

The best fit values to Equation 9.1 with the near sideband data were $A = 1.21$ and $B = 1.66$, leading to a most likely signal box expectation of 18 data events. The 95% CL upper and lower bounds to the signal box expectation from the sideband data fit were 42 and 0, respectively. In this case, these fit values are not shown in the θ_{NuMI} plot for this mock data test since the range in expected signal box values overlaps with the background-only expectation for the signal box from Monte Carlo.




Figure B.5: Same as Figure B.2, for Mock Data Test with Background Only, Throw 2.

The signal box was opened and revealed 14 data events, an observation above the background-only expectation with a discovery sensitivity p-value of 6×10^{-6} . This p-value would correspond to a z-value of 4.34 for a one-sided Gaussian distribution centered at 0 with a width of 1. This result thus constitutes evidence for new physics with $\sigma = 4.34$.

B.6 Mock Data Test: Background + ALP Signal with $m_a = 400$, $f_a = 5000$ GeV, Throw 1

This mock data set was sampled from the expectation for background + alps with model parameters $m_a = 400$, $f_a = 5000$ GeV (and the associated running coupling c_μ). This is a heavy QCD axion model benchmark that could lead to $\mathcal{O}(100\text{s} - 1000\text{s})$ of events in the near sideband and signal box. This is the regime where the fit to sideband data is useful for predicting the signal box observation, and checking if the (real) signal box data is compatible with the fit and therefor with the heavy QCD axion model.

The best fit values to Equation 9.1 with the near sideband data were $A = 1.01$ and $B = 140.70$, leading to a most likely signal box expectation of 1352 data events. The 95% CL upper and lower bounds to the signal box expectation from the sideband data fit were 1589 and 1111, respectively.




Figure B.6: Same as Figure B.3, for Mock Data Test with Background Only, Throw 2.

These fit values do not overlap the background-only expectation for the signal box from Monte Carlo, so they are shown in the θ_{NuMI} plot for this mock data test.

The signal box was opened and revealed 1354 data events, an observation far above the background-only expectation with a discovery sensitivity p-value equal to zero (given the precision obtained from simulating 1,000,000 toy signal box values). If a scenario like this is encountered with the real data, more signal box toy values should be thrown so that a non-zero z value can be computed. In any case, such a scenario would constitute extremely strong evidence for BSM physics. Note that the signal box data here is within the expected range of values determined from the near sideband data fit, bolstering this discovery as consistent with the heavy QCD axion model.




Figure B.7: Same as Figure B.1, for Mock Data Test with Background Only, Throw 3.




Figure B.8: Same as Figure B.2, for Mock Data Test with Background Only, Throw 3.




Figure B.9: Same as Figure B.3, for Mock Data Test with Background Only, Throw 3.




Figure B.10: Same as Figure B.1, but for mock data corresponding to background + ALP Signal with $m_a = 400$, $f_a = 25000$ GeV, Throw 1.




Figure B.11: Same as Figure B.2, but for mock data corresponding to background + ALP Signal with $m_a = 400$, $f_a = 25000$ GeV, Throw 1.




Figure B.12: Same as Figure B.3, but for mock data corresponding to background + ALP Signal with $m_a = 400$, $f_a = 25000$ GeV, Throw 2




Figure B.13: Same as Figure B.2, but for mock data corresponding to background + ALP Signal with $m_a = 400$, $f_a = 25000$ GeV, Throw 2




Figure B.14: Same as Figure B.1, but for mock data corresponding to background + ALP Signal with $m_a = 400$, $f_a = 5000$ GeV, Throw 1. The range of predicted signal box values from the Likelihood fit to near sideband data (see Figure B.15) is shown in this θ_{NuMI} plot, since it does not overlap with the background-only expectation predicted by the Monte Carlo. Notably, the mock data value in the signal box is consistent with the range of values predicted by the fit to the near sideband data.




Figure B.15: Same as Figure B.2, but for mock data corresponding to background + ALP Signal with $m_a = 400$, $f_a = 5000$ GeV, Throw 1.