A Search for Sterile Neutrinos at the NOνA Far Detector

A DISSERTATION PRESENTED
BY
GARETH KAFKA
TO
THE DEPARTMENT OF PHYSICS

IN PARTIAL FULFILLMENT OF THE REQUIREMENTS
FOR THE DEGREE OF
DOCTOR OF PHILOSOPHY
IN THE SUBJECT OF
PHYSICS

HARVARD UNIVERSITY
CAMBRIDGE, MASSACHUSETTS

MAY 2016

©2016 – Gareth Kafka
all rights reserved.

Thesis advisor: Professor Gary Feldman Gareth Kafka

A Search for Sterile Neutrinos at the NOνA Far Detector

Abstract

NOνA is the current United States flagship long-baseline neutrino experiment designed to study the properties of neutrino oscillations. It consists of two functionally identical detectors each located 14.6 mrad off the central axis from the Fermilab NuMI neutrino beam. The Near Detector is located 1 km downstream from the beam source, and the Far Detector is located 810 km away in Ash River, Minnesota. This long baseline, combined with the ability of the NuMI facility to switch between nearly pure neutrino and anti-neutrino beams, allows NOνA to make precision measurements of neutrino mixing angles, potentially determine the neutrino mass hierarchy, and begin searching for CP violating effects in the lepton sector. However, NOνA can also probe more exotic scenarios, such as oscillations between the known active neutrinos and new sterile species.

This thesis showcases the first search for sterile neutrinos in a $3+1$ model at NOνA. The analysis presented searches for a deficit in the rate of neutral current events at the Far Detector using the Near Detector to constrain the predicted spectrum. The comparison between the observed and predicted spectra is translated into a measurement of the expanded PMNS mixing angles, θ_{34} and θ_{24} elements, and the matrix elements $|U_{\mu 4}|^2$ and $|U_{\tau 4}|^2$, assuming $0.05 < \Delta m^2_{41} < 0.5\,\text{eV}^2$. This analysis was performed using data taken between February 2014 and May 2016 corresponding to 6.69×10^{20} protons on target. No deficit of events was observed, leading only to upper limits on the sterile mixing angles, $\theta_{34} < 32°$ and $\theta_{24} < 20°$, and on the matrix elements, $|U_{\mu 4}|^2 < 0.12$ and $|U_{\tau 4}|^2 < 0.27$, all at the 90% confidence level. These measurements are consistent with the no sterile neutrino hypothesis. At the end of the dissertation there is a short discussion of future sensitivity improvements using a larger dataset.

Contents

Abstract iii

List of Figures ix

List of Tables xi

Dedication xi

Acknowledgments xiii

1 A Brief History of Neutrinos 1
- 1.1 Introduction . . . 1
- 1.2 First Detection of Neutrinos . . . 2
- 1.3 Evidence of Neutrino Oscillations . . . 2
- 1.4 Possible Evidence of Sterile Neutrinos . . . 4

2 Theory of Neutrino Oscillations 8
- 2.1 The PMNS Matrix . . . 8
- 2.2 Vacuum Oscillations . . . 9
- 2.3 Standard 3-Flavor Oscillations . . . 13
- 2.4 Matter Effects . . . 16
- 2.5 Current Measurements . . . 19
- 2.6 Sterile Neutrinos . . . 21
- 2.7 Sterile Matter Effect . . . 36

3 The NOνA Experiment 39
- 3.1 The NuMI Beam . . . 39
- 3.2 Off-Axis Detectors . . . 42
- 3.3 The NOνA Detectors . . . 45
  - 3.3.1 Far Detector . . . 51
  - 3.3.2 Near Detector . . . 52

4 Experiment Simulation 55
- 4.1 Flux Simulation . . . 56
- 4.2 Detector Simulation . . . 57
  - 4.2.1 Neutrino Interactions . . . 57
  - 4.2.2 Photon Propagation . . . 61
  - 4.2.3 Electronic Readout . . . 62
  - 4.2.4 Simulation Miscellany . . . 63
- 4.3 Studying Simulation Files . . . 64

5 Event Reconstruction 67
5.1 Calibration . . . 67
5.1.1 Relative Calibration . . . 68
5.1.2 Absolute Calibration . . . 72
5.2 Reconstruction Chain . . . 73
5.2.1 Event Slicing . . . 75
5.2.2 Hough Transform Line Finding . . . 78
5.2.3 Elastic Arms Vertex Reconstruction . . . 79
5.2.4 Fuzzy K-Means Clustering . . . 82
5.2.5 Kalman Tracking . . . 85
5.3 Particle Identification Algorithms . . . 87
5.3.1 ν_e and ν_μ CC Specific PIDs . . . 88
5.3.2 Cosmic Rejection PID . . . 89
5.3.3 CVN, the All Purpose PID . . . 90

6 Neutral Current Event Selection 96
6.1 A Note on Preselection . . . 97
6.2 Data Quality . . . 97
6.3 Event Quality . . . 99
6.4 Fiducial Volume and Containment . . . 99
6.4.1 Far Detector . . . 101
6.4.2 Near Detector . . . 102
6.5 NC Selection . . . 104
6.6 Cosmic Rejection . . . 106
6.7 Summary . . . 107

7 Neutral Current Spectrum Prediction 109
7.1 ND Decomposition . . . 110
7.2 Extrapolation . . . 111
7.3 Far Detector Prediction . . . 114

8 Systematic Uncertainty Analysis 117
8.1 Beam . . . 118
8.2 Birks-Chou Light Yield Simulation . . . 119
8.3 Calibration . . . 121
8.4 GENIE Simulation . . . 122
8.5 ND Containment . . . 126
8.6 ND Rock Event Contamination . . . 126
8.7 ND Data/MC Difference and CC Background . . . 128
8.8 MC Statistics . . . 129
8.9 Noise Model . . . 130
8.10 Overall Normalization . . . 131
8.11 Systematic Uncertainty Summary . . . 131

9 Analysis Results 133
9.1 Fitting Method . . . 134
9.2 Robustness Studies . . . 135
9.3 Sideband Studies . . . 140
9.4 Results . . . 142

10 Future Improvements and Conclusions 151
10.1 Future Improvements . . . 151
10.2 Conclusion . . . 153

References 161

List of Figures

1.1 SNO Result 5

2.1 Standard Model Neutrino Interaction Diagrams 9
2.2 Neutrino Mass Splitting Schematic 14
2.3 MSW Effect Interactions 16
2.4 First Measurement of θ_{13} from Daya Bay 20
2.5 Bi-Probability Plots for ν_e Appearance 21
2.6 First Measurement of δ by NOνA 22
2.7 Oscillation Probabilities for Values of Δm^2_{41} 28
2.8 Oscillation Probabilities for Values of θ_{34} 29
2.9 Oscillation Probabilities for Values of θ_{24} 30
2.10 Oscillation Probabilities for Values of θ_{23} 31
2.11 Oscillation Probabilities for Values of δ_{24} 32
2.12 Oscillation Probabilities for Values of θ_{14} 33
2.13 Oscillation Probabilities for Values of θ_{13} 34
2.14 Oscillation Probabilities for Values of δ_{13} 35
2.15 The Sterile Matter Effect 38

3.1 Fermilab Accelerator Complex 40
3.2 Beam Power vs Time 41
3.3 NuMI Beam Schematic 42
3.4 Off-Axis Flux and Energy 43
3.5 Beam Energy and Composition 44
3.6 Probability of ν_e Appearance and Expected Event Rates for NOνA 44
3.7 A PVC Extrusion 45
3.8 A Detector Module 46
3.9 A Detector Cell 47
3.10 An APD 48
3.11 Quantum Efficiency of APDs and Light Intensity Output of WLS fibers 49
3.12 The NOνA Far Detector 51
3.13 The NOνA Near Detector 53
3.14 Schematic of the NuMI Beam line 53

4.1 Model of the Second Focusing Horn 56
4.2 Data/MC Comparison With and Without MEC Events 59
4.3 ND Neutrino Interaction Vertex Distribution 60
4.4 Simulation of Photon Transport 62
4.5 Identifying Missing Simulated Events in the FD 66
4.6 Identifying Missing Simulated Events in the ND 66

5.1 Threshold and Shielding Correction Factors . . . 69
5.2 Attenuation Fits . . . 71
5.3 Relative Calibration Results . . . 71
5.4 Detector Response to Stopping Cosmic Muons vs Distance to Track End . . . 73
5.5 Absolute Calibration Results . . . 74
5.6 An Example NOνA Event . . . 76
5.7 An Example NOνA Event with Slicer Applied . . . 77
5.8 An Example Neutrino Candidate Slice . . . 77
5.9 Hough Line Generation from a Hough Map . . . 79
5.10 An Example Neutrino Candidate Slice with Hough Lines . . . 80
5.11 An Example Neutrino Candidate Slice with a Reconstructed Vertex . . . 83
5.12 An Example Neutrino Candidate Slice with Prongs . . . 86
5.13 An Example Neutrino Candidate Slice with Tracks . . . 88
5.14 Cosmic Rejection BDT Output Distributions . . . 91
5.15 Inception Module Schematic . . . 92
5.16 CVN Architecture Schematic . . . 93
5.17 Feature Maps from Representative Events . . . 95

6.1 Energy Spectra After Data Quality Cuts . . . 98
6.2 Energy Spectra After Event Quality Cuts . . . 100
6.3 Fiducial and Containment Variable Distributions . . . 101
6.4 Energy Spectra After Fiducial and Containment Cuts, FD . . . 102
6.5 Energy Spectra After Fiducial and Containment Cuts, ND . . . 103
6.6 NC Selection Variable Distributions . . . 104
6.7 Energy Spectra After NC Selection Cuts . . . 105
6.8 LID and RemID Distributions at the FD . . . 105
6.9 Cosmic Rejection Variable Distributions . . . 107
6.10 Energy Spectra After All Cuts . . . 108

7.1 ND Decomposition . . . 111
7.2 F/N Ratio and ND Data/MC Ratio for NC Component Extrapolation . . . 113
7.3 Effect of the Extrapolation on the NC Component Distribution . . . 114
7.4 NC Component Flavor Proportions . . . 115
7.5 Predicted Component Distributions . . . 116

8.1 Beam Systematic Uncertainty Envelopes . . . 120
8.2 Birks-Chou Shifted FD Predictions . . . 121
8.3 Shifted FD Predictions Due to Miscalibration at Both Detectors . . . 123
8.4 Shifted FD Predictions Due to Miscalibration at One Detector . . . 123
8.5 GENIE Systematic Uncertainty Envelopes . . . 126
8.6 Shifted FD Predictions from Extrapolation of Halves of the ND . . . 127
8.7 ND Rock Contamination Shifted Spectra . . . 128
8.8 ND Data/MC Energy Spectrum Comparison . . . 129
8.9 Systematic Uncertainty Due to ND Data/MC Discrepancy and CC Background Uncertainty 130
8.10 ND Data/MC Spectra with Systematics . . . 132

9.1 ND Data/MC Comparison: Energy Distribution . . . 136
9.2 ND Data/MC Comparison: Fiducial and Containment Variable Distributions . . . 137
9.3 ND Data/MC Comparison: NC Selection Variable Distributions . . . 138
9.4 ND Data/MC Comparison: Cosmic Rejection Variable Distribution . . . 138
9.5 Angle Sensitivities for Shifted Energy Spectra . . . 139
9.6 Angle Sensitivity With and Without Profiling δ_{24} . . . 140
9.7 High Energy Sideband . . . 141
9.8 Low CVN Sideband . . . 142
9.9 Mid Cosmic BDT Sideband . . . 143
9.10 FD Data and MC Energy Distribution . . . 144
9.11 FD Data/MC Comparison: Fiducial Variable Distributions . . . 144
9.12 FD Data/MC Comparison: CVN Distribution . . . 145
9.13 FD Data/MC Comparison: Cosmic Rejection Variable Distributions . . . 145
9.14 FD Data Vertex Distributions . . . 146
9.15 Two Dimensional θ_{34} vs θ_{24} Contours . . . 147
9.16 One Dimensional θ_{34} and θ_{24} Slices . . . 147
9.17 FC Corrected Significance Surface and 68% and 90% Contours . . . 149
9.18 FC Corrected One Dimensional θ_{34} and θ_{24} Slices . . . 150

10.1 Projected Sensitivity for θ_{34} After 18×10^{20} POT. . . . 152

List of Tables

2.1 Best Fit Parameters for Three Neutrino Oscillation Model . . . 21
2.2 Four Flavor Fixed Oscillation Parameters . . . 27

3.1 Scintillator Solution Summary . . . 47
3.2 Detector Properties . . . 54

5.1 Slicing Algorithm Free Parameters . . . 76
5.2 Results from Elastic Arms Vertex Fits . . . 82

6.1 Assumed Oscillation Parameters . . . 97
6.2 Beam Quality Cuts . . . 98
6.3 Data Quality Cuts . . . 98
6.4 Event Table: Data Quality Cuts . . . 98
6.5 Event Quality Cuts . . . 100
6.6 Event Table: Event Quality Cuts . . . 100
6.7 Implicit Event Quality Cuts . . . 100
6.8 FD Fiducial and Containment Cuts . . . 102
6.9 Event Table: Fiducial and Containment Cuts, FD . . . 102
6.10 ND Fiducial and Containment Cuts . . . 103
6.11 Event Table: Fiducial and Containment Cuts, ND . . . 103
6.12 NC Selection Cuts . . . 104
6.13 Event Table: NC Selection Cuts . . . 105
6.14 Final Selection Cuts . . . 106
6.15 Event Table: FD Cosmic Rejection Cuts . . . 106
6.16 Event Table: FD Selection Cuts . . . 108
6.17 Event Table: ND Selection Cuts . . . 108

7.1 Extrapolation Method and Rate Summary . . . 116

8.1 Beam Systematic Uncertainties . . . 119
8.2 Birks-Chou Systematic Uncertainties . . . 121
8.3 Calibration Systematic Uncertainties . . . 122
8.4 GENIE Systematic Uncertainties . . . 125
8.5 ND Containment Systematic Uncertainties . . . 127
8.6 ND Data/MC and CC Background Uncertainties . . . 130
8.7 Systematic Uncertainty Summary . . . 132

9.1 Fixed Parameters and Values for Fitting . . . 134
9.2 FD Event Rates for Shifted Energy Spectra . . . 139
9.3 FD Event Rates for Selected Values of δ_{24} . . . 140

9.4 Sideband Event Rates 141

To my father, from whom my love of logic was born.
To my wife, Kristen, for an unyielding love I look forward to sharing every day.

Acknowledgments

There are an incredible number of people who helped me along the path to earning a doctorate. I wish I could recognize and thank everyone here, but the list of people who have supported me personally and professionally throughout my life is impossibly long. What follows is certainly an incomplete list; to everyone else, know that I appreciate the presence and impact you have had in my life.

First and foremost, I would like to thank my research advisor, Gary Feldman, for your advice and support throughout my time at Harvard. Your input, understanding, and guidance helped me grow and find success. Thank you to my other committee members, Melissa Franklin and Howard Georgi. I am grateful for your physics advice and for the positive environment you helped to shape in the physics department. I would also like to thank my original graduate academic advisor, Matthew Schwartz, for your support and encouragement through my initial years as a graduate student.

Thank you to the NOνA collaboration, whose work produced the wonderful experiment I was able to be part of. In particular, thank you to the NC/Sterile group conveners, Alex Sousa and Louise Suter, for your guidance and support, especially during the analysis push. Thank you to the other contributors to the NC disappearance analysis, Adam Aurisano, Aaron Markowitz, and Shaokai Yang. The analysis had one other major contributor who I believe deserves special recognition, Gavin Davies. Thank you for your part in the analysis, being a friend, allowing me to vent my frustrations, and for all your help before and during my work on the NC disappearance analysis; I owe you a great debt of gratitude. Thank you also to my fellow collaborators Michael Baird, Xuebing Bu, Joseph Lozier, Evan Niner, Dan Pershey, Fernanda Psihas, Justin Vasel, for your friendship.

The post docs I worked with in the Harvard neutrino group were truly instrumental to my graduate work. Thank you to Raphael Schroeter and Ruth Toner; you both truly cared about my success and I was able to achieve it with your guidance and support. I will call out Alex Sousa again here-thank you for your

part in providing a positive research experience (and for sharing your fantastic knowledge of restaurants in Naperville) while we overlapped at Harvard.

There are several other people in the Harvard Department of Physics who deserve recognition. Thank you to Jacob Barandes for your advice throughout the graduate process. Thank you, Lisa Cacciabaudo and Carol Davis, for your open doors and tasty treats. Many graduate students shared late nights with me as we worked through problem sets. Thank you specifically to Mike Goldman who worked with me through several courses and has remained a good friend. Thank you to my current and former Palfrey house denizens, Stephan Chan, Tomo Lazovich, Chris Rogan, Emma Tolley, and Ann Wang. I will fondly remember our Taco Tuesdays, especially the ones on the Felipe's roof deck.

I wish to thank two people from my time as an undergraduate at CWRU, my advisor, Glenn Starkman, and my senior project advisor, Robert "Doc" Brown. Both of you encouraged me personally and professionally, and I owe much of my success to the mentorship, opportunities, and encouragement you provided me.

Last, but certainly not least, is my family. Thank you to my mother and father, Shelly and Kris Kafka, for all of your encouragement through my entire life. You always believed in me and pushed me to achieve greatness. Both of you have had tremendous influence on how I approach the world, and I cannot imagine how I would have accomplished as much without your impact and love. Thank you to my brother Wesley for inspiring me to be the best version of myself possible. I want always to be a good role model for you, and it brings me endless pride to see what you have accomplished so far. To my adopted brothers, Scott and Cody, thank you for always welcoming me into a home away from home. I am now a fully fledged wizard. Thank you to Karen and Glenn Berry, my mother-in-law and father-in-law, for always being so welcoming and supportive. Finally, I extend a most heartfelt thanks to my wife, Kristen. You provided comfort on my worst days and amplified the triumphs of my best days, always having faith in me. Your constant love and support gave me the strength to complete my degree.

1

A Brief History of Neutrinos

1.1 Introduction

The neutrino was first postulated by Wolfgang Pauli as a possible explanation for the continuous spectrum of electrons emitted from nuclear β decay [1]. This decay was originally thought to be the emission of an electron from an atom, resulting in a different nucleus, via the process,

$$N \to N' + e \tag{1.1}$$

where N and N' are the parent and daughter nuclei, respectively. In a two body decay such as this, the momenta and energies of the outgoing particles are exactly constrained. Pauli's new particle explained the continuous spectrum of electron energy via a modified decay process:

$$N \to N' + e + \nu \tag{1.2}$$

where ν is the outgoing neutral particle. Pauli's original proposal called the new particle the neutron, but this name was later used to name the massive neutral nucleon discovered by Chadwick in 1932 [2]. Three years after Pauli's idea, Fermi proposed a model for nuclear β decay that included the new particle, which he coined the neutrino, or little neutral one [3].

1.2 First Detection of Neutrinos

Twenty years passed from Fermi's model proposal before neutrinos were discovered experimentally. Fred Reines and Clyde Cowan made the discovery by placing a detector near a nuclear reactor as a source of neutrinos and observing inverse β decay [4, 5]. The neutrinos observed were anti-electron neutrinos, thus the following was the observed process.

$$p + \bar{\nu}_e \rightarrow n + e^+ \tag{1.3}$$

Reines earned the Nobel Prize in Physics in 1995 for the detection of the neutrino.

In 1962, the muon neutrino was discovered at Brookhaven National Laboratory using the first neutrino beam [6] in a scheme still used in neutrino experiments today. The beam was generated by colliding protons with a target, producing pions that decayed into muons and muon neutrinos. The resultant beam then passed through thick steel, absorbing everything but the neutrinos. Leon Lederman, Melvin Schwartz, and Jack Steinberger won the Nobel Prize in Physics in 1988 for the discovery of the muon neutrino.

The last generation of neutrino, the tau neutrino, was known to exist by 1982 [7] as measurements of the τ decay were inconsistent with the existence of only two neutrinos. It was discovered at Fermilab by the DONUT collaboration in 2000 [8].

1.3 Evidence of Neutrino Oscillations

Pontecorvo first postulated neutrino oscillations between neutrinos and anti-neutrinos, analogous to $K^0/\bar{K}^0$ oscillations, in 1957 [9]. Nothing came of the proposal immediately, but the idea was later re-

vived and modified to solve the solar neutrino problem. The physics community initially viewed neutrino oscillations with skepticism and believed the experiments to be flawed, but over time oscillations have become an unmistakable and accepted phenomenon.

The solar neutrino problem was born from a large discrepancy between the theoretical and observed number of neutrinos produced by the sun. Neutrinos were used as a study for solar models because photons take a thousand years to escape the dense nuclear plasma to the surface of the sun, but neutrinos are unimpeded. The models, which have been confirmed today, describe a somewhat complicated chain of nuclear reactions, many of which produce neutrinos. Each individual process contributes a neutrinos in a different energy spectrum, but all of the neutrinos are created as electron neutrinos.

The experimental observations and theoretical predictions were both published in 1968. Ray Davis designed an experiment underground in the South Dakota Homestake mine consisting of a tank of an ultra pure chlorine cleaning solution capable of neutrino capture via the process

$$\nu_e + {}^{37}\mathrm{Cl} \rightarrow {}^{37}\mathrm{Ar} + e^- \tag{1.4}$$

The argon atoms could be collected and counted for a direct measurement of the neutrino flux [10]. Meanwhile, John Bahcall precisely calculated the expected neutrino flux [11], and the observed rate was found to be about one third of the predicted rate. Pontecorvo revived his theory with the modification of allowing ν_e to ν_μ oscillations [12], but the idea was still not taken seriously and it was another 20 years before the solar neutrino problem was confirmed.

Beginning in 1989, multiple experiments with different methodologies confirmed the solar neutrino problem. Kamiokande, a water Cherenkov detector, measured a rate deficit in 1989 [13]. Two experiments measured solar neutrinos via the reaction

$$\nu_e + {}^{71}\mathrm{Ga} \rightarrow {}^{71}\mathrm{Ge} + e^- \tag{1.5}$$

and measured a similar deficit, SAGE in 1991 [14] and GALLEX in 1992 [15]. With results from three different experimental methods all showing similar rate deficits, the solar neutrino problem could no longer

be relegated to an experimental error.

Evidence soon emerged for oscillations with atmospheric neutrinos as well. These neutrinos are produced when cosmic rays collide with particles in the atmosphere and decay, predominantly via the following channels.

$$\begin{aligned} \pi^{+/-} &\rightarrow \mu^{+/-} + \nu_\mu/\bar{\nu}_\mu \\ \mu^{+/-} &\rightarrow e^{+/-} + \nu_e/\bar{\nu}_e + \bar{\nu}_\mu/\nu_\mu \end{aligned} \tag{1.6}$$

Thus, the expected ratio of muon family neutrinos to muon family neutrinos was expected to be 2. Kamiokande measured this ratio in 1992 and found the ratio to be much closer to 1 [16]. Furthermore, the ratio seemed to be dependent on zenith angle, with the measurement being nearly 2 for neutrinos coming from directly overhead, and dropping as the angle increased. Super-Kamiokande (or Super-K), the successor to Kamiokande, improved upon this measurement in 1998 [17], providing the most definitive evidence of neutrino oscillations to that point.

A resolution to the solar neutrino problem did not have to wait much longer with detector technologies capable of discerning different neutrino interaction types. SNO was designed as a heavy water (D_2O) Cherenkov detector experiment to be sensitive to both the flux of electron neutrinos and the flux of all neutrinos. In 2002, it released results for these measurements, finding what was then the expected deficit in electron neutrino flux, but a total flux consistent with the standard solar model, see figure 1.1 [18]. With this result, neutrino oscillations were confirmed, and subsequent experiments now measure oscillation parameters with precision.

1.4 Possible Evidence of Sterile Neutrinos

Several experiments have reported results that can be explained by oscillations between more than three neutrinos. The frequency of oscillations is controlled by the mass splitting between neutrino mass states, Δm^2, and will be considered in more detail in chapter 2. Historically, solar and atmospheric neutrino oscillations were discovered first and were well explained by oscillations between the standard model's three

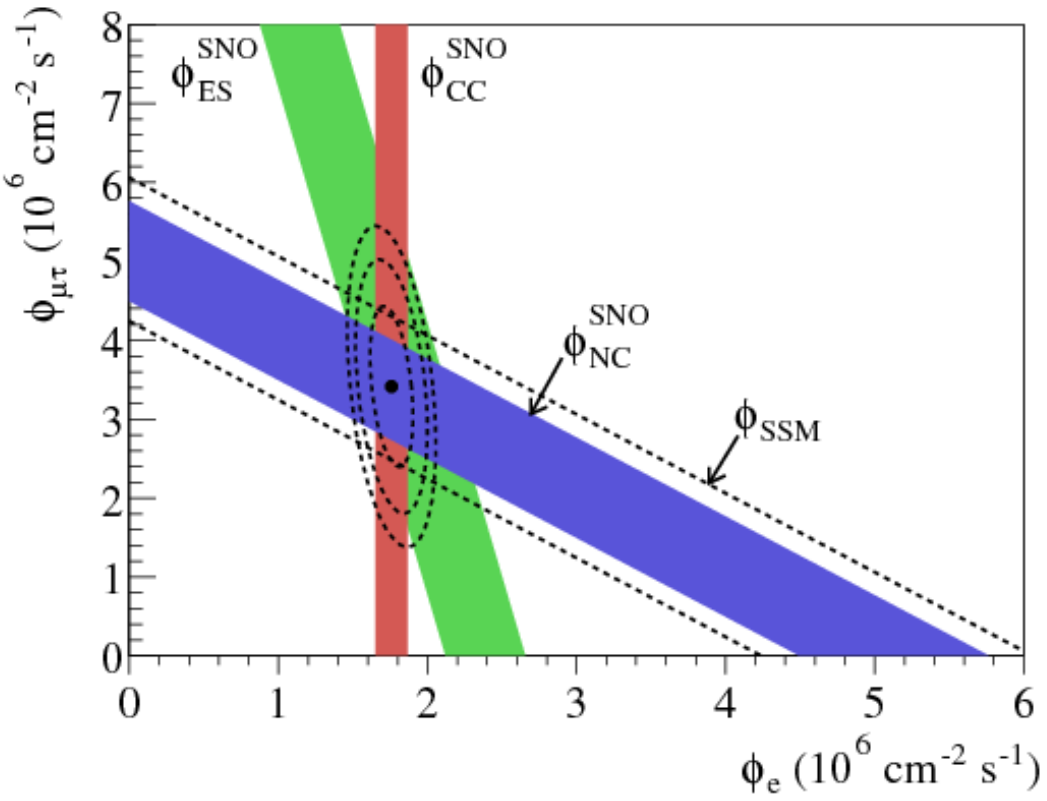


Figure 1.1: The measurement of different event rates at SNO [18]. The red band represents ν_e CC interactions with the deuterium neutron, an interaction only sensitive to electron neutrinos. The blue band represents neutral current scattering off of the deuterium nucleus, an interaction sensitive to the total neutrino flux. The green band represents elastic scattering of the neutrino off the deuterium electron, an interaction sensitive to all neutrino flavors, but not completely independent of neutrino flavor. The dashed straight lines represent the flux prediction by the standard solar model. The point represents the best fit for the flux of electron neutrinos and the flux for the combined muon and tau neutrinos.

active neutrinos, and with three neutrinos, there are only two independent mass splittings. Furthermore, it is known that there can only be three active neutrinos. However, there have been some experimental results that cannot be explained by oscillations based on the solar and atmospheric mass splitting scales, but they can be explained by oscillations based on a third independent mass splitting.

The number of active neutrinos is constrained by measurements of the width of the Z boson. LEP has measured the number of active neutrinos to be 2.984 ± 0.008 [19], so the discoveries of the ν_e, ν_μ, and ν_τ leave no room for new active neutrinos. Strictly speaking, there could be other active neutrinos if they had mass greater than half the mass of the Z boson so the Z could not decay to them, but while current evidence hinting at a sterile neutrino would suggest a mass splitting much larger than the solar and

atmospheric scales, it is still much smaller than the mass of the Z.

The first evidence for an additional neutrino came from the Liquid Scintillator Neutrino Detector, or LSND, in 1995. This experiment searched for $\bar{\nu}_e$ appearance in a $\bar{\nu}_\mu$ beam. When it found an excess of events, it reported a measurement of a mass splitting between neutrino states Δm^2 of $O(1\,\text{eV}^2)$ [20], a measurement incompatible with the mass splittings measured in both the atmospheric and solar oscillation experiments, suggesting the addition of at least one more neutrino. However, based on the result from LEP, this new neutrino cannot couple to the Z boson, hence the suggestion for a sterile neutrino.

Searches for sterile neutrinos can be categorized in one of three channels: $\overset{(}{\bar{\nu}}_\mu \rightarrow \overset{(}{\bar{\nu}}_e$ appearance, $\overset{(}{\bar{\nu}}_e$ disappearance, and $\overset{(}{\bar{\nu}}_\mu$ disappearance. The LSND result was from the appearance channel. The MiniBooNE experiment at Fermilab searched for both ν_e appearance in a ν_μ beam and $\bar{\nu}_e$ appearance in an $\bar{\nu}_\mu$ beam. While they first reported no event excess in 2007 [21], their more results show excesses in both modes [22] that could be consistent with some sterile neutrino models. Not all experiments have found evidence of sterile neutrinos, however. The most recent results from KARMEN in 2002 [23] and NOMAD in 2003 [24], both measuring the appearance channel, showed no evidence of oscillations at the same mass scale as LSND.

Experiments measuring the $\overset{(}{\bar{\nu}}_e$ disappearance channel have also reported some indication for a neutrino at the same mass scale as the appearance channel. The so-called Gallium anomaly comes from a measured deficit of events in GALLEX [25] and Sage [26] when placing a radioactive source inside the detectors for detector response measurements. On the other hand, reactor experiments such as Bugey [27] reported no oscillations and RENO [28] and Daya Bay [29] reported only oscillations consistent with standard three flavor oscillations. Furthermore, T2K, an accelerator experiment not unlike NOνA, performed a short baseline analysis at its Near Detector and also found no evidence of sterile oscillations at 95% confidence.

To date, experiments probing the $\overset{(}{\bar{\nu}}_\mu$ disappearance channel have universally shown no evidence of sterile neutrinos. This includes an atmospheric neutrino measurement by Super-K [30] and a long baseline analysis by MINOS [31]. This list of experimental evidence for sterile neutrinos (or lack thereof) is not intended to be exhaustive, but it should be clear that there is not yet a scientific consensus on the existence of sterile oscillations despite the variety of past and current searches.

NOνA is a long baseline neutrino experiment with a near detector, thus capable of performing both short baseline and long baseline sterile neutrino analyses; this dissertation focuses on the long baseline. It measures events generated by the NuMI beam, a mostly pure ν_μ beam, and thus measures the ν_μ disappearance channel. (In the future, the beam will switch to $\bar{\nu}_\mu$ allowing for measurement in the $\bar{\nu}_\mu$ disappearance channel.) The analysis presented here searches for a deficit in the number of neutral current events at the far detector, using the near detector data to constrain the predicted spectrum. Neutral currents are insensitive to the flavors of the standard 3 active neutrinos, so a rate deficit would point to the existence of oscillations between active and sterile neutrinos.

2

Theory of Neutrino Oscillations

The idea of neutrino oscillations was first proposed by Pontecorvo in 1957 [9], but his proposal described oscillations between neutrinos and anti-neutrinos. In 1962, after the discovery of the muon neutrino, Maki, Nakagawa, and Sakata proposed the theory that described oscillations between neutrino flavors due to differing neutrino flavor and mass eigenstates [32]. This chapter describes the modern formalism in detail and uses natural units where $\hbar = c = 1$, except where otherwise noted.

2.1 The PMNS Matrix

In the Standard Model, neutrinos only interact via the W and Z bosons as shown by the Feynman diagrams in figure 2.1. From these diagrams, it is clear that neutrinos always interact in a definite flavor eigenstate, $|\nu_\alpha\rangle$. Furthermore, when a neutrino is produced from a W boson, the flavor is always determined by the

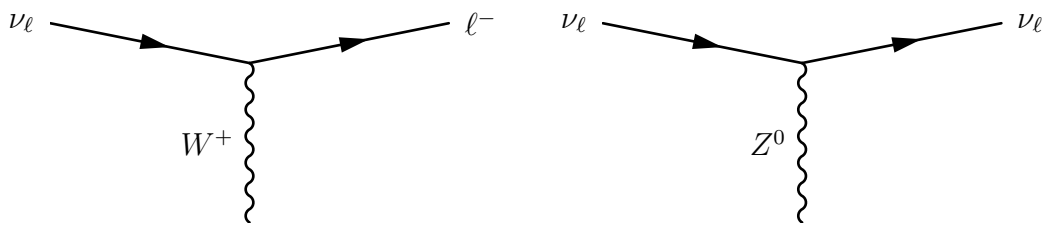


Figure 2.1: Standard Model Weak interactions involving a neutrino. Left: Charged Current interaction. Right: Neutral Current interaction.

associated charged lepton shown in equation 2.1.

$$\begin{pmatrix} \nu_e \\ e \end{pmatrix}, \quad \begin{pmatrix} \nu_\mu \\ \mu \end{pmatrix}, \quad \begin{pmatrix} \nu_\tau \\ \tau \end{pmatrix} \tag{2.1}$$

On the other hand, neutrinos propagate through spacetime with a definite mass, $|\nu_i\rangle$ an eigenstate of the free Hamiltonian. The flavor states can be written as a superposition of the mass states via

$$|\nu_\alpha\rangle = \sum_{i=1}^{n} U^*_{\alpha i} |\nu_i\rangle , \tag{2.2}$$

where n is the number of neutrinos, and U is the unitary PMNS matrix, named after Pontecorvo, Maki, Nakagawa, and Sakata. The PMNS matrix is unitary, and would reduce to the identity matrix if neutrinos did not oscillate between flavor states. Yet since it does provide the mechanism for flavor transitions, it can be thought of as analogous to the quark sector CKM matrix.

2.2 Vacuum Oscillations

In this section, the basics of neutrino oscillations are developed by considering oscillations in a vacuum. The neutrinos are treated as plane waves, as in [33], with the assumption that the neutrino is actually

localized in space put in by hand. A careful, rigorous analysis treating the neutrinos as wave packets in [34] reproduces the same results.

Consider a neutrino in a state of definite flavor α at time $t = 0$, $|\nu(0)\rangle = |\nu_\alpha\rangle$. This state is in a superposition of mass eigenstates. The time evolution of this neutrino is simply the time evolution of the individual mass states. In a vacuum, this adds a phase factor to each mass state.

$$|\nu_\alpha(t)\rangle = \sum_i U^*_{\alpha i} e^{-i(E_i t - \mathbf{p_i}\cdot\mathbf{x})} |\nu_i\rangle \tag{2.3}$$

With the neutrino at position $\mathbf{x} = L$ at time t, the dot product evaluates to $\mathbf{p_i} \cdot \mathbf{x} = p_i L$. Eq 2.3 can then simplified by making use of the fact that neutrinos are ultra-relativistic, allowing for several assumptions. First, the time, t, is replaced by the distance, L. Next, the energy of each mass state is approximated to be the same energy, $E_i = E$. Last, the momentum is expanded as $p_i = \sqrt{E^2 - m_i^2} \approx E - m_i^2/2E$. With these assumptions, equation 2.3 simplifies as:

$$|\nu_\alpha(L)\rangle = \sum_i U^*_{\alpha i} e^{-i m_i^2 L/2E} |\nu_i\rangle \,. \tag{2.4}$$

The mass eigenstate inside the sum is then re-expressed in terms of flavor eigenstates using the inverse of equation 2.2 and unitarity of U.

$$|\nu_\alpha(L)\rangle = \sum_{\alpha'} \sum_i U^*_{\alpha i} U_{\alpha' i} e^{-i m_i^2 L/2E} |\nu'_\alpha\rangle \,. \tag{2.5}$$

Equation 2.5 can then be used to find the probability that the original neutrino in flavor state α has transitioned (or survived) as flavor state β. First, the matrix element $\langle\nu_\alpha|\nu_\beta(L)\rangle$ is computed.

$$\langle\nu_\alpha|\nu_\beta(L)\rangle = \sum_{\beta'} \sum_i U_{\beta' i} U^*_{\beta i} e^{-i m_i^2 L/2E} \langle\nu_\alpha|\nu'_\beta\rangle = \sum_i U_{\alpha i} U^*_{\beta i} e^{-i m_i^2 L/2E} \tag{2.6}$$

The last equality in equation 2.6 follows from the orthogonality of individual flavor eigenstates. The prob-

ability of the flavor transition is then the square of this matrix element.

$$P(\nu_\alpha \to \nu_\beta) = |\, \langle \nu_\alpha | \nu_\beta(L) \rangle \,|^2 = \sum_{i,j} U_{\alpha i} U^*_{\beta i} U_{\beta j} U^*_{\alpha j} e^{-i(m_i^2 - m_j^2)L/2E} \tag{2.7}$$

It is standard to rewrite the mass squared difference as $\Delta m^2_{ij} \equiv m_i^2 - m_j^2$. Equation 2.7 is then manipulated using the properties of unitary matrices.

$$\begin{aligned} P(\nu_\alpha \to \nu_\beta) &= \sum_{i,j} U_{\alpha i} U^*_{\beta i} U_{\beta j} U^*_{\alpha j} + \sum_{i,j} U_{\alpha i} U^*_{\beta i} U_{\beta j} U^*_{\alpha j} (e^{-i\Delta m^2_{ij} L/2E} - 1) \\ &= \delta_{\alpha\beta} + \sum_{i,j} U_{\alpha i} U^*_{\beta i} U_{\beta j} U^*_{\alpha j} (e^{-i\Delta m^2_{ij} L/2E} - 1) \end{aligned} \tag{2.8}$$

The remaining summed term is further simplified making use of two facts. When $i = j$, the complex phase is 0 as $\Delta m^2_{ii} = 0$, and thus these terms vanish. Second, the terms with $i < j$ are complex conjugates of those with $i > j$, and $z + z^* = 2\Re(z)$ for any complex number z.

$$P(\nu_\alpha \to \nu_\beta) = \delta_{\alpha\beta} + 2 \sum_{i>j} \Re \left[U_{\alpha i} U^*_{\beta i} U_{\beta j} U^*_{\alpha j} (e^{-i\Delta m^2_{ij} L/2E} - 1) \right] \tag{2.9}$$

Both pieces of this term are split into their real and imaginary parts, and simplified using the trigonometric identity $\cos 2\theta - 1 = -2\sin^2\theta$. Defining $\mathcal{U} \equiv U_{\alpha i} U^*_{\beta i} U_{\beta j} U^*_{\alpha j} (e^{-i\Delta m^2_{ij} L/2E} - 1)$ and $\phi \equiv \Delta m^2_{ij} L/2E$:

$$\Re(\mathcal{U}) = \Re \left[U_{\alpha i} U^*_{\beta i} U_{\beta j} U^*_{\alpha j} (e^{-i\Delta m^2_{ij} L/2E} - 1) \right] \tag{2.10}$$

$$= \Re \left\{ \left[\Re(U_{\alpha i} U^*_{\beta i} U_{\beta j} U^*_{\alpha j}) + i\Im(U_{\alpha i} U^*_{\beta i} U_{\beta j} U^*_{\alpha j}) \right] \left[-2\sin^2(\phi/2) - i\sin\phi \right] \right\} \tag{2.11}$$

$$= -2\Re(U_{\alpha i} U^*_{\beta i} U_{\beta j} U^*_{\alpha j}) \sin^2(\phi/2) + \Im(U_{\alpha i} U^*_{\beta i} U_{\beta j} U^*_{\alpha j}) \sin\phi \tag{2.12}$$

Inserting the expression from equation 2.12 into equation 2.9, we find:

$$P(\nu_\alpha \to \nu_\beta) = \delta_{\alpha\beta} \quad - \quad 4\sum_{i>j} \Re(U_{\alpha i}U^*_{\beta i}U_{\beta j}U^*_{\alpha j}) \sin^2 \Delta_{ij} \\ + \quad 2\sum_{i>j} \Im(U_{\alpha i}U^*_{\beta i}U_{\beta j}U^*_{\alpha j}) \sin 2\Delta_{ij}, \tag{2.13}$$

where $\Delta_{ij} \equiv \Delta m^2_{ij} L/4E = 1.267\Delta m^2_{ij} L\,(\text{km})/E\,(\text{GeV})$. It can now be seen that the distance the neutrino travels, its energy, and the different mass splittings all affect the frequency of oscillation. Ideally, neutrino oscillations would be studied by having neutrinos with a fixed energy profile (preferably monoenergetic) and varying the baseline. However, neutrino detectors are incredibly large, so in practice the baseline is fixed and the oscillation probability is studied as a function of neutrino energy.

For the case of survival probability, $\alpha = \beta$ and equation 2.13 simplifies further. The imaginary piece from equation 2.13 drops out, as

$$\Im(U_{\alpha i}U^*_{\alpha i}U_{\alpha j}U^*_{\alpha j}) = \Im(|U_{\alpha i}|^2|U_{\alpha j}|^2) = 0. \tag{2.14}$$

The survival probability is then given by:

$$P(\nu_\alpha \to \nu_\alpha) = 1 - 4\sum_{i>j} |U_{\alpha i}|^2|U_{\alpha j}|^2 \sin^2 \Delta_{ij}. \tag{2.15}$$

Due to the combined influence of mass splitting, oscillation baseline, and neutrino energy on the oscillation probability, it is often the case that only one term contributes to the sums in equations 2.13 and 2.15. The two neutrino approximation can be instructive in this instance. For this model, the mixing matrix simplifies to the two dimensional rotation matrix:

$$U = \begin{pmatrix} \cos\theta & \sin\theta \\ -\sin\theta & \cos\theta \end{pmatrix}. \tag{2.16}$$

As this matrix is entirely real, the imaginary piece of equation 2.13 drops out. Plugging the matrix elements into the remaining term directly and simplifying slightly, we find the following forms for the survival and

appearance probabilities.

$$P(\nu_\alpha \to \nu_\alpha) = 1 - \sin^2 2\theta \sin^2\left(\frac{\Delta m^2 L}{4E}\right) \tag{2.17}$$

$$P(\nu_\alpha \nrightarrow \nu_\alpha) = \sin^2 2\theta \sin^2\left(\frac{\Delta m^2 L}{4E}\right) \tag{2.18}$$

From these equations it is clear that the mixing matrix parameters control the amplitude of neutrino oscillations. For small angles, most neutrinos will not change flavor, while larger angles can cause most of the neutrinos to change flavor. The case where $\theta = 45°$ is called maximal mixing as at specific baseline lengths the probability of oscillation becomes 1.

2.3 Standard 3-Flavor Oscillations

The Standard Model includes three neutrinos, so the PMNS matrix is 3×3 in this picture. Explicitly expanding equation 2.2, U takes the following form:

$$\begin{pmatrix} \nu_e \\ \nu_\mu \\ \nu_\tau \end{pmatrix} = \begin{pmatrix} U_{e1} & U_{e2} & U_{e3} \\ U_{\mu 1} & U_{\mu 2} & U_{\mu 3} \\ U_{\tau 1} & U_{\tau 2} & U_{\tau 3} \end{pmatrix} \begin{pmatrix} \nu_1 \\ \nu_2 \\ \nu_3 \end{pmatrix}. \tag{2.19}$$

The PMNS matrix can be parametrized in terms of 3 real mixing angles, θ_{ij} and a complex phase, δ, called the CP phase. Following the convention from the Particle Data Group [35], the expanded matrix takes the form

$$\begin{aligned} U &= \begin{bmatrix} c_{12}c_{13} & s_{12}c_{13} & s_{13}e^{-i\delta} \\ -s_{12}c_{23} - c_{12}s_{23}s_{13}e^{i\delta} & c_{12}c_{23} - s_{12}s_{23}s_{13}e^{i\delta} & s_{23}c_{13} \\ s_{12}s_{23} - c_{12}c_{23}s_{13}s^{i\delta} & -c_{12}s_{23} - s_{12}c_{23}s_{13}e^{i\delta} & c_{23}c_{13} \end{bmatrix} \\ &= \begin{bmatrix} 1 & 0 & 0 \\ 0 & c_{23} & s_{23} \\ 0 & -s_{23} & c_{23} \end{bmatrix} \begin{bmatrix} c_{13} & 0 & s_{13}e^{-i\delta} \\ 0 & 1 & 0 \\ -s_{13}e^{i\delta} & 0 & c_{13} \end{bmatrix} \begin{bmatrix} c_{12} & s_{12} & 0 \\ -s_{12} & c_{12} & 0 \\ 0 & 0 & 1 \end{bmatrix} \end{aligned} \tag{2.20}$$

where $c_{ij} \equiv \cos\theta_{ij}$ and $s_{ij} \equiv \sin\theta_{ij}$.

With three neutrinos, the expanded forms of equations 2.13 and 2.15 can still balloon into unwieldy messes. Fortunately, based on current knowledge of the mass splittings, it is usually the case that only one mass splitting scale matters and other terms can be dropped. Figure 2.2 shows a schematic of the mass splittings. For historic reasons, Δm^2_{21} is known as the solar mass splitting and the larger mass splitting is called the atmospheric mass splitting. The atmospheric mass splitting is about 30 times the solar mass splitting. The sign of the solar mass splitting is known, while that of the atmospheric mass splitting is not. A positive value of Δm^2_{32} is called the normal hierarchy; a negative value is called the inverted hierarchy.

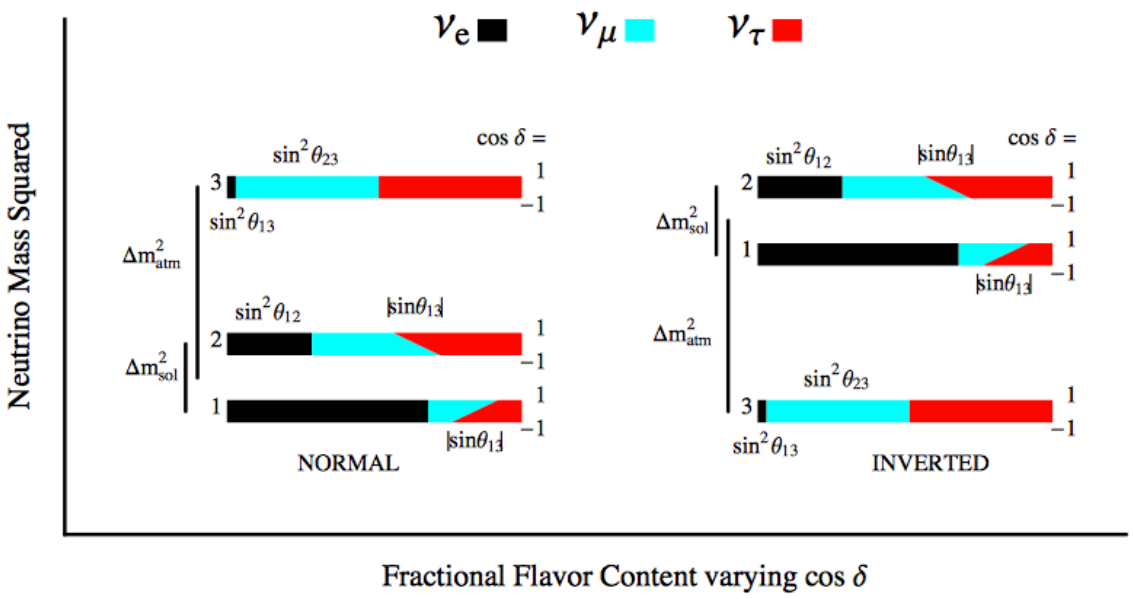


Figure 2.2: A schematic of the mass splittings between the three known neutrino mass states and how much they couple to each of the flavor states [36].

Two oscillation probabilities that are of interest to NOνA are the muon neutrino survival probability and electron neutrino appearance from a muon neutrino beam. Since $|\Delta m^2_{21}|$ is so much smaller than $|\Delta m^2_{32}|$, the solar oscillation baseline is much longer, thus the oscillation probability is first dominated by terms containing Δm^2_{32}. This is the case for NOνA. Furthermore, the probability can be simplified by making making the assumption that $|\Delta m^2_{32}| \approx |\Delta m^2_{31}|$. Under these conditions, the survival probability

of muon neutrinos is calculated as follows:

$$
\begin{aligned}
P(\nu_\mu \to \nu_\mu) &\approx 1 - 4|U_{\mu 3}|^2(|U_{\mu 1}|^2 + |U_{\mu 2}|^2)\sin^2\Delta_{32} && (2.21)\\
&\approx 1 - 4s_{23}^2(1-s_{13}^2)(c_{23}^2 + s_{23}^2 s_{13}^2)\sin^2\Delta_{32} && (2.22)\\
&\approx 1 - 4s_{23}^2 c_{23}^2 \sin^2\Delta_{32} + 4s_{23}^2 s_{13}^2(c_{23}^2 - s_{23}^2)\sin^2\Delta_{32} && (2.23)\\
&= 1 - \sin^2 2\theta_{23}\sin^2\Delta_{32} + 4\sin^2\theta_{23}\sin^2\theta_{13}\cos^2 2\theta_{23}\sin^2\Delta_{32} && (2.24)
\end{aligned}
$$

Between equations 2.22 and 2.23, the term proportional to s_{13}^4 was dropped using the current knowledge that s_{13}^2 is small [35]. Note that if θ_{13} were 0, then equation 2.24 would reduce to equation 2.17, the two neutrino survival probability.

The full 3 flavor electron neutrino appearance from muon neutrino oscillation probability is often written in the form [37]:

$$
P(\overset{(-)}{\nu}_\mu \to \overset{(-)}{\nu}_e) = P_{atm} + 2\sqrt{P_{atm}}\sqrt{P_{sol}}\left(\cos\delta\cos\Delta_{32} \overset{(+)}{(-)} \sin\delta\sin\Delta_{32}\right) + P_{sol} \qquad (2.25)
$$

where

$$
\begin{aligned}
\sqrt{P_{atm}} &\equiv \sin\theta_{23}\sin 2\theta_{13}\sin\Delta_{32} && (2.26)\\
\sqrt{P_{sol}} &\equiv \cos\theta_{23}\sin 2\theta_{12}\sin\Delta_{21} && (2.27)
\end{aligned}
$$

where the approximation $|\Delta m_{32}^2| \approx |\Delta m_{31}^2|$ has been made and higher order terms of s_{13}^2 been dropped. For an experiment at a short enough baseline such as NOνA, the P_{sol} term is negligible as it depends on a higher order term of the solar mass splitting. The cross term is also not the dominant effect as it also depends upon the solar mass splitting, but it demonstrates interesting behavior. The $\cos\delta$ term is CP conserving, but the $\sin\delta$ term exhibits CP violation. This is why δ is called the CP violating phase angle.

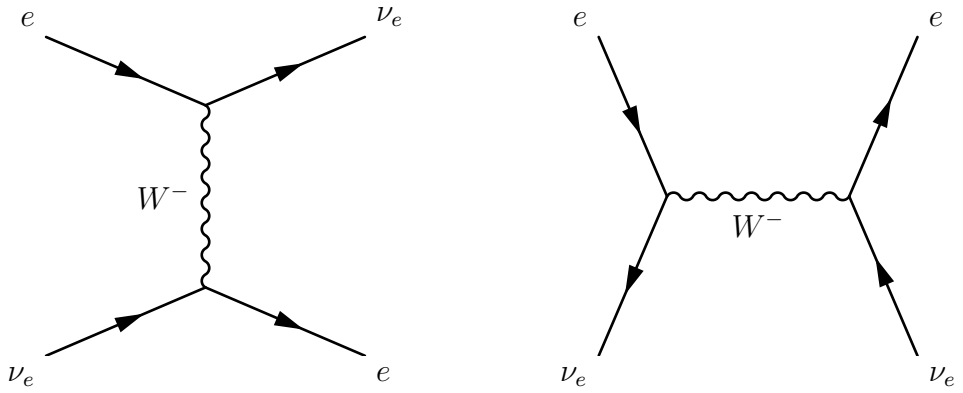


Figure 2.3: CC coherent forward scattering interactions involved in the MSW effect. Left: Scattering of electron neutrinos on electrons. Right: Scattering of anti-electron neutrinos on electrons.

2.4 Matter Effects

So far, the oscillation formalism has been developed only considering neutrinos in a vacuum. However, most neutrino oscillation experiments involve neutrinos traveling through matter, be it the Sun or the Earth. This affects the oscillation probabilities in a process called the Mikheyev-Smirnov-Wolfenstein effect, or MSW effect. The phenomenon was first proposed by Wolfenstein in 1978 [38]; Mikheyev and Smirnov built upon that work in 1985 [39] as a possible solution for the solar neutrino problem.

The MSW effect is the CC coherent forward scattering of neutrinos off of the electrons in ordinary matter, a channel only available to electron flavor neutrinos and anti-neutrinos. Figure 2.3 illustrates the interactions. The electrons contribute an additional potential term, $V_e = \pm\sqrt{2}G_F N_e$, where G_F is Fermi's constant, N_e is the electron number density, the positive sign is for neutrinos, and the negative for anti-neutrinos. Neutrinos also forward scatter off the neutrons and protons in matter via neutral current interactions, but this only provides an overall phase as all neutrino flavors participate in these interactions equally. The matter induced potential adds an additional term to the Schrödinger equation, affecting the time evolution of the flavor states and thus changing the oscillation probabilities.

The following derivation will consider the MSW effect in the case of two neutrino flavors. The time

evolution of the flavor states is written as follows:

$$i\begin{pmatrix}\nu_e\\ \nu_\mu\end{pmatrix} = \left[U\begin{pmatrix}\frac{m_1^2}{2E} & 0\\ 0 & \frac{m_2^2}{2E}\end{pmatrix}U^\dagger + \begin{pmatrix}\pm V_e & 0\\ 0 & 0\end{pmatrix}\right]\begin{pmatrix}\nu_e\\ \nu_\mu\end{pmatrix} \tag{2.28}$$

Inserting the 2 flavor PMNS matrix from equation 2.16, applying some trigonometric identites, and dropping common diagonal terms, equation 2.28 simplifies to

$$i\begin{pmatrix}\nu_e\\ \nu_\mu\end{pmatrix} = \frac{1}{4E}\begin{pmatrix}-\Delta m_{21}^2\cos 2\theta \pm 4EV_e & \Delta m_{21}^2\sin 2\theta\\ \Delta m_{21}^2\sin 2\theta & \Delta m_{21}^2\cos 2\theta\end{pmatrix}\begin{pmatrix}\nu_e\\ \nu_\mu\end{pmatrix}. \tag{2.29}$$

The diagonal terms are dropped because they can be absorbed by the phase convention of the neutrino states. This Hamiltonian can be re-diagonalized with another unitary transformation, $H_M = U_M^\dagger H U_M$, with the following results:

$$H_M = \frac{1}{2}\begin{pmatrix}-\frac{\Delta m_M^2}{2E} & 0\\ 0 & \frac{\Delta m_M^2}{2E}\end{pmatrix} \tag{2.30}$$

$$U_M = \begin{pmatrix}\cos\theta_M & \sin\theta_M\\ -\sin\theta_M & \cos\theta_M\end{pmatrix}, \tag{2.31}$$

where

$$\sin 2\theta_M \equiv \frac{\sin 2\theta}{A_M} \tag{2.32}$$

$$\Delta m_M^2 \equiv \Delta m_{21}^2 A_M \tag{2.33}$$

$$A_M \equiv \sqrt{\left(\cos 2\theta \mp \frac{2EV_e}{\Delta m_{21}^2}\right)^2 + \sin^2 2\theta}, \tag{2.34}$$

and now the negative sign in A_M is for neutrinos and the positive sign for anti-neutrinos. As the electron number density goes to 0, so too does V_e and the vacuum solution is recovered.

From the form of this solution, it can be seen that the Hamiltonian takes the same form as that in

vacuum oscillations, but with modified effective masses. Likewise, U_M has the same form as the 2 neutrino PMNS matrix, so θ_M can be considered the effective mixing angle. In the absence of neutrino oscillations (when $\theta = 0$), matter effects cannot "create" them. However, even for small angles θ, the matter effect can create a resonant effect pushing the effective mixing angle, θ_M, maximally to $45°$. This occurs when the term in parenthesis in the definition of A_M is 0 (equation 2.34).

$$N_e^{res} = \frac{\Delta m_{21}^2 \cos 2\theta}{2\sqrt{2} G_F E} \tag{2.35}$$

In the case of 3 neutrinos, the same procedure is followed to diagonalize the Hamilton and obtain effective values for the various oscillation parameters. The effects are considerably more complicated, but the general effect is the same–matter changes the effective neutrino mass and alters the oscillation probability curves differently for neutrinos and anti-neutrinos. Under the same conditions that were used to calculate $P(\overset{(-)}{\nu}_\mu \to \overset{(-)}{\nu}_e)$ in section 2.3, the results can be simplified to a few basic replacements [40].

$$P(\overset{(-)}{\nu}_\mu \to \overset{(-)}{\nu}_e) = P_{atm}^M + 2\sqrt{P_{atm}^M}\sqrt{P_{sol}^M}\left(\cos\delta\cos\Delta_{32} \overset{(+)}{(-)} \sin\delta\sin\Delta_{32}\right) + P_{sol}^M \tag{2.36}$$

This is exactly the same form as equation 2.25. The interesting effects are seen with how P_{atm}^M and P_{sol}^M differ from their respective vacuum counterparts.

$$\sqrt{P_{atm}^M} \equiv \sin\theta_{23}\sin 2\theta_{13}\frac{\sin(\Delta_{31} - aL)}{\Delta_{31} - aL}\Delta_{31} \tag{2.37}$$

$$\sqrt{P_{sol}} \equiv \cos\theta_{23}\sin 2\theta_{12}\frac{\sin(aL)}{aL}\Delta_{21} \tag{2.38}$$

Here, $a \equiv \pm G_F N_e/\sqrt{2}$ where the positive sign is for neutrinos and the negative sign for anti-neutrinos. For the Earth, $|a| \approx 1/3500$ km.

The combined effect that appears in equations 2.36, 2.37, and 2.38 due to the presence of matter plays an interesting role in the search for CP violation. The MSW effect by itself mimics CP violation as it alters oscillation probabilities for neutrinos and anti-neutrinos differently. Depending on the value of δ that nature has chosen, the differences in oscillation probabilities due to the CP violation angle and the MSW

effect can either compound or cancel out.

2.5 Current Measurements

Most of the free parameters in the PMNS matrix have been measured by various solar, atmospheric, accelerator, and reactor neutrino experiments. However, any given neutrino experiment does not have sensitivity to all of the oscillation parameters. Instead, experiments are sensitive to specific angles based on their baseline and the energies of the neutrinos they observe. Solar neutrino experiments, such as GALLEX, SAGE, Super-K, and SNO, measure neutrinos with energies on the order of several MeV after a very long baseline, and are most sensitive to θ_{12} and Δm^2_{21}. Due to the strong MSW effect within the Sun, solar neutrino experiments also determined the ordering of mass states ν_1 and ν_2; ν_2 is defined as the heavier state. Atmospheric neutrino experiments, such as Super-K, SNO, and MINOS, measure neutrinos generated by cosmic ray collisions with the Earth's atmosphere, and are sensitive to θ_{23} and Δm^2_{32}.

Reactor neutrino experiments, such as Chooz, Double Chooz, RENO, and Daya Bay, measure $\bar{\nu}_e$ generated by nearby nuclear reactors. Like solar neutrinos, reactor neutrinos have energies on the order of a few MeV. By measuring these neutrinos with a short baseline ($O(1km)$), the 2 neutrino approximation is valid, so the oscillation probability can be approximated as:

$$P(\bar{\nu}_e \to \bar{\nu}_e) \approx 1 - \sin^2 2\theta_{13} \sin^2 \Delta_{31}. \tag{2.39}$$

Daya Bay made the first nonzero measurement of θ_{13} in 2012, reporting a value of $\sin^2 2\theta_{13} = 0.092 \pm 0.016\,(\text{stat}) \pm 0.005\,(\text{syst})$ after taking just 55 days of data [41]. This result excluded a zero value for θ_{13} at 5.2σ and is shown in figure 2.4. Since that result, the limits have only continued to improve, and the leading measurement still comes from Daya Bay.

Accelerator neutrino experiments, such as MINOS, T2K, and NOνA, begin with a beam of nearly pure $(\bar{\nu}_\mu)$ and search for both a disappearance of $(\bar{\nu}_\mu)$ and appearance of other neutrino flavors. These experiments are sensitive to θ_{13}, θ_{23}, Δm^2_{32}, and δ. The experiments that have the largest matter effect are the most sensitive to δ and the mass hierarchy.

Global fits to the combined data of these (and other) neutrino experiments have been performed and

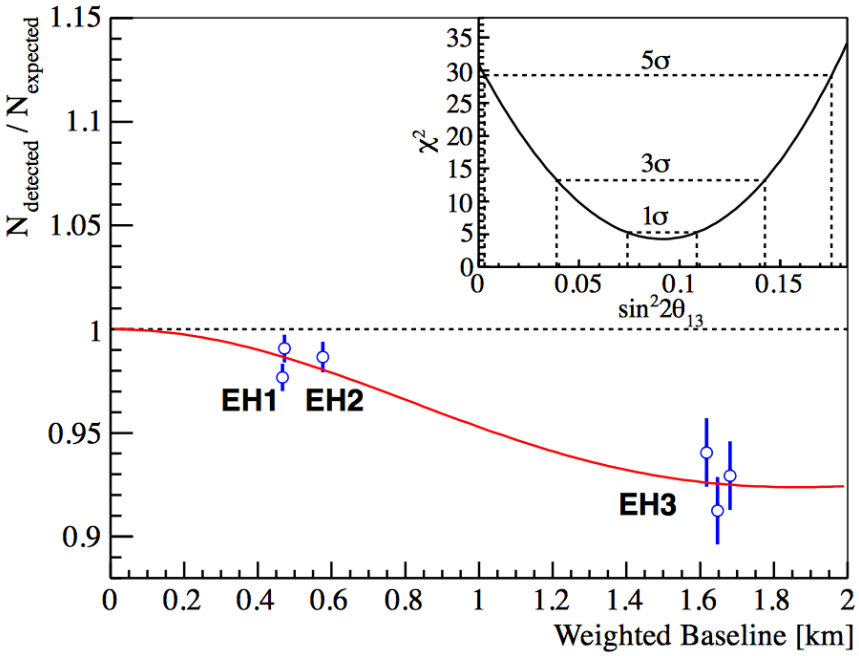


Figure 2.4: First measurement of θ_{13} from Daya Bay [41]. The points show the ratio of observed to expected events assuming $\theta_{13} = 0$. Each point is a measurement at a different detector. The inset in the upper corner shows the χ^2 value vs value of $\sin^2 2\theta_{13}$, and excludes $\theta_{13} = 0$ at greater than 5σ.

summarized in [35, 42]; the best fit values are shown in table 2.1. While most of the parameters have been measured with good precision, there are still a few lingering questions. From the table it is clear that a much better measurement on the CP violation angle is needed. The mass hierarchy still needs to be definitively measured as well. The other main question is whether θ_{23} is maximal, and if not, whether it is in the lower or upper octant.

Current and next generation reactor experiments have a great outlook to answer these outstanding questions. Making use of the MSW effect, the ν_e and $\bar{\nu}_e$ appearance results from experiments like T2K and NOνA could simultaneously measure δ, the mass hierarchy, and θ_{23} octant. The prospects for NOνA to make these measurements are shown in figure 2.5. The first analysis results from NOνA [43, 44] were published after taking about 10% of the experiments design statistics and already show promise. The NOνA measurement for δ, shown in figure 2.6, provide a hint toward the normal hierarchy and eliminate portions of δ space at 90% confidence.

Parameter		Best-Fit ($\pm 1\sigma$)	3σ Range
Δm^2_{21} [10^{-5} eV^2]		$7.54^{+0.26}_{-0.22}$	$6.99 - 8.18$
$\lvert\Delta m^2\rvert$ [10^{-3} eV^2]	NH	2.43 ± 0.06	$2.23 - 2.61$
	IH	2.38 ± 0.06	$2.19 - 2.56$
$\sin^2\theta_{12}$		0.308 ± 0.017	$0.259 - 0.359$
$\sin^2\theta_{23}$	NH	$0.437^{+0.033}_{-0.023}$	$0.374 - 0.628$
	IH	$0.455^{+0.039}_{0.031}$	$0.380 - 0.641$
$\sin^2\theta_{13}$	NH	$0.0234^{+0.0020}_{-0.0019}$	$0.0176 - 0.0295$
	IH	$0.0240^{+0.0019}_{-0.0022}$	$0.0178 - 0.0298$
δ/π (2σ range)	NH	$1.39^{+0.38}_{-0.27}$	$(0.00 - 0.16) \oplus (0.86 - 2.00)$
	IH	$1.31^{+0.29}_{-0.33}$	$(0.00 - 0.02) \oplus (0.70 - 2.00)$

Table 2.1: Current status of best fit oscillation parameters, from [35, 42]. The last column shows the allowed values within a 3σ range, with the exception of δ, which is shown at a 2σ range. This is because the current global best fit for δ still allows the full range from 0 to 2π at 3σ. NOνA should vastly improve the limits on δ.

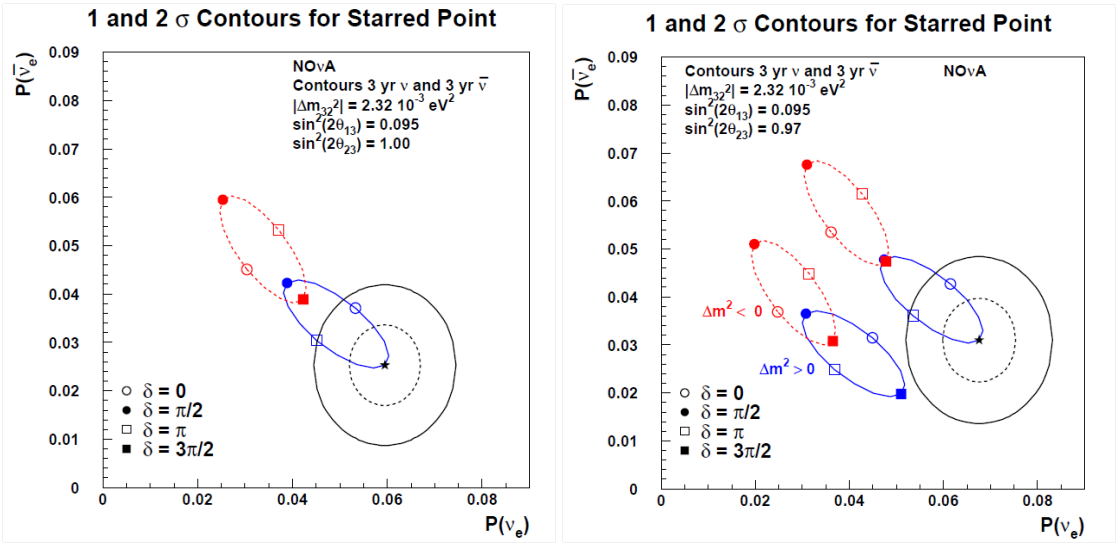


Figure 2.5: Probability of ν_e appearance versus $\bar{\nu}_e$ appearance at NOνA. The blue ellipses are for the normal hierarchy; the red ellipses are for the inverted hierarchy. The starred points show a possible measurement NOνA could make. The matter effect can either constructively or destructively combine with the CP violation effect. A larger matter effect, further separates the two mass hierarchy ellipses. This corresponds to neutrinos passing through more matter. On the left, θ_{23} is assumed to be 45° for maximal mixing, purely showcasing the interference between the matter and CP violation effects. On the right, θ_{23} is non-maximal, showing how the dependence on θ_{23} affects both ellipses in the same way.

2.6 Sterile Neutrinos

The analysis presented in this dissertation considers sterile neutrinos in a $3+1$ model. This adds a fourth neutrino flavor state, ν_s, and mass state, ν_4, and the PMNS matrix becomes a 4×4 matrix. Given that

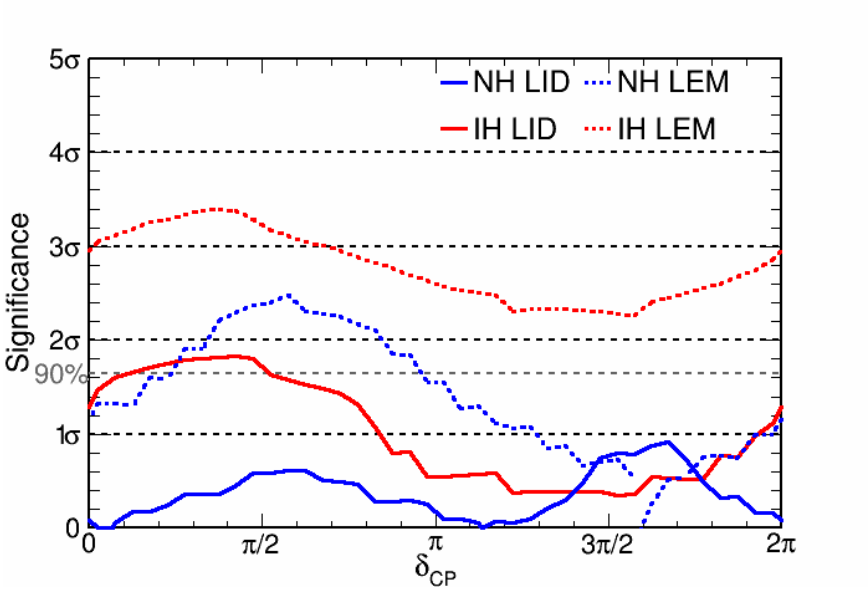


Figure 2.6: First measurement of δ from NOνA [43]. The plot shows the significance of the difference between the observed and predicted number of events as a function of delta. NOνA used a primary and secondary selection technique, the primary (secondary) technique is shown as the solid (dotted) line. The secondary selection disfavors the inverted hierarchy for all values of δ.

standard three flavor oscillations explain solar, atmospheric, and long-baseline oscillations so well, this fourth state must be mostly sterile, i.e., $|U_{\alpha 4}| \ll 1$ for $\alpha \in \{e, \mu, \tau\}$.

The experimental evidence discussed in section 1.4 suggests a mass splitting of $\Delta m^2 \sim O(1\,\text{eV}^2)$. If the fourth neutrino mass state were the lightest, then all of the other mass states would have absolute masses $m \gtrsim 1\,\text{eV}$. However, the sum of the active neutrino masses is well constrained by cosmology; the most recent Planck results reported $\sum m_\nu < 0.23\,\text{eV}$ [45]. Consequently, it is assumed that the additional neutrino mass state is heavier than the other three. Furthermore, since $\Delta m^2_{4i} \gg |\Delta m^2_{ji}|$ for $i, j < 4$, it is assumed that $\Delta_{41} \approx \Delta_{42} \approx \Delta_{43}$.

As mentioned above, adding a fourth neutrino state expands the PMNS matrix into a 4×4 matrix. It can be written as the product of individual (complex) rotation matrices, as in the second line of equation 2.20. This adds an additional 3 real rotation angles and 2 complex phases. It is a matter of convention on how exactly to add these mixing angles into the PMNS matrix, but as physical results cannot be affected by a particular parametrization, it is best to choose the most convenient one. Following the parametrization

in [46], U takes the form:

$$U = R_{34}C_{24}R_{23}R_{14}C_{13}C_{12}, \tag{2.40}$$

where R_{ij} is a real matrix parametrized by a single, real mixing angle, θ_{ij}, and C_{ij} is a complex matrix parametrized by a real mixing angle θ_{ij} and complex phase δ_{ij}. This parametrization is convenient because as was the case in section 2.3, the solar mass scale Δ_{21} is negligible, thus mixing angles parametrized by C_{12} drop out of the oscillation probabilities. Based on this, equation 2.40 can be expanded in terms of the mixing angles:

$$U = \begin{bmatrix} U_{e1} & U_{e2} & c_{14}s_{13}e^{-i\delta_{13}} & s_{14} \\ U_{\mu 1} & U_{\mu 2} & -s_{14}s_{13}e^{-i\delta_{13}}s_{24}e^{-i\delta_{24}} + c_{13}s_{23}c_{24} & c_{14}s_{24}e^{-i\delta_{24}} \\ U_{\tau 1} & U_{\tau 2} & -s_{14}c_{24}s_{34}s_{13}e^{-i\delta_{13}} - c_{13}s_{23}s_{34}s_{24}e^{i\delta_{24}} + c_{13}c_{23}c_{34} & c_{14}c_{24}s_{34} \\ U_{s1} & U_{s2} & -s_{14}c_{24}c_{34}s_{13}e^{-i\delta_{13}} - c_{13}s_{23}c_{34}s_{24}e^{i\delta_{24}} - c_{13}c_{23}s_{34} & c_{14}c_{24}c_{34} \end{bmatrix} \tag{2.41}$$

Above, only the columns that appear in the oscillation probabilities have been expanded. Before using these expanded forms, however, it is useful to write the oscillation probability formulae in terms of the matrix elements first.

For NOνA, which starts with a nearly pure ν_μ beam, the two most important oscillation probabilities are the ν_μ survival probability, $P(\nu_\mu \to \nu_\mu)$ and the probability that a ν_μ will not oscillate to a sterile neutrino, $1 - P(\nu_\mu \to \nu_s)$. Starting with equations 2.15 and 2.13, these probabilities are calculated as follows, using the unitarity of U to remove elements found in its the first two columns.

$$\begin{aligned} P(\nu_\mu \to \nu_\mu) \approx 1 &- 4|U_{\mu 4}|^2(|U_{\mu 1}|^2 + |U_{\mu 2}|^2 + |U_{\mu 3}|^2)\sin^2\Delta_{41} \\ &- 4|U_{\mu 3}|^2(|U_{\mu 1}|^2 + |U_{\mu 2}|^2)\sin^2\Delta_{31} \\ \approx 1 &- 4|U_{\mu 4}|^2(1 - |U_{\mu 4}|^2)\sin^2\Delta_{41} \\ &- 4|U_{\mu 3}|^2(1 - |U_{\mu 3}|^2 - |U_{\mu 4}|^2)\sin^2\Delta_{31} \end{aligned} \tag{2.42}$$

$$
\begin{aligned}
1 - P(\nu_\mu \to \nu_s) = 1 \quad &+ \quad 4\sum_{i>j} \Re(U_{\mu i} U^*_{si} U_{sj} U^*_{\mu j}) \sin^2 \Delta_{ij} \\
&- \quad 2\sum_{i>j} \Im(U_{\mu i} U^*_{si} U_{sj} U^*_{\mu j}) \sin 2\Delta_{ij} \\
\approx 1 \quad &+ \quad 4\Re(C_{41,42,43}) \sin^2 \Delta_{41} + 4\Re(C_{31,32}) \sin^2 \Delta_{31} \\
&- \quad 2\Im(C_{41,42,43}) \sin 2\Delta_{41} - 2\Im(C_{31,32}) \sin 2\Delta_{31}
\end{aligned} \tag{2.43}
$$

In equation 2.43, $C_{41,42,43}$ and $C_{31,32}$ are sums of products of PMNS matrix elements.

$$
\begin{aligned}
C_{41,42,43} &= U_{\mu 4} U^*_{s4} (U_{s1} U^*_{\mu 1} + U_{s2} U^*_{\mu 2} + U_{s3} U^*_{\mu 3}) \\
&= U_{\mu 4} U^*_{s4} (-U_{s4} U^*_{\mu 4}) \\
&= -|U_{\mu 4}|^2 |U_{s4}|^2
\end{aligned} \tag{2.44}
$$

Note that this term is real, thus the term with $\Im(C_{41,42,43})$ in equation 2.43 drops out.

$$
\begin{aligned}
C_{31,32} &= U_{\mu 3} U^*_{s3} (U_{s1} U^*_{\mu 1} + U_{s2} U^*_{\mu 2}) \\
&= -U_{\mu 3} U^*_{s3} (U_{s4} U^*_{\mu 4} + U_{s3} U^*_{\mu 3}) \\
&= -|U_{\mu 3}|^2 |U_{s3}|^2 - U_{\mu 3} U^*_{\mu 4} U^*_{s3} U_{s4} \\
&= -|U_{\mu 3}|^2 |U_{s3}|^2 - \Re(U_{\mu 3} U^*_{\mu 4} U^*_{s3} U_{s4}) - i\Im(U_{\mu 3} U^*_{\mu 4} U^*_{s3} U_{s4}) \\
&= -|U_{\mu 3}|^2 |U_{s3}|^2 - \Re(U^*_{\mu 3} U_{\mu 4} U_{s3} U^*_{s4}) + i\Im(U^*_{\mu 3} U_{\mu 4} U_{s3} U^*_{s4})
\end{aligned} \tag{2.45}
$$

The last equality in 2.45 was done so that all terms in the oscillation probability are subtracted from 1. Using these results, equation 2.43 becomes

$$
\begin{aligned}
1 - P(\nu_\mu \to \nu_s) \approx 1 \quad &- \quad 4|U_{\mu 4}|^2 |U_{s4}|^2 \sin^2 \Delta_{41} \\
&- \quad 4(|U_{\mu 3}|^2 |U_{s3}|^2 + \Re(U^*_{\mu 3} U_{\mu 4} U_{s3} U^*_{s4})) \sin^2 \Delta_{31} \\
&- \quad 2\Im(U^*_{\mu 3} U_{\mu 4} U_{s3} U^*_{s4}) \sin 2\Delta_{31}.
\end{aligned} \tag{2.46}
$$

Equations 2.42 and 2.46 are the probabilities that a ν_μ will survive as a ν_μ and survive as one of the active neutrino flavors, respectively. These can now be written in terms of the mixing angles. However, several approximations make the expressions much more bearable. First, based on the best value reported in section 2.5, the value of s_{13} is taken as 0. This has the fortunate side effect of removing any dependence on δ_{13} as well. Second, the value of s_{14} is also taken as 0, based on reactor and other constraints, e.g. [29, 47]. Note that in the relevant matrix elements, s_{13} and s_{14} always appear together as a product, further justifying the approximation. The matrix element products that appear in the oscillation probabilities are calculated below.

$$
\begin{aligned}
|U_{\mu 4}|^2 &= s_{24}^2 && (2.47)\\
|U_{s4}|^2 &= c_{24}^2 c_{34}^2 && (2.48)\\
|U_{\mu 3}|^2 &= s_{23}^2 c_{24}^2 && (2.49)\\
|U_{s3}|^2 &= c_{23}^2 s_{34}^2 + s_{23}^2 s_{24}^2 c_{34}^2 + 2 c_{23} s_{23} s_{24} c_{34} s_{34} \cos\delta_{24} && (2.50)\\
U_{\mu 3}^* U_{\mu 4} U_{s3} U_{s4}^* &= -s_{23} c_{24}^2 s_{24} c_{34} (s_{23} s_{24} c_{34} + c_{23} s_{34} e^{-i\delta_{24}}) && (2.51)\\
\Re(U_{\mu 3}^* U_{\mu 4} U_{s3} U_{s4}^*) &= -s_{23} c_{24}^2 s_{24} c_{34} (s_{23} s_{24} c_{34} + c_{23} s_{34} \cos\delta_{24}) && (2.52)\\
\Im(U_{\mu 3}^* U_{\mu 4} U_{s3} U_{s4}^*) &= c_{23} s_{23} c_{24}^2 s_{24} c_{34} s_{34} \sin\delta_{24} && (2.53)
\end{aligned}
$$

These results can be substituted into equations 2.42 and 2.46.

$$
\begin{aligned}
P(\nu_\mu \to \nu_\mu) \approx 1 &- 4|U_{\mu 4}|^2 (1 - |U_{\mu 4}|^2) \sin^2 \Delta_{41}\\
&- 4|U_{\mu 3}|^2 (1 - |U_{\mu 3}|^2 - |U_{\mu 4}|^2) \sin^2 \Delta_{31}\\
\approx 1 &- 4 s_{24}^2 (1 - s_{24}^2) \sin^2 \Delta_{41}\\
&- 4 s_{23}^2 c_{24}^2 (1 - s_{23}^2 c_{24}^2 - s_{24}^2) \sin^2 \Delta_{31}\\
\approx 1 &- 4 c_{24}^2 s_{24}^2 \sin^2 \Delta_{41} - 4 c_{24}^4 s_{23}^2 (1 - s_{23}^2) \sin^2 \Delta_{31}\\
\approx 1 &- \sin^2 2\theta_{24} \sin^2 \Delta_{41} - c_{24}^4 \sin^2 2\theta_{23} \sin^2 \Delta_{31} && (2.54)
\end{aligned}
$$

$$
\begin{aligned}
1 - P(\nu_\mu \to \nu_s) \approx 1 \quad &- \quad 4|U_{\mu 4}|^2 |U_{s4}|^2 \sin^2 \Delta_{41} \\
&- \quad 4(|U_{\mu 3}|^2 |U_{s3}|^2 + \Re(U^*_{\mu 3} U_{\mu 4} U_{s3} U^*_{s4})) \sin^2 \Delta_{31} \\
&- \quad 2\Im(U^*_{\mu 3} U_{\mu 4} U_{s3} U^*_{s4}) \sin 2\Delta_{31} \\
\approx 1 \quad &- \quad 4 c^2_{24} s^2_{24} c^2_{34} \sin^2 \Delta_{41} \\
&- \quad 4 \left[s^2_{23} c^2_{24} (c^2_{23} s^2_{34} + s^2_{23} s^2_{24} c^2_{34} + 2 c_{23} s_{23} s_{24} c_{34} s_{34} \cos \delta_{24}) \right. \\
&\qquad \left. - s_{23} c^2_{24} s_{24} c_{34} (s_{23} s_{24} c_{34} + c_{23} s_{34} \cos \delta_{24}) \right] \sin^2 \Delta_{31} \\
&- \quad 2 c_{23} s_{23} c^2_{24} s_{24} c_{34} s_{34} \sin \delta_{24} \sin 2\Delta_{31} \\
\approx 1 \quad &- \quad c^2_{34} \sin^2 2\theta_{24} \sin^2 \Delta_{41} \\
&- \quad 4 \left[c^2_{23} s^2_{23} c^2_{24} s^2_{34} - s^2_{23} c^2_{24} s^2_{24} c^2_{34} (1 - s^2_{23}) \right. \\
&\qquad \left. + c_{23} s_{23} c^2_{24} s_{24} c_{34} s_{34} (2 s^2_{23} - 1) \cos \delta_{24} \right] \sin^2 \Delta_{31} \\
&- \quad \frac{1}{4} c_{24} \sin 2\theta_{23} \sin 2\theta_{24} \sin 2\theta_{34} \sin \delta_{24} \sin 2\Delta_{31} \\
\approx 1 \quad &- \quad c^2_{34} \sin^2 2\theta_{24} \sin^2 \Delta_{41} - c^2_{24} (s^2_{34} - s^2_{24} c^2_{34}) \sin^2 2\theta_{23} \sin^2 \Delta_{31} \\
&- \quad \frac{1}{4} c_{24} \sin 2\theta_{23} \sin 2\theta_{24} \sin 2\theta_{34} \sin \delta_{24} \sin 2\Delta_{31}
\end{aligned} \tag{2.55}
$$

In the last equality, the term proportional to $\cos \delta_{24}$ was dropped noting that $2s^2_{23} - 1 \approx 0$. Even if it is not exact, there is already an additional factor of $s_{24} s_{34}$. Both s_{24} and s_{34} must be small in order to satisfy $|U_{\mu 4}|, |U_{\tau 4}| \ll 1$, thus the term proportional to $s_{24} s_{34} (2 s^2_{23} - 1)$ can be considered higher order than the remaining terms.

The effect of each of the oscillation parameters can be seen in the figures below. Each page shows the ν_μ survival probability followed by the active neutrino survival probability and uses the exact oscillation probabilities. Figures 2.7-2.11 show how the mixing parameters that appear in equations 2.54 and 2.55 affect the oscillation probabilities. Figures 2.12-2.14 are similar but show the much smaller effects of the parameters that dropped out of these equations due to approximations. (Solar mixing parameters are not shown.) The individual plots show different probabilities curves while varying a single oscillation parameter over

a set of representative values; the mixing parameters that are not varied in a particular figure are held fixed to the values in table 2.2. In the figures, the bottom axis shows the neutrino L/E, thus the plots show oscillation probabilities at both the near and far detectors. The top axis shows the neutrino energy at each NOνAdetector, the dashed line separates the Near Detector (ND) on the left from the Far Detector (FD) on the right. Finally, the grey bands show the neutrino energies each detector is sensitive to; a darker color means a higher intensity of neutrinos at that energy.

Oscillation Parameter	Value
ρ	$2.84\,\mathrm{g/cm^3}$
Δm^2_{21}	$7.53 \times 10^{-5}\,\mathrm{eV^2}$
Δm^2_{32}	$2.44 \times 10^{-3}\,\mathrm{eV^2}$
Δm^2_{41}	$0.5\,\mathrm{eV^2}$
$\sin^2 2\theta_{12}$	0.846
$\sin^2 2\theta_{13}$	0.085
θ_{23}	$\pi/4$
θ_{14}	0
θ_{24}	10°
θ_{34}	35°
δ_{13}	0
δ_{24}	0

Table 2.2: The 4 flavor oscillation parameters assumed for studying the effects of different parameters on the oscillation probabilities.

For the particular analysis presented in the rest of this dissertation, it is assumed that there are no Near Detector oscillations. Thus, the analysis will make measurements of θ_{24}, θ_{34}, and these will be translated into measurements of $|U_{\mu 4}|^2$ and $|U_{\tau 4}|^2$, and the results will be valid in the range $0.05 < \Delta m^2_{41} < 0.5\,\mathrm{eV}^2$.

To make the above measurements, the analysis presented here considers the rate of Neutral Current (NC) events. The rate of Charged Current (CC) events is obviously affected by neutrino oscillations. Given that standard three flavor oscillations already impact the number of CC events, it would be a difficult task to decouple oscillations between the active neutrino flavors and a fourth, sterile flavor. On the other hand, NC interactions are completely insensitive to the active neutrino flavors, and hence to oscillations between them. In the standard three flavor model, then, the total rate of NC events is only determined by the over-

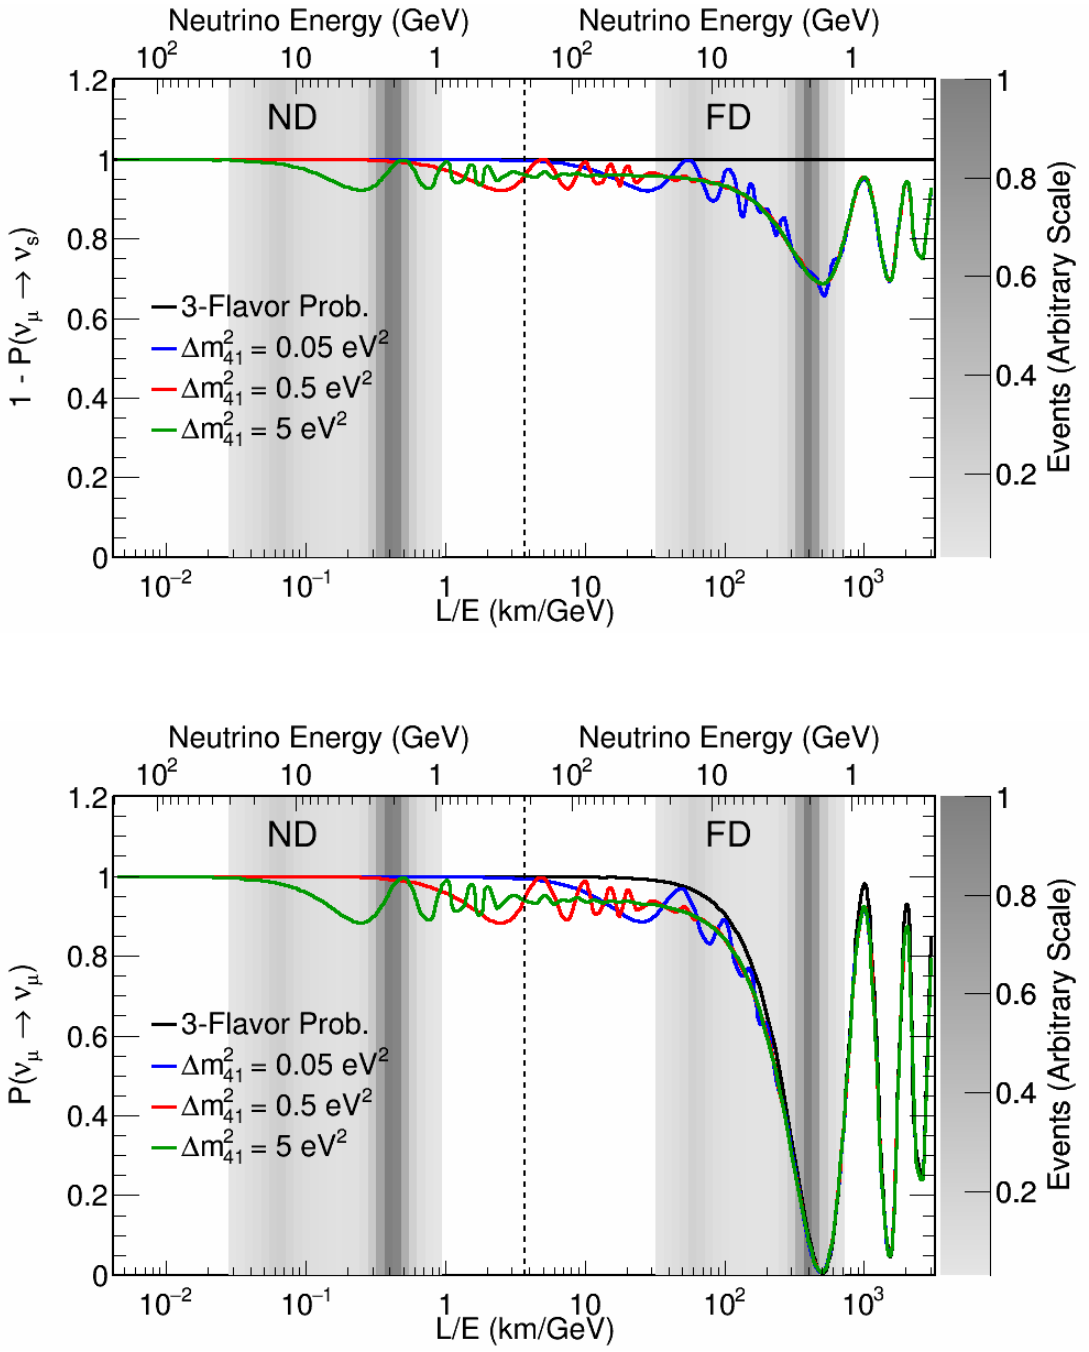


Figure 2.7: Total active neutrino survival probability (top) and ν_μ survival probability (bottom) for different values of Δm^2_{41}. Above $\Delta m^2_{41} = 0.5\,\text{eV}^2$, oscillations start to affect the ND. The grey bands show the neutrino energy spectra observed by NOνA, with darker colors representing a higher intensity of neutrinos. Parameters not varied in the figures are held fixed to values shown in table 2.2.

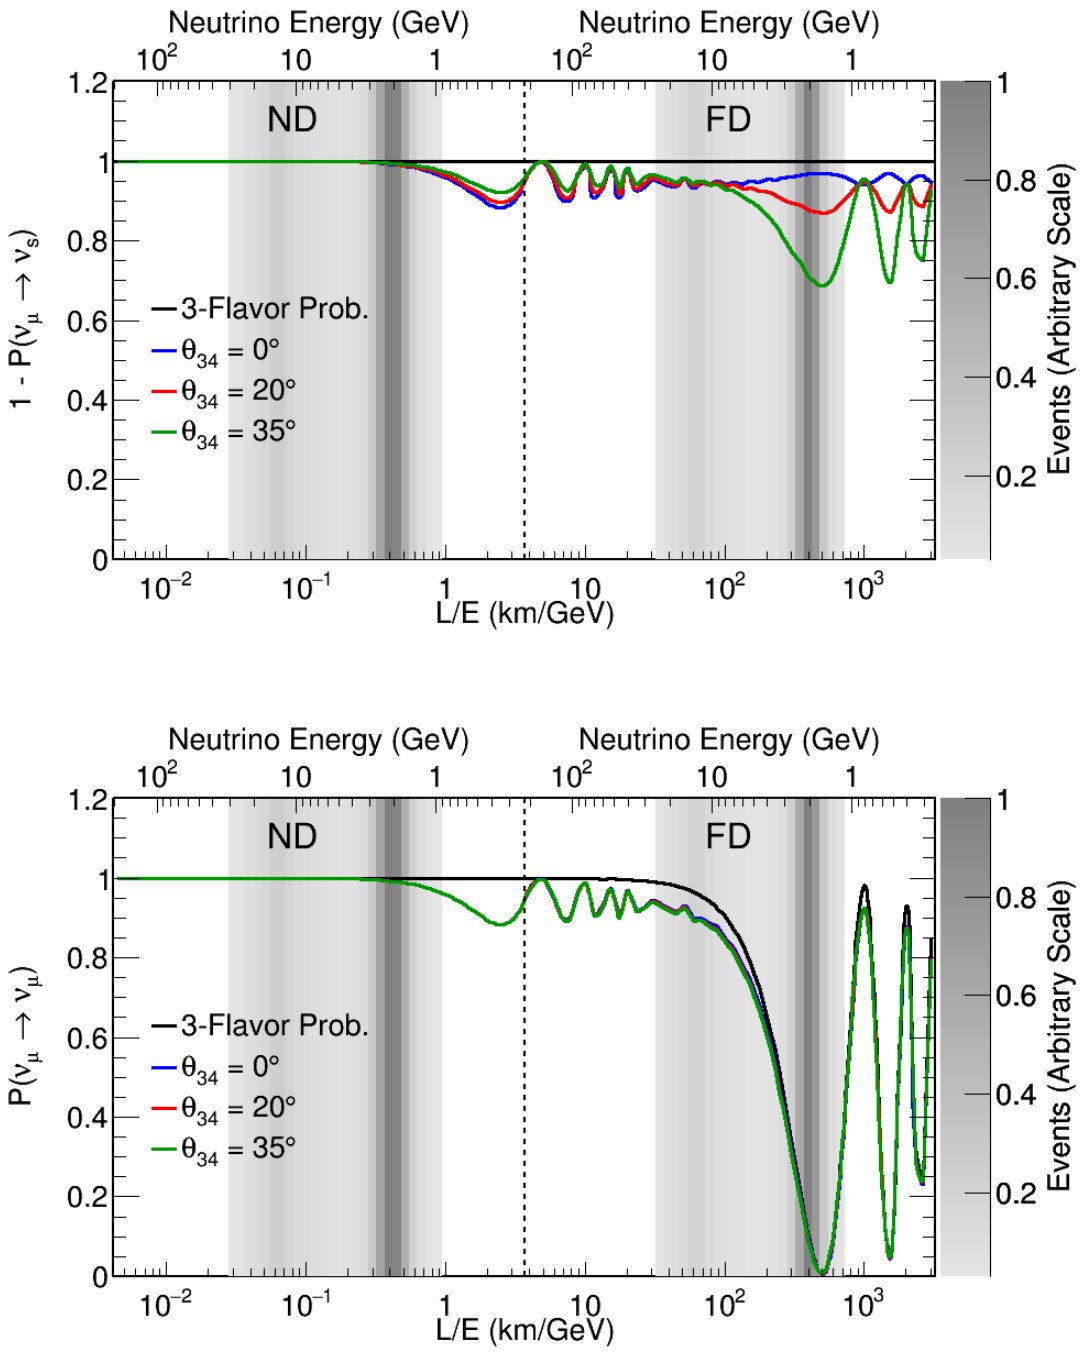


Figure 2.8: Total active neutrino survival probability (top) and ν_μ survival probability (bottom) for different values of θ_{34}. The grey bands show the neutrino energy spectra observed by NOνA, with darker colors representing a higher intensity of neutrinos. Parameters not varied in the figures are held fixed to values shown in table 2.2.

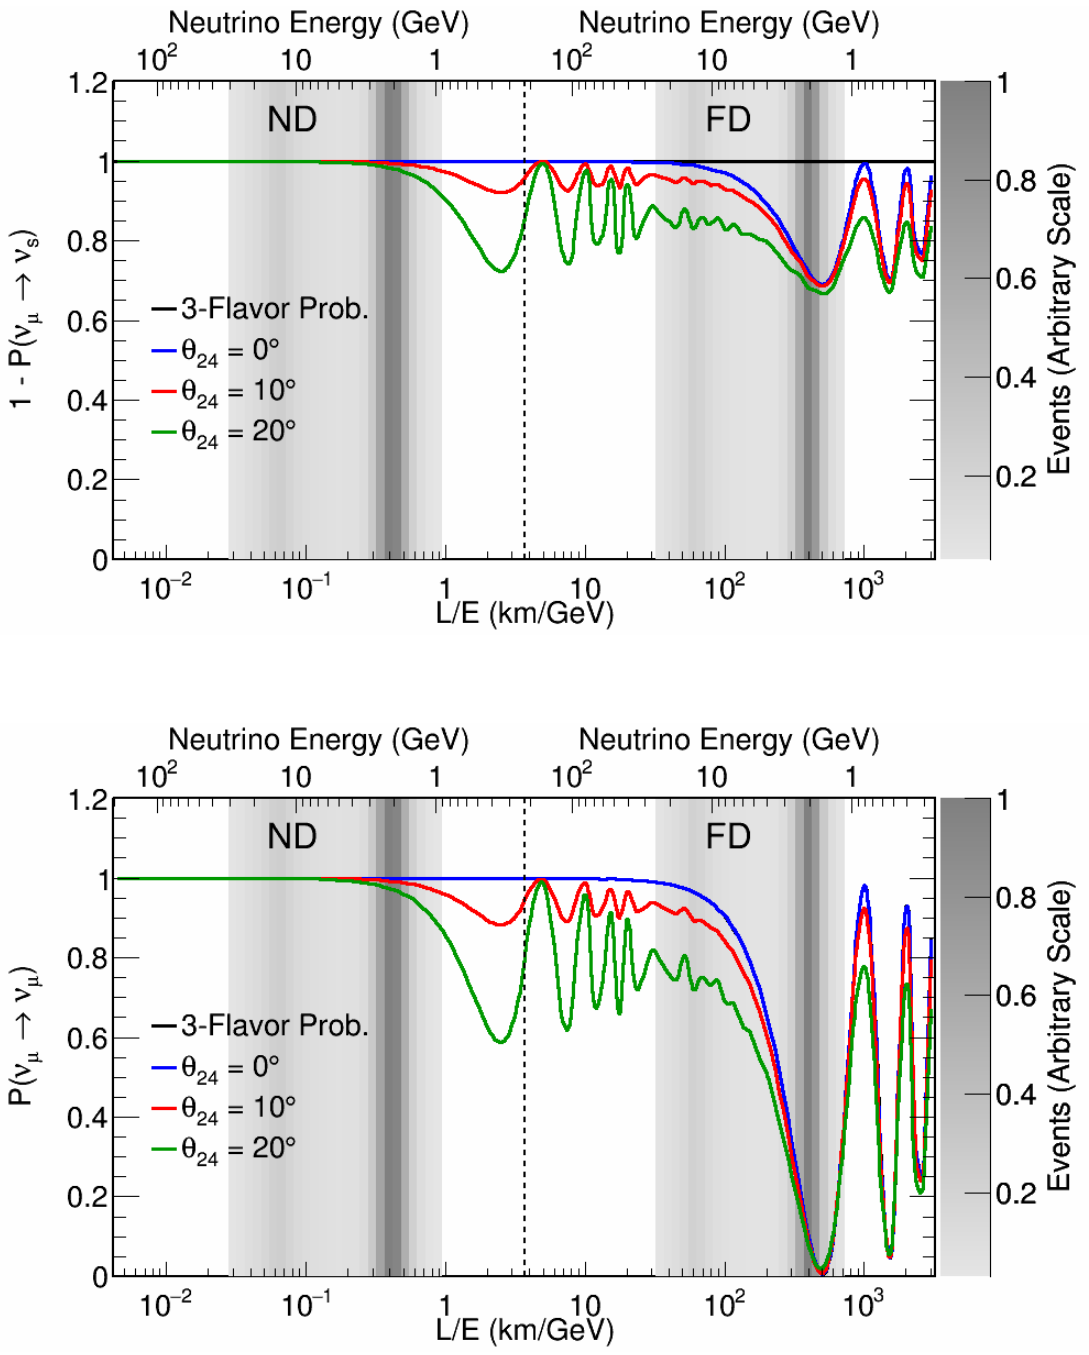


Figure 2.9: Total active neutrino survival probability (top) and ν_μ survival probability (bottom) for different values of θ_{24}. The grey bands show the neutrino energy spectra observed by NOνA, with darker colors representing a higher intensity of neutrinos. Parameters not varied in the figures are held fixed to values shown in table 2.2.

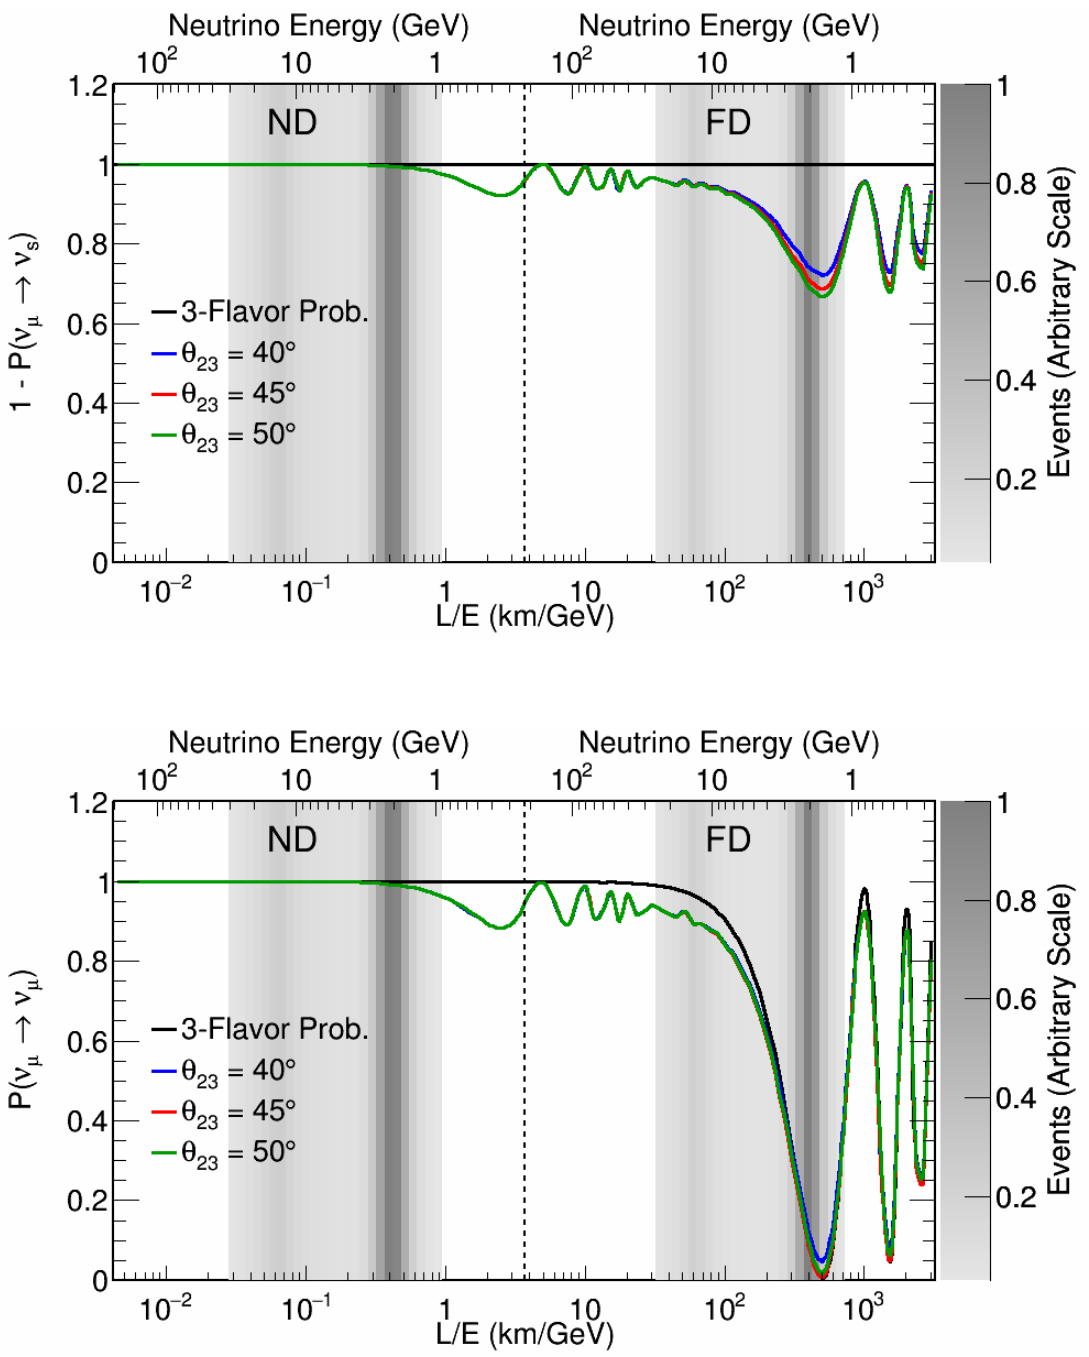


Figure 2.10: Total active neutrino survival probability (top) and ν_μ survival probability (bottom) for different values of θ_{23}. Shifts in θ_{23} have a modest effect at the NOνA peak energy. The grey bands show the neutrino energy spectra observed by NOνA, with darker colors representing a higher intensity of neutrinos. Parameters not varied in the figures are held fixed to values shown in table 2.2.

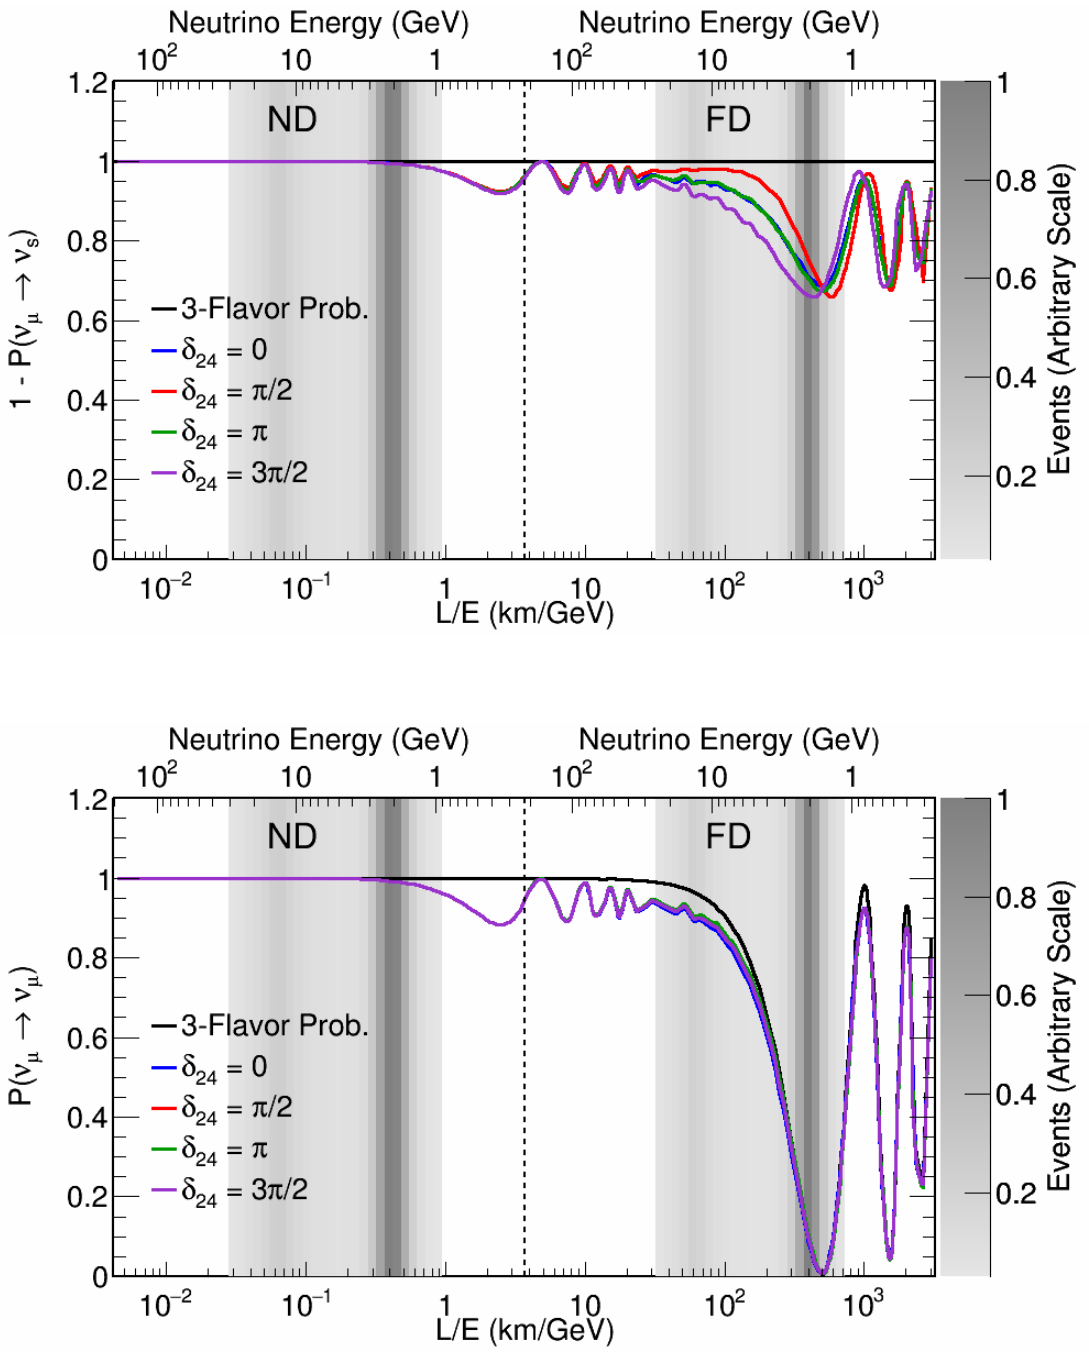


Figure 2.11: Total active neutrino survival probability (top) and ν_μ survival probability (bottom) for different values of δ_{24}. The grey bands show the neutrino energy spectra observed by NOνA, with darker colors representing a higher intensity of neutrinos. Parameters not varied in the figures are held fixed to values shown in table 2.2.

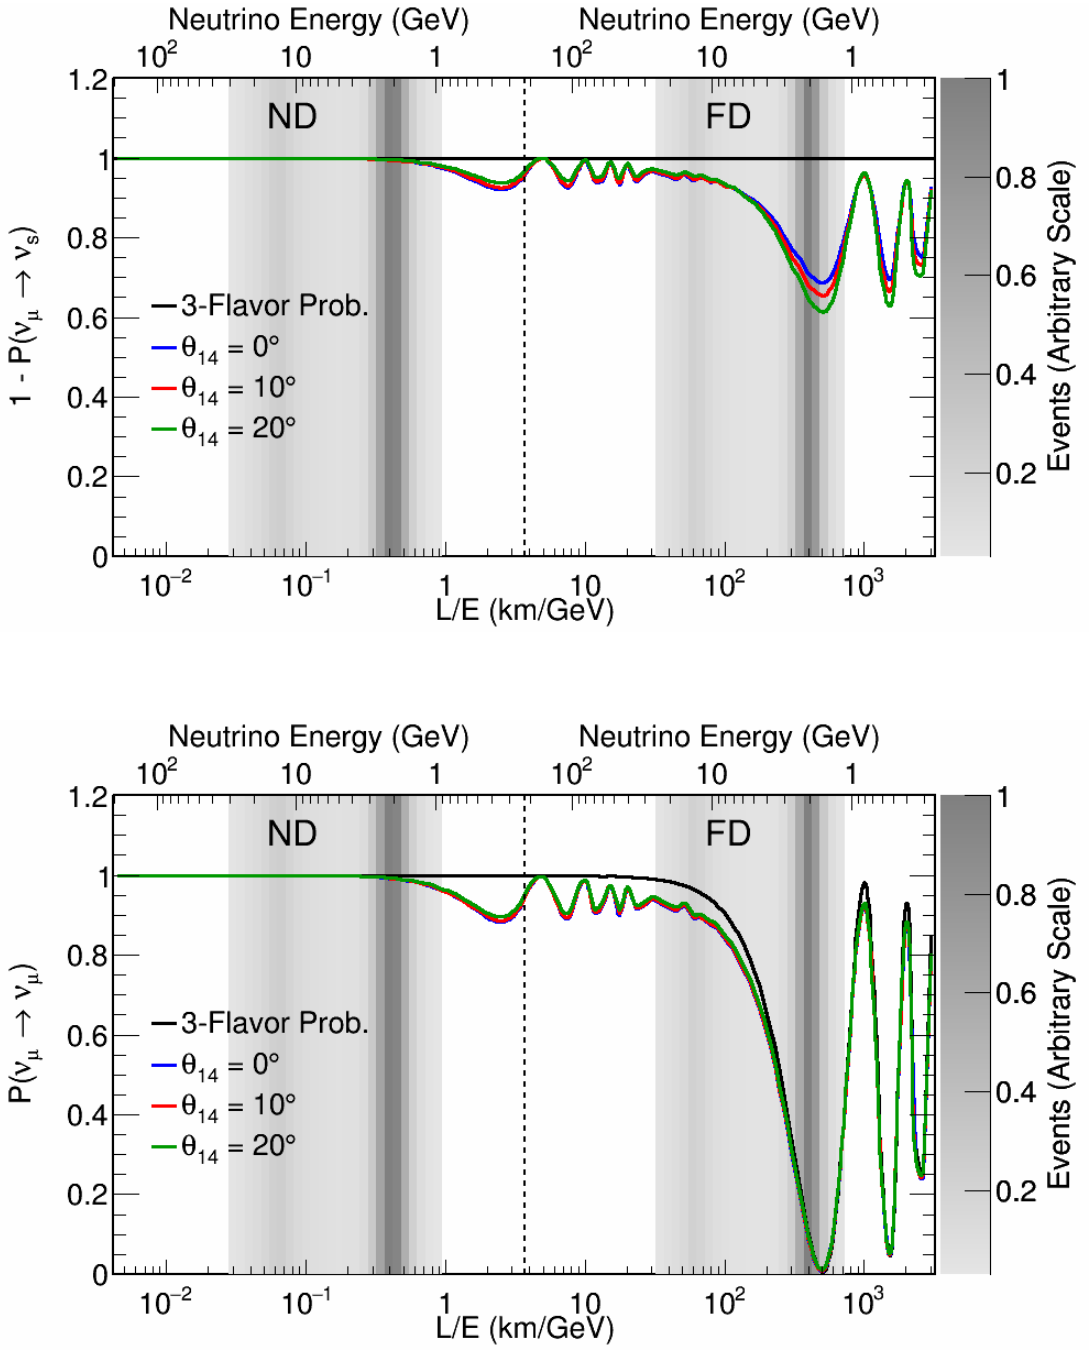


Figure 2.12: Total active neutrino survival probability (top) and ν_μ survival probability (bottom) for different values of θ_{14}. The values shown are well outside of experimental bounds, but even these large shifts have only a small effect on the oscillation probabilities that is largely degenerate with θ_{23}. The grey bands show the neutrino energy spectra observed by NOνA, with darker colors representing a higher intensity of neutrinos. Parameters not varied in the figures are held fixed to values shown in table 2.2.

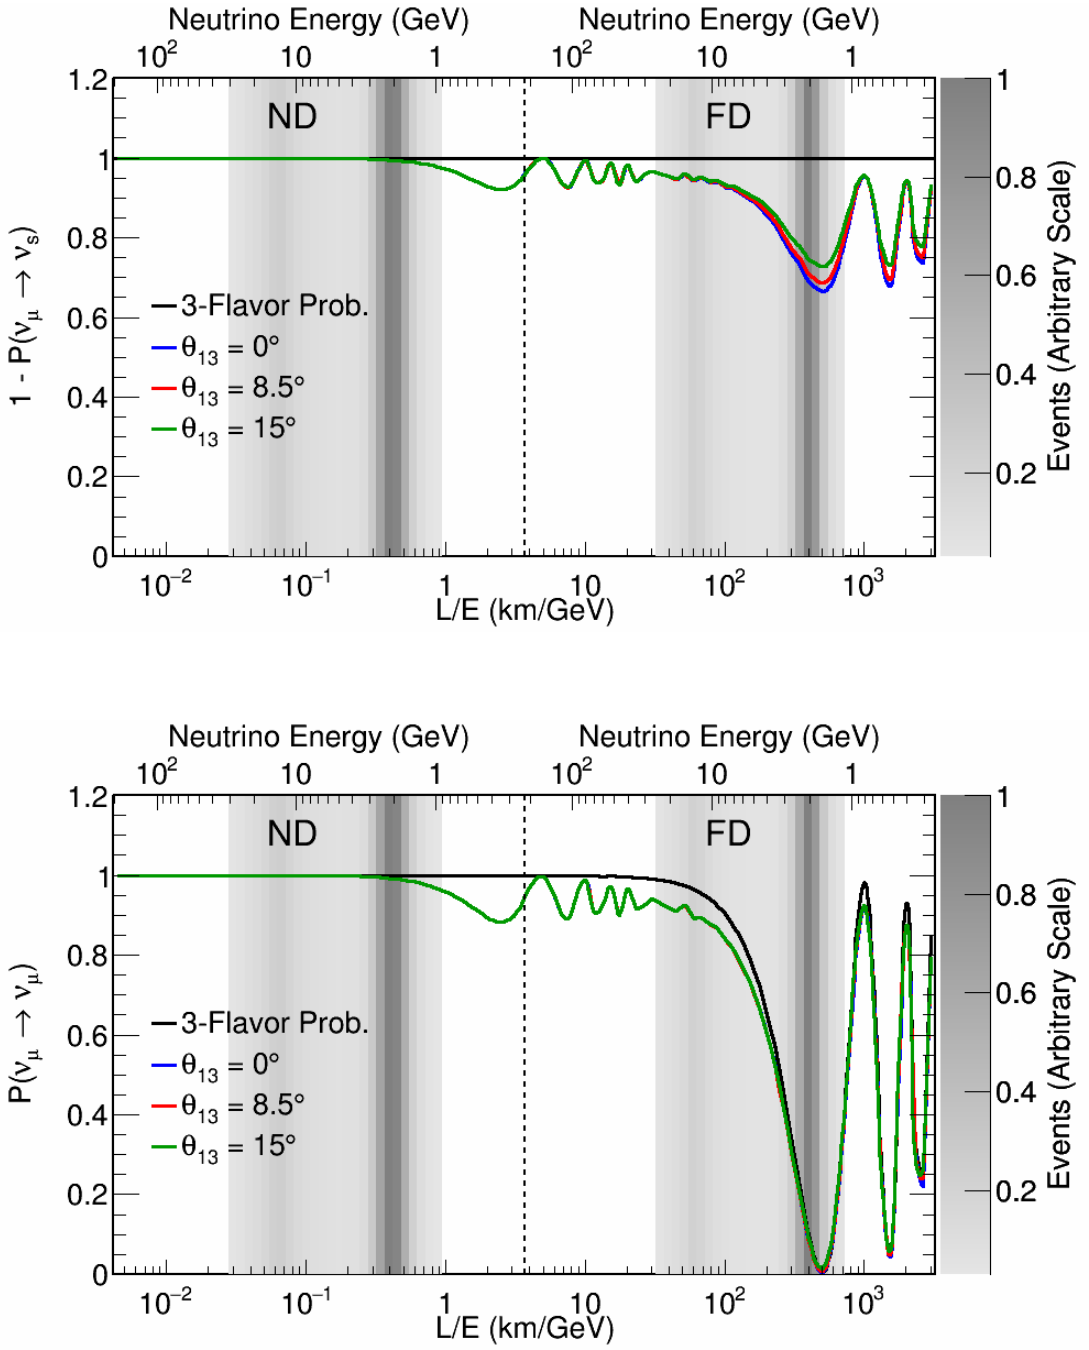


Figure 2.13: Total active neutrino survival probability (top) and ν_μ survival probability (bottom) for different values of θ_{13}. The values shown are well outside of experimental bounds, but even these large shifts have only a small effect on the oscillation probabilities that is largely degenerate with θ_{23}. The grey bands show the neutrino energy spectra observed by NOνA, with darker colors representing a higher intensity of neutrinos. Parameters not varied in the figures are held fixed to values shown in table 2.2.

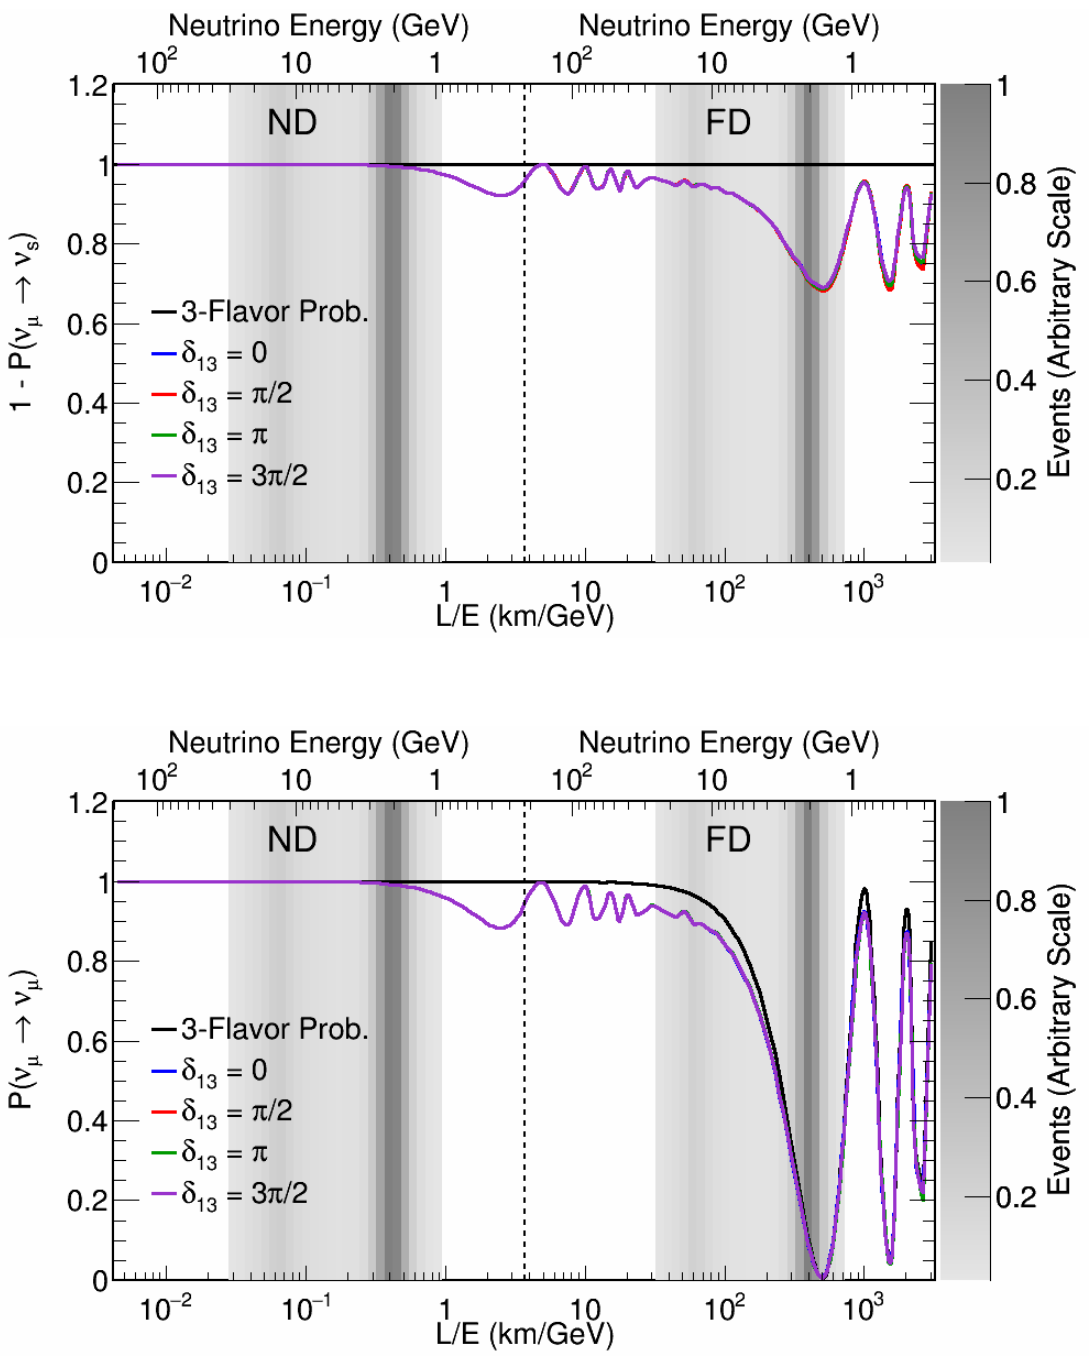


Figure 2.14: Total active neutrino survival probability (top) and ν_μ survival probability (bottom) for different values of δ_{13}. The effects of δ_{13} are negligible on the oscillation probabilities. The grey bands show the neutrino energy spectra observed by NOνA, with darker colors representing a higher intensity of neutrinos. Parameters not varied in the figures are held fixed to values shown in table 2.2.

all flux of neutrinos. Yet sterile neutrinos do not interact via the weak force at all, so if oscillations occur between active and sterile flavor states, this would also manifest in the NC rate as a deficit of events. Therefore, a NC event deficit would be a clear indication of physics outside the standard three neutrino flavor model, and this dissertation will interpret that in the context of oscillations to an extra, sterile, neutrino flavor state.

2.7 Sterile Matter Effect

In the standard 3 flavor neutrino model, matter effects the oscillation probabilities due to coherent CC forward scattering of electron (anti-)neutrinos from electrons in matter, adding an effective potential $V_e = \pm\sqrt{2}G_F N_e$ to the Hamiltonian. Following the formalism in reference [48], this is written in matrix notation as

$$V_{CC} = \pm\sqrt{2}G_F N_e \text{diag}(1, 0, ...), \tag{2.56}$$

where the ellipses denote that the remaining elements are 0. Of course, all neutrino flavors undergo coherent NC forward scattering off all particles in the earth, but this does not affect 3 flavor oscillation probabilities because it affects the neutrino flavors identically. This is no longer true when adding a sterile neutrino, and this effect must now be accounted for. The effective potentials due to the electrons and protons in matter cancel each other out due to equal densities, but the effect of the neutrons remains. This is adds an another piece to the effective potential.

$$V_{NC} = \mp\frac{\sqrt{2}}{2}G_F N_n \text{diag}(1, 1, 1, 0, ...). \tag{2.57}$$

The total effective potential is

$$V = \sqrt{2}G_F \left[\pm N_e \text{diag}(1, 0, 0, 0, ...) \mp \frac{1}{2}N_n \text{diag}(1, 1, 1, 0, ...)\right]. \tag{2.58}$$

The time evolution of the flavor states is written analogously to equation 2.28.

$$i\nu_\alpha = \left[\frac{1}{2E} U \,\mathrm{diag}(m_i^2, ...) U^\dagger + V\right] \nu_\alpha \tag{2.59}$$

As in section 2.4, the oscillation probabilities can be worked out by re-diagonalizing the Hamiltonian and finding the effective masses and effective mixing angles. However, with more than 2 flavors this process becomes quite challenging and tedious to do by hand. Instead, the exact effect can be seen by plotting the active neutrino survival probability for muon neutrinos in vacuum, for muon neutrinos in matter, and for muon anti-neutrinos in matter. These probabilities, and the ratios of the matter affected probabilities to the vacuum probability, are shown in figure 2.15.

The maximum size of the effect is less than 2% and occurs at about 1 GeV. At the peak sensitivity region for the NOνA Far Detector, the effect is even smaller at about 1%. To contrast, the MSW effect for ν_e appearance effects the oscillation probabilities by $\sim 20\%$ in NOνA [37]. Nonetheless, the small sterile matter effect is included in the analysis presented here.

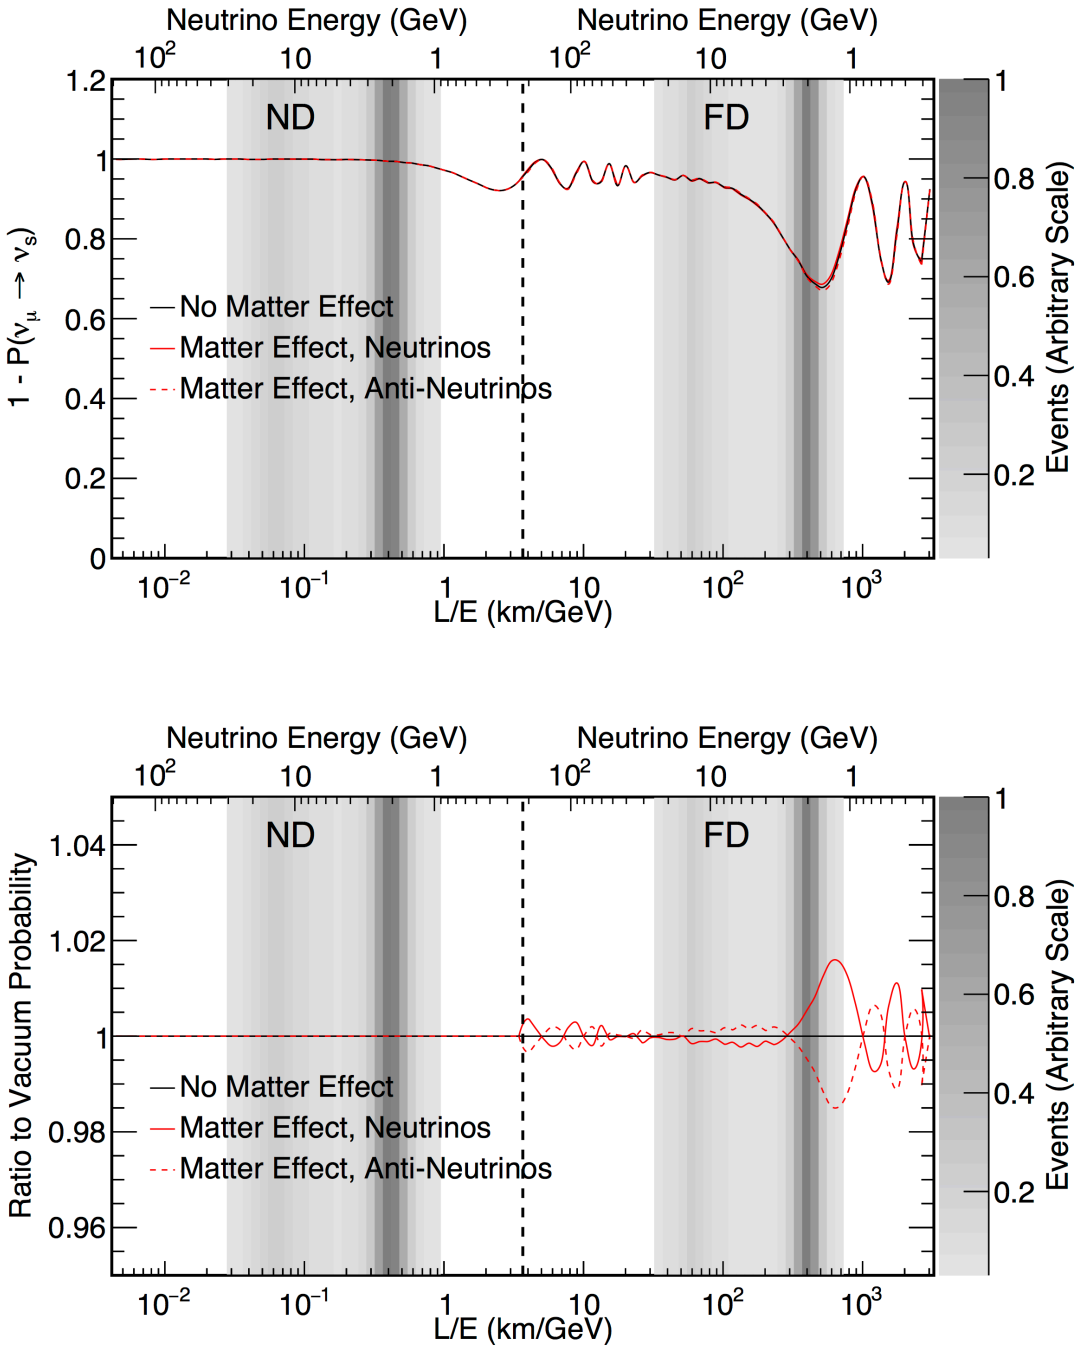


Figure 2.15: Total active neutrino survival probability demonstrating the size of the sterile matter effect. The probabilities are shown in the top figure, the ratios of the matter affected probabilities to the vacuum probability is shown in the bottom figure. The black line shows the survival probability in vacuum. The solid (dashed) red line shows the survival probability for muon (anti-)neutrinos in matter. The grey bands show the neutrino energy spectra observed by NOνA, with darker colors representing a higher intensity of neutrinos. Oscillation parameters are held fixed to values shown in table 2.2.

3

The NOνA Experiment

The NOνA experiment, or the NuMI Off-Axis ν_e Appearance experiment, is a long baseline, two detector, neutrino oscillation experiment designed to measure ν_μ to ν_e appearance from the Fermilab NuMI beam. The two functionally identical NOνA detectors are situated 14.6 mrad off-axis of the beam center, with its Near Detector (ND) located approximately 1 km downstream of the beam target source and its Far Detector (FD) located 810 km from the source near Ash River, MN. The main purpose of the ND is to constrain the neutrino beam energy and composition by measuring the beam very close to its source, i.e., before the neutrinos have had a chance to oscillate. The FD then measures the oscillated neutrinos. This chapter describes the NuMI beam and NOνA detectors in more detail.

3.1 The NuMI Beam

The NuMI beam is a ν_μ beam generated at Fermi National Accelerator Laboratory, or Fermilab, in Batavia, IL, and the source of neutrinos for the NOνA experiment. The neutrino beam is created by accelerating

protons to 120 GeV, colliding them with a graphite target, and allowing the products to decay to neutrinos.

The NuMI beam originates with the 120 GeV protons, and these are accelerated in the complex shown in figure 3.1. First, negatively charged hydrogen atoms are accelerated to 400 MeV in the Linac. These atoms next enter the Booster synchrotron where the electrons are removed and the protons are accelerated to 8 GeV. The output of from the booster is 13 bunches, 12 which are extracted into the Recycler Ring, each with approximately 4×10^{12} protons. In a procedure called slip stacking, additional bunches are used to double the intensity of each bunch. The first bunches in the Recycler are decelerated slightly while 6 new bunches enter the ring. As the two sets of bunches have slightly different energies, they slip relative to each other. When two bunches overlap they are captured with a special RF pulse, creating 6 larger bunches that are extracted into the Main Injector (MI) and accelerated to 120 GeV. Finally, these protons are directed to a target for neutrino production. At this point, the 6 bunches, or the spill, total about 5×10^{13} protons. With a cycle time of 1.333 s, the NuMI beam reaches 700 kW, which is the most powerful beam in the world [49].

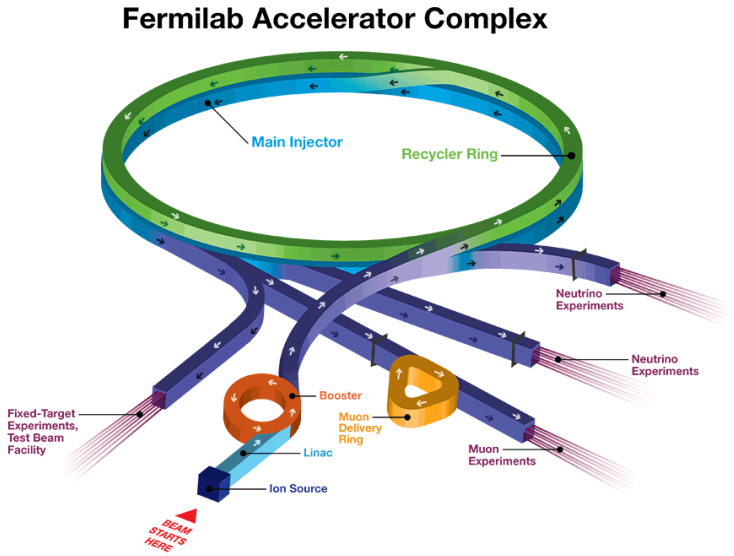


Figure 3.1: Schematic of the Fermilab accelerator complex.

In truth, the above is the design goal of the NuMI beam upgrades. The NuMI beam was originally de-

signed for the MINOS experiment [50], and it was operating at ∼200 kW with no slip stacking at the time of writing the NOνA Technical Design Report (TDR) in 2007 [49]. Since then, the Fermilab Accelerator Division has been steadily ramping up the beam power, as shown in figure 3.2. At the time of writing of this dissertation, the NuMI beam was in the midst of its 2016 summer shutdown. Before the beam was powered down, it most recently was running stably at about 560 kW with 6+4 slip stacking, or doubling the proton intensity in just 4 of 6 bunches [51, 52]. However, the NuMI beam briefly reached its design goal of 700 kW with full 6+6 slip stacking during a test on June 13, 2016 [53].

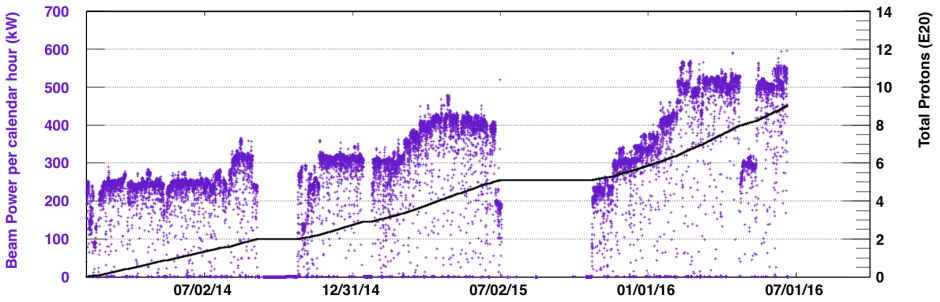


Figure 3.2: Beam power vs time since NOνA began collecting data. The average beam power is shown as purple points, and the integrated POT is shown in the black curve.

The generation of the actual neutrino beam begins with the extraction of the 120 GeV protons from the MI into the complex shown as a schematic in figure 3.3. The protons are directed 3.3° downward into the earth so the NuMI beam points directly toward the MINOS FD in Soudan, MN. First the protons collide with atoms inside a long, thin target. The target consists of 47 graphite 'fins' that are each 20.0 mm long, 6.4 mm wide, and spaced 0.3 mm apart for a total length of 95.4 cm. The beam size at the target is 1 mm $\times$ 1 mm, and the length of the target corresponds to ∼2 interaction lengths for the protons [50]. NOνA performs its data accounting by summing the protons on target, or POT. The interactions in the target produce a large number of secondary hadrons consisting mainly of charged pions with a smaller contribution of kaons. These particles next pass through two magnetic, parabolic horns. A current of 200 kA is passed through the horns creating a $1/r$ magnetic field that focuses a one sign of the charged particles and defocuses the opposite sign. The current in the horns can be reversed, which allows the NuMI beam to run in either neutrino or antineutrino mode. When positively charged particles are focused, they decay

into neutrinos, and similar for negatively charged particles and antineutrinos. Next, the focused charged particles travels through a 675 m long, 1 m radius decay pipe where the hadrons decay into neutrinos and tertiary charged leptons [49]. Finally, everything passes through a hadron monitor and absorber followed by three muon monitors spaced between sections of rock, leaving only the neutrinos to reach the ND hall.

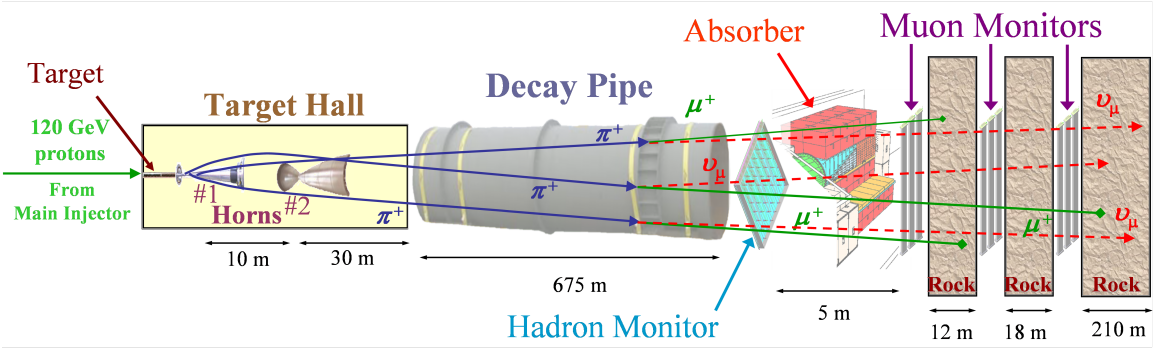


Figure 3.3: Schematic of the NuMI beam.

As a final note for this section, all of the data analyzed in this dissertation was taken during neutrino mode running, so the anti-neutrino mode is not considered further.

3.2 Off-Axis Detectors

The NOνA experiment was designed to have its detectors located 14.6 mrad off-axis from the NuMI beam, and this greatly affects the observed neutrino energy spectrum. The main decay mode for charged pions (99.98770%) and kaons (63.56%) is the two body decay into a muon and muon-neutrino [35]. In the center of mass frame of the decaying meson, this is a deterministic, isotropic decay. However, the lab frame is highly boosted resulting in the following neutrino flux and energy spectrum for a detector with cross section area A located a distance L from the decay point.

$$\phi_\nu = \left(\frac{2\gamma_m}{1+\gamma_m^2\theta^2}\right)^2 \frac{A}{4\pi L^2} \tag{3.1}$$

$$E_\nu = \left(1-\frac{m_\mu^2}{m_m^2}\right)\frac{E_m}{1+\gamma_m^2\theta^2} \tag{3.2}$$

Above, θ is the angle between the neutrino and meson, and the subscript m refers to the decaying meson, i.e., E_m is the meson energy, m_m is the meson mass, and γ_m is E_m/m_m. The effects of an off-axis location is shown in figure 3.4. There are two competing effects that occur. While the flux of neutrinos is drastically reduced, the neutrino energy is highly constrained, with a net result of a very narrow band beam of neutrinos.

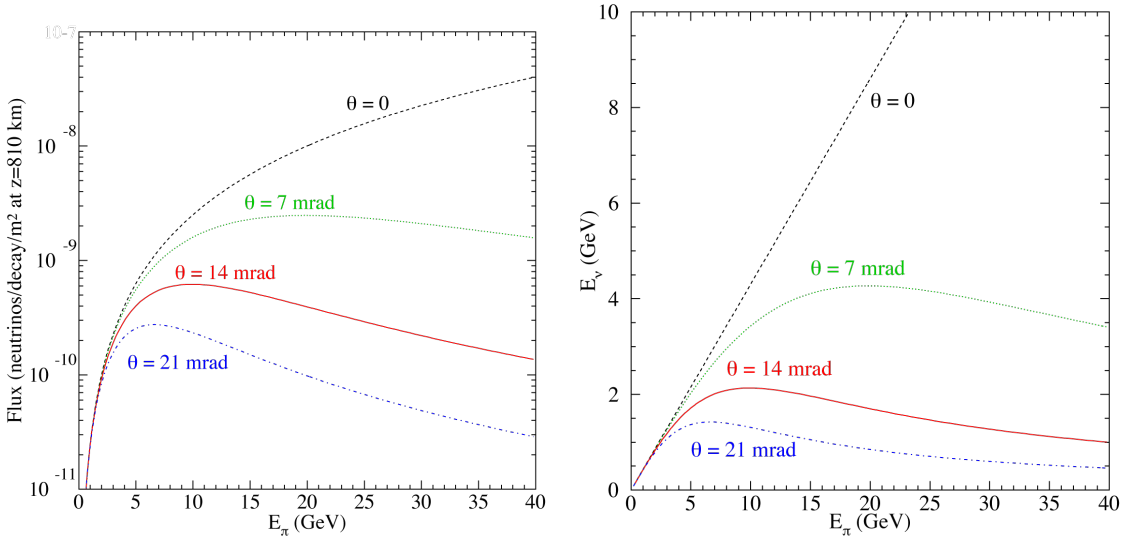


Figure 3.4: The neutrino flux (left) and energy (right) as a function of the decaying pion energy for several different off-axis angles. The on axis values are shown in black, and the red curves show the values for NOνA. Plots from [49]

The narrow energy band result is shown in figure 3.5, which compares the energy spectrum for several off-axis angles and shows composition of the neutrino beam observed by NOνA, both for the FD and without considering oscillations. The small contamination of $\bar{\nu}_\mu$, ν_e, and $\bar{\nu}_e$ components come from several effects. The focusing horns discussed in the previous section are unable to defocus the most energetic wrong sign mesons, which mostly contributes to the $\bar{\nu}_\mu$ component through the decay $\pi^- \rightarrow \bar{\nu}_\mu + \mu^-$. In the decay pipe, some of the muons from pion decay also decay, contributing to both the $\bar{\nu}_\mu$ and ν_e components. While most K^+ decay to ν_μ, a small fraction (5.07% [35]) decay via $K^+ \rightarrow \pi^0 + e^+ + \nu_e$. There is a very small $\bar{\nu}_e$ component that is largely due to K_L decay.

NOνA was designed to be off-axis in an attempt to optimize the amount of observed ν_e appearance, as the experiment name would suggest. At the NOνA FD, with $L = 810$ km, the neutrino energy range ob-

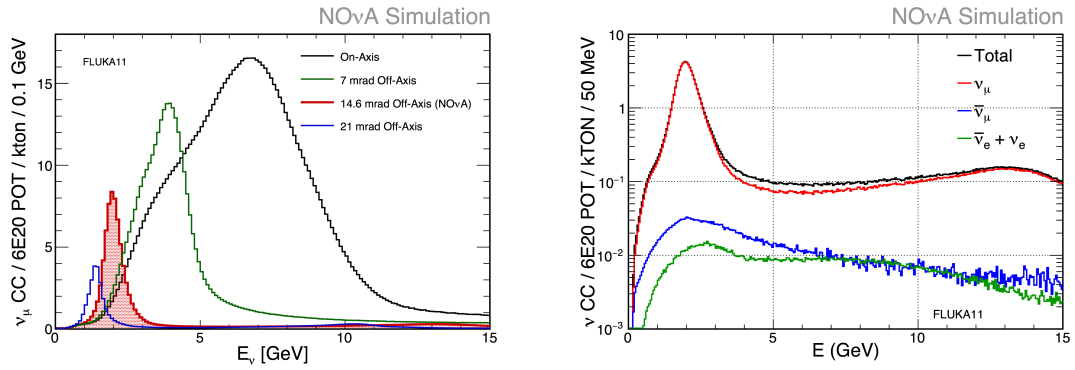


Figure 3.5: The neutrino flux (left) and energy (right) as a function of the decaying pion energy for several different off-axis angles. The on axis values are shown in black, and the red curves show the values for NOνA.

served due to the particular off-axis angle falls right within the first oscillation maximum for ν_e appearance. Figure 3.6 shows the oscillation probability for ν_e appearance and the event predictions for the dominant components, both as a function of the true neutrino energy. The narrow band beam and first oscillation peaks match quite well, allowing for the best possible measurement of ν_e appearance, one of NOνA's two main analyses [43]. Furthermore, the probability of ν_μ disappearance (shown only in the event component predictions) highly suppresses the number of expected ν_μ CC events, allowing for precision measurements of ν_μ disappearance, the other main NOνA analysis [44].

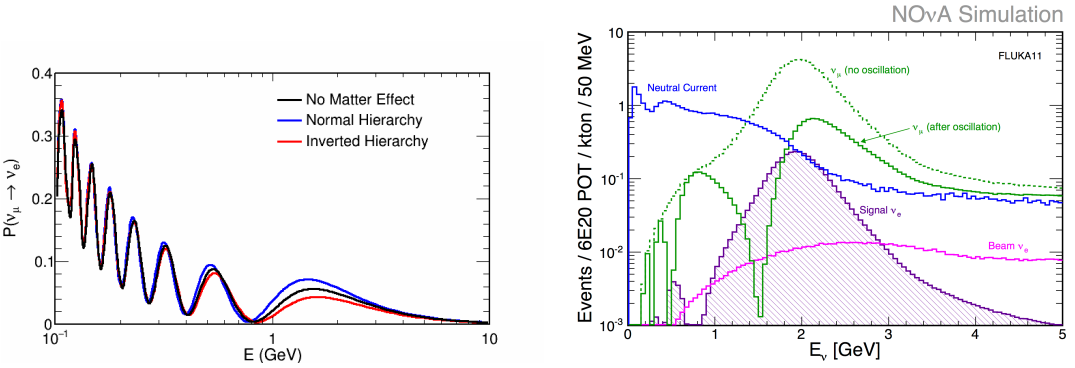


Figure 3.6: Left: the probability of ν_e appearance. The probabilities are calculated assuming standard three flavor oscillations using the relevant values shown in table 2.2, except for δ_{CP} which is here set to $3\pi/2$. Right: The expected number of events as a function of true neutrino energy (with the exception of NCs), for individual components. The NC component is shown as a function of true visible energy $E_{vis} \equiv E_\nu \times y_{Bj}$, where y_{Bj} is the fractional energy loss Bjorken scaling variable.

As with the main NOνA analyses, the off-axis nature of the experiment has both pros and cons when

searching for sterile oscillations via NC disappearance. Obviously a reduced neutrino flux is not great when searching for neutrino interactions. However, the CC backgrounds are highly suppressed due to oscillations, even with the ν_e appearance. This makes the selection of a relatively pure sample of NC events at the FD a much easier task.

3.3 The NOνA Detectors

The NOνA Detectors are highly segmented, liquid scintillating calorimeters. They consist of PVC plastic cells filled with a liquid scintillating solution, and the active scintillating material makes about 65% of the detectors by mass. The two detectors are functionally identical in order to reduce the effect of systematic errors, with the ND used to constrain the neutrino beam flux and energy and the FD used to measure the oscillated neutrino spectrum. This section first describes the general detector design, and later focuses on the specific details of the individual detectors. The information contained here is largely a summary of the NOνA TDR [49], which contains a much more detailed account of the detectors.

The basic structure of the NOνA detectors is a set of PVC plastic cells arranged together as planes and alternating in orientation. The cross section of the cells is approximately rectangular with dimensions of 4 cm $\times$ 6 cm. The cell length spans the entire height or width of the particular detector, depending on the orientation of the given plane. The PVC is extruded in groups of 16 cells, as shown in figure 3.7. Two extrusions are glued together to create a 'module,' shown in figure 3.8. Many modules are glued together side by side to form a single plane. The planes are oriented such that the NuMI beam direction is approximately parallel to the plane normal vector.

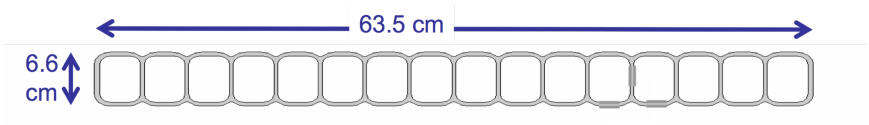


Figure 3.7: A schematic of a PVC cell extrusion. The PVC is extruded along the axis into the page with a length equal to the detector height or width. Figure from [49].

The local detector coordinate systems consider the NuMI beam direction as the Z axis, height (or altitude) as the Y axis, and the X axis is the other remaining direction. As the NuMI beam points approximately to the north, the direction of positive X happens to point west. (The beam travels in the positive

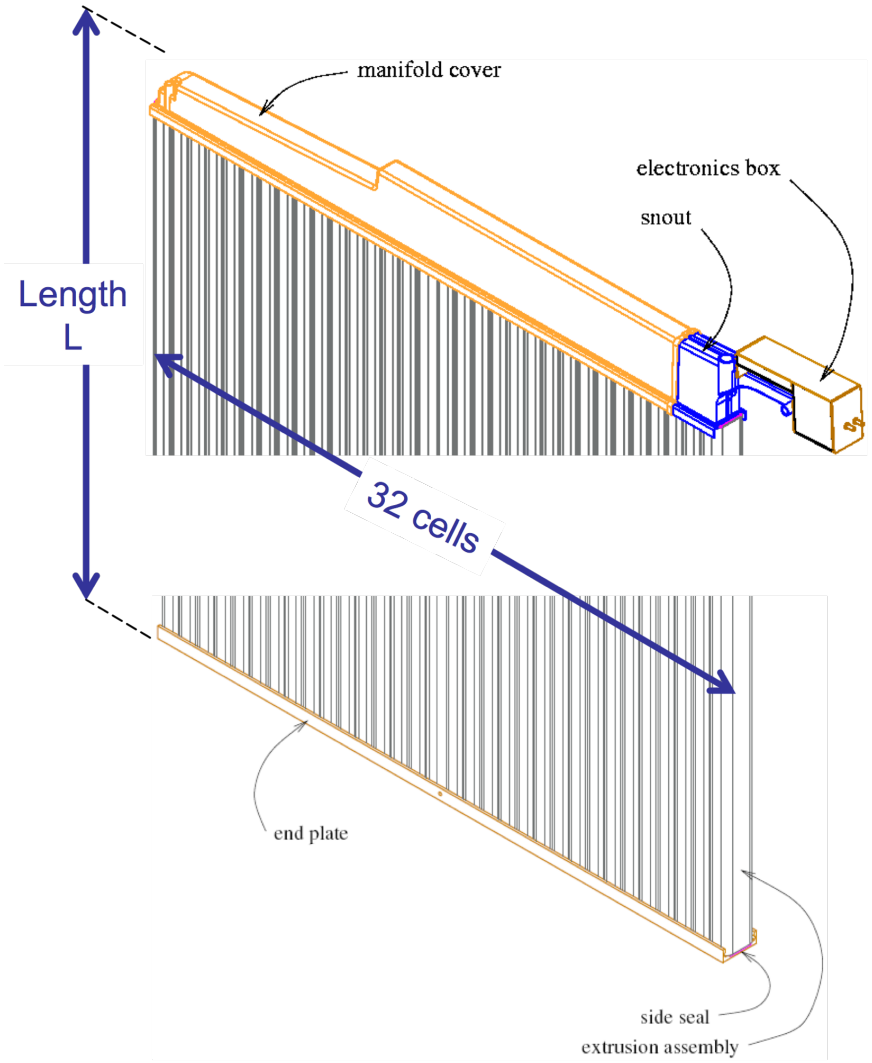


Figure 3.8: A schematic of a module. A full plane is made up of modules connected side by side. Figure from [49].

Z direction, and the positive Y direction points to higher elevation.) In this frame, a given plane can either locate activity in an XZ plane if the cells are in the vertical orientation, or a YZ plane if the cells are in the horizontal orientation. For this reason, the planes alternate between horizontal and vertical orientations to give a full three dimensional view of detector activity.

Each cell is filled with a liquid scintillator solution that consists of several different compounds, each with a specific function. Pseudocumene (1,2,4-trimethylbenzene) is the scintillant used in the solution. The solution contains two wave shifters, PPO (2,5-diphenyloxazole) and bis-MIS (1,4-bis(2-methylstyryl) benzene). A small amount of Stadis-425 is used in the solution as an antistatic agent, and a small amount

of tocopherol (Vitamin E) is added as an antioxidant. Mineral oil provides the solvent for each of the above compounds and makes up the majority of the solution. Table 3.1 summarizes the solution contents and lists their mass fractions. The resultant mixture produces scintillation light in the low wavelength UV region, and shifts this to the violet-blue region.

Compound	Purpose	Mass Fraction
Mineral Oil	Solvent	95.8%
Pseudocumene	Scintillant	4.1%
PPO	Waveshifter	0.091%
bis-MSB	Waveshifter	0.0013%
Stadis-425	Anti-static	0.0003%
Vitamin E	Antioxidant	0.0010%

Table 3.1: A summary of each component in the scintillation solution. Adapted from [49].

To collect the scintillation light, the inside of each cell is coated with titanium dioxide and a loop of a plastic wavelength shifting (WLS) fiber rests inside each cell. The TiO_2 is used as a reflective agent to maximize the scintillation light collection. A fluorescent dye in the WLS fiber absorbs this scintillation light and shifts it to the blue-green region. This light is then transmitted through the fiber to electronics for readout. A complete schematic of a cell is shown in figure 3.9.

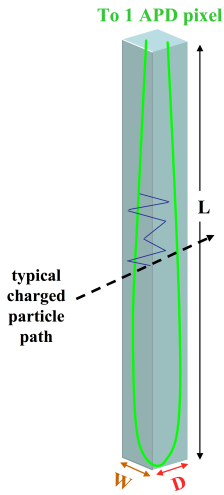


Figure 3.9: A schematic of the inside of a cell. The width and depth are approximately $W \approx 4$ cm and $D \approx 6$ cm, respectively. The green curve is the WLS fiber, and the zigzagging blue line is a path light might take from emission by the scintillator through absorption by the WLS fiber. Figure from [49].

The first component in data readout is collection and amplification by a Hamamatsu avalanche photodiode (APD). A single APD collects the light from all 32 cells in a module, collecting light from both ends of the WLS fiber. Figure 3.10 shows an APD and the ends of the WLS fibers set in a mount to be placed next to the APD. APDs were chosen over PMTs due to their high quantum efficiency, especially for green light. This is shown in figure 3.11 against the light output of the WLS fibers.

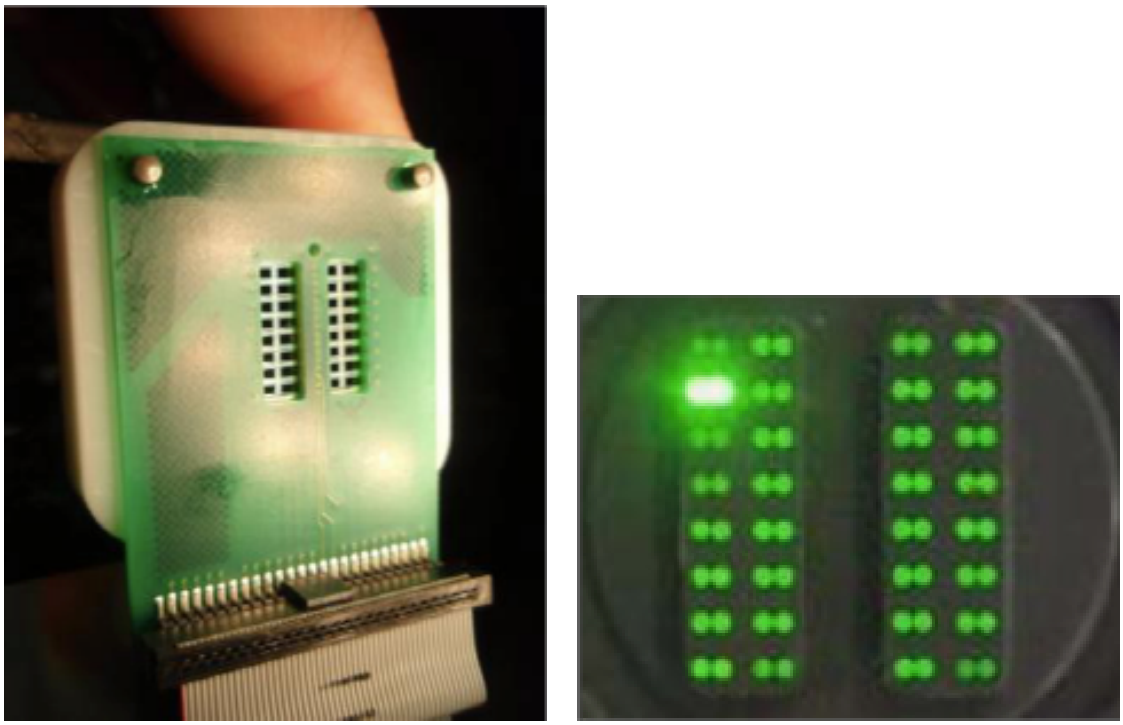


Figure 3.10: Left: An APD. Right: The ends of 32 WLS fibers that will be pressed against the APD face shown on the left.

Each APD is connected directly to a thermoelectric cooler and housed in an electronics box as shown in figure 3.8. The cooler keeps the operating temperature for the APD at -15° C to reduce noise. The APDs are held at a voltage of about 425 V to run at a gain of 100, though exact voltages are set individually. In October 2015, the gain for the FD APDs was increased to 150 to be more sensitive to light from the far ends of the fibers. The thermoelectric cooler is part of a larger front-end board (FEB) that is also housed inside the same electronics box as the APD. The FEB thus provides the cooling and voltage for the APD, but it also digitizes the signals from the APD.

Besides cooling the APDs, the FEBs also read the output from the APDs and prepare the signal for

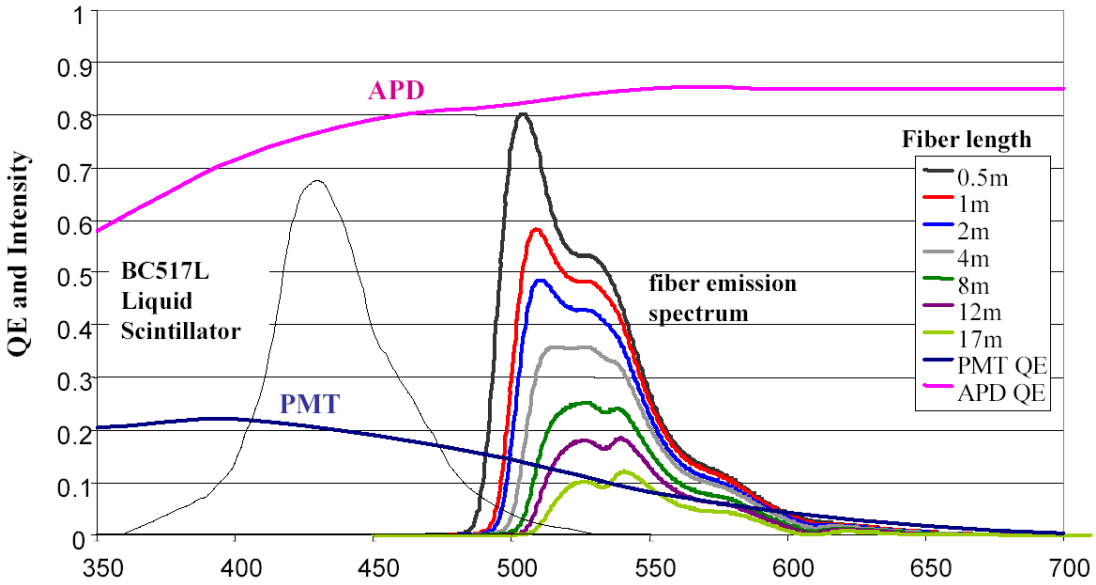


Figure 3.11: Quantum efficiency of APDs as a function of light wavelength. For comparison, the QE of PMTs is shown in navy. The light output of the liquid scintillator is shown in the thin, black curve. The light output for the WLS fibers is also shown for a variety of different lengths through the fiber. The intensities are in arbitrary units. While the blue light is almost fully attenuated through the fiber, the green light is only modestly attenuated. Figure from [49].

data storage. This part of the FEB consists of three main chips, an application specific integrated circuit (ASIC), an analog to digital converter (ADC), and a field programmable gate array (FPGA) [54, 55]. The ASIC continuously integrates and shapes each channel from the APD, with separate, programmable, exponential rise and fall times [56]. The rise time at the Far (Near) Detector is set to 460 (140) ns; the fall time is 7000 (4500) ns [37]. Each channel is periodically sampled in round robin fashion based on 16 MHz clock, or every 62.5 ns. For the FD (ND) FEBs, 4 (16) groups of 8 (2) channels are sampled and multiplexed based on this clock, so each channel is sampled every 500 (125) ns [56]. The multiplexed signals from the ASIC are next digitized by ADCs. The output from the ADCs are lastly fed to the FPGA, which determines whether a pulse is signal and extracts the pulse height. A pulse is considered signal if the quantity $ADC_i - ADC_{i-3}$ [57] is above a threshold set individually for each channel. To set these thresholds, traces are collected in *digital oscilloscope* (DSO) mode to find the baseline signal and noise spectrum. The threshold is set at 4σ greater than the RMS of the noise [58]. The samples associated with the pulses above threshold are the final readout.

The FEB output can be configured in single point or multipoint readout. In single point mode, only the pulse height $ADC_i - ADC_{i-3}$ mentioned above is read out. In multipoint mode, all baseline subtracted samples, $s_i \ldots s_{i-3}$, are read out, where s_i is the first sample above threshold. The pulse shape is fit offline for better timing resolution using all of this information. The ND was always run in multipoint mode; the FD was run in single point mode until October 2014 [37].

The electronics for the NOνA detectors were installed in 64 plane sections called diblocks. The signals from (up to) 64 FEBs are read out by a single data concentrator module (DCM). Due to the alternating orientations of the planes, and thus the modules, electronics boxes are found on either the top or one side of the detector. Consequently, a given DCM does not read out one FEB per plane, but rather (up to) two FEBs per plane. The two FEBs are always from adjacent modules, and the modules always share the same X or Y coordinate range based on the module orientation. This modular structure is important because it allowed for data collection at the FD before the detector was fully commissioned.

Data stored by the DCMs are stored in 50 μs blocks called microslices, which are transferred to a data buffer farm in larger 5 ms blocks called millislices. Data is stored in the farm for up to 20 minutes before being erased. During this time, different trigger decisions can be applied to keep the data for different purposes such as recording beam activity, searching for good calibration candidates, supernova detection, or searching for exotic physics like monopoles. Data that pass a trigger decision are written to disk and permanently stored.

Detector activity that could be due to a beam event is captured within a 500 μs window. These events are specifically triggered by a signal from Fermilab, corrected for time of flight at the FD. The event window is centered around the $\sim$10 μs beam spill, which allows for considerable drift in the timing systems without loss of data [37]. Other 500 μs event windows without beam candidates are stored as cosmic trigger data that are useful for predicting the cosmic background event rate.

Data are collected by organizing these events into larger run and subrun file structures. Subruns are populated with events until either a maximum file size is reached, or after an hour of data taking has elapsed, whichever comes first. Runs are comprised of 64 subruns for the FD, and 24 subruns for the ND, and are numbered in succession such that larger numbered runs correspond to later times.

3.3.1 Far Detector

The NOνA FD is situated 810 km from the beam target near Ash River, MN. Both the horizontal and vertical planes consist of 12 modules, or 384 cells, for a total cross sectional area of about 15.2 m $\times$ 15.2 m. The fully commissioned detector has 14 diblocks (data taking began with as few as 4), or 896 planes, with a total length of about 59.6 m. This means that the detector reads out 344,064 individual cells, or channels. The FD weighs 14 kt, which is about 65% scintillator and 35% PVC plastic.

The detector itself is inside a building that is at the surface, mostly dug underground and partially above ground. The site is at an elevation of about 1220 ft and thus subject to a large background of cosmic rays. A large amount of overburden is used to mitigate this effect. 2.5 ft of precast concrete covers the entire detector and is capped with 1.5 ft of concrete cast in place. Finally, 6 in of barite, or barium sulfate, is placed above the concrete and composed of loosely packed rocks. Barite was chosen because it is a high Z material with a high stopping power for cosmic ray photons. Together, the concrete and barite provide around 14 radiation lengths for the cosmic rays. Figure 3.12 shows the FD and the building inside which it is located.



Figure 3.12: Left: The NOνA FD. The beam comes from the back of the photograph and out of the page. The red and yellow structure is a piovter that was used to rotate and put the detector blocks into position and is not part of the detector. A person can be seen halfway up on the right of the detector for scale. Right: The FD building. The detector is beneath the rocky section on the right of the photograph, which only partially shows the barite overburden.

Even with the overburden, the FD is still subject to cosmic rays with a rate of 150 kHz [59]. This is the main reason why NOνA must use a trigger from the Fermilab accelerator to signal a beam event, but also why taking cosmic trigger data is so important. These data are used to develop and train cuts designed to remove cosmic background. Later, a predicted background is made from an independent sample, namely the data regions in the NuMI trigger events that fall safely outside the smaller beam window.

3.3.2 Near Detector

The NOνA ND is situated 1015 m from the beam target, on site at Fermilab. It is composed of two different sections, a fully active region much like the FD and a muon catcher. In the fully active region, each plane consists of 3 modules, or 96 cells, for a total cross sectional area of about 3.8 m $\times$ 3.8 m. The region is made up of 192 planes with a length of about 12.7 m and corresponds to 3 diblocks for the electronics. The NuMI beam first passes through this fully active region, but since it is so much shorter than the FD, the muon catcher is necessary to be able to accurately measure muon energies. For this section, planes of steel, which have an excellent stopping power for muons, are interspersed between the alternating planes. The energy loss of muons passing through steel is well known allowing for accurate event reconstruction. This region is specifically made of a horizontal plane, a 10 cm steel plane, and a vertical plane, repeated 11 times for an additional 22 active planes. However, for economic reasons the steel and planes in this region were recycled from a prototype near detector that was built using the original design of 3 modules $\times$2 modules. Thus, the planes in the muon catcher are as wide as the fully active region, but only 2/3 as tall, at about 2.6 m. In total, the ND is 15.8 m in length, has 20,192 cells, and 214 active planes. It weighs 290 t, of which 130 t is scintillator, 81 t is PVC plastic, and 78 t is steel [60]. Figure 3.13 shows the ND from both beam directions, and table 3.2 summarizes the properties of both detectors.

Unlike the FD, the ND is not a surface detector. It is located 105 m below the surface in an underground cavern adjacent to the MINOS ND, as depicted in figure 3.14. This means that the ND is not subject to a cosmic ray rate like FD. However, given the close proximity to the beam source, the ND can observe multiple neutrino events in every beam spill. This leads to the potential for pileup; this is why the electronics for the ND were designed to sample the data 4 times faster than the FD.



Figure 3.13: The NOνA ND. Left: View from the side where the beam enters the detector. The red structure is a bookend and not part of the active detector. A person is seen on the right for scale. Right: View from the side where the beam exits the detector. The red and grey structure above the muon catcher is for worker access to this detector region and not part of the active detector.

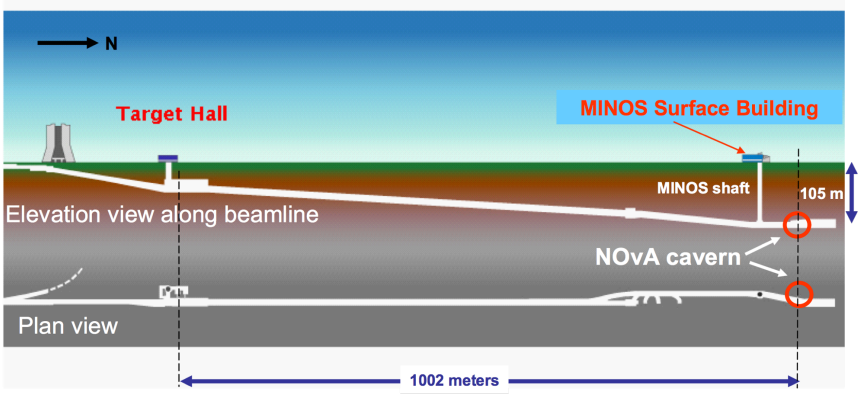


Figure 3.14: A schematic of the NuMI beam line, showing the location of the ND hall. Figure from [49].

Detector Property	FD	ND
Detector Weight (kt)	14	0.3
Number of Diblocks	14	3 (+Muon Catcher)
Number of Planes	896	214
Number of Cells	344064	20192
Detector Length (m)	59.6	15.8
Detector Width (m)	15.8	4.2
Detector Height (m)	15.8	4.2 (2.8)
Cell Depth (cm)	5.6	5.6
Cell Width (cm)	3.6	3.6
Cell Length (m)	15.2	3.8 (3.8H, 2.6V)

Table 3.2: A summary of the physical properties of both detectors. For the ND, a single value means the property is consistent across the entire detector. Otherwise, values in parenthesis are for the muon catcher only. Note that the detector width and height exceed the cell lengths due to inclusion of the end caps on either side of the planes.

4

Experiment Simulation

Simulated data, or Monte Carlo (MC), plays a very important role in the NOνA analyses. The FD only expects $O(10^2)$ events per year, not nearly enough statistics to fine tune things like selection cuts. Furthermore, it is important that actual parameter measurements use an independent sample of data from that used to train and tune an anlysis. MC solves both of these problems by being explicitly different from the real data and having only computer memory limit statistics.

The full simulation chain is very long and comprises multiple steps, which reduces complexity and provides more opportunity for validation. At every step, the information from previous pieces of the simulation chain are kept so it is always possible to reproduce future results and trace any errors that may occur. The simulation chain consists of two main components, simulation of the beam and the detector response. This chapter discusses both of these areas in more detail, and finishes with a discussion of packages designed to study and validate the results.

4.1 Flux Simulation

The first main simulation segment is the simulation of the NuMI beam, or the flux simulation. This section starts with the interaction of protons in the target and ends with experiment independent flux files full of neutrino rays that most importantly contain the flavor, direction, energy, and momentum of the neutrinos.

The flux simulation is performed using the FLUGG package, an interface between FLUKA [61, 62] and Geant4 [63, 64]. The version used for the analysis in this dissertation was FLUGG `2009.3`, combining `FLUKA2011` and Geant4 `v4.9.6.p03(c)`. The FLUKA package is designed to simulate particle interactions and was used to simulate the proton interactions in the target. Geant4 is a toolkit for simulation of particles propagating through matter, and is used in this context to propagate the target interaction products through a detailed model of the NuMI beam line. This model of the beam includes all of the elements discussed in section 3.1, starting from the target hall all the way through the rock before the detector halls. The second focusing horn, a representative component of the geometry model, is shown in figure 4.1.

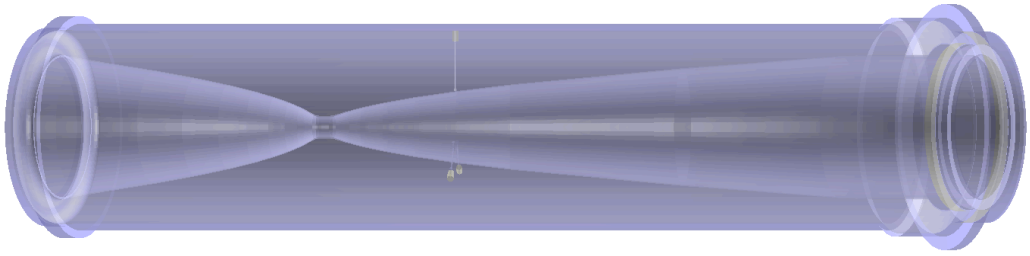


Figure 4.1: Visualization of the geometry used to model the second horn. This image was made during efforts to correctly position the three spider supports seen near the center of the schematic. Figure from [65].

The output of the FLUGG simulation is a set of flugg format flux files. These include the information described at the beginning of this section, but have several other key features as well. Simulating the particle products of every target interaction is extremely expensive on computing resources, so not every particle is tracked. Instead, similar neutrinos are often dropped in lieu of a single entry that is given a greater *importance weight*, one of two weights stored in the flux files. The information about the neutrino parent is also stored in the files, including the particle type, energy, momentum, and decay point. This affords later simulation steps the ability to "re-decay" a neutrino ray to a specific location and calculate the relative

probability that this would occur. The latter value is stored as a *propagation weight*. Retaining the parent information also allows for event reweighting based on hadronic model studies.

The next piece of the flux simulation is largely a repackaging of the flugg format flux files into a unified format, called Dk2Nu [66]. FLUGG is certainly not the only flux generator, and NOνA is not the only experiment that uses the NuMI beam line. The Dk2Nu format is unified in the sense that it has the same format regardless of flux generator, and valid for any experiment using the NuMI beam.

Dk2Nu format flux files (or simply dk2nu files) are one of the main outputs from the flux simulation. However, one last step is taken to further simplify computation time for later simulation steps. The entries in the dk2nu files are sampled over a 'window' that shadows the detector of interest, resulting only in a set of neutrino rays that could create interactions within the detector. The output files from this final step are in the GSimpleNtpFlux format, or more simply referred to as gsimple files [67].

4.2 Detector Simulation

The second main section in the simulation chain is the detector simulation. This section begins with neutrino interactions, models many different detector response effects, and finishes with files of simulated raw hits, or the signals read out by the detector electronics. The latter portion of this section closely follows the NOνA Simulation Technical Notes [68, 69], which provide more detail and further references.

4.2.1 Neutrino Interactions

Neutrino interactions are modeled using the GENIE event generator product [70, 71]. For the analysis in this dissertation, version `v2.10.4` was used. GENIE requires as input flux files, a model of the detector geometry, and cross section information, and from this determines if and where a neutrino interaction occurs, the type of interaction, and the kinematics of the interaction products. The generator uses a sophisticated model of the nucleus that allows for quasi-elastic, resonant, or deep inelastic scattering, taking into account any intranuclear scattering that may occur after the initial interaction takes place. The output of a GENIE interaction is a list of primary particles that escape the target nucleus and the relevant kinematic variables that describe each entry.

There is growing evidence [72, 44] that there is a type of nuclear interaction involving a quasi-bound nucleon-nucleon pair. This interaction is coined the meson exchange current, or MEC, and is included as an empirical model in GENIE. These events largely occur in an energy region between quasi-elastic (QE) and resonant (RES) events, vanishing by 5 GeV where QE scattering from free nucleons model external data well. While the inclusion of MEC interactions does not bring the data and MC into complete agreement, it substantially improves the discrepancy; see figure 4.2. Consequently, an additional systematic uncertainty is evaluated for the MEC scale, discussed in chapter 8. Furthermore, data used to constrain the empirical model in GENIE was only available for CC interactions, so there are no NC MEC events in the MC [73].

In truth, the actual 'tuned' MEC model is applied as a weight on an event level, and there are a couple other effects applied in a similar fashion. The effect of long-range nuclear charge screening, or RPA (random phase approximation) suppression was studied by creating a shifted spectrum using a set of event weights. This shifted spectrum was translated into a systematic uncertainty, assessed in chapter 8. The single non-resonant pion production rate was decreased also using event weights to a level consistent with the most recent data studies [73].

The GENIE output is next fed back to Geant4 for particle propagation through the detector geometry [63, 64]. It models the primary particle interactions within the detector, including energy deposition and secondary interactions. The physical processes considered and modeled are somewhat configurable via the choice of the physics list. At this point in the simulation, it is much more important to accurately track the detailed interactions, so particles are not dropped as in the flux simulation; instead, a more precise model was used here. The specific physics list used was `QGSP_BERT_HP`, where QGSP specifies that high energy hadrons (> 10 GeV) are modeled with the quark gluon string model, BERT speficies that lower energy hadrons (< 10 GeV) are modeled with the Bertini cascade model, and HP toggles the usage of a high precision neutron model that tracks low energy neutrons (< 20 MeV) [68]. The output from Geant4 at this stage is a list of particles involved in the detector interaction and a full suite of information about each one, called the MC truth. The information that is most important to the next simulation step is a list of energy depositions and positions.

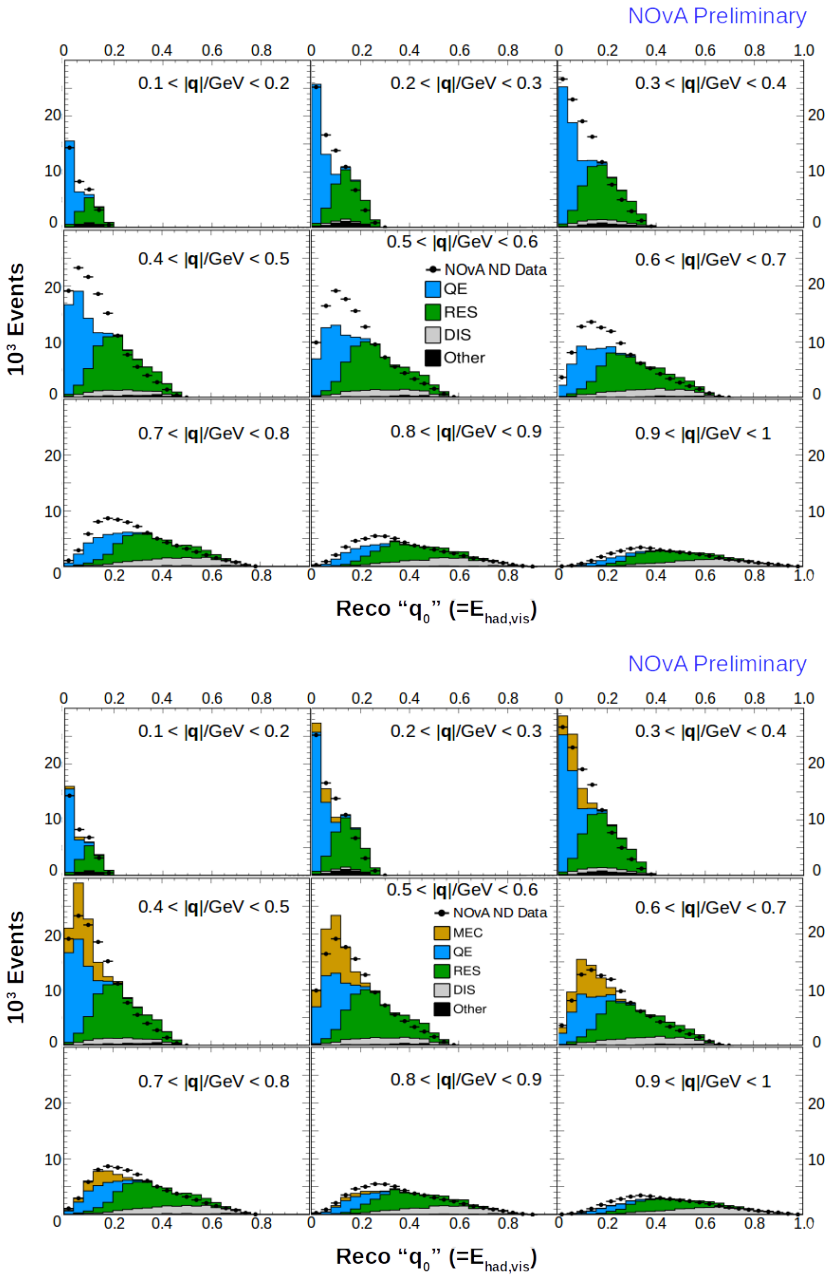


Figure 4.2: Distributions of the reconstructed visible hadronic energy for different 'slices' of estimated three momentum transfer, with (bottom) and without (top) MEC events. The events used for these distributions come entirely from ND selected ν_μ CC events. Plots from [74].

It is worth noting that the simulated events for the Near and far Detectors are generated somewhat differently. The FD expects a very low event rate per beam spill, while the ND expects multiple events per spill. As a result, the FD events are generated one at a time with the simulation calculating the POT needed for the event occur, and ND events are generated assuming a constant POT/spill and a variable number of events. Furthermore, the high rate of interactions at the ND means that it observes many muons that originate from neutrino interactions that begin in the rock outside of the detector. Most of these so called rock events do not result in any detector activity, so tracking them wastes valuable computing time. To deal with this, the ND detector events and single rock events are generated separately, and only the rock events that cause detector activity are kept. The final ND event records are made by overlaying rock events on the detector events. The number of rock singles to overlay is drawn from a Poisson distribution with a mean determined from data [75]. A top down view of the neutrino interaction vertices at the ND is shown in figure 4.3 and shows the level of detail used to model the geometry of the ND hall.

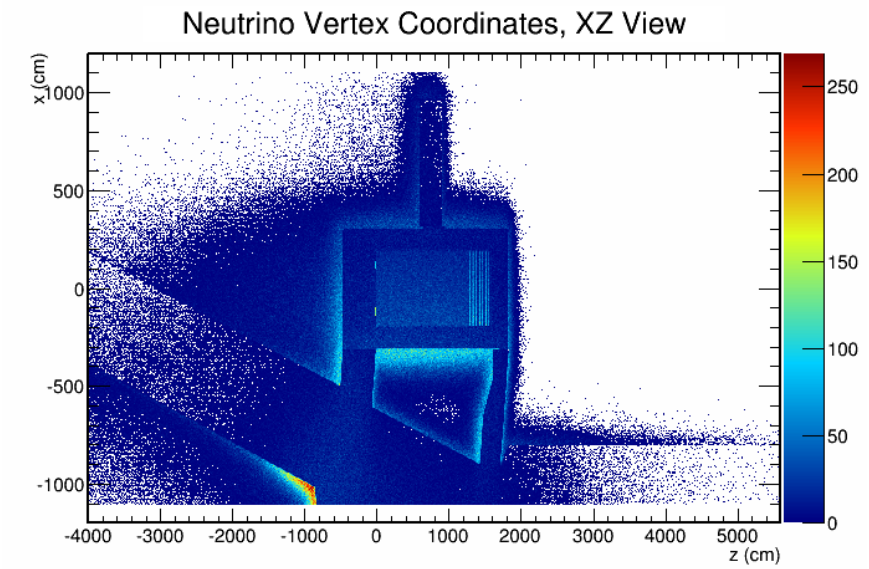


Figure 4.3: Distribution of neutrino interaction vertices in the XZ view at the ND hall. Several features of the geometry model can be seen. The detector sits between ± 200 cm in X and $0 < Z < 1600$ cm. The muon catcher steel planes are easily identified at the back of the detector as the line segments with slightly enhanced interaction rates. The two small blips of greater activity at ± 100 cm in X and 0 cm in Z are the bookend keeping the detector in place. The wide, sloped region extending from 0 on the Z axis to 0 on the X axis is an excavated pathway.

Geant4 also handles the next piece of the simulation, the conversion between energy deposition and

light yield. For low energies, the rate of light produced in liquid scintillator is proportional to the energy deposition, but the light yield begins to quench for particles with high enough energies. The Birks-Chou Law, encapsulated within a Geant4 module, is used to model the relationship between scintillator light yield rate, $\frac{dL}{dx}$, and particle energy deposition rate, $\frac{dE}{dx}$ [76].

$$\frac{dL}{dx} = L_0 \frac{\frac{dE}{dx}}{1 + k_B \frac{dE}{dx} + k_C \left(\frac{dE}{dx}\right)^2} \tag{4.1}$$

Above, the constants L_0, k_B, and k_C are dependent on the scintillator material and had to be estimated for NOνA as no measurement existed for the particular materials used in the experiment. A study was performed comparing the energy deposition at the end of proton tracks in the ND for both data and MC to find parameters that would generate agreement between the two [77]. The results of the study were $k_B = 0.04\,\text{cm/MeV}$ and $k_C = -0.0005\,(\text{cm/MeV})^2$. The output from this simulation step is a set of energy depositions now encoded as light yield called FLSHits, or fiber in liquid scintillator hits.

4.2.2 Photon Propagation

The next part of the simulation handles the capture of the scintillation photons, their transport through the WLS fiber to the electronics, and their conversion into photoelectrons after collection by the APD. The detector is treated as uniform for this section, with the assumption that any individual differences between cells would be removed by the downstream calibration. Individual photons are not traced from emission to the electronics due to its expensive computation time, despite the ability of Geant4 to perform this simulation. Rather, a template that parametrizes the photon transport are created and individual results are drawn from this.

The first template is a collection rate of photons by the WLS fibers as a function of time from scintillator emission and distance along the cell length from scintillator emission. The template is constructed using a ray tracing algorithm that considers a photon to be captured when it intersects with a fiber, and is shown in figure 4.4. The mean number of photons per energy deposited is a tunable parameter that is constrained using cosmic ray data. The scintillator emission spectrum and cell wall reflectivity are modeled as a function of wavelength based on data from NOνA's quality control database. The reflectivity is also shown in figure

4.4.

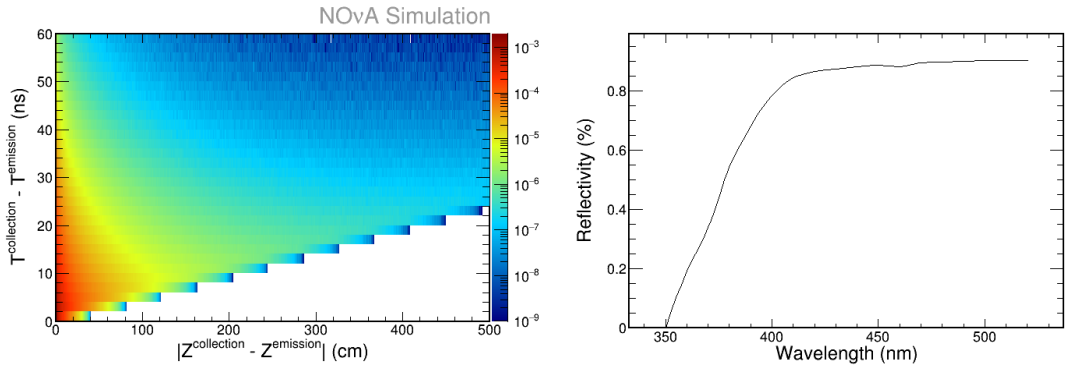


Figure 4.4: Distributions used during the simulation of photon transport. Left: collection rate as a function of time and distance travelled from scintillator emission. Right: cell wall reflectivity as a function of wavelength.

Next, the photons are transported through the fiber to the electronics. Since the fiber is looped and both ends are read out, half of the photons are transported to each end. The number of photons that reach the fiber ends is attenuated by applying an attenuation curve derived from quality control tests as a function of distance from the fiber end. Another template is constructed using a ray tracer that models transport time as a function of distance. The actual transport times are drawn from this template instead of handling the photon capture angles individually.

At this point in the simulation, there is a number of photons that reach the APD as a function of time. This is converted to photoelectrons (PEs) assuming a flat quantum efficiency of 0.85 for every APD. The exact number of PEs collected is drawn from a log-normal distribution with a mean of the expected number of PEs to account for the repeated random charge avalanche process within the APD. The final output of the photon transport simulation is a set of objects that encodes the number of PEs as a function of time, called PhotonSignals.

4.2.3 Electronic Readout

The last piece of the detector simulation models the electronics readout. This section starts with the number of PEs collected by the APDs and simulates the effects of the electronics to generate the raw signals equivalent to data. Before the PE signal from the APD is input to the circuitry simulation, an excess noise

factor is applied. The excess noise PE signal is drawn from a log-normal distribution, the limit of the product of many random variables, motivated by the fact that the actual PE signal comes from many random charge multiplications. Next, three chips on the FEBs are simulated, an application-specific integrated circuit (ASIC), an analong to digital converter (ADC), and a field-programmable gate array (FPGA).

The ASIC pre-amplifies and shapes the PE signal. A data driven effect is included at this point called APD sag, or the tendency when one APD channel observes a large hit for other channel baselines to slightly drop. The recovery from the drop can appear like real activity causing all APD pixels to register hits. Another form of noise is modeled at this point as the sum of two Gaussian Markov chains. This is done to simulate the high correlation of noise traces between adjacent samples, and is tuned to best match pedestal detector scans.

The shaped pulse is next input to the ADC, which performs two functions. The signal is first converted into an integer value, up to a maximum threshold that would saturate the ADC. These values are then sampled on a periodic basis to create a simulated time series of channel outputs. The FPGA then searches for peaks above threshold using a dual correlated sampling trace, defined from $ADC_i - ADC_{i-3}$, where ADC_j is the jth output from the ADC. A final type of noise is considered for otherwise empty cells, drawn probabilistically from templates made from histograms of ADC differences from data noise hits.

The hits above threshold are the final output of the detector simulation and are recorded as a vector of objects called RawDigits. This is the lowest level object that has the same format for both data and MC; however, the additional truth information is always stored for the MC for validation and further studies.

4.2.4 Simulation Miscellany

In real data, readout did not always occur from every cell in the full detector. Full diblocks sometimes were missing due to original commissioning or repair work. Within individual diblocks, single channels can be masked off for too high noise levels or too low data rates. The diblock configuration for data was stored by run, and the channel masks were stored by subrun. To most accurately reflect the detector running conditions, the MC was generated matching the configurations in data, or weighting the MC configurations simulated by POT.

The FD MC files generated do not include any realistic oscillation probabilities; instead, three sets of files are produced. The first set assumes that there are no oscillations, called nonswap files. These files are consist of the same percentage of ν_μ and ν_e generated naturally in the NuMI beam. The simulation does not generate any ν_τ so these events are entirely absent from nonswap files. The second set of files, called fluxswap, completely oscillate all ν_μ to ν_e, and vice versa. Finally, the third set of files, called tauswap, completely oscillate all ν_μ and ν_e to ν_τ. All three files include the same rate of NC events. Oscillation weights as a function of true energy can be applied individually to each interaction channel at a later stage, which allows for realistic conditions that can also be tuned without generating entirely new MC.

4.3 Studying Simulation Files

The benefits of retaining all of the information described in the previous sections is that it allows for quicker studies and validation. Truth information can be used to test the idealized proof of concept of a new analysis technique, and studies involving alternative flux models but no detector effects can be run without the time and computation expensive process of running the full simulation and reconstruction. Two packages were developed for these purposes, FluxReader and MCCheckOut.

FluxReader is a package designed to read and make distributions from Dk2Nu flux files [78]. Like Dk2Nu files, FluxReader is experiment independent and can be used for any experiment that uses the NuMI beam line. Its main design is to create sets of spectra from the Dk2Nu files, where each spectrum in the set plots the same variable(s) but shows the distribution for a unique set of neutrino flavor, parent, applied cross section, and detector location. Users are able to configure nearly everything inside macros, thus eliminating the need for recompiling code.

FluxReader makes clever use of several existing capabilities within other packages to generate its sets of spectra. Every individual neutrino ray in the flux files is re-decayed to each detector configured in a particular FluxReader macro, so each flux file can populate spectra for both the NOνA ND and FD. The package is able to apply cross sections by reading in cross section splines from GENIE, since flux files by definition do not simulate detector interactions. The user is able to configure what cross sections are applied, with different spectra generated for each. In both of these cases, any additional histograms added to a particular

set are handled automatically alleviating the need for user generated repetition of error-prone code.

The strategy for evaluating beam related systematic uncertainties was not to generate entirely new MC with systematically shifted quantities. Instead, only shifted Dk2Nu files were created, and ratios of shifted to nominal spectra as a function of true neutrino energy were generated for each neutrino flavor (also split by sign) and provided as event weights for the fully simulated MC. These studies were performed using the FluxReader framework [79]. The specific systematics are discussed in detail in chapter 8.

While the systematic studies performed using FluxReader are the most directly related to the analysis in this dissertation, it has facilitated many other studies, from which a representative sample is briefly discussed below. The newest version of Fluka was compared against the previous version used in the NOνA simulation, a collaboration internal study necessary to help trace major simulation differences that might arise between major MC productions [80]. A proposal to run the NuMI beam with various different horn currents was supported with studies utilizing FluxReader [81]. This test was eventually carried out, and the FluxReader distributions were used for the MC predicted energy spectra compared to the data [82]. Another study considered the effect of increasing the spacing between target fins to allow pions to more easily escape before scattering in a secondary target interaction [83].

MCCheckOut is a package designed for validation of various sections of the NOνA simulation. It creates distributions of relevant variables for each step of the simulation to search for changes between different MC productions. Figure 4.3 shown above is one of the standard plots created by MCCheckOut. Between major productions, modules are run that compare distributions for the MC truth, FLS hits, simulated POT, and RawDigits. MCCheckOut is typically run by a single person who pushes the results to an online webpage so that a wider audience can look for and comment on simulation differences.

Ideally, only expected changes from simulation improvements are seen, but in reality, unexpected and unintended differences sometimes appear. It is vitally important to find and correct these mistakes before valuable computation time is wasted generating erroneous MC files. Several problems were discovered and fixed while this author was in charge of running simulation validation, two of which are briefly mentioned below. At one point, the window over which neutrino rays were sampled for FD interactions was shortened without taking into account the 3° angle between the beam direction and detector. This led to

an absence of events in a section of the detector, as shown in figure 4.5 [84]. A similar problem occurred later when a code change caused the location GENIE used to start generating neutrino interactions at the ND to be pushed too far forward in Z. The result was an absence of events at the front of the detector, and specifically in the steel bookend, as shown in figure 4.6 [85]. MCCheckOut is clearly a valuable and necessary package for the MC simulation production.

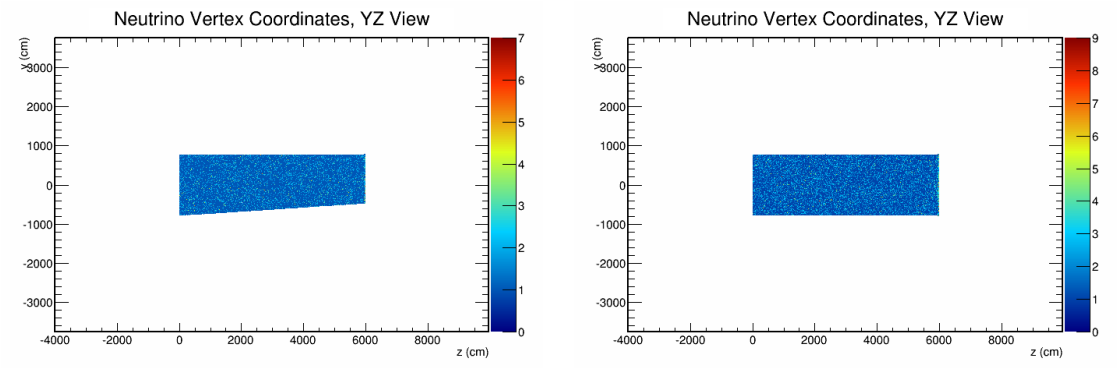


Figure 4.5: Distribution of neutrino interaction vertices in the YZ view at the FD. The left plot shows a region of detector with no events due to a shortened flux window. The right plot shows the corrected distribution after the flux window was lengthened.

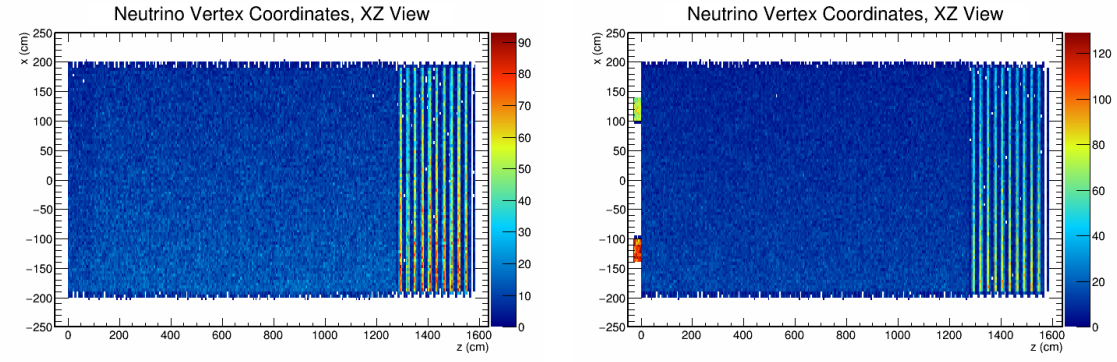


Figure 4.6: Distribution of neutrino interaction vertices in the XZ view at the ND. Rock events were removed in these distributions. There are no interaction in the steel bookend in the left plot. On the right, the simulation was corrected and events can be seen within the bookend.

5

Event Reconstruction

The most basic form of raw data collected is a vector of hits above threshold. The MC simulation described in chapter 4 also outputs events in this format. However, these objects by themselves are not very useful; instead, a certain level of reconstruction is required before real physics can be studied. The first major step in this process is to apply calibration so that the hits can be translated into a set of energy depositions, consistent throughout and across both detectors. Any number of algorithms can then be applied to create new objects or search for features, including tracks, the event vertex, or particle identifiers (PIDs). This chapter describes the calibration and elements of the reconstruction chain relevant to the NC disappearance analysis.

5.1 Calibration

The purpose of calibration is to ensure a uniform detector response throughout each and across both detectors. This is done in two major steps, a relative and absolute calibration. The relative calibration ac-

counts for threshold effects and attenuation across a single cell. It is designed to create a uniform response throughout a cell and across a single detector. The absolute calibration creates a scale factor for each detector to convert the calibrated PE scale from the relative calibration into an energy unit. This section follows the SA notes in reference [86].

5.1.1 Relative Calibration

The relative calibration is designed to convert the PE signal output from the electronics into a calibrated unit, such that two equal signals from any two detector locations mean equal true energy deposited. This part of the calibration accounts for threshold effects and attenuation in the WLS fibers, outputting a corrected PE value, the PECorr.

Hits from through-going cosmic ray muons, or muons that enter and exit the detector without stopping, are used for the relative calibration. The WindowTrack algorithm, a fast algorithm that fits straight lines through hits [87], is used to produce 3D tracks from the cosmic ray events, and only those with a successful reconstruction are used. Within these events, only tricell hits are used for the calibration procedure. A tricell hit is defined as a hit in cell i within a given plane that also has hits in cells $i+1$ and $i-1$. Under special circumstances (low statistics, too many dead neighboring cells) different hits are used for a particular cell. The path length of the cosmic traveling through the cell and the distance from readout are calculated for all of the selected hits in a cell. The distance from readout is labeled W, an alias for either X or Y, such that $W = 0$ is the center of a cell and positive values of W are closer to the readout. From this information, individual histograms of the average PE/cm vs W are constructed for each cell. The relative calibration procedures apply corrections to these histograms.

The first effect handled by the relative calibration is of threshold and shielding. Thresholds refer to the issue that an energy deposition may not register for a hit at all, as opposed to simply being attenuated, if there are not enough photons that reach the APD. Shielding refers to the tendency of the detector mass to alter the average signal of a minimum ionizing particle, or MIP, as a function of distance to the readout. Both of these effects would bias the set of hits used by the calibration by preferably selecting hits with greater numbers of photons, in turn underestimating the true energy deposited in the cell. To account for

this, a correction factor is applied for each cell,

$$T = \frac{PE}{\lambda} \cdot \frac{E_{True}}{E_{MIP}} \tag{5.1}$$

where T is correction factor, PE is the number of simulated photons that the electronics register, λ is the number of photons that would be seen without fluctuations, E_{True} is the true energy deposited in the scintillator, and E_{MIP} is the energy that would be deposited based only on the particle path length through the cell. The ratio on the left accounts for the threshold correction since λ is only dependent on the simulated threshold level, and the ratio on the right accounts for the shielding correction since E_{MIP} is only depenedent on the path length. Two dimensional histograms of the correction factor as a function of the cell number and distance from the readout are made for each detector and view, then fit with a polynomial to remove noise. These histograms are then used to correct the corresponding data and MC. Figure 5.1 shows examples of this correction factor used for the FD.

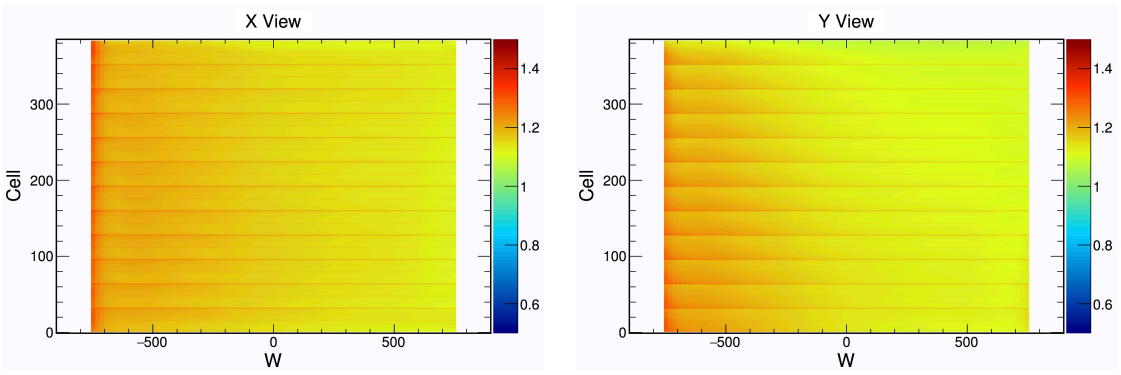


Figure 5.1: Correction factor for threshold and shielding effects at the FD as a function of cell number and distance from electronic readout. Cells in the X view are shown on the left, Y view on the right.

The next part of the relative calibration is the general form of the attenuation correction. Tricell hits are grouped by cell and fit to a double exponential to consider both short and long path light. For data, individual fits are performed for every single cell. In MC, a fit is performed on the group of all cells in a particular view and at the same location in a plane due to much lower simulated cosmic event statistics.

The double exponential has the form

$$y = C + A\left(\exp\left(\frac{W}{X}\right) + \exp\left(-\frac{L+W}{X}\right)\right) \tag{5.2}$$

where C, A, and X are the free parameters, and L is the full length of the cell. X is the cell attenuation length. The fit excludes hits in the regions closest and furthest to the readout. It includes the range $[-750, 750]$ at the FD, $[-150, 150]$ for the fully active region of the ND and for Y view muon catcher cells, and $[-150, 50]$ for X view muon catcher cells, with all numbers in cm. These central regions are chosen to exclude the significant rolloff regions at the end of the cells, which are handled differently.

The final step in the relative calibration handles the attenuation correction at the ends of the cells and the residuals from the fit above in the central part of the cells. First, both of these regions are fit with a single LOWESS curve using a tricube weight,

$$w_i = \begin{cases} \left(1 - \left|\frac{W - W_i}{\sigma}\right|^3\right)^3 & \text{for } |W - W_i| < \sigma \\ 0 & \text{for } |W - W_i| \geq \sigma \end{cases} \tag{5.3}$$

where W is a local point on the curve, W_i is the ith neighbor of the local point, w_i is the weight on W_i, and $\sigma = 30$ cm is the range of neighbors that affect the value of W. W is then the weighted mean of W_i. Next, a simple line of the form $y = mW + c$ is fit to a collection of the neighbor points around a point, W, to give the corrected response, y. The full LOWESS curve is approximated by linear interpolation between 20 points calculated using this procedure. The full attenuation calibration combines the results of the linear interpolation with the result of the double exponential above. Figure 5.2 shows examples of the fully fit attenuation curves used in FD data calibration.

The relative calibration combines the results of all three steps to create a uniform detector response as a function of distance to readout. Figure 5.3 shows the MC before and after calibration is applied to assess the performance of the full procedure. The regions with deviations at the end of the cells ultimately are not used for the analysis.

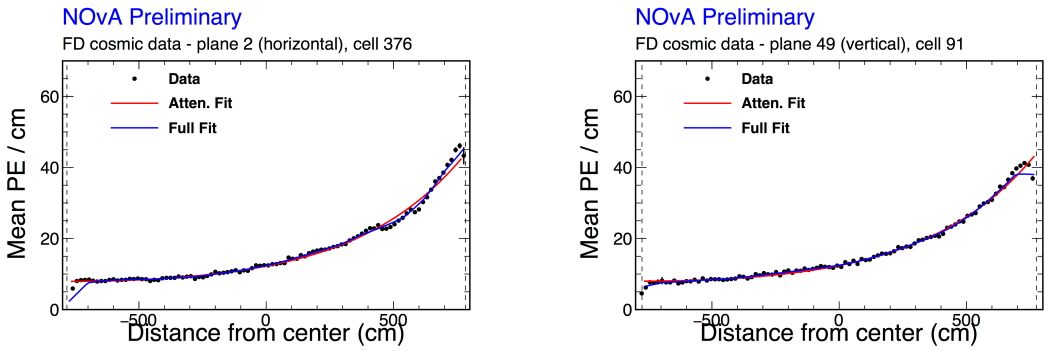


Figure 5.2: Fits to the average detector response for a cell in the FD. The red curves show just the result from the double exponential fit, and the blue curves show the full fit that also includes the LOWESS correction. The left plot is an example of a horizontal view cell, the right plot is an example of a vertical view cell.

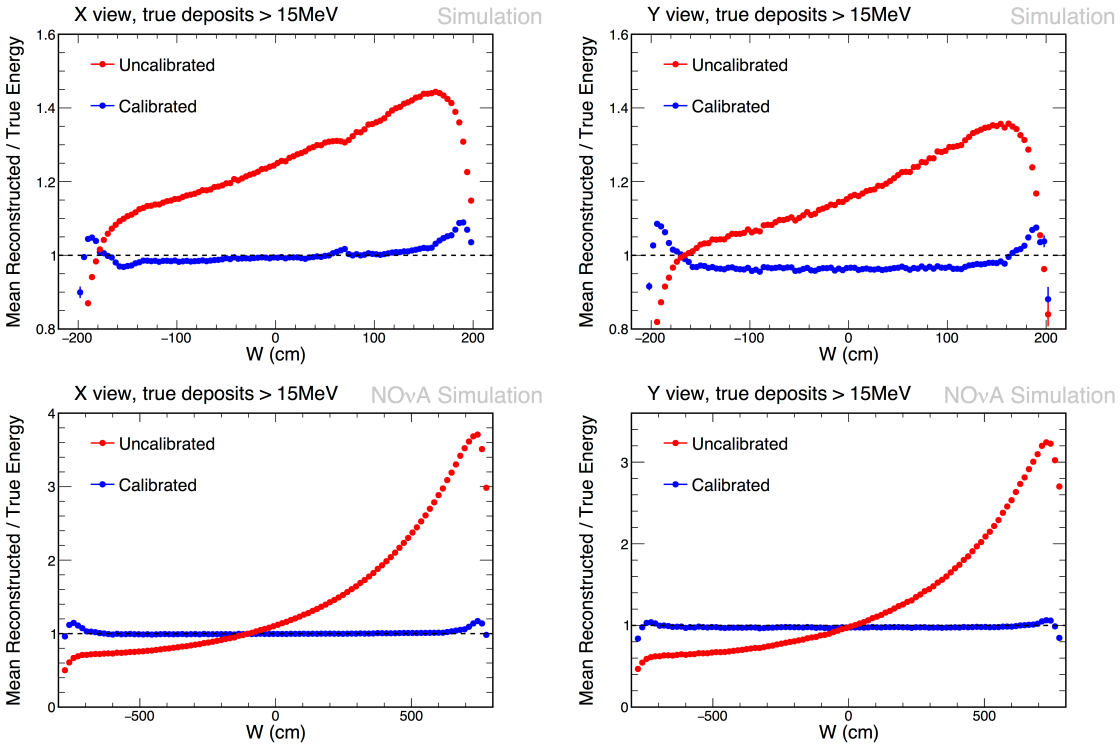


Figure 5.3: Ratios of mean reconstructed to true energy as a function of distance to readout to assess the performance of the relative calibration. The red points show the ratios before calibration; the blue points after. The left column shows calibration of the X view cells, the right column shows the Y view cells, the top row shows the ND relative calibration results, and the bottom shows the FD results.

5.1.2 Absolute Calibration

The absolute calibration is performed after the relative calibration to convert the corrected PE signal into an energy value. Much like the relative calibration, the absolute calibration uses a sample of tricell hits from cosmic ray muons. However, in this case the energy deposited in the cell must be known in order to convert the PE signal into an energy. The energy is known either from external sources [88] or the Bethe-Block equation [35],

$$\left\langle -\frac{dE}{dx} \right\rangle = Kz^2 \frac{Z}{A} \frac{1}{\beta^2} \left[\frac{1}{2} \ln \frac{2m_e c^2 \beta^2 \gamma^2 W_{max}}{I^2} - \beta^2 - \frac{\delta(\beta\gamma)}{2} \right], \tag{5.4}$$

where $K \equiv 4\pi N_A r_e^2 m_e c^2$, r_e is the classical electron radius, $m_e c^2$ is the electron rest energy, z is the charge number of the incident particle, Z and A are the atomic number and mass of the stopping material, β and γ are calculated for the incident particle, W_{max} is the maximum energy transfer to an electron in a single collision, I is the mean excitation energy of electrons in the stopping material, and $\delta(\beta\gamma)$ is the density effect correction to ionization energy loss in the stopping material. With known energy deposition and number of PEs, determining the calibration energy scale is a simple procedure.

The particular energy used for the absolute calibration is the minimum energy deposition of muons through liquid scintillator. For NOνA, the liquid scintillator is well approximated as chains of polyethylene, or $(C_2H_4)_n$, for which muons have a minimum average energy loss of 2.079 MeV cm^2/g [88]. The density of the scintillator was measured to be 0.8617 g/cm^3 [89], giving an average energy loss per unit length of 1.792 MeV/cm. From simulation, it was determined that this minimum energy deposition occurs between $100 - 200$ cm from the end of a muon track, with only a 1.8% variation throughout that region. Figure 5.4 shows the distributions of the relative calibration corrected response as a function of distance from the muon track end for FD data and MC and the true energy deposition for the MC.

The absolute calibration energy scale is determined from distributions of tricell hits from stopping muons that occur between 100 and 200 cm from the end of the muon track. These hits are used to make one dimensional muon energy unit (MEU) distributions of the relative calibration corrected detector response for data and MC called MEU_{reco}, and the true energy deposition for MC called $\text{MEU}_{\text{truth}}$. The

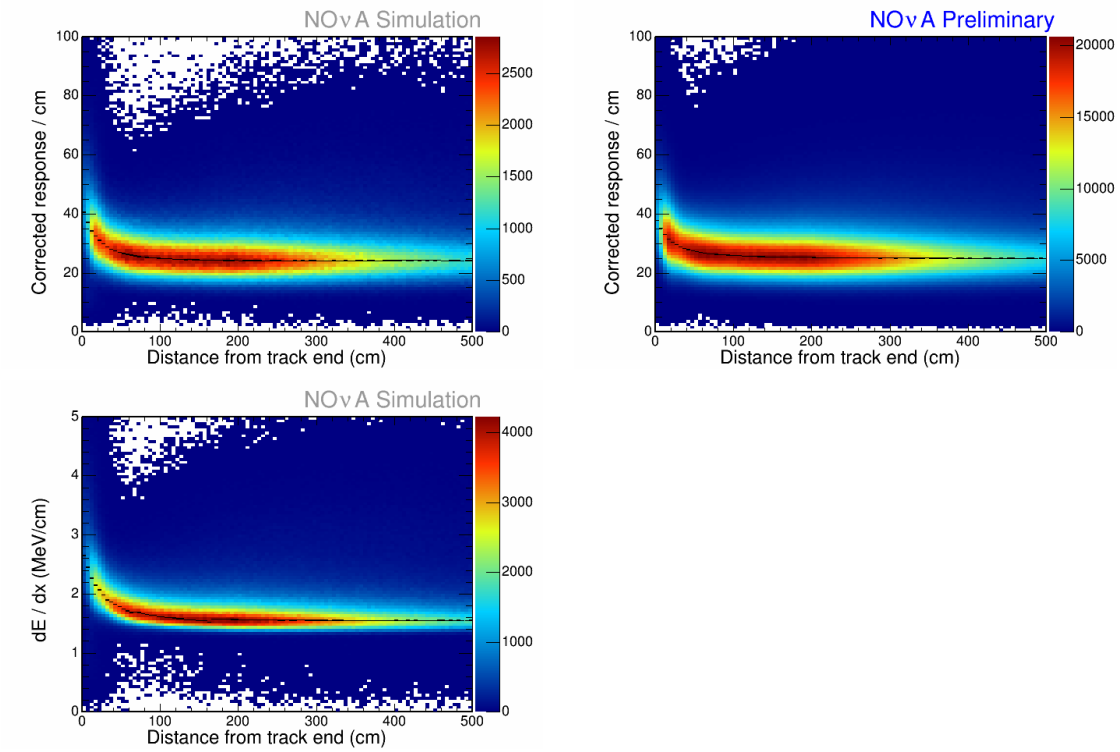


Figure 5.4: Distribution of tricell hits from cosmic ray muon tracks at the FD as a function of distance to the end of the muon track. The top row shows the response in PECorr for both data and MC. The bottom left plot shows the true energy deposited for the MC.

calorimetric energy scale is then taken as the mean of the MEU_{truth} distribution over the mean of the MEU_{reco} distribution. Figure 5.5 shows the distributions of the corrected detector response before the absolute calibration is applied, and the energy distributions after the absolute calibration is applied.

At this point the calibration procedure is complete. The calibration constants output by the relative and absolute calibrations are stored in a database so that the raw PE signal in any RawDigit object can be converted into an energy value.

5.2 Reconstruction Chain

The first step in the NC disappearance analysis is the selection of a relatively pure sample of NC events. However, the relatively raw calibrated hits are a far cry from the complete events necessary for the analysis. A set of reconstruction algorithms are thus applied to take individual hits and group them into more complete objects, often providing extra information along the way. The reconstruction procedures be-

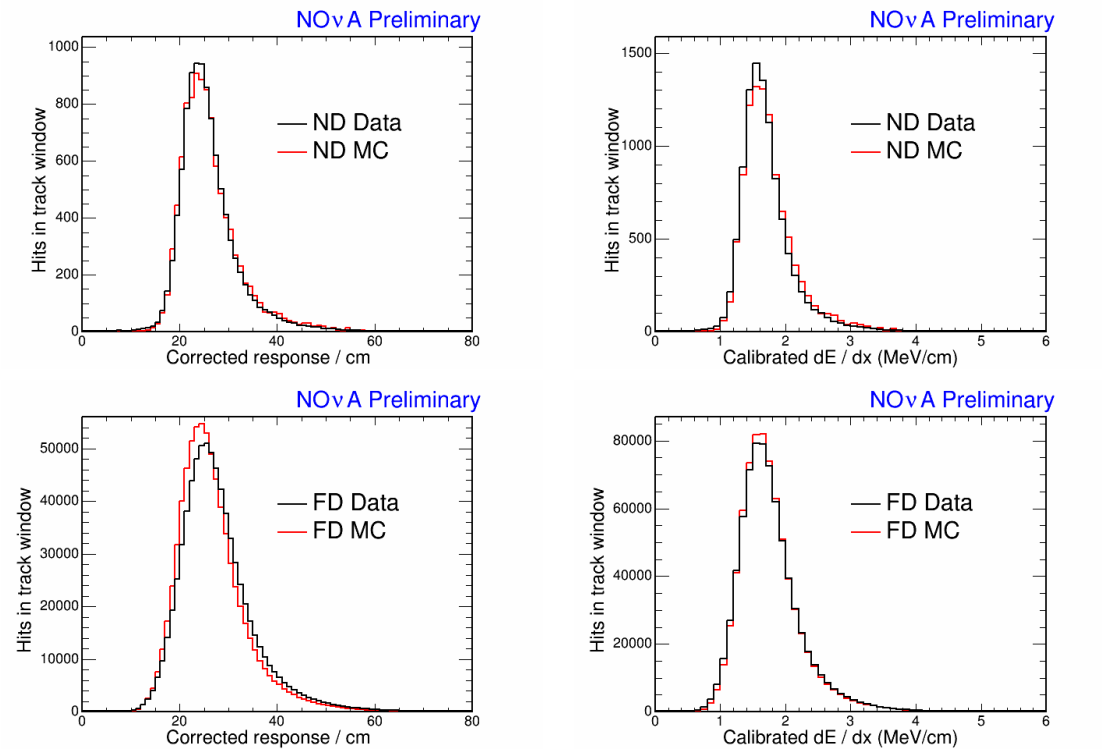


Figure 5.5: Distributions of the tricell hits from cosmic ray muons $100-200$ cm from the end of the track. The left column shows the relative calibration corrected detector response before the absolute calibration is applied. The right column shows the absolute calibrated energy deposition. The top row shows results for the ND; the bottom shows the FD.

gin with the basic grouping of hits in spacetime, and become as complex as applying machine learning algorithms to separate events into different components. The NC disappearance analysis did not have an independent reconstruction chain from the two main NOνA analyses. Rather, the information used for the NC selection discussed later in chapter 6 was a mixture from the ν_e appearance and ν_μ disappearance analyses. The reconstruction chains for these analyses were developed with different philosophies; the ν_e reconstruction chain was designed to identify shower like objects and is focused on clustering particles, white the ν_μ reconstruction chain was designed to identify long tracks. The remainder of this section discusses the reconstruction components from these reconstruction chains that are relevant to the analysis described in this dissertation.

5.2.1 Event Slicing

The first part step in reconstruction is the clustering of hits into separate events, or slices. This is done based on the Density-Based Spatial Clustering of Application with Noise (DBSCAN) algorithm [90]. The algorithm as applied to NOνA is described in detail in reference [91]; the main points are summarized here.

The DBSCAN algorithm uses a score function to compute a distance between neighboring points. A threshold is set to determine whether two neighbors are close. Points that have more than a set number of neighbors within a that distance are labeled *core points*, and the close neighbors of the core points that are not themselves considered core points are labeled *border points*. Clusters are formed by iterating over the hits in an event window, computing whether a point is a core point, then expanding to the neighbors of the core point, and moving to the next cluster when the current one is entirely bounded by border points. At the end of this procedure, slices with more than 3 hits in each view are made into a 'physics' slice, and any hits not assigned to a cluster are placed into a 'noise' slice.

The score function used between two hits in the slicing algorithm is

$$\epsilon = \left(\frac{\Delta T - |\Delta\vec{r}/c|}{T_{res}}\right)^2 + \left(\frac{\Delta Z}{D_{pen}}\right)^2 + \left(\frac{\Delta XY}{D_{pen}}\right)^2 + \left(\frac{PE_{pen}}{PE}\right)^5, \tag{5.5}$$

T_{res} is the timing resolution of the hits added in quadrature, D_{pen} is a distance penalty, PE is the number of photoelectrons in both hits added in quadrature, and PE_{pen} is a penalty on the number of photoelectrons. Hits that are in the same view vs opposite views are handled slightly differently. For hits in the same view, $\Delta\vec{r}$ is calculated in two dimensions. For hits in opposite views, $\Delta\vec{r}$ is one dimensional, ΔXY is 0, and D_{pen} is replaced with a separate, smaller, opposite view plane penalty. T_{res} was set individually for each hit during the timing calibration, discussed in reference [86]. The exponent on the PE term was set at 5 as the PE spectrum for noise falls as $PE^{-2.5}$.

The free parameters and their values are summarized in Table 5.1. The parameters were tuned and the performance of the algorithm was measured based on two metrics, completeness, the percentage of energy deposited in the scintillator from a physics interaction included in a slice, and purity, the percentage of energy in a slice that came from a single physics interaction. Values were tuned separately for the ND

and FD. Notably, PE_{pen} was set to 0 for both detectors, effectively removing this term from the distance calculation.

Parameter	ND	FD
ϵ, minimum score required to be considered close neighbors	5.0	2.0
Minimum close neighbors required to be considered core point	4	4
D_{pen}, distance penalty	75.0	100.0
Opposite view plane penalty	8	4
PE_{pen}, penalty on the PE total	0	0

Table 5.1: A summary of the free parameters used in the slicing algorithm, and the tuned values used for each detector. Table adapted from [91].

An example event is shown in figure 5.6. Figure 5.7 shows the same event after slicer was applied, and figure 5.8 zooms in on the neutrino candidate event.

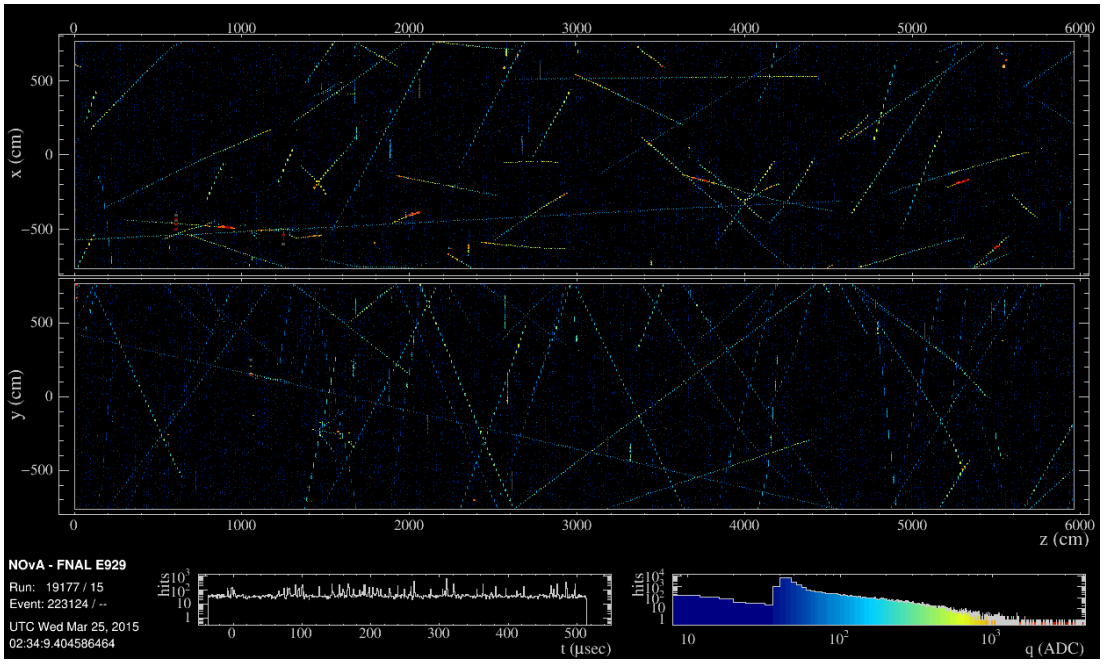


Figure 5.6: An example NOνA event in the FD. The upper panel shows the XZ view and the bottom panel shows the YZ view. The color of the hits relates to the charge deposited.

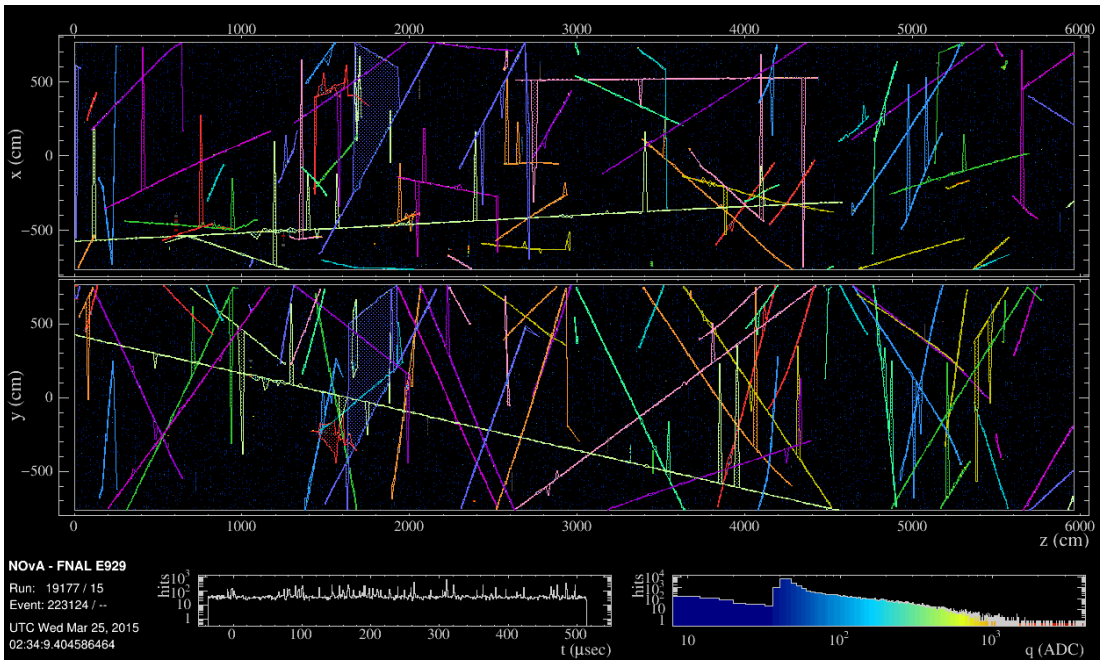


Figure 5.7: The same NOνA event as in figure 5.6 after slicer was applied. Colors only denote different slices. 65 different slices were made for this event, with 1 candidate neutrino interaction located at approximately $(500, -300, 1500)$ cm.

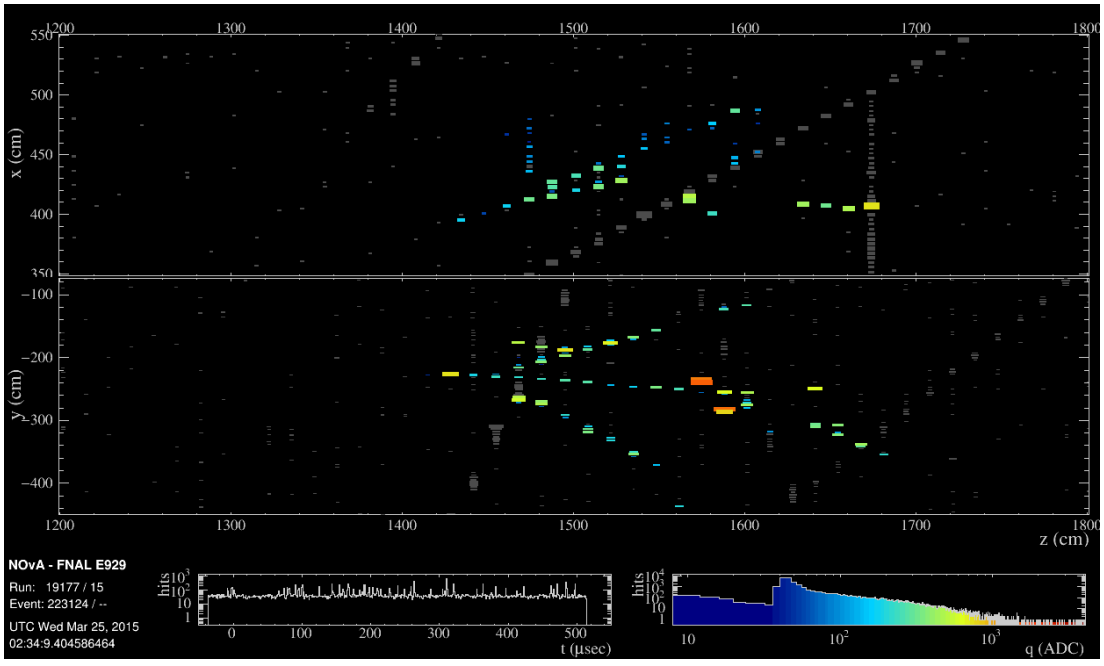


Figure 5.8: Zoomed view of the neutrino candidate slice from figures 5.6 and 5.7. Activity in other slices was grayed out.

5.2.2 Hough Transform Line Finding

The ν_e reconstruction chain begins by searching for line like features in the slice using a modified Hough Transform [92]. The algorithm as applied to NOνA is described in references [91, 93]; it is summarized here.

The algorithm begins by constructing a Hough map with the parameters of lines between pairs of points. A separate map is constructed for the XZ and YZ views. The lines are parametrized in polar coordinates, (ρ, θ), with ρ the perpendicular distance from the origin and θ the angle away from the X axis. For each line calculated, the Hough map is filled with a Gaussian smeared weight based on the line parameters,

$$w = \exp\left(-\frac{(\rho-\rho_0)^2}{2\sigma_\rho^2}\right)\exp\left(-\frac{(\theta-\theta_0)^2}{2\sigma_\theta^2}\right) \tag{5.6}$$

$$\sigma_\rho = \frac{3}{\sqrt{12}}[\text{cm}] \tag{5.7}$$

$$\sigma_\theta = \frac{3\,\text{cm}}{d\sqrt{6}} \tag{5.8}$$

where d is the distance between the two points in cm. There are two restrictions on the pairs of points that are evaluated. Any pairs of points that are farther than $\sqrt{15000}$ cm apart are not evaluated together to limit the total number of contributions and limit contributions with very low uncertainties in θ. Pairs of points that have the same X or Y coordinate are not evaluated together unless they are farther than $\sqrt{15000}/4$ cm apart to avoid the creation of spurious horizontal or vertical lines. Finally, the Hough map is smoothed using a Gaussian smoothing weight over non-zero bins.

The algorithm then iterates over peaks in the Hough map to find and label valid lines. The highest peak is considered first. A 7×7 square of bins centered on the peak is used to find an average (ρ, θ), with the bin values weighted by the inverse distance from the central bin. The resultant parameters are recorded along with the size of the peak. The Hough map is then regenerated without hits within 6 cm of the line, with the exception of the most upstream and downstream hits. This process is then repeated to find the next line. For all iterations after the first, the current line is tested against previous lines. If the new line is within 15 cm in ρ and 0.02 rad in θ, the line is not added to the list of valid Hough lines, but the hits associated with the line are still removed. The algorithm terminates when there are no remaining peaks, or

when there are 10 Hough lines, whichever comes soonest. An example of the algorithm generating Hough lines on a MC event is shown in figure 5.9.

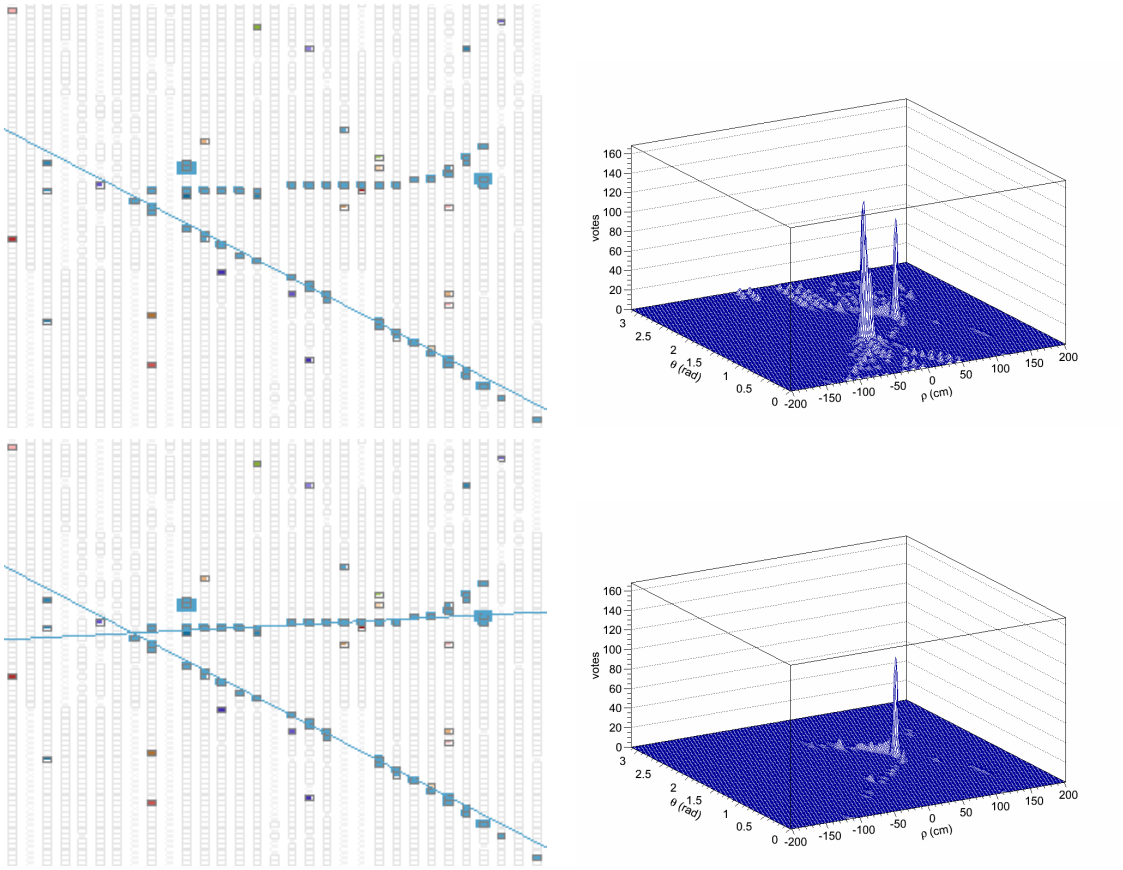


Figure 5.9: An example of the Hough transform algorithm generating lines from an event. The top row, left panel shows an example MC event from which a Hough map is generated and shown on the right. The strongest peak is used to generate the Hough line shown on the event in the top left panel. The Hough map is then regenerated without hits associated with the first line as shown in the bottom right panel. From this map, the strongest peak is used to generate the second Hough line, shown on the bottom left panel.

Figure 5.10 shows the Hough lines that were created for the neutrino candidate slice shown from figure 5.8.

5.2.3 Elastic Arms Vertex Reconstruction

The output of Hough Transform is the input to the Elastic Arms algorithm that is used to reconstruct the original neutrino interaction vertex. Elastic arms, also called deformable templates, is described in references [94, 95, 96, 97] and is designed to find outgoing tracks with a known vertex. The algorithm as

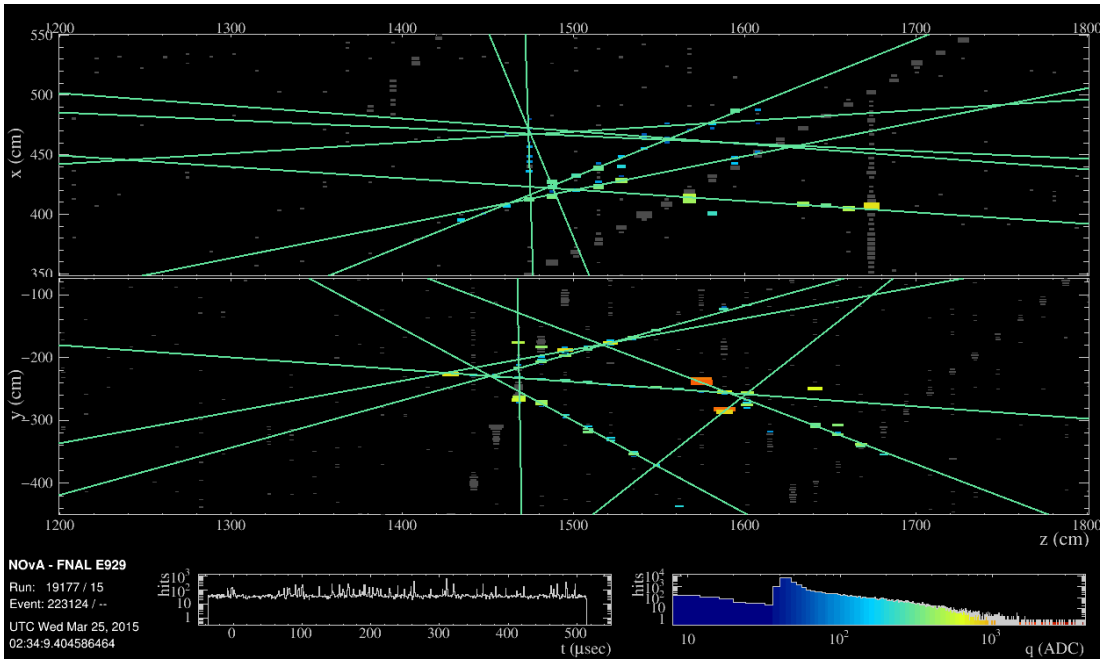


Figure 5.10: The neutrino candidate slice from figure 5.8 with Hough lines drawn.

applied to NOνA was modified to account for the unknown vertex, is fully described in reference [98], and summarized here.

The Elastic Arms algorithm attempts to fit a set of arms, or tracks that originate at a single point, to the hits in the slice. Each arm, a, is parametrically described in spherical coordinates according to

$$\begin{aligned} x(s) &= x_0 + s \sin\theta_a \cos\phi_a \\ y(s) &= y_0 + s \sin\theta_a \sin\phi_a \\ z(s) &= z_0 + s \cos\theta_a \end{aligned} \tag{5.9}$$

where (x_0, y_0, z_0) is the neutrino event vertex, s is the distance from the vertex, and (θ_a, ϕ_a) are the standard spherical angles. The algorithm fits for values of $(x_0, y_0, z_0, \vec{\theta}, \vec{\phi})$ by minimizing an energy cost function,

$$E = \sum_{i=1}^{N}\sum_{a=1}^{M} V_{ia} M_{ia} + \lambda \sum_{i=1}^{N} \left(\sum_{a=1}^{M} V_{ia} - 1 \right)^2 + \frac{2}{\lambda_v} \sum_{a=1}^{M} D_a, \tag{5.10}$$

for an event with N hits and M arms, where M_{ia} is the distance of hit i to arm a, V_{ia} is the strength of

association of hit i to arm a, D_a is a penalty on the distance between the vertex and the first hit on arm a, and λ and λ_v control the penalty of the second and third terms, respectively. M_{ia} is specifically computed as the perpendicular distance to the arm, unless the hit is in a backward direction with respect to the arm.

$$M_{ia} = \begin{cases} \left(\frac{d_{ia}^{perp}}{\sigma_i}\right)^2 & \text{Standard} \\ \left(\frac{d_i^{vtx}}{\sigma_i}\right)^2 & \text{Hit in backward direction, } \frac{d_i^{vtx}}{\sigma_i} \leq 1 \\ \left(\frac{d_i^{vtx}}{\sigma_i}\right)^4 & \text{Hit in backward direction, } \frac{d_i^{vtx}}{\sigma_i} > 1 \end{cases} \tag{5.11}$$

Above, σ_i is a measure of the spatial resolution of the cell, and set to $\sigma_i = 3/\sqrt{12}$ cm. V_{ia} is calculated from

$$V_{ia} = \frac{e^{-\beta M_{ia}}}{e^{-\beta\lambda} + \sum_{b=1}^{M} e^{-\beta M_{ib}}}, \tag{5.12}$$

where β can be interpreted as a temperature via $\beta = 1/T$. For any hit, the sum of the association to all arms is bounded between 0 and 1, $0 < \sum_{a=1}^{M} V_{ia} \leq 1$, with the difference from 1 representing the probability that a hit is considered noise by this algorithm. This gives λ the interpretation as the distance compared to M_{ia} that a hit has a 50% chance as being considered noise.

The penalty on the distance between the vertex and the first hit on the arm is not part of the general algorithm and is specific to NOνA. Since there are two views, D_a combines the distance to the first hit in both views, $D_a = d_a^{xz} + d_a^{yz}$. To calculate these distances, only hits with an above average association to an arm are used, as technically each hit has an association to every arm. This term in equation 5.10 was partially motivated by photons. The likelihood that a photon will travel a distance d before converting into an e^+e^- pair is proportional to $\exp(-d/\lambda_v)$, where $\lambda_v \approx (7/9)X_0$ and X_0 is the radiation length of the material the through which the photon is traveling. For the NOνA detectors, $X_0 \approx 39$ cm, or 6 planes. This photon conversion length leads to an energy function error of

$$\chi^2 = -2\ln\mathcal{L} = 2\frac{d}{\lambda_v}, \tag{5.13}$$

hence the form of the penalty on this term in equation 5.10.

Before searching for the best fit vertex and arm directions, the algorithm first considers the best seed from a set of vertex positions and directions. The number of arms is set as the larger of the number of quality Hough lines from the two views and generates a list of possible vertex seeds and arm direction seeds based on a sorted list of the hits in z and from the Hough lines themselves. The algorithm then searches for the seed that minimizes the energy cost function. Once this is chosen, the best fit vertex and directions are found by minimizing equation equation 5.10. Since this function is dependent on many parameters and can have many local minima, the algorithm starts with low β or high T, allowing hits to have moderate association with many arms. The temperature is gradually lowered while searching for the best fit parameters, pushing hits to their correct arms in the process.

The results of the fit for the vertex position are summarized in table 5.2. It was found that the final arm directions were not very reliable, even while the vertex position was. Therefore, the vertex is the only output used from the Elastic Arms algorithm. Figure 5.11 shows the vertex reconstructed for the neutrino candidate slice shown from figures 5.8 and 5.10.

Interaction	Δx		Δy		Δz		3D Resolution	
	Mean	FWHM	Mean	FWHM	Mean	FWHM	Mean	FWHM
ν_e CC	0.04	4.11	−0.02	4.44	0.73	7.38	10.65	17.20
ν_μ CC	0.11	4.43	−0.24	4.61	−0.61	8.16	11.56	17.08
NC	0.35	8.22	−0.77	7.84	0.47	9.50	28.72	30.80

Table 5.2: Difference between the true and reconstructed vertex for Elastic Arms vertex fits to FD MC events. All numbers shown are in cm. Table adapted from [37].

5.2.4 Fuzzy K-Means Clustering

The reconstructed vertex output by Elastic Arms is used as the input to a clustering algorithm that determines the number of prongs that are in the event, or tracks with an angular spread that originate at a single point, and the hits that belong to each prong. This algorithm is a possibilistic Fuzzy-K means clustering algorithm, described in detail in references [99, 37] and summarized here.

The clustering algorithm is a Fuzzy-K Means algorithm extended to deal with an unknown number of clusters and noise. The classic Fuzzy-K Means [100, 101] takes an event origin, or vertex, and assumes that

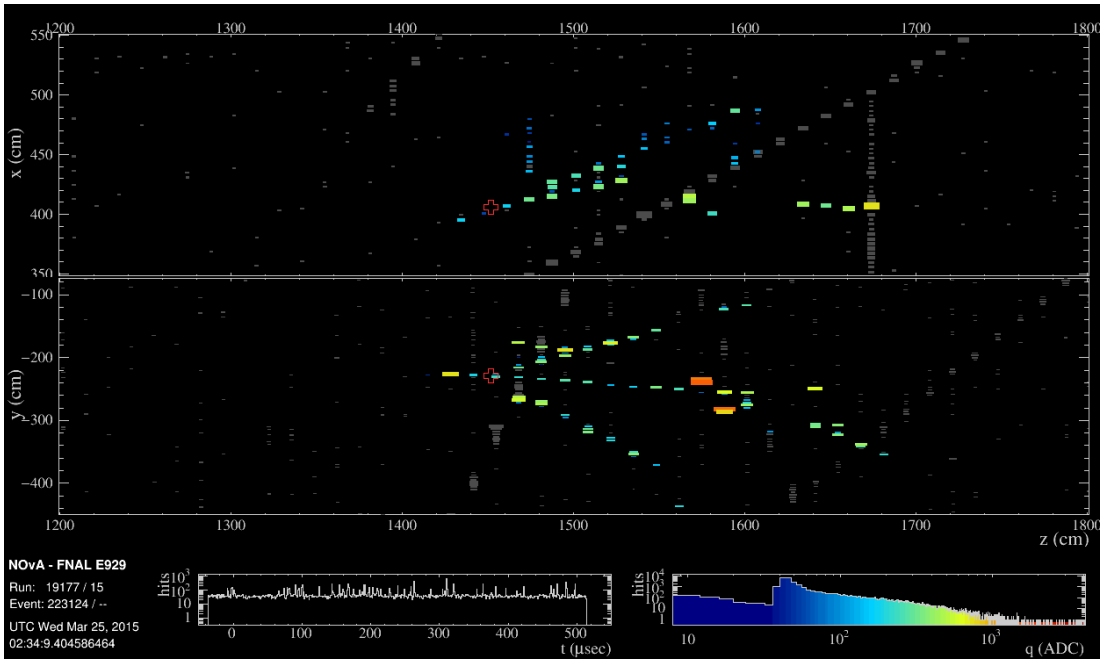


Figure 5.11: The neutrino candidate slice from figure 5.8 with a reconstructed vertex drawn as a red cross. The vertex was found to be slightly downstream of all activity at the intersection point of 3 Hough lines. In reality, the vertex was most likely at the most upstream hit and a π^0 traveled downstream briefly before decaying to two photons that both quickly converted. However, the vertex was still quite close to the most upstream hit and this was considered a good reconstruction.

all activity radiating outward with tracks appearing as peaks in angular space. It searches for the optimal placement of k clusters, or prongs, and associates activity, or hits, to them. Each hit is allowed to be associated to multiple prongs leading to the fuzzy nature of the prongs. However, they are subject to the constraint that all association strengths add to 1, meaning that every hit, including noise, will be associated to at least one prong. An extension of this algorithm [102] removes the association normalization constraint and allows for a changing number of clusters. This allows noise hits to remain unassociated with any cluster, changing the probabilistic nature of the prong association to a possibilistic association.

The algorithm first proceeds by seeding the prongs, calculating the association of the hits to the prongs, then updating the prong centers. The prongs are seeded in angular space based on an activity density matrix,

$$w_k = \sum_{i=1}^{n} e^{-\left(\frac{\theta_k - \theta_i}{\sigma_i}\right)^2} \tag{5.14}$$

with 360 equally spaced bins from $-\pi$ to π. The hit uncertainty, σ_i, is given by

$$\sigma_i = \frac{1.745}{d} + 0.0204 + 0.000173d, \quad d < 5\,\text{m}, \tag{5.15}$$

where d is the distance from the vertex to the hit. Equation 5.15 was modeled after the behavior of multiple scattering for muons but found to work well for all event types. This equation is used until $d = 5\,\text{m}$, at which point the uncertainty is held constant. Once the density matrix is populated, the prongs are seeded by the densest bins. With the prongs seeded, the algorithm proceeds with calculating the membership of each hit. First, the angular distance between the hit and a cluster center is calculated,

$$d_{ij} = \left(\frac{\theta_j - \theta_i}{\sigma_j}\right)^2, \tag{5.16}$$

where j refers to the j^{th} hit and i refers to the i^{th} cluster center. Next, the algorithm assigns membership to the cluster,

$$\mu_{ij} = e^{-\frac{m d_{ij}\sqrt{c}}{\beta}}, \tag{5.17}$$

where m is the degree of fuzziness for the clusters, c is the current number of clusters, and β is a measure of the expected spread in the clusters. Both m and β are tunable parameters. When $m = 0$, the algorithm reduces to a hard clustering model where each hit belongs to a separate cluster. In NOνA, m is set to 2 and β to 4. Finally, the cluster centers are updated.

$$\theta_i' = \theta_i + \frac{\sum_{j=1}^{n} \frac{\mu_{ij}^m}{\sigma_j^2}(\theta_j - \theta_i)}{\sum_{j=1}^{n} \frac{\mu_{ij}^m}{\sigma_j^2}} \tag{5.18}$$

This loop continues until none of the cluster centers update by more than a specified tolerance.

The next major component of the algorithm performs checks on the prongs and may modify the total number. If any cluster centers are found to have converged to the same value, the cluster centers are re-seeded. The lower peak in the density matrix that led to a duplicate prong center is replaced with the next lowest peak, and the above procedure is repeated. If there are no valid peaks to replace the original seed, the

total number of prongs is reduced by one but the algorithm does not repeat the above procedure. Once the cluster centers are stable and there are no duplicates, all the hits are tested for cluster membership. All hits that have $< 1\%$ membership to all clusters are used to populate a new density matrix and an extra cluster is added, seeded with the highest density bin from the new density matrix. All of the cluster centers are then updated through the above procedure, and this full process is repeated until all hits have at least 1% membership to a cluster or when there are 7 total prongs.

Lastly, the 2D prongs are matched across views to form 3D prongs. To do this, prongs from the XZ and YZ views that overlap for at least one plane are paired and their energy distributions are compared using a Kuiper test [103]. Specifically, the test considers the difference in the cumulative fractions of deposited energy between the two views as a function of distance through the prong, s.

$$\begin{aligned} D^{+}(s) &= E^{XZ}(s) - E^{YZ}(s) \\ D^{-}(s) &= E^{YZ}(s) - E^{XZ}(s) \end{aligned} \tag{5.19}$$

The final score of the Kuiper test maximizes each difference, and sums the result.

$$K = \max(D^{+}(s)) + \max(D^{-}(s)) \tag{5.20}$$

The prongs with the lowest score are matched, all other pairs that include either of the 2D prongs are removed, and the process continues until there are no remaining possible matches.

Figure 5.12 shows the prongs created for the neutrino slice from figure 5.8.

5.2.5 Kalman Tracking

The ν_μ reconstruction chain begins by using a Kalman filter to identify track like features and from hits in the slice. The algorithm outputs reconstructed objects called Kalman tracks. Reference [104] describes the algorithm in full detail; it is summarized here.

The KalmanTrack algorithm uses a Kalman filter to build up tracks from hits in a slice. The Kalman filter is used to estimate the state of a measurable system, x, given a set of k measurements, z of the system.

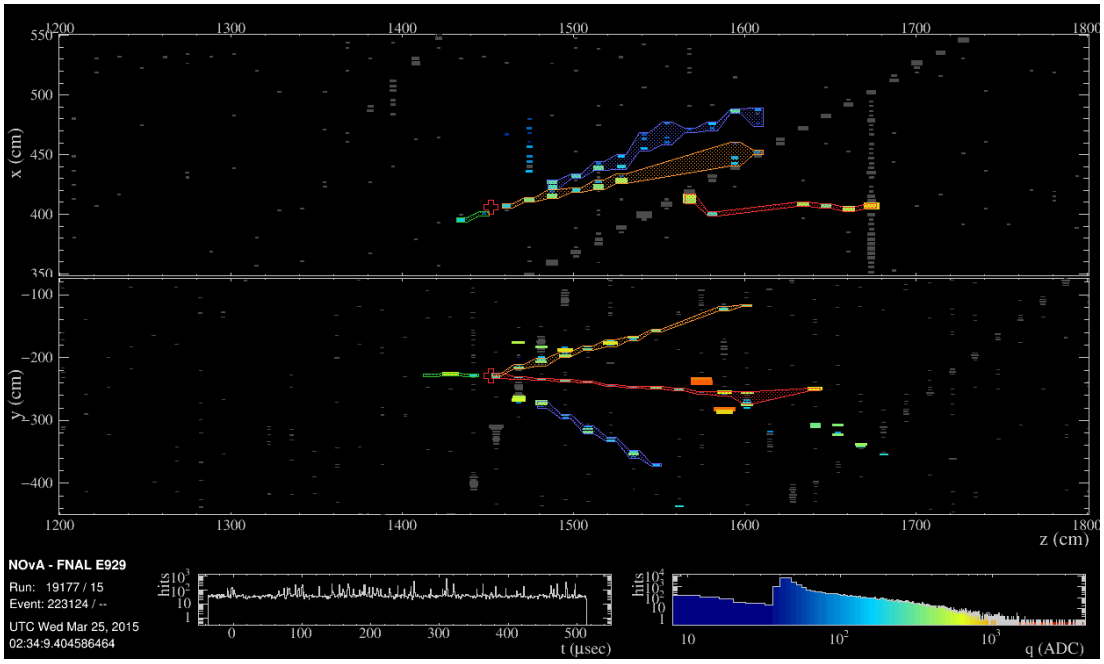


Figure 5.12: The neutrino candidate slice from figure 5.11 with prongs drawn.

It is assumed that the system is affected by a set of inputs, u, and system noise, w, while the measurements have an independent, inherent noise, v. Note the difference and independence between w, the error in the system process, and v, the error in the measurement process.

In the NOνA detectors, the Kalman filter prediction, x, is the position of a charged particle. The path of the particle is expected to be mostly straight since the algorithm was designed for tracking muons. The system noise, w, is taken as the result of small angle multiple scattering drawn from a normal distribution with width

$$\theta_0 = \frac{13.6\,\text{MeV}}{\beta p} z \sqrt{x/X_0}\,(1 + 0.038\ln(x/X_0))\,, \tag{5.21}$$

which is about 5.4 mrad for 2 GeV muons passing through one plane. Large angle scattering is handled separately at a later stage. The measurements of the particle position, z, come from the cell locations within the detector, and the measurement noise, v, is represented using the finite dimensions of the NOνA detector cell. Initially, the algorithm only reconstructs tracks in a single view; it merges tracks from both views in a later step.

The algorithm begins each iteration by constructing track seeds from pairs of hits separated by fewer

than 4 cells. To ensure that these track seeds are not created from hits belonging to multiple particles, the seeds are created and the tracking begins at the downstream end of the detector. For each seed, a line that connects the hits is calculated, and the track is then extrapolated to the next upstream plane based on the line output. Any hits that are within a given tolerance are added to the track, and the line parameters are updated with the new information. The algorithm repeats this process until no new hits are found within 4 planes of the previous hit. At this point, the track propagation reverses and proceeds downstream to find any hits that may have been missed from the original upstream propagation. Once the filter finishes tracking all of the track seeds, the track with the largest number of hits is saved, and the algorithm starts a new iteration on all of the remaining, unassociated hits. The algorithm terminates when no it can no longer reconstruct a track.

The final step of the tracking algorithm merges tracks from the XZ and YZ views. This process assumes that the start and stop endpoints for the opposing views are at roughly the same plane. For planes that have at least one plane of overlap, a metric for matching is calculated according to

$$S = \frac{\Delta\text{Start Plane} + \Delta\text{Stop Plane}}{\text{N Overlap Planes}}. \tag{5.22}$$

Aside from checking single pairs of tracks, the merging step also checks if two tracks in a single view can be joined together to improve the score for the three dimensional matching. This step is where large scattering can be handled. The set of tracks that produce the lowest score are matched as a 3D Kalman track, and the merging process is repeated on the remaining tracks. The algorithm terminates when it is unable to match any other 2D tracks.

Figure 5.13 shows the prongs created for the neutrino slice from figure 5.8.

5.3 Particle Identification Algorithms

Unlike the other reconstruction algorithms, Particle Identification (PID) algorithms attempt to directly the particle interaction type. PIDs generally take some amount of both raw and reconstructed information and output one or more probability distributions for event classification. The algorithms developed for the main NOνA analyses are discussed here, with most emphasis placed on those used for the NC

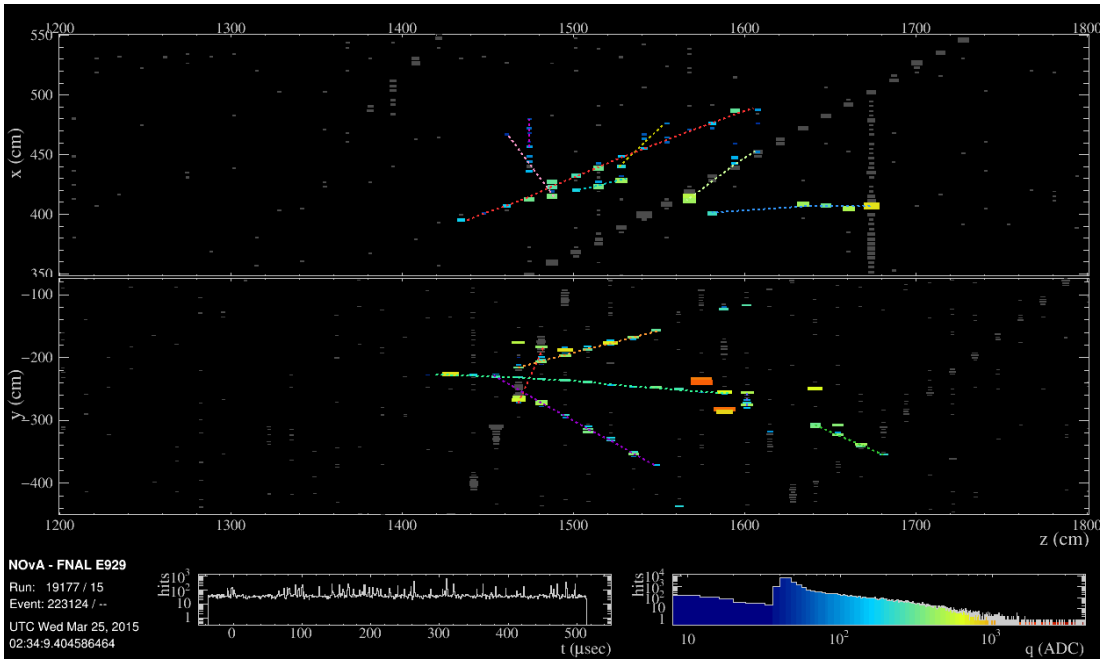


Figure 5.13: The neutrino candidate slice from figure 5.11 with tracks drawn. Tracks are not as accurate for NC events and are better suited for eliminating cosmics and ν_μ CCs.

disappearance analysis.

5.3.1 ν_e and ν_μ CC Specific PIDs

The ν_e appearance and ν_μ disappearance analyses developed independent PIDs specifically trained to identify ν_e CC and ν_μ CC events, respectively. The ν_e analysis developed two PIDs that had equal performance, the Likelihood Identifier (LID) and Library Event Matching (LEM) algorithms, and the ν_μ analysis developed the Reconstructed Muon Identifier (RemID). These were only used briefly as a selection cross check in chapter 6, so they are only briefly described. Technical notes with more detail are provided for each.

LEM, described in reference [105], compares an event to a library of known simulated events to find the best matches. The charge distribution of the trial event and a library event are aligned by the vertex and compared using a metric inspired by electrostatics. The algorithm compares the trial event to 77 million library events, and the 1000 best candidate matches are used to construct five physically based variables, a weighted fraction of signal matches, mean hadronic y, mean matched charge fraction, match energy difference, and enriched match fraction. Lastly, these variables and the calorimetric energy of the trial

event are input into a decision tree.

LID, described in reference [106], uses an artificial neural network to test the most energetic shower against electron, muon, proton, neutron, pion (charged and neutral), and photon hypotheses. The main inputs to this PID are the longitudinal and transverse energy deposition, but it also considers the fraction of event energy contained in the primary shower, the invariant mass of all event showers, the energy contained within eight planes of the vertex, the distance between the event vertex and shower start point, and the angle of the shower away from the beam direction. Note that the shower object mentioned above is a reconstructed object downstream of the Fuzzy K-Means Clustering algorithm, but these were not used for the NC disappearance analysis.

ReMID, discussed in detail in reference [107], is specifically designed to identify muons from charged pions. The PID is trained using a $k-$Nearest Neighbor classification on four input variables for a reconstructed track, the track energy deposition log likelihood, a scattering log likelihood, the track length, and the fraction of planes passed by the track that are uncontaminated by vertex or hadronic activity.

5.3.2 Cosmic Rejection PID

Aside from ReMID, the ν_μ disappearance analysis also developed a cosmic rejection PID [108]. This algorithm was originally tuned to distinguish cosmic events from ν_μ CC events, but was found to have similar separation power for the NC disappearance analysis as well. A boosted decision tree (BDT) was trained using a considerable number of inputs, largely quantities derived from the Kalman tracks best identified as a muon track by ReMID (referred to as 'the track' below). These inputs were:

- The angle of the track away from the neutrino beam.
- Y direction of the track.
- Maximum Y position of any hit activity.
- Number of hits in the track.
- Length of the track.

- Projected forward distance of the track to the detector edge.
- Projected backward distance of the track to the detector edge.
- Total scattering angle of the track over the total length of the track.
- Number of hits in the slice.
- Calorimetric energy in the slice.
- Minimum Y position of any hit activity.
- Total number of 3D Kalman tracks.

The BDT was trained separately for different data taking periods, and the final two items listed were only used for training in later data periods. The separate tunings were done to accommodate the changing detector conditions. This PID was only trained and filled for FD events. The output score distribution from the later periods is shown in figure 5.14.

5.3.3 CVN, the All Purpose PID

CVN, discussed in detail in reference [109], does not attempt to identify a single event type; rather, it gives the most likely event type among all possibilities. CVN, or convolutional visual network, is the specific application of a CNN, or convolutional neural network, a deep learning technique that has had success in the field of computer vision and specifically in image processing and classification. CVN was developed as a classifier for NOνA as the events in each view are essentially images of physics processes.

One of the advantages of CNNs is that they do not require a known set of features as inputs; instead, the CNN itself finds and trains on the image features. CNNs use a stack of kernels within a convolutional layer to transform an image in a variety of ways, where each kernel can be thought of as finding a particular feature. For a layer f that acts on an $n \times m$ set of pixels from an image g, the image is transformed via

$$(f * g)_{p,q,r} = \sum_{i=1}^{n} \sum_{j=1}^{m} \sum_{k=1}^{c} f_{i,j,k,r} g_{p+i,q+j,k}, \tag{5.23}$$

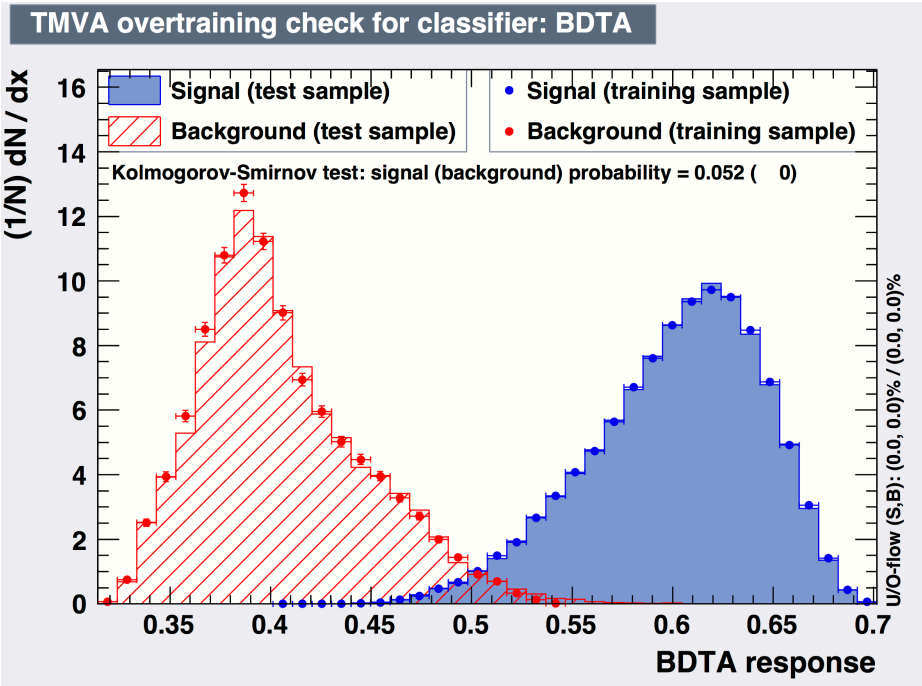


Figure 5.14: The cosmic rejection BDT output distributions for later data taking periods that encompass the majority of the data. The training sample is only ν_μ CCs and the background sample is only cosmic trigger data events. Other beam events tend to have a similar distribution to the signal.

where p and q are the output pixel indices, r is the output channel, k is the input channel, and c is the total number of channels. For a traditional image, channels refer to the RGB values of the pixel. In NOνA, the pixels are the individual cells and the channels are the two views of the event. The energy deposited in a cell is encoded into an 8-bit integer with 256 possible values much like the intensities of a standard color are encoded. After the kernel is applied to the full image, the output is called a feature map, and a separate map is created and stored from each of the r kernels in the convolutional layer.

CNNs contain many layers that perform different functions. Convolutional layers nearest to the input image tend to output feature maps that exaggerate specific features while later layers become more holistic and topological. Other layers down-sample the image using a pooling technique that replaces an $n \times m$ subset of pixels with either the maximum pixel value (max pooling) or the average pixel value (average pooling). Finally, a process called local response normalization (LRN) is used that normalizes the result of a given pixel in a feature map relative that pixel in adjacent maps to avoid local minima.

The different layers are chained together into a particular architecture. CVN was built as a modified

version of the GoogLeNet [110] architecture using the Caffe [111] framework. One of the features of the GoogLeNet architecture used in several layers for CVN is the inception module, shown in figure 5.15. This module performs several parallel feature extractions at different scales before concatenating the results into a single output layer. The full architecture of CVN is shown in figure 5.16. The XZ and YZ views are split and passed through two independent, equivalent branches. The output from these branches are only joined toward the end of the chain. The output of the final CVN inception module layer is passed into a traditional network for classification, where the event type outputs are normalized to 1 to be interpreted as probabilities.

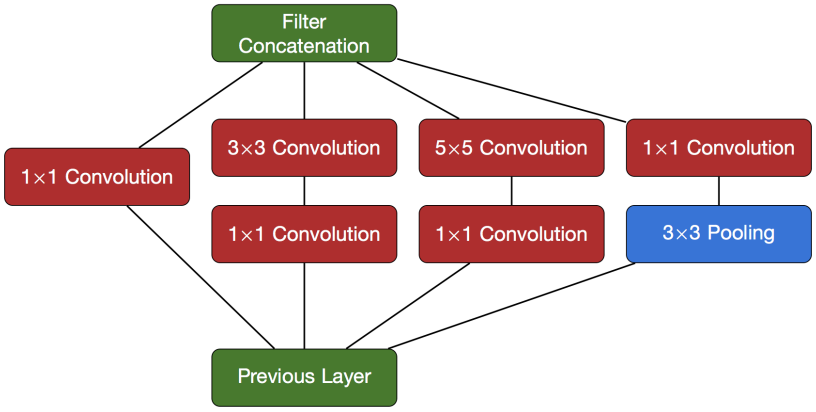


Figure 5.15: Schematic of the Inception module.

CNNs are trained by using a set of images with known classifications and by optimizing a set of weights for each feature to minimize a loss function. This was done for CVN in two steps [112], first training over a set of MC events, and second adding in cosmic data. The MC events were split into 13 types, $\{\nu_\mu$ CC, ν_e CC, ν_τ CC$\} \times \{$QE, RES, DIS, Other$\}$ and NC. After the training weights for these event types stabilized, a sample of cosmic trigger data was added to train the PID with the full suite of event types. The cosmics were added in a separate step to prevent the PID from simply learning to reject data. Instead of using the entire detector view as the initial input, a smaller 100 plane $\times$ 80 cell view is constructed with the most upstream plane lining up with the first hit and the cells centered on the median hit for training, testing, and evaluation.

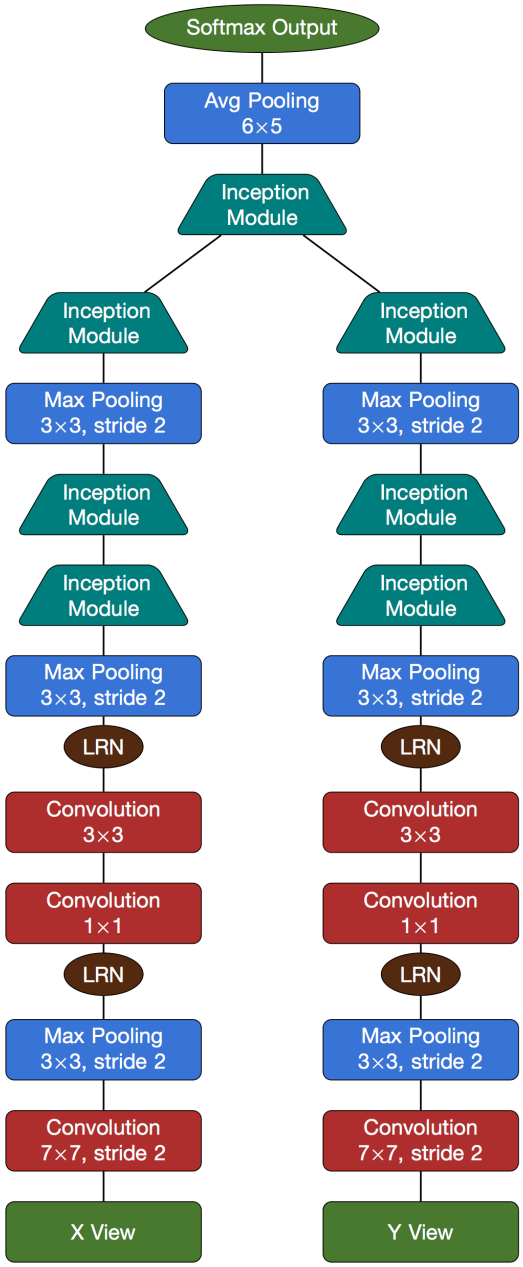


Figure 5.16: Schematic of the CVN architecture. The algorithm starts at the bottom and moves upward.

A true NC, ν_e CC, and ν_μ CC event are each shown with a corresponding feature map collection in figure 5.17. The maps shown are from the output of the first inception module for the YZ view. Three particular features are highlighted for each of these events that appear to correspond to traditional features. The green feature map shows activity for the ν_μ CC event as a long, track like feature. Unsurprisingly, this feature map shows little activity for the other two events. The blue map shows a more diffuse pattern for the ν_e CC event, but is almost entirely absent for the ν_μ CC and NC events. This feature map thus appears to be picking out the characteristic shower development of electrons. Finally, the pink map shows moderate activity for all events, with a small peak of activity in the center of the left edge. This feature map seems to be sensitive to the hadronic activity near the event vertex. To reiterate, these feature maps are not the inputs into the final classification layer, but to another inception module layer. As the layers become deeper, the extracted features become more abstract. Distributions of the NC classifier are shown in section 6.5 on NC event selection, section 9.2 on ND data/MC comparisons, and section 9.4 on the final results.

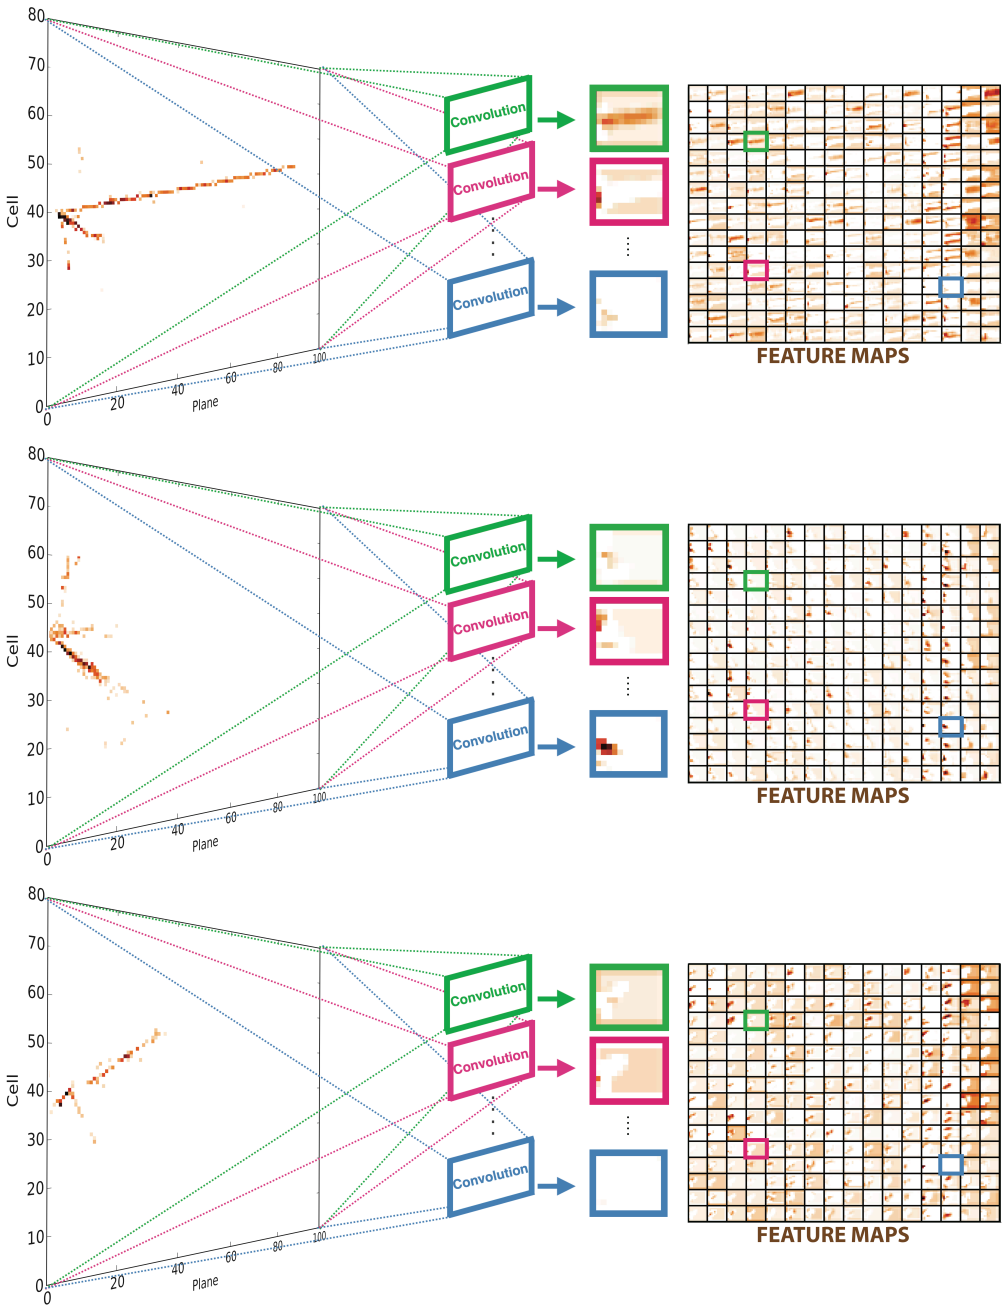


Figure 5.17: Feature maps from representative events. Top: True ν_μ CC. Middle: True ν_e CC. Bottom: True NC.

6

Neutral Current Event Selection

The NC disappearance analysis is performed by comparing the observed FD spectrum of NC events to the predicted number of events after extrapolation. This requires a largely pure sample of NC signal events. There are three main backgrounds for this sample, ν_e CC interactions, ν_μ CC interactions, and cosmic ray events. This chapter describes the full process for selecting a sample of NC events and eliminating backgrounds, briefly discussing data quality and event quality metrics, and fully detailing the fiducial, containment, NC selection, and cosmic rejection.

The analysis specific cuts were chosen based on the official second analysis (SA) files. These files were produced in tag `R16-03-03`, and NuS group specific DeCAF files were produced in tag `S16-04-08`. These files included MEC events, and the corresponding Tufts CC weight was applied to all (applicable) events [73]. For normalization simplicity, only 14 diblock MC and data were studied. After applying this criterion, all beam spectra were scaled to 6×10^{20} POT, and the cosmic spectra were scaled to 120 s. The cuts were set assuming the 3 flavor oscillation parameters listed in table 6.1. The remainder of this chapter

describes each level of event cut and their affect on the ND and FD spectra.

Oscillation Parameter	Value
ρ	$2.84\,\mathrm{g/cm^3}$
Δm^2_{21}	$7.53 \times 10^{-5}\,\mathrm{eV^2}$
$\sin^2 2\theta_{12}$	0.846
Δm^2_{32}	$2.37 \times 10^{-3}\,\mathrm{eV^2}$
θ_{23}	$\pi/4$
$\sin^2 2\theta_{13}$	0.085
δ	0

Table 6.1: The 3 flavor oscillation parameters assumed for studying the NC selection cuts

6.1 A Note on Preselection

The cuts discussed in this chapter were trained on cosmic trigger data files. These files had a cosmic veto motivated by the ν_e appearance analysis baked directly into the files. The cuts used are described in reference [113]. The actual cosmic prediction will came from the NuMI out of time side band which will had a more generalized cosmic veto.

6.2 Data Quality

Data quality cuts were developed well before the first NOνA results to ensure proper data taking conditions, and all analysis groups apply them as standard. These cuts are applied per spill, and spills that fail these cuts are not included in POT accounting. The cuts can be categorized into three main groups, beam quality, data quality, and timing. The beam quality cuts were studied and set as described in reference [114]. To be included, spills must meet all of the criteria listed in table 6.2. Two event quality cuts were applied to data from each detector. These cuts were motivated and set as described in references [115, 116, 117], and are summarized in table 6.3. Finally, a timing cut is applied to cosmic data to ensure that the data is not too close to the edge of the data taking window. For cosmic events within a given 500 μs trigger window, only events between 25 μs $< t <$ 475 μs are kept. Table 6.4 shows the number of events that pass these data quality cuts, and figure 6.1 shows the energy spectra at the ND and FD.

Beam Quality Parameter	Minimum	Maximum
Spill POT	2.00×10^{12}	
Horn Current	-202 kA	-198 kA
Beam X and Y position on target	0.02 mm	2.00 mm
Beam X and Y width	0.57 mm	1.58 mm
Time to nearest beam spill		0.5 s

Table 6.2: Beam quality cuts applied to each spill to ensure proper data taking conditions. Taken from reference [114].

Data Quality Parameter	Detector	Metric for Spill to Pass
Number of Missing DCMs	ND	$= 0$
Lights On Effect Hit Fraction	ND	≤ 0.45
Missing DCMs from LiveGeometry	FD	$= 0$
DCM Edge Match Fraction	FD	> 0.2

Table 6.3: Data quality cuts applied to each spill to ensure proper data taking conditions.

Cut Level	NC	ν_μ CC	ν_e CC	Cosmic
FD:				
Data Quality	337.0	230.6	58.5	23.4×10^6
ND ($\times 10^3$):				
Data Quality	11930	82594	1011	

Table 6.4: The number of events that pass the data quality cuts, at both detectors.

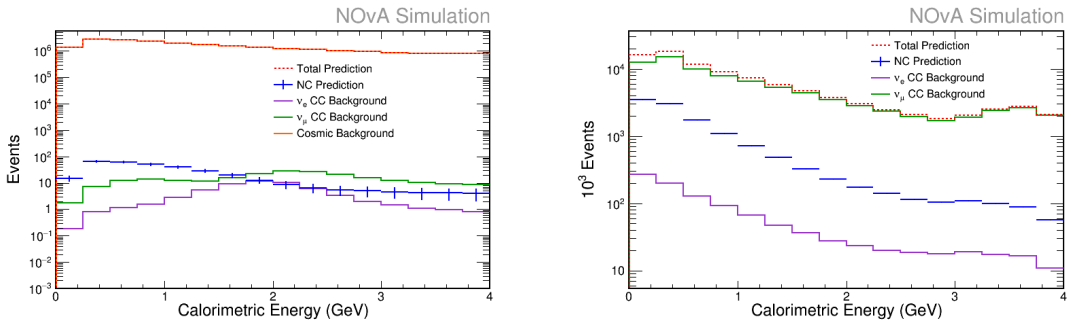


Figure 6.1: Energy spectra after Data Quality cuts for the FD (left) and ND (right).

6.3 Event Quality

Event quality cuts are the first applied to individual events and ensure that there is enough reconstructed information to be properly analyzed and that there were no obvious reconstruction failures. As with the data quality cuts above, these cuts were taken directly from the ν_e [118, 119] and ν_μ [120] analyses. Two of the cuts within this suite simply require the presence of a reconstructed vertex and a shower object. These reconstructed objects are used more extensively in later stages, so the event quality cuts make sure they are available. Other quantities considered are the number of hits per plane, distance between the reconstructed vertex and leading shower, and number of contiguous planes. Events with a high number of hits per plane are cut to remove so called 'FEB Flashers,' an event most often triggered by high energy cosmic rays. Likewise, events with a low number of contiguous planes are most often associated with very vertical cosmic rays. Finally, the a large distance between a vertex and leading shower is considered a reconstruction failure. These cuts are summarized with the exact parameters used for the cuts in table 6.5. The number of events before and after the event quality cuts are listed in table 6.6, and figure 6.2 shows the energy spectra of the events that pass these cuts.

It was discovered after these cuts were frozen that the requirement for the presence of at least one shower object implicitly imposed several other cuts. These cuts are thus already applied in the spectra shown in figure 6.2, but they are listed separately in table 6.7. Since the number of hits cut was not the same at both detectors, an explicit cut is made later in the cut flow as discussed in section 6.5.

6.4 Fiducial Volume and Containment

Fiducial volume and containment cuts are applied to remove events originating outside of the detector and to ensure that events originating inside of the detector do not have activity that escapes the detector. The fiducial volume cut is a cut on the location of the reconstructed neutrino vertex. The containment cut considers the leading prong of an event and cuts the event if the start or stop point is too close to a detector wall. The specific cuts were set separately at each detector.

Event Quality Metric	Metric for Event to Pass
Number of vertex objects	> 0
Number of shower objects	> 0
Number of hits per plane	< 8
Number of contiguous planes	> 2
Distance between vertex and leading shower	< 100 cm

Table 6.5: Quality cuts applied to individual events to ensure properly reconstructed quantities.

Cut Level	NC	ν_μ CC	ν_e CC	Cosmic
FD:				
Data Quality	337.0	230.6	58.5	23.4×10^6
+Event Quality	210.6	112.0	54.5	0.340×10^6
ND ($\times 10^3$):				
Data Quality	11930	82594	1011	
+Event Quality	7233	45251	592	

Table 6.6: The number of events before and after application of event quality cuts, at both detectors.

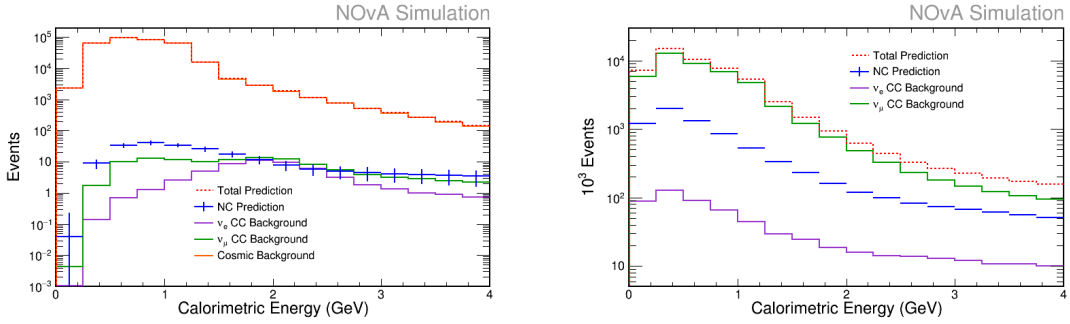


Figure 6.2: Energy spectra after Event Quality cuts for the FD (left) and ND (right).

Event Quality Metric	Metric for Event to Pass	Detector
Number of hits	> 20	FD
Number of hits	< 200	FD
Number of hits	> 10	ND
Length of Longest Prong	< 500 cm	ND, FD

Table 6.7: Event quality cuts applied implicitly due to the requirement of at least one shower object.

6.4.1 Far Detector

For the FD, the fiducial and containment cuts have the largest effect on the cosmic background. Since cosmic ray events originate above the detector, there are many more cosmic events with reconstructed vertices in the upper portion of the detector. The fiducial volume cut in Y is thus applied in an asymmetric way to eliminate more of the cosmic background. The fiducial volume cut on Z (the beam direction) is also asymmetric to account for cosmic rays entering the back of the detector hall where there is a smaller amount of rock overburden to shield these events. The fiducial volume cut on X is also applied slightly asymmetrically for a slight performance gain. Events that pass these fiducial cuts must still pass the containment cuts. For the FD, the leading prong must neither start nor stop within 10 cm of any detector face. Figure 6.3 shows the distributions of these variables before application of the fiducial or containment cuts. Table 6.9 summarizes the event rates before and after applying fiducial and containment cuts; figure 6.4 show the energy spectra of events that pass these cuts.

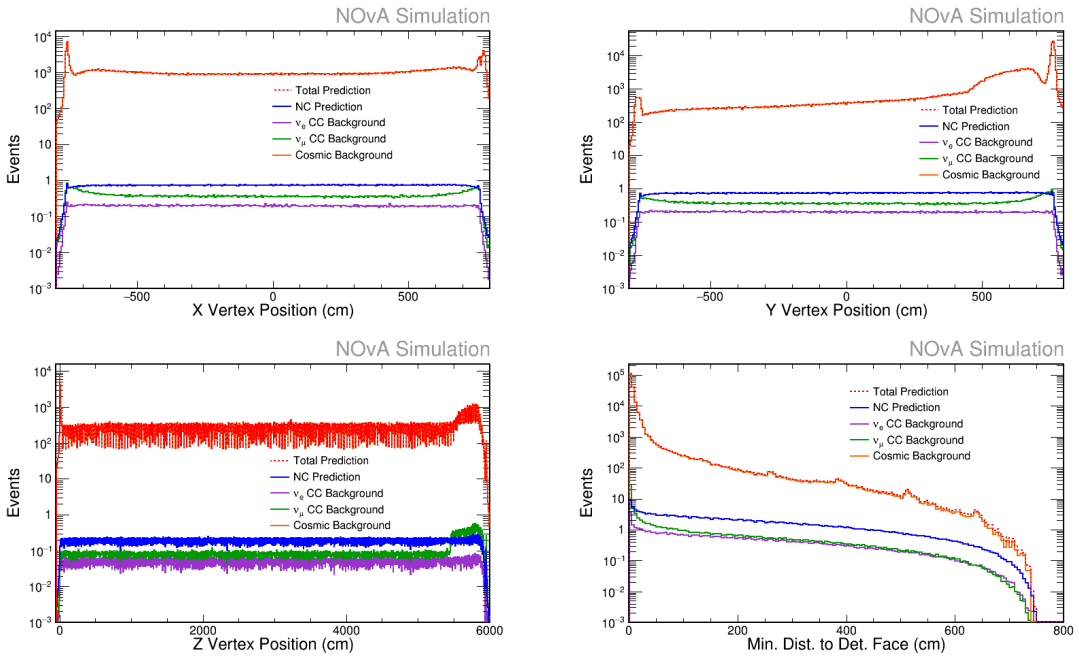


Figure 6.3: Distributions of variables used for fiducial and containment cuts at the FD. The top left, top right, and bottom left show the distributions of the reconstructed vertex. The bottom right shows the minimum distance between the start or stop of the leading prong to any detector face.

Reconstructed Quantity	Metric for Event to Pass
Reconstructed Vertex X Coordinate	$-680\,\text{cm} \leq vtxX \leq 650\,\text{cm}$
Reconstructed Vertex Y Coordinate	$-720\,\text{cm} \leq vtxY \leq 500\,\text{cm}$
Reconstructed Vertex Z Coordinate	$50\,\text{cm} \leq vtxZ \leq 5450\,\text{cm}$
Leading prong start/stop distance to detector face	$> 10\,\text{cm}$

Table 6.8: Fiducial and containment cuts applied to events in the FD.

Cut Level	NC	ν_μ CC	ν_e CC	Cosmic
...Event Quality	210.6	112.0	54.5	339600
+ Fiducial	132.0	46.7	34.0	26210
+ Containment	129.6	42.7	33.3	4324

Table 6.9: The number of events before and after application of fiducial and containment cuts at the FD.

6.4.2 Near Detector

While the ND does not have a large cosmic background to eliminate, there are many events that originate in the rock outside of the detector that can leak into the detector. Furthermore, the small size of the ND means that many events have activity that escapes outside of the detector. The fiducial and containment cuts at the ND are designed to combat both of these effects. The fiducial cuts on the X and Y coordinates of the reconstructed vertex were applied symmetrically, with a modestly large cut to remove events that originate in the rock outside the detector. The vertex cut on Z also cuts a large portion of the detector to cut rock events that leak into the front of the detector. The containment cuts at the ND are more strict

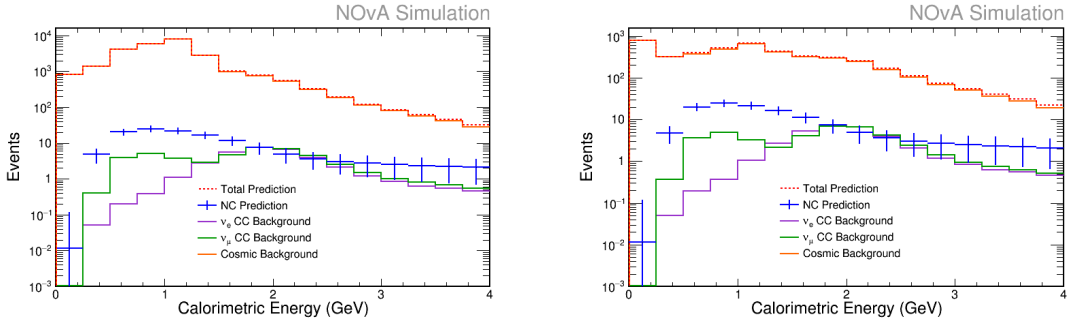


Figure 6.4: Energy spectra after fiducial and containment cuts for the FD. On the left, the cuts applied are data quality, event quality, and fiducial. The right plot also includes the containment cut.

than the FD. The leading prong is required to be greater than 25 cm from each detector face, and every track in the event is also subject to this cut. For the cut on each track, the back 'face' of the detector is considered to be the last plane of the fully active portion of the detector, effectively cutting events with activity in the muon catcher. All of these cuts are summarized in table 6.10, table 6.11 lists the event rates before and after the cuts are applied, and figure 6.5 shows the energy spectra of events that pass these cuts.

Reconstructed Quantity	Metric for Event to Pass
Reconstructed Vertex X Coordinate	$-100\,\text{cm} \leq vtxX \leq 100\,\text{cm}$
Reconstructed Vertex Y Coordinate	$-100\,\text{cm} \leq vtxY \leq 100\,\text{cm}$
Reconstructed Vertex Z Coordinate	$200\,\text{cm} \leq vtxZ \leq 1000\,\text{cm}$
Leading prong start/stop point, distance to detector face	$> 25\,\text{cm}$
Each track start/stop point, distance to detector face	$> 25\,\text{cm}$

Table 6.10: Fiducial and containment cuts applied to events in the ND.

Cut Level	NC	ν_μ CC	ν_e CC
...Event Quality	7233	45251	592
+ Fiducial	570.7	1397.8	21.7
+ Containment	328.0	379.4	11.0

Table 6.11: The number of events before and after application of fiducial and containment cuts at the ND. Here, containment refers to cuts on the distance of the leading prong and each track to the closest detector face. All numbers in this table are $\times 10^3$.

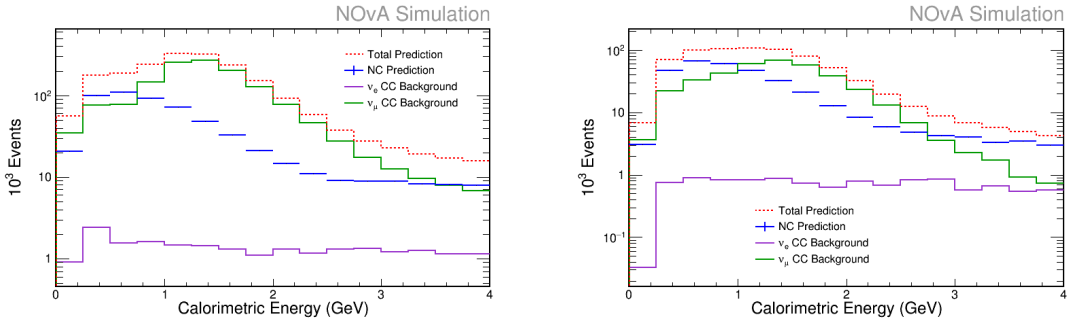


Figure 6.5: Energy spectra after fiducial and containment cuts for the ND. On the left, the cuts applied are data quality, event quality, and fiducial. The right plot also includes the containment cuts.

6.5 NC Selection

All of the cuts discussed to this point ensure that the remaining sample of events are well reconstructed, contained events. These events should all be representative of their respective types, i.e., NC, ν_e CC, and ν_μ CC. The NC selection cuts are designed to select a largely pure sample of NC events from amongst the CC backgrounds. The CVN, LEM, LID, and ReMID PID distributions all have power to complete this goal. Studies showed that CVN was sufficient on its own and that combining it with the other PIDs did not improve results.

A cut on the optimized NC classifier score is the main piece of the NC selection. However, due to the implicit cuts on the number of hits discussed in section 6.3, an explicit cut on the number of hits at both detectors is made here as well. The distributions of CVN and the number of hits at the FD are shown in figure 6.6. The cut on CVN was chosen by eye and set to keep events above 0.2. Table 6.12 summarizes the cuts used, table 6.13 lists the event rates before and after the selection cuts are applied, and figure 6.7 shows the energy distributions of events that pass these cuts.

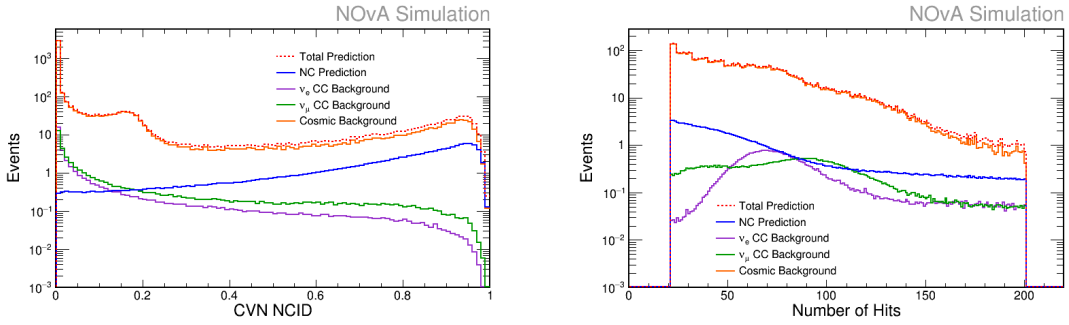


Figure 6.6: Distributions of variables used for NC selection at the FD.

Selection Parameter	Metric for Event to Pass
CVN NC distribution score	≥ 0.2
Number of Hits	$\geq 20, < 200$

Table 6.12: NC selection cuts to reject CC events, leaving a relatively pure sample of NC events.

Chronologically, CVN was developed later than the other PIDs and the analysis was originally developed using LID and ReMID. While they were ultimately replaced by CVN as the sole PID used in the

Cut Level	NC	ν_μ CC	ν_e CC	Cosmic
FD:				
...Containment	129.6	42.7	33.3	4324
+NC Selection	123.4	11.5	6.0	643.8
ND ($\times 10^3$):				
...Containment	328.0	379.4	11.0	
+NC Selection	263.9	124.4	3.5	

Table 6.13: The number of events before and after application of NC selection cuts, at both detectors.

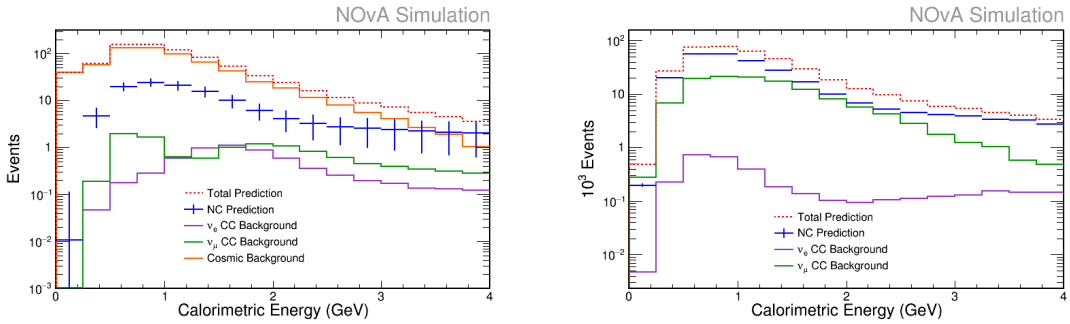


Figure 6.7: Energy spectra after NC selection cuts for the FD (left) and ND (right).

analysis, they nevertheless provided a valuable check that CVN was behaving sensibly. The PID distributions for ReMID and LID were studied after applying the full cut flow to ensure that it was removing largely the same events, and if there were any other minor performance increases to be gained. Figure 6.8 shows these distributions. Based on the results, it was decided that CVN provided all of the necessary separation power and so RemID and LID were not used in the analysis.

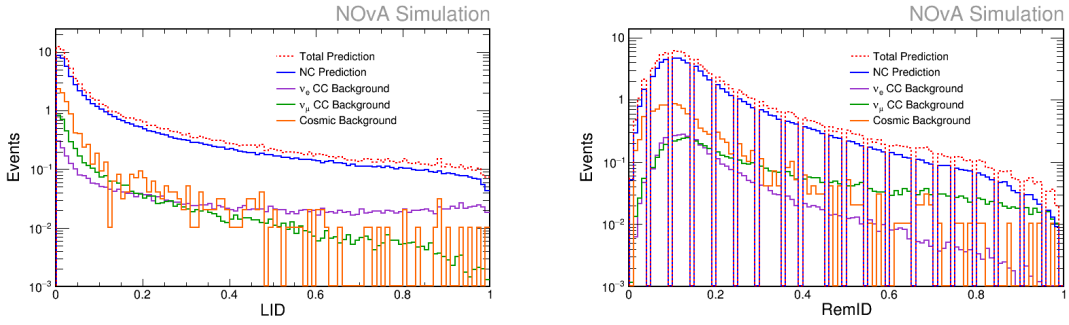


Figure 6.8: LID and RemID distributions at the FD after all other cuts applied.

6.6 Cosmic Rejection

The final set of cuts applied to the events is cosmic rejection. While there are essentially no cosmic events at the ND, the cuts are still applied at the ND when possible to maintain similar spectra at both detectors. The main variable to discriminate against cosmics is the output of the BDT described in section 5.3.2. Another powerful discriminant is the fraction of total transverse energy deposition versus total event energy deposition, originally developed for the ν_e appearance analysis [119]. The other cuts used for cosmic rejection are a cut on low average energy per hit, a harsher cut on the start/stop distance of the leading prong to the top of the detector, and removing events where the leading shower has a greatest likelihood of being a muon. The distributions of these variables at the FD are shown in figure 6.9 and table 6.14 summarizes the cosmic rejection with the values used for the actual cuts. Table 6.15 lists the event rates before and after the cosmic rejection cuts are applied, and figure 6.10 shows the energy distributions of events that pass these cuts.

Selection Parameter	Metric for Event to Pass
Cosmic BDT Score	> 0.5 (FD Only)
Particle Transverse Momentum Fraction	< 0.8
LID Algorithm μ Identifier	False
Average Energy per Hit	> 18 MeV (FD), > 9 MeV (ND)
Distance of Leading Prong to Detector Top	≥ 480 cm (FD Only)
Calorimetric Energy	> 0.5 GeV, < 4.0 GeV

Table 6.14: Final selection cuts applied to all events.

Cut Level	NC	ν_μ CC	ν_e CC	Cosmic
FD:				
...NC Selection	123.4	11.5	6.0	643.8
+ Cosmic Rejection	65.0	5.0	3.7	11.0
ND ($\times 10^3$):				
...NC Selection	263.9	124.4	3.5	
+ Cosmic Rejection	198.3	77.5	2.6	

Table 6.15: The number of events after each cut level at the FD.

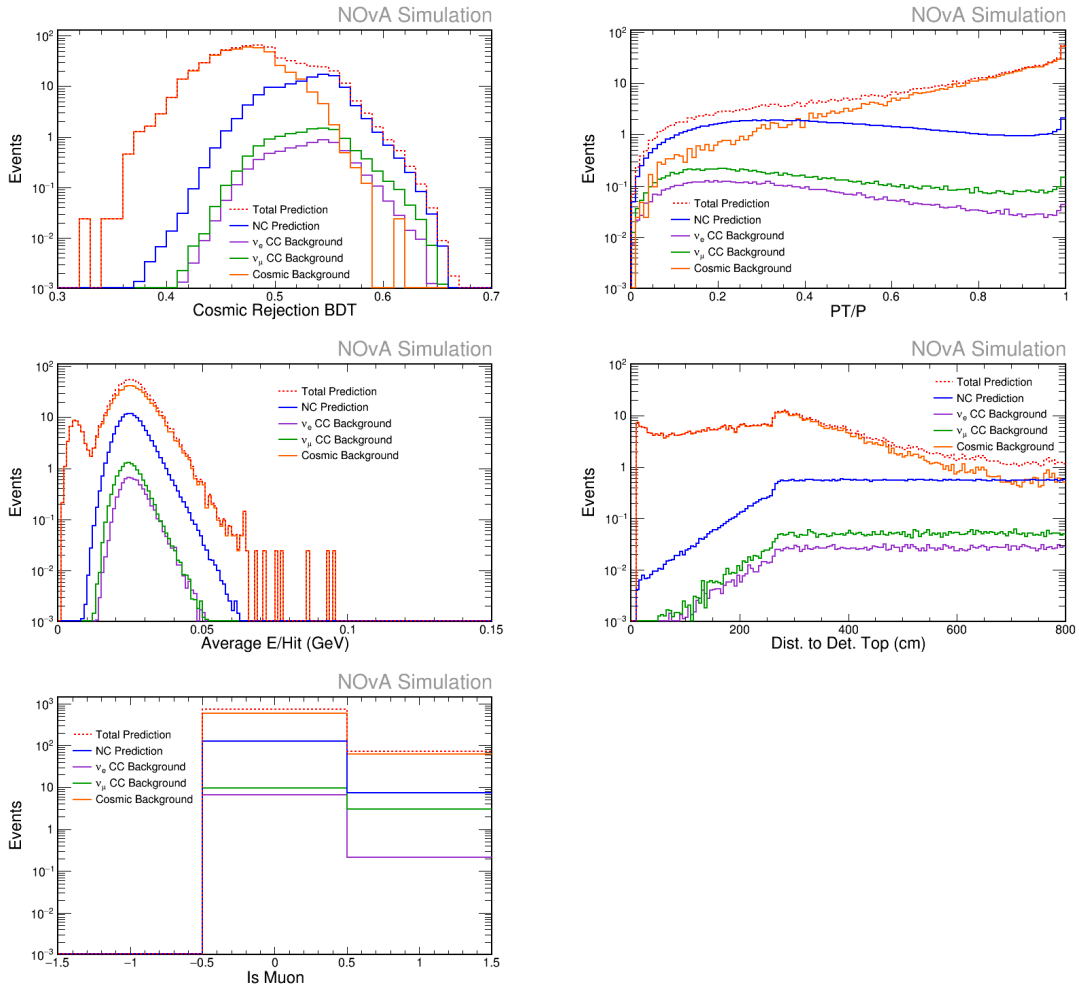


Figure 6.9: Distributions of variables used for cosmic rejection at the FD.

6.7 Summary

The final selection combines all of the previous sections with explicit energy cuts applied at both low and high energy. A low energy cut rejecting events below 0.5 GeV was used to reject very low energy events that could be subject to threshold events, and a high energy cut rejecting events above 4 GeV was placed due to diverging ND and FD selection efficiencies above that point. These energy cuts are listed in table 6.15. The event rates at each level of cut are summarized for the FD in table 6.16, and in table 6.17 for the ND. The final energy spectra are shown in figure 6.10.

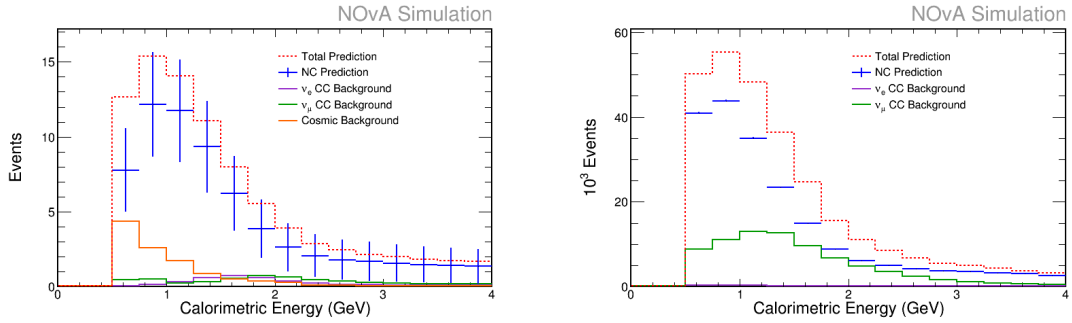


Figure 6.10: Energy spectra after all cuts for the FD (left) and ND (right).

Cut Level	NC	ν_μ CC	ν_e CC	Cosmic
Data Quality	337.0	230.6	58.5	23.42×10^6
+ Event Quality	210.6	112.0	54.5	339600
+ Fiducial	132.0	46.7	34.0	26210
+ Containment	129.6	42.7	33.3	4324
+ NC Selection	123.4	11.5	6.0	643.8
+ Cosmic Rejection	65.0	5.0	3.7	11.0

Table 6.16: The number of events after each cut level at the FD.

Cut Level	NC	ν_μ CC	ν_e CC
Data Quality	11930	82594	1011
+Event Quality	7233	45251	592
+ Fiducial	570.7	1397.8	21.7
+ Containment	328.0	379.4	11.0
+ NC Selection	263.9	124.4	3.5
+ Cosmic Rejection	198.3	77.5	2.6

Table 6.17: The number of events after each cut level at the ND. All numbers in this table are $\times 10^3$.

7

Neutral Current Spectrum Prediction

In order to extract measurements of oscillation parameters, the measured FD spectrum must be compared to a predicted distribution. The FD MC by itself does not constitute a good prediction due to various uncertainties like the beam flux and interaction cross sections. A more robust prediction is crafted by combining information from the ND data with the ND and FD MC. This chapter describes the bulk of that procedure for the FD NC disappearance analysis. The methods presented use the NC event selection discussed in chapter 6 to predict the FD energy spectrum, unless otherwise noted.

Prediction of the FD energy spectrum consists of three main steps called decomposition, extrapolation, and prediction. Decomposition is a technique to separate ND data into the various interaction types. Extrapolation uses the ND data, ND MC, and FD MC to create an unoscillated FD spectrum with enough information retained to apply oscillation weights in bins of true energy. The prediction step applies the oscillations and returns a spectrum in bins of reconstructed energy, the final FD prediction.

The analysis techniques described here were performed in the CAFAna software framework [121] that

analyzes Common Analysis Format files, or CAFs. The design principles of CAFAna recognize that all neutrino analyses have to perform essentially this same analysis chain. A given task might have multiple implementations, but while the inner details are different, the number and type of end products are the same. Thus CAFAna has a fixed general analysis chain with the ability to easily swap specific implementations of any given piece of the chain. The remainder of this chapter describes the methods used for the NC disappearance analysis.

7.1 ND Decomposition

The ND decomposition is the process of splitting ND data into a set of component spectra. Each interaction type, neutrino flavor, and neutrino sign has a different set of oscillation weights, so the data must be split by these categories at some stage before oscillations are applied. The CAFAna design choice was to make this split in the decomposition to allow the option for each component to extrapolate differently. Since there are no $(\bar{\nu}_\tau)$ CC interactions at the ND, the decomposition returns 5 component spectra, NC, ν_μ CC, $\bar{\nu}_\mu$ CC, ν_e, and $\bar{\nu}_e$ CC.

A number of approaches are possible for performing the decomposition. For this analysis, the data was decomposed proportional to the MC; i.e., the total normalization of the data was respected for each energy bin, but the percentage of events attributed to each component was set to that found in the MC. Figure 7.1 shows the results of this decomposition using the standard NC selection. Some data driven techniques have been developed [122, 123], but they were not explored as possibilities for the NC disappearance analysis. Other techniques like the varied horn current, a generalized version of the horn on/off technique, are not feasible for NOνA. In this method, the beam composition is different if the horn current is changed, and this fact is used to create a system of equations to solve for the NC and ν_μ CC beam components.

$$N^a_{NC} = \frac{r_{\nu_\mu} N^a - N^b + \left(r_{\nu_e} - r_{\nu_\mu}\right) N^a_{\nu_e}}{r_{\nu_\mu} - r_{NC}} \tag{7.1}$$

$$N^a_{\nu_\mu} = \frac{r_{NC} N^a - N^b + \left(r_{\nu_e} - r_{NC}\right) N^a_{\nu_e}}{r_{NC} - r_{\nu_\mu}} \tag{7.2}$$

Above, a and b refer to the two horn current configurations, N^x is the total number of observed events in

the energy bin, $N^a_{\nu_e}$ is the number of MC ν_e CC events in the energy bin, and $r_x \equiv N^b_x/N^a_x$ is calculated from the MC. Studies showed that since NOνA is off-axis, the event ratios r were too similar in even the horn on/off case, meaning that the system of equations is not truly independent and the technique is not viable [124].

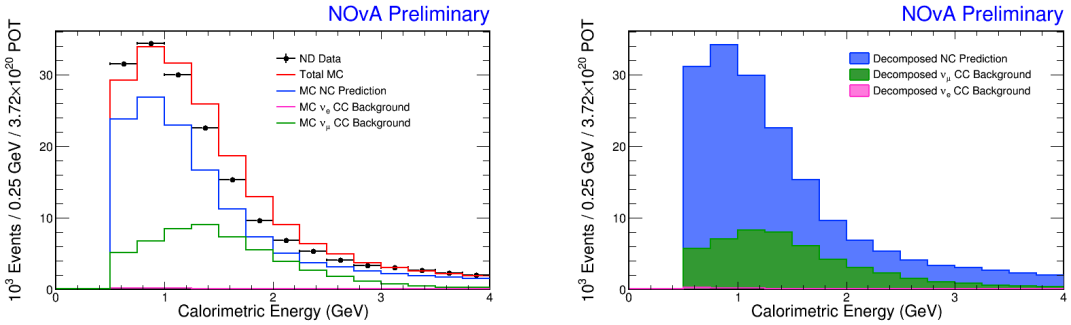


Figure 7.1: ND data decomposed into individual components. Left: the total data, total MC, and each MC component (not split by sign). Right: a stacked histogram of the decomposed data.

7.2 Extrapolation

The next part of the analysis chain extrapolates the ND data to the FD using both ND and FD MC. The output of the extrapolation is each oscillation component in bins of reconstructed energy (also called calorimetric energy) and true energy. The advantage of applying oscillation weights at a separate, later stage is that the extrapolation does not need to be repeated when applying many different oscillations as in a fit. There are two main types of extrapolations, indirect and direct. Indirect extrapolation involves measuring the ND data spectrum and using the data to tune the MC simulation. Direct extrapolation applies MC techniques to the measured ND spectra. For the analysis presented here, a simple Far Over Near (F/N) direct extrapolation method was used.

The basic premise of the extrapolation is to apply a F/N ratio to the data.

$$\mathrm{FD}^{Pred} = \mathrm{ND}^{Data}\frac{\mathrm{FD}^{MC}}{\mathrm{ND}^{MC}} = \mathrm{FD}^{MC}\frac{\mathrm{ND}^{Data}}{\mathrm{ND}^{MC}} \tag{7.3}$$

This method takes advantage of the fact that systematic effects cancel to first order due to the use of functionally identical detectors. Consequently, the F/N ratio mainly encodes the differences between the two

detector spectra due to flux kinematics. Note that the F/N ratio is equivalent to applying a ND data/MC ratio to the FD MC, as shown in the second equality of equation 7.3.

The actual process is a bit more complicated than this, and the results depend on which variable the F/N ratio is measured in and what selection cuts are applied at the ND. For components with sufficient statistics and energy resolution, the F/N ratio can be applied in bins of true energy. This requires more MC information and extra error checking. As the statistics or energy resolution degrades, the F/N ratio is applied only in bins of calorimetric energy. For the components with the least information, the FD MC is taken as the component prediction. Each component is extrapolated separately using one of these three methods. For appearance components, $\nu_\alpha \rightarrow \nu_\beta$, it is important to apply a ν_α-like selection to the ND spectra. This provides the most accurate correction to the FD MC as can be seen from the ND data/MC interpretation of equation 7.3. Of course, the desired output is a number of events selected as NC-like, so the NC selection cuts are always applied at the FD.

The goal of the extrapolation is to predict the number of events from each component that are selected as NC-like in bins of true and calorimetric energy, $\mathrm{FD}^{Pred}_{\alpha\rightarrow\beta}(S_{NC}; E^R_i, E^T_j)$, where S_x denotes x-like event selection, E^R_i denotes the ith row of calorimetric energy, and E^T_j denotes the jth row of true energy. This is represented as a two dimensional histogram. In the simplest case the MC is taken as the prediction,

$$\mathrm{FD}^{Pred}_{\alpha\rightarrow\beta}(S_{NC}; E^R_i, E^T_j) = \mathrm{FD}^{MC}_{\alpha\rightarrow\beta}(S_{NC}; E^R_i, E^T_j), \tag{7.4}$$

where FD^{MC} denotes a quantity calculated from the MC. This method is used for the smallest CC background components with the lowest statistics, $\overset{(-)}{\nu}_\mu \rightarrow \overset{(-)}{\nu}_\tau$, $\overset{(-)}{\nu}_e \rightarrow \overset{(-)}{\nu}_\mu$, $\overset{(-)}{\nu}_e \rightarrow \overset{(-)}{\nu}_\tau$, $\bar{\nu}_\mu \rightarrow \bar{\nu}_\mu$, and $\bar{\nu}_e \rightarrow \bar{\nu}_e$. For components with better statistics, the FD MC is corrected in bins of calorimetric energy.

$$\mathrm{FD}^{Pred}_{\alpha\rightarrow\beta}(S_{NC}; E^R_i, E^T_j) = \frac{\mathrm{FD}^{MC}_{\alpha\rightarrow\beta}(S_{NC}; E^R_i, E^T_j) \cdot \mathrm{ND}^{Data}_{\alpha}(S_{NC}; E^R_i)}{\mathrm{ND}^{MC}_{\alpha}(S_{NC}; E^R_i)} \tag{7.5}$$

This method is used for the NC signal, ν_μ and ν_e CC survival background components. For survival components ($\beta = \alpha$), the NC selection is applied to the event spectra at both detectors. For appearance components, the events in the FD are due to neutrinos of a different flavor than those at the ND. Thus

the NC event selection is still applied to the FD spectrum, but at the ND a different selection method is applied to select a pure sample of the unoscillated neutrinos. The remaining two components, $\overset{(-)}{\nu}_\mu \rightarrow \overset{(-)}{\nu}_e$, are both appearance components from muon flavor neutrinos, so a ν_μ-like event selection is applied at the ND, specifically the cut `kNumuND` described in reference [125]. Furthermore, these components have enough ND statistics to apply the F/N ratio in bins of true energy.

$$\mathrm{FD}^{Pred}_{\alpha\rightarrow\beta}(S_{NC}; E^R_i, E^T_j) = \frac{\mathrm{FD}^{MC}_{\alpha\rightarrow\beta}(S_{NC}; E^R_i, E^T_j) \cdot \mathrm{ND}^{Pred}_{\alpha}(S_{\nu_\mu}; E^T_j)}{\mathrm{ND}^{MC}_{\alpha}(S_{\nu_\mu}; E^T_j)} \tag{7.6}$$

$$\mathrm{ND}^{Pred}_{\alpha}(S_{\nu_\mu}; E^T_j) = \sum_k \frac{\mathrm{ND}^{MC}_{\alpha}(S_{\nu_\mu}; E^R_k, E^T_j) \cdot \mathrm{ND}^{Data}_{\alpha}(S_{\nu_\mu}; E^R_k)}{\mathrm{ND}^{MC}_{\alpha}(S_{\nu_\mu}; E^R_k)} \tag{7.7}$$

Equation 7.7 can be interpreted as a two step process, where the combination of terms inside the sum normalizes the calorimetric energy columns to the data values, then sums across the true energy row to find the ND true energy predicted value. Equations 7.5 and 7.6 then have the same structure, where ND^{Pred} is used as the data input. Note that for equations 7.5, 7.6, and 7.7, if the ND^{MC} value is 0 the extrapolation falls back to using the base FD MC for the prediction, but only for the particular bin that would otherwise require division by 0. Figure 7.2 shows the F/N ratio for the NC component, as well as the ND data/MC for the alternative interpretation of equation 7.3. Figure 7.3 shows how the extrapolation changes the predicted distribution of the NC component.

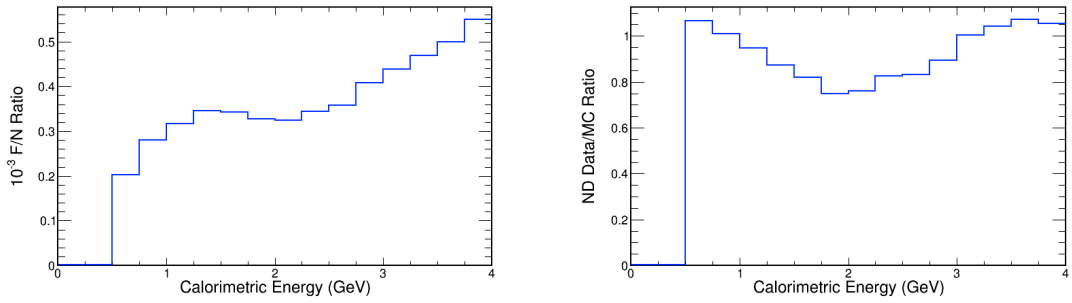


Figure 7.2: The left figure shows the F/N ratio for the NC component extrapolation. Equation 7.3 can also be interpreted as applying a ND data/MC ratio to the FD MC; this ratio is shown in the right figure.

In reality, only the neutrino flavor affects the oscillations, and not whether the neutrino interacts via a

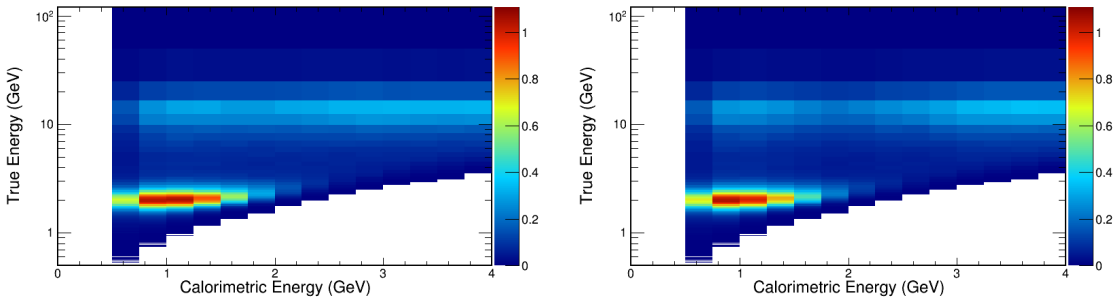


Figure 7.3: These plots show how the extrapolation affects the predicted NC component distribution. Left: FD base MC. Right: Extrapolated spectrum. Both plots are scaled to 6×10^{20} POT and have the same Z axis range.

CC or NC interaction. For the standard three flavor model, the number of NC interactions is not affected by oscillations, so it is okay to treat the NCs as any other component within the extrapolation. The NC disappearance analysis built its extrapolation upon the existing ModularExtrap class [126] that handles the NC component in this fashion. However, when sterile neutrinos are added to the model, oscillations do affect the NC rate and thus the neutrino flavor breakdown of the NC interactions does matter. So long as the final neutrino flavor is an active one, the exact flavor does not matter. Therefore, it suffices to split the NCs based on their four initial flavors, $\overset{(-)}{\nu}_\mu$ and $\overset{(-)}{\nu}_e$. This was done by extrapolating the NCs as described above and splitting the FD result with flavor proportions constructed from the FD MC.

$$\mathrm{FD}^{Pred}_{NC,\alpha}(S_{NC}; E^R_i, E^T_j) = \mathrm{FD}^{Pred}_{NC}(S_{NC}; E^R_i, E^T_j) \frac{\mathrm{FD}^{MC}_{NC,\alpha}(S_{NC}; E^R_i)}{\mathrm{FD}^{MC}_{NC}(S_{NC}; E^R_i)} \tag{7.8}$$

Figure 7.4 shows the flavor breakdown of the NC component. It was convenient to use this method based on what already existed within ModularExtrap, but it also avoided the need to construct raw spectra of the number of neutrinos per flavor using cross section information with large uncertainties.

Table 7.1 summarizes the method used to extrapolate each component of the FD prediction.

7.3 Far Detector Prediction

Oscillations weights are applied in the final step of the analysis chain in the prediction step. The output of the extrapolation is required to contain true energy information as the oscillations are applied in bins of true energy. The predicted event spectrum for a given component is then calculated by summing over the

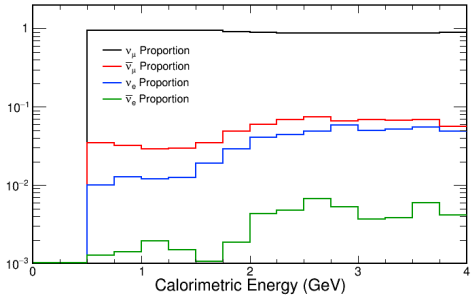


Figure 7.4: Proportions of the NC component from each neutrino flavor. The proportions are calculated based on the initial neutrino flavor, i.e., before oscillations are considered.

bins of true energy, returning a one dimensional spectrum in calorimetric energy bins.

$$\mathrm{FD}^{Pred}_{\alpha\to\beta}(S_{NC};E^R_i) = \sum_j \mathrm{FD}^{Pred}_{\alpha\to\beta}(S_{NC};E^R_i,E^T_j)\cdot P(\nu_\alpha\to\nu_\beta,E^T_j) \tag{7.9}$$

The NC component was handled as four separate components based on the initial neutrino flavor. For each of these, the probability of oscillating into any active neutrino flavor was applied.

$$\mathrm{FD}^{Pred}_{NC,\alpha}(S_{NC};E^R_i) = \sum_j \mathrm{FD}^{Pred}_{NC,\alpha}(S_{NC};E^R_i,E^T_j)\cdot P(\nu_\alpha\to\nu_{Active},E^T_j) \tag{7.10}$$

$$P(\nu_\alpha\to\nu_{Active}) \equiv \sum_{\ell\in\{e,\mu,\tau\}} P(\nu_\alpha\to\nu_\ell) = 1 - P(\nu_\alpha\to\nu_s) \tag{7.11}$$

The final event spectrum is simply the sum over all components, signal and background.

$$\mathrm{FD}^{Pred}(S_{NC};E^R_i) = \sum_{\{Comp\}} \mathrm{FD}^{Pred}_{Comp}(S_{NC};E^R_i) \tag{7.12}$$

The total event rates for each component, based on the oscillation parameters in table 6.1, are summarized in table 7.1. This table also shows a comparison between the rates predicted by the base MC and the extrapolation. Figure 7.5 shows all of the predicted component distributions based on the extrapolation, with the exception of cosmics.

Component	F/N Method	FD MC Prediction	Extrapolated Prediction
NC	Calorimetric Energy Bins	65.19	60.61
$\nu_\mu \rightarrow \nu_\mu$	Calorimetric Energy Bins	4.98	4.47
$\nu_\mu \rightarrow \nu_e$	True Energy Bins	2.88	2.83
$\nu_\mu \rightarrow \nu_\tau$	Base FD MC	0.32	0.32
$\bar{\nu}_\mu \rightarrow \bar{\nu}_\mu$	Base FD MC	0.10	0.10
$\bar{\nu}_\mu \rightarrow \bar{\nu}_e$	True Energy Bins	0.043	0.044
$\bar{\nu}_\mu \rightarrow \bar{\nu}_\tau$	Base FD MC	0.046	0.046
$\nu_e \rightarrow \nu_\mu$	Base FD MC	0.010	0.010
$\nu_e \rightarrow \nu_e$	Calorimetric Energy Bins	0.77	0.72
$\nu_e \rightarrow \nu_\tau$	Base FD MC	0.0010	0.0010
$\bar{\nu}_e \rightarrow \bar{\nu}_\mu$	Base FD MC	0.0002	0.0002
$\bar{\nu}_e \rightarrow \bar{\nu}_e$	Base FD MC	0.036	0.036
$\bar{\nu}_e \rightarrow \bar{\nu}_\tau$	Base FD MC	0.0001	0.0001

Table 7.1: Extrapolation method used for each component. The final two columns show the total number of predicted event rates after oscillations are applied as described in section 7.3. These numbers used all diblock configurations and scaled to 6.69×10^{20} POT as described at the beginning of chapter 9. All components other than NC are CC backgrounds.

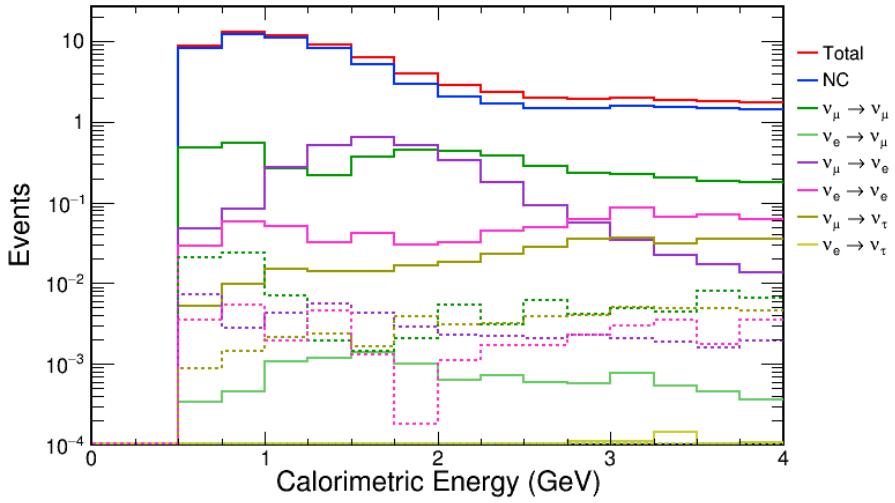


Figure 7.5: FD predicted distributions of all 13 beam components, assuming oscillation parameters listed in table 6.1. The dashed line of a given color shows the predicted anti-neutrino component that corresponds to the component listed in the legend.

8

Systematic Uncertainty Analysis

As with any analysis, the NC disappearance analysis is sensitive to a number of systematic effects. NOνA was designed with two functionally identical detectors, so that Near Detector data can be used to constrain or correct the Far Detector prediction. Since many effects such as beam and cross section uncertainties affect the spectra at both detectors in a similar or the same way, this two detector technique leads to the reduction of these systematic errors. Other effects, such as Near Detector rock event contamination, require a data driven technique to quantify.

The general technique for analyzing systematic uncertainties was to run the full extrapolation chain and generate a predicted spectrum with and without a systematic effect applied. As in chapter 6, the prediction assumed the same three flavor oscillation parameters listed in table 6.1, and for normalization simplicity, only 14 diblock MC and data were studied. All spectra were then scaled to 6×10^{20} POT. Each given systematic effect was used to shift the MC simulation at one or both detectors as appropriate. The resulting difference between the shifted and nominal spectra was quantified as a systematic uncertainty.

The systematic effects analyzed included uncertainties arising from the beam, GENIE simulation, Birks-Chou light yield simulation, calibration, bias from the ND containment, ND rock event contamination, ND data/MC spectrum and hadronic energy differences, noise model, MC statistics, and overall normalization. The rest of this chapter discusses each of these effects in greater detail.

8.1 Beam

The NOνA MC simulation involves a fully detailed model of the NuMI beam process in an attempt to create the most realistic MC possible, but systematic errors can result from any mismatch between simulation and reality. The NOνA Beam Working Group performed studies to assess the effect that uncertainties in the simulation can have on the neutrino flux [79]. These studies included the effects of incorrectly modeling various parts of the beam transport and the effects of uncertainties in hadron production arising from fixed target experiments.

To quantify the systematic uncertainty caused by these beam uncertainties, a sample flux was generated using a systematic shift and compared to the nominal flux via a simple ratio. Results were generated separately for each neutrino flavor and for each detector. The ratios were used to modify the MC from the full simulation used as inputs to the extrapolated prediction. Finally, the shifted FD prediction was compared to the nominal prediction, with any differences being taken as the systematic uncertainty. At the end of this process, the individual uncertainties were added in quadrature. Uncertainties were calculated for the following systematics:

- Beam position on target varied by $\pm 0.5\,mm$ in X
- Beam position on target varied by $\pm 0.5\,mm$ in Y
- Beam spot size varied by $\pm 0.2\,mm$ in both X and Y
- Target position varied by $+2\,mm$ in Z
- Horn current varied $\pm 1\,kA$
- Horn 1 position varied by $\pm 2\,mm$ in both X and Y
- Horn 2 position varied by $\pm 2\,mm$ in both X and Y

- Horn magnetic field model changed from linear to exponential distribution
- Changes to MC beamline simulation
- Comparison between versions of FLUKA
- Comparison between FLUKA and G4NuMI
- Comparison between FLUKA and NA49

Hadron production uncertainties were combined in quadrature before being provided as weights, so this is evaluated as single systematic uncertainty. The percentage difference due to each systematic is shown in table 8.1, and the full uncertainty envelopes for NC signal and background are shown in figure 8.1. The beam systematics had an overall effect of 3.4% on the NC signal and 3.6% on the background.

Systematic	NC Difference (%)	Background Difference (%)
Beam Position, X	0.5	0.4
Beam Position, Y	0.2	0.2
Beam Spot Size	0.2	0.3
Target Position	0.04	0.1
Horn Current	0.1	0.1
Horn 1 Position	0.6	0.8
Horn 2 Position	0.2	0.2
Horn B Field	0.05	0.05
Geometry Simulation Change	0.4	1.2
FLUKA Versions	0.04	0.4
FLUKA v G4NuMI	0.4	1.2
FLUKA v NA49	3.0	2.6
Combined	3.4	3.6

Table 8.1: The percentage difference between the shifted and nominal predictions for the number of FD events due to beam systematics.

8.2 Birks-Chou Light Yield Simulation

As discussed in section 4.2.1 the MC simulation employs the Birks-Chou Law (equation 4.1) to model the relationship between scintillator light yield and energy deposition. This empirical law depends on several constants dependent on the particular scintillator material. These constants were estimated for the

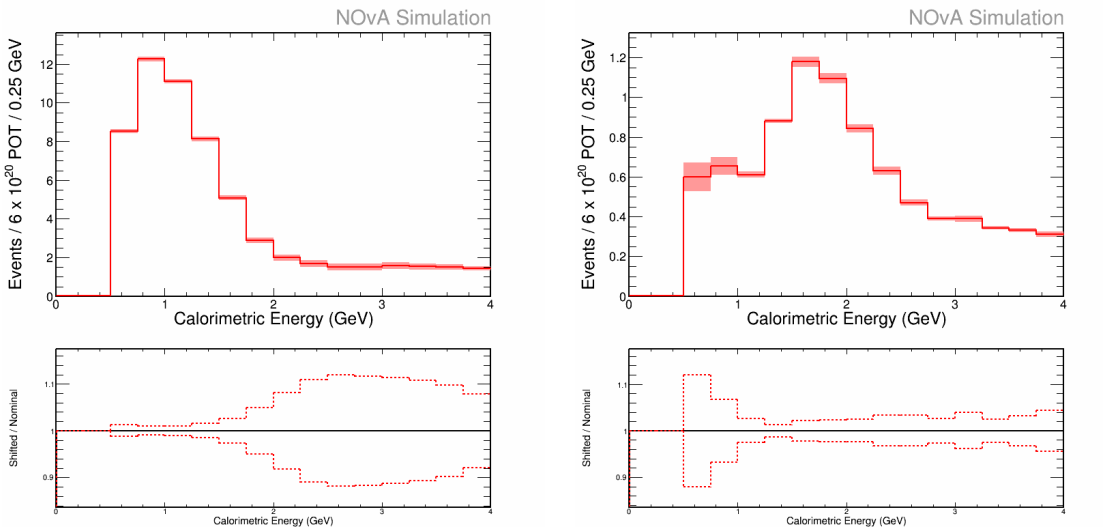


Figure 8.1: Beam systematics uncertainty envelope on the NC signal (left) and background (right) event spectra, after extrapolation. The envelope was calculated by adding in quadrature the larger of $|+1\sigma|$ and $|-1\sigma|$ for each individual systematic.

NOνA scintillator via a data driven study [77]. The results of the study were $k_B = 0.04$ cm/MeV and $k_C = -0.0005\,(\text{cm/MeV})^2$.

The systematic uncertainty based on the Birks-Chou light yield simulation was quantified by comparing the nominal FD prediction to predicted spectra using alternative Birks-Chou parameter constants. The values reported in the study from reference [77] were much larger than other typical measurements, so two MC samples were generated with more traditional values, one with $k_B = 0.01$ cm/MeV called BirksB, the other with $k_B = 0.02$ cm/MeV called BirksC, both with $k_C = 0$. Shifted FD event spectra were predicted by extrapolating the same set of ND data as the nominal prediction, but using the MC with alternative Birks-Chou model parameters. The uncertainty was taken as the percentage difference between the nominal and shifted predictions. Table 8.2 shows the percentage differences from both MC samples; figure 8.2 shows the shifted event spectra compared to nominal. Instead of combining the individual uncertainties in quadrature, the shifted sample with the larger overall difference from the BirksC model taken as the systematic uncertainty, which had a 2.4% effect on the NC signal, and a 1.8% effect on the background.

Model	NC Difference (%)	Background Difference (%)
BirksB	0.5	1.2
BirksC	2.4	1.8

Table 8.2: The percentage difference between the shifted and nominal predictions for the number of FD events due to extrapolation using MC with alternative Birks-Chou model parameters.

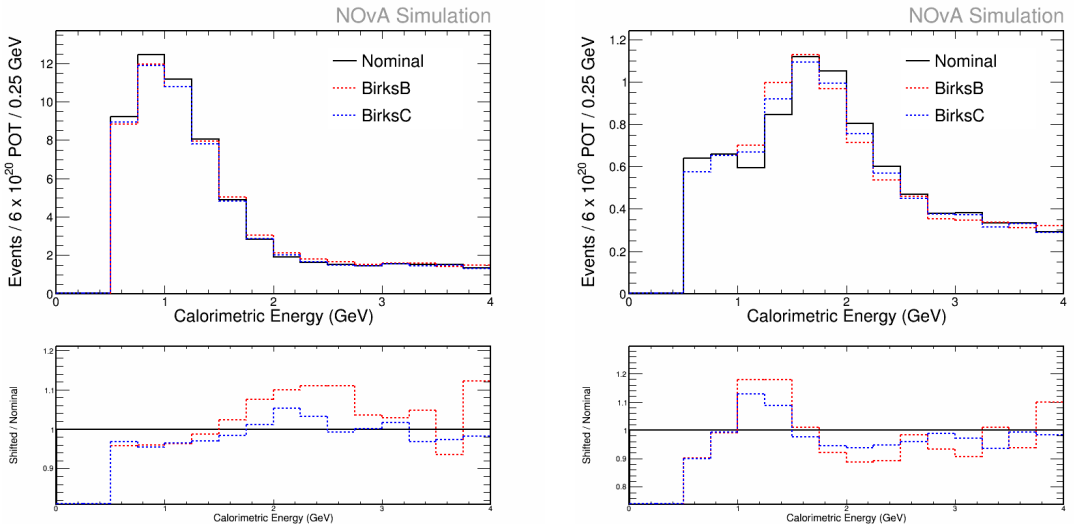


Figure 8.2: Shifted FD predictions due to extrapolation of ND data with MC using different Birks-Chou parameter values. The NC signal spectrum is on the left; the background spectrum is on the right.

8.3 Calibration

The calibration procedure is designed to make a constant energy response both across each detector and between the two, but any problem can introduce a systematic error. This error was evaluated by studying various MC samples with an engineered miscalibration. The effects studied included a miscalibration that varied as a function of the cell length and an overall scale miscalibration. The functional miscalibration was studied separately between the X and Y views of the detector. The scale miscalibration was applied as a 5% effect both up and down at each detector. These effects are motivated by studying data/MC comparisons for various samples [86, 127, 128] and were updated based on the most recent round of data [129].

The systematic uncertainty was evaluated in two regimes. The first applied the same miscalibration to the MC at both detectors, and a shifted prediction was then generated using these shifted as inputs into the extrapolation. This procedure was followed for each type of miscalibration, and all of these systematics

were added in quadrature at the end. The overall scale miscalibration was also applied as a miscalibration to a single detector to measure the effect of a missed absolute calibration. A shifted prediction was then generated as above, but with this procedure only the maximum overall effect was used as the systematic uncertainty that was then added in quadrature with the systematics from above. Table 8.3 shows the percentage difference for each of the calibration systematics. The effects of the calibration systematics on the NC signal and CC background spectra are shown in figures 8.3 and 8.4 when applied to both detectors and one detector, respectively. The overall uncertainty was 5.8% on the NC signal and 6.0% on the background.

Miscalibration	NC Difference (%)	Background Difference (%)
Miscalibration applied to both detectors:		
Sloped X	0.5	0.5
Sloped Y	1.6	1.1
Flat Scale Up	1.6	1.2
Flat Scale Down	3.5	2.3
Miscalibration applied to one detector:		
Flat Scale Up, ND	2.2	3.8
Flat Scale Down, ND	4.3	5.5
Flat Scale Up, FD	0.6	5.6
Flat Scale Down, FD	0.6	2.4

Table 8.3: The percentage difference between the shifted and nominal predictions for the number of FD events due to deliberately applied miscalibrations. For the flat scale miscalibration applied to one detector, the largest overall effect from the flat scale down at the ND was taken as the systematic.

8.4 GENIE Simulation

Neutrino interactions in the NOνA simulation are generated using GENIE [71], a generator that involves a detailed physics modeling of cross sections, hadronization, and final state interactions. GENIE includes a plethora of parameters that alter individual physics input quantities, with the parameters themselves acting as systematic uncertainties, either dialing up or down a particular quantity by a standard deviation recommended by the GENIE authors.

The systematic uncertainties from physics modeling were evaluated by using the GENIE parameters as

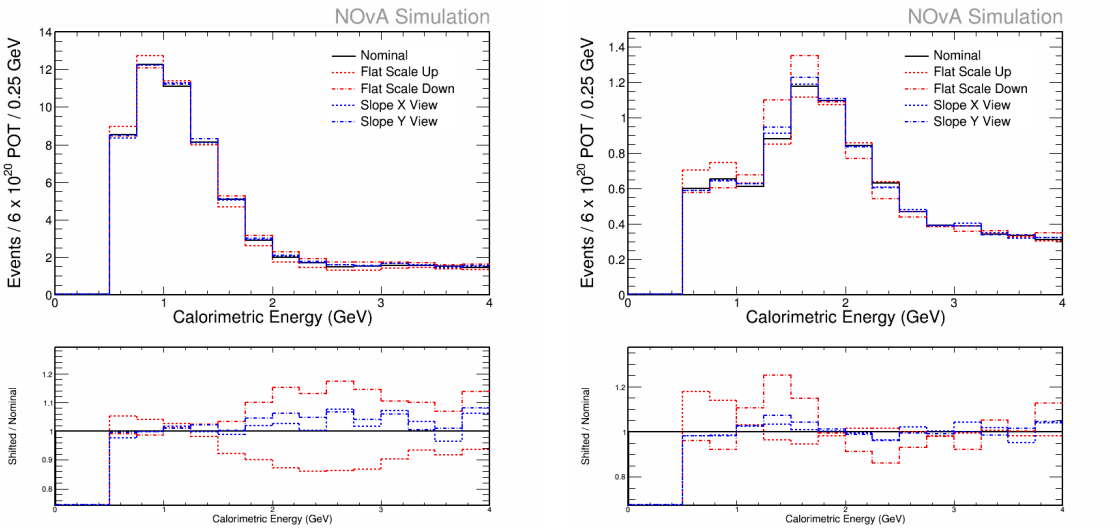


Figure 8.3: Shifted FD predictions due to miscalibration applied to both detector MC samples. The NC signal spectrum is on the left; the background spectrum is on the right.

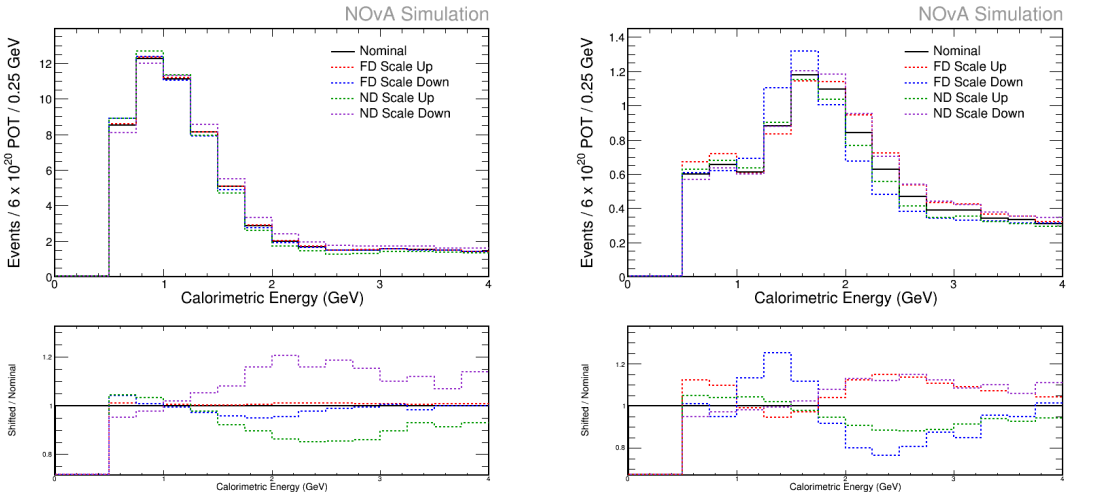


Figure 8.4: Shifted FD predictions due to miscalibration applied to one detector MC sample. The NC signal spectrum is on the left; the background spectrum is on the right.

event reweights. Nominal MC was produced using the parameters and weights as provided by the GENIE authors. These files also included a table of weights to modify an events 'worth' based on how much the different GENIE parameters shifted. For each parameter provided by GENIE, a shifted prediction was produced by applying the event reweights associated with the given parameter stored in the table to the nominal MC events. The percentage difference was taken as the systematic uncertainty for the particular parameter. Table 8.4 lists all of the parameters considered for the systematic study, and the associated uncertainty. Figure 8.5 shows the final systematic uncertainty envelope. The combined systematic uncertainty, adding the individual contributions in quadrature, was 1.6% for the NC signal and 4.8% for the background.

Parameter Description and Standard Deviation		NC Diff. (%)	Bkg. Diff. (%)
Axial mass for NC elastic	$\pm 25\%$	0.4	0.00
Strange axial form factor η for NC elastic	$\pm 30\%$	0.02	0.00
Normalization factor for CCQE	$^{+25\%}_{-15\%}$	0.00	1.3
Axial mass CCQE distribution shape	$^{+25\%}_{-15\%}$	0.0	0.4
CCQE Pauli suppression (via changes in Fermi level k_F)	$\pm 35\%$	0.00	0.6
Choice of CCQE vector form factors (BBA05 $\leftrightarrow$ Dipole)	D	0.00	0.3
Axial mass for CC resonance neutrino production	$\pm 20\%$	0.00	2.2
Vector mass for CC resonance neutrino production	$\pm 10\%$	0.00	1.1
Axial mass for NC resonance neutrino production	$\pm 20\%$	0.7	0.00
Vector mass for NC resonance neutrino production	$\pm 10\%$	0.2	0.00
Axial mass for CC and NC coherent pion production	$\pm 50\%$	0.4	0.3
Nuclear size parameter controlling π absorption in RS model	$\pm 10\%$	0.4	0.3
Non-resonance background in νp CC1π reactions	$\pm 50\%$	0.00	0.4
Non-resonance background in νp CC2π reactions	$\pm 50\%$	0.00	0.9
Non-resonance background in νn CC1π reactions	$\pm 50\%$	0.00	1.9
Non-resonance background in νn CC2π reactions	$\pm 50\%$	0.00	0.9
Non-resonance background in νp NC1π reactions	$\pm 50\%$	0.2	0.00
Non-resonance background in νp NC2π reactions	$\pm 50\%$	0.2	0.00
Non-resonance background in νn NC1π reactions	$\pm 50\%$	0.2	0.00
Non-resonance background in νn NC2π reactions	$\pm 50\%$	0.1	0.00
Non-resonance background in $\bar{\nu} p$ CC1π reactions	$\pm 50\%$	0.00	0.1
Non-resonance background in $\bar{\nu} p$ CC2π reactions	$\pm 50\%$	0.00	0.02
Non-resonance background in $\bar{\nu} n$ CC1π reactions	$\pm 50\%$	0.00	0.01
Non-resonance background in $\bar{\nu} n$ CC2π reactions	$\pm 50\%$	0.00	0.02

Parameter Description and Standard Deviation		NC Diff. (%)	Bkg. Diff. (%)
Non-resonance background in $\bar{\nu}p$ NC1π reactions	$\pm 50\%$	0.02	0.00
Non-resonance background in $\bar{\nu}p$ NC2π reactions	$\pm 50\%$	0.02	0.00
Non-resonance background in $\bar{\nu}n$ NC1π reactions	$\pm 50\%$	0.02	0.00
Non-resonance background in $\bar{\nu}n$ NC2π reactions	$\pm 50\%$	0.02	0.00
A_{HT} higher-twist parameter in BY model scaling variable ξ_w	$\pm 25\%$	0.01	0.2
B_{HT} higher-twist parameter in BY model scaling variable ξ_w	$\pm 25\%$	0.02	0.3
C_{V1u} u valence GRV98 PDF correction parameter in BY model	$\pm 30\%$	0.02	0.1
C_{V2u} u valence GRV98 PDF correction parameter in BY model	$\pm 40\%$	0.01	0.1
Pion transverse momentum (p_T) for $N\pi$ states in AGKY	D	0.01	0.02
Pion Feynman x (x_F) for $N\pi$ states in AGKY	D	0.01	0.02
Pion angular distribution in $\Delta \rightarrow \pi N$ (isotropic $\leftrightarrow$ RS)	D	0.1	0.3
Branching ratio for radiative resonance decays	$\pm 50\%$	0.02	0.01
Branching ratio for single-η resonance decays	$\pm 50\%$	0.1	0.2
Nucelon mean free path (total rescattering probability)	$\pm 20\%$	0.2	0.4
Nucleon charge exchange probability	$\pm 50\%$	0.1	0.2
Nucleon elastic reaction probability	$\pm 30\%$	0.1	0.4
Nucleon inelastic reaction probability	$\pm 40\%$	0.1	0.2
Nucleon absorption probability	$\pm 20\%$	0.05	0.2
Nucleon π-production probability	$\pm 20\%$	0.04	0.1
π mean free path (total rescattering probability)	$\pm 20\%$	0.4	0.7
π charge exchange probability	$\pm 50\%$	0.1	0.2
π elastic reaction probability	$\pm 10\%$	0.1	0.1
π inelastic reaction probability	$\pm 40\%$	0.6	0.3
π absorption probability	$\pm 20\%$	0.6	0.4
π π-production probability	$\pm 20\%$	0.03	0.1
MEC event scale	$+50\%$	0.00	1.6
RPA event weight	Varied	0.00	1.1
Combined		1.6	4.8

Table 8.4: The systematic uncertainty, in percentage difference, for each GENIE systematic parameter. The description and standard deviations come from reference [71, 73]. Parameters with a standard deviation marked as 'D' are discrete and thus do not have a standard deviation estimate.

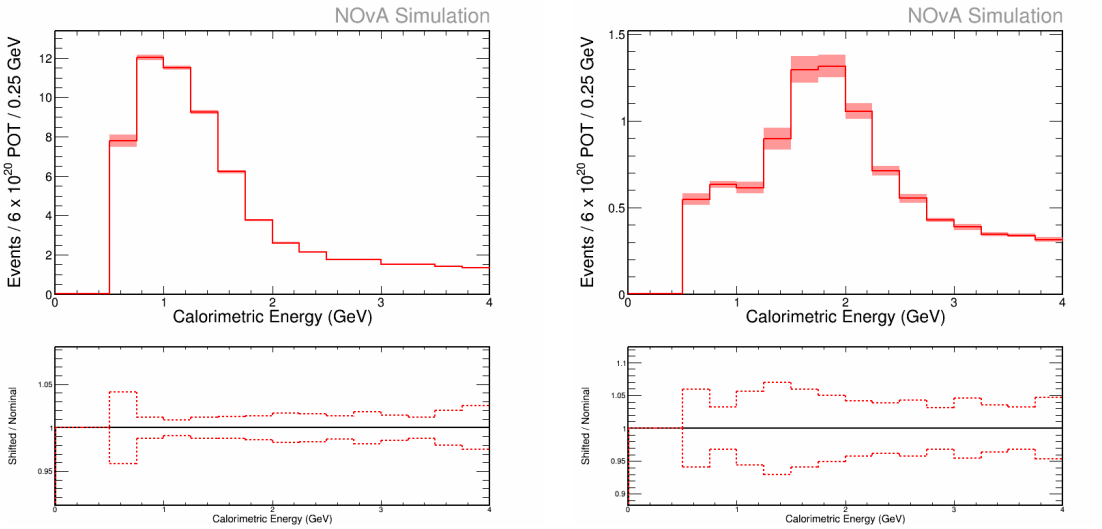


Figure 8.5: GENIE systematics uncertainty envelope on the NC signal (left) and background (right) event spectra, after extrapolation. The envelope was calculated by adding in quadrature the larger of $|+1\sigma|$ and $|-1\sigma|$ for each individual systematic.

8.5 ND Containment

The ND does not see an effective point source of neutrinos like the FD due to its proximity to the beam source. As a result, the neutrino flux is not uniform across the detector, and so the energy spectrum of neutrinos seen by the two NOνA detectors is slightly different. To study the effect this has on the extrapolated prediction, multiple predicted FD spectra were generated using subsamples of the ND. The fiducial volume of the ND was split in half along each axis, and split into an inner and outer half (the overall containment criteria was left in tact). The extrapolation was performed using each of these 'half detectors.' The uncertainty was taken as the percentage difference from the shifted prediction to the nominal. Table 8.5 shows the results from each of these extrapolations and figure 8.6 show the shifted spectra. Like the Birks-Chou systematic, the largest overall difference was taken as the systematic uncertainty for a 1.0% effect on the NC signal and a 0.6% effect on the background.

8.6 ND Rock Event Contamination

The MC simulation does include neutrino interactions that occur in the rock that surrounds the ND. These events often leak into the detector volume, and while most of them are cut away by fiducial and con-

ND Half	NC Difference (%)	Background Difference (%)
West (+X)	0.3	0.01
East (-X)	0.2	0.1
Top (+Y)	0.3	0.01
Bottom (-Y)	0.3	0.01
Front (Low Z)	0.1	0.3
Back (High Z)	0.1	0.4
Inner (Low |X|, |Y|)	0.8	0.02
Outer (High |X|, |Y|)	1.0	0.6

Table 8.5: The percentage difference between the shifted and nominal predictions for the number of FD events after extrapolation using half of the fiducial volume at the ND.

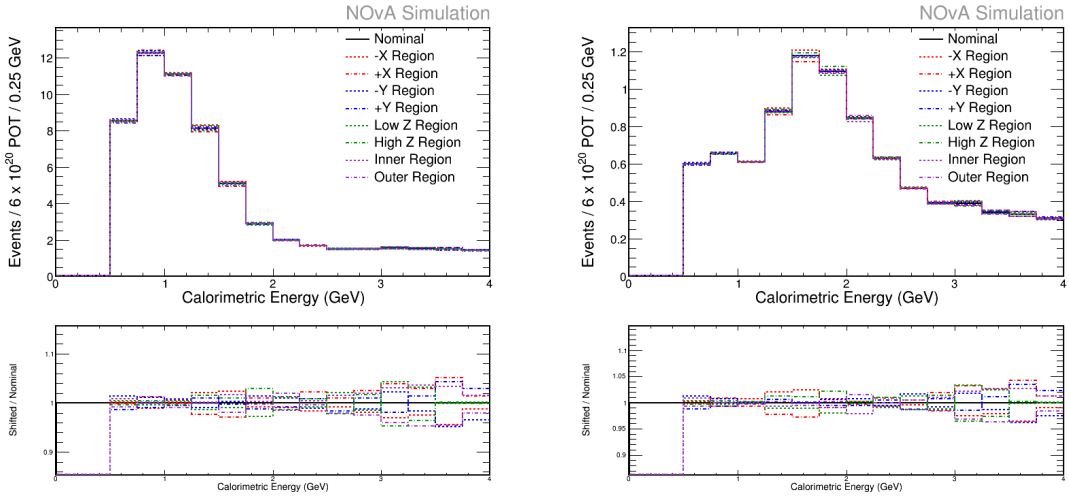


Figure 8.6: Shifted FD predictions after extrapolation using only half of the ND fiducial volume. The NC signal spectrum is on the left; the background spectrum is on the right.

tainment cuts, there are some that remain. Those events that do remain cannot be reconstructed properly as their origins are outside of the detector.

The systematic error that is incurred due to the rock event contamination was estimated by predicting the FD event spectrum from an extrapolation with rock events removed by MC truth and comparing to the nominal predicted spectrum. The events were only removed from the ND MC sample, requiring that the true neutrino vertex was inside the detector to remain. The shifted spectra are shown in figure 8.7. This systematic amounted to an overall 4.1% shift on the NC signal spectrum and a 1.7% shift on the background spectrum.

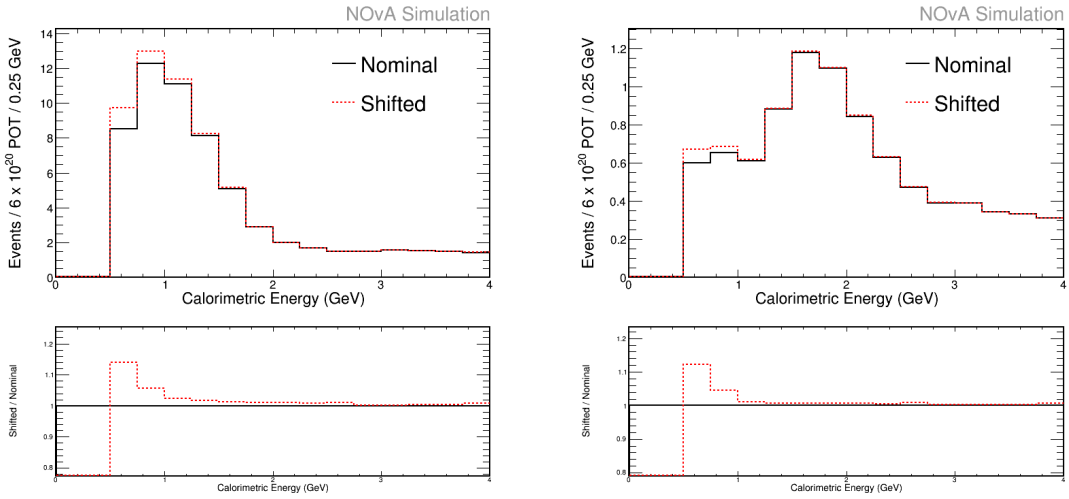


Figure 8.7: Shifted vs nominal spectra for the NC signal (left) and background (right). The shifted spectra are extrapolated after removal of rock events in the ND MC by truth.

8.7 ND Data/MC Difference and CC Background

To assess the results over the MC simulation, data was compared to MC at the ND, and the event spectra were found to be an imperfect match; see figure 8.8. The FD prediction is made by decomposing ND data into NC, ν_μ CC and ν_e CC components proportional to the amounts in the MC, so a data/MC difference can cause a systematic effect. This error was treated as a hybrid decomposition and CC background error, and as such, the uncertainty was quantified using a hybrid method.

For the first NOνA ν_μ disappearance and ν_e appearance analyses, the two groups used different methods to estimate their background uncertainties. The ν_μ analysis noted that their backgrounds were small and thus placed a 100% uncertainty on each [44]. The ν_e analysis assigned any discrepancy between ND data and MC to each of the three ND background components, used each shifted decomposition to perform an extrapolation, compared the shifted predictions to nominal, and took the shift with the largest difference as the systematic uncertainty [43].

The NC disappearance analysis used a combination of the above techniques to evaluate the ND data/MC and CC background systematic uncertainty. The background from the beam ν_e component was calculated to be small, so this component was allowed to vary by $\pm 100\%$. The shifted spectrum was used for extrapolation, and the percentage difference between the nominal and shifted predictions was taken as the

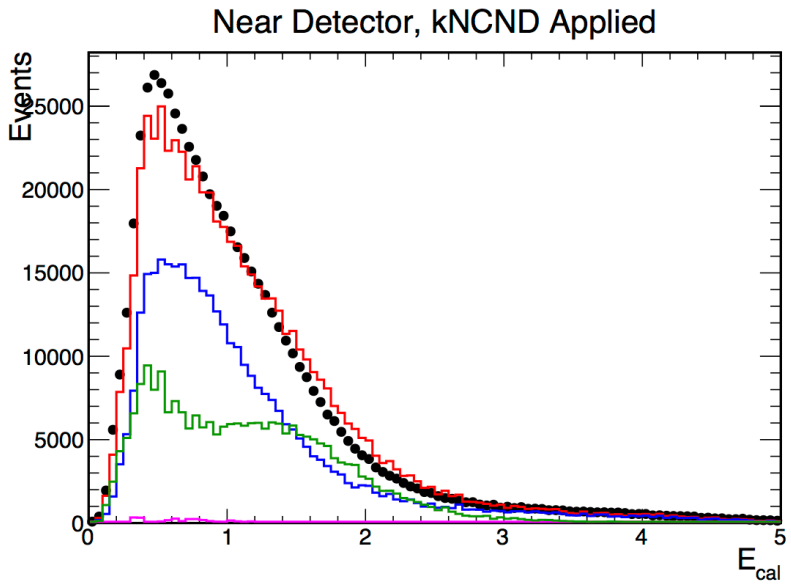


Figure 8.8: Energy spectra for data and MC at the ND, from reference [130].

uncertainty. For both the NC and ν_μ CC components, the uncertainty was taken following the same procedure as the ν_e analysis. Results from both these two studies were then combined in quadrature. Shifting the beam ν_e component up and down produced the same net change in event numbers, albeit in opposite directions. Assigning the data/MC difference to the ν_μ CC component caused the larger overall difference, compared to assignment to the NC component. The overall uncertainty was thus calculated by combining in quadrature the uncertainties from assigning the data/MC difference to the ν_μ CC component and shifting the beam ν_e component by $|100\%|$, for a result of a 7.0% effect on the NC signal and a 10.4% effect on the CC background. Table 8.6 summarizes the results from these studies, figure 8.9 shows the shows the changes in the predicted spectrum due to each systematic shift.

8.8 MC Statistics

In a perfect world, there would be enough MC statistics that this section would be unnecessary, but alas, a perfect world this is not. To estimate the systematic uncertainty due to MC statistics, the MC was split into five uniformly sized samples, and each was used to extrapolate the same set of ND data. One of the resultant predicted FD spectra was labeled as nominal, and the other four were compared to this. The bin

Evaluation Method	NC Difference (%)	Background Difference (%)
Beam ν_e $+100\%$	0.00	8.2
Beam ν_e -100%	0.00	8.2
ND Data/MC difference assigned to NC component	5.2	7.7
ND Data/MC difference assigned to ν_μ CC component	7.0	6.4
Overall	7.0	10.4

Table 8.6: Systematic uncertainty covering ND data/MC discrepancy and CC backgrounds. The percentage differences were calculated by shifting the ND decomposition using the listed evaluation method, performing an extrapolation, and comparing the predicted results to nominal.

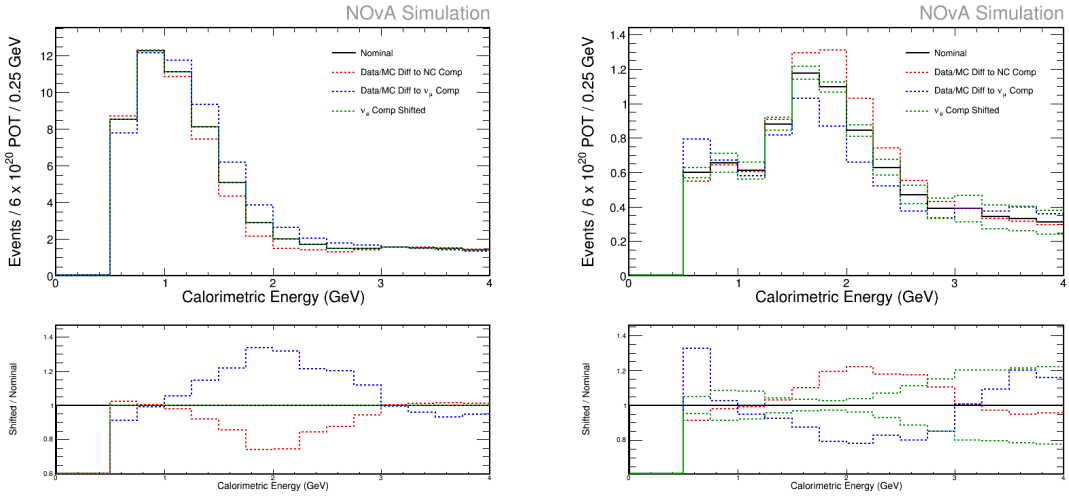


Figure 8.9: Shifted FD predictions from studies to quantify ND data/MC difference and CC background uncertainties.

by bin differences were taken as the uncertainty between the nominal spectrum and the 'shifted' spectrum. To come up with an overall uncertainty, all of the uncertainties were added in quadrature and the result was divided by the square root of the number of samples, or $\sqrt{4} = 2$. The result was a 2.0% uncertainty on the NC signal spectrum and a 4.8% uncertainty on the background spectrum.

8.9 Noise Model

A sample of MC was generated for the FD using a different model of noise to study the systematic uncertainty caused by the noise model. To quantify the uncertainty, the predicted FD event spectra from the

nominal and alternative MC samples were compared against each other after applying normal three flavor oscillation weights. The difference seen was 0.9% for the NC signal, and 0.7% for the background. As the shifted sample was only generated for the FD MC and the uncertainty was comparatively small, this systematic was added into the overall normalization.

8.10 Overall Normalization

Several independent effects contributed to an overall normalization systematic uncertainty. A 0.5% uncertainty on the POT counting came from a small difference in the two toroids that determine the POT in a spill [79]. Uncertainties in the masses of the various parts of the NOνA detectors contributed another uncertainty of 0.7% [131]. A study of the reconstruction efficiency between ND data and MC showed a 2.9% difference between the number of events that get fully reconstructed [132], which was taken directly as a contribution to the normalization uncertainty. Finally, studies of the interaction rate after overlaying cosmic trigger data over MC beam interactions suggested a rate difference of 3.7%. These four effects and the noise systematic were combined in quadrature and constituted a systematic uncertainty of 4.9% on the NC signal and 4.9% on the background.

8.11 Systematic Uncertainty Summary

Table 8.7 shows a summary of all of the systematics, as well as an overall uncertainty. The overall uncertainty was calculated by summing the uncertainty from each row in quadrature. The final systematic uncertainty on the NC signal is 12.2% and the uncertainty on the background is 15.3%.

All of the systematics discussed to this point were also calculated at the ND. Many of these, including the Beam, Birks-Chou, and GENIE, are much larger without mitigation through extrapolation. The combined effects were studied to be certain that any data/MC discrepancy was covered under the full ND systematic uncertainty band. The data/MC spectra are shown in figure 8.10. Note that the systematic uncertainty that assesses any data/MC difference to a single component does not contribute to the full band. The technique used for this systematic did not change the total number of events since both the nominal and shifted spectra respected the normalization from the data. As the systematic uncertainty band was

Systematic	NC Difference (%)	Background Difference (%)
Beam	3.4	3.6
Birks-Chou	2.4	1.8
Calibration	5.8	6.0
GENIE	1.6	4.8
ND Containment	1.0	0.6
ND Rock Contamination	4.1	1.7
ND Data/MC	7.0	10.4
MC Statistics	2.0	4.8
Overall Normalization	4.9	4.9
Combined	12.2	15.3

Table 8.7: A summary of the individual systematic uncertainties for the NC disappearance analysis. The uncertainties are percentage differences between the nominal and shifted predicted spectra.

calculated using all events, by construction this data driven systematic effect does not have an effect on the total uncertainty band.

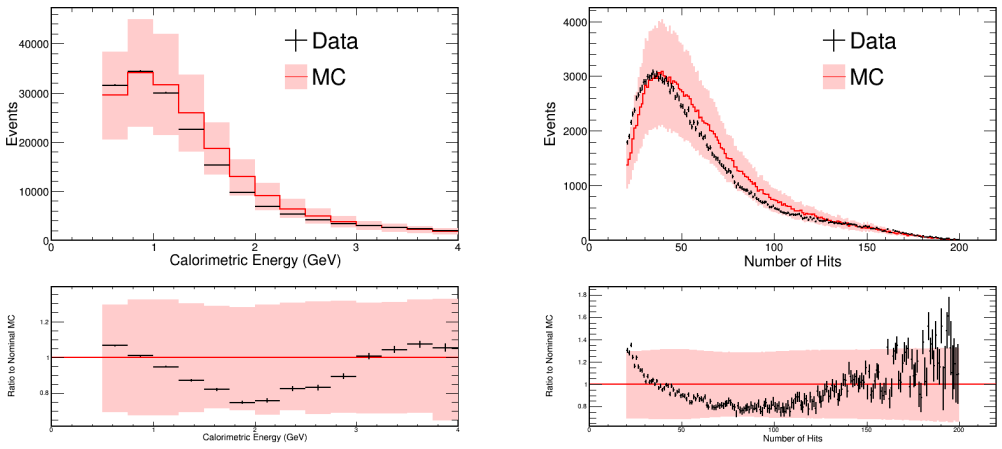


Figure 8.10: ND Data/MC spectra with full systematic uncertainty band. The calorimetric energy spectrum is shown on the left, the number of hits spectrum is shown on the right.

9

Analysis Results

The analysis presented in this dissertation was designed to make several measurements. As this dissertation was written alongside the first NOνA NC analysis, one primary goal of the analysis was to demonstrate the successful ability of NOνA to identify NC events and measure the ratio of observed to predicted NC events in a 3 flavor hypothesis, R_{NC}.

$$R_{NC} \equiv \frac{N_{Obs} - N_{Pred}^{Bkg}}{N_{Pred}^{NC}} \tag{9.1}$$

The other (arguably more important) goal was to start contributing to the global data on the sterile mixing parameters by extracting measurements on the mixing angles θ_{24}, θ_{34} and matrix elements $|U_{\mu 4}|^2$ and $|U_{\tau 4}|^2$.

Data for this analysis were collected between February 2014 and May 2016. For the FD, this corresponds to 6.69×10^{20} POT, or 6.05×10^{20} full detector equivalent POT. For the ND, 3.72×10^{20} POT was collected.

The first section in this chapter describes the fitting methods used to make the mixing angle and matrix element measurements. Before presenting the results, the next two sections discuss preliminary tests to demonstrate satisfactory performance of the analysis. The final section in this chapter presents the ultimate results.

9.1 Fitting Method

It was decided to make rate only measurements for the first NOνA NC disappearance analysis. Thus, even though the extrapolation and prediction were performed in 250 MeV bins of calorimetric energy, only the integrated event totals were input to the fitting framework.

Fitting was performed allowing some parameters to float, some held to current best fit values, and others set to 0 due to insensitivity. As mentioned above, a major analysis goal was to measure θ_{24} and θ_{34}, so these angles were allowed to float between $0°$ and $45°$. Values outside of this range are either equivalent to this range through redefinition of the angles, or already highly disfavored through previous experiments and in a region difficult for fitting due to many local minima. θ_{23} was allowed to float with a Gaussian constraint with mean $45.8°$ and standard deviation $3.2°$, the best fit measurement by T2K [133]. All of the remaining parameters were held fixed to the values listed in table 9.1.

Parameter	Value
Δm^2_{21}	7.53×10^{-5} eV^2
Δm^2_{32}	2.37×10^{-3} eV^2
$\sin^2 2\theta_{12}$	0.846
$\sin^2 2\theta_{13}$	0.085
θ_{14}	0
δ_{13}	0
δ_{24}	0
δ_{14}	0

Table 9.1: The parameters held fixed and the values they were set at for fitting.

The best fit values were found by calculating the minimum χ^2 value and allowing individual systematics

to float with a penalty,

$$\chi^2 = 2\left(N_{Pred} - N_{Obs} + N_{Obs}\ln\frac{N_{Obs}}{N_{Pred}}\right) + \left(\frac{\theta_{23}-\mu_{23}}{\sigma_{23}}\right)^2 + \sum_{i=1}^{N}\left(\frac{\sigma_i^{BF}}{\sigma_i}\right)^2, \qquad (9.2)$$

where μ_{23} and σ_{23} are the mean and standard deviation used in the Gaussian constraint on the value of θ_{23}, σ_i^{BF} is the number of standard deviations the ith systematic error is shifted in the best fit, and σ_i is one standard deviation for the ith systematic error. Equation 9.2 is the standard formula for calculating χ^2 with penalty terms, where the first term is the basic value for χ^2, and the second and third terms add a unit for each shift of θ_{23} and the systematic errors by their respective variances. For one and two dimensional χ^2 surfaces, the remaining angles not shown were profiled.

9.2 Robustness Studies

Before diving head first into the FD data, several studies were performed to test the robustness of the analysis. These included data/MC comparisons of ND distributions, the effect of shifting event energies, and the effect of δ_{24} on the event distributions.

The first necessary analysis benchmark was ND data/MC comparisons. Figure 9.1 shows the event energy distributions and figures 9.2, 9.3, and 9.4 show the distributions for all of the variables used for selection discussed in chapter 6.

While most of the distributions differ by at most a small normalization, a few have an apparent shift, notably the energy and number of hits distributions. Discrepancies in these particular distributions were actually expected due to a known mismodeling of NC interactions. As discussed in section 4.2.1, the newly simulated MEC event type was only added for CC events. Furthermore, there are large uncertainties in the cross sections for the resonant and DIS events. Improvements to these cross section models and inclusion of NC MEC events could largely explain the data/MC differences seen in the energy and number of hits distributions.

Two sanity checks were performed to ensure that the data/MC difference would not negatively impact the analysis. First, the ND data/MC energy distribution was shown with a full systematic uncertainty

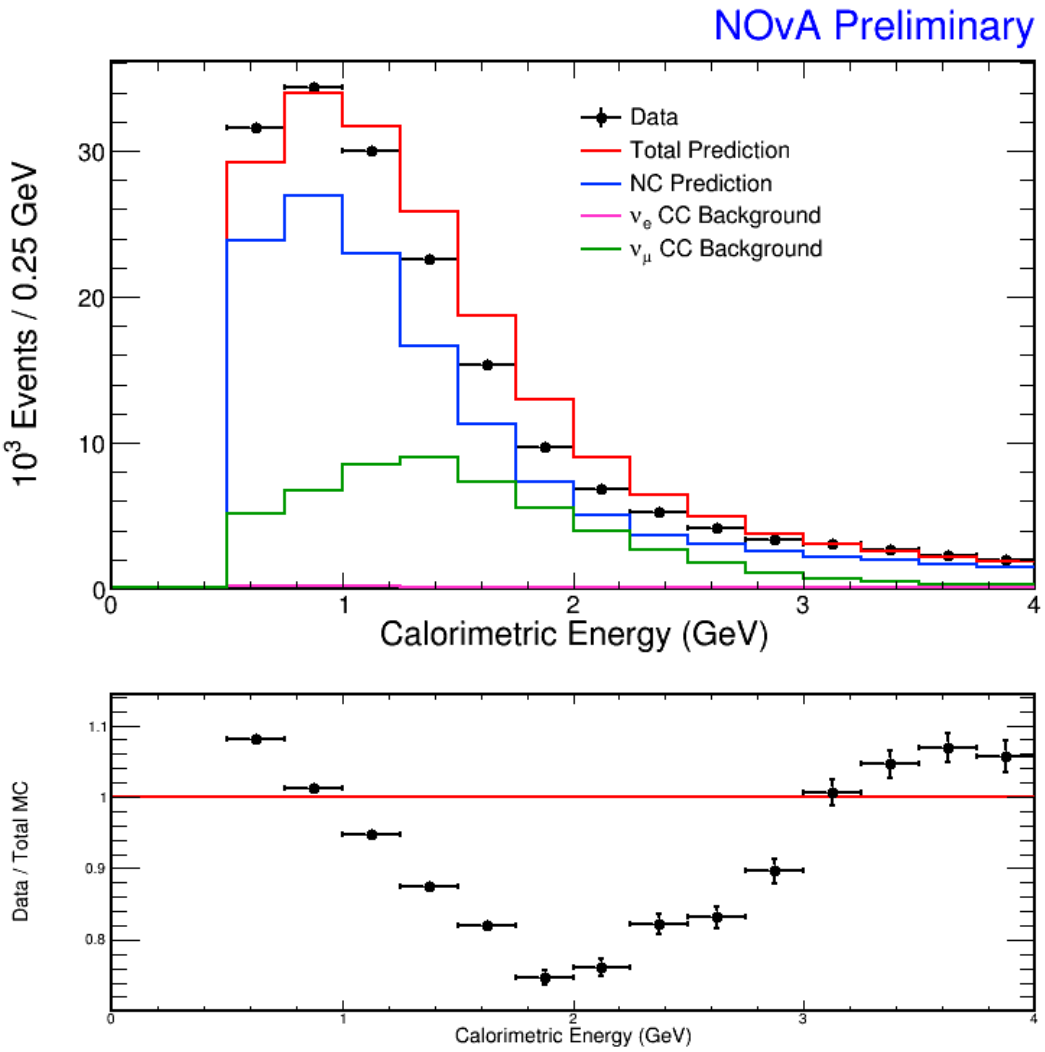


Figure 9.1: ND data/MC comparison of the calorimetric energy.

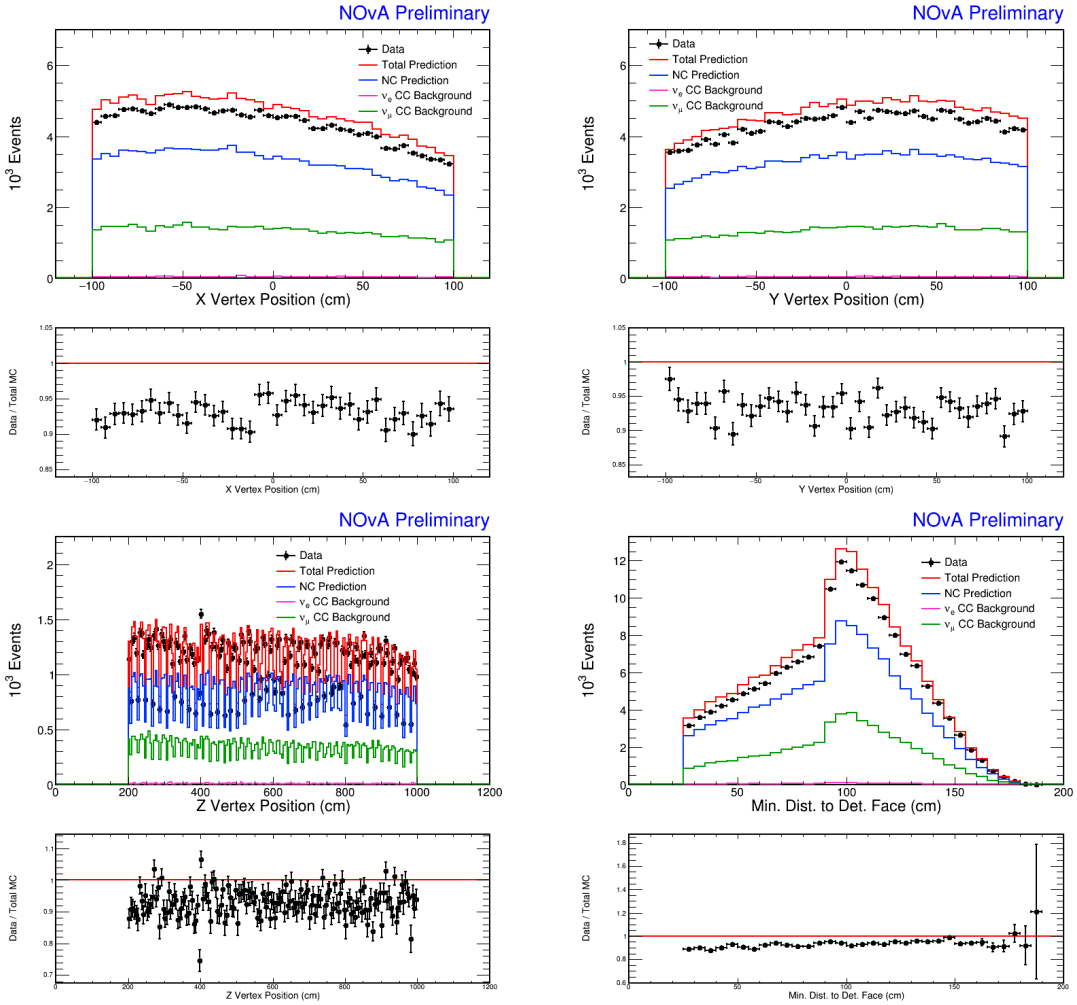


Figure 9.2: ND data/MC comparison of the reconstructed vertex position and containment distributions.

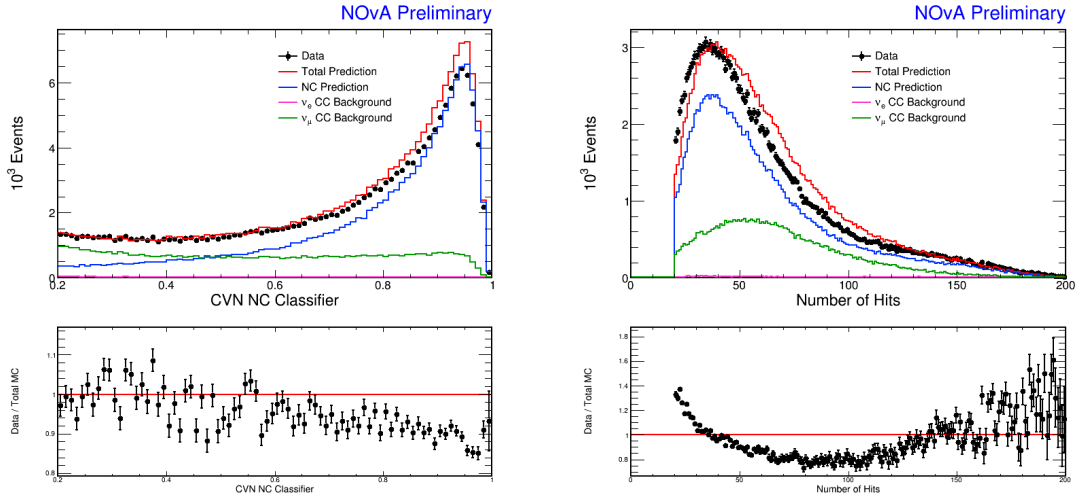


Figure 9.3: ND data/MC comparison of the CVN and number of hits distributions.

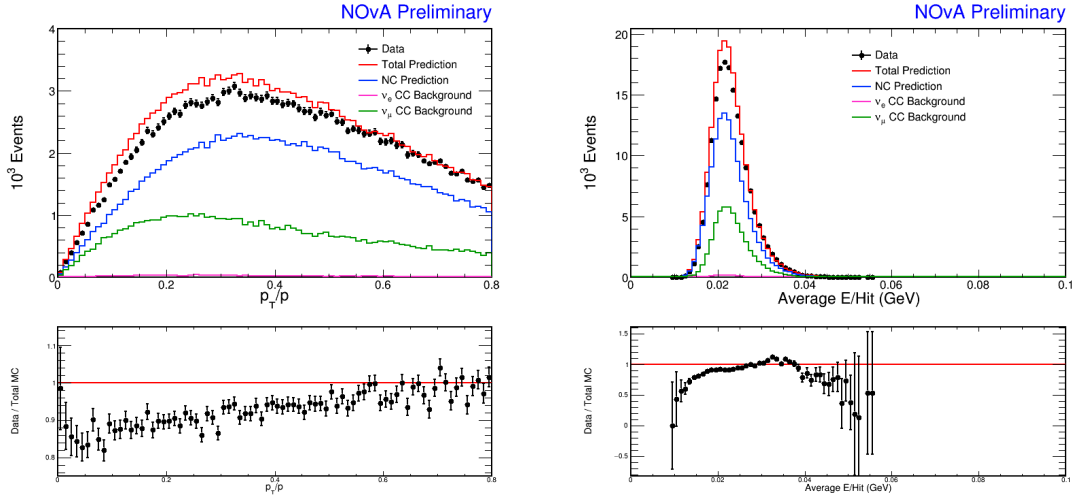


Figure 9.4: ND data/MC comparison of particle transverse momentum and energy per hit distributions.

band to show the difference is covered by systematics. This was shown in figure 8.10. Secondly, the effect of shifting either the data or MC in relation to the other on the predicted event rate and 1D angle sensitivities was studied. This was done by shifting the data up or the MC down by the ratio of the means of the energy distributions for MC and data, $1.39\,\text{GeV}/1.33\,\text{GeV} \approx 4.5\%$ and also for a much larger shift of 10%. The event counts are shown in table 9.2 and the angle sensitivities are shown in figure 9.5. Due to the nature of a counting experiment, since the overall event rates are essentially unchanged, the effect of these rather large shifts is negligible. As a result, it was decided to take the data and MC as is and push for model improvements for future analyses.

Shift	All	NC	ν_μ CC	ν_e CC	Cosmic	FOM
Nominal	83.5	60.6	4.6	3.6	14.3	6.633
Data 4.5% Up	84.6	61.0	4.7	3.7	14.9	6.634
Data 10% Up	86.2	61.4	4.9	3.7	15.8	6.612
MC 4.5% Down	85.1	61.8	4.8	3.7	14.3	6.702
MC 10% Down	87.2	63.4	5.2	3.8	14.3	6.792

Table 9.2: The number of predicted events at the FD after applying an energy shift to data or MC.

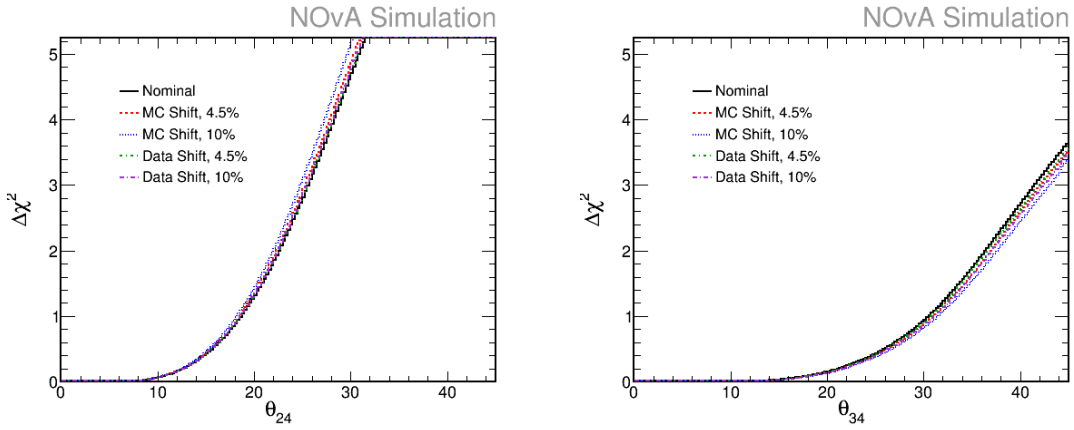


Figure 9.5: Angle sensitivities after applying an energy shift to either MC or data. Left: θ_{24}, Right: θ_{34}. The fits are statistics only and do not include any systematic uncertainties.

The final robustness study considered the effect of δ_{24} on the analysis. Table 9.3 shows the predicted number of events for various values of δ_{24}. For $\theta_{34} = 35°$, $\theta_{24} = 10°$, changing δ_{24} from 0 to one of its extrema can have up to an 8% effect on the absolute event rate. Comparing to the results from the energy

shifting studies, this could affect the sensitivity to the real mixing angles by $2-3^\circ$ at the 90% confidence level. However, the observed event rate must be below the predicted number of events under the 3 flavor model in order to have any effect, otherwise the best fit values for the real mixing angles go to 0 and the effect of varying δ_{24} vanishes. Indeed, when the fitting procedure was tested against a spectrum created under the 3 flavor model, there was no appreciable difference in the fit results between setting δ_{24} to 0 or profiling it, as shown in figure 9.6. Thus, δ_{24} remained set to a single value to remove a fit variable, saving computation time.

δ_{24}	All	NC	ν_μ CC	ν_e CC	Cosmic
0	64.3	41.7	4.8	3.1	14.3
$\pi/2$	59.5	37.4	4.5	2.9	14.3
π	64.5	42.4	4.3	3.0	14.3
$3\pi/2$	69.4	46.7	4.5	3.1	14.3

Table 9.3: The number of predicted events at the FD for selected values of δ_{24}. Oscillations parameters were held fixed to the values listed in table 9.1, with $\theta_{23} = 45^\circ$, $\theta_{34} = 35^\circ$, and $\theta_{24} = 10^\circ$.

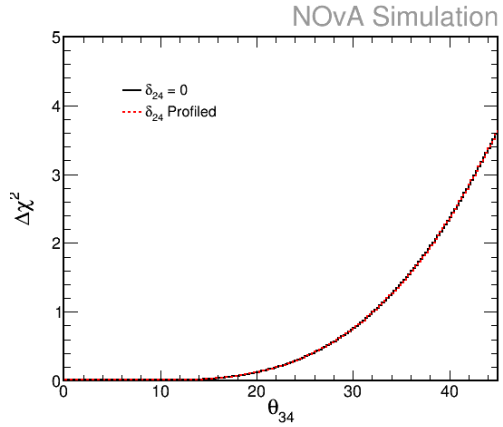


Figure 9.6: Sensitivity to θ_{34} when profiling δ_{24} versus leaving it set to 0. The prediction was fit to a spectrum generated under the 3 flavor model. The sensitivity to θ_{24} shows the same level of agreement between the two fit methods. The fits are statistics only and do not include any systematic uncertainties.

9.3 Sideband Studies

Three different sidebands were studied to test the performance of the analysis, a high energy sideband, a low CVN sideband, and mid cosmic BDT sideband. The predicted and observed event rates for each

sideband are shown in table 9.4.

Sideband	Data	Total MC	NC	ν_μ CC	ν_e CC	Cosmic
High Energy	15	8.2	6.5	0.8	0.5	0.1
Low CVN	35	32.7	3.7	5.5	19.3	4.0
Mid BDT	17	14.5	6.0	0.5	0.4	7.5

Table 9.4: The observed and predicted events at the FD for each sideband.

The high energy sideband applied the standard selection and considered events between 4 and 6 GeV, a region chosen due to its high purity of NC events. The predicted and observed event distributions are shown in figure 9.7, 8.2 events were predicted and 15 ± 4 were observed in data. While the observed rate is slightly high, the low statistics means that the observation is within 2σ of the prediction. Furthermore, the discrepancy is largely driven by a single bin rather than a systematic offset. This result was thus interpreted as validation for the general analysis procedure.

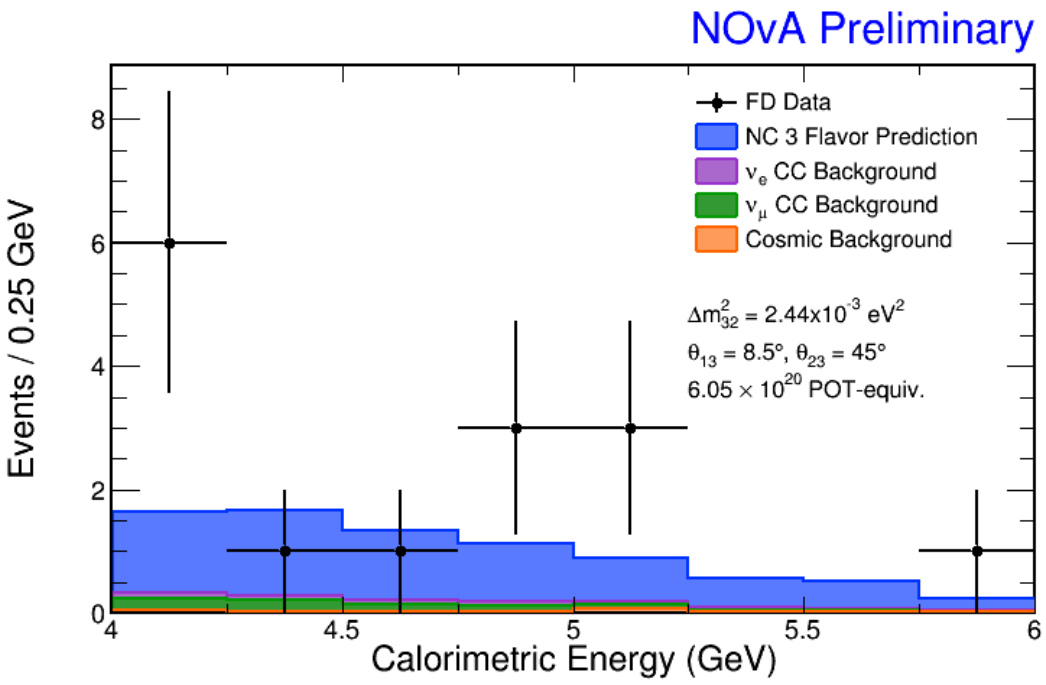


Figure 9.7: The observed and predicted FD event rates for the high energy sideband.

The low CVN sideband considered events that fail the CVN cut, or those with a CVN NC score below 0.2. This region was used as validation for the CVN selector. The predicted and observed event distribu-

tions are shown in figure 9.8, 32.7 events were predicted and 35 ± 6 were observed in data. The overall agreement in both shape and rate for this sideband provided great confidence in the performance of CVN.

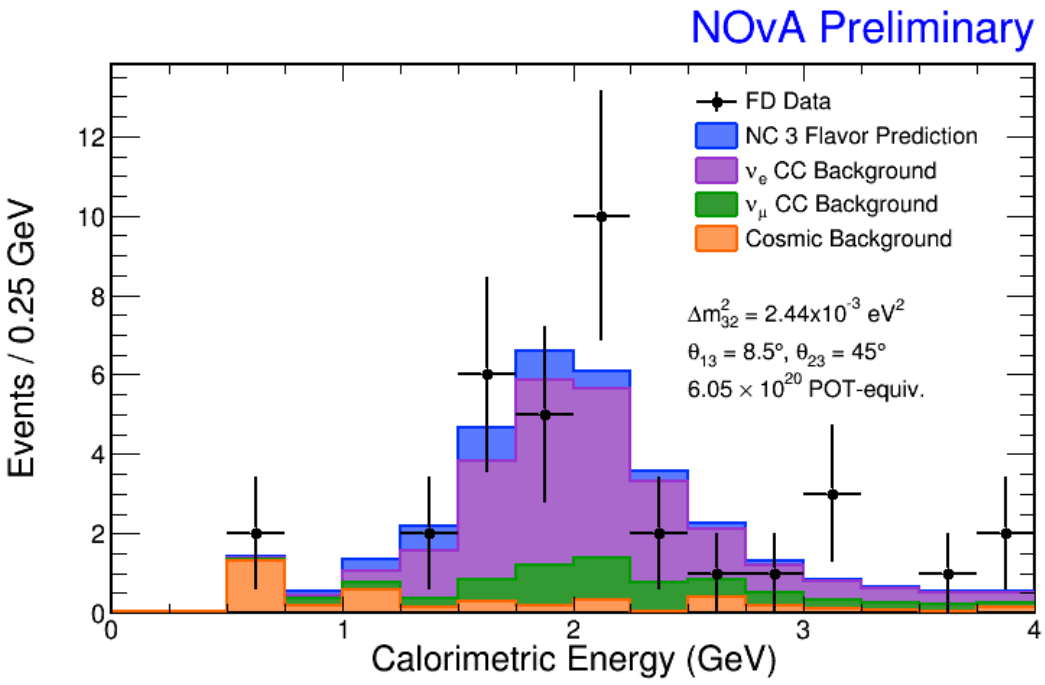


Figure 9.8: The observed and predicted FD event rates for the low CVN sideband.

The mid cosmic BDT sideband considers events with a BDT score between 0.42 and 0.5, a region which fails the standard selection cuts but still has NC events. The predicted and observed event distributions are shown in figure 9.9, 14.5 events were predicted and 17 ± 4 were observed in data. This sideband also showed excellent agreement in both shape and rate, providing further validation for the general analysis procedure, and specifically for the main cosmic rejection variable.

9.4 Results

This analysis predicted that there would be $83.5 \pm 0.8(\text{stat.})^{+10.9}_{-7.2}(\text{syst.})$ NC-like events selected in the FD data. 95 ± 9.7 were observed, corresponding to a measurement of $R_{NC} = 1.19 \pm 0.16(\text{stat.})^{+0.10}_{-0.14}(\text{syst.})$ If active neutrinos mixed with sterile neutrinos, the number of observed events would be depleted, causing R_{NC} to be less than one. Thus, the measurement of R_{NC} is consistent with the no sterile mixing

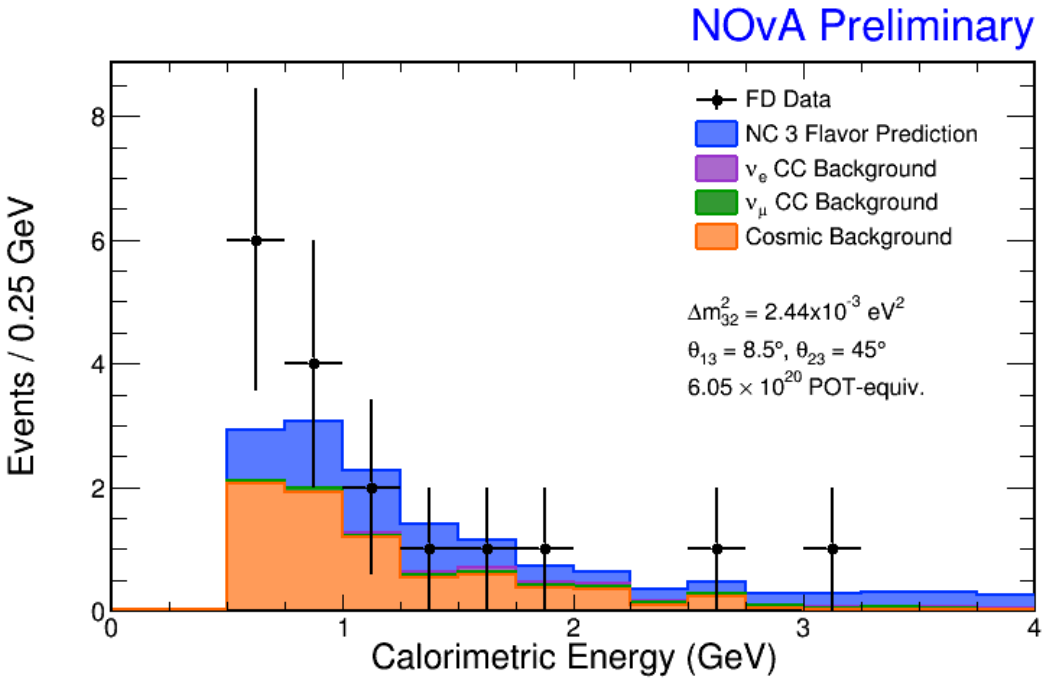


Figure 9.9: The observed and predicted FD event rates for the mid cosmic BDT sideband.

hypothesis.

The energy distribution of the predicted and observed events is shown in figure 9.10. The χ^2 of this distribution is 23.31 for 14 bins, or 1.665/bin. To better assess the understanding of the observed data, data/MC comparisons distributions were studied for many of the variables used for selection, including regions cut from the analysis. These can be seen in figures 9.11, 9.12, and 9.13. All of the distributions showed excellent agreement, especially within the CVN classifier. Lastly, the event vertices were well dispersed throughout the entire FD volume, as shown in figure 9.14.

Oscillation parameter measurements were extracted from a fit to the data, using the methods described in section 9.1. The two dimensional 68% and 90% confidence levels are shown in figure 9.15. One dimensional slices of θ_{24} and θ_{34} were constructed from the full fit by profiling over the remaining angles and can be seen in figure 9.16. The best fit values were within an arc second of 0°, and from the one dimensional slices the mixing angles were limited to $\theta_{34} < 35°$ and $\theta_{24} < 21°$ at 90% confidence. The θ_{34} limit is already competitive with the MINOS 2011 result [31], also shown in figure 9.15.

The data were also analyzed using a Feldman-Cousins (FC) statistical correction [134]. The basic pre-

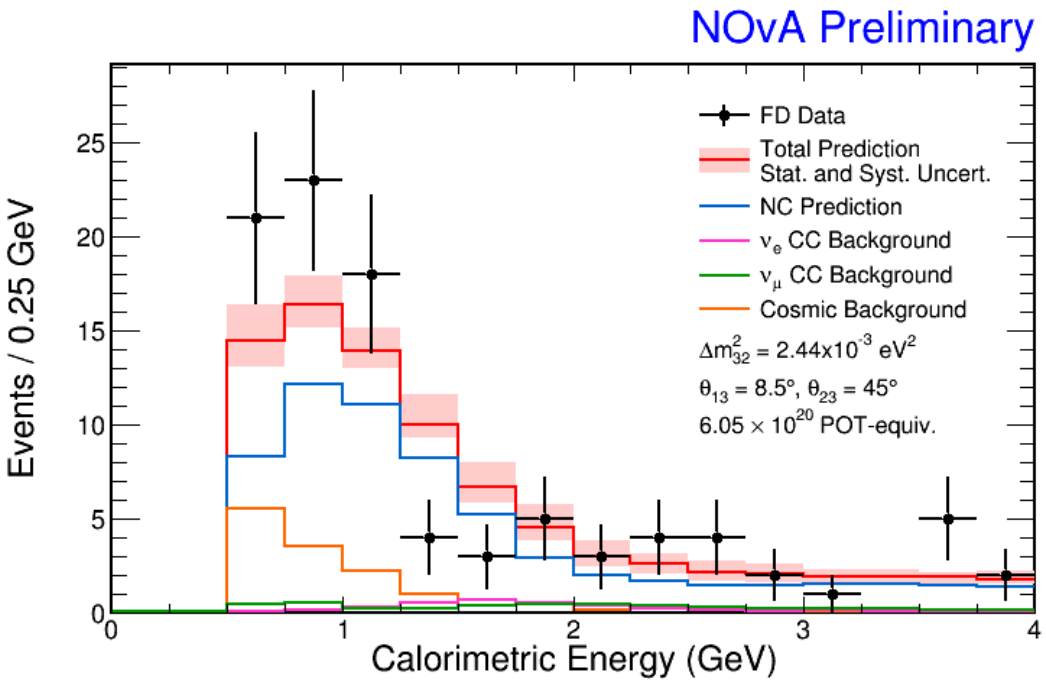


Figure 9.10: Calorimetric energy distribution of events predicted and observed in the FD.

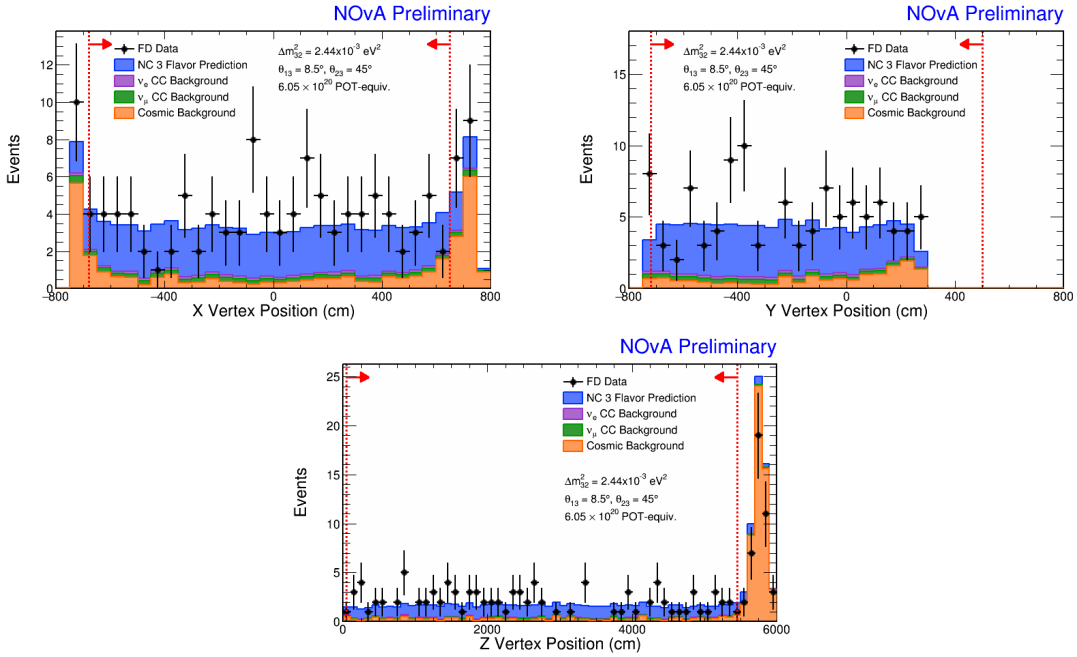


Figure 9.11: FD data/MC comparison of the reconstructed vertex position distributions. The dashed lines and arrows indicate the regions kept for the analysis.

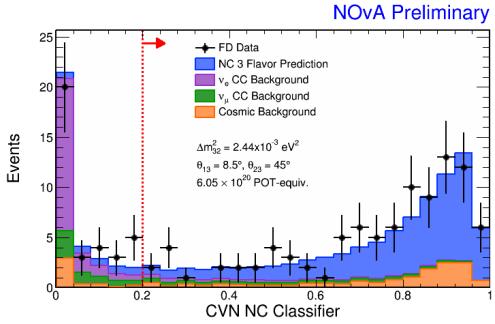


Figure 9.12: FD data/MC comparison of the CVN distribution. The dashed line and arrow indicate the region kept for the analysis.

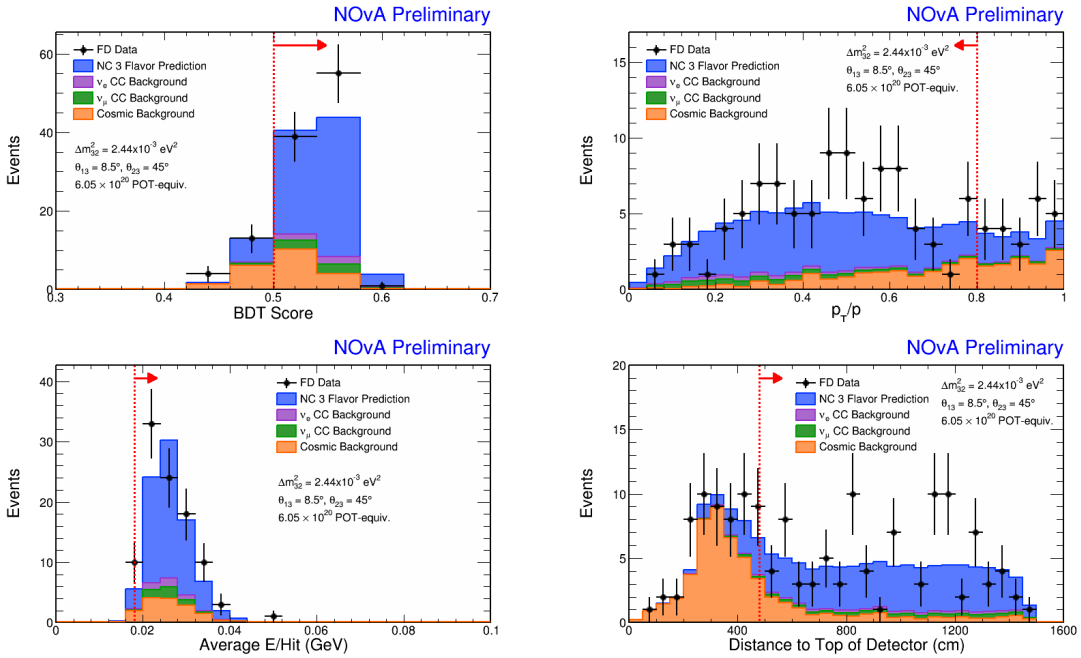


Figure 9.13: FD data/MC comparisons of the variables used for cosmic rejection. The dashed lines and arrows indicate the regions kept for the analysis.

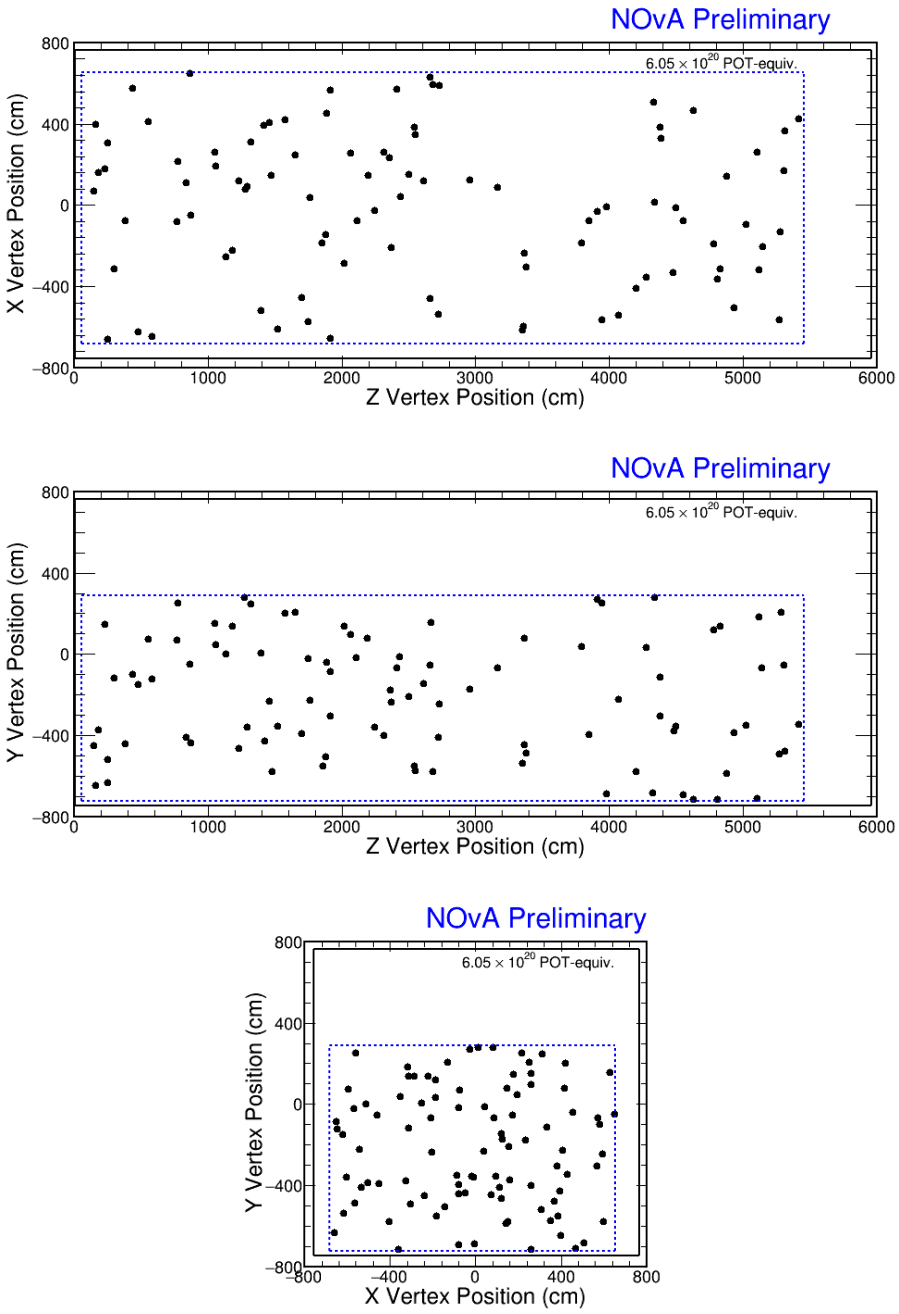


Figure 9.14: Distributions of the selected data event vertices in the detector. The solid black outline shows the full detector volume; the dashed blue line shows the fiducial and contained region. Top: XZ view. Middle: YZ view. Bottom: XY view.

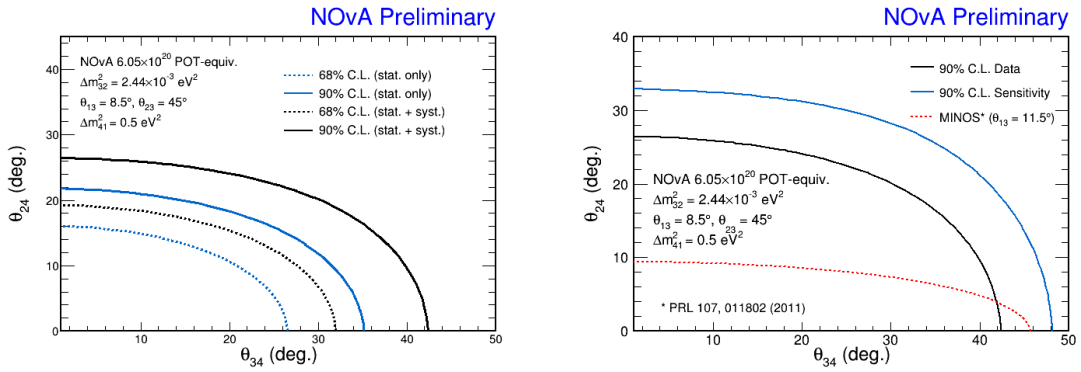


Figure 9.15: Two dimensional contours of θ_{34} vs θ_{24}. The left figure shows the 68% and 90% confidence limits. The right figure shows the 90% confidence limit with a comparison to the MINOS 2011 results [31].

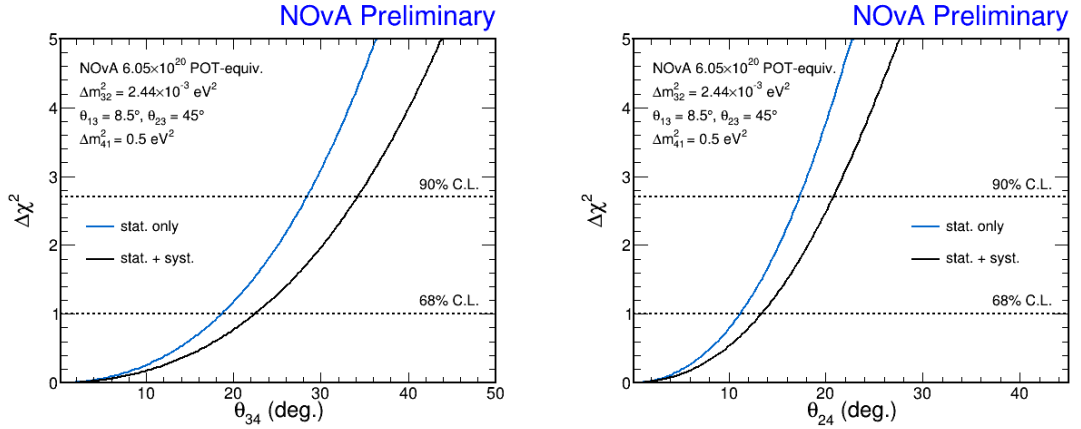


Figure 9.16: One dimensional slices of θ_{34} and θ_{24}. These distributions were constructed by profiling over the remaining angles from the fit. Left: θ_{34}. Right: θ_{24}.

mise of this technique is to extract confidence intervals from an ordered likelihood ratio, instead of directly from $\Delta\chi^2$ values. The FC method performs particularly well near physical boundaries, such as angles close to maximum or minimum. The likelihood ratio, R, is calculated from

$$R_i = P\left(x_i|x_e\left(\vec{\theta}\right)\right)/P\left(x_i|x_e\left(\vec{\theta}_{BF}\right)\right), \tag{9.3}$$

where P is the Poisson probability of observing x_i events based on an expectation of x_e, and x_e is a function of the parameters $\vec{\theta}$. Holding to a fixed $\vec{\theta}$, all R_i are calculated and ordered from highest to lowest. The exclusion significance, c, of $\vec{\theta}$ is then calculated from

$$c = \sum_{i|R_i=R_{Max}}^{i|R_i=R_{Data}} P\left(x_i|x_e\left(\vec{\theta}\right)\right). \tag{9.4}$$

In practice, all R_i are not necessary, which is important as x_i goes from 0 to infinity. Instead, two facts are used to limit the number of calculations. First, the max rank is found using the fact that the Poisson probability peaks when the observed rate matches the mean. Second, since the probability summation stops when $R_i < R_{Data}$, the R_i are calculated moving outward from x_{Max} until this condition is met.

Nuisance parameters, including θ_{23} and the systematics, are handled by including and profiling them in equations 9.3 and 9.4 [135]. Specifically,

$$R_i = P\left(x_i|x_e\left(\vec{\theta},\vec{\lambda}_{BF(Data)}\right)\right)/P\left(x_i|x_e\left(\vec{\theta}_{BF},\vec{\lambda}_{BF}\right)\right), \tag{9.5}$$

where in the numerator, $\vec{\lambda}_{BF(Data)}$ is found by fitting to the actual data, and $\vec{\lambda}_{BF}$ in the denominator is calculated by fitting to the possible observation x_i. In other words,

$$\vec{\lambda}_{BF(Data)} = \vec{\lambda}|\chi^2\left(x_{Data},\vec{\theta},\vec{\lambda}\right) = \chi^2_{Min}\left(x_{Data},\vec{\theta}\right), \tag{9.6}$$

$$\left(\vec{\theta}_{BF},\vec{\lambda}_{BF}\right) = \left(\vec{\theta},\vec{\lambda}\right)|\chi^2\left(x_i,\vec{\theta},\vec{\lambda}\right) = \chi^2_{Min}(x_i). \tag{9.7}$$

The confidence interval is calculated using the best fit of the nuisance parameters to the data,

$$c = \sum_{i|R_i=R_{Max}}^{i|R_i=R_{Data}} P\left(x_i|x_e\left(\vec{\theta}, \vec{\lambda}_{BF(Data)}\right)\right), \tag{9.8}$$

where R_i is that which includes the nuisance parameters calculated from equation 9.5.

As mentioned at the end of section 9.2, the value of δ_{24} was held fixed. The FC corrected curves were calculated using values of $\delta_{24} = 0$, $\pi/2$, and $3\pi/2$. The most conservative results were obtained from $\delta_{24} = \pi/2$, so this value was used for the final measurements.

The results of the FC corrections are shown in figure 9.17, which shows the significance surface and compares the FC corrected 68% and 90% contours to their uncorrected counterparts. The massive improvement is largely an artifact of the assumed numbers of degrees of freedom (DOF). The uncorrected curves were constructed varying θ_{34} and θ_{24} in equation 9.2, generating a $\Delta\chi^2$, and interpreting the result assuming a separate DOF for each angle. For 2 DOF, the 68% contour was thus calculated where $\Delta\chi^2 = 2.30$, and the 90% contour at $\Delta\chi^2 = 4.61$. Given that the fits had a single measured number as input, using 2 DOFs was dubious, though it only resulted in an overly conservative limit. On the other hand, the FC corrected contour only considered a single DOF by construction.

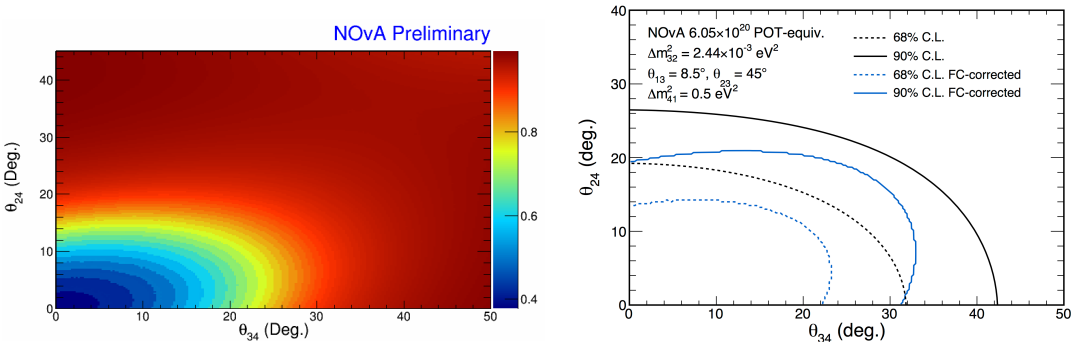


Figure 9.17: Left: the FC corrected significance surface. The $(\theta_{34}, \theta_{24})$ values are included in a c% confidence interval if the significance value is $\leq c$. Right: the 68% and 90% FC contours as compared to the uncorrected contours.

The two dimensional surface was profiled to create one dimensional slices of the two angles. These are shown in figure 9.18 with a comparison to the uncorrected versions. Here, the results are much more

consistent, as the uncorrected one dimensional slices only show a single DOF.

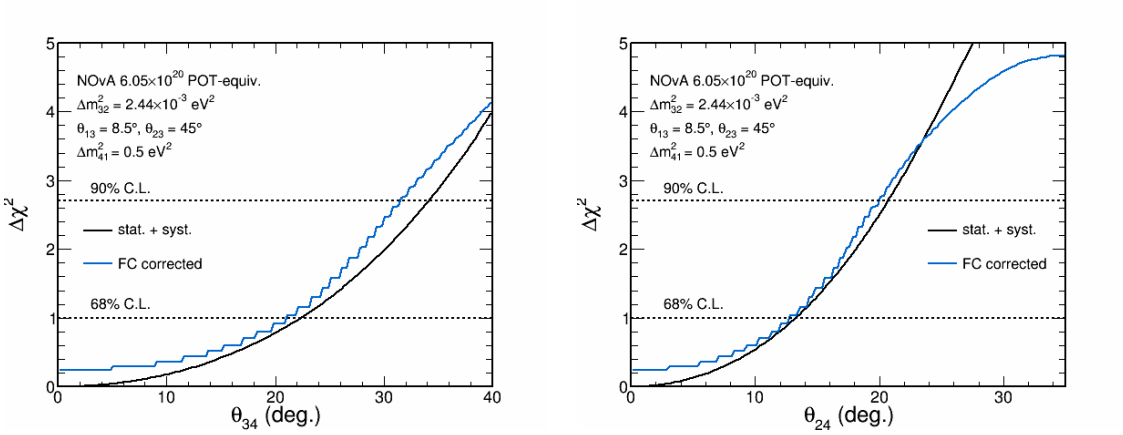


Figure 9.18: One dimensional slices of θ_{34} and θ_{24}, with the uncorrected slice in black and the FC corrected slice in blue. Left: θ_{34}. Right: θ_{24}.

The final results of the analysis, based on the FC corrected surface, placed upper limits on the mixing angles and matrix elements. The sterile mixing angle limits were measured at $\theta_{34} < 32°$ and $\theta_{24} < 20°$ at the 90% confidence level. These values were converted into measurements of the matrix elements via

$$\begin{aligned} |U_{\mu 4}|^2 &= \sin^2\theta_{24} \\ |U_{\tau 4}|^2 &= \cos^2\theta_{24}\sin^2\theta_{34} \end{aligned} \tag{9.9}$$

where $\theta_{14} = 0$ is assumed. The limits on the matrix elements were measured to be $|U_{\mu 4}|^2 < 0.12$ and $|U_{\tau 4}|^2 < 0.27$ at the 90% confidence level.

10

Future Improvements and Conclusions

10.1 Future Improvements

There are a number of improvements that could be made for this analysis. First and foremost, more data will mean stronger limits. The data used in this analysis represents less than 20% of the total planned exposure, so the measurements will naturally become stronger with time. Figure 10.1 shows how the θ_{34} limit will improve with 18×10^{20} POT, using the same rate only fit and assuming a modest 50% reduction in the cosmic background. Even under these assumptions, NOνA will match the current best measurement from Super-K [30] of $\theta_{34} < 25°$ (90% C.L.).

Probably the most important advance for the analysis itself would be the improvement of the cross section information, including the use of data to tune existing models and to include NC MEC events. In fact, the ND data/MC shape discrepancy was the driving force behind using the more robust rate only analysis, as opposed to a rate and shape analysis. The cross section model work, which is already underway, should improve this data/MC disagreement, at which point it will be natural to move to an analysis that

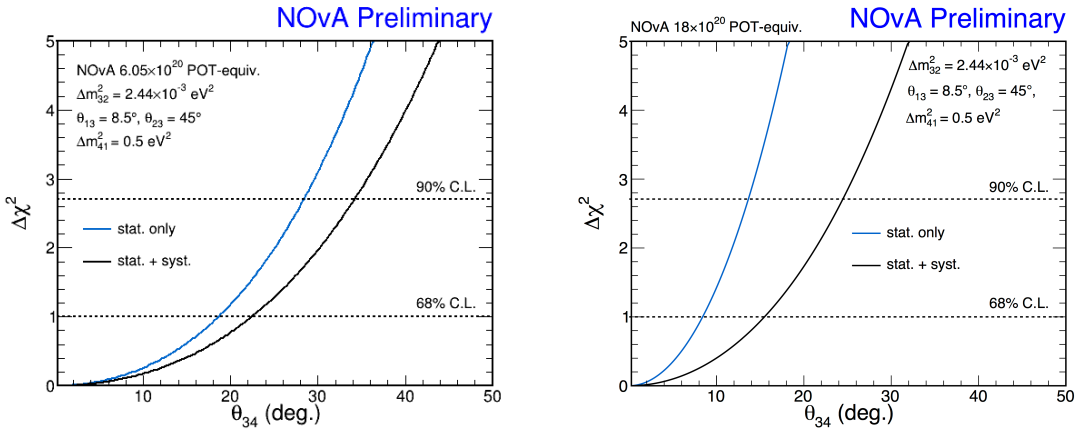


Figure 10.1: The projected sensitivity for θ_{34}. The left figure shows the measurement from this analysis. The right figure shows the projected measurement after 18×10^{20} POT, and assuming a 50% reduction in the cosmic background. The projected 90% confidence measurement is an upper limit slightly below 25°, the current best measurement on θ_{34} from Super-K [30]. The right figure is from reference [59].

includes shape and rate information. In fact, the functionality to perform a shape and rate analysis already exists and has been demonstrated, so this is the lowest hanging fruit for analysis improvements. The shape information will also provide an important handle on how δ_{24} affects the FD spectrum.

The neutral current selection has a number of places it could be improved. As mentioned in several sections of chapter 6, it was discovered after selection cuts were frozen that many of the reconstructed objects used for selection had hidden preselections applied. The prongs and tracks used for containment are prime examples of this, so there are plans to move to a more reconstruction free containment. The cosmic rejection BDT also requires a reconstructed track, so there are ongoing efforts to perform the cosmic rejection in a way that relaxes this constraint, even if this means handling 0 track events differently. Any overall improvements to the cosmic rejection could also translate into signal gains. The very harsh fiducial volume and containment cut at the top of the FD was made to remove cosmic events, so overall cosmic rejection improvements could relax these cuts and allow more signal events into the selected sample.

The largest systematic uncertainty came from the ND data/MC discrepancy. Future analyses should improve this in two ways. First, the inclusion of NC MEC events and improvement of the cross section models should naturally improve the ND data/MC shape agreement. Even a modest 50% improvement on this systematic would translate into overall systematic uncertainty reductions of 13% for the NC sig-

nal and 19% for the CC background. The other improvement would come from choosing a data driven decomposition method. This would require that the systematic be evaluated in a different way as well.

One other major analysis improvement would be to increase the range of Δm^2_{41} to higher values. It can be seen from figure 2.7 that this could affect the ND event rate. Consequently, this would require performing a joint oscillation parameter fit to the ND and FD data.

10.2 Conclusion

The NC disappearance analysis presented in this dissertation searched for mixing between active and sterile neutrinos in a $3+1$ model. 6.69×10^{20} POT were collected at the FD, or 6.05×10^{20} full detector equivalent POT. The results of this analysis were consistent with the no sterile neutrino mixing hypothesis and placed upper limits on the allowed values of the mixing angles, $\theta_{34} < 32^\circ$ and $\theta_{24} < 20^\circ$, and matrix elements, $|U_{\mu 4}|^2 < 0.12$ and $|U_{\tau 4}|^2 < 0.27$. There are many known avenues for future improvements to the analysis, including model improvements, moving to a shape and rate analysis, gaining signal efficiency from different reconstructed objects, and performing a joint ND and FD fit. With only modest improvements, NOνA will set competitive limits on the best fit value of θ_{34} and $|U_{\tau 4}|^2$ in a few years time. This is just the beginning of an exciting NOνA analysis!

References

[1] W. Pauli. Letter to a physicists' gathering at tubingen, 1930.

[2] J. Chadwick. Possible existence of a neutron. *Nature*, 192:312, 1932.

[3] E. Fermi. Versuch einer theorie der β-strahlen. *Z. Phys.*, 88:161–177, 1934.

[4] F. Reines and C. L. Cowan. Detection of the free neutrino. *Phys. Rev.*, 92:830–831, 1953.

[5] C. L. Cowan, F. Reines, F. B. Harrison, H. W. Kruse, and A. D. McGuire. Detection of the free neutrino: A confirmation. *Science*, 124:103–104, 1956.

[6] G. Danby et al. Observation of high-energy neutrino reactions and the existence of two kinds of neutrinos. *Phys. Rev. Lett.*, 9:36–44, 1962.

[7] G. Feldman. The lepton spectrum. *AIP Conf. Proc.*, 81:280, 1982.

[8] DONUT Collaboration, K. Kodama et al. Observation of tau neutrino interactions. *Phys. Lett. B*, 504:218–224, 2001.

[9] B. Pontecorvo. Mesonium and anti-mesonium. *Sov. Phys. JETP*, 6:429, 1957.

[10] R. Davis, Jr., D. S. Harmer, and K. C. Hoffman. Search for neutrinos from the sun. *Phys. Rev. Lett.*, 20:1205–1209, 1968.

[11] J. N. Bahcall, N. A. Bahcall, and G. Shaviv. Present status of the theoretical predictions for the ^{37}Cl solar-neutrino experiment. *Phys. Rev. Lett.*, 20:1209–1212, 1968.

[12] V. Gribov and B. Pontecorvo. Neutrino astronomy and lepton charge. *Phys. Lett. B*, 28:493–496, 1969.

[13] K.S. Hirata et al. Observation of B-8 solar neutrinos in the Kamiokande-II detector. *Phys. Rev. Lett.*, 63:16–19, 1989.

[14] A. I. Abazov et al. Search for neutrinos from the sun using the reaction $^{71}\mathrm{Ga}\,(\nu_e, e^-)\,^{71}\mathrm{Ge}$. *Phys. Rev. Lett.*, 67:3332–3335, 1991.

[15] P. Anselmann et al. Solar neutrinos observed by GALLEX at Gran Sasso. *Phys. Lett. B*, 285:376–389, 1992.

[16] K.S. Hirata et al. Observation of a small atmospheric ν_μ/ν_e ratio in Kamiokande. *Phys. Lett. B*, 280:146–152, 1992.

[17] Y. Fukuda et al. Evidence for oscillation of atmospheric neutrinos. *Phys. Rev. Lett.*, 81:1562–1567, 1998.

[18] Q. R. Ahmad et al. Direct evidence for neutrino flavor transformation from neutral-current interactions in the sudbury neutrino observatory. *Phys. Rev. Lett.*, 89:011301, 2002.

[19] ALEPH Collaboration, DELPHI Collaboration, L3 Collaboration, OPAL Collaboration, SID Collaboration, LEP Electroweak Working Group, SID Electroweak Group, Heavy Flavour Group. Precision electroweak measurements on the Z resonance. *Phys. Rep.*, 427:218–224, 2006.

[20] C. Athanassopoulos et al. Candidate events in a search for anti-muon-neutrino –> anti-electron-neutrino oscillations. *Phys. Rev. Lett.*, 75:2650–2653, 1995.

[21] A. A. Aguilar-Arevalo et al. A search for electron neutrino appearance at the $\Delta m^2 \sim 1$ eV^2 scale. *Phys. Rev. Lett.*, 98:231801, 2007.

[22] A. A. Aguilar-Arevalo et al. Improved search for $\bar{\nu}_\mu \to \bar{\nu}_e$ oscillations in the MiniBooNE experiment. *Phys. Rev. Lett.*, 110:161801, 2013.

[23] B. Armbruster et al. Upper limits for neutrino oscillations $\bar{\nu}_\mu \to \bar{\nu}_e$ from muon decay at rest. *Phys. Rev. D*, 65:112001, 2002.

[24] P. Astier et al. Search for $\nu_\mu \to \nu_e$ oscillations in the NOMAD experiment. *Phys. Lett. B*, 570:19–31, 2003.

[25] F. Kaether et al. Reanalysis of the GALLEX solar neutrino flux and source experiments. *Phys. Lett. B*, 685:47–54, 2010.

[26] J. Abdurashitov et al. Measurement of the response of a Ga solar neutrino experiment to neutrinos from an ^{37}Ar source. *Phys. Rev. C*, 73:045805, 2006.

[27] B. Achkar et al. Search for neutrino oscillations at 15, 40 and 95 meters from a nuclear power reactor at bugey. *Nucl. Phys. B*, 434:503–532, 1995.

[28] J. Ahn et al. Observation of reactor electron antineutrino disappearance in the reno experiment. *Phys. Rev. Lett.*, 108:191802, 2012.

[29] F. P. An et al. Search for a light sterile neutrino at Daya Bay. *Phys. Rev. Lett.*, 113:141802, 2014.

[30] K. Abe et al. Limits on sterile neutrino mixing using atmospheric neutrinos in Super-Kamiokande. *Phys. Rev. D*, 91:052019, 2015.

[31] P. Adamson et al. Active to sterile neutrino mixing limits from neutral-current interactions in minos. *Phys. Rev. Lett.*, 107:011802, 2011.

[32] Z. Maki, M. Nakagawa, and S. Sakata. Remarks on the unified model of elementary particles. *Prog. Theor. Phys.*, 28:870–880, 1962.

[33] J. Rich. Quantum mechanics of neutrino oscillations. *Phys. Rev. D*, 48:4318–4325, 1993.

[34] B. Kayser. On the quantum mechanics of neutrino oscillation. *Phys. Rev. D*, 24:110–116, 1981.

[35] K. A. Olive et al. Review of particle physics. *Chin. Phys. C*, 38:090001, 2014.

[36] O. Mena and S. Parke. Unified graphical summary of neutrino mixing parameters. *Phys. Rev. D*, 69:117301, 2004.

[37] E. Niner. *Observation of Electron Neutrino Appearance in the NuMI Beam with the NOνA Experiment*. PhD thesis, Indiana University, 2015.

[38] L. Wolfenstein. Neutrino oscillations in matter. *Phys. Rev. D*, 17:2369–2374, 1978.

[39] S. P. Mikheev and A. Y. Smirnov. Resonance enhancement of oscillations in matter and solar neutrino spectroscopy. *Sov. J. Nucl. Phys.*, 42:913–917, 1985.

[40] H. Nunokawa, S. Parke, and J. W. F. Valle. CP violation and neutrino oscillations. *Prog. Part. Nucl. Phys.*, 60:338–402, 2008.

[41] F. P. An et al. Observation of electron-antineutrino disappearance at Daya Bay. *Phys. Rev. Lett.*, 108:171803, 2012.

[42] F. Capozzi et al. Status of three-neutrino oscillation parameters, circa 2013. *Phys. Rev. D*, 89:093018, 2014.

[43] P. Adamson et al. First measurement of electron-neutrino appearance in NOvA. *Phys. Rev. Lett.*, 116:151806, 2016.

[44] P. Adamson et al. First measurement of muon-neutrino disappearance in NOvA. *Phys. Rev. D*, 93:051104, 2016.

[45] P. A. R. Ade et al. Planck 2015 results. XIII. Cosmological parameters. arXiv:1502.01589 [astro-ph.CO].

[46] J. Kopp, P. A. N. Machado, M. Maltoni, and T. Schwetz. Sterile neutrino oscillations: The global picture. *J. High Energy Phys.*, 2013:050, 2013.

[47] A. Palazzo. Estimate of θ_{14} independent of the reactor antineutrino flux determinations. *Phys. Rev. D*, 85:077301, 2012.

[48] A. Strumia and F. Vissani. Neutrino masses and mixings and ... https://arxiv.org/abs/hep-ph/0606054.

[49] D. S. Ayers et al. NOνA technical design handbook. Technical report, Fermilab, Batavia, IL, 2007.

[50] NuMI technical design handbook. Technical report, Fermilab, Batavia, IL, 2004.

[51] P. Vahle. New results from NOνA. NOνA Internal Document 15688.

[52] M. Strait. Operations plenary June 2016. NOνA Internal Document 15489.

[53] J. Zalesak. Operations status. NOνA Internal Document 15654.

[54] N. Felt, J. N. Oliver, and T. Zimmerman. Front-End board. NOνA Internal Document 4088.

[55] N. Felt and J. N. Oliver. FEBs for the NOvA Far Detector. NOνA Internal Document 8172.

[56] T. Zimmerman. NOvA APD readout chip v2. NOνA Internal Document 4371.

[57] N. Felt and J. N. Oliver. NOvA FEB programming manual. NOνA Internal Document 4662.

[58] A. Waldron. Lowering readout thresholds. NOνA Internal Document 9671.

[59] G. Davies. NOvA NC JETP seminar (draft). NOνA Internal Document 15831.

[60] J. Cooper. NOνA Near Detector mass. NOνA Internal Document 11906.

[61] T. T. Böhlen et al. The FLUKA code: Developments and challenges for high energy and medical applications. *Nucl. Data Sheets*, 120:211–214, 2014.

[62] A. Ferrari et al. FLUKA: A multi-particle transport code. *CERN-2005-10, INFN/TC_05/11, SLAC-R-773*, 2005.

[63] S. Agostinelli et al. Geant4–a simulation toolkit. *Nucl. Instrum. Methods Phys. Res., Sect. A*, 506:250–303, 2003.

[64] J. Allison et al. Geant4 developments and applications. *IEEE Trans. Nucl. Sci.*, 53:270–278, 2006.

[65] G. Kafka. Spider support fix/update. NuMI-X Internal Document 39.

[66] R. Hatcher. Proposal for a unified "flux" N-tuple format. MINOS Internal Document 9070.

[67] R. Hatcher. The NOνA simulation chain. NOνA Internal Document 6941.

[68] A. Aurisano. The NOνA detector simulation. NOνA Internal Document 13577.

[69] A. Aurisano and J. Musser. Detector simulation technote. NOνA Internal Document 15338.

[70] C. Andreopoulos et al. The GENIE neutrino Monte Carlo generator. *Nucl. Instrum. Methods Phys. Res., Sect. A*, 614:87–104, 2010.

[71] C. Andreopoulos et al. The GENIE neutrino Monte Carlo generator: Physics and user manual. http://arxiv.org/abs/1510.05494.

[72] P. A. Rodriques et al. Identification of nuclear effects in neutrino-carbon interactions at low three-momentum transfer. *Phys. Rev. Lett.*, 116:071802, 2016.

[73] H. Gallagher et al. GENIE tune and uncertainty band for second analysis. NOνA Internal Document 15214.

[74] J. Wolcott. Blessed plots for SA GENIE tune. NOνA Internal Document 15481.

[75] G. Davies and R. Murphy. Updated ND rock rate. NOνA Internal Document 13247.

[76] C. N. Chou. The nature of the saturation effect of fluorescent scintillators. *Phys. Rev.*, 87:904–905, 1952.

[77] D. Pershey. Birks-chou parameter selection. NOνA Internal Document 13158.

[78] G. Kafka. The FluxReader framework. https://cdcvs.fnal.gov/redmine/attachments/download/19714/FluxReader_TN.pdf.

[79] H. Duyang et al. A technote describing the derivation and size of NuMI flux uncertainties used in the second NOvA analyses. NOνA Internal Document 15296.

[80] M. Del Tutto. Fluka 11.2c.0 validation. NOνA Internal Document 13164.

[81] J. Cooper et al. Special run proposal. NOνA Internal Document 13451.

[82] M. Del Tutto. Special runs - horn current study. NOνA Internal Document 13523.

[83] G. Brunetti. Target studies with FLUGG. NOνA Internal Document 12556.

[84] A. Sousa, J. Musser, and G. Kafka. Detector simulations status. NOνA Internal Document 12263.

[85] J. Musser. Detsim summary. NOνA Internal Document 14740.

[86] C. Backhouse et al. Calibration technotes. NOνA Internal Document 13579.

[87] B. Rebel. Window tracking update. NOνA Internal Document 11806.

[88] D. E. Groom, N. V. Mokhov, and S. I. Striganov. Muon stopping power and range tables 10 MeV–100 TeV. *Atom. Data Nucl. Data Tabl.*, 78:183–356, 2001.

[89] S. Mufson. Density of NOνA liquid scintillator at 69° F. NOνA Internal Document 11886.

[90] M. Ester et al. A Density-Based algorithm for discovering clusters in large spatial databases with noise. pages 226–231. AAAI Press, 1996.

[91] M. Baird. *An Analysis of Muon Neutrino Disappearance from the NuMI Beam Using an Optimal Track Fitter*. PhD thesis, Indiana University, 2015.

[92] L. Fernandes and M. Oliveira. Real-time line detection through an improved Hough transform voting scheme. *Pattern Recogn.*, 41:299–314, 2008.

[93] M. Baird. Tech note for the Multi-Hough transform. NOνA Internal Document 8241.

[94] M. Gyulassy and M. Harlander. Elastic tracking and neural network algorithms for complex pattern recognition. *Comput. Phys. Commun.*, 66:31–46, 1991.

[95] M. Ohlsson, C. Peterson, and A. L. Yuille. Track finding with deformable templates - the elastic arms approach. *Comput. Phys. Commun.*, 71:77–98, 1992.

[96] M. Ohlsson. Extensions and explorations of the elastic arms algorithm. *Comput. Phys. Commun.*, 77:19–32, 1993.

[97] R. Frühwirth and A. Strandlie. Track fitting with ambiguities and noise: A study of elastic tracking and nonlinear filters. *Comput. Phys. Commun.*, 120:197–214, 1999.

[98] M. Messier. Vertex reconstruction based on elastic arms. NOνA Internal Document 7530.

[99] E. Niner. Technical note of FuzzyKVertex reconstruction. NOνA Internal Document 7648.

[100] J. C. Dunn. A fuzzy relative of the ISODATA process and its use in detecting compact well-separated clusters. *J. Cybern.*, 3:32–57, 1973.

[101] J. C. Bezdek. *Pattern Recognition with Fuzzy Objective Function Algorithms*. Plenum Press, New York, 1981.

[102] R. Krishnapuram and J. M. Keller. A possibilistic approach to clustering. *IEEE Trans. Fuzzy Syst.*, 1:98–110, 1993.

[103] N. H. Kuiper. Tests concerning random points on a circle. *P. K. Ned. Akad. Wetensc. A*, 63:38–47, 1960.

[104] N. Raddatz. KalmanTrack technical note. NOνA Internal Document 13545.

[105] C. Backhouse. The library event matching nue PID – technote. NOνA Internal Document 13588.

[106] K. Sachdev, J. Bian, and E. Niner. Likelihood based nue identifier (LID). NOνA Internal Document 13590.

[107] N. J. Raddatz. ReMID technical note. NOνA Internal Document 11206.

[108] K. Bays. Cosmic rejection technical note. NOνA Internal Document 11205.

[109] A. Aurisano et al. A convolutional neural network neutrino event classifier. *JINST*, 11:P09001, 2016.

[110] C. Szegedy et al. Going deeper with convolutions. `http://arxiv.org/abs/1409.4842`.

[111] Y. Jia et al. Caffe: Convolutional architecture for fast feature embedding. `http://arxiv.org/abs/1408.5093`.

[112] A. Aurisano. CVN for NC selection. NOνA Internal Document 15349.

[113] K. Sachdev. Tech note: Fd selection cuts for second analysis. NOνA Internal Document 15387.

[114] L. Goodenough and S. R. Phan-Budd. Technical note on the NOvA beam monitoring for 2015 summer analysis. NOνA Internal Document 13572.

[115] X. Bu and K. Sachdev. Spill level data quality- technical note. NOνA Internal Document 12437.

[116] K. Bays. LiveGeometry tech note. NOνA Internal Document 11470.

[117] S. Lein. DCM edge metric. NOνA Internal Document 13527.

[118] C. Backhouse et al. Near detector data/MC comparisons for the nue analysis. NOνA Internal Document 13587.

[119] J. Bian and T. Xin. Note of selection cuts for Nue first analysis. NOνA Internal Document 13592.

[120] K. Bays et al. Executive summary for the first numu CC analysis. NOνA Internal Document 13641.

[121] C. Backhouse et al. The CAFAna framework. NOνA Internal Document 9222.

[122] E. Catano-Mur. BEN decomposition technote. NOνA Internal Document 15392.

[123] D. Pershey. MichelDecomp for the SA. NOνA Internal Document 15395.

[124] G. Kafka. Decomposition by varied horn current. NOνA Internal Document 13424.

[125] L. Suter. NuMu ND selection tech note. NOνA Internal Document 13212.

[126] J. Lozier. ModularExtrap technical note. NOνA Internal Document 12563.

[127] D. Pershey. ME check with SA rock muons in ND. NOνA Internal Document 14959.

[128] G. Davies. pi0 invariant mass in SA files. NOνA Internal Document 15005.

[129] K. Matera et al. Calibration status for second analysis. NOνA Internal Document 15127.

[130] A. Markowitz. Data MC comparison for NC analysis, updated with SA files. NOνA Internal Document 15151.

[131] N. Raddatz. Fiducial mass systematic. NOνA Internal Document 13237.

[132] A. Markowitz. Reconstruction efficiency study summary. NOνA Internal Document 15447.

[133] K. Abe et al. Measurement of neutrino oscillation in appearance and disappearance channels by the T2K experiment with 6.6×10^{20} protons on target. *Phys. Rev. D*, 91:072010, 2015.

[134] G. J. Feldman and R. D. Cousins. Unified approach to the classical statistical analysis of small signals. *Phys. Rev. D*, 57:3873–3889, 1998.

[135] G. Feldman. Notes on the inclusion of nuisance parameters in the unified approach. NOνA Internal Document 15884.