

CONSTRAINTS ON NEUTRINO OSCILLATION PARAMETERS FROM NEUTRINOS
AND ANTINEUTRINOS WITH MACHINE LEARNING

Micah Groh

Submitted to the faculty of the University Graduate School

in partial fulfillment of the requirements

for the degree

Doctor of Philosophy

in the Department of Physics,

Indiana University

March 2021

Accepted by the Graduate Faculty, Indiana University, in partial fulfillment of the requirements
for the degree of Doctor of Philosophy.

Doctoral Committee

Mark Messier, PhD

Rex Tayloe, PhD

Scott Wissink, PhD

Charles Horowitz, PhD

19 February 2021

Copyright © 2021

Micah Groh

Acknowledgments

First, thank you to my family whose love for science has led me to where I am today. Thank you for the endless support and fun.

Thank you to Mark Messier, my advisor for this thesis, who always found time to help and guide me even when no time existed. I learned much about neutrinos, plotting, writing styles, and much more from our discussions which will serve me through my career. Thank you to the other members of my committee, Rex Tayloe, Scott Wissink, and Charles Horowitz. Also, thank you to Jon Urheim who was not my advisor, but provided guidance on neutrinos, DUNE, and new events in high energy physics.

Thank you to those I worked with at Indiana University. Thank you to post-docs Denver Whittington and Gavin Davies who provided me guidance and insight throughout graduate school particularly when I was just starting on NOvA. Thank you to graduate students Justin Vasel, Bruce Howard, and Teresa Lackey who were always willing to help me when I was stuck. Also, thank you to the other members of the group who I have gotten to know over the years and have had numerous conversations with both in and out of work: Ryan Murphy, Ashley Back, Erica Smith, Maria Manrique, Reed Bowles, and Chris Macias.

Thank you to the members of the physics department staff at Indiana University who have helped me through many essential tasks throughout graduate school: Anne-Marie Foley, Anne Sophie Stringer, Brendan Devine, and Nelson Batalon.

Thank you to the many people I worked with on NOvA and got to know while developing this thesis: Nitish Nayak, Abhilash Dombara, Shiqi Yu, Liudmila Kolupaeva, Biswaranjan Behera, and Reddy Gandrajula. Thank you to the NOvA working group conveners who guided me on the numerous projects I worked on as a graduate student: Karl Warburton, Evan Niner, Michael Baird, Alex Radovic, and Matt Strait. Also, thank you to the senior members of the NOvA

collaboration who have provided me with support or help over the years: Alex Himmel, Louise Suter, Peter Shanahan, Patricia Vahle, and Alec Habig. Particular thanks goes to the members of Young NOvA. Thank you for giving me the chance to serve as your leader and for the fun outside of meetings.

Finally, I would like to thank Fernanda Psihas who took the time out of writing her own dissertation to teach me everything she knew, guided me every step of the way, and has been my friend ever since. I'll never be able to repay you.

Thank you to the support staff for the Carbonate Computing Cluster at Indiana University and the Wilson Cluster at Fermilab which were instrumental for the development of the algorithms used in this thesis. This research was supported in part by Lilly Endowment, Inc., through its support for the Indiana University Pervasive Technology Institute.

Thank you to the staff at the Fermilab accelerator division who monitor the NuMI beam 24/7 and made the data collection for this thesis possible. This work was funded by a grant from the United States Department of Energy. Thank you for your continued funding of scientific research.

Micah Groh

CONSTRAINTS ON NEUTRINO OSCILLATION PARAMETERS FROM NEUTRINOS AND
ANTINEUTRINOS WITH MACHINE LEARNING

NOvA is a two detector, long baseline neutrino oscillation experiment measuring the oscillations of muon neutrinos from the NuMI neutrino beam over a baseline of 810 km. The experiment uses four oscillation channels, $\nu_\mu \rightarrow \nu_\mu$, $\bar{\nu}_\mu \rightarrow \bar{\nu}_\mu$, $\nu_\mu \rightarrow \nu_e$, and $\bar{\nu}_\mu \rightarrow \bar{\nu}_e$, with a peak neutrino energy of 1.8 GeV. This dissertation describes the analysis of these channels using a dataset of 13.6×10^{20} protons on target neutrino beam mode and 12.5×10^{20} protons on target antineutrino beam mode. The analysis makes use of improvements in the treatment of systematic uncertainties and machine learning techniques to reconstruct neutrino interactions. A technique for decorrelating systematic errors using principle component analysis was utilized to reduce and optimize neutrino cross section and beam related uncertainties. The improved machine learning algorithms make use of convolutional neural networks for neutrino event classification, particle classification, and instance segmentation. The selection of neutrino signal events utilizing the neutrino event classifier shows an efficiency of 63% for the selection of electron neutrinos in neutrino beam mode and 75% for electron antineutrinos in antineutrino beam mode. Using this algorithm, 82 appearing electron neutrino candidates and 33 appearing electron antineutrino candidates were observed with expected backgrounds of 26.8 and 14.0 respectively. In addition, 211 surviving muon neutrino candidates and 105 muon antineutrino candidates were identified with a purity of more than 96% using the same neutrino event classifier. Fitting these data to the three flavor neutrino oscillation model, using constraints on θ_{12} , θ_{13} , and Δm_{12}^2 from solar and reactor neutrino experiments, the oscillation parameters are measured to be $\sin^2\theta_{23} = 0.57_{-0.03}^{+0.04}$, $\Delta m_{32}^2 = +2.41 \pm 0.07 \times 10^{-3} \text{ eV}^2$, and $\delta_{\text{CP}} = 0.82_{-0.87}^{+0.27}\pi$ with a preference for the normal neutrino mass hierarchy. Leading systematic uncertainties for these measurements come from detector calibration and neutrino interaction models.

Contents

1	Introduction	1
2	Neutrino Interactions and Oscillations	6
2.1	The Standard Model	6
2.2	Neutrino Interactions	7
2.3	Neutrino Oscillations	11
2.3.1	Muon Neutrino Disappearance	13
2.3.2	Electron Neutrino Appearance	14
2.3.3	Current Parameters and Open Questions	15
2.3.4	Matter Effects	16
2.4	Current Experiments	19
3	The NOvA Experiment	21
3.1	The NuMI Beam	21
3.2	NOvA Detectors	25
3.3	Data Acquisition System	30
3.4	Timing System	31
3.5	Neutrino Simulation	36
3.6	Detector Calibration	37
4	Event Reconstruction	41
4.1	Neutrino Interaction Topologies	41
4.2	Isolating Neutrino Interactions	44
4.3	Resolving the Interaction Point	47
4.4	Clustering Particles	50

4.5	Particle Tracking	54
4.5.1	Kalman Filter Tracking	55
4.5.2	Break Point Fitter	56
5	Machine Learning Techniques	59
5.1	Convolutional Neural Networks	59
5.2	Event Classification	62
5.3	Particle Classification	70
5.3.1	π^0 Mass Peak	74
5.4	Instance Aware Semantic Segmentation	75
6	Neutrino Oscillation Analysis	84
6.1	Energy Estimation	84
6.1.1	Electron Neutrino Energy	84
6.1.2	Muon Neutrino Energy	86
6.2	Event Selection	89
6.2.1	Electron Neutrino Selection	89
6.2.2	Muon Neutrino Selection	92
6.3	Far Detector Prediction	96
6.3.1	Signal Extrapolation	97
6.3.2	Background Extrapolation	100
6.3.3	Measurement of Cosmic Ray Backgrounds	102
6.4	Fitting and Sensitivities	106
7	Systematic Uncertainties	107
7.1	Systematics in the Analysis	107
7.2	Sources of Systematic Uncertainties	109

7.2.1	Detector Calibration	109
7.2.2	Detector Response	110
7.2.3	Neutrino Cross Sections	111
7.2.4	Beam Flux	114
7.2.5	Neutron Uncertainty	114
7.2.6	Lepton Reconstruction	114
7.2.7	Near-Far Uncorrelated	116
7.3	Uncertainties by Principal Component Decomposition	117
7.4	Uncertainty Interpolation	122
7.5	Fitting with Systematics	123
8	Experimental Results and Discussion	126
8.1	Results	126
8.2	Neutrino Oscillation Parameter Constraints	127
8.3	Comparison to Other Experiments	131
8.4	The Future of NOvA	132
8.5	Conclusions	134
	Bibliography	136
	Curriculum Vitae	

List of Figures

1.1	β Decay Spectrum	2
1.2	Neutrino Interactions 43 Years Apart	4
2.1	The Standard Model	7
2.2	Neutrino Interaction Modes.	9
2.3	Neutrino Cross Sections	9
2.4	Neutrino Oscillation Ratio	14
2.5	Neutrino Mass Ordering	16
2.6	ν_e Biprobability	18
3.1	The NuMI Beamline	23
3.2	The Off-Axis Spectrum	24
3.3	Projected Neutrino Events	25
3.4	Daily POT Recorded by NOvA	26
3.5	A Detector Cell and APD	27
3.6	A Map of NOvA	28
3.7	The NOvA Detectors	29
3.8	Overview of the DAQ	31
3.9	A Cosmic Trigger Event	32
3.10	Timing System Schematic	33
3.11	Duel Correlated Sampling	35
3.12	Passage of particles through Matter	38
3.13	Cell Attenuation Correction	39
3.14	Absolute Energy Calibration	40

4.1	NOvA Event Topologies	42
4.2	Timing Event Display	45
4.3	Near Detector Pileup	47
4.4	The Hough Transform and Elastic Arms Vertex	49
4.5	Elastic Arms Vertex Resolution	50
4.6	Prong View Matching	52
4.7	K-Means Prong Performance	54
4.8	Reconstructed Track Direction	55
4.9	Break Point Fitter Schematic	56
4.10	Break Point Fitter Track	57
5.1	Elastic Arms Failure	60
5.2	Overview of a CNN	61
5.3	NOvA Pixel Map	63
5.4	Separable Convolutions	64
5.5	CVN Architecture	66
5.6	Event Classification Matrix	69
5.7	Neutrino tSNE	70
5.8	Prong Classification Inputs	71
5.9	Prong Classification Matrix	72
5.10	Context Enriched Improvement	73
5.11	π^0 Mass Peak	75
5.12	Mask R-CNN Diagram	76
5.13	Instance Segmentation Event Display	79
5.14	Reconstructed True Particles	80
5.15	Prong Quality Comparison	81

5.16	Event Complexity	82
5.17	Instance Segmentation Classification Matrix	82
6.1	Detector Energy Response	85
6.2	ν_e Energy	86
6.3	ν_e Energy Resolution	87
6.4	Spline Fits	88
6.5	ν_μ Energy Resolution	89
6.6	ν_e Event Selection	90
6.7	ν_μ Event Selection	94
6.8	Extrapolation Diagram	96
6.9	ND Event Spectra	97
6.10	Transverse Momentum Spectra	99
6.11	ND Background Decomposition	102
6.12	Predicted ν_e Spectra	103
6.13	Predicted ν_μ Spectra	104
7.1	Systematics Extrapolation Comparison	109
7.2	Neutron Prongs	115
7.3	Beam Flux PCA	121
7.4	Cross Section PC 0	122
7.5	Systematic Interpolation	123
7.6	Effects of Systematics	124
8.1	ν_e Data Septra	127
8.2	ν_e Event Range	128
8.3	ν_μ Data Septra	129

8.4	ν_μ Oscillation Ratio	130
8.5	Oscillation Fit Contours	131
8.6	Atmospheric Measurements	132
8.7	Comparison to T2K	133
8.8	Future Sensitivities	134

Chapter 1

Introduction

Neutrinos are among the most elusive particles in nature. They rarely interact with other forms of matter and instead many travel through the universe never to be observed. However, they are also the most abundant massive particle in the universe. For this reason, the effect they can have on physical phenomena is significant despite their small interaction rate. Neutrinos show themselves as ghost particles emitted during nuclear decay to the main form of energy released when star goes supernova. Neutrinos may hold the answer to many questions about the formation and evolution of the universe.

Unfortunately, their small interaction rate also means studying neutrinos is a technological and scientific challenge. Not only do they rarely interact, but there are large uncertainties in the rates and processes through which they interact in current models. Experiments aiming to uncover the properties of neutrinos often employ massive detectors, intense sources of neutrinos, or long enough exposures in order to observe enough neutrinos to make statistically significant measurements. Despite these challenges, the field of experimental particle physics is dedicated to learning as much as possible about these elusive particles and how they relate to the development of the universe. The field has made tremendous progress since neutrinos were first observed, but many open questions remain.

Neutrinos were first postulated in 1930 by Wolfgang Pauli [1] in order to explain the missing energy in β decay. At the time, β decay was believed to be a two body process, what we now recognize as $p^+ \rightarrow n + e^+$, and, thus, the daughter positron would have to have a discrete energy value. Instead, it was found to have the continuous energy spectrum shown in Fig. 1.1. Pauli was the first to propose that a third particle was also being emitted carrying the missing energy. To be undetected, the particle would need to be electrically neutral, and for the spectrum to extend to




Figure 1.1: β Decay Spectrum. The energy carried by the electron in β decay. If the process were two body, the electron would have a fixed energy of E_{\max} . Instead a neutrino carries away some of the energy.

the maximum available energy, the particle would need to be very light or even massless.

Enrico Fermi developed a formal theory of β decay in 1933 [2]. He named the third particle “neutrino”, meaning “little neutral one”.

More than 20 years would pass from Pauli’s prediction before anyone definitively observed a neutrino. In 1956, Clyde Cowan and Fred Reines observed a neutrino through inverse β decay [3], $\bar{\nu} + p^+ \rightarrow n + e^+$. They observed the scintillation light from a positron together with the delayed photon emitted through neutron capture from the interactions of neutrinos generated by a nearby nuclear reactor (though they initially considered a nuclear bomb as a source of neutrinos). This experiment, known as Project Poltergeist, discovered the first “flavor” of neutrinos, the electron neutrino.

Today, we have evidence from the width of Z decays that there are three flavors of light, active neutrinos [4], one corresponding to each of the charged leptons: electron neutrino, muon neutrino, and tau neutrino. When a neutrino interacts, it is always associated with the corresponding lepton in charged current interactions or another neutrino in neutral current interactions. The second

flavor of neutrinos, the muon neutrino, was observed in 1962 at Brookhaven National Lab [5]. The final flavor, the tau neutrino, wasn't observed until 2000 by the DONUT experiment at Fermi National Accelerator Laboratory [6].

In the mid 60s, little was known about the properties of the neutrino. They were known to be by-products of nuclear reactions, including those that power stars like the sun. John Bahcall made the first prediction of the rate of neutrinos that should be coming from the sun. He and Ray Davis developed the Homestake Experiment to experimentally measure this rate using 100000 gallons of tetrachloroethylene. However, the measured rate was less than a third of the predicted value, a discrepancy that became known as the solar neutrino problem.

The Super-Kamiokande experiment was the first to observe evidence of what is known as neutrino oscillations in 1998 using atmospheric muon neutrinos produced by cosmic rays in the Earth's atmosphere interacting in a water cherenkov detector [7]. Under neutrino oscillations, a neutrino produced in one flavor can interact as a different flavor after some time. What Bahcall and Davis didn't realize was that two thirds of their solar neutrinos were oscillating to muon or tau neutrinos, which their detector was not sensitive to. The phenomenon of neutrino oscillations was further confirmed by the SNO experiment in 2001 using neutrinos definitively coming from the sun [8], finally resolving the solar neutrino problem.

One of the most interesting consequences of neutrino oscillations is that neutrinos must have mass, a property not predicted by the Standard Model. Neutrino oscillations are, thus, the only experimentally confirmed evidence of physics beyond the Standard Model.

All neutrino detectors never directly observe a neutrino, rather they detect the particles created by a neutrino interacting within the detector medium. The early neutrino experiments only counted coincidence rates of these particles emitted from the neutrino interaction. The first images of neutrino interactions came from spark and bubble chambers during the 60s and 70s [9]. The first neutrino recorded by a bubble chamber came from the detector at Argonne National Laboratory in




Figure 1.2: Neutrino Interactions 43 Years Apart. Left: The first muon neutrino interaction recorded in the 12 foot bubble chamber detector at Argonne National Laboratory, 13 November 1970. Right: The first neutrino interaction recorded in the 14000 ton segmented NOvA detector, 12 November 2013, just one day short of 43 years later.

1970 and can be seen on the left in Fig. 1.2. It depicts a muon neutrino charged current interaction producing a final state pion in $\nu_\mu + p^+ \rightarrow \mu^- + p^+ + \pi^+$.

This image of a neutrino interaction allowed for a full reconstruction of the properties of the incident neutrino. Bubble chambers are magnetized so any charged particles passing through them curve in a direction depending on their charge and with a curvature proportional to their momentum. In this case, the negatively charged muon is clearly curving in the opposite direction from the positively charge pion and proton. The curvature can be used to determine the momentum of each of the three particles. The common origin of the particle trajectories is where the neutrino passed into the detector undetected before interacting with a proton. Finally, the presence of the muon is a clear sign that the flavor was a muon neutrino.

On the right of Fig. 1.2 is the first neutrino recorded by the NOvA detector, just one day less than 43 years later. This event depicts a neutral current interaction of $\nu + p^+ \rightarrow p^+ + \pi^0$ and a subsequent $\pi^0 \rightarrow \gamma + \gamma$. The amount of energy deposited is shown by the color of each pixel. The point on the left is where the neutrino interacted with a proton to produce a π^0 . The π^0 decayed into two photons which produces the gappy electromagnetic showers that follow.

The reconstruction of particle properties has evolved over time as new technologies in both detector design and computing have developed. Many neutrino detection mechanisms have been successfully employed to analyze the elusive neutrinos. Some detectors use cherenkov radiation, cones of light emitted by particles going faster than light through a medium such as water or ice. Others use ionization electrons, which drift in an electric field towards a detection plane. Each technology requires a different set of tools and techniques to accurately identify and reconstruct neutrinos that interact within the detector.

The bubble chamber event would have been identified and annotated by a human scanning through hundreds of such images hoping to spot the elusive neutrinos. The NOvA event was processed by a computer algorithm using various methods to group the hits and identify each of the individual particles. In recent decades, thanks to advancements in high performance computing, new software, and hardware upgrades, millions of such events can be processed in a fraction of the time and with greater accuracy than a human could. Today, developments in artificial intelligence and machine learning have allowed experiments to push the boundaries of the physics that can be achieved. NOvA is one experiment utilizing advances in these new techniques to further our understanding of the properties of neutrinos which will be discussed in this work.

This dissertation describes the measurements of neutrino oscillations in the $\nu_\mu \rightarrow \nu_\mu$, $\bar{\nu}_\mu \rightarrow \bar{\nu}_\mu$, $\nu_\mu \rightarrow \nu_e$, and $\bar{\nu}_\mu \rightarrow \bar{\nu}_e$ oscillation channels and the associated systematic uncertainties, as well as improvements in reconstruction using machine learning techniques. I begin by describing our current understanding of neutrino interactions and the phenomenon of neutrino oscillations in Chap. 2. Chapter 3 describes the experimental design and setup of NOvA. Chapters 4 and 5 describe the reconstruction algorithms used in the analysis and how machine learning has expanded the capabilities of the experiment. The description of how neutrino oscillation parameters are measured is in Chap. 6 and the treatment of systematic uncertainties for the analysis is in Chap. 7. Finally, Chap. 8 describes the experimental results from the NOvA experiment.

Chapter 2

Neutrino Interactions and Oscillations

This chapter will give an overview of the theoretical models needed to understand neutrino interactions and oscillations. Focus will be given to processes relevant to the NOvA experiment and the results in this dissertation.

This chapter will use natural units with $\hbar = c = 1$. Many steps in the derivations will be skipped as they are not insightful and can be read about in other texts such as Ref. [10].

2.1 The Standard Model

The Standard Model (SM) is a quantum field theory which describes the interactions between fundamental particles. It includes a description of the electromagnetic force, strong nuclear force, and weak nuclear force. There are currently seventeen known fundamental particles, of which five are bosons which mediate the forces. The remaining twelve are fermions which are the building blocks of matter. Figure 2.1 shows the particles making up the standard model. The SM has 19 free parameters, including the masses of the fermions, which uniquely determines particle dynamics.

The discovery of the Higgs boson in 2012 [11] completed the SM, it was the last particle predicted by the SM. Today, the SM is regarded as one of the most successful theories in modern particle physics. It makes predictions about numerous physical phenomena, many of which have been confirmed experimentally. Current experiments are making precision measurements of SM processes or searching for exotic phenomena [12].

Neutrinos are a subset of the leptons which only interact via the weak force, as they have no electric charge and no color charge. They come in three flavors: ν_e , ν_μ , or ν_τ corresponding to each of the charged leptons. The most interesting characteristic of neutrinos is their ability to change into a different flavor, a process known as neutrino oscillations. Neutrino oscillations are currently




Figure 2.1: The Standard Model. The standard model of particle physics describes the dynamics and interactions of the seventeen known fundamental particles.

the only experimentally confirmed evidence of physics beyond the SM in a laboratory.

One of the most significant consequences of neutrino oscillations is that neutrinos must have non-zero mass [13], a property in opposition with the SM where neutrinos have no mass. The current best constraints on the neutrino mass come from cosmological data on the sum of the neutrino masses $\sum m_{\nu_i} < 120 \text{ meV}$ [14], which is more than six orders of magnitude smaller than the electron mass. Direct measurements of the neutrino mass come from β decay detection with current limits at $\sum m_{\nu_i} < 800 \text{ meV}$ [15]. The minute mass allows for interesting mechanisms for mass scales [16] and could answer fundamental questions about particle physics [17].

2.2 Neutrino Interactions

The total number of leptons from each of the three generations of leptons is observed to be the same in the initial and final state of a neutrino interaction, where anti-leptons count as -1. This is the conservation of lepton flavor which means that a neutrino interaction must have a product which is either a neutrino with the same flavor as the original or the corresponding charged lepton.

Neutrinos interact by the weak nuclear force in either Neutral Current (NC) or Charged Current (CC). NC interactions are mediated by the exchange of the neutral Z^0 boson. The outgoing lepton is another neutrino of the same flavor. NC interactions have the common form:

$$\nu_\alpha + N \rightarrow \nu_\alpha + X \quad (2.1)$$

$$\bar{\nu}_\alpha + N \rightarrow \bar{\nu}_\alpha + X \quad (2.2)$$

where N is a nucleon, either a proton or neutron, α is any lepton flavor e , μ , or τ , and X is some collection of hadrons which depends on the details of the interaction. The first interaction is for neutrinos and the second for antineutrinos.

CC interactions are mediated by the charged W^\pm boson. In this case, lepton flavor is conserved by the corresponding charged lepton in the final state:

$$\nu_\alpha + N \rightarrow \ell_\alpha^- + X \quad (2.3)$$

$$\bar{\nu}_\alpha + N \rightarrow \ell_\alpha^+ + X \quad (2.4)$$

where ℓ_α is a charged lepton and the remaining notation is the same as before.

Whether the neutrino interaction is CC or NC determines the lepton in the final state. The hadrons in the final state are determined by the nature of the neutrino interaction with the nucleon. The three most common interaction types or modes are Quasi-Elastic (QE), Resonant Pion Production (RES), and Deep Inelastic Scattering (DIS). These three modes are depicted in Fig. 2.2. Other interaction modes such as Coherent scattering (COH), which involves scattering off an entire nucleus, and Meson Exchange Currents (MEC), which involves interactions with a nucleon pair, are less common. Interactions of neutrinos on electrons are very rare and can be safely neglected for this dissertation¹. The total neutrino cross section is the sum of the cross section for each

¹The cross section ratio with electrons to that with nucleons is about $m_e/m_N \approx 5 \times 10^{-4}$ at 1 GeV.




Figure 2.2: Neutrino Interaction Modes. Feynman diagrams depicting the three most common neutrino interaction modes: quasi-elastic (left), resonant pion production (center), and deep inelastic scattering (right). These are depicted as ν_μ CC interactions mediated by a W boson with an outgoing μ in the final state and different hadronic systems depending on the interaction mode.




Figure 2.3: Neutrino Cross Sections. Cross sections for each interaction mode as a function of neutrino energy for neutrinos (left) and antineutrinos (right).

individual interaction mode.

Neutrino experiments study a wide range of neutrino energies from MeV [18] to TeV [19] scales. Each interaction mode has its own energy dependence for the neutrino cross section, shown in Fig. 2.3. At small neutrino energies $E_\nu < 1$ GeV QE interactions are the most common. The cross section for RES peaks between 1 GeV to 3 GeV. At higher energies $E_\nu > 10$ GeV DIS interactions are completely dominant. For experiments studying neutrinos in the few GeV region, such as NOvA, all interaction modes are relevant and require careful study and consideration for reconstruction tools.

QE scattering is the scattering of a neutrino off a nucleon in either CC or NC processes. The final state is a single lepton, a charged lepton in CC or a neutrino in NC, and a single nucleon. The QE models are attributed to the formalism by Llewellyn Smith [20]. The momentum of the final state lepton is a function of Q^2 , the square of the four momentum transferred to the nucleon.

At neutrino energies around 2 GeV, the neutrino excites resonant states in the nucleon. RES interactions result in a single baryon and a meson via the intermediate hadronic excited state. There are several possible resonant states, but the most common is the Δ baryon which will decay to a nucleon and a single pion. RES interactions are predicted based on the model constructed by Rein and Sehgal [21]. The model includes the scattering, charge exchange, and absorption of pions as they traverse the nucleus.

At larger neutrino energies, the de Broglie wavelength is small compared to the size of the nucleon and the neutrino can penetrate deep within the nucleon to interact with a constituent quark. In low energy DIS interactions, the quark is knocked out of the nucleus and will hadronize, but at higher energies the quark is shattered emitting various hadrons. DIS interactions are modelled by the formalism from Bodek and Yang [22]. At extremely high neutrino energies $E_\nu > 100$ GeV DIS interactions are completely dominant with cross section linearly proportional to the neutrino energy.

MEC interactions have a peak cross section in the energy region between QE and RES interactions. Here, the neutrino scatters off a nucleon pair in some correlated state in an otherwise QE-like interaction. The nucleon pair can be a pair of neutrons nn (protons pp) for neutrinos (antineutrinos) or a neutron-proton pair for either case. Several theoretical models exist for MEC, but the one used in this analysis is in Ref. [23]. MEC interaction model parameters are not well constrained and have seen growing interest for inclusion in neutrino generators to explain discrepancies seen in cross section measurement data [24].

Cross sections are known to be small for all modes in the few GeV region, $\sigma \approx 10^{-38}$ cm², as

shown in Fig. 2.3. However, they are not well constrained with uncertainties around 20% [25]. These large uncertainties pose a challenge to modern neutrino experiments and often constitute one of the largest systematic uncertainties in precision neutrino measurements.

2.3 Neutrino Oscillations

The neutrino states in the SM in Fig. 2.1 have a definite flavor as electro-weak eigenstates, but do not have a definite mass. This means that the electroweak states, $|\nu_\alpha\rangle$, can be written as a linear combination of mass eigenstates, $|\nu_i\rangle$:

$$|\nu_\alpha\rangle = \sum_{i=1}^3 U_{\alpha i}^* |\nu_i\rangle \quad (2.5)$$

where α is a lepton flavor state e , μ , or τ and i is one of three mass states. U is a unitary rotation matrix from the mass eigenstate basis to the flavor eigenstate basis. It is called the Pontecorvo-Maki-Nakagawa-Sakata (PMNS) matrix named after Bruno Pontecorvo who originally proposed the idea [26] and the other three who cemented the idea of neutrino oscillations [27].

Each mass state evolves in time according to the Schrödinger equation as a plane wave:

$$|\nu_i(t)\rangle = e^{-iE_i t} |\nu_i(0)\rangle \quad (2.6)$$

where $E_i = \sqrt{p^2 + m_i^2}$ are the energy eigenvalues. Intuitively, a neutrino is produced with a definite flavor, but each mass state evolves independently accumulating a phase difference, and a different neutrino flavor may be detected at a later time.

Since neutrino energies are typically much larger than their masses we can expand $E_i = p + \frac{m_i^2}{2p}$.

The oscillation probability is then given by:

$$P_{\alpha \rightarrow \beta}(t) = |\langle \nu_\beta(t) | \nu_\alpha(0) \rangle|^2 = \left| \sum_{i=1}^3 U_{\alpha i}^* U_{\beta i} e^{-i \frac{m_i^2}{2p} t} \right|^2 \quad (2.7)$$

where the common phase ipt has been dropped. In the ultrarelativistic case, $t = L$ where L is the distance travelled and $p = E$ so we can write $\frac{t}{p} = \frac{L}{E}$. Using the unitarity of U , Eq. 2.7 can be expanded as:

$$P_{\alpha \rightarrow \beta}(L, E) = \delta_{\alpha\beta} - 4 \sum_{i>j} \mathcal{R}[U_{\alpha i}^* U_{\alpha j} U_{\beta i} U_{\beta j}^*] \sin^2 \left(\frac{\Delta m_{ij}^2 L}{4E} \right) \quad (2.8)$$

$$+ 2 \sum_{i>j} \mathcal{I}[U_{\alpha i}^* U_{\alpha j} U_{\beta i} U_{\beta j}^*] \sin^2 \left(\frac{\Delta m_{ij}^2 L}{2E} \right) \quad (2.9)$$

where $\Delta m_{ij}^2 \equiv m_i^2 - m_j^2$. Thus, neutrino oscillations are not sensitive to the absolute neutrino mass, but can make measurements of the difference in masses. Experimentally, Δm_{12}^2 has been determined to be much smaller than Δm_{32}^2 or Δm_{31}^2 . It is common to assume $\Delta m_{32}^2 = \Delta m_{31}^2$ to simplify the oscillation probabilities.

It is often useful to parameterize the PMNS matrix in terms of three mixing angles, θ_{12} , θ_{13} , and θ_{23} , and a complex phase δ_{CP} . In the expanded form:

$$U = \begin{pmatrix} 1 & 0 & 0 \\ 0 & c_{23} & s_{23} \\ 0 & -s_{23} & c_{23} \end{pmatrix} \begin{pmatrix} c_{13} & 0 & s_{13}e^{-i\delta_{\text{CP}}} \\ 0 & 1 & 0 \\ -s_{13}e^{i\delta_{\text{CP}}} & 0 & c_{13} \end{pmatrix} \begin{pmatrix} c_{12} & s_{12} & 0 \\ -s_{12} & c_{12} & 0 \\ 0 & 0 & 1 \end{pmatrix} \quad (2.10)$$

where $c_{ij} = \cos \theta_{ij}$ and $s_{ij} = \sin \theta_{ij}$. δ_{CP} is a phase which behaves differently for neutrinos and antineutrinos allowing for charge-parity (CP) violation. The expanded form divides the PMNS matrix into three components each dependent on a different mixing angle.

Thus, all neutrino oscillations can be determined by the three mixing angles, the three mass splittings, and δ_{CP} based on the energy and distance travelled by the neutrino. There are two neutrino oscillation channels that are relevant to the NOvA experiment, $\nu_\mu \rightarrow \nu_\mu$ and $\nu_\mu \rightarrow \nu_e$, which will be discussed in detail here.

2.3.1 Muon Neutrino Disappearance

The first channel of interest is $\nu_\mu \rightarrow \nu_\mu$, often referred to as ν_μ disappearance since the search is for some initial ν_μ to “disappear” or oscillate to a different flavor. Starting from Eq. 2.8, there is a significant simplification when $\alpha = \beta$ since the products of the PMNS elements are real and the last term can be ignored.

$$P_{\mu \rightarrow \mu} = 1 - \sin^2 2\theta_{23} \sin^2 \Delta_{31} + 4 \sin^2 \theta_{13} \sin^2 \theta_{23} \cos 2\theta_{23} \sin^2 \Delta_{31} \quad (2.11)$$

where $\Delta_{ij} = \frac{\Delta m_{ij}^2 L}{4E}$. The term proportional to $\sin^2 \Delta_{12}$ has been dropped since it involves a much longer oscillation length than $\sin^2 \Delta_{31}$.

For the results in this dissertation, the full oscillation probability is used, but it is instructive to simplify further. Since θ_{13} is the smallest of the three mixing angles, $P_{\mu \rightarrow \mu}$ is dominated by the term proportional to $\sin^2 2\theta_{23}$ and the probability is rewritten under the assumption $\Delta m_{32}^2 = \Delta m_{31}^2$:

$$P_{\mu \rightarrow \mu} = 1 - \sin^2 2\theta_{23} \sin^2 \Delta_{32}. \quad (2.12)$$

In this form, it is clear how to interpret the mixing angles and mass squared differences. The mixing angles are related to oscillation amplitudes in terms of trigonometric functions: $\sin^2 2\theta_{23}$. The mass squared differences are oscillation frequencies in terms of the ratio L/E : $\sin(\Delta m_{32}^2 \frac{L}{E})$. Neutrino experiments usually have a fixed L and sample different energies. This is shown pictorially in Fig. 2.4.

The ν_μ disappearance channel depends on a limited number of oscillation parameters, in particular it has no dependence on δ_{CP} . The SM assumes Lorentz invariance which means that CPT, the combination of charge (C), parity (P), and Time Reversal (T) is a conserved symmetry. The CP transformation interchanges neutrinos with antineutrinos $\nu_\alpha \rightarrow \nu_\beta \xrightarrow{CP} \bar{\nu}_\alpha \rightarrow \bar{\nu}_\beta$. The T transformation reverses the initial and final states $\nu_\alpha \rightarrow \nu_\beta \xrightarrow{T} \nu_\beta \rightarrow \nu_\alpha$. Clearly in the case of a survival

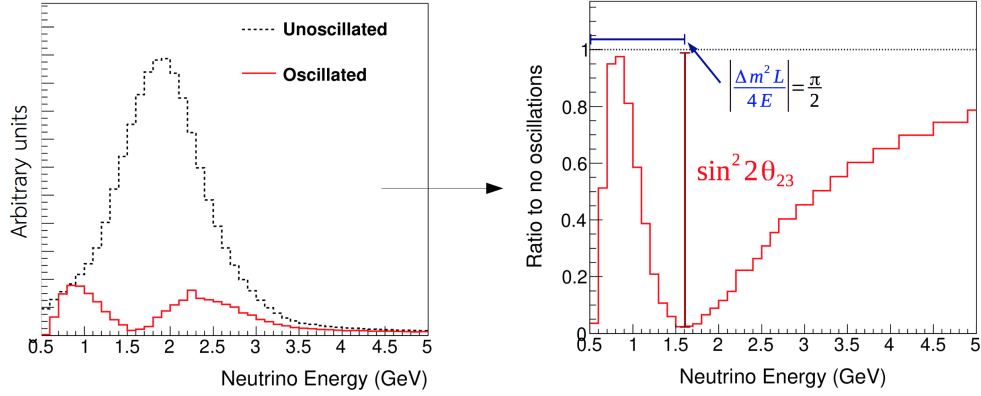


Figure 2.4: Neutrino Oscillation Ratio. Left: An arbitrary neutrino energy distribution under the assumption with and without neutrino oscillations. Right: The ratio of the oscillated distribution to the unoscillated distribution. The amplitude of oscillation is given by the mixing angle θ_{23} . The frequency of oscillation and the point of maximum oscillation is determined by Δm_{32}^2 .

probability where $\alpha = \beta$, CPT symmetry implies:

$$P(\nu_\alpha \rightarrow \nu_\alpha) = P(\bar{\nu}_\alpha \rightarrow \bar{\nu}_\alpha) \quad (2.13)$$

and only appearance measurements where $\alpha \neq \beta$ can probe the value of δ_{CP} .

2.3.2 Electron Neutrino Appearance

The other channel of interest is $\nu_\mu \rightarrow \nu_e$ where some ν_e have “appeared” from the initial ν_μ . This channel depends on all three mixing angles and δ_{CP} . Once again, starting from Eq. 2.8:

$$P_{\mu \rightarrow e} = P_{\text{atm}} + P_{\text{sol}} + 2\sqrt{P_{\text{atm}}P_{\text{sol}}}(\cos \Delta_{32} \cos \delta_{\text{CP}} \mp \sin \Delta_{32} \sin \delta_{\text{CP}}) \quad (2.14)$$

with

$$P_{\text{atm}} = \sin^2 \theta_{23} \sin^2 2\theta_{13} \sin^2 \Delta_{31} \quad (2.15)$$

$$P_{\text{sol}} = \cos^2 \theta_{23} \cos^2 \theta_{13} \sin^2 2\theta_{12} \sin^2 \Delta_{21} \quad (2.16)$$

Parameter	Best Fit	Hierarchy, Octant
Δm_{12}^2	$(7.53 \pm 0.18) \times 10^{-5} \text{eV}^2$	
Δm_{32}^2	$(-2.56 \pm 0.04) \times 10^{-3} \text{eV}^2$	IH
$\Delta m_{\bar{3}2}^2$	$(2.51 \pm 0.05) \times 10^{-3} \text{eV}^2$	NH
$\sin^2 \theta_{12}$	0.307 ± 0.013	
$\sin^2 \theta_{13}$	$(2.12 \pm 0.08) \times 10^{-2}$	
$\sin^2 \theta_{23}$	$0.421^{+0.033}_{-0.025}$	IH, LO
$\sin^2 \theta_{23}$	$0.592^{+0.023}_{-0.030}$	IH, UO
$\sin^2 \theta_{23}$	$0.417^{+0.025}_{-0.028}$	NH, LO
$\sin^2 \theta_{23}$	$0.597^{+0.024}_{-0.030}$	NH, UO

Table 2.1: The current best fit values of the neutrino oscillation mixing angles and mass squared differences. The third column shows whether the measurement was made for the inverted or normal hierarchy and the upper or lower octant. Values taken from Ref. [28].

where the $-$ is for neutrinos and the $+$ for antineutrinos since $\delta_{\text{CP}} \rightarrow -\delta_{\text{CP}}$ [10].

Equation 2.14 gives the complete oscillation probability for $\nu_\mu \rightarrow \nu_e$ in a vacuum. However, consideration must be given to the passage of neutrinos through matter and discussion of these equations will be left until after the full derivation.

2.3.3 Current Parameters and Open Questions

The three mixing angles and mass splittings have all been measured to better than 10% precision. Of most importance is that all mixing angles are known to be non-zero. The current values are shown in Table 2.1.

δ_{CP} is a cyclic phase linked to CP violation. Values of 0 or π are CP conserving and values of $\pi/2$ or $3\pi/2$ are maximal CP violating. Measurements are not well constrained, but there is evidence which rules out the CP conserving values [29].

Δm_{12}^2 is positive by definition,² but the sign of Δm_{32}^2 is not currently known. There are two possibilities: the positive case known as the Normal Hierarchy (NH) and the negative case known as the Inverted Hierarchy (IH). The hierarchies are also sometimes referred to as the orderings (NO and IO) and will be used interchangeably in this dissertation. The neutrino mass states, by

²This choice is associated with a question of whether $\theta_{12} < \pi/4$ which is now known to be true.




Figure 2.5: Neutrino Mass Ordering. The two possible neutrino mass orderings depending on the sign of Δm_{32}^2 . The inverted ordering (left) with $m_2 > m_1 > m_3$ and the normal ordering (right) with $m_3 > m_2 > m_1$.

convention, are ordered by decreasing ν_e content, so m_1 has the most ν_e content and m_3 has the least. The normal hierarchy is so called since the neutrino mass eigenstates would follow a normal order with $m_3 > m_2 > m_1$ making the state with the most ν_e the lightest just as $m_\tau > m_\mu > m_e$. The two hierarchies are shown in Fig. 2.5.

θ_{23} is the largest of the three mixing angles and may even be maximal. Equation 2.12 shows that the ν_μ disappearance probability is maximal when $\sin^2 2\theta_{23} = 1$ or $\theta_{23} = \pi/4$. However, whether θ_{23} is exactly this maximal value, slightly less than maximal, known as the lower octant (LO), or slightly more than maximal, known as the upper octant (UO), is an open question. Notably, a maximal θ_{23} would imply a symmetry between ν_μ and ν_τ in the ν_3 mass eigenstate. The ν_μ disappearance channel is not sensitive to the octant, but the ν_e appearance channel, which has dependencies on $\sin^2 \theta_{23}$ rather than $\sin^2 2\theta_{23}$, can resolve the degeneracy.

2.3.4 Matter Effects

Even though neutrinos have low cross sections, they are still affected by passing through large quantities of matter, such as the Earth. There is a coherent forward scattering amplitude of

neutrinos with electrons. This can be by NC interactions, but this affects all flavors identically and can be ignored. However, only ν_e will forward scatter by CC interactions. Thus, there is an additional potential which affects only the ν_e component when passing through matter which changes the effective mass of the neutrinos and, thus, the time evolution. Any oscillations involving ν_e will need to be adjusted to account for this.

The effective potential from matter is introduced into the hamiltonian as an additional term

$$H = U \begin{pmatrix} \frac{m_1^2}{2E} & 0 & 0 \\ 0 & \frac{m_2^2}{2E} & 0 \\ 0 & 0 & \frac{m_3^2}{2E} \end{pmatrix} U^\dagger + \begin{pmatrix} \pm\sqrt{2}G_F N_e & 0 & 0 \\ 0 & 0 & 0 \\ 0 & 0 & 0 \end{pmatrix} \quad (2.17)$$

where G_F is the Fermi constant and N_e is the number density of electrons in the medium.³ The $+$ is for neutrinos and the $-$ is for antineutrinos. The new Hamiltonian can be re-diagonalized to obtain a new set of neutrino mass eigenstates.

With the new Hamiltonian, Eq. 2.14 is unchanged, but the definition of P_{atm} and P_{sol} need to be updated [10]:

$$P_{\text{atm}} = \sin^2 \theta_{23} \sin^2 2\theta_{13} \frac{\sin^2(\Delta_{31} \mp aL)}{(\Delta_{31} \mp aL)^2} \Delta_{31}^2 \quad (2.18)$$

$$P_{\text{sol}} = \cos^2 \theta_{23} \cos^2 \theta_{13} \sin^2 2\theta_{12} \frac{\sin^2 \mp aL}{(\mp aL)^2} \Delta_{21}^2 \quad (2.19)$$

where $a = G_F N_e / \sqrt{2}$ and the $-$ is for neutrinos and the $+$ for antineutrinos.

With the full oscillation probability in hand, it is instructive to compute some values to see the effect of the different parameters on the oscillation probability. Figure 2.6 shows the $\nu_\mu \rightarrow \nu_e$ and $\bar{\nu}_\mu \rightarrow \bar{\nu}_e$ oscillation probabilities under different assumptions with experimental parameters from the NOvA experiment, $L = 810$ km and $E = 2$ GeV. Varying δ_{CP} traces out a clockwise (counter-clockwise) ellipse for the NH (IH). The ellipses for the NH and IH have been split from the diagonal,

³ N_e is assumed to be constant for this derivation. This is reasonable for neutrinos traversing the Earth's crust, but may not be for other media such as neutrinos exiting the Sun's core.

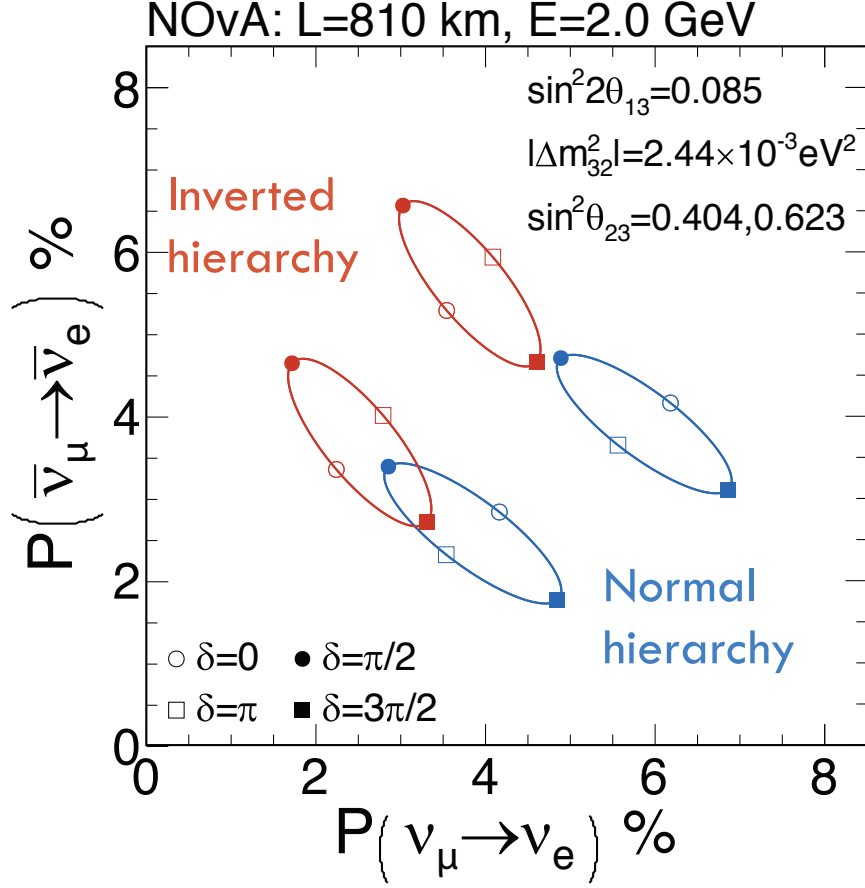


Figure 2.6: ν_e Biprobability. The $\nu_\mu \rightarrow \nu_e$ (X-axis) and $\bar{\nu}_\mu \rightarrow \bar{\nu}_e$ (Y-axis) oscillation probabilities. The baseline and neutrino energy are chosen from the NOvA experiment and the oscillation parameters are generic values. The four ellipses are combinations of the NH or IH and UO or LO. The NH and IH are separated due to the matter effect in the oscillation probability. The bottom left (top right) two ellipses are constructed with θ_{23} in the LO (UO). The ellipses are traced by changing the value of δ_{CP} from 0 to 2π .

where the neutrino and antineutrino probability are equal, by the matter effect. Moving from the LO to the UO has the effect of enhancing or diminishing the oscillation probability. Thus, if the two oscillation probabilities can be measured, which of the four ellipses tells the NH or IH and the LO or UO. The position on the ellipse depends on the value of δ_{CP} .

The matter effect changes all three mass eigenstates and would have an effect on other oscillation channels including $\nu_\mu \rightarrow \nu_\mu$. However, for experiments in the energy range of NOvA the effects are much less than 1% [30].

2.4 Current Experiments

Neutrino experiments are divided into different categories based on their neutrino source, neutrino energy, and baseline. The ratio L/E , typically in km/GeV, is the experimental property which determines which oscillation parameters an experiment is sensitive to. Each of the three matrices in the expanded PMNS matrix, Eq. 2.10, depends on a different mixing angle and is associated with a different category of experiments. These are atmospheric/accelerator, reactor, and solar neutrino experiments respectively.

Atmospheric experiments measure the flux of neutrinos produced in the decay of pions from cosmic rays. The flux is measured both for neutrinos above the detector ($L \approx 10$ km) and neutrinos which traversed the earth from below ($L \approx 10^4$ km). The most notable atmospheric neutrino experiment was the Super-Kamiokande (SuperK) experiment in Japan [31]. SuperK observed a deficit of neutrinos coming from below which was attributed to the oscillation of ν_μ into ν_τ since there was only a small change in the ν_e flux compared to models [7]. Atmospheric experiments are mainly sensitive to θ_{23} and Δm_{32}^2 .

Accelerator experiments also study the atmospheric oscillation parameters, but instead by producing a beam of neutrinos from a particle accelerator. There is more control over the neutrino baseline, energy, and flavor composition of the neutrino source with this method than any other. Typical experiments have energies in the range 10^{-1} GeV to 10^1 GeV. Baselines vary, with short baseline experiments having $L \approx 1$ km and longer baselines varying from 10^2 km to 10^3 km. The short baseline experiments are sensitive to high frequency neutrino oscillations with large Δm^2 . Notable short baseline experiments include LSND [32] and the upcoming SBN program [33]. Long baseline experiments have L/E with sensitivity to the atmospheric oscillation parameters. Through appearance channels, these experiments can also measure δ_{CP} , the mass hierarchy, and the θ_{23} octant. NOvA is one example of a long baseline experiment. Other notable experiments include T2K [34], MINOS [35], and the upcoming DUNE experiment [36]. The DUNE experiment has a

baseline of 1300 km which provides a very large matter effect to neutrino oscillations. Figure 2.6 shows a degeneracy between IH with $\delta_{\text{CP}} = 3\pi/2$ and NH with $\delta_{\text{CP}} = \pi/2$ which will be broken by the increased baseline in DUNE.

Reactor experiments measure the disappearance of $\bar{\nu}_e$ produced by β -decay within nuclear reactors. This gives sensitivity primarily to θ_{13} . The measurement of θ_{13} is significant as it was the last of the three angles to be measured and all three need to be non-zero for CP violation to be possible. Reactor experiments typically have baselines of $L \approx 1$ km and energies in the range of $E = 10^{-3}$ GeV to 10^{-2} GeV. In this region of L/E , the oscillation probability greatly simplifies to

$$P(\bar{\nu}_e \rightarrow \bar{\nu}_e) = 1 - \sin^2 2\theta_{13} \sin^2 \Delta_{31} \quad (2.20)$$

which allows for precision measurements of θ_{13} and it is now the most precisely measured of the three mixing angles. The most notable reactor experiments are Double Chooz [37], Daya Bay [38], and RENO [39]

The final category of experiments are solar neutrino experiments which detect ν_e produced by fusion within the Sun. Due to matter effects within the core of the Sun, the ν_e flavor state is nearly identical to the ν_2 mass state, so the flux from the Sun observed at Earth is $|\nu_2\rangle = \sin\theta_{12}|\nu_e\rangle + \cos\theta_{12}|\nu_X\rangle$. Today, we know to expect less than one-third of the predicted ν_e flux since $\sin^2\theta_{12} = 0.307$, but this was the source of the original solar neutrino problem discussed in Chap. 1. The SNO experiment [8] definitively solved the solar neutrino problem by detecting the NC interactions from any flavor, which matched the prediction, not just the ν_e CC interactions. Other notable solar experiments are Kamiokande [40], SAGE [41], and GALLEX [42].

Chapter 3

The NOvA Experiment

NOvA is a two detector, long baseline, neutrino oscillation experiment [43]. NOvA is primarily sensitive to four oscillation channels of neutrinos and antineutrinos: $\bar{\nu}_{\mu} \rightarrow \bar{\nu}_{\mu}$, known as muon neutrino (antineutrino) disappearance, and $\bar{\nu}_{\mu} \rightarrow \bar{\nu}_{e}$, known as electron neutrino (antineutrino) appearance. The design of the experiment is motivated by three main physics goals:

- Precision measurements of the atmospheric mixing angle, θ_{23} . θ_{23} is a measurement of the relative amounts of ν_{μ} and ν_{τ} in the ν_3 mass state.
- Determination of the neutrino mass hierarchy. This is equivalent to measuring the sign of Δm_{32}^2 and resolves which of the three neutrino mass states is the heaviest.
- Search for CP violation. Chapter 2 describes how neutrino oscillations can probe the δ_{CP} parameter.

The following sections summarize the experimental setup of NOvA and introduce some common terms which will be used throughout the remainder of this dissertation. The neutrino beam source, the design of the two detectors, and the data acquisition system used to collect data will be discussed. The chapter concludes with a discussion of the simulation chain for neutrino interactions in NOvA and the calibration of the detector.

3.1 The NuMI Beam

The neutrino source used by the NOvA experiment is the Neutrinos at the Main Injector (NuMI) beamline from Fermi National Accelerator Laboratory (Fermilab). The NuMI beam was originally designed for the MINOS experiment [44] which began data collection in 2005.

The beam is produced by extracting protons accelerated to 120 GeV by the Main Injector [45]. The protons are grouped into packets of $\sim 4 \times 10^{12}$ protons called bunches. The NuMI beam pro-

duces a pulsed stream of protons with a spill of 12 bunches taking 10 μs and a total cycle time of $\sim 1.4\text{s}$, with the exact time depending on the status of the accelerator and the list of currently active experiments.

The collision of these protons with a graphite target produces many different charged mesons, primarily pions and kaons, which then decay to produce neutrinos:

$$\pi^\pm \rightarrow \mu^\pm + \bar{\nu}_\mu^{(-)} \quad (3.1)$$

$$K^\pm \rightarrow \mu^\pm + \bar{\nu}_\mu^{(-)}. \quad (3.2)$$

The mesons enter a decay pipe which is 675 m long, approximately the distance needed for a 10 GeV pion to decay. Following the decay pipe, an absorber stops the residual hadronic components of the beam. The beam and detector are separated by 240 m of rock with a series of muon detectors to monitor beam quality. For a complete description of the NuMI beamline see [46]. Note that muons primarily decay by

$$\mu^\pm \rightarrow e^\pm + \bar{\nu}_e^{(-)} + \bar{\nu}_\mu^{(-)} \quad (3.3)$$

which produces an intrinsic ν_e contamination to the primarily ν_μ beam when the muons are not absorbed.

A pair of focusing horns placed after the target increase the acceptance of hadrons into the decay region. These horns have a current running through them producing a magnetic field which acts to focus one sign of the mesons and defocus the other. By reversing the current in the horn and, thus, the magnetic field, the beam can be configured to focus the positively charged mesons with a Forward Horn Current (FHC) and thus producing ν_μ or focusing negatively charged mesons with a Reverse Horn Current (RHC) and producing $\bar{\nu}_\mu$. The NuMI beamline is depicted in Fig. 3.1.

The flux and energy of the resulting neutrinos from the meson decay in flight at a distance z




Figure 3.1: The NuMI Beamline. 120 GeV protons extracted from the Main Injector collide with a graphite target (left of image). The produced mesons are focused into a beam before decaying into muons and muon neutrinos. The absorber and rock act to remove the muons from the beam, leaving primarily muon neutrinos. The magnetic field in the horn can be reversed to produce primarily muon antineutrinos [46].

in a detector of area A are given by

$$\phi = \left(\frac{2\gamma}{1 + \gamma^2\theta^2} \right)^2 \frac{A}{4\pi z^2}, \quad (3.4)$$

$$E_\nu = E_{(\pi,K)} \frac{1 - \frac{m_\mu^2}{m_{(\pi,K)}^2}}{1 + \gamma^2\theta^2} \quad (3.5)$$

where θ is the angle between the meson and outgoing neutrino and $\gamma = \frac{E_{(\pi,K)}}{m_{(\pi,K)}}$ is the relativistic Lorentz factor. These equations show that the neutrino energy range is narrow for a wide range of initial meson energies for small deflection angles from the beam axis. Neutrino experiments can utilize this by positioning detectors slightly off-axis to constrain the energy spectrum. The NOvA detectors are located 14.6 mrad off the primary beam axis which has a peak energy around 2 GeV. Neutrino energies and the observed spectrum in NOvA are shown in Fig. 3.2.

There are several benefits of having the detector off-axis. First, about five times more neutrinos are produced around 2 GeV than in the off-axis spectrum. This is ideal since the maximum oscillation probability of $\nu_\mu \rightarrow \nu_e$ at 810 km is at an energy of ~ 1.6 GeV, thus increasing the number of ν_e events observed. Second, backgrounds are reduced by the narrow spectrum which reduces the flux of background high energy neutrinos. High energy NC events can often mimic the ν_e CC




Figure 3.2: The Off-Axis Spectrum. Left: The energy of neutrinos produced at angle θ relative to the pion direction. Right: The ν_μ spectrum expected at 810 km from the beam source at different off-axis angles. The NOvA detectors were chosen to be 14.6 mrad off-axis from the beam to constrain the energy spectrum to peak near the maximum oscillation probability for $\nu_\mu \rightarrow \nu_e$.

signal through $\pi^0 \rightarrow \gamma + \gamma$ decays. Since the departing neutrino carries some of the energy out of the detector, these NC events can be misidentified as signal ν_e CC events in the energy region of interest. However, the narrow spectrum means these often have energies outside the signal region. Another background of concern is the intrinsic ν_e in the beam from three-body muon and kaon decays. These events have a much broader spectrum than the narrow ν_e appearance signal. Thus, proper energy reconstruction can be used to reject background events from outside the energy region of interest.

As mentioned above, the two horns can be configured to focus either charge of the produced mesons. FHC mode produces a beam of primarily ν_μ , while RHC mode produces a beam of primarily $\bar{\nu}_\mu$. Fig. 3.3 shows the expected composition of events in both modes in the far detector corrected for flux and cross sections. For FHC, the wrong sign contamination from $\bar{\nu}_\mu$ is about 1.7% and the intrinsic ν_e and $\bar{\nu}_e$ background is about 0.6%. For RHC, the wrong sign contamination from ν_μ is about 11.3%. The larger contamination is due to the smaller cross section of antineutrinos compared to neutrinos.




Figure 3.3: Projected Neutrino Events. The projected events of $(\bar{\nu})_{\mu}$ and $(\bar{\nu})_{e}$ of the beam in the forward (left) or reverse (right) horn current mode. The counts have been corrected for the flux of the beam and neutrino cross sections in the detector, but no oscillations have been applied. The reverse horn current spectrum has a larger background from the wrong sign of neutrinos due to the reduced cross section of antineutrinos.

The beam exposure recorded by NOvA is measured in the quantity Protons On Target (POT), the number of protons that collide with the target. This dissertation reports an analysis of 13.6×10^{20} POT in FHC mode and 12.5×10^{20} POT in RHC beam mode. Fig. 3.4 shows the history of POT recorded by NOvA. NOvA has maintained more than 98% uptime for recording beam POT in recent years.

3.2 NOvA Detectors

The NOvA experiment consists of two detectors: a near detector (ND) located 1 km from the beam source and a second far detector located 810 km from the beam source. The ND is used to detect neutrinos before any oscillations have occurred and to constrain systematic uncertainties which impact the experiment. The FD is used to detect neutrinos after oscillations have occurred. The FD is much larger in size than the ND, but otherwise the two detectors are functionally equivalent. The two detector design with common technology helps to cancel many systematic uncertainties in the neutrino flux, neutrino cross-sections, and detector performance. The treatment of systematic uncertainties will be described in Chap. 7.

Both detectors consist of extruded PVC cells filled with liquid scintillator. The main particle de-

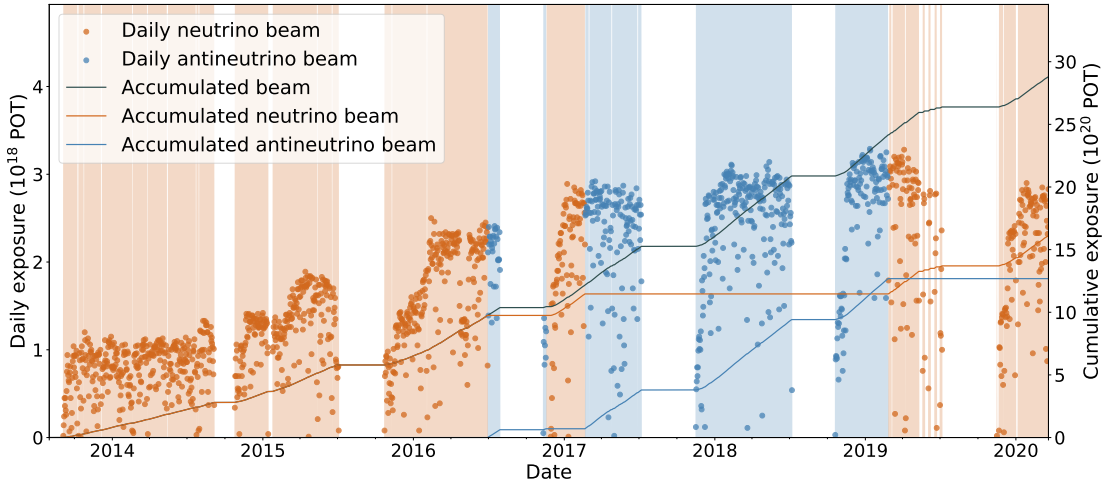


Figure 3.4: Daily POT Recorded by NOvA. The points show the daily POT recorded by the NOvA experiment since 2014. The orange regions are time periods where the beam was in neutrino mode and the blue regions are in antineutrino beam mode.

tection mechanism is scintillation light produced by charged particles in the liquid scintillator. The solution is 94.6% mineral oil, 5.2% pseudocumene (scintillator), 0.1% PPO (wavelength shifter), 0.002% bis-MSB (wavelength shifter), 0.001% stadis-425 (anti-static), and 0.001% vitamin-E (anti-oxidant) by mass [47]. The solution produces light in the ultra-violet spectrum and shifts it to the visible region of 400 nm to 450 nm, violet light. The emitted light reflects off the surface of the PVC of each cell which is coated in TiO_2 . Each cell is equipped with a wavelength shifting fiber loop which absorbs the emitted light and shifts the wavelength to the blue-green, 450 nm to 600 nm. Both ends of the fiber are read by a pair of photodetectors on an Avalanche Photodiode (APD) which has a quantum efficiency of 85% for green light. High quantum efficiency is needed to have long cells in the FD, but still record a minimum of 20 photoelectrons (PE). Each APD has enough photodetectors for 32 cells. The APDs are kept at -15°C to reduce thermal noise and have dry nitrogen gas pumped into their enclosure to prevent buildup of ice. A picture of an APD is shown in Fig. 3.5.

The cells themselves are $3.9\text{ cm} \times 6.6\text{ cm}$ with a wall thickness of 2 mm to 5 mm depending




Figure 3.5: A Detector Cell and APD. Left: The ends of the wave length shifting fibers within each cell. Each pair of ends is a single fiber with 32 fibers in total. Center: The front face of an APD which will be pressed against the fiber ends. Right: A schematic of a detector cell. Scintillation light is emitted when a charged particle passes through the material. The light is carried to the APD by a wavelength shifting fiber loop.

on needs for structural integrity. The cells are extruded in length to 15.8 m (3.8 m) for the FD (ND). Modules of 32 cells with a common APD are glued together side by side to form planes of 15.8 m \times 15.8 m (3.8 m \times 3.8 m) for the FD (ND). The planes are then glued back to back in alternating vertical and horizontal orientations to give two orthogonal views of the energy depositions within the detector. A diagram of one cell is shown on the right in Fig. 3.5. The detector is 65% active scintillator and 35% PVC by mass.

The detector is constructed from low Z materials, both PVC and mineral oil are carbon based, to aid in the identification of electrons. The materials have an average radiation length of $X_0 = 35.8$ cm and Moliere radius of $R_M = 11$ cm. The Moliere radius is sufficient to distinguish wide electromagnetic showers from non-showering particles. The radiation length is the average distance a photon will travel before pair producing and initiating a shower aiding in the electron identification. This will be discussed in more detail in Chap. 4.



Figure 3.6: A map of NOvA. A map showing the locations of the two detectors used by NOvA. The near detector is located on-site at Fermilab and the far detector is located near Ash River, Minnesota separated by 810 km.

The interaction length of the medium is $\lambda_I = 90$ cm and the collision length is $\lambda_C = 60$ cm. λ_I and λ_C characterize the propagation of hadrons in the detector [48].

A map showing the location of the two detectors is shown in Fig. 3.6 and pictures of the detectors themselves are shown in Fig 3.7. The ND is located 105 m underground on the Fermilab site within the MINOS underground area. It consists of 20192 cells and 192 planes, the last 22 of which are interlaced with 10 cm thick steel sheets to aid in the containment of muons. For financial reasons, these final planes are also about two thirds the height of the rest of the detector. The total dimensions are $3.8 \text{ m} \times 3.8 \text{ m} \times 16 \text{ m}$ with a total mass of 300 tons. Neutrinos enter the detector at a downward angle of about 3° .

The far detector is located on the surface near Ash River, Minnesota and the Canadian border, the furthest site from the beam source accessible within the United States. It consists of 896 planes, in sets of 64 called a diblock, and 344064 cells. The total dimensions of the far detector



Figure 3.7: The NOvA Detectors. The NOvA near (far) detector on the left (right).

are $15.8\text{ m} \times 15.8\text{ m} \times 60\text{ m}$ with a total mass of 14kton. Just as in the ND, neutrinos enter the detector at an angle, but now upwards at about 3° .

The detector is capable of operating with only part of the detector active. Data taking began with only 4 of the 14 diblocks and even after construction of the full detector was completed, some periods are missing one or more diblocks. The 13.6×10^{20} POT recorded by the experiment for FHC beam mode stated previously has been corrected from the total POT, 14.2×10^{20} POT, for the fraction of active detector at the time of recording. Almost all RHC data has been recorded with a complete detector and the 12.5×10^{20} POT in RHC beam mode only has a trivial correction. For the remainder of this dissertation, 13.6×10^{20} POT and 12.5×10^{20} POT for FHC and RHC beam mode respectively will be used without further clarification.

Since the far detector is located on the surface, it has a cosmic-ray flux of about 150 kHz, which are backgrounds to the neutrino analyses. Cosmic-ray muons can appear as backgrounds to the selection of ν_μ and cosmic-ray photons and neutrons can appear as a background to a ν_e selection. To shield from these backgrounds, the detector enclosure is surrounded by an overburden composed of 120 cm of concrete and 14 cm of barite. Barite is a high-Z material which will stop

many cosmic-ray photons before they can enter the detector. In total, the overburden provides roughly 12 radiation lengths of material to reduce the photon background.

3.3 Data Acquisition System

The Data Acquisition (DAQ) system is responsible for converting signals from the APDs into blocks of data that can be analysed offline. The detector and DAQ are designed to be capable of continuous readout with 100% livetime.

The signal output from each APD is read and digitized by a Front-End Board (FEB) with each FEB servicing one APD (and thus 32 cells). The FEB converts the PE count to a digital signal (ADC). A group of 64 FEBs transmit timestamped data to a single Data Concentrator Module (DCM). The DCM combines together data from each FEB into 50 μ s blocks called a microslice. There are a total of 168 DCMs on the FD, 12 for each diblock, and 14 DCMs for the ND. The data from the DCM is sent to a buffer farm, with temporary storage, where each microslice can be processed for trigger decisions, whether to keep or throw away the microslice. The buffer farm can store up to 20 minutes of data depending on the number of active nodes and the number of triggers that need processing. If a microslice is kept, the set of coincident microslices from all DCMs for the deciding trigger, called an event, are sent to a data logger which puts the event into permanent storage where it can be analysed offline. An overview of the DAQ is shown in Fig. 3.8.

The DAQ is capable of processing several types of triggers.

- Clock Triggers: triggers with a specific time interval.
- Signal Triggers: the daq receives an external trigger to readout data.
- Data-Driven Triggers (DDT): look for specific conditions within the detector.

The triggers which will be used within this dissertation are below along with other triggers of note.

- DDActivity Trigger. A DDT in the ND which reads out 50 μ s of data whenever there is activity. Used to calibrate the ND and monitor running conditions.

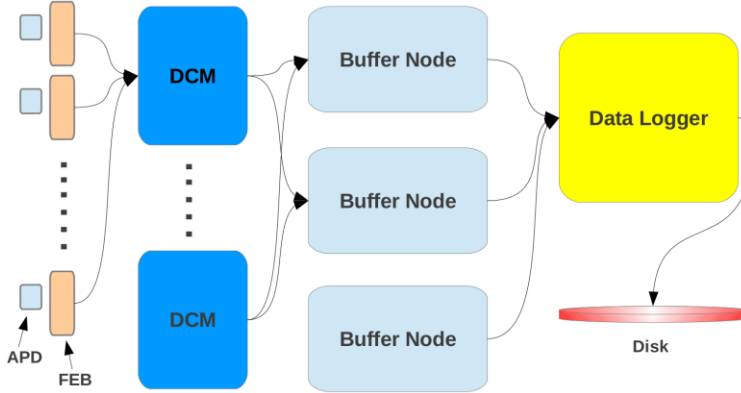


Figure 3.8: Overview of the DAQ. Dataflow from the APDs to permanent storage. Triggers are processed within the buffer nodes.

- Cosmic Trigger. A 10 Hz clock trigger in the FD which reads out $550 \mu\text{s}$ of data. There is too much activity, particularly from cosmic rays, in the FD for an activity trigger. Used as minimally biased data to calibrate the FD and to estimate cosmic backgrounds for neutrino analyses.
- NuMI Trigger. A signal trigger on the beam spill from the accelerator at Fermilab. The readout is $550 \mu\text{s}$ centered on the $10 \mu\text{s}$ beam spill.
- DDEnergy Trigger. A DDT in the FD which triggers when the total ADC in the detector has exceeded a set threshold. Used to look for high energy cosmic events.
- Supernova Trigger. A DDT designed to look for supernova candidates in both detectors. NOvA is also capable of triggering on a signal from the Supernova Early Warning System (SNEWS) [49].

An example of an event recorded from the cosmic trigger stream is shown in Fig. 3.9.

3.4 Timing System

The timing system is a crucial component of the DAQ. Proper timing in the detectors is necessary both externally with the neutrino beam clock to detect beam spills and internally to maintain synchronization between all detector electronics. As has been discussed in the previous sections,




Figure 3.9: A Cosmic Trigger Event. Each hit shown here represents the charge deposited within a single cell by the color scale. The top (bottom) half of the image shows hits recorded from the cells with vertical (horizontal) orientation. All of the activity seen in the event is from cosmic rays.

the timing of the neutrino beam varies between 1 s to 1.5 s depending on accelerator and experiment conditions. For this reason, a simple clock trigger is not sufficient to record events from the neutrino beam. In addition, the FD has significant activity from cosmic rays, so a data-driven activity based trigger is also not sufficient (the FD sees about 1-2 neutrinos per day compared to the 150 kHz of cosmic activity). Instead, when a beam spill is generated from Fermilab, it is time stamped by the NOvA clock, corrected for time of flight ($2700.56 \mu\text{s}$ to the FD), and a network packet is sent to both detectors to read out the corresponding window of time from the DAQ buffer, $550 \mu\text{s}$ centered on the $10 \mu\text{s}$ beam spill. For this reason, both detectors must be synced in time to each other and an external wall clock at the accelerator in order to properly trigger on the NuMI beam signal. Further, in order to ensure the possibility of reconstruction across the entire FD volume and reduce the pileup of events in the ND, individual components at each detector must be synced to within 10 ns of each other.

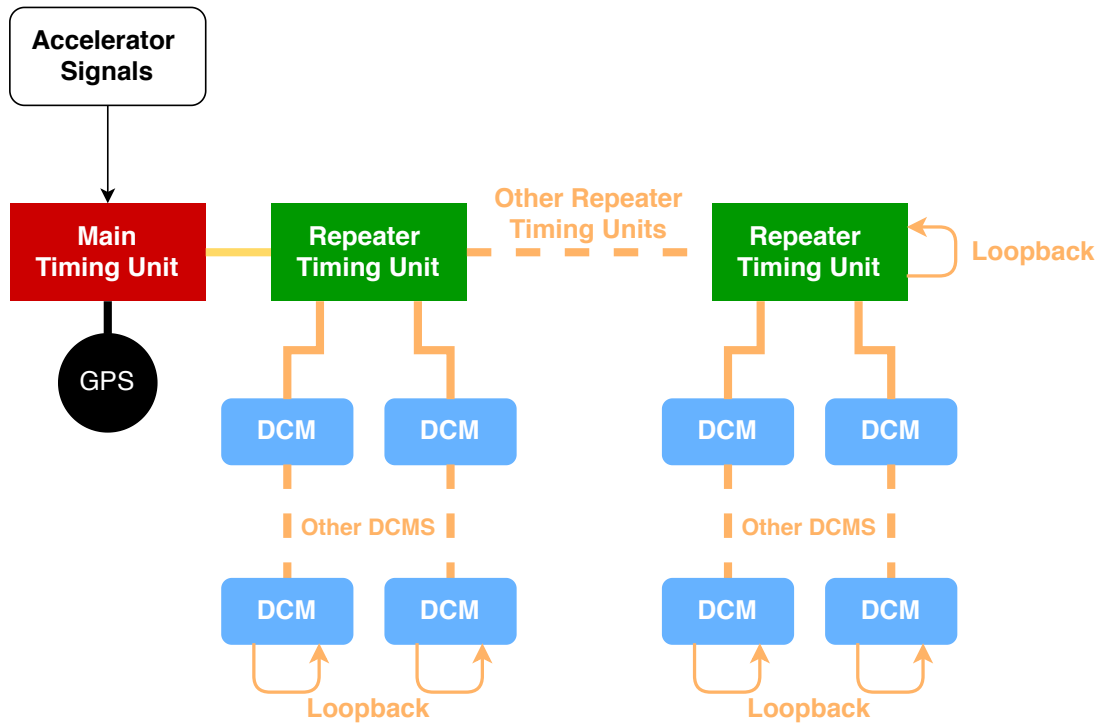


Figure 3.10: Timing System Schematic. The timing system uses a distributed system to sync detector components to one another. The main timing unit (left of image) is connected to a GPS to keep time and send signals and commands to the other components of the system. A backbone of repeater timing units are connected to the individual components of the detector.

The timing system is controlled by a chain of Timing Distribution Units (TDU) controlled by a main unit. The main TDU is connected to a GPS which keeps time for the experiment and relays commands to a series of repeater TDUs connected to the individual detector components. Each repeater has two subchains: one for 6 DCMs along the top of the detector and one for 6 DCMs along the side of the detector at the FD. Each chain in the system is terminated by a loopback used for calibration. The DCMs fan out to the 64 FEBs to which they are connected. A diagram of the system is shown in Fig. 3.10. Most components are connected in a daisy chain of copper cables with 4 LVDS lines to carry the clock, commands, sync signal, and sync return. The one exception is the connection between the main TDU on the surface and the first repeater TDU underground at the ND which is connected by a fiberoptic cable spanning the distance between the surface and the detector.

The NOvA clock is a 64 MHz clock derived in a phase lock loop from the 10 MHz clock on the GPS receiver. Every TDU, DCM, and FEB has its own clock register which is driven by the main TDU. The NOvA epoch is defined as the number of these 64 MHz clock cycles since 00:00 1 January 2010 GMT. Therefore, one unit of NOvA time corresponds to 15.625 ns. The registers used are 56 bits, allowing for a validity of 35.7 years.

In the event that GPS lock loss occurs, the 10 MHz clock remains stable to 2 parts per billion per day. Both detectors also have access to a separate timing reference GPS, the Timing Calibration Reference (TCR). The TCR acts as a reference to monitor clock drift.

In order to synchronize the time registers between all the components, each TDU must learn the delays in the copper wires to the next component. The main TDU initiates the learning process by computing a delay offset value, which is arbitrary but must be larger than the maximum time to the end of any chain. The main TDU then sends a sync command down the chain. Upon receiving the command, each component clears its current time and begins a counter which stops when the return sync echo signal is received from the loopback at the end of the chain. Each repeater TDU then stores half of the counter, the time of flight to the end of the chain. Each repeater TDU also keeps a counter for the time of flight down each DCM branch. The main TDU stores the total delay down the TDU chain plus the initial delay offset value called the “early sync” time, the time before a GPS 1 second boundary such that a sync can be issued and complete on the 1 second boundary. Note that for technical reasons, DCMs learn their delay value but are unable to store it locally as there is no non-volatile memory. Timing corrections to DCMs must be applied to data offline [50]. In addition, FEBs do not learn their delay at all. Instead, strict requirements were placed on the length of cable between an FEB and its DCM.

Once the delays are learned the system can be synchronized at anytime. The main TDU compares the current time, the early sync time, and the upcoming 1 second boundaries to determine a time in the future for which the sync can complete. The future time is sent to all components

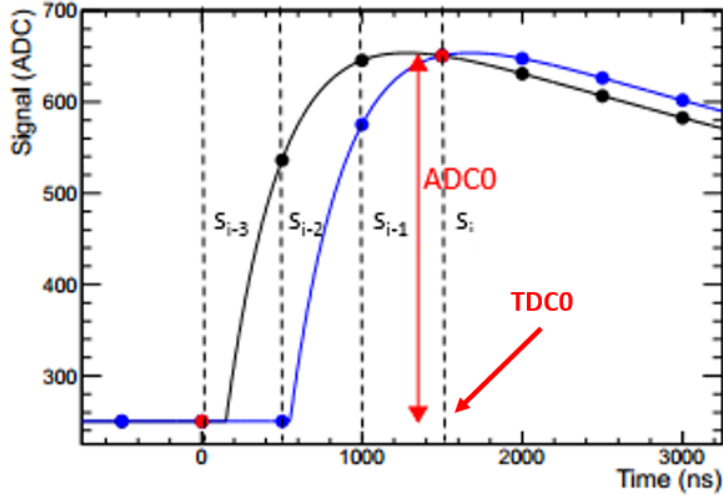


Figure 3.11: Duel Correlated Sampling. In single-point timing, these two pulse shapes would have the same ADC and TDC values. In multi-point timing, the pulse shape is fit using the four samples to determine the maximum pulse height.

in the system. A sync command is then sent down the chain. Upon receiving the command, each component begins a delayed buffer corresponding to its learned delay value from the end of the chain. When the sync signal reaches the end of the chain, all detector components exit the buffer simultaneously and begin counting from the new NOvA time.

With the detector components synced, the time register in each FEB is used to time stamp the 32 channels it monitors. The sampling rate of the APD is multiplexed down to 2 MHz (8 MHz) at the FD (ND). The FEB shapes the pulse signal from the fiber in each channel with a 460 ns rise time and 7000 ns fall time. Since the ND samples four times as often, the rise time is 140 ns and fall time is 4500 ns. The readout from each channel is triggered by a Duel Correlated Sampling (DCS) algorithm. DCS compares the ADC value at the current sample, s_i , to three samples prior, s_{i-3} . If the difference in the quantities is above a threshold set for each channel, the FEB reads out the current sample with the current time as the TDC and the difference in ADC as the charge. This method is known as single-point timing and was used in the FD until August 2014.

Single-point timing has generally low resolution and can often assign two very different pulse shapes the same ADC and TDC values, as shown in Fig. 3.11. The ND and the FD after August

2014 use multi-point timing. The readout is triggered in the same way, but instead of using the TDC and ADC of s_i , the 4 samples in the pulse are fit to precisely determine the maximum height and location of the pulse. Multi-point timing shows better timing resolution which is critical for separating interactions in the ND [51].

3.5 Neutrino Simulation

Simulations help to quantify our understanding of the underlying physics in particle detectors. They will be used extensively throughout the next chapters of this dissertation to construct reconstruction tools and evaluate algorithm performance. NOvA uses a number of different software packages to simulate the NuMI beamline, neutrino interactions, and the propagation of particles through the various media that make up the detector and beamline. All of these tools rely on probabilistic Monte Carlo (MC) simulation methods to approximate particle interactions.

The NOvA simulation begins with protons from the accelerator. The FLUGG [52] and FLUKA [53] packages are used to simulate the production of hadrons within the target. The propagation of the hadrons down the beamline is carried out by a GEANT4 implementation of the NuMI beamline [54] including the target, horns, and decay pipe. The result is a file containing the beam composition at neutrino production: neutrino flavors, energy, and production location. Separate flux files are made for both the forward and reverse horn current modes as well as at the ND and FD. The hadron production is weighted using the Package to Predict the Flux (PPFX) [55] to alter the resulting neutrino flux based on measurements of hadron production in hadron-nucleus collisions from external experiments. PPFX also constrains the uncertainties associated with hadron production and the beamline.

For the FD, no oscillation assumptions are made in the simulation. Instead, several different flux files are generated based on the original:

- Nonswap - The original NuMI beam flux at the FD location.

- Fluxswap - The NuMI beam flux where all ν_μ (ν_e) in the beam have been replaced by ν_e (ν_μ). This is equivalent to a 100% oscillation probability for $\nu_\mu \rightarrow \nu_e$.
- Tau - All neutrinos in the beam are replaced by ν_τ . These files are used to explore the high energy backgrounds from ν_τ and NC interactions.

Together these fluxes allow for the construction of any flux that might result from oscillations from any set of oscillation parameters.

The GENIE package [56] is used to determine if a given neutrino interacts and the process through which it interacts using models of neutrino interaction cross sections. GENIE outputs 4-vectors for each secondary particle produced in the interaction. The default GENIE models are found to be discrepant with NOvA data and are corrected based on NOvA's measurements of neutrino cross sections, external measurements, and observed discrepancies with the data [57].

Once again, GEANT4 is used to handle the propagation of these particles through the detector geometry and also through the detector enclosures or surrounding rock if needed. Finally, NOvA internal packages, PhotonTransport and ReadoutSim [58], model the propagation of light through the cells and wavelength shifting fiber into the APD and output a digitized waveform. The result is a file with the same format as data generated from the DAQ, but with additional truth level information that can be utilized in the development of analysis tools and the application of physics models.

In the FD, real cosmic-ray data events from the cosmic trigger are overlaid with the neutrino interaction to add real detector backgrounds and noise. In the ND, neutrinos interacting in the rock around the detector are simulated independently and then overlaid with the neutrino event.

3.6 Detector Calibration

Before the detector data can be analyzed, the ADC recorded in each cell must be individually calibrated to an absolute energy scale. The conversion of ADC to an energy must correct for




Figure 3.12: Passage of Particles through Matter. The energy losses of a muon as it traverses a copper medium. At NOvA’s energies, muons are primarily minimally ionizing until they stop with energy depositions well described by the Bethe-Bloch equation. Figure recreated from [48].

several effects: the light yield of the scintillator, attenuation of light in the wavelength shifting fiber, and the quantum efficiency of the APDs. The calibration of each cell is done in two parts. First, the ADC is converted to an estimate for the number of PE seen in the fiber accounting for the attenuation along the length of the cell [59]. Second, the corrected PE is scaled to physical units of energy [60].

Both steps take advantage of the well understood energy depositions of muons, modelled by the Bethe-Bloch equation [48], see Fig. 3.12. Through-going muons provide a long region of uniform energy depositions which can be used for the attenuation correction, and stopping muons provide a region well modelled by the Bethe-Bloch equation to measure an absolute energy. These muons are provided in great abundance by cosmic rays. The FD samples the 150 kHz of cosmic rays traversing the detector using the cosmic trigger and the ND samples about 5 kHz picked up by the DDActivity trigger.

Cell hits in cosmic-ray events with activity in both adjacent cells in the same plane are selected.




Figure 3.13: Cell Attenuation Correction. An attenuation curve for a cell in the ND (left) or FD (right). The readout is at the positive end of the x-axis. The data is fit to a double exponential due to the bi-directional nature of the fibers. The blue curve shows the corrected PE/cm.

These combinations of three consecutive hits are known as tricells. Tricells in through-going cosmic rays allow for precise computation of the muon path length within the cell and, thus, the dE/dx within the cell. These cells are corrected in PE/cm vs distance from the readout, as shown in Fig. 3.13. The fit function used in Fig. 3.13 is given by

$$y = C + A \left(e^{x/X} + e^{-(3L/2+x)/X} \right) \quad (3.6)$$

where L is the length of the cell, x is the distance measured from the center of the cell, and X is the attenuation length. Light propagates in both directions in the fibers, so the second exponential accounts for the long path to the readout. The behavior of the fit function at the edges accounts for light reflection off the ends of the cells and is corrected for empirically. This fit is made for each cell to correct for the attenuation.

The absolute calibration of corrected PE to GeV is also carried out using tricells from cosmic-ray muons. Stopping muons are selected in the detector and the dE/dx behavior at the end of the track is compared to the Bethe-Bloch expectation. The end of track behavior is shown in Fig. 3.14. This measurement is done in both data and simulation where the corrected PE/cm has




Figure 3.14: Absolute Energy Calibration. The charge deposition behavior at the end of cosmic-ray muon tracks is used to set an absolute energy scale for each cell. The black points show the mean energy lost per centimeter as a function of distance from the end of the track. This curve should be well described by the Bethe-Bloch equation.

been matched to the data. The truth information from the simulation is then used to set a scale factor from corrected PE/cm to GeV. This correction is done for all cells in the detector to yield a fully calibrated energy for each hit.

Chapter 4

Event Reconstruction

Reconstruction algorithms take the calibrated hit energies and times from each detector cell and use them to estimate physics quantities to be used in an analysis such as particle identities and momenta. On NOvA, many of these tasks are solved by employing tools from machine learning, which is the focus of Chap. 5. This chapter will focus on the traditional reconstruction methods which utilize the space-time relationship between detector hits and knowledge of the physics processes common to particle interactions and propagation.

4.1 Neutrino Interaction Topologies

Trained physicists are talented at reconstructing individual events by hand [61]. Interactions recorded by the first particle detectors would have been reconstructed entirely using such strategies. However, today's detectors produce million of events, far too much data to be reasonably analyzed by humans. Computer algorithms developed for reconstruction tasks should be able to achieve comparable or better performance than humans.

While developing reconstruction tools, it is important to consider the types and topologies of events that need reconstructing. Figure 4.1 shows examples of candidate neutrino interactions in the NOvA detectors which have already been isolated from other activity in the event.

The top-left event shown is a ν_e CC interaction. As mentioned in Chap. 2, neutrino interactions are lepton flavor conserving and produce the corresponding charged lepton, in this case an electron. Above a critical energy, about 100 MeV for NOvA, electrons lose energy primarily by Bremsstrahlung radiation, the emission of energy as photons when a charged particle is deflected by another charged particle. These photons then pair produce into an e^+e^- pair which can also emit Bremsstrahlung photons. The result is an electromagnetic shower with length and width

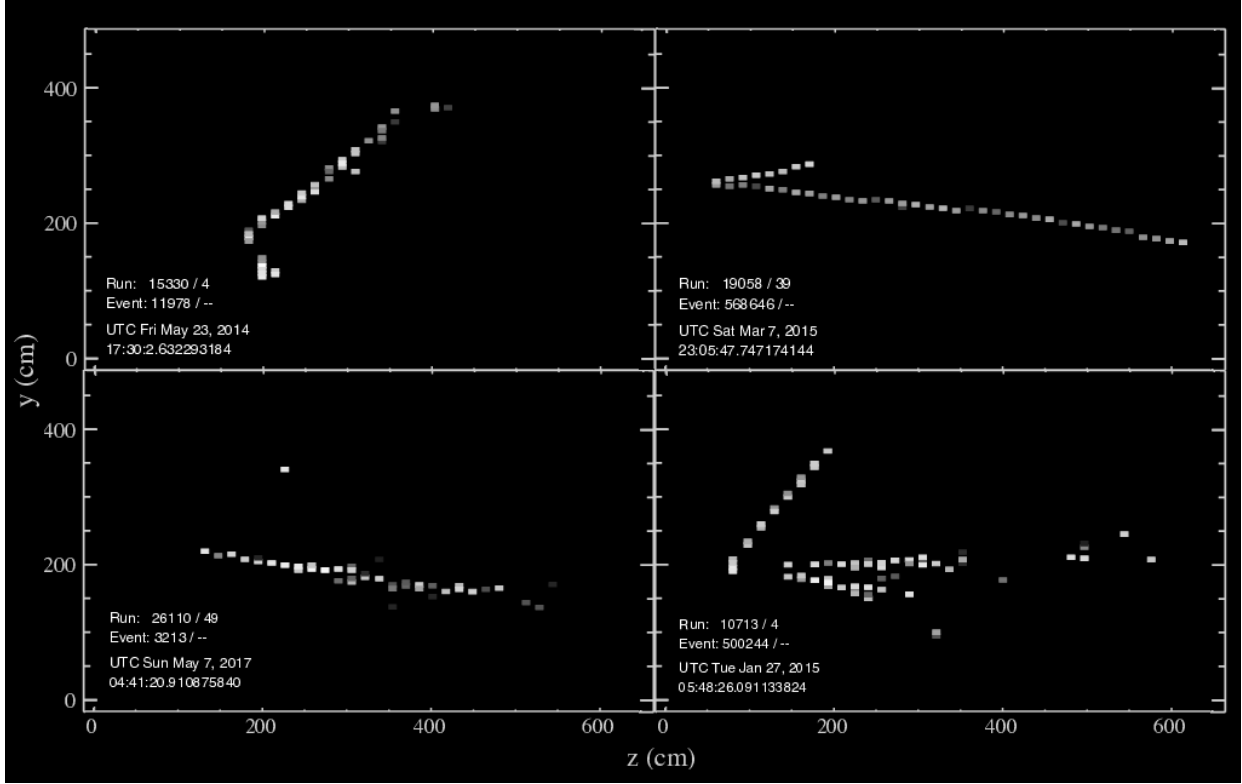


Figure 4.1: NOvA Event Topologies. One detector view of candidate data neutrino interactions in the NOvA detectors. The grayscale corresponds to the ADC in each cell. Top Left: An electron neutrino interaction producing an electron and proton. Top Right: A muon neutrino interaction producing a muon and proton. Bottom Left: An electron antineutrino interaction producing a positron and neutron. Bottom Right: A neutral current interaction. Reconstruction algorithms need to be capable of handling a variety of different event topologies.

characterized by the radiation length and Moliere radius of the medium, see Sec. 3.2. It is worth noting that at some points the energy of the shower could be entirely in photons, creating gaps in the shower, a feature relevant for reconstruction tools. This interaction also produced hadronic activity, in this case a single proton which appears as a heavily ionizing track.

The top-right event is a ν_μ CC interaction. The muon mass is greater than the electron and, thus, the primary energy loss is from ionization. Muons appear as long straight tracks with approximately equal energy deposited along its length. However, muons have unique behavior at the end of their track characterized in the Bethe-Bloch equation, which is why they are well suited for calibration purposes, before decaying into a Michel electron.

The bottom-left event is a $\bar{\nu}_e$ CC interaction. In NOvA, antineutrino interactions look very similar to their neutrino counterparts. The characteristic electromagnetic shower is still present, however this time induced by a positron rather than an electron (by charge conservation). Note that NOvA has no mechanism to directly identify the charge of an individual particle, such as the magnetic fields employed by other experiments [62]. However, antineutrino interactions do have distinct topological differences from their neutrino counterparts. Neutrinos frequently interact with a neutron to produce a proton and antineutrinos frequently interact with a proton to produce a neutron. Unlike protons, neutrons are neutrally charged and so are not visible in the detector scintillator until they interact with another nucleus. Neutrons will either collide with a nucleus and knock out a proton or be captured and re-emit photons at a fixed energy. Regardless of the process, neutrons usually appear as few hit clusters separated by on average one interaction length from the interaction vertex.

The bottom-right event is a NC interaction. In NC interactions, the departing lepton is another neutrino and there is no handle on what its flavor was. Instead, they leave only some hadronic activity in the form of mesons (mostly pions), protons, and neutrons. Charged pions appear as heavily ionizing tracks, though often longer than that of protons. High energy charged pions can often imitate shorter muon tracks. Neutral pions have a lifetime of 8.4×10^{-17} s and decay into a pair of photons with branching ratio 0.99. Photons will travel, on average, $\frac{9}{7}X_0$ before pair producing and initiating an electromagnetic shower. This photon electromagnetic shower imitates that of an electron, particularly when the two photons are overlapping, except for the separation of the shower from the interaction vertex which serves as the main mechanism for identifying them. Thus, NC interactions are the primary background of interest when identifying ν_μ CC or ν_e CC events which will be demonstrated several times through the following chapters.

A ν_τ CC interaction is not shown. The peak energy of NOvA's beam is less than 2 GeV which is well below the kinematic energy threshold for tau production at 3.5 GeV. However, when a ν_τ

does interact via a CC interaction, the resulting τ decays after traveling a very short distance into either an electron (and two neutrinos) with a branching ratio of 0.1782, a muon with branching ratio 0.1739, or a number of charged pions and neutral pions. In this way, ν_τ CC interactions mimic the topologies of the above interactions.

Note that the events in Fig. 4.1 are primarily quasi-elastic interactions. At a peak energy centered around 2 GeV, NOvA also detects resonance and deep inelastic scattering interactions as shown in Chap. 2. These must also be considered when developing reconstruction tools.

Throughout the remainder of this dissertation, ν_e , ν_μ , and ν_τ will refer to their corresponding CC interactions and NC will refer to NC interactions of any flavor unless otherwise specified.

4.2 Isolating Neutrino Interactions

The NOvA DAQ reads out all activity within 550 μs events for NuMI triggers or cosmic triggers. In the ND, multiple neutrino interactions occur within the detector in each 10 μs beam spill. In the FD, neutrino interactions from the beam need to be separated from the surrounding activity, such as cosmic rays or detector noise. Therefore, the first step in reconstructing neutrino interactions is to separate them from each other and the rest of the activity within the detector. Figure 4.2 shows an example detector readout from the ND with several neutrino interactions. The figure shows how the NOvA timing system already gives reasonable separation of the interactions.

Previous NOvA analyses utilized a clustering algorithm based on DBSCAN [63] called Slicer4D to group together hits belonging to independent interactions into clusters known as slices, with each slice containing all hits from a single neutrino interaction or cosmic ray [64]. This algorithm has been successful, but has a few common failure modes. The most notable is that it is possible for two clearly independent slices to be merged together when there are “bridge” hits between the two. This effect is particularly problematic in the ND where multiple neutrino interactions occur in a short time window causing pile-up in the number of slices at high beam intensity. NOvA now uses




Figure 4.2: Timing Event Display. A detector readout from the ND depicting several neutrino interactions in a single beam spill. The hits are colored by time. Timing alone already gives reasonable separation of the different neutrino interactions.

a Time Density Slicer (TDSlicer) [65] which attempts to solve this problem by identifying centroid hits that make the base for individual clusters.

The centroid hits are identified by finding local maxima in the density of hits [66]. The euclidean distance between a pair of points in the same detector view is defined by

$$d_{ij} = ||\Delta t| - R/c| \quad (4.1)$$

where Δt is the time difference between two points, R is the spatial distance between points, and c is the speed of light. A local density score and isolation score are assigned to each hit based on d_{ij} for each detector view. The density score is computed as

$$\rho_i = \sum_j e^{-\frac{d_{ij}^2}{\tau^2}} \quad (4.2)$$

where τ is a scale parameter related to the timing resolution for each hit, set to 16 ns in the ND and 80 ns in the FD. The isolation score is defined as

$$\delta_i = \min_{j; \rho_j > \rho_i} \frac{d_{ij}}{\tau} \quad (4.3)$$

which is the scaled euclidean distance between the point i and closest point with a greater density score. The hit with the greatest density score is assigned a score equal to the greatest possible separation between two hits in a single beam spill. Centroids are then identified as hits with $\rho > 3(10)$ and $\delta > 8(6)$ in the ND (FD).

Starting from the centroids, 3D clusters in each view, xzt and yzt , are constructed by finding a minimally connected spanning tree using Prim's algorithm [67]. Hits outside any cluster are added to the nearest cluster based on the distance to any hit within the cluster. There is a minimum distance threshold of $8(5)\tau$ in the ND (FD) above which hits will not be clustered.

Finally, the clusters are merged across the two views into the final set of 4D slices. The average zt values are measured for all possible pairs of clusters from different views. The best pair of clusters are merged iteratively until all paired clusters have been joined together. The remaining unmatched clusters are considered noise.

The performance of the slices is tuned based on two metrics, the completeness and purity:

$$\text{completeness} = \frac{\text{Energy from True Particle in Slice}}{\text{Total Energy from True Particle}}, \quad (4.4)$$

$$\text{purity} = \frac{\text{Energy from True Particle in Slice}}{\text{Total Energy in Slice}}. \quad (4.5)$$

The completeness measures how much energy from a particle made it into the slice. The completeness is sometimes also called the efficiency of the cluster. The purity measures how much energy from other particles made it into the slice. TDSlicer achieves an average completeness of 0.9705(0.8489) and purity of 0.9869(0.9939) on FD(ND) neutrino interactions. As mentioned at

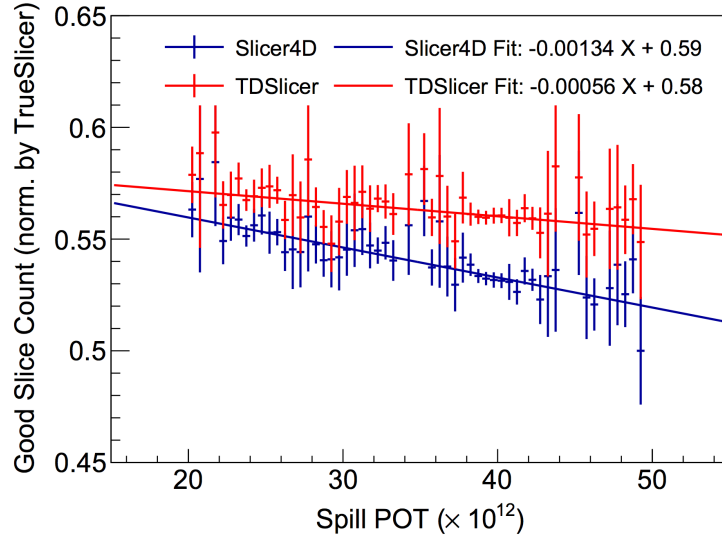


Figure 4.3: Near Detector Pileup. The fraction of good slices produced by TDSlicer and Slicer4D normalized by the number of true slices as the beam intensity increases. The solid lines show a linear fit demonstrating the pileup effects in both algorithms at high beam intensities. TDSlicer shows more good slices for all beam intensities with improved performance at higher intensities.

the start of this section, of particular interest is how the performance of the algorithm in the ND scales with the intensity of the neutrino beam to reduce pileup effects. The fraction of “good” slices, those with completeness and purity better than 0.9, versus the beam intensity is shown in Fig. 4.3.

Throughout the remainder of this dissertation, event and slice will be used interchangeably to refer to clusters containing a single physics interaction.

4.3 Resolving the Interaction Point

The next stage of the reconstruction is to identify the point where the neutrino interacted within the detector, the interaction vertex. All energy deposited in an event is assumed to originate from a single point, which is mostly true for neutrino interactions. This ignores any secondary interactions such as hard scatters or decays of particles after the neutrino interaction.

NOvA uses a vertexing algorithm known as elastic arms which assumes all particle tracks, the

arms, emanate from the vertex [68]. To generate an initial set of arms, a hough transform [69] is performed on all hits in a slice in each view. The hough transformation is designed to identify the major linear features in an event. Every pair of hits in a view is fit with a line and parameterized in a space (ρ, ϕ) such that

$$\rho = x \cos \phi + y \sin \phi \quad (4.6)$$

where (x, y) are the line points in cartesian space, ρ is the nearest distance of the line to the origin, and ϕ is the angle between ρ and the positive x-axis. Each pair is weighted by a gaussian smear such that a smaller weight is given to hit pairs that are further apart or at a much different angle from the neutrino momentum which reduces the effect from hit pairs on different particle tracks. To reduce computations, hit pairs further than some configurable distance apart are not considered. In addition, to prevent overproducing horizontal lines, hits with the same x or y coordinate are only paired if they are a configurable distance apart.

In this new hough space, collinear hit pairs appear as peaks. The predominant linear features of the event are identified as bins in (ρ, ϕ) space with more hit pairs than the average across the entire hough map. To identify the exact values of (ρ, ϕ) , the 7×7 set of bins centered on the peak bin are averaged together weighted by their content and separation from the peak bin. An example event with the final hough lines is shown in Fig. 4.4.

The hough lines are used to construct 3D lines from similar peaks then used as the seed arms to the minimization of an energy function

$$E = \sum_{i=1}^N \sum_{a=1}^M V_{ia} M_{ia} + \lambda \sum_{i=1}^N \left(\sum_{a=1}^M V_{ia} - 1 \right)^2 + \frac{2}{\lambda_\nu} \sum_{a=1}^M D_a \quad (4.7)$$

where M_{ia} is the distance between the i th hit and the a th arm, V_{ia} is the the strength of the association, D_a is the distance between the vertex and the first hit along the arm, and λ and λ_ν control the strength of each term. The first term measures the goodness of fit between the hits and




Figure 4.4: The Hough Transform and Elastic Arms Vertex. A candidate ν_μ event with the predominant hough lines drawn. The yellow cross indicates the fitted elastic arms vertex.

the arms. The second and third term are penalties for hits that are not associated with any arm and arms where the first hit is a large distance from the vertex, tuned to the conversion distance for photons. V_{ia} is given by

$$V_{ia} = \frac{e^{-\beta M_{ia}}}{e^{-\beta\lambda} + \sum_{b=1}^M e^{-\beta M_{ib}}} \quad (4.8)$$

where $e^{-\beta M_{ia}}$ measures the likelihood that a hit belongs to the arm and $e^{-\beta\lambda}$ is the likelihood that a hit is noise. The equation is normalized such that a hit will have a likelihood of belonging to any arm and the remaining likelihood is noise. The parameter β can be interpreted as a temperature $\beta = 1/T$ where at higher temperatures arms will have more associated hits.

The vertex is chosen to minimize the total energy in Eq. 4.7 by simulated annealing [70]. The performance of the vertexing is shown in Fig. 4.5 and an example event with the vertex drawn is in Fig. 4.4. The best performance is seen on QE events where the simple, two particle topology has an interaction point that is easier to identify. The reconstructed vertex is within 20 cm of the true




Figure 4.5: Elastic Arms Vertex Resolution. The difference between the reconstructed vertex by elastic arms and the true neutrino interaction vertex. The events are separated by QE (red) and non-QE (blue) interactions and all interactions (green). The left shows ν_μ interactions and the right shows ν_e interactions.

vertex for 90% of events.

4.4 Clustering Particles

With the interaction vertex identified, the next step is to group hits within the event into directed candidate particle clusters known as prongs. Prongs make the basis for future analysis including particle identification and energy estimation.

The event vertex is used to cast the problem of clustering as one dimensional using a possibilistic fuzzy k-means algorithm [71, 72]. Here, possibilistic means that each hit can be unclustered with some probability for being noise and fuzzy means a hit can belong to multiple clusters. The k refers to the arbitrary number of clusters.

Prongs are first constructed for each detector view then matched across views. Each hit is identified by an angle from the positive z-axis. In this angular space, particle tracks appear as peaks of deposited energy at some common angle.

The algorithm proceeds iteratively by adding new prong centers and updating the membership

of each hit to each prong, μ_{ij} , given by

$$\mu_{ij} = \exp\left(-\frac{m\sqrt{a}d_{ij}}{\beta}\right) \quad (4.9)$$

where m is the allowed total membership of hits to clusters, a is the number of clusters, β is a normalization describing the allowed width of clusters, and d_{ij} is the angular separation between the hit and cluster. The distance d_{ij} is normalized by an uncertainty determined by the multiple scattering of muons [51] to allow for particle scattering further from the vertex. After each iteration, the cluster centers are updated until an allowed tolerance of 1×10^{-7} radians is found for the update.

A number of cleanup steps are needed to produce the final set of prongs. It is possible for the algorithm to identify the same location twice for a candidate prong. In this case, the duplicate is removed and the iterations are repeated with a new prong seed until all hits have membership to a prong or 7 prongs are produced. Prongs within 30° of one another are merged if 66% of the hits in the smaller prong have membership to the larger prong. This merging is intended to join together the fringes of electromagnetic showers. Finally, the algorithm is invariant under displacements directly away from the vertex. In the case of photons, for example from π^0 decays, few hits near the vertex would be merged with the photon hits despite the large gap between the two. If the gap begins within two planes of the vertex and is larger than 25 cm, the two regions are separated. Further from the vertex, the distance is increased to 75 cm for separation.

The final set of 2D prongs in each view need to be matched into full 3D prongs. A matrix of possible prong pairs is constructed. Prong pairs are only considered if they overlap endpoints by at least one plane. The pair of 2D prongs are temporarily matched into a candidate 3D prong so that information across both views can be used in the attenuation correction for calibrated energy. With the calibrated energy, the integrated energy along the 3D path length of the prong is identified for each view. A measure of similarity K between the two energy distributions is computed using a

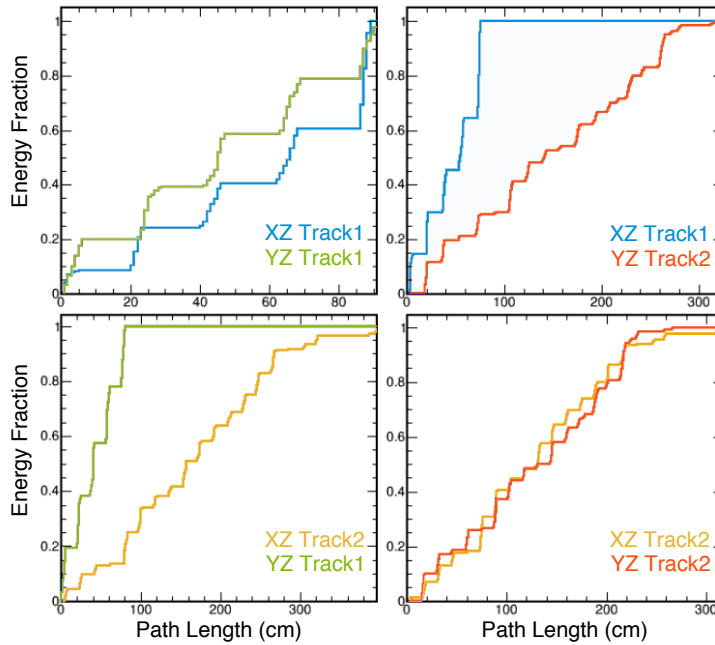



Figure 4.6: Prong View Matching. Top: An event display with the 2D prongs produced by the k-means algorithm. Bottom: The energy distributions for each possible candidate 3D prong. The two distributions for each candidate are compared using a kuiper test. In this case, XZ track 2 and YZ track 2 have the closest match and will be merged, then XZ track 1 and YZ track 1 have the next best match (and are the only remaining prongs) and are merged.

Kuiper Test [73] with

$$K = D^+ + D^- \quad (4.10)$$

$$D^+ = \max(E^{XZ}(s) - E^{YZ}(S)) \quad (4.11)$$

$$D^- = \max(E^{YZ}(s) - E^{XZ}(S)) \quad (4.12)$$

where $E^{XZ(YZ)}(s)$ is the fraction of energy a path length s along the prong. In this way, $D^{+(-)}$ is the maximum deviation between the two distributions from above (below). Because of the planar geometry of the detector, there are natural jumps in the two distributions. The path length in one view is allowed to shift by ± 12 cm in steps of 0.5 cm and the lowest score is kept.

With the scores for all possible pairs computed, the pair with the lowest score (the closest matching) are combined and the two prongs are removed from the list of remaining 2D prongs. This is repeated until all prongs have been matched. There is no minimum similarity score for a match to be made. However, there will be one view where remaining 2D prongs are unmatched. This most commonly happens when two particles are overlapping in one view producing one prong, but distinct in the other producing two prongs. The single prong from the view will be matched to one of the two from the other view. The spare prong will either be incorrectly matched to another prong in the event or be left unmatched. These prongs remain 2D and are carried through the analysis in this form. An example of the matching process for a single event is shown in Fig. 4.6.

The performance of the algorithm is evaluated on contained ν_e interactions in the FD shown in Fig. 4.7. As with the slicing algorithm, the performance metrics of interest are the completeness and purity, defined in Eq. 4.4, of the prongs. For true lepton energies above 0.5 GeV the reconstructed prong achieves a completeness above 90% and purity of 80%. The performance of this algorithm will be examined further when compared to an alternative in Chap. 5.




Figure 4.7: K-Means Prong Performance. Completeness (left) and purity (right) of prongs in contained ν_e events between between 0.5 and 3.5 GeV in the FD versus the visible energy in the electron. QE events are shown in red, non-QE events in blue, and all events in green.

4.5 Particle Tracking

Particle tracking is the process of determining the precise trajectory of particles traversing the detector. The goal of a tracking algorithm is to reconstruct the true position of a particle given the positions of hits in the detector. This process relies on the understanding of how particles traverse the medium.

NOvA currently employs two parallel tracking algorithms: one based on kalman filters and one based on break point fitting. Both algorithms rely on identifying points where a particle undergoes small angle multiple scattering. The scatter of a particle can be approximated, when the angle is small, by a gaussian with width given by

$$\Delta\theta = \frac{13.6 \text{ MeV}}{\beta p} \sqrt{\frac{x}{X_0}} \left[1 + 0.038 \ln \left(\frac{x}{X_0} \right) \right] \quad (4.13)$$

where β is the particle velocity in units of c , p is the particle momentum, and x is the length along the trajectory. For a 2 GeV muon in the NOvA detector, $\Delta\theta$ is 5.4×10^{-3} rad.




Figure 4.8: Reconstructed Track Direction. The cosine of the angle between the reconstructed initial track direction using Kalman filter (left) or BPF (right) and the true track direction for muons in true ν_μ events. The X-axis starts from $\cos \theta = 0.995$ or $\theta = 5.7^\circ$.

4.5.1 Kalman Filter Tracking

The kalman tracker identifies a set of tracks in each 2D detector view of the neutrino slice [74]. The algorithm starts from track seeds which are hit pairs separated by less than 4 cells in the slice. Starting from the seeds, track clusters are constructed by iteratively adding nearby hits to the cluster. The algorithm estimates possible track positions based on the current track position and slope from current hits. A score is assigned to each nearby hit based on the likelihood of being part of the current track from Eq. 4.13. Hits with high scores are added to the track and the position and slope are updated before proceeding to the next hit.

Kalman filters are used to fit the track clusters to true particle positions [74, 75]. In general, Kalman filters solve the problem of estimating the true value of a state at discrete steps based on the prior state of the system.

Once the tracks are constructed and fit, the two dimensional tracks need to be merged across views. A score S is assigned to each possible track pair by

$$S = \frac{\Delta z_{\text{start}} - \Delta z_{\text{stop}}}{\text{overlap}} \quad (4.14)$$




Figure 4.9: Break Point Fitter Schematic. The measurements are the locations of detector hits in 3D prongs. Scattering planes are identified where the particle track has deviated sufficiently from the extrapolated track at the previous plane. Each plane adds a scatter of an angle given by Eq. 4.15.

where $\Delta z_{\text{start}(\text{stop})}$ is the difference in the z coordinate of the beginning (end) of the track and the overlap is the range of z coordinates common to both tracks. Any tracks with no overlap will not be matched. Track pairs with the lowest scores are merged and then removed from the pool of available tracks. This process repeats iteratively until no more matches can be made.

The angle between the reconstructed initial track direction and the true track direction for muons is shown in Fig. 4.8.

4.5.2 Break Point Fitter

Instead of creating new track clusters, the break point fitter (BPF) algorithm constructs particle tracks from the already constructed 3D prongs [76]. The procedure is depicted in Fig. 4.9. Starting from the end of the track, the deposited energy through each cell and the length of materials (PVC or oil) traversed are measured as a function of the distance d . A new scattering plane is added when one of two conditions are met:

1. d exceeds $2X_0$ or




Figure 4.10: Break Point Fitter Track. Left: The end of a muon track in the NOvA detector. The gray boxes indicate the cells within the detector. The green dashed line shows the reconstructed track under the muon hypothesis. Right: The locations of scattering planes under the muon hypothesis (blue) and the proton hypothesis (red). Both hypotheses show more scattering planes near the end of the track where the lower energy particles are more likely to scatter.

2. the transverse scattering of the particle $d\Delta\theta$ exceeds 3 mm.

The first condition ensures that there is always at least one scattering plane along a reasonable length of track. This condition means at least one scattering for every seven detector planes. The second condition ensures that the linear extrapolation between two scattering planes is accurate to within the minimum resolvable distance in the NOvA detector, the width of the plastic cell walls. If one of the conditions are met, a scattering plane is added half a step upstream.

With the scattering planes found, the transverse location of the particle, measured along the z direction, is given by

$$\xi_i = a + bz_i + \sum_{J=1}^M \alpha_J (z_i - Z_J) \Theta(z_i - Z_J) \quad (4.15)$$

where ξ_i is the location at the i^{th} measurement, a and b are the intercept and slope of the initial track, α_J is the scattering angle (assumed to be small) at the J^{th} scattering plane, and the heaviside step function Θ is used to account only for previous scattering angles upstream of the current measurement. The measurements are the locations of hits in the detector. A set of a , b , and α_J uniquely identify the particle trajectory. With the trajectory identified, the energy and momentum can be measured by integrating dE/dx along the particle path through each of the different detector materials.

For every prong, three different tracks are reconstructed based on different particle hypotheses: muon, proton, or pion. Since the particles differ in mass and energy depositions, each hypothesis will result in a unique track with a unique energy and momentum measurement. The end of a muon track is shown in Fig 4.10.

The angle between the reconstructed initial track direction and the true track direction for muons is shown in Fig. 4.8. BPF shows a 17% improvement in the average angular reconstruction compared to the Kalman filter.

Chapter 5

Machine Learning Techniques

Machine Learning (ML) refers to any computer algorithm which improves with experience. Training examples are fed through the algorithm and then a set of weights or training parameters are modified to improve the overall performance. ML techniques are viable candidates for solving reconstruction problems in particle physics experiments where data, either real detector data or simulated MC, are plentiful. These techniques have proven successful in applications for particle classification, estimation of physics quantities, reconstruction, and more [77].

This chapter will describe the ML algorithms utilized by the NOvA experiment. However, it will not cover how ML algorithms learn; please see [78] for a mathematical description of loss functions and back propagation, essential components of the learning process.

5.1 Convolutional Neural Networks

The most common forms of ML are simple networks such as Boosted Decision Trees (BDT) or Artificial Neural Networks (ANN) which have contributed to particle physics discoveries for the last 20 years [77]. However, in recent years interest has turned to Convolutional Neural Networks (CNN) which offer a number of advantages over more traditional methods of ML.

One of the most beneficial aspects of CNNs is they can often be developed as end-to-end algorithms. The previous chapter described the reconstruction chain used by NOvA, where each tool used as input the output of the previous stage. This causes stacking inefficiencies where misreconstruction in one stage effects the performance of all subsequent algorithms that depend on it. An example is shown in Fig. 5.1 where the vertex was placed incorrectly and, as a result, the prongs were misreconstructed. Using physical quantities measured by these tools as input to a BDT or ANN will result in similar inefficiencies. CNNs can often be designed to use the simplest

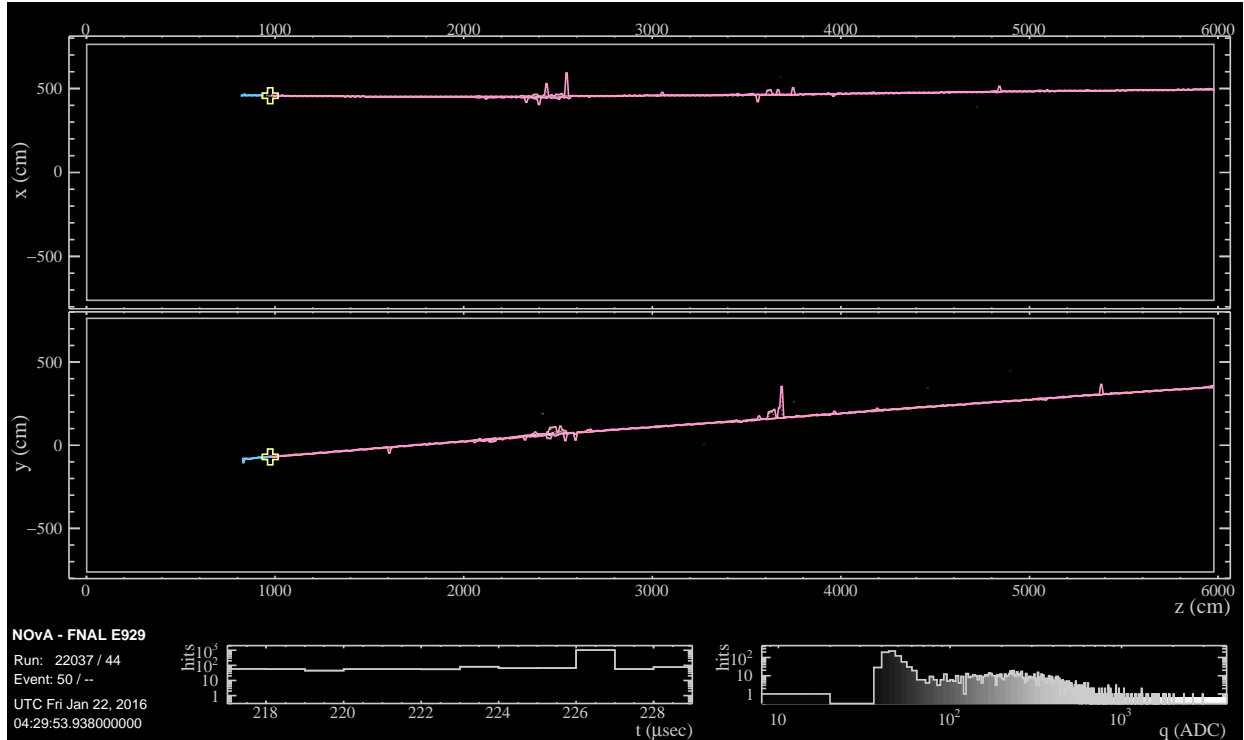


Figure 5.1: Elastic Arms Failure. An example ν_μ QE event where the elastic arms vertex, shown by the yellow cross, is pulled downstream along the muon track. This results in the muon being split into two separate prongs, the pink and blue clusters, on opposite sides of the vertex.

possible input, in the case of NOvA this is usually the slice, which isolates the reconstruction task from the performance of other algorithms.

CNNs operate by convolving an image with a kernel, a matrix which acts as a weighted sum of pixel values in the image. Traditional kernels are used in image processing to sharpen, emboss, detect edges, or any number of other applications. In CNNs, the kernel sizes are determined by the user, but the values in the kernels are learned such that the network can extract the optimal set of features to use as discriminators in the task being solved. This is much different from BDTs or ANNs where the user determines what set of inputs to use.

A set of kernels are grouped into a single convolution layer. These layers make the main component of a CNN. The output of a convolution layer is often called a feature map where each channel is generated from a single kernel. The other main component are pooling layers which




Figure 5.2: Overview of a CNN. The input to a CNN is an image or some set of data that has been recast as a two dimensional image. Convolution layers convolve a number of kernels with the image to extract a set of discriminating features. Each kernel produces one channel in the output. Pooling layers group together a set of pixels to reduce the overall size of the image. By applying successive convolution and pooling layers, features of different sizes can be extracted by the network. The final pooling layer reduces the image to a one dimensional set of values referred to as the feature vector. The final component is a fully connected layer with the final output in purple.

group together a set of pixels based on a given function, usually the maximum or average of the pixel values. Pooling is useful to reduce the computational complexity of the network. A series of convolution and pooling layers make up the core components of a CNN that extracts features from an image. After each convolutional layer is a non-linear¹ activation function, which aids in the learning of complex patterns [79]. The schematic of a general CNN is shown in Fig. 5.2. As a CNN learns, the weights in all kernels and the fully connected layers are modified together to optimize algorithm performance.

When developing CNNs for use in a particle physics experiment, careful consideration must be given in constructing the input. CNNs rely on the topological relationship between nearby pixels in an image, however not all detectors read out information directly in this way. Detector readout must be cast to an image-like structure in a way that preserves the event topology [80]. On NOvA, the detector is already arranged in a grid of cells for each view. Thus, the calibrated energy in each cell will be realized as a single pixel value so that events can be analyzed in the same style

¹If a network consists only of linear functions, there would be no reason to have multiple layers since the composition of two linear functions is itself another linear function. There would be no higher level learning by the network.

as in image recognition by CNNs. Note that analyzing events in this way produces two image-like matrices, one for each detector view.

5.2 Event Classification

The most important component of the analysis described in this dissertation is the identification of the neutrino flavor, which is one of the two quantities needed to measure the oscillation probability (the other being the energy). A siamese CNN called the Convolutional Visual Network (CVN) is used to classify NOvA events as ν_e , ν_μ , NC, or cosmic ray. NOvA has used CVN in all oscillation analyses since 2016 [81] at which time it showed a 30% increase in effective exposure compared to traditional identification methods [82, 51].

NOvA events are broadcast to two dimensional arrays, called pixel maps, where each cell in the detector corresponds to one element of the array. The pixel maps are 200 planes long (100 planes in each view) and 80 cells tall. The first plane of the pixel map is the upstream most plane with a hit and such that the following 8 plane window has more than 4 hits. This arrangement prevents the pixel map from starting too early due to outlier hits upstream from the core of the event. The center cell is the median cell of all hits in the event determined independently for each view. One might consider using the reconstructed vertex as the starting point, but this would add unnecessary dependencies into the network.

The values in the array are determined by the calibrated energy for each hit. The energy is scaled so the values can be cast as 8-bit unsigned chars with values between 0 and 255, much like a gray scale image. Hits with energy above 278 MeV are saturated at the maximum possible value of 255. This representation offers a factor of eight savings in the size of the pixel map compared to a floating point representation [83] without significant loss in the precision. The result of this process is a pair of 100×80 arrays with values corresponding to the hit energies for each view of the event. An example pixel map and a comparison between the two representations is shown in




Figure 5.3: NOvA Pixel Map. The X-view (left) and Y-view (center) pixel map for a ν_μ event. The color corresponds to the energy of a hit using 8-bit precision. On the right is a comparison between 8-bit precision (red) and floating point precision (blue). Both representations are saturated at 278 MeV, but very few hits have energies above this.

Fig. 5.3.

The network was optimized for both performance and evaluation time. A number of network architectures were considered, but the final choice was a modified version of mobilenet v2 [84]. Mobilenets are a class of CNNs which use separable convolutions to create large, but efficient, CNNs optimized for running on small devices. Separable convolutions divide a traditional kernel convolution into two steps: a depthwise convolution to convolve the individual channels in an image and a pointwise convolution to convolve the channels together. This results in a 10-100 times speedup, depending on the exact shape of the convolution, allowing many convolutions to be stacked together and improving the overall performance of the network. A diagram of a separable convolution is shown in Fig. 5.4.

The separable convolutions are utilized within bottleneck layers. The first step in the bottleneck is a 1×1 convolution expansion layer that increases the number of channels in the feature map. This is followed by the depthwise convolution across every channel. A second 1×1 convolution acts as both the pointwise convolution in the separable convolution and as a compression layer to reduce the number of channels in the output feature map. The two 1×1 convolutions can be viewed as compression and decompression steps to keep the memory footprint of the network light. Only the most efficient convolution, the depthwise convolution, is performed when the number of




Figure 5.4: Separable Convolutions. A separable convolution separates a traditional convolution into two separate steps. First, a convolution kernel is applied to each channel individually. The output is the same size as the input. Second, the channels are convolved together using a 1×1 convolution. An activation function is often applied between the two steps.

channels in the feature map is large.

Mobilenets also utilize residual connections [85] and squeeze-excite blocks [86]. Residual connections were designed to solve the vanishing gradient problem when training very large CNNs. The residual connection acts as a shortcut by which the input to a layer can pass to the output of the layer, usually by element-wise addition between the output and the shortcut. Squeeze-excite blocks allow a network to learn relationships between the many channels of the feature map in a content aware way. The activation function used throughout the network is relu6 [87]:

$$f(x) = \begin{cases} 6, & x \geq 6 \\ x, & 0 < x < 6 \\ 0, & x \leq 0 \end{cases} \quad (5.1)$$

which is applied to every element of the feature map after each layer.

The final feature map in the network is pooled to a single value for each channel. A fully connected layer with 1024 nodes provides a rich set of features on which to perform the network task. In the case here for event classification on NOvA, the 1024 nodes are connected to four nodes, each representing one of the possible classification targets for the network, ν_e , ν_μ , NC, or cosmic ray. The output of the network is a set of four normalized scores between 0 and 1 from a softmax activation function:

$$\sigma(y_i) = \frac{e^{y_i}}{\sum_j e^{y_j}} \quad (5.2)$$

where y_i are the outputs from the network before the activation function and $\sigma(y_i)$ are the corresponding softmax scores. The scores can be interpreted as a probability for belonging to the corresponding class.

The simplest approach to using this network for the classification task would be to overlay the two views of the event as two channels to a single input of the network. However, there is not expected to be any spatial correlation between the two views which the network could utilize. Instead, the first layers of the network are separated to make a two-tower, siamese CNN where each tower takes as input one of the views of the event. With this design, the network can learn features unique to each view. For example, cosmic rays are predominantly vertical in the Y-view.

Compared to the original mobilenet v2 architecture, the number of layers was reduced and the number of convolutional kernels within each layer was optimized for both performance and runtime. In addition, pooling layers were added to the network in place of strided convolutions. Pooling layers were found to train faster and have higher performance than the original method. The full CVN network architecture used for event classification is shown in Fig. 5.5.

The training sample consisted of about 2.5 million ν_e , 2.5 million ν_μ , 2.5 million NC, and 800 thousand cosmic ray events. All ν_τ were removed since they are so rarely observed in the data. The neutrino interactions were taken from the nonswap, fluxswap, and tau MC files. The cosmic

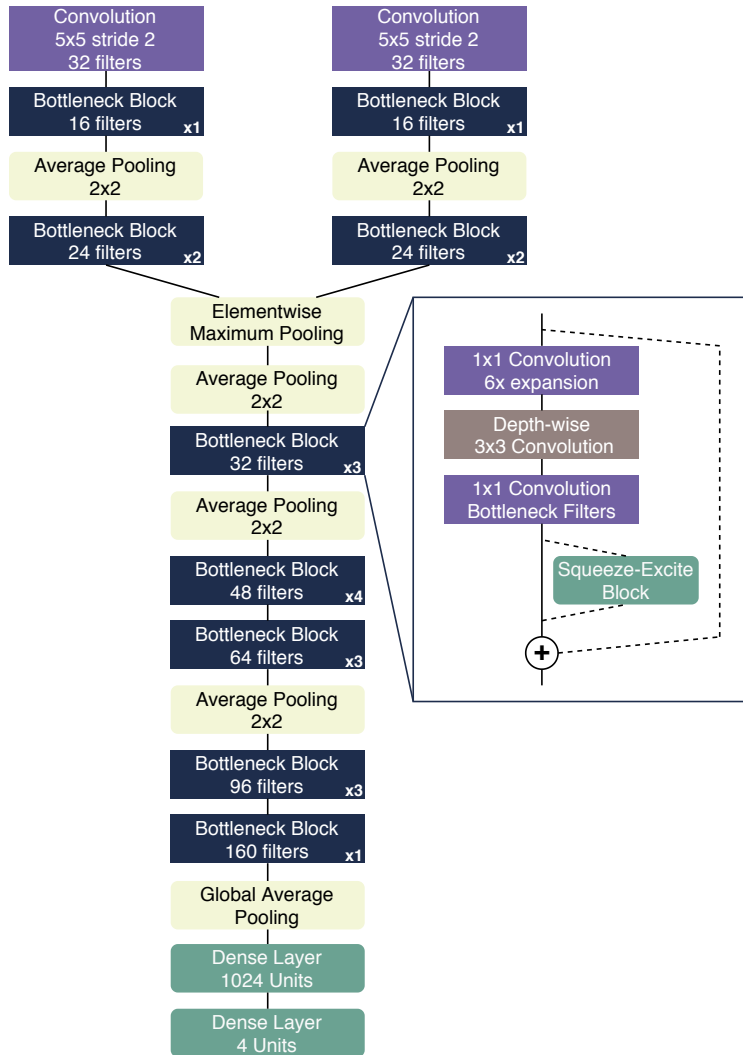


Figure 5.5: CVN Architecture. The network architecture based on mobilenet v2 used for neutrino event classification on NOvA. The first layers are separated so the network learns features unique to each view. The feature map from each view is joined together using elementwise maximum pooling. A series of bottleneck layers are used as the main feature extractors. The residual connection is shown by the dashed line circumventing the bottleneck layers. The final 4 bottleneck layers also include a squeeze-excite block. Every few layers, the size of the feature maps are reduced by a factor of 4 by pooling layers. The output of the network is a set of 4 normalized scores for each of ν_e , ν_μ , NC, and cosmic ray.

ray events are selected from events outside of the $10\ \mu\text{s}$ beam window of the simulated neutrino events. Note that these are real data cosmic ray events, not simulated. It was found that cosmic ray events were much easier to identify and so were purposefully limited to being 10% of the total training sample. 20% of the full sample was randomly separated to evaluate the performance of the network independently from the events being trained on.

The sample had a loose preselection applied to remove some of the most obvious cosmics. A veto was used which removes events which touch two opposite walls of the detector and the transverse momentum fraction from the beam was required to be less than 0.95. A number of other preselection cuts were considered, particularly whether to require that events be contained, but no significant change in performance was found so the most general training sample was used.

Before being trained on by the network, all values in a pixel map were scaled by a random value from a normal distribution with standard deviation of 0.1, $\mathcal{N}(1, 0.1)$, which is larger than NOvA's absolute calibration uncertainty discussed in Chap. 7. The intent was to reduce the networks dependence on the absolute energy scale in the calibration, one of the largest systematic uncertainties. However, it was found that the scaling also improved the performance of the network, a result commonly seen with other forms of data augmentation, the process of randomly manipulating images while training neural networks [88]. Note that the scaling was not applied when evaluating the network; the performance optimizations only considered the performance on the central value.

As mentioned in Chap. 4, there are topological differences between neutrino and antineutrino events. For this reason, two separate networks are trained for event classification: one trained on FHC MC and a second trained on RHC MC. Each network can optimize for identifying the neutrinos and corresponding backgrounds in each sample. A 10% improvement in the selection of $\bar{\nu}_e$ was found from the RHC trained network compared to the FHC trained network.

The simplest method for evaluating the network is considering the output category with the largest score as the reconstructed identity of the event. The reconstructed and true identity of

each event can be used to construct a two dimensional classification matrix. These matrices are shown in Fig. 5.6. As shown in the figure, the network achieves 90.8% (93.5%) and 87.4% (90.0%) efficiency for selecting ν_μ and ν_e respectively in FHC (RHC) with no additional selections applied. The largest misclassification of ν_μ and ν_e are as NC events with 10.9% (9.2%) and 7.0% (5.2%) misclassification for ν_e and ν_μ respectively in FHC (RHC).

Another method to evaluate the network is to examine the feature vector, the second-to-last layer of the network, which consists of 1024 values which represent the features extracted by the CNN. Of interest is if events with similar topologies have similar features extracted by the network. This can be visualized using t-Distributed Stochastic Neighbor Embedding (tSNE), a method of dimensionality reduction [89]. The tSNE method is a non-linear transformation from the 1024-dimensional space down to 2-dimensional space in a way that preserves euclidean distance between nearby points. Separation between events in this 2D space corresponds to topological separation in the high dimensional space. The tSNE for 20000 events is shown in Fig. 5.7. As shown in the figure, some overlap is seen between topologically similar categories such ν_e and NC, but little overlap is seen between topologically distinct categories such as ν_e and ν_μ . The cosmic events have been almost entirely separated from the neutrino interactions.

The network evaluation time is also critically important when comparing several networks. Due to the computational complexity of CNNs, they are well suited for efficiency on GPUs [90]. This network was trained on Nvidia Tesla P100 GPUs at the Carbonate large-memory, computer cluster at Indiana University. However, NOvA's production environment currently uses CPUs exclusively and consideration must be given to the CPU evaluation time, particularly as the amount of data and simulation created by the experiment grows. The mobilenet based network described here is 40% faster than the network architecture used by NOvA previously [81]. In addition, this network was the first network used by NOvA trained in tensorflow [91] and keras [92] which is nearly three times faster than the caffe [93] networks used previously. In total, the new network processes an




Figure 5.6: Event Classification Matrix. Events in the evaluation sample are used to fill a matrix with the X-axis showing the true identity and the Y-axis show the largest scoring category output from the network. The top (bottom) plots are for FHC (RHC). Left: The plots are normalized so that each column sums to one. The diagonal shows the efficiency for selecting each category and the off-diagonal values shows how events are misclassified. Right: The plots are normalized so that each row sums to one. The diagonal shows the purity of each selected sample and the off-diagonal shows the backgrounds to each selection.




Figure 5.7: Neutrino tSNE. The tSNE transformation of events classified by the CVN network. The color shows the true identity of each event.

event in 0.5 s which is 85% faster than networks used in previous NOvA production campaigns.

5.3 Particle Classification

Another technique employing CNNs is that of classifying individual particles produced by the neutrino interaction [94]. Identifying particles is useful for particle energy estimation and identification of final states for measuring neutrino interaction cross sections. Particle classification is done using another CNN to classify the 3D prongs produced by the reconstruction chain in Chap. 4.

The inputs to the network are constructed in the same way as the pixel maps used for CVN. The position of hits in the slice are used to determine the boundaries of the pixel map and the energy values of hits are scaled and used to fill two dimensional arrays. The hits corresponding to each view of the prong of interest are used to fill two arrays. In addition, the X and Y-view of the entire event, the same maps used for event classification, are also used. This provides the network with particle interaction context and will be discussed further in this section. The four inputs for an example event can be seen in Fig. 5.8.



Figure 5.8: Prong Classification Inputs. The 4 arrays used as input to the network for particle classification. The left two inputs are the X-view of the prong and full event respectively. The right two inputs are the Y-view of the prong and full event respectively.

The prong classification architecture is identical to mobilenet v2 [84], but the first three layers have been duplicated to make a four-tower, siamese architecture. Two of the towers take as input the two views of the prong being classified and two towers take as input the two views of the entire event. The output of the network is a normalized score between 0 and 1 for each of five possible particle types: electron, muon, charged pion, proton, or photon. Many other particles are produced in neutrino interactions, for example π^0 , but these five particles are the ones that deposit energy in the detector and can be directly observed.

The training sample was constructed from the FD MC. The entire event, including the prong, are required to be contained within the volume of the detector. A cut on the purity of the prong was included to balance realistic looking particles with prongs that have a clear identity. The cut varied per particle depending on the average purity of that particle and the number of that particle in the training sample. These cuts are: 50% purity for muons and photons, 40% purity for electrons, and 35% purity for charged pions and protons. Finally, prongs longer than 5 m were removed to reduce overall computation time. More than 95% of prongs longer than 5 m are muons and can either be classified as muons or be easily identified using other techniques [95]. Finally, the selected sample underrepresents particles that are produced less frequently in neutrino interactions, particularly charged pions, so prongs were selected at random from each category to construct a balanced sample of each particle type. About 400000 of each particle were used to train the network.

As before, the simplest method to evaluate the network performance is to construct classification




Figure 5.9: Prong Classification Matrix. Prongs in the evaluation sample are used to fill a matrix with the X-axis showing the true identity and the Y-axis show the largest scoring category output from the network. Left: The diagonal shows the efficiency for selecting each particle and the off-diagonal values shows how particles are misclassified. Right: The diagonal shows the purity of each selected sample and the off-diagonal shows the backgrounds to each selection.

matrices for the evaluation sample. The two matrices can be seen in Fig. 5.9. The network achieves more than 80% efficiency and nearly 80% purity for every particle type with no additional selections. In particular, the performance on charged pions is significantly improved compared to previous iterations of the network [81]. In previous iterations, charged pions represented a very small fraction of the total training sample. When the fraction is too small, the network will preferentially give higher scores to the other categories since that results in the best performance across the entire sample.

As seen in Fig. 5.6, the largest background fractions are between electrons and photons and also between protons and charged pions. Electrons and photons both induce electromagnetic showers which can be difficult to distinguish. Protons and charged pions both produce highly ionizing, short tracks in the detector. For this reason, it is common in some applications to sum the electron and photon score into an electromagnetic score and the proton and charged pion score into a hadronic score. This treats the network as a three label classifier with electromagnetic, muon, and hadronic as the possible categories.




Figure 5.10: Context Enriched Improvement. The elementwise difference in the efficiency and purity classification matrix between the network trained with the full event as context and the network trained with only the prong information. On the left (right), the blue quantities on the diagonal show the improvement in efficiency (purity) and the red on the off-diagonal elements shows the reduction in misclassification (backgrounds) for each category.

The inclusion of the entire event in the training provides the network with contextual information in the training. For example, electrons and photons can be distinguished by the separation of a photon induced shower from the interaction vertex². To evaluate how the entire event effects the network, a second version of the classifier was trained with only two towers. The training proceeded identically to the original network, but only the two views of the prong were trained on. The classification matrix for this two view network was evaluated and the delta was taken from the four view network in Fig. 5.9 to show the benefit of the inclusion of the event context. The resulting values can be seen in Fig. 5.10. As seen in the figure, improvements are made in the efficiency and purity of every particle type. The photon-electron contamination improves by 5% for electrons identified as photons and 3% for photons identified as electrons. The efficiency of selecting charged pions is improved by 11%.

²The vertex is not included in the training, but the network could learn to find this point.

5.3.1 π^0 Mass Peak

It is important to compare ML algorithms between data and MC to ensure they are not learning features which are unique to the sample they were trained on, in particle physics experiments this is usually the MC. The π^0 mass peak is a standard candle in particle physics for gauging the energy response of a detector to electromagnetic energy. The most common decay mode for π^0 is to two photons which will have an invariant mass of 135 MeV, the well measured mass of π^0 . The mass peak from these photons is used on NOvA as a cross check of the absolute energy calibration and the reconstruction of the event.

The invariant mass can be reconstructed from the energy and direction of the two identified photons. In this case with two massless particles, the invariant mass, $M_{\gamma\gamma}$, can be computed as

$$M_{\gamma\gamma} = \sqrt{2E_{\gamma 1}E_{\gamma 2}(1 - \cos\theta_{\gamma\gamma})} \quad (5.3)$$

where $E_{\gamma 1}$ and $E_{\gamma 2}$ are the energies of the two photons and $\theta_{\gamma\gamma}$ is the opening angle between the two photons.

To measure the energy of a photon, the total calibrated energy of all hits in the prong is summed to give a visible energy. The visible energy is corrected by a factor of 1.57 for the average dead material in the detector, which is measured to be 35.8% from a simulation of photon showers. The dead material correction accounts for energy which was deposited in the detector plastic. The cosine of the opening angle is measured as the dot product of the unit vectors pointing in the direction of the two prongs.

π^0 events are selected from the ND by requiring that an event have exactly two reconstructed prongs and be contained within the volume of the detector, but outside the muon catcher. The two prongs are required to have no more than 4 consecutive planes of hits and have a photon score from the particle classifier greater than 0.6. The resulting mass peak for these events can be seen

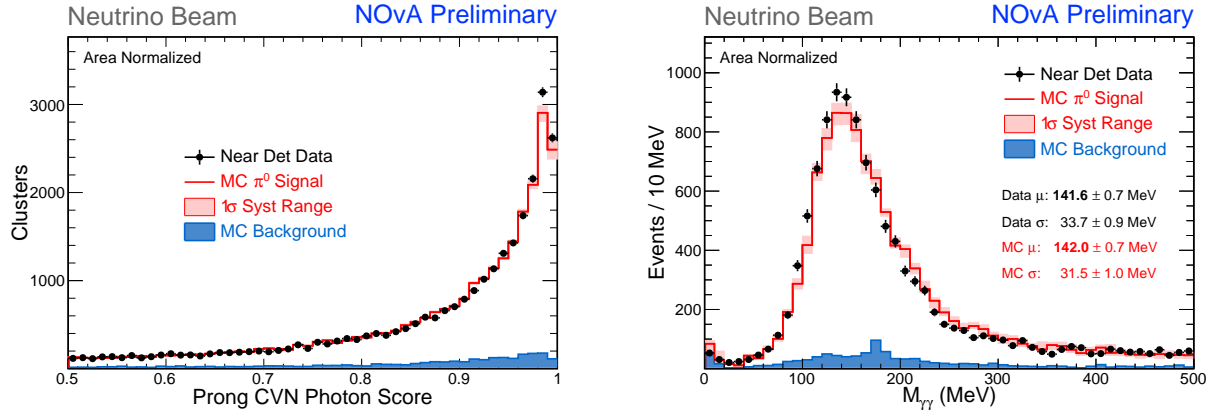


Figure 5.11: π^0 Mass Peak. ND FHC data (black) and MC (red), which has been normalized to the same area as the data, of selected π^0 events. The blue region shows selected MC events which do not have a π^0 . The shaded red region indicated NOvA’s cross section uncertainties on these events which will be discussed in Chap. 7. Left: The distribution of photon scores from the particle classifier network for each prong in the event. Right: The measured invariant mass of photon pairs.

in Fig. 5.11. The figure shows good agreement between the mean mass in data and MC. Events with $M_{\gamma\gamma} > 250$ MeV are primarily due to events with hadronic activity clustered into the photon prongs.

This selection based on the particle classifier has been compared to a more traditional selection based on the $\frac{dE}{dx}$ of the particles to remove highly ionizing particles and a nearest neighbors technique to omit muons [96]. The score threshold at 0.6 used above was matched to give the same efficiency as the traditional selection. This technique shows a 60% reduction in backgrounds.

5.4 Instance Aware Semantic Segmentation

One of the advantages of ML algorithms is that they can be designed to depend on only the most basic of reconstructed objects, the slice in the case of CVN. However, the particle classifier identifies the particle contribution to reconstructed prongs. Thus, its performance is dependent on the quality of those prongs, which is itself dependent on the reconstructed vertex.

An end-to-end algorithm which clusters particle contributions and identifies them, starting from the slice with no other dependencies, could show significant improvements. The field of




Figure 5.12: Mask R-CNN Diagram. The architecture proceeds in two stages. The first stage extracts a set of features from the input image. Then, an object detection network locates bounding boxes which, ideally, each contain all pixels from a single object. The second stage applies a position and size correction to the bounding box, classifies the object within the bounding boxes, and identifies which pixels belong to the object. Figure reproduced from [100].

computer vision calls such an algorithm instance aware semantic segmentation or simply instance segmentation, which combines two computer vision methods into a single algorithm. The first is object detection [97], an algorithm for finding the locations of objects within an image. The second is semantic segmentation [98], a pixel-wise classification technique.

The Mask R-CNN [99] algorithm for instance segmentation was adapted for use by NOvA. A diagram of the network architecture is shown in Fig. 5.12.

The network has two inputs: an input pixel map for the slice and a set of anchors to act as candidate bounding boxes. Only one view of the event is evaluated by the network at a time. The pixel map used is similar to those used for the event and particle classifier. However, the pixel map is upscaled by a factor of three in width and height so that each hit in the event translates to nine pixels in the input array. Upscaling greatly improves the networks ability to reconstruct particles

with only one or two hits. A second channel is added to the input array which is identical, but is fixed at a value of 255 for each hit. This channel enhances the network’s ability to locate hits with low energies and improves the completeness of the final particle clusters. Finally, it is beneficial in the feature extractor, described below, for the input to have dimensions with powers of two. The downstream most 15 planes are removed and the input is padded with zeros. The resulting final shape is $256 \times 256 \times 2$.

The first element of the architecture is a CNN based on a residual network with 50 layers [85]. Whenever the input shape is reduced, either by pooling or a strided convolution, the feature map is saved to act as one of five levels of the feature map. The five levels are used in a feature pyramid network [101] which upsamples one level of the feature map to be element-wise added with the level above. The result is a set of five feature maps which capture features of different sizes.

The second input to the network are the anchors, boxes in the image space which will be used to make bounding boxes. Three anchors are generated for each pixel in each level of the feature map. These anchors come from three different width/height ratios such that there is a square anchor, an anchor with half the width and double the height, and an anchor with double the width and half the height. The anchor shapes for each feature level are 8×8 , 16×16 , 32×32 , 64×64 , or 128×128 . In total, 16368 anchors are used.

The feature maps are then used to assign a score between 0 and 1 for each anchor using a region proposal network (RPN) [102]. A score of 1 signifies that that anchor contains a particle which will be reconstructed. The RPN also applies a correction to the size and position of each anchor to make bounding boxes for each particle.

Given the number of anchors analyzed, it is likely that multiple anchors are found for each particle. Any pair of anchors with large overlap³ will have the anchor with lower score removed from consideration for the remainder of the algorithm. At this step, there is ideally one, and only

³The overlap is defined by the ratio of the interesting area, the area common to both anchors, to the union, the total area of both anchors. This is often called the intersection over union (IOU).

one, anchor for each particle in the event which will make the bounding boxes used in the second stage of the architecture.

The second stage does three steps in parallel. The first is to classify the particle contained within each bounding box. The second is to correct the size and shape of the bounding once again based on the identified class. The third is to assign a score to each pixel within the bounding box based on how likely it is to have been produced by a particle of that type. The optimization of the network while learning is to improve all three of these tasks.

There are seven possible classes to be identified in this algorithm. Five are the same used in the particle classifier: electron, muon, photon, proton, or charged pion. A bounding box can also be classified as background, usually the result of clustering only detector noise. The final category consists of any other particles. Since this algorithm is capable of making clusters out of very few hits, it is probable that some are primarily from delta rays or kaons which will make up most of this category.

To conserve memory, the bounding boxes are reshaped to a 28×28 pixel mask within the network. The use of masks greatly reduces the computational load of the network when the bounding boxes are expected to be large. It also provides a fixed shape as the intermediate input to the component of the network which computes the mask scores.

Thus, the network has three outputs: the set of scores for each class, a set of bounding boxes, and a mask for each bounding box. The bounding boxes give coordinates for each particle in the event and the pixel scores identify which hits should be included into a cluster. The masks must be reshaped to the size of the bounding box and then placed on the appropriate location in the event to act as the initial set of clusters. These initial 2D clusters must go through post-processing to create the final set of prongs. First, hits outside of the pixel map boundary must be included into the particle clusters. It is assumed that any hits outside the pixel map will not belong to any new particles. Thus, Prim's algorithm, described in Chap. 4, can be used to iteratively add




Figure 5.13: Instance Segmentation Event Display. A ν_μ event with the prongs produced by instance segmentation overlaid. All hits make it into an appropriate cluster. The yellow cluster at the end of the muon track in the X-view is a clustered Michel electron.

any unclustered hits to the nearest cluster up to 80 cm away. In addition, the step which removes overlapping anchors will sometimes let through an anchor that does not contain a unique particle, particularly when it identifies a different class for the spare anchor. Electromagnetic showers with a large gap will often produce two bounding boxes, one box containing the entire shower and a second containing just the subset downstream of the gap. Whenever this is seen, the smaller of the two bounding boxes is merged with the larger bounding box. The final step is to merge the results for each view using the Kuiper test, also described in Chap. 4. The prong directions are defined from the reconstructed vertex by elastic arms.

Instance segmentation offers a number of clear advantages over the k-means clustering and the single particle classifier. It is an end-to-end algorithm with no reconstruction dependencies which would affect its performance. It is capable of reconstructing particles resulting from secondary




Figure 5.14: Reconstructed True Particles. The difference in the number of true particles that have their own reconstructed cluster between the k-means algorithm and instance segmentation.

activity in an event, such as a muon decaying into a michel electron. The traditional reconstruction is limited to reconstructing particles with some minimum energy deposited in the detector, but instance segmentation can reconstruct particles even with only one hit. An example event display depicting prongs produced by the algorithm is shown in Fig. 5.13.

The most interesting improvement in the particle clustering performance is in the number of reconstructed clusters. As mentioned, instance segmentation is capable of producing clusters even with only one hit, so many additional true particles make it into their own cluster. The number of additional true particles in a cluster compared to the traditional reconstruction is shown in Fig. 5.14. As shown in the figure, instance segmentation reconstructs, on average, almost three additional true particles in each event.

The efficiency and purity of the prongs is shown in Fig. 5.15. Instance segmentation shows an improvement in the average prong efficiency from 81% to 84% and an improvement in the average purity from 66% to 80%, a 21% improvement in purity. In particular, the fraction of prongs with 100% purity has been greatly improved at a comparable efficiency. In addition, improvement is seen in every particle type. The purity plots in Fig. 5.15 show a discontinuity at 0.5. Prongs labels




Figure 5.15: Prong Quality Comparison. Top: The efficiency (left) and purity (right) of reconstructed prongs using instance segmentation and the traditional k-means reconstruction. Bottom: The purity of prongs using instance segmentation for true particles. From left to right: electron, muon, photon, proton, and charged pion.

are identified by the highest energy contributor, so a purity less than 0.5 is only possible with three particles contributing to a prong. The steeper discontinuity at this point shows that fewer particles are usually being clustered together; this is particularly evident in the photon purity distribution.

It is important to consider the performance of the algorithm under a variety of event topologies. As an event becomes more “complex”, the difficulty in reconstructing particles increases. Figure 5.16 shows the average purity of clusters separated by the number of primaries produced in the neutrino interaction. Here, the number of primaries is being used as a proxy for the complexity of the event. Not only does instance segmentation produce higher quality prongs for all topologies, but the improvements are more significant for more complex events.

To evaluate the performance of the particle classification, we once again consider the highest score output by the network and construct classification matrices. These can be seen in Fig. 5.17.




Figure 5.16: Event Complexity. The average purity of prongs reconstructed by instance segmentation and k-means separated by the number of primaries produced in the neutrino interaction.




Figure 5.17: Instance Segmentation Classification Matrix. Prongs reconstructed by instance segmentation are used to fill a matrix with the X-axis showing the true identity and the Y-axis show the largest scoring category output from the network. Left: The diagonal shows the efficiency for selecting each particle and the off-diagonal values shows how particles are misclassified. Right: The diagonal shows the purity of each selected sample and the off-diagonal shows the backgrounds to each selection.

A new feature in these plots is the confusion between protons and photons. While high energy protons and photons are topologically very different, at low energies they are very similar, often producing only a single hit, increasing the challenge of the classification. These single hit clusters are particularly common from neutron interactions which will either collide with a nucleus to make

a proton or be captured and emit several low energy photons. Also seen in the plot is the increased misclassification between charged pions and protons compared to the single particle classifier. The training sample was balanced to have approximately equal numbers of ν_e , ν_μ , and NC, but was not balanced for the individual particle types. Charged pions are produced more rarely in neutrino interactions and as such, the network did not fully learn the topology. Future iterations of this algorithm should either balance or weight the sample to remedy this.

Chapter 6

Neutrino Oscillation Analysis

This chapter describes the analysis of four neutrino oscillation channels: $\nu_\mu \rightarrow \nu_\mu$, $\bar{\nu}_\mu \rightarrow \bar{\nu}_\mu$, $\nu_\mu \rightarrow \nu_e$, and $\bar{\nu}_\mu \rightarrow \bar{\nu}_e$. The selection and analysis of ν_e and ν_μ events to generate a FD prediction of the number of events to expect in each channel are done in parallel. The expectation and systematic uncertainty pulls are fit to the observed FD data to make a measurement of the neutrino oscillation parameters.

When relevant, figures and numbers in this chapter are scaled to the POT recorded by the analysis, 13.6×10^{20} POT FHC and 12.5×10^{20} POT RHC. When oscillations are needed, parameters measured by NOvA's last analysis [103] are used: $\Delta m_{32}^2 = 2.48 \times 10^{-3} \text{ eV}^2$, $\sin^2 \theta_{23} = 0.56$, and $\delta_{\text{CP}} = 0.0$. The remaining parameters are determined from other experiments and recent global fits [28]: $\Delta m_{12}^2 = 7.53 \times 10^{-5} \text{ eV}^2$, $\sin^2 2\theta_{12} = 0.851$, and $\sin^2 2\theta_{13} = 0.085$. The baseline is 810 km with an average density of 2.84 g/cm^3 based on Earth models [104].

6.1 Energy Estimation

Accurate reconstruction of the neutrino energy is critical for measuring neutrino oscillation parameters since the oscillation probability depends directly on the energy. The energy of the neutrino cannot be measured directly and instead must be estimated using the momenta of the products in the neutrino interaction.

6.1.1 Electron Neutrino Energy

The NOvA detector response is known to differ between hadronic and electromagnetic energy depositions [105]. The e/h ratio measures the ratio of the detector charge response to events that are entirely electromagnetic to events that are entirely hadronic. Figure 6.1 shows the response




Figure 6.1: Detector Energy Response. The NOvA detector has a different response to depositions of electromagnetic and hadronic energy. The e/h is measured to be 1.26 ± 0.02 based on the two ends of this distribution.

for different electromagnetic energy fractions in the MC and shows that for the NOvA detector $e/h = 1.26 \pm 0.02$.

Due to inefficiencies in the reconstruction, the e/h fraction cannot be used directly as a correction in the energy estimation. However, the estimator can be designed to fit for hadronic and electromagnetic energy separately to account for the e/h response. The energy estimation is fit using ν_e MC events with a basic selection, but events are weighted to produce a flat true energy distribution [106]. The flattened energy reduces bias with respect to the position of the beam peak in the MC.

All 3D prongs, constructed by the fuzzy k-means algorithm, in a selected event are identified as being either electromagnetic or hadronic using the single particle classifier from Chap. 5. Prongs are classified as electromagnetic if the sum of their electron and photon scores are greater than 0.5. All other energy in the event outside those prongs, including 2D prongs and hits outside of prongs, are identified as hadronic.

The average true neutrino energy in each bin is shown in Fig. 6.2. The total electromagnetic and hadronic energy are used to fit a quadratic function to the average neutrino energy. The resulting




Figure 6.2: ν_e Energy. The color shows the average true neutrino energy in each bin of electromagnetic and hadronic energy from FHC MC.

function is:

$$E_{\text{reco}} = \frac{1}{1 + 0.04} (1.02E_{\text{EM}} + 1.11E_{\text{Had}} + 1.44 \times 10^{-3}E_{\text{EM}}^2 + 0.11E_{\text{Had}}^2) \quad (6.1)$$

where E_{EM} is the total electromagnetic energy and E_{Had} is the total hadronic energy. The term in front corrects the energy to have a mean energy bias of 0.

The above process of constructing the 2D distribution of electromagnetic energy and hadronic energy and then fitting the energy function is done for both FHC and RHC. This gives an energy estimator for both cases with energy resolution shown in Fig. 6.3. As shown in the figure, the ν_e ($\bar{\nu}_e$) energy estimator achieves a resolution of 10.3% (9.1%) in FHC (RHC).

6.1.2 Muon Neutrino Energy

The ν_μ energy estimator sums energy contributions from the muon and hadronic energy:

$$E_{\text{reco}} = E_\mu + E_{\text{had}}. \quad (6.2)$$




Figure 6.3: ν_e Energy Resolution. The ν_e energy resolution in FHC (left) and RHC (right). The mean and RMS come directly from the histogram.

True ν_μ events are selected using a basic preselection requiring events be contained and well reconstructed [107]. The reconstructed kalman track belonging to the muon is identified using a BDT [95] trained on four variables:

- $\frac{dE}{dx}$ log-likelihood,
- scattering log-likelihood,
- track Length,
- and track hadronic fraction,

which were chosen due to their individual separation of muons from other particles in ν_μ events.

The energy of the muon produced in the event can be reconstructed from its total length, as one might expect from application of the Bethe-Bloch equation. Once again, the equation is not applied directly due to reconstruction inefficiencies in the tracking. Instead, the muon track length is fit to match its energy. The true muon energy is then fit to the reconstructed kalman track length using a piecewise-linear, spline fit. The distributions between the two are shown in Fig. 6.4. The points where the linearity changes are connected to the energy losses in the Bethe-Bloch equation.

Once the muon energy is calculated, the remaining true energy of the ν_μ is fit to the visible hadronic energy in the event, the energy not connected to the muon track, to find the total hadronic energy. The hadronic fit is also shown in Fig. 6.4.




Figure 6.4: Spline Fits. The ν_μ energy estimator separates the contributions of the muon (left) and hadronic (right) to the energy. The track length of the muon is fit by a spline to the true muon energy. The visible hadronic energy is fit by another spline to the remaining energy in the event.

The spline fits are repeated separately for FHC and RHC as well as for the ND and FD. In the ND, special consideration is needed for muons entering the muon catcher [107].

Due to the different reconstruction methods, the energy resolution of the muon, by track length, is about 2.5% while the energy resolution of the hadronic energy, by calorimetry, is about 25%. The ν_μ sample is separated into four quartiles with a quarter of the events in each based on the fraction of energy which is hadronic, $f_{\text{had}} = \frac{E_{\text{had}}}{E_{\text{reco}}}$. This gives the first f_{had} quartile, where most of the energy is in the muon, the best energy resolution and the fourth f_{had} quartile the worst energy resolution. As will be seen later, this also has the effect of moving backgrounds to the fourth quartile. The energy resolution for each quartile is shown in Fig. 6.5. The estimator achieves a resolution of 4.2% (4.0%), 6.5% (5.8%), 8.2% (7.2%), 10.5% (9.3%) in the first, second, third, and fourth quartiles respectively in FHC (RHC). The RHC energy estimator has better performance in every quartile as expected from the reduced hadronic energy in those events.

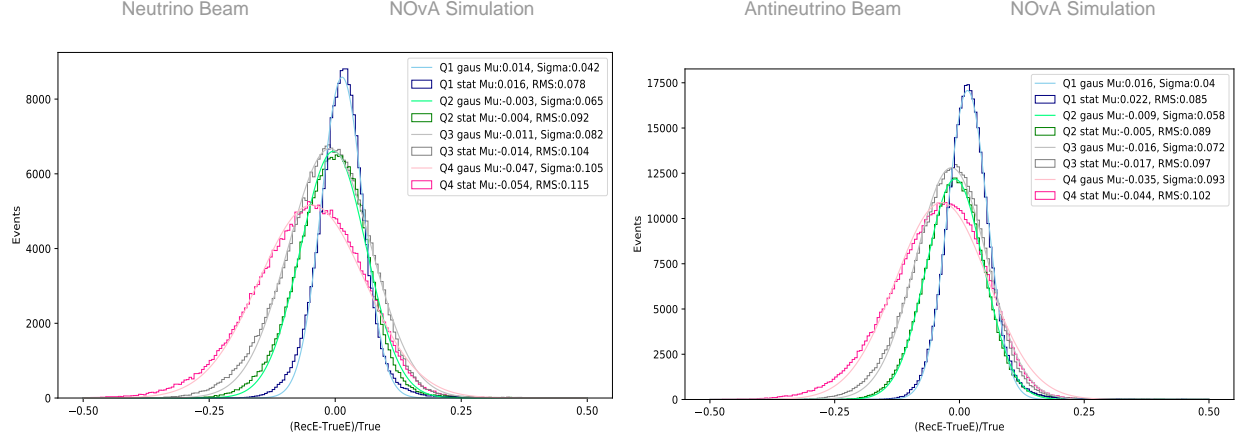


Figure 6.5: ν_μ Energy Resolution. The ν_μ energy resolution for FHC (left) and RHC (right) separated by f_{had} quartiles. The first quartile has the best energy resolution and the last quartile has the worst energy resolution.

6.2 Event Selection

6.2.1 Electron Neutrino Selection

The ν_e event selection is optimized to select the appearing $\bar{\nu}_e^{(-)}$ from $\bar{\nu}_\mu^{(-)} \rightarrow \bar{\nu}_e^{(-)}$ oscillations in FHC (RHC) mode which make up one set of the signal events for the analysis. A flow chart of the ν_e selection can be seen in Fig. 6.6.

An analysis mask is applied to veto cosmic events from the entire sample before any reconstruction is done. A first veto removes events with hits on two opposite walls of the detector. A second veto using a CNN analyses $16 \mu\text{s}$ windows of the data. The output of the CNN is a score from 0 to 1 where 1 signifies that there is a neutrino-like event in that time window [108]. Time windows with scores less than 0.2 are determined to not have a neutrino-like event and are removed.

The data quality cuts ensure a clean sample and remove any major issues with the beam or detector running state. Only periods of data taking with at least four contiguous active diblocks are used, a large enough region for most neutrinos. A set of beam quality cuts ensures the beam was in a consistent running state while data was being collected. In addition, events with reconstruction failures or APD saturation effects are removed. The event must have a reconstructed vertex, at




Figure 6.6: ν_e Event Selection. A flow chart showing each stage of the selection of ν_e events.

least one reconstructed prong, and that less than eight hits in the slice are in the same plane.

The preselection looks for events with the ν_e -like topology. The number of hits in the slice is required to be between 30 and 150, the length of the longest prong in the event is required to be between 1 m and 5 m, and the reconstructed energy is required to be between 1 GeV and 4 GeV. In addition, all of the reconstructed prongs are required to be contained within the detector: at least 63 cm from the top, 18 cm from the front and back, and 12 cm from the east and west wall and the bottom of the detector. The increased distance from the top, front, and back faces is to reduce the number of preselected cosmic events. The selection uses the start and end points of the prongs,

rather than the extent of the slice, to reduce loss of signal events from individual hits far removed from the core of the event.

A specially trained BDT removes as many remaining cosmic rays as possible [109]. The BDT was trained on:

- the number of hits in the event,
- the asymmetry in the number of hits at the start and end of the event,
- the fraction of the neutrino momentum transverse to the beam,
- the distance of all prongs in the event to each wall of the detector,
- the fraction of energy in the identified electron,
- and the width of the reconstructed electromagnetic shower.

The BDT was trained independently on FHC and RHC events passing the preselection to compare the cosmic rays to the differing neutrino topologies. The output is a score from 0 (cosmic-like) to 1 (ν_e -like). Events with scores greater than 0.49 (0.47) are kept in FHC (RHC).

The main ν_e identifier to select ν_e signal events is the CVN discussed in Chap. 5. Events with ν_e scores greater than 0.84 (0.85) are kept for FHC (RHC). These events make up the core ν_e sample. Core sample events are separated into two bins: low CVN scores from 0.84 to 0.97 (0.85 to 0.97) and high CVN scores from 0.97 (0.97) to 1 in FHC (RHC). The low and high CVN bins are further divided into bins of 0.5 GeV width.

Events failing the preselection or cosmic rejection cuts are reconsidered for selection if they have very high CVN scores. These events are on the “periphery” of the core sample and, as such, make up what is termed the peripheral sample. The most common failure in the core sample selection is the containment cut of the prongs. For this reason, no energy reconstruction is used for these events as accounting for energy outside the detector is challenging and not possible with the current energy estimators.

A second, special BDT was trained to remove cosmics from the peripheral sample. The training variables here were:

- the number of hits in the slice,
- the asymmetry in the number of hits at the start and end of the event,
- the fraction of momenta in the horizontal and vertical directions,
- the distance to the top of the detector,
- and the X, Y, and Z coordinates of the reconstructed vertex.

Once again, the BDT was trained separately for FHC and RHC. Peripheral sample events are selected if they have peripheral BDT scores greater than 0.61 (0.6) and CVN scores greater than 0.97 or BDT scores greater than 0.57 (0.56) and CVN scores greater than 0.995 in FHC (RHC).

The cut values used for the BDT and CVN scores, as well as the boundaries of the low and high CVN bins in the core sample, were optimized to maximize a Figure of Merit (FOM). The FOM used was:

$$\text{FOM} = \sum_{i=1}^N \frac{S_i}{\sqrt{S_i + B_i + (\sigma B_i)^2}} \quad (6.3)$$

where N is the total number of bins (13 in total), S_i is the number of signal ν_e events, B_i is the number of non- ν_e neutrinos and cosmic events, and σB_i accounts for an estimated 11% uncertainty in the number of background events. Tables 6.1 and 6.2 show the estimated number of selected signal and background events at each stage of the selection. The FHC (RHC) selection achieves a combined efficiency of 63% (75%) between the core and peripheral samples. The main backgrounds come from ν_e inherent to the beam and NC interactions. The ν_e in the beam make an irreducible background to the ν_e selection; in fact the high CVN bin has more than 95% purity of ν_e from any source.

6.2.2 Muon Neutrino Selection

The ν_μ event selection is optimized to select the surviving $\bar{\nu}_\mu$ from $\bar{\nu}_\mu \rightarrow \bar{\nu}_\mu$ oscillations in FHC (RHC) mode which make up the other set of the signal events for the analysis. A flow chart of the ν_μ selection can be seen in Fig. 6.7.

Cut Level	ν_e	Beam Background	Cosmics	Efficiency	Purity
No Cut	96.8	2.4×10^3	3.0×10^6	1.0	0.0
Data Quality	93.2	2.2×10^3	5.1×10^5	0.96	0.0
Preselection	71.4	1.0×10^3	2.6×10^4	0.74	0.0
Cosmic Rejection	63.5	262	353	0.66	0.09
CVN Cut	52.0	11.4	17.0	0.54	0.74
Reverse Core	38.7	1.1×10^3	4.5×10^5	0.40	0.0
CVN Cut	8.6	4.3	1.6	0.09	0.59

Table 6.1: The estimated number of selected signal and background at each stage of the ν_e event selection in FHC. The top are the cuts used in the core sample selection and the bottom is the peripheral sample selection. The ν_e column is the signal, appearing ν_e events. The beam background is the total background from beam related sources. The right shows the purity and efficiency of the selection at each stage.

Cut Level	$\bar{\nu}_e$	Beam Background	Cosmics	Efficiency	Purity
Veto	27.0	1.0×10^3	1.9×10^6	1.0	0.0
Data Quality	26.4	693	3.3×10^5	0.98	0.0
Preselection	20.9	455	1.6×10^4	0.77	0.0
Cosmic Rejection	19.3	87.8	296	0.72	0.05
CVN Cut	17.3	9.0	0.589	0.64	0.64
Reverse Core	8.6	461	2.9×10^5	0.32	0.0
CVN Cut	3.01	2.6	1.0	0.11	0.45

Table 6.2: The estimated number of selected signal and background at each stage of the ν_e event selection in RHC. The top are the cuts used in the core sample selection and the bottom is the peripheral sample selection. The $\bar{\nu}_e$ column is the signal, appearing $\bar{\nu}_e$ events. The beam background is the total background from beam related sources. The right shows the purity and efficiency of the selection at each stage.

The same analysis mask and data quality cuts in the ν_e selection are used for the ν_μ selection. The one additional requirement is that there is a reconstructed kalman track in the event to be used for the ν_μ energy estimator.

All reconstructed prongs in the event are required to be 60 cm from the top, 18 cm from the front and back, and 12 cm from the east, west, and bottom of the detector. Because of the importance of the reconstructed track length in the energy estimation, special containment constraints are used for the identified muon track. There must be at least 2 (3) planes of active detector between the start (end) of the kalman track and the front (back) of the detector. The forward and backward projection of the track must pass through 6 cells before reaching any edge of the detector.




Figure 6.7: ν_μ Event Selection. A flow chart showing each stage of the selection of ν_μ events.

Cosmic rejection for the ν_μ selection is also done using a BDT trained separately for FHC and RHC [110]. The BDT was trained using properties of the kalman track identified as the muon:

- the cosine of the angle with the beam direction and the vertical axis,
- the length of the track,
- the maximum vertical coordinate of the track,
- the distance of the track from each detector wall,
- fraction of hits in the track from the entire event,
- and the fraction of the track momentum transverse to the beam direction.

As before, the output score is between 0 (cosmic-like) and 1 (ν_μ -like) with a score greater than 0.45 being used for the selection.

Cut Level	ν_μ	Beam Background	Cosmics	Efficiency	Purity
Veto	602	730	2.5×10^6	1.0	0.0
Data Quality	569	652	4.7×10^5	0.95	0.0
Containment	249	434	7.0×10^3	0.41	0.03
Cosmic Rejection	239	385	1.5×10^3	0.40	0.11
CVN Cut	201	15.7	6.0	0.33	0.96

Table 6.3: The estimated number of selected signal and background at each stage of the ν_μ event selection in FHC. The ν_μ column is the signal, surviving ν_μ events. The beam background is the total background from beam related sources. The right shows the purity and efficiency of the selection at each stage.

Cut Level	$\bar{\nu}_\mu$	Beam Background	Cosmics	Efficiency	Purity
Veto	193	382	1.6×10^6	1.0	0.0
Data Quality	186	346	3.0×10^5	0.96	0.0
Containment	86.7	200	4.4×10^3	0.45	0.02
Cosmic Rejection	84.9	169	635	0.44	0.10
CVN Cut	79.5	26.8	1.7	0.41	0.97

Table 6.4: The estimated number of selected signal and background at each stage of the ν_μ event selection in RHC. The ν_μ column is the signal, surviving ν_μ events. The beam background is the total background from beam related sources. The right shows the purity and efficiency of the selection at each stage.

The identified muon track must have a score from the muon BDT of more than 0.3 and the energy of the event must be less than 5 GeV. The main identifier for ν_μ events is the output ν_μ score from CVN which is required to be greater than 0.8. The cuts were optimized to maximize the FOM in Eq. 6.3, but only in the region between 1 GeV to 2 GeV which contains the region of maximum oscillations.

The selected events are divided into four equally sized quartiles of f_{had} determined by the energy estimator. Each quartile is further divided into bins of energy with width varying from 0.1 GeV in the region between 1 GeV to 2 GeV up to 1 GeV width in the region above 4 GeV. The small bin width gives increased sensitivity in the region with the maximum disappearance probability around 1.6 GeV. Tables 6.3 and 6.4 show the estimated number of selected signal and background events at each stage of the selection. The FHC (RHC) selection achieves a purity of 96% (97%). The main backgrounds come from cosmic rays and NC interactions.




Figure 6.8: Extrapolation Diagram. The extrapolation of $\nu_\mu \rightarrow \nu_\mu$ from the ND to the FD. The ND data is extrapolated by a F/N ratio from the MC in bins of true energy.

Since the scores output from CVN sum to unity and the cuts used for both ν_μ and ν_e are more than 0.5, there is no chance for an event to be selected by both the ν_μ and ν_e selection and the two datasets are distinct.

6.3 Far Detector Prediction

The prediction of events in the FD is constructed by extrapolating the ND data to the FD. The purpose of the extrapolation is to use the ND data to make a data-driven prediction of the number of events to expect in the FD. An overview of the extrapolation is shown in Fig. 6.8. A brief summary will be given here with more details in the subsections below.

Selected events in the ND data and MC in reconstructed energy bins are migrated to true energy bins using a reconstructed-to-true unfolding matrix constructed from the MC. Similarly, selected MC events in the FD are converted to true energy so that a F/N ratio can be computed for each true energy bin. The F/N ratio is applied to the ND data to give a spectrum of FD events in true energy. At this point, oscillations can be applied, usually either $P_{\nu_\mu \rightarrow \nu_\mu}(E_{\text{true}})$ or $P_{\nu_\mu \rightarrow \nu_e}(E_{\text{true}})$ where E_{true} is the center of a bin of true energy. Finally, the reconstructed-to-true unfolding matrix is inverted to convert the FD true energy distribution to a reconstructed energy distribution.

The F/N extrapolation procedure is carried out for each neutrino component in the ND. Note




Figure 6.9: ND Event Spectra. Spectra showing FHC events passing the ν_e and ν_μ selections in the ND. The ν_μ spectrum is used to construct the signal ν_μ and ν_e FD spectra. The ν_e spectrum is used to construct the ν_e background spectrum. The backgrounds for each selection are shown.

that the signal in the ND for both ν_μ disappearance and ν_e appearance are ν_μ events. Events passing the ν_e selection in the ND are used to construct the background to the ν_e selection at the FD. The ND selections were not discussed in the last section, but they are identical to the FD selection with modified containment cuts for the size of the detector and reduced cosmic rejection cuts since the ND is naturally underground with a low cosmic-ray rate. The energy range on the ν_e selection is extended to 0.5 GeV to 4.5 GeV. Most importantly, the cuts on the muon BDT for the ν_μ selection and CVN for the ν_e and ν_μ selections are identical between the two detectors. A spectrum of ND events with data and MC are shown in Fig. 6.9.

6.3.1 Signal Extrapolation

The signal extrapolation is carried out by the ν_μ in the ND N_{α, S_μ}^{Data} and N_{α, S_μ}^{MC} [111]. In this notation, N indicates a ND spectrum, S_μ indicates events passing the ν_μ selection, α is one neutrino component of the selection, and *Data* or *MC* indicates the spectrum is derived from the data or MC. For the signal extrapolation, the composition of the data and MC are assumed to be the same, in other words the number of backgrounds in the data is equal to the number of backgrounds in

the MC. The data must first be translated from bins of reconstructed energy to bins of true energy:

$$N_{\alpha, S_\mu}^{Pred}(E_i^T) = \sum_k \frac{N_{\alpha, S_\mu}^{Data}(E_k^R) \cdot N_{\alpha, S_\mu}^{MC}(E_i^T, E_k^R)}{N_{\alpha, S_\mu}^{MC}(E_k^R)} \quad (6.4)$$

where $E^{T(R)}$ indicates a true (reconstructed) energy, the superscript *Pred* now indicates that this is a data-driven, predicted spectrum, and the sum is over the bins of reconstructed energy. This sum is equivalent to applying the reconstructed-to-true unfolding matrix.

The predicted ND spectrum can then be extrapolated to the FD. For the FD ν_μ :

$$F_{\alpha \rightarrow \beta, S_\mu}^{Pred}(E_i^T, E_j^R) = \frac{N_{\alpha, S_\mu}^{Pred}(E_i^T) \cdot F_{\beta, S_\mu}^{MC}(E_i^T, E_j^R)}{N_{\alpha, S_\mu}^{MC}(E_i^T)} \quad (6.5)$$

where F indicates a FD spectrum, $\alpha \rightarrow \beta$ indicates a neutrino flavor transition, and i and j are bins of true and reconstructed energy, respectively. The flavor transition is not the oscillation probability, but rather the F/N ratio for a given transition at 100% oscillation probability, independent of the oscillation parameters. Recall that the nonswap FD MC files are generated with $P(\bar{\nu}_\mu \rightarrow \bar{\nu}_\mu) = 100\%$ for the ν_μ signal extrapolation and the fluxswap files are generated with $P(\bar{\nu}_\mu \rightarrow \bar{\nu}_e) = 100\%$ for the ν_e signal extrapolation discussed below.

For ν_μ signal, the transition, $\alpha \rightarrow \beta$, is either $\nu_\mu \rightarrow \nu_\mu$ or $\bar{\nu}_\mu \rightarrow \bar{\nu}_\mu$ in both FHC and RHC. The Wrong Sign (WS), the $\bar{\nu}_\mu$ in FHC or the ν_μ in RHC, are treated as signal in the ν_μ analysis. However, the probability can now be applied and summed across bins of true energy to give the FD prediction in bins of reconstructed energy:

$$F_{\alpha \rightarrow \beta, S_\mu}^{Pred}(E_j^R) = \sum_i F_{\alpha \rightarrow \beta, S_\mu}^{Pred}(E_i^T, E_j^R) \cdot P_{\alpha \rightarrow \beta}(E_i^T) \quad (6.6)$$

where P is the oscillation probability for the transition. Summing over the two possible transitions $\alpha \rightarrow \beta$ gives the total signal prediction.




Figure 6.10: Transverse Momentum Spectra. The ND data and MC and FD MC of ν_e (left) and ν_μ (right) FHC events. The ND and FD $|p_T|$ differ due to the different geometries between the detectors.

The signal extrapolation for ν_e is identical, but uses the ν_e selection, S_e , at the FD:

$$F_{\alpha \rightarrow \beta, S_e}^{Pred}(E_i^T, E_j^R) = \frac{N_{\alpha, S_\mu}^{Pred}(E_i^T) \cdot F_{\beta, S_e}^{MC}(E_i^T, E_j^R)}{N_{\alpha, S_\mu}^{MC}(E_i^T)} \quad (6.7)$$

and

$$F_{\alpha \rightarrow \beta, S_e}^{Pred}(E_j^R) = \sum_i F_{\alpha \rightarrow \beta, S_e}^{Pred}(E_i^T, E_j^R) \cdot P_{\alpha \rightarrow \beta}(E_i^T) \quad (6.8)$$

where the reconstructed energy for the FD now comes from the ν_e energy estimator. The signal transition for ν_e is $\nu_\mu \rightarrow \nu_e$ for FHC and $\bar{\nu}_\mu \rightarrow \bar{\nu}_e$ for RHC. Here, the WS ν_e are treated as backgrounds.

The two detectors are functionally equivalent, but have a major difference in their size. This difference has the effect of sculpting the energy distribution in the ND where muons with large opening angles from the beam cannot be contained [112]. The transverse momentum from the beam, $|p_T|$, is illustrated in Fig. 6.10 for both detectors.

Both ν_μ and ν_e signal extrapolation are separated into three quantiles of $|p_T|$ with equal number

of events in each. The extrapolation is done independently for each $|p_T|$ quantile, but the three quantiles are summed after the extrapolation procedure. As will be shown in Chap. 7, this extrapolation method reduces systematic uncertainties related to differences between the two detectors.

The ν_μ signal extrapolation is done independently for each f_{had} quartile. The f_{had} quartiles remain separate to be fit independently. The quartile extrapolation is not used for the ν_e signal extrapolation.

Finally, the ν_e peripheral sample has no corresponding selection in the ND, and so is determined directly from the FD MC in the same style as the minor backgrounds in the following section.

6.3.2 Background Extrapolation

The major backgrounds in the ν_e prediction are $\nu_e \rightarrow \nu_e$, $\bar{\nu}_e \rightarrow \bar{\nu}_e$, $\nu_\mu \rightarrow \nu_\mu$, $\bar{\nu}_\mu \rightarrow \bar{\nu}_\mu$, and NC \rightarrow NC. The other possible ν_e backgrounds and all ν_μ backgrounds are considered minor backgrounds since they are much less common than the major backgrounds. The background extrapolation is done in two different ways for major backgrounds and minor backgrounds.

For the extrapolation of major backgrounds, the composition of the data and MC are not assumed to be the same and the data must be decomposed using two empirical methods [113]. The two methods are Beam Electron Neutrino (BEN) decomposition and Michel Decomposition. The decomposition addresses the data/MC discrepancy seen in the ν_e events in Fig. 6.9. Note that the fraction of neutrinos and antineutrinos within each component is still assumed to be correct in the simulation.

BEN decomposition measures the number of ν_e inherent in the beam that come from either pion or kaon decays in the beam. The ND ν_μ selection provides a sample of contained ν_μ events which are expected to primarily come from the decay of charged pions in the beam which are well modeled. The discrepancy between data and MC in the energy distribution of this sample is used to construct a weight for pion events to correct the pion component of the ν_e spectrum.

The ND ν_μ selection is repeated with the containment cuts reversed, which selects a sample of uncontained events with both pion and kaon decay daughters. The pion component is corrected using the previously determined weights and the remaining discrepancy is attributed to the kaon component.

Michel decomposition resolves the ν_μ and NC component of the ν_e sample. Muons in ν_μ events decay into Michel electrons, which are identified using a special Michel electron finder [114]. This gives a handle on the ratio of ν_μ events to NC events. The contained and uncontained distributions for BEN, the Michel distribution, and a corrected ν_e distribution are shown in Fig. 6.11. Together, BEN and Michel decomposition determine the composition of the ND FHC data with the ν_e selection.

The RHC data lacks sufficient statistics to do a data-driven decomposition. Instead, the fraction of each component is assumed to be correct, but not the overall normalization. So, each component is scaled proportionally to match the number of data events within each bin.

The major ν_e backgrounds are all extrapolated from events passing the ν_e selection at the ND:

$$F_{\alpha \rightarrow \beta, S_e}^{Pred}(E_i^T, E_j^R) = \frac{N_{\alpha, S_e}^{Data}(E_j^R) \cdot F_{\beta, S_e}^{MC}(E_i^T, E_j^R)}{N_{\alpha, S_e}^{MC}(E_j^R)} \quad (6.9)$$

where N_{α, S_e}^{Data} is determined from the decomposition. Note that the ND events are not migrated to true energy. These events are misidentified and as such the ν_e energy estimator is not expected to have sensible performance. In addition, S_e is used for all major backgrounds so that misidentified events in the ND are used to constrain the rates of misidentified events in the FD.

The minor backgrounds are not extrapolated and are instead determined directly from the FD MC:

$$F_{\alpha \rightarrow \beta, S_\alpha}^{Pred}(E_i^T, E_j^R) = F_{\alpha \rightarrow \beta, S_\alpha}^{MC}(E_i^T, E_j^R) \quad (6.10)$$

where S_α is either S_μ for ν_μ backgrounds or S_e for ν_e backgrounds.




Figure 6.11: ND Background Decomposition. Top: The contained (left) and uncontained (right) samples used to determine the inherent ν_e component of the beam with BEN decomposition. Bottom Left: The distribution of Michel electrons used to resolve the ν_μ and NC background to the ν_e sample. Bottom Right: The fully decomposed and corrected ν_e sample. The data and corrected MC perfectly match by construction.

6.3.3 Measurement of Cosmic Ray Backgrounds

The NOvA FD records nearly 150 kHz of cosmic rays which is clearly independent of the ND. Thus, the rate of cosmic-ray backgrounds to the selection can be measured directly from the FD data. The background is determined from two sources. The number of events is determined from the timing sidebands of the NuMI trigger stream. The beam window plus a $10 \mu\text{s}$ gap on each side are omitted from the $550 \mu\text{s}$ NuMI trigger readout. The NuMI sidebands do not have enough statistics to properly determine the shape, so the shape of the predicted spectrum is determined from the




Figure 6.12: Predicted ν_e Spectra. The complete prediction of ν_e signal (purple and green), beam background (pink), and cosmic (blue) for FHC (left) and RHC (right). The prediction is divided into samples of low CVN score, high CVN score, and the peripheral sample.

cosmic trigger data stream instead. The selected events from both samples, either the ν_μ or ν_e event selections, are used to find selected backgrounds for both analyses.

The distribution from the cosmic trigger stream is scaled to match the total number of events from the NuMI timing sidebands. Then, the total livetime analyzed in the timing sidebands is scaled to match the livetime within the beam window, approximately $10 \mu\text{s}$ for each beam spill. This gives an estimate of the number of events expected to have been selected within the beam window.

The signal, background, and cosmic prediction together make a complete estimation of the number of events to expect in each bin of reconstructed energy in the FD when examining the data. The complete ν_e prediction is shown in Fig. 6.12 with a breakdown by component in Table 6.5 and the ν_μ prediction in Fig. 6.13 with corresponding breakdown in Table 6.6.




Figure 6.13: Predicted ν_μ Spectra. The complete prediction of ν_μ signal (purple and green), beam background (grey), and cosmic (blue) for FHC (top) and RHC (bottom). The prediction is divided into the four quartiles of hadronic energy fraction.

		Low PID	High PID	Peripheral	Total
FHC	App. ν_e	11.77	49.46	10.25	71.49
	App. $\bar{\nu}_e$	0.17	0.64	0.20	1.02
	Beam $\nu_e + \bar{\nu}_e$	2.39	7.70	3.53	13.63
	$\nu_\mu + \bar{\nu}_\mu$	0.76	0.19	0.14	1.09
	$\nu_\tau + \bar{\nu}_\tau$	0.19	0.30	0.13	0.62
	NC	4.26	1.36	0.28	5.89
	Cosmic	1.32	0.22	1.68	3.22
	Total	20.86	59.87	16.21	96.95
RHC	App. $\bar{\nu}_e$	2.03	13.67	2.81	18.52
	App. ν_e	0.47	1.73	0.51	2.71
	Beam $\nu_e + \bar{\nu}_e$	0.96	4.61	2.01	7.58
	$\nu_\mu + \bar{\nu}_\mu$	0.12	0.05	0.06	0.22
	$\nu_\tau + \bar{\nu}_\tau$	0.09	0.21	0.08	0.37
	NC	1.29	0.50	0.11	1.91
	Cosmic	0.42	0.14	0.97	1.53
	Total	5.38	20.91	6.55	32.84

Table 6.5: The predicted number of appearing ν_e and $\bar{\nu}_e$ signal and background events in FHC (top) and RHC (bottom). The predictions are divided into the core low and high PID and the peripheral sample. The far right column has the totals.

		Q1	Q2	Q3	Q4	Total
FHC	Surv. ν_μ	41.60	39.96	42.55	48.99	173.10
	Surv. $\bar{\nu}_\mu$	5.31	2.72	2.27	1.26	11.56
	NC	0.10	0.16	0.40	1.81	2.47
	Cosmic	0.36	0.36	0.89	3.50	5.09
	Other	0.14	0.14	0.16	0.33	0.77
	Total	47.52	43.34	46.27	55.88	193.00
RHC	Surv. $\bar{\nu}_\mu$	19.98	16.78	15.60	14.11	66.47
	Surv. ν_μ	3.50	4.74	6.09	9.76	24.08
	NC	0.02	0.04	0.09	0.63	0.78
	Cosmic	0.06	0.03	0.11	0.73	0.94
	Other	0.06	0.08	0.09	0.12	0.35
	Total	23.62	21.66	21.98	25.36	92.62

Table 6.6: The predicted number of surviving ν_μ and $\bar{\nu}_\mu$ signal and background events in FHC (top) and RHC (bottom). The predictions are divided into the four quartiles of hadronic energy fraction. The far right column has the totals.

6.4 Fitting and Sensitivities

Once the FD data is collected, the prediction is fit to the data. The fit is the minimization of a Poisson log-likelihood:

$$-2\mathcal{L}(\bar{\theta}) = -2 \sum_{i=1}^N \left[E_i(\bar{\theta}) - O_i + O_i \ln \frac{O_i}{E_i(\bar{\theta})} \right] \quad (6.11)$$

where $\bar{\theta}$ is a set of oscillation parameters, i is a bin of reconstructed energy being used in the fit, E_i is the predicted number of events in the bin which depends on the oscillation parameters, and O_i is the observed number of data events in the bin. The best fit point is the set of oscillation parameters that minimizes the log-likelihood.

Nearly all ν_e analysis bins and ν_μ analysis bins near the oscillation maximum are subject to Poissonian fluctuations. In addition, there is a physical boundary at $\sin^2 \theta_{23} = 0$ and $\sin^2 \theta_{23} = 1$ and a cyclical boundary in δ_{CP} which must be considered when fitting in the oscillation phase space. To achieve accurate confidence intervals, Feldman-Cousins corrections are applied [115] which will be included in all results presented in Chap. 8.

Chapter 7

Systematic Uncertainties

Systematic uncertainties characterize how well neutrino models and the NOvA detection mechanism are understood. The two detector design of the NOvA experiment helps to reduce or even cancel many sources of systematic uncertainties due to the equivalent technology used in both. This chapter will describe the treatment and sources of systematic uncertainties used in the measurement of neutrino oscillation parameters by NOvA.

The POT and oscillation parameters used to create the figures are identical to the previous chapter.

7.1 Systematics in the Analysis

Systematic uncertainties are accounted for in the analysis by deliberate changes to the ND and FD MC. The entire analysis is repeated using the shifted MC, energy estimation, event selection, and extrapolation, to achieve a new, systematically shifted prediction. The shifts can be introduced in three possible ways:

- **Reweighting.** For systematics where the effect alters the event rate, effected events can have their weight changed when filling the event spectra.
- **Modifying Reconstructed Variables.** In cases where the systematic has a measureable effect on the reconstruction, the quantity is manually changed based on the systematic.
- **New Simulation.** For all other systematics, the entire simulation must be redone with the systematic shift applied at the generation step.

For most systematics, the $\pm 1\sigma$ and $\pm 2\sigma$ shifts are evaluated unless otherwise noted. When resimulation is needed, only the $\pm 1\sigma$ are evaluated for computational reasons.

Notice that in Sec. 6.3 all extrapolations from the ND to the FD are of the form:

$$F^{Pred} = \frac{N^{Data} \cdot F^{MC}}{N^{MC}} \quad (7.1)$$

which can be trivially rewritten as

$$F^{Pred} = N^{Data} \frac{F^{MC}}{N^{MC}} = \frac{N^{Data}}{N^{MC}} F^{MC} \quad (7.2)$$

where in the first form the ND data is being extrapolated to the FD by a F/N ratio from the MC and in the second form the FD MC is being corrected by Data/MC discrepancies in the ND. In the first rewritten form, it is clear that increasing both the ND and FD MC by the same proportion will cancel and have no effect on the resulting prediction. Similarly, in the second form it is clear that data/MC discrepancies in the ND will cancel as long as the discrepancy is identical in the FD. The original form is used to demonstrate both of these simultaneous effects.

Systematic uncertainties that are correlated between the ND and FD will cancel and have a very small effect on the resulting measurements. It is only the differences in the effect between the two detectors that survives in the measurement. Since the two detectors use the same technology, many systematic uncertainties cancel, particularly those related to the neutrino flux and neutrino cross sections.

Figure 7.1 shows the 1σ uncertainty of seven systematic categories on the expected number of signal ν_e ($\bar{\nu}_e$) events in FHC (RHC) beam mode both with and without the extrapolation. The systematic uncertainty related to lepton reconstruction is increased from the extrapolation in $|p_T|$ quantiles, but reduces the uncertainty in the other categories. The systematic uncertainty in the number of ν_e ($\bar{\nu}_e$) events is reduced from 14% (11.5%) to 4% (3.5%) in FHC (RHC) using the extrapolation.




Figure 7.1: Systematics Extrapolation Comparison. The 1σ systematic uncertainty in the number of expected ν_e events in FHC (left) and the $\bar{\nu}_e$ events in RHC (right). Within each category, the effect of each systematic is combined in quadrature. The blue bar shows the uncertainty by only considering the FD MC. The red bar shows the uncertainty after extrapolating from the ND to the FD.

7.2 Sources of Systematic Uncertainties

NOvA considers more than 100 systematic uncertainties which can be separated into seven general categories shown in Fig. 7.1. Each category will be discussed here with consideration of which systematics are correlated or uncorrelated between detectors.

7.2.1 Detector Calibration

Systematics related to the calibration come in three types: an absolute energy uncertainty, a shape uncertainty, and a detector aging uncertainty [116]. These systematics are evaluated by repeating the simulation with changes made at the calibration step.

A number of methods of gauging an absolute energy scale have been discussed in this dissertation including the dE/dx of stopping muons, the Michel electron energy spectrum, and the π^0 mass peak and several exist which have not been shown yet. None of these have perfect agreement between data and MC, but the most discrepant is the dE/dx of protons. Candidate protons in the ND exhibit a 5% disagreement between data and MC. It is difficult to examine the disagreement in the FD due to the lack of neutrino statistics. So, the 5% disagreement is taken as an absolute calibration

uncertainty and a relative calibration uncertainty. The absolute uncertainty is treated as fully correlated (both detectors shifted in the same direction) and the relative uncertainty is treated as fully anticorrelated (both detectors shifted in opposite directions) between the two detectors.

The energy response of the detector is different at the ends of the cells compared to the middle, which is accounted for imperfectly in the calibration. A change to the calibration is introduced to cover the discrepancy seen in the reconstructed and true energy of events near the edges of the detector. The effect of this systematic is expected to be correlated between the ND and FD.

The number of hits observed in the FD is seen to decrease at a rate of about 0.2%/year. The effect is not fully understood, but is attributed to the aging of the mineral oil resulting in a reduced light yield and some hits not exceeding the trigger threshold. To cover this, the light production model is reduced by 4.5%/year to match the reduction seen in the number of hits. The calibration is redone to correct the reduced yield to match the original energy response, but with some small fraction of hits missing. The direction of this systematic is clearly one-directional, so only the $+1\sigma$ shift is evaluated.

7.2.2 Detector Response

The light recorded by the detector has two sources: the scintillation light and Cherenkov radiation. The detector response systematics address uncertainties in the light production model from both sources. Once again, they are evaluated by repeating the simulation.

The light level model is tuned independently for each detector. As such, the size of the systematic differs for each detector and they are treated as two uncorrelated systematics. In the ND, the light production is shifted $\pm 10\%$ in both X and Y-views. In the FD, it is shifted $\pm 16\%$ in the X-view and an inverse $\mp 6\%$ in the Y-view. With the light production modified, the calibration is also redone to match the original energy response.

A second systematic uncertainty in the light model addresses the data/MC discrepancy in

the dE/dx of protons noted above. The light level model parameters are modified to reduce the discrepancy from 5% to $< 1\%$, but without changing the detector response to muons. The changes are primarily to the production of Cherenkov light. Because this is clearly one directional, only the $+1\sigma$ uncertainty is measured.

7.2.3 Neutrino Cross Sections

The neutrino cross section and Final State Interaction (FSI) uncertainties come from reweight parameters in the GENIE simulation, called “knobs”, which can be applied to our simulated neutrino interactions based on truth quantities. The uncertainties are based on the cross section central value tune with input from theoretical models and empirical evidence from NOvA and other experiments. There are nearly 80 total uncertainties in this category, for more details see Ref. [117].

Uncertainties are included for all neutrino interaction modes. The uncertainties for all modes come from GENIE knobs and the cross section central value tune, though the knobs are altered in some cases to match theoretical models or experimental needs.

- Quasi-Elastic Scattering. GENIE includes an uncertainty on the normalization and shape of the z-expansion-based axial form factor model. The normalization uncertainty is $+20/-15\%$. GENIE indicates that the remaining shape knobs are expected to be correlated, so NOvA developed independent weights using BDTs which are varied along with the normalization.

Long-range interactions between nucleons modify the kinematics of QE interactions which are modelled by random phase approximation [118]. The effect to is to significantly suppress the cross section at low Q^2 , where Q is the four momentum transferred to the nucleus, and slightly enhance at high Q^2 . The MINERvA experiment has constructed a set of nuclear model weights [119] which are used as systematic uncertainties to characterize how well the effect is understood.

- Resonant Production. An excess of observed RES events compared to prediction suggests

that the RES cross section should be suppressed at low Q^2 . GENIE has included this suppression into its new models, but an uncertainty is included in this analysis which considers whether more or less suppression is needed. A parameterization of the effect from the MINOS experiment which increases the suppression slightly compared to GENIE was used to constrain the allowed range [120]. The difference between the NOvA MC and the MINOS range is taken as the uncertainty.

- Deep Inelastic Scattering. The normalization uncertainty for DIS interactions in GENIE is 50% below a hadronic invariant mass of $W = 2 \text{ GeV}$ for one or two pion production, but there are no corresponding uncertainties above this invariant mass or with more than two produced pions. This DIS normalization uncertainty is treated as $32 (\{0\pi, 1\pi, 2\pi, > 2\pi\} \times \{\text{CC, NC}\} \times \{\nu, \bar{\nu}\} \times \{\text{interaction on neutron, proton}\})$ independent 50% uncertainties up to a larger invariant mass of $W < 3 \text{ GeV}$ covering all types of DIS interactions. In addition, the unphysical change from 50% to 0% uncertainty is removed by linearly scaling the uncertainty from 50% at $W = 3 \text{ GeV}$ to 5% at $W = 5 \text{ GeV}$ where the 5% for large W is suggested by external measurements [121].

In DIS interactions, there is some range where the interacting quark will propagate through the nuclear medium before hadronizing known as the formation zone. In GENIE, the default values are 0.342 fm for pions and 2.3 fm for nucleons both with a 50% uncertainty. In practice, there are not enough individual events with long ranges to properly reweight to large shifts. Instead, the two uncertainties are treated as fully correlated and the two are shifted together in a single formation zone uncertainty.

- Meson Exchange Currents. The normalization uncertainty for MEC interactions is determined by comparison to two different theoretical MEC models. The alternative models are rescaled so that the cross section matches that of the used model at high energy, $E_\nu = 10 \text{ GeV}$. The ratio of the rescaled models to the GENIE model define an envelope which gives an

energy-dependent, asymmetric systematic uncertainty.

The shape uncertainty for MEC interactions is evaluated by modifying the non-MEC MC to be more QE-like, more low energy events, or more RES-like, (more high energy events). Using the modified MC, the central value MEC tune is re-evaluated and the difference is taken as a systematic uncertainty.

The fraction of nucleon pairs, np or nn (pp) for neutrinos (antineutrinos), in MEC events in the model used is

$$\frac{np}{np + nn(pp)} = 0.69(0.66)_{-5\%}^{+15\%} \quad (7.3)$$

for neutrinos (antineutrinos).

- ν_μ and ν_e Differences. ν_μ and ν_e interactions are very similar, but differ primarily due to the masses of the corresponding charged lepton that is produced. Uncertainties are included to allow for radiative corrections and second class currents in the cross sections.
- ν_τ Interactions. The ν_τ cross section is poorly constrained by global data. One measurement is from OPERA of $(1.2 \pm 0.6) \times \sigma_{\text{SM}}$ [122] where σ_{SM} is the cross section predicted by the standard model used as the default in GENIE. The change in the central value from the standard model and the 50% uncertainty are combined in quadrature to give 63%. Finally, the 63% is applied to the roughly 90% of the ν_τ background in the ν_e and ν_μ samples which are CC interactions and rounded up. This results in a 60% uncertainty in the total ν_τ cross section used in the analysis.
- Final State Interactions. The FSI model used for this analysis has no associated GENIE uncertainties. Instead, a set of custom weights was developed using different GENIE variations where the probability of particles reinteracting within the nucleus was altered. Comparisons of this dataset with those from a similar model used by T2K resulted in a set of four FSI uncertainties: one for the pion mean free path (30%), and three which treat correlated variations in pion absorption (up to 40%), pion charge exchange (up to 20%), and pion quasielastic

scattering (up to 30%) cross sections.

7.2.4 Beam Flux

A number of uncertainties are considered to address the production and propagation of hadrons through the NuMI target and decay pipe. Events are reweighted based on the effect of the uncertainty measured by changes to the simulation.

The production of hadrons is tuned by the PPFX, which also informs the associated systematic uncertainties.

Uncertainties exist with the hadron propagation through the NuMI horn, target, and beam pipe. These include the horn (± 3 mm transverse to the beam) and target (± 7 mm parallel to the beam) position, the horn current ($\pm 1\%$), the position of the beam on the target (± 1 mm), and the size of the beam spot (± 0.2 mm) [123]. Since the NOvA detectors are off-axis the effects of these systematics are greatly reduced and are well correlated between detectors.

7.2.5 Neutron Uncertainty

The detector response to fast neutrons, roughly 1 MeV or more of kinetic energy, is not well understood using existing models [124]. The modelling of neutrons is important for the $\bar{\nu}_\mu$ analysis where the hadronic component consists primarily of neutron activity. An excess of events in low energy prongs in MC compared to the data is observed in a neutron-enriched sample of prongs from $\bar{\nu}_\mu$ events shown in Fig. 7.2. To compensate for this, energy is omitted from the hadronic component of the energy by randomly removing 47% of all truth level neutron prongs with energy less than 20 MeV from the energy calculation. The fraction was tuned to remove the observed discrepancy.

7.2.6 Lepton Reconstruction

The reconstruction of muon tracks is an important component of the analysis. The track length is the main component of ν_μ energy estimation and the track angle is used to measure the neutrino




Figure 7.2: Neutron Prongs. The total energy of prongs from a neutron enriched sample. Prongs were selected from $\bar{\nu}_\mu$ events with at least 20 cm between the vertex and the prong start, fewer than 6 hits, and photon scores from the particle classifier less than 0.8. A clear excess in the MC is seen in the ratio plot at the lowest energies.

$|p_T|$ used in the signal extrapolation. However, the track reconstruction depends on the properties of the medium through which it is traversing. Consideration is given to muons from ν_μ interactions in the FD and ND as well as separate consideration of muons that enter the muon catcher at the ND.

The systematic uncertainty on the track length is divided into five parts [125]. Three are uncorrelated errors between the detectors which are treated independently. These are 0.15% in the FD, 0.13% in the ND, and 0.48% in the muon catcher. The uncorrelated errors are primarily driven by uncertainties in the mass accounting of the detector. The fourth is a correlated 0.74% error which combines uncertainties from the Fermi density effect and fundamental uncertainties in the GEANT simulation.

The final part is an uncertainty addressing long-staying neutrons in the ND. It has been observed that over the course of the beam spill, slow neutrons from early in the beam spill produce excess hits near the end of the beam spill. These neutrons are not accurately simulated in the MC [126]. Reconstructed muon track lengths near the end of the beam spill were found to be 0.46 cm longer

in the main detector and 1.3 cm longer in the muon catcher.

7.2.7 Near-Far Uncorrelated

This category covers several uncertainties which are explicitly uncorrelated between the ND and the FD.

The first covers the acceptance of ν_e signal in the FD. The ν_e signal is extrapolated by the ν_μ in the ND which causes a number of differences in kinematic distributions due to detector sizes and the difference in CVN performance between ν_μ and ν_e . The known differences are extrapolated by the $|p_T|$ quantile extrapolation, but the unknown differences are covered by an explicit systematic. The ND ν_μ data and MC are reweighted to match the ν_e FD MC kinematics in Q^2 , $\cos \theta_{\text{beam}}$, and $|p_T|$. The reweighted ND distributions are extrapolated to the FD and the largest difference between the three new predictions and the nominal prediction is taken as a systematic. The acceptance is measured separately for FHC and RHC.

In the ND, the intensity of the NuMI beam influences the reconstruction and selection efficiency of neutrino interactions. At high intensities, more neutrino interactions occur and it's possible for multiple interactions to be reconstructed into a single slice. To evaluate this effect, a single neutrino interaction is overlaid with real data or MC events and the selection efficiency for that neutrino is compared between the two samples. The difference in selection efficiency is treated as a normalization systematic in the ND.

The ND FHC ν_e decomposition is reliant on the efficiency of clustering and selecting Michel electrons from the ND data. A shift is applied which varies the selection efficiency by $\pm 5\%$. Rather than altering the reconstruction, 5% of candidates are randomly removed or samples that weren't selected are artificially included in the MC sample.

In the FD, there is a small rate of neutrino interactions in the surrounding rock which enter the detector and are selected. Roughly 0.5% (0.6%) of events in ν_e and 0.01% (0.05%) of events

in ν_μ for FHC (RHC) are expected to come from interactions in the rock. Most of the ν_e rock events are selected in the peripheral sample. These events are treated as background since they are not contained within the detector fiducial volume. Rock events are difficult to constrain and measure, particularly for the ν_e peripheral sample, so a 100% uncertainty is applied to the number of predicted rock events.

The cosmic prediction is measured directly from the data, so there is no systematic error associated with it. However, the statistical error on the number of cosmic events does influence the analysis. Within each analysis bin, the number of cosmic events is varied by $\pm 1\sigma$ from the upper (lower) bound of the 68.2% interquartile range with mean equal to the original cosmic prediction.

The FD normalization is varied by $\pm 0.82\%$ due to the uncertainty in the FD mass and a not yet understood failure mode in the detector simulation reducing the number of events. There is also a 0.55% uncertainty in the total recorded POT for both FHC and RHC.

7.3 Uncertainties by Principal Component Decomposition

Each systematic uncertainty is included along with the oscillation parameters when performing fits. The more systematics that need to be considered, the more parameters used in the fit and the longer the fit takes to converge. In addition, if there are correlations between different systematics, they can no longer be treated as one-dimensional gaussian errors. This is particularly important for the hadron propagation beam systematics from PPFX and many of the nearly 80 cross section systematics whose individual effects on the predictions are small, but collectively need to be included in the analysis.

To avoid computational costs, previous oscillation analyses either ignored these smaller systematics, which clearly underestimates their contribution, or summed their effects in quadrature into a single systematic, which overestimates the uncertainty due to possible correlations. A method of dimensionality reduction and decorrelation is needed to properly use these systematics without

significantly increasing the computing resources needed for fitting.

One method of dimensionality reduction is Principal Component Analysis (PCA) [127]. PCA is a linear transformation of a dataset to a new coordinate system of the same dimension. In this new coordinate system, the greatest variance, by projection, in the data lies along the first axis. The second axis is chosen orthogonal to the first with the next largest component of the remaining variance. This repeats until the entire dimensionality of the dataset is covered. The axes are known as the Principal Components (PC).

Often, the first few PCs account for a majority of the total variation in the data. Using this knowledge, the first few principal components alone can be used and the remaining ignored. In this way, PCA can be used to greatly reduce the dimensionality of a dataset while still capturing its largest features. In addition since the PCs are orthogonal, the variance along each PC is decorrelated from all other PCs.

PCA is typically carried out by diagonalizing the covariance matrix of a dataset. Given a vector of random variables \bar{X} , the covariance matrix has elements given by the covariance of two variables:

$$\Sigma_{ij} = \text{cov}(X_i, X_j) = E(X_i X_j) - E(X_i)E(X_j) \quad (7.4)$$

where Σ is the covariance matrix, X_i is one of the random variables in the vector \mathbf{X} , and $E(X)$ denotes the expected value. If the two variables are uncorrelated, the covariance is zero. With this construction, the covariance matrix is a symmetric matrix and, thus, the eigenvalues are real and the eigenvectors are orthogonal.

If these eigenvectors are normalized, the result is a basis of orthonormal vectors. In addition, the eigenvectors can be sorted by decreasing eigenvalue. These sorted pairs of eigenvectors and

eigenvalues are now the PCs of the data with explained variance given by

$$EV_i = \frac{\lambda_i}{\sum_{j=0}^N \lambda_j} \quad (7.5)$$

where EV_i is the fraction of the total variance of the PC and λ_i is the eigenvalue of the PC. PC shifts are then given by $\sqrt{\lambda_i}v_i$, where v_i is one of the orthonormal eigenvectors. The behavior of the entire covariance matrix can be recreated by the PC shifts.

With regards to NOvA systematics, these constructed PCs would be uncorrelated shifts with an effect proportional to the eigenvalue of the PC. Each PC can now be used as a parameter in the oscillation fit.

Two separate sets of PCA are used. The first is for all the beam flux systematics, the hadron production and propagation uncertainties in the beam. The second is for the cross section systematics. However, which cross section systematics are the largest is of interest for future cross section tuning, so only the uncertainties with the smallest effect on the nominal prediction are used in the PCA.

For the beam PCA, the data is separated into 16 samples of $\{\nu_\mu, \nu_e\} \times \{\nu, \bar{\nu}\} \times \{\text{ND, FD}\} \times \{\text{FHC, RHC}\}$. Each sample is divided into 40 bins from 0 GeV to 10 GeV for a total of 640 bins. Since the interest in the systematics is the effect on the F/N ratios used in the extrapolation, the FD samples are divided by the ND samples. The ND samples are kept to maintain a complete basis, but suppressed by a factor of 10^{-6} so the PCA converges to the effect on the F/N ratios. The PPFX hadron production uncertainties are combined with the hadron propagation uncertainties varied by a random amount to create a total of 2000 beam flux “universes”. The variation within each bin caused by all the universes defines one random variable which can be used to construct the covariance matrix shown in Fig. 7.3. The off-diagonal elements of the matrix show clear covariance between many bins as a result of correlated effects among the beam flux systematics. Diagonalization of the covariance matrix results in a set of eigenvalues and orthonormal eigenvectors. Since

there are 640 bins, there are 640 eigenvalues which are also shown in Fig. 7.3. The ND is rescaled to its original weight and the F/N is multiplied back to just the FD values. The PCs now define systematic shifts which can be applied to each neutrino based on its true sign, flavor, and energy in each detector.

The set of “small” cross section systematics to incorporate in the PCA are chosen based on the χ^2 between the $\pm 1\sigma$ shifted prediction and the nominal prediction. This results in about 20 “large” cross section systematics which are used directly in the oscillation fit and the remaining are small systematics. A GENIE universe is constructed by varying each of these knobs by a random amount chosen from a normal distribution, $\mathcal{N}(0, 1)$. In total, 1000 unique universes are created to be used to construct the covariance matrix. The data is separated into 28 samples in 20 bins of true energy from 0 GeV to 10 GeV for a total of 560 bins. The list of samples with the ν_μ selection are:

- Signal ν_μ QE
- Other signal ν_μ
- Background NC

and list of samples with the ν_e selection are:

- Signal ν_e QE
- Other signal ν_e
- Background ν_μ
- Background NC

and for each of $\{\text{FHC}, \text{RHC}\} \times \{\text{FD}, \text{ND}\}$. The covariance matrix is constructed and diagonalized identically to the beam flux PCA, including the suppression of the ND and using the F/N values. The first PC produced is shown in Fig. 7.4 which defines a new systematic to be used in the oscillation fit.

All the PCs cover the entire space of the universes, but not all PCs are needed as seen in Fig. 7.3. For the GENIE PCA, only 12 PCs were needed to cover 95% of the RMS for 97% of analysis bins.




Figure 7.3: Beam Flux PCA. Top: The covariance between each F/N bin used in 2000 universes for the beam flux PCA. The ND bins, not shown, are suppressed by a factor of 10^{-6} . Bottom: The eigenvalues sorted by decreasing value resulting from diagonalizing the covariance matrix, shown on a log-scale.




Figure 7.4: Cross Section PC 0. The first principle component from the cross section PCA. The shift is the weight which is applied to the nominal spectrum based on its sample and energy bin. Here, N_α (F_α) corresponds to the ν_α selection in the ND (FD). The top half shows shifts for FHC and the bottom for RHC. The $\pm 1\sigma$ shift is shown in red (blue). This PC has an explained variance of 45% and mainly effects the NC backgrounds in both the ν_e and ν_μ selections.

To cover the remaining variance, the GENIE PCs are scaled up by 35%. Similarly, for the beam PCA 5 PCs scaled up by 25% are needed. The 12 GENIE PCs and 5 beam PCs are incorporated in the oscillation fit along with all other systematics in the previous section.

7.4 Uncertainty Interpolation

As mentioned in Sec. 7.1, each shifted prediction is only evaluated for particular shifts. In some cases more could be evaluated, but in most, particularly the calibration systematics, it would be computationally infeasible to create so many new simulation files. So, within each bin, the ratio of the shifted predictions to the nominal are fit using a cubic polynomial. The fit is usually to




Figure 7.5: Systematic Interpolation. An examples of the systematic interpolation function fit for two ν_e analysis bins by an example cross section systematic. The X-axis is a given systematic shift measured in the uncertainty σ of that systematic. The Y-axis shows ratios of the shifted prediction to the nominal prediction. The black points are evaluated and then fit by the interpolation function to determine the effect of intermediate shifts.

five points: $\{-2\sigma, -1\sigma, 0\sigma, +1\sigma, +2\sigma\}$ ¹, where 0σ is just the nominal prediction. In cases where the shift is one sided, such as the calibration drift or Cherenkov uncertainties, the -2σ and -1σ shifts are fixed to the nominal. The fitted function can then be used to interpolate the effect the systematic uncertainty has on the prediction beyond just the evaluated shifts. An example for one systematic is shown in Fig 7.5.

7.5 Fitting with Systematics

Systematics are included as nuisance parameters in the fit. Equation 6.11 is expanded to:

$$-2\mathcal{L}(\bar{\theta}, \bar{\delta}) = -2 \sum_{i=1}^N \left[E_i(\bar{\theta}, \bar{\delta}) - O_i + O_i \ln \frac{O_i}{E_i(\bar{\theta}, \bar{\delta})} \right] + \sum_{j=1}^S \frac{\delta_j^2}{\sigma_j^2} \quad (7.6)$$

where δ_j is the nuisance of systematic j and σ_j is its uncertainty. The predicted number of events in each bin E_i now also depends on the systematic shifts which are fit together with all of the oscillation parameters. The set of systematic shifts which minimizes the log-likelihood are referred to as the systematic “pulls”.

¹Five points clearly overconstrains the fit of a cubic. Instead, a cubic polynomial is fit to each rolling set of four points and then matched smoothly between them. The resulting function is continuous and twice differentiable by construction.




Figure 7.6: Effects of Systematics. The increase in quadrature of the 68% confidence interval by including a set of systematic uncertainties into the oscillation fit. Shown for θ_{23} (top left), Δm_{32}^2 (top right), and δ_{CP} (bottom). The final two rows of each show the uncertainty from all systematics and the statistical uncertainty in the measurement. The systematic groups are ordered by decreasing effect for each measurement.

The effect of each systematic on the oscillation measurements can be determined by fitting the prediction to fake data. The fit is done by profiling over a chosen oscillation parameter, all other parameters are free, once without the systematic and then again with the systematic. The difference in quadrature in the 68% confidence interval can determine the effect of that systematic on that measurement. The fake data is generated to match the nominal prediction at the chosen set of oscillation parameters. As a result, the systematic pull terms $\bar{\delta}$ will all be 0 at the best fit point. Away from the best fit point, the systematic pulls are fit along with the oscillation parameters to match the fake data as closely as possible. The results are shown in Fig. 7.6.

Figure 7.6 shows that NOvA's measurements are currently statistically limited, particularly for

the determination of δ_{CP} where selecting as many ν_e as possible is critical for resolving CP violation. As more data are collected, the measurements of Δm_{32}^2 and θ_{23} will become systematically limited and future improvements will incorporate better understanding and treatment of the systematic uncertainties.

Chapter 8

Experimental Results and Discussion

With the total prediction of events for every analysis bin and the effect of every systematic on those predictions, the FD data can be unveiled. This chapter presents the results of fitting the data from 13.6×10^{20} POT in FHC beam mode and 12.5×10^{20} POT in RHC beam mode to the three flavor neutrino oscillation model. The results of this analysis were first presented in July 2020 [128].

8.1 Results

The resulting ν_e energy spectra with data are shown in Fig. 8.1 and the breakdown by analysis component is shown in Table 8.1. In total 82 (33) candidate ν_e ($\bar{\nu}_e$) events are selected with an expected background of 26.83 (13.98) in FHC (RHC). The number of events observed in RHC and corresponding uncertainties represents more than 4σ evidence of $\bar{\nu}_e$ appearance [103], the most significant observation of electron antineutrino appearance ever.

The total expected FHC events ranges from 52 in the IH, LO, and $\delta_{\text{CP}} = \pi/2$ to 110 in the NH, UO, and $\delta_{\text{CP}} = 3\pi/2$. The total expected RHC events ranges from 25 in the NH, LO, and $\delta_{\text{CP}} = 3\pi/2$ to 45 in the IH, UO, and $\delta_{\text{CP}} = \pi/2$. These ranges with the observed data are shown in Fig. 8.2. The structure of Fig. 8.2 is akin to Fig. 2.6, but the probabilities have been replaced with an expected event count. As seen in the figure, the observed result is in the degenerate region between the two hierarchies.

The resulting ν_μ energy spectra with data are shown in Fig. 8.3 and the breakdown by component is shown in Table 8.2. In total 211 (105) candidate ν_μ ($\bar{\nu}_\mu$) events are selected with an expected background of 8.2 (2.1) in FHC (RHC). The expectation with no oscillations is 1156.1 (488.1) in FHC (RHC) showing unambiguous evidence of neutrino oscillations. The ratio of the data to no oscillations is shown in Fig. 8.4. Note the region of maximum oscillations which provides




Figure 8.1: ν_e Data Spectra. The ν_e prediction (purple) and observed data (black points) in FHC (left) and RHC (right). The prediction is made using the best fit to the data. The shaded histograms represent the 1σ systematic uncertainties. The component breakdown is in Table 8.1.

Component	FHC	RHC
WS	1.04	2.25
Beam $\nu_e/\bar{\nu}_e$	14.08	7.29
NC	6.21	2.21
$\nu_\mu/\bar{\nu}_\mu$	1.73	0.36
$\nu_\tau/\bar{\nu}_\tau$	0.52	0.32
Beam Background	23.69	12.43
Cosmic Background	3.13	1.55
Total Background	$26.83^{+1.46}_{-1.62}$	$13.98^{+0.76}_{-0.84}$
Total Signal	$58.96^{+2.50}_{-2.54}$	$19.20^{+0.64}_{-0.66}$
Total Expected	$85.79^{+2.96}_{-3.05}$	$33.18^{+1.08}_{-1.23}$
Total Observed	82	33

Table 8.1: The expected number of ν_e events in FHC and RHC. The error bars are the 1σ systematic uncertainty. The final row is the observed number of candidate data events.

the most sensitivity to θ_{23} and Δm_{32}^2 as discussed in Chap. 2.

8.2 Neutrino Oscillation Parameter Constraints

We vary the neutrino oscillations parameters and systematics pull terms until the predicted number of events give a best fit to the observed data, yielding measurements of the oscillation parameters.

In the fits, Δm_{32}^2 , θ_{23} , and δ_{CP} are unconstrained and are the main results which will be discussed.

The remaining parameters are constrained based on measurements from other experiments [28].

$\Delta m_{12}^2 = 7.53 \times 10^{-5} \text{ eV}^2$ and $\sin^2 2\theta_{12} = 0.851$ are fixed in the fits. The final value $\sin^2 2\theta_{13} =$

NOvA Preliminary




Figure 8.2: ν_e Event Range. The range of possible events for FHC (X-axis) and RHC (Y-axis) under different mass hierarchy, θ_{23} octant, and δ_{CP} values. The black point shows the observed data.

0.085 ± 0.003 is allowed to float within the uncertainty.

Fitting the data in the four ν_μ quartiles and the three ν_e PID ranges in both FHC and RHC in Figs. 8.1 and 8.3 using the log-likelihood in Eq. 7.6 results in the following values:

- $\Delta m_{32}^2 = (2.41 \pm 0.07) \times 10^{-3} \text{ eV}^2$
- $\sin^2 \theta_{23} = 0.57^{+0.04}_{-0.03}$
- $\delta_{CP} = 0.82^{+0.27}_{-0.87} \pi$

where Δm_{32}^2 is in the NH with a preference of 1.0σ and θ_{23} is in the UO with a preference of 1.2σ .

Figure 8.5 shows the confidence intervals for these measurements.

As discussed in Chap. 7, the systematic uncertainties are incorporated into the fit as nuisance parameters which pull on the prediction. The largest pulls in this result come from the ND light




Figure 8.3: ν_μ Data Spectra. The ν_μ prediction (purple) and observed data (black points) in FHC (top) and RHC (bottom) for each of the four f_{had} quartiles. The prediction is made using the best fit to the data. The shaded histograms represent the 1σ systematic uncertainties. The component breakdown is in Table 8.2.

Component	FHC	RHC
Beam Background	3.25	1.12
Cosmic Background	4.96	0.95
Total Background	8.21	2.07
WS	12.64	26.10
Total Signal	214.1	103.4
Total Expected	$222.3^{+23.5}_{-22.4}$	$105.4^{+10.6}_{-10.1}$
Total Observed	211	105

Table 8.2: The expected number of ν_e events in FHC and RHC. The error bars are the 1σ systematic uncertainty. The final row is the observed number of candidate data events.




Figure 8.4: ν_μ Oscillation Ratio. The ratio of the predicted number of events (purple) and observed data to the expectation with no neutrino oscillations. The FHC (RHC) ν_μ ($\bar{\nu}_\mu$) ratio is shown on the left (right). The shaded histogram shows the 1σ systematic uncertainties.

level uncertainty $+1.5\sigma$, the MEC nucleon fraction $+0.7\sigma$, and the absolute calibration -0.5σ . The remaining systematic pulls are all less than 0.5σ .

As seen in Fig. 8.2, this measurement is in the degenerate region with weak preferences for the hierarchy and octant. However, strong constraints can be set for the outlier regions. The region in the NH with $\delta_{CP} = 3\pi/2$ is rejected at 2σ and the region in the IH with $\delta_{CP} = \pi/2$ is rejected at more than 3σ .




Figure 8.5: Oscillation Fit Contours. Two dimensional contours of oscillation fit to the observed data under the assumption of the NH (top) and IH (bottom). The black star represents the best fit point for the data. The color shades represent increasing significances. Left: $\sin^2 \theta_{23}$ (X-axis) compared to Δm_{32}^2 (Y-axis). Right: δ_{CP} (X-axis) compared to $\sin^2 \theta_{23}$ (Y-axis).

8.3 Comparison to Other Experiments

The value of the atmospheric parameters, θ_{23} and Δm_{32}^2 , are measured by several accelerator and atmospheric neutrino experiments observing ν_μ disappearance. Comparison of 90% confidence limits of NOvA and other experiments are shown in Fig. 8.6. As seen in the figure, the best fit point by NOvA is within the limits of all five shown experiments.

While many experiments have measured the atmospheric parameters, very few have sensitivity to δ_{CP} and the octant of θ_{23} . The most similar experiment to NOvA is the T2K [34] experiment which has a similar L/E to NOvA, but at a shorter baseline ($L = 295$ km, $E \approx 0.7$ GeV) and measures ν_μ disappearance and ν_e appearance. The shorter baseline and lower energy give reduced sensitivity to the mass hierarchy, but increased sensitivity to CP violation. The 90% confidence




Figure 8.6: Atmospheric Measurements. 90% confidence limits of the atmospheric oscillation parameters set by NOvA and other experiments. The black point is the best fit oscillation point by NOvA.

limits for NOvA and T2K are shown in Fig. 8.7 based on their most recent results [129]¹. Both results prefer the NH and UO, but disagree on the value of δ_{CP} . The best fits disagree outside the 90% confidence limits.

NOvA and T2K have plans to execute a joint fit of data from both experiments [130] which will break degeneracies within each individual experiment. The key challenge is this joint fit is correlating systematic uncertainties, particularly neutrino interaction models, in the fit between the two experiments.

8.4 The Future of NOvA

The NuMI beam will continue operating until the main injector is turned off for upgrades around 2025, at which time a new beam line will be installed and there will be an increase to the beam power for the future DUNE experiment [131]. Some of the future beam upgrades may be applicable

¹T2K uses a different convention for the range of values of δ_{CP} . NOvA uses $[0, 2\pi)$ and T2K uses $[-\pi, \pi)$. The figure here uses the NOvA convention.




Figure 8.7: Comparison to T2K. The 68% and 90% confidence limits set by NOvA (blue regions) and T2K (black lines) for δ_{CP} and $\sin^2 \theta_{23}$. The best fit point for NOvA (T2K) is indicated by the star (square).

to NOvA allowing the beam power to increase to up to 900 kW [132]. Before the main injector is turned off, NOvA will record 50% of the remaining data in FHC beam mode and 50% in RHC beam mode, totalling about 31.5×10^{20} POT FHC and 31.5×10^{20} POT RHC. The expected final sensitivities for measuring CP violation and resolving the mass hierarchy are shown in Fig. 8.8 based on the current measurement. NOvA has the ability to resolve the hierarchy at more than 3σ if the parameters are favorable.

Figure 8.8 assumes current analysis techniques remain unchanged. Possible improvements have been proposed to nearly every element of the analysis, many incorporating ML techniques to enhance the physics capabilities. The ML algorithms described in this dissertation were primarily based on classification tasks, but they are also capable of estimating continuous, physical quantities. A CNN energy estimator for ν_e has been developed which shows a significant improvement in energy resolution over the current method [133]. A new energy estimator for ν_μ events has also been made which incorporates recurrent neural networks to analyze the reconstructed prongs and




Figure 8.8: Future Sensitivities. The significance for NOvA to measure CP violation (left) or resolve the mass hierarchy (right) using projected data to the end of NOvA. The significance is shown for assumptions of the NH (blue) and IH (red) as well as for the range of δ_{CP} values.

shows a 20% improvement in effective POT for the atmospheric parameter measurements [134]. Further enhancements would incorporate the improved prong reconstruction by the instance segmentation method for energy estimation, cross section measurements, and reduction in systematic uncertainties.

Figure 7.6 shows that NOvA’s current measurements are statistically limited, but as more data is collected, the importance of improving our understanding of systematic uncertainties increases. NOvA is currently operating a small detector in a test beam at Fermilab [135]. The test beam program will provide a library of tagged particles to improve understanding of the detector response and calibration as well as validate reconstruction methods. The data are currently being analyzed and incorporated into the detector calibration and response systematics. Future improvements for the implementation of the neutron systematic, the acceptance systematics, and several others are under investigation as well [136].

8.5 Conclusions

This dissertation has described the analysis of neutrino data collected by the NOvA experiment and the resulting constraints set on the neutrino oscillation parameters. Neutrino oscillations are in a

precision era with the mixing angles and mass splittings all having been measured. The remaining questions are the neutrino mass hierarchy, the octant of θ_{23} , and if there is CP violation in neutrino oscillations. NOvA is making contributions to all of these questions.

Key to NOvA's measurements are the reduction of uncertainties in the analysis. This includes statistical uncertainties by improving our neutrino event selection and systematic uncertainties by improving our understanding of neutrino models and the detector. NOvA has incorporated ML techniques into the oscillation analysis for event classification, particle identification, and more to accomplish these goals. This thesis contributed to the development of the ML tools used in the analysis, improved the treatment of systematic uncertainties, and bettered NOvA's understanding of how systematic uncertainties affect the results.

Bibliography

- [1] Wolfgang Pauli. “Letter to the physical society of Tubingen”. In: (1930).
- [2] Enrico Fermi. “Tentativo di una teoria dell’emissione dei raggi beta”. In: *Ric. Sci.* 4 (1933), pp. 491–495.
- [3] C. L. Cowan et al. “Detection of the free neutrino: A Confirmation”. In: *Science* 124 (1956), pp. 103–104. DOI: 10.1126/science.124.3212.103.
- [4] S. Schael et al. “Precision electroweak measurements on the Z resonance”. In: *Phys. Rept.* 427 (2006), pp. 257–454. DOI: 10.1016/j.physrep.2005.12.006. arXiv: hep-ex/0509008 [hep-ex].
- [5] G. Danby et al. “Observation of High-Energy Neutrino Reactions and the Existence of Two Kinds of Neutrinos”. In: *Phys. Rev. Lett.* 9 (1962), pp. 36–44. DOI: 10.1103/PhysRevLett.9.36.
- [6] K. Kodama et al. “Observation of tau neutrino interactions”. In: *Phys. Lett.* B504 (2001), pp. 218–224. DOI: 10.1016/S0370-2693(01)00307-0. arXiv: hep-ex/0012035 [hep-ex].
- [7] Y. Fukuda et al. “Evidence for oscillation of atmospheric neutrinos”. In: *Phys. Rev. Lett.* 81 (1998), pp. 1562–1567. DOI: 10.1103/PhysRevLett.81.1562. arXiv: hep-ex/9807003 [hep-ex].
- [8] B. Aharmim et al. “Electron energy spectra, fluxes, and day-night asymmetries of B-8 solar neutrinos from measurements with NaCl dissolved in the heavy-water detector at the Sudbury Neutrino Observatory”. In: *Phys. Rev. C* 72 (2005), p. 055502. DOI: 10.1103/PhysRevC.72.055502. arXiv: nucl-ex/0502021 [nucl-ex].

- [9] Donald A. Glaser and David C. Rahm. “Characteristics of Bubble Chambers”. In: *Phys. Rev.* 97 (2 Jan. 1955), pp. 474–479. DOI: 10.1103/PhysRev.97.474. URL: <https://link.aps.org/doi/10.1103/PhysRev.97.474>.
- [10] Hiroshi Nunokawa, Stephen J. Parke, and Jose W.F. Valle. “CP Violation and Neutrino Oscillations”. In: *Prog. Part. Nucl. Phys.* 60 (2008), pp. 338–402. DOI: 10.1016/j.pnpnp.2007.10.001. arXiv: 0710.0554 [hep-ph].
- [11] Georges Aad et al. “Observation of a new particle in the search for the Standard Model Higgs boson with the ATLAS detector at the LHC”. In: *Phys. Lett. B* 716 (2012), pp. 1–29. DOI: 10.1016/j.physletb.2012.08.020. arXiv: 1207.7214 [hep-ex].
- [12] Eva Halkiadakis, George Redlinger, and David Shih. “Status and Implications of Beyond-the-Standard-Model Searches at the LHC”. In: *Ann. Rev. Nucl. Part. Sci.* 64 (2014), pp. 319–342. DOI: 10.1146/annurev-nucl-102313-025632. arXiv: 1411.1427 [hep-ex].
- [13] Mark D. Messier. “Evidence for neutrino mass from observations of atmospheric neutrinos with Super-Kamiokande”. PhD thesis. Boston U., 1999.
- [14] Nathalie Palanque-Delabrouille et al. “Neutrino masses and cosmology with Lyman-alpha forest power spectrum”. In: *JCAP* 11 (2015), p. 011. DOI: 10.1088/1475-7516/2015/11/011. arXiv: 1506.05976 [astro-ph.CO].
- [15] Susanne Mertens. “KATRIN: Recent Results and Future Perspectives”. The XXIX International Conference on Neutrino Physics and Astrophysics. 2020.
- [16] Goran Senjanovic. “See-saw and grand unification”. In: *SEESAW25: International Conference on the Seesaw Mechanism and the Neutrino Mass*. Jan. 2005, pp. 45–64. DOI: 10.1142/9789812702210_0004. arXiv: hep-ph/0501244.
- [17] Susanne Mertens. “Direct Neutrino Mass Experiments”. In: *J. Phys. Conf. Ser.* 718.2 (2016), p. 022013. DOI: 10.1088/1742-6596/718/2/022013. arXiv: 1605.01579 [nucl-ex].

- [18] M. Sajjad Athar, Shakeb Ahmad, and S.K. Singh. “Neutrino nucleus cross sections for low energy neutrinos at SNS facilities”. In: *Nucl. Phys. A* 764 (2006), pp. 551–568. DOI: 10.1016/j.nuclphysa.2005.09.017. arXiv: nucl-th/0506046.
- [19] M.G. Aartsen et al. “Measurement of the multi-TeV neutrino cross section with IceCube using Earth absorption”. In: *Nature* 551 (2017), pp. 596–600. DOI: 10.1038/nature24459. arXiv: 1711.08119 [hep-ex].
- [20] C.H. Llewellyn Smith. “Neutrino Reactions at Accelerator Energies”. In: *Phys. Rept.* 3 (1972), pp. 261–379. DOI: 10.1016/0370-1573(72)90010-5.
- [21] Dieter Rein and Lalit M. Sehgal. “Neutrino Excitation of Baryon Resonances and Single Pion Production”. In: *Annals Phys.* 133 (1981), pp. 79–153. DOI: 10.1016/0003-4916(81)90242-6.
- [22] A. Bodek and U.K. Yang. “Modeling neutrino and electron scattering inelastic cross-sections in the few GeV region with effective LO PDFs TV Leading Order”. In: *2nd International Workshop on Neutrino-Nucleus Interactions in the Few GeV Region*. Aug. 2003. arXiv: hep-ex/0308007.
- [23] Juan Nieves et al. “2p2h Excitations, MEC, Nucleon Correlations and Other Sources of QE-like Events”. In: *JPS Conf. Proc.* 12 (2016), p. 010002. DOI: 10.7566/JPSCP.12.010002.
- [24] Teppei Katori. “Meson Exchange Current (MEC) Models in Neutrino Interaction Generators”. In: *AIP Conf. Proc.* 1663.1 (2015). Ed. by H. Da Motta, Jorge G. Morfin, and M. Sakuda, p. 030001. DOI: 10.1063/1.4919465. arXiv: 1304.6014 [nucl-th].
- [25] Teppei Katori and Marco Martini. “Neutrino–Nucleus Cross Sections for Oscillation Experiments”. In: *J. Phys. G* 45.1 (2018), p. 013001. DOI: 10.1088/1361-6471/aa8bf7. arXiv: 1611.07770 [hep-ph].

- [26] B. Pontecorvo. “Inverse beta processes and nonconservation of lepton charge”. In: *Sov. Phys. JETP* 7 (1958), pp. 172–173.
- [27] Ziro Maki, Masami Nakagawa, and Shoichi Sakata. “Remarks on the unified model of elementary particles”. In: *Prog. Theor. Phys.* 28 (1962), pp. 870–880. DOI: 10.1143/PTP.28.870.
- [28] M. Tanabashi et al. “Review of Particle Physics”. In: *Phys. Rev. D* 98.3 (2018), p. 030001. DOI: 10.1103/PhysRevD.98.030001.
- [29] K. Abe et al. “Constraint on the matter–antimatter symmetry-violating phase in neutrino oscillations”. In: *Nature* 580.7803 (2020). [Erratum: *Nature* 583, E16 (2020)], pp. 339–344. DOI: 10.1038/s41586-020-2177-0. arXiv: 1910.03887 [hep-ex].
- [30] Evgeny K. Akhmedov et al. “Series expansions for three flavor neutrino oscillation probabilities in matter”. In: *JHEP* 04 (2004), p. 078. DOI: 10.1088/1126-6708/2004/04/078. arXiv: hep-ph/0402175.
- [31] Y. Fukuda et al. “The Super-Kamiokande detector”. In: *Nucl. Instrum. Meth. A* 501 (2003). Ed. by V.A. Ilyin, V.V. Korenkov, and D. Perret-Gallix, pp. 418–462. DOI: 10.1016/S0168-9002(03)00425-X.
- [32] Geoffrey B. Mills. “Neutrino oscillation results from LSND”. In: *Nucl. Phys. B Proc. Suppl.* 91 (2001). Ed. by J. Law, R.W. Ollerhead, and J.J. Simpson, pp. 198–202. DOI: 10.1016/S0920-5632(00)00941-5.
- [33] Pedro AN Machado, Ornella Palamara, and David W Schmitz. “The Short-Baseline Neutrino Program at Fermilab”. In: *Ann. Rev. Nucl. Part. Sci.* 69 (2019), pp. 363–387. DOI: 10.1146/annurev-nucl-101917-020949. arXiv: 1903.04608 [hep-ex].
- [34] K. Abe et al. “The T2K experiment”. In: *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment* 659.1

- (2011), pp. 106–135. ISSN: 0168-9002. DOI: <https://doi.org/10.1016/j.nima.2011.06.067>. URL: <http://www.sciencedirect.com/science/article/pii/S0168900211011910>.
- [35] P. Adamson et al. “Improved search for muon-neutrino to electron-neutrino oscillations in MINOS”. In: *Phys. Rev. Lett.* 107 (2011), p. 181802. DOI: 10.1103/PhysRevLett.107.181802. arXiv: 1108.0015 [hep-ex].
- [36] B. Abi et al. “Long-baseline neutrino oscillation physics potential of the DUNE experiment”. In: *Eur. Phys. J. C* 80.10 (2020), p. 978. DOI: 10.1140/epjc/s10052-020-08456-z. arXiv: 2006.16043 [hep-ex].
- [37] Y. Abe et al. “Indication of Reactor $\bar{\nu}_e$ Disappearance in the Double Chooz Experiment”. In: *Phys. Rev. Lett.* 108 (2012), p. 131801. DOI: 10.1103/PhysRevLett.108.131801. arXiv: 1112.6353 [hep-ex].
- [38] F.P. An et al. “Observation of electron-antineutrino disappearance at Daya Bay”. In: *Phys. Rev. Lett.* 108 (2012), p. 171803. DOI: 10.1103/PhysRevLett.108.171803. arXiv: 1203.1669 [hep-ex].
- [39] J.K. Ahn et al. “Observation of Reactor Electron Antineutrino Disappearance in the RENO Experiment”. In: *Phys. Rev. Lett.* 108 (2012), p. 191802. DOI: 10.1103/PhysRevLett.108.191802. arXiv: 1204.0626 [hep-ex].
- [40] K. S. Hirata et al. “Observation of B-8 Solar Neutrinos in the Kamiokande-II Detector”. In: *Phys. Rev. Lett.* 63 (1989), p. 16. DOI: 10.1103/PhysRevLett.63.16.
- [41] Raymond Davis Jr., Don S. Harmer, and Kenneth C. Hoffman. “Search for neutrinos from the sun”. In: *Phys. Rev. Lett.* 20 (1968), pp. 1205–1209. DOI: 10.1103/PhysRevLett.20.1205.
- [42] P. Anselmann et al. “Solar neutrinos observed by GALLEX at Gran Sasso.” In: *Phys. Lett.* B285 (1992), pp. 376–389. DOI: 10.1016/0370-2693(92)91521-A.

- [43] D.S. Ayres et al. “The NOvA Technical Design Report”. In: (Oct. 2007). DOI: 10.2172/935497.
- [44] I. Ambats et al. “The MINOS Detectors Technical Design Report”. In: (Oct. 1998).
- [45] “The Fermilab Main Injector Technical Design Handbook”. In: (Aug. 1994). DOI: 10.2172/1127909.
- [46] P. Adamson et al. “The NuMI Neutrino Beam”. In: *Nucl. Instrum. Meth. A* 806 (2016), pp. 279–306. DOI: 10.1016/j.nima.2015.08.063. arXiv: 1507.06690 [physics.acc-ph].
- [47] S. Mufson et al. “Liquid scintillator production for the NOvA experiment”. In: *Nucl. Instrum. Meth. A* 799 (2015), pp. 1–9. DOI: 10.1016/j.nima.2015.07.026. arXiv: 1504.04035 [physics.ins-det].
- [48] D. Groom and S. Klein. “Passage of particles through matter”. In: *European Physical Journal C - EUR PHYS J C* 15 (Mar. 2000), pp. 163–173. DOI: 10.1007/BF02683419.
- [49] Pietro Antonioli et al. “SNEWS: The Supernova Early Warning System”. In: *New J. Phys.* 6 (2004), p. 114. DOI: 10.1088/1367-2630/6/1/114. arXiv: astro-ph/0406214.
- [50] Evan Niner. *TDU/DCM Delay Calibration Tech Note*. NOvA Internal Document, DocDb 11568.
- [51] Evan David Niner. “Observation of Electron Neutrino Appearance in the NuMI Beam with the NOvA Experiment”. PhD thesis. Indiana U., 2015. DOI: 10.2172/1221353.
- [52] A. Aurisano et al. “The NOvA simulation chain”. In: *J. Phys. Conf. Ser.* 664.7 (2015), p. 072002. DOI: 10.1088/1742-6596/664/7/072002.
- [53] Giuseppe Battistoni et al. “The FLUKA code: Description and benchmarking”. In: *AIP Conf. Proc.* 896.1 (2007). Ed. by M. Albrow and R. Raja, pp. 31–49. DOI: 10.1063/1.2720455.

- [54] S. Agostinelli et al. “GEANT4—a simulation toolkit”. In: *Nucl. Instrum. Meth. A* 506 (2003), pp. 250–303. DOI: 10.1016/S0168-9002(03)01368-8.
- [55] L. Aliaga et al. “Neutrino Flux Predictions for the NuMI Beam”. In: *Phys. Rev. D* 94.9 (2016). [Addendum: *Phys.Rev.D* 95, 039903 (2017)], p. 092005. DOI: 10.1103/PhysRevD.94.092005. arXiv: 1607.00704 [hep-ex].
- [56] C. Andreopoulos et al. “The GENIE Neutrino Monte Carlo Generator”. In: *Nucl. Instrum. Meth. A* 614 (2010), pp. 87–104. DOI: 10.1016/j.nima.2009.12.009. arXiv: 0905.2517 [hep-ph].
- [57] M. A. Acero et al. “Adjusting Neutrino Interaction Models and Evaluating Uncertainties using NOvA Near Detector Data”. In: (2020). arXiv: 2006.08727 [hep-ex].
- [58] Adam Aurisano. *The NOvA Detector Simulation*. NOvA Internal Document, DocDb 13577.
- [59] Chris Backhouse and Alexander Radovic. *The Attenuation and Threshold Correction of the NOvA Detectors*. NOvA Internal Document, DocDb 13579.
- [60] Luke Vinton. *Calorimetric Energy Scale Calibration of the NOvA Detectors*. NOvA Internal Document, DocDb 13579.
- [61] Mark Messier. *Hand scanning*. NOvA Internal Document, DocDb 38313.
- [62] M. Aleksa et al. “Measurement of the ATLAS solenoid magnetic field”. In: *JINST* 3 (2008), P04003. DOI: 10.1088/1748-0221/3/04/P04003.
- [63] Martin Ester et al. “A Density-Based Algorithm for Discovering Clusters in Large Spatial Databases with Noise”. In: *Proceedings of the Second International Conference on Knowledge Discovery and Data Mining*. KDD’96. Portland, Oregon: AAAI Press, 1996, pp. 226–231.
- [64] Biswaranjan Behera, Gavin Davies, and Fernanda Psihas. “Event Reconstruction in the NOvA Experiment”. In: *Meeting of the APS Division of Particles and Fields*. 2017. arXiv: 1710.03772 [physics.ins-det].

- [65] Dan Pershey et al. *TDSlicer Technote*. NOvA Internal Document, DocDb 27689.
- [66] Alex Rodriguez and Alessandro Laio. “Clustering by fast search and find of density peaks”. In: *Science* 344.6191 (2014), pp. 1492–1496. ISSN: 0036-8075. DOI: 10.1126/science.1242072. eprint: <https://science.sciencemag.org/content/344/6191/1492.full.pdf>. URL: <https://science.sciencemag.org/content/344/6191/1492>.
- [67] R. C. Prim. “Shortest connection networks and some generalizations”. In: *The Bell System Technical Journal* 36.6 (1957), pp. 1389–1401. DOI: 10.1002/j.1538-7305.1957.tb01515.x.
- [68] Mark Messier. *Vertex Reconstruction Based on Elastic Arms*. NOvA Internal Document, DocDb 7530.
- [69] Allam Shehata Hassanein et al. *A Survey on Hough Transform, Theory, Techniques and Applications*. 2015. arXiv: 1502.02160 [cs.CV].
- [70] S. Kirkpatrick, C. D. Gelatt, and M. P. Vecchi. “Optimization by simulated annealing”. In: *SCIENCE* 220.4598 (1983), pp. 671–680.
- [71] R. Krishnapuram and J.M. Keller. “A possibilistic approach to clustering”. In: *IEEE Trans. Fuzzy Syst.* 1 (1993), pp. 98–110.
- [72] J.C. Bezdek. *Pattern Recognition with Fuzzy Objective Function Algorithms*. New York: Plenum Press, 1981.
- [73] Nicolaas H. Kuiper. “Tests concerning random points on a circle”. In: *Indagationes Mathematicae (Proceedings)* 63 (1960), pp. 38–47. ISSN: 1385-7258. DOI: [https://doi.org/10.1016/S1385-7258\(60\)50006-0](https://doi.org/10.1016/S1385-7258(60)50006-0). URL: <http://www.sciencedirect.com/science/article/pii/S1385725860500060>.
- [74] Nicholas Raddatz. *Kalman Track Technical Note*. NOvA Internal Document, DocDb 13545.
- [75] Greg Welch and Gary Bishop. *An Introduction to the Kalman Filter*. 1995.

- [76] Michael David Baird. “An Analysis of Muon Neutrino Disappearance from the NuMI Beam Using an Optimal Track Fitter”. PhD thesis. Indiana U., 2015. DOI: 10.2172/1223262.
- [77] Fernanda Psihas et al. “A Review on Machine Learning for Neutrino Experiments”. In: (Aug. 2020). arXiv: 2008.01242 [physics.comp-ph].
- [78] Jianqing Fan, Cong Ma, and Yiqiao Zhong. *A Selective Overview of Deep Learning*. 2019. arXiv: 1904.05526 [stat.ML].
- [79] Chigozie Nwankpa et al. *Activation Functions: Comparison of trends in Practice and Research for Deep Learning*. 2018. arXiv: 1811.03378 [cs.LG].
- [80] Alexander Radovic et al. “Machine learning at the energy and intensity frontiers of particle physics”. In: *Nature* 560.7716 (2018), pp. 41–48. DOI: 10.1038/s41586-018-0361-2. URL: <https://doi.org/10.1038/s41586-018-0361-2>.
- [81] Fernanda Psihas. “Measurement of Long Baseline Neutrino Oscillations and Improvements from Deep Learning”. PhD thesis. Indiana U., 2018. DOI: 10.2172/1437288.
- [82] C. Backhouse and R.B. Patterson. “Library Event Matching event classification algorithm for electron neutrino interactions in the NO ν A detectors”. In: *Nucl. Instrum. Meth. A* 778 (2015), pp. 31–39. DOI: 10.1016/j.nima.2015.01.017. arXiv: 1501.00968 [physics.ins-det].
- [83] A. Aurisano et al. “A Convolutional Neural Network Neutrino Event Classifier”. In: *JINST* 11.09 (2016), P09001. DOI: 10.1088/1748-0221/11/09/P09001. arXiv: 1604.01444 [hep-ex].
- [84] Mark Sandler et al. *MobileNetV2: Inverted Residuals and Linear Bottlenecks*. 2018. arXiv: 1801.04381 [cs.CV].
- [85] Kaiming He et al. *Deep Residual Learning for Image Recognition*. 2015. arXiv: 1512.03385 [cs.CV].
- [86] Jie Hu et al. *Squeeze-and-Excitation Networks*. 2017. arXiv: 1709.01507 [cs.CV].

- [87] Andrew G. Howard et al. “MobileNets: Efficient Convolutional Neural Networks for Mobile Vision Applications”. In: *CoRR* abs/1704.04861 (2017). arXiv: 1704.04861. URL: <http://arxiv.org/abs/1704.04861>.
- [88] Connor Shorten and Taghi M. Khoshgoftaar. “A survey on Image Data Augmentation for Deep Learning”. In: *Journal of Big Data* 6.1 (2019), p. 60. DOI: 10.1186/s40537-019-0197-0. URL: <https://doi.org/10.1186/s40537-019-0197-0>.
- [89] L.J.P. van der Maaten and G.E. Hinton. “Visualizing High-Dimensional Data Using t-SNE”. In: (2008).
- [90] Sebastian Raschka, Joshua Patterson, and Corey Nolet. *Machine Learning in Python: Main developments and technology trends in data science, machine learning, and artificial intelligence*. 2020. arXiv: 2002.04803 [cs.LG].
- [91] Martín Abadi et al. *TensorFlow: Large-Scale Machine Learning on Heterogeneous Systems*. Software available from tensorflow.org. 2015. URL: <http://tensorflow.org/>.
- [92] Francois Chollet et al. *Keras*. 2015. URL: <https://github.com/fchollet/keras>.
- [93] Yangqing Jia et al. “Caffe: Convolutional Architecture for Fast Feature Embedding”. In: *Proceedings of the 22nd ACM International Conference on Multimedia*. MM ’14. Orlando, Florida, USA: Association for Computing Machinery, 2014, pp. 675–678. ISBN: 9781450330633. DOI: 10.1145/2647868.2654889. URL: <https://doi.org/10.1145/2647868.2654889>.
- [94] F. Psihas et al. “Context-Enriched Identification of Particles with a Convolutional Network for Neutrino Events”. In: *Phys. Rev. D* 100.7 (2019), p. 073005. DOI: 10.1103/PhysRevD.100.073005. arXiv: 1906.00713 [physics.ins-det].
- [95] Joshua Porter. *ReMid retraining 2019 technote*. NOvA Internal Document, DocDb 42277.
- [96] Micah Groh and Fernanda Psihas. *Technote: Pi0 Mass Peak in RHC and FHC Data using Prong CVN*. NOvA Internal Document, DocDb 23817.

- [97] Licheng Jiao et al. “A Survey of Deep Learning-Based Object Detection”. In: *IEEE Access* 7 (2019), pp. 128837–128868. ISSN: 2169-3536. DOI: 10.1109/access.2019.2939201. URL: <http://dx.doi.org/10.1109/ACCESS.2019.2939201>.
- [98] Irem Ulku and Erdem Akagunduz. *A Survey on Deep Learning-based Architectures for Semantic Segmentation on 2D images*. 2020. arXiv: 1912.10230 [cs.CV].
- [99] Kaiming He et al. *Mask R-CNN*. 2018. arXiv: 1703.06870 [cs.CV].
- [100] Xiang Zhang. *Simple Understanding of Mask RCNN*. Apr. 2018. URL: <https://alittlepain833.medium.com/simple-understanding-of-mask-rcnn-134b5b330e95>.
- [101] Tsung-Yi Lin et al. *Feature Pyramid Networks for Object Detection*. 2017. arXiv: 1612.03144 [cs.CV].
- [102] Shaoqing Ren et al. *Faster R-CNN: Towards Real-Time Object Detection with Region Proposal Networks*. 2016. arXiv: 1506.01497 [cs.CV].
- [103] M. A. Acero et al. “First Measurement of Neutrino Oscillation Parameters using Neutrinos and Antineutrinos by NOvA”. In: *Phys. Rev. Lett.* 123.15 (2019), p. 151803. DOI: 10.1103/PhysRevLett.123.151803. arXiv: 1906.04907 [hep-ex].
- [104] C.G.L. Bassin, Gabi Laske, and Guy Masters. “The current limits of resolution for surface wave tomography in North America. EOS Trans AGU 81:F897”. In: *Eos* 81 (Jan. 2000), F897.
- [105] Fernanda Psihas and Stanley Wojcicki. *EM energy calibration update*. NOvA Internal Document, DocDb 14413.
- [106] Shiqi Yu. *Ana2020 Nue Energy Estimator Technote*. NOvA Internal Document, DocDb 43814.
- [107] Wei Mu. *NuMu Energy Estimator Technote*. NOvA Internal Document, DocDb 43757.

- [108] Alexander Himmel, Adam Moren, and Evan Niner. *Early Permanent Cosmic Rejection with a CNN*. NOvA Internal Document, DocDb 42766.
- [109] Reddy Pratap Gandrajula and Lyudmila Kolupaeva. *Cosmic rejection BDT for ν_e appearance analysis*. NOvA Internal Document, DocDb 42473.
- [110] Thomas Warburton. *NuMu Cosmic Rejection BDT Technote*. NOvA Internal Document, DocDb 42358.
- [111] Joseph Lozier. *ModularExtrap Technical Note*. NOvA Internal Document, DocDb 12563.
- [112] Aaron Mislivec. *Tech Note: Near-to-Far Extrapolation in Transverse Momentum for the NOvA 2020 3-Flavor Analysis*. NOvA Internal Document, DocDb 44401.
- [113] Miranda Elkins. *BEN+Michel Decomposition Tech Note*. NOvA Internal Document, DocDb 43712.
- [114] Daniel Pershey. *MEFinder Tech Note*. NOvA Internal Document, DocDb 14789.
- [115] Gary J. Feldman and Robert D. Cousins. “A Unified approach to the classical statistical analysis of small signals”. In: *Phys. Rev. D* 57 (1998), pp. 3873–3889. DOI: 10.1103/PhysRevD.57.3873. arXiv: physics/9711021.
- [116] Lisa W Koerner and Ryan J Nichol. *Calibration Tech Note*. NOvA Internal Document, DocDb 43820.
- [117] Kirk Bays et al. *Cross-section Tuning Tech Note*. NOvA Internal Document, DocDb 43962.
- [118] J. Nieves, J. E. Amaro, and M. Valverde. “Inclusive quasielastic charged-current neutrino-nucleus reactions”. In: *Phys. Rev. C* 70 (5 Nov. 2004), p. 055503. DOI: 10.1103/PhysRevC.70.055503. URL: <https://link.aps.org/doi/10.1103/PhysRevC.70.055503>.
- [119] Richard Gran. “Model Uncertainties for Valencia RPA Effect for MINERvA”. In: (May 2017). arXiv: 1705.02932 [hep-ex].

- [120] P. Adamson et al. “Study of quasielastic scattering using charged-current $\nu\mu$ -iron interactions in the MINOS near detector”. In: *Phys. Rev. D* 91.1 (2015), p. 012005. DOI: 10.1103/PhysRevD.91.012005. arXiv: 1410.8613 [hep-ex].
- [121] Janet M. Conrad, Michael H. Shaevitz, and Tim Bolton. “Precision measurements with high-energy neutrino beams”. In: *Rev. Mod. Phys.* 70 (1998), pp. 1341–1392. DOI: 10.1103/RevModPhys.70.1341. arXiv: hep-ex/9707015.
- [122] N. Agafonova et al. “Final Results of the OPERA Experiment on ν_τ Appearance in the CNGS Neutrino Beam”. In: *Phys. Rev. Lett.* 120.21 (2018). [Erratum: Phys.Rev.Lett. 121, 139901 (2018)], p. 211801. DOI: 10.1103/PhysRevLett.120.211801. arXiv: 1804.04912 [hep-ex].
- [123] Linda Cremonesi. *Beam Systematic Uncertainties*. NOvA Internal Document, DocDb 17608.
- [124] Ryan B. Patterson and Matthew J. Wetstein. *Neutron Systematic Tech Note*. NOvA Internal Document, DocDb 43811.
- [125] Matthew Strait. *Tech note: Muon Energy Scale Systematic*. NOvA Internal Document, DocDb 20816.
- [126] Michael Baird and Thomas Warburton. *Neutrons Within Beam Spill*. NOvA Internal Document, DocDb 31144.
- [127] Ian Jolliffe. “Principal Component Analysis”. In: *International Encyclopedia of Statistical Science*. Ed. by Miodrag Lovric. Berlin, Heidelberg: Springer Berlin Heidelberg, 2011, pp. 1094–1096. ISBN: 978-3-642-04898-2. DOI: 10.1007/978-3-642-04898-2_455. URL: https://doi.org/10.1007/978-3-642-04898-2_455.
- [128] Alex Himmel. “New Oscillation Results from the NOvA Experiment”. The XXIX International Conference on Neutrino Physics and Astrophysics. 2020.

- [129] Patrick Dunne. “Latest Neutrino Oscillation Results from T2K”. The XXIX International Conference on Neutrino Physics and Astrophysics. 2020.
- [130] Ryan Patterson et al. *Development of a joint oscillation analysis by the NOvA and T2K collaborations*. Snowmass 2021 Letter of Interest. 2020.
- [131] R. Acciarri et al. “Long-Baseline Neutrino Facility (LBNF) and Deep Underground Neutrino Experiment (DUNE): Conceptual Design Report, Volume 2: The Physics Program for DUNE at LBNF”. In: (Dec. 2015). arXiv: 1512.06148 [physics.ins-det].
- [132] “The PIP-II Reference Design Report”. In: (June 2015). Ed. by Valeri Lebedev. DOI: 10.2172/1365571.
- [133] Pierre Baldi et al. “Improved Energy Reconstruction in NOvA with Regression Convolutional Neural Networks”. In: *Phys. Rev. D* 99.1 (2019), p. 012011. DOI: 10.1103/PhysRevD.99.012011. arXiv: 1811.04557 [physics.ins-det].
- [134] Dmitrii Torbunov. “Neutrino Energy Reconstruction with Recurrent Neural Networks at NOvA”. Neutrino Physics and Machine Learning. 2020.
- [135] David Duenas Tonguino et al. “Design and Operation of a Charged Particle Beamline for the NOvA Test Beam Experiment”. The XXIX International Conference on Neutrino Physics and Astrophysics. 2020.
- [136] Michael Baird, Louise Suter, and Jeremy Wolcott. *Executive summary for 2020 3-flavor analysis review*. NOvA Internal Document, DocDb 44422.

MICAH GROH
mcgroh@iu.edu

EDUCATION	Ph.D. Physics , Indiana University Supervisor: Mark Messier, Rudy Professor of Physics	March 2021
	B.S. Physics and Mathematics , University of Cincinnati Supervisor: Alex Sousa, Professor of Physics	May 2015
HONORS & AWARDS	Graduate Student Travel Award , Indiana University Monetary award from the Indiana University College of Arts and Sciences to present at the Reconstruction and Machine Learning in Neutrino Experiments conference at DESY.	September 2019
	Presentation Award , 12 th International Neutrino Summer School Awarded for presentation at the international neutrino summer school on using Super-Kamiokande data to measure neutrino oscillation parameters and the eccentricity of the earth.	August 2019
	Outstanding Candidate Award , Indiana University Monetary award as an outstanding graduate candidate of the Indiana University physics department.	July 2015
	Junior Achievement Award , University of Cincinnati Awarded a monetary prize for exceptional scholarship in physics by the University of Cincinnati.	April 2014
SELECT PUBLICATIONS	Co-author on 28 papers, see list at the end of CV. Select publications: <ul style="list-style-type: none">■ Fernanda Psihas, Micah Groh, Christopher Tunnell, and Karl Warburton. A Review on Machine Learning for Neutrino Experiments. <i>Int. J. Mod. Phys. A</i>, 35(33):2043005, 2020■ F. Psihas, E. Niner, M. Groh, R. Murphy, A. Aurisano, A. Himmel, K. Lang, M.D. Messier, A. Radovic, and A. Sousa. Context-Enriched Identification of Particles with a Convolutional Network for Neutrino Events. <i>Phys. Rev. D</i>, 100(7):073005, 2019■ M.A. Acero et al. First Measurement of Neutrino Oscillation Parameters using Neutrinos and Antineutrinos by NOvA. <i>Phys. Rev. Lett.</i>, 123(15):151803, 2019	
LEADERSHIP EXPERIENCE	Young NOvA President Elected to represent the students and post-doctoral researchers on the experiment to the rest of the collaboration.	2019-2020
	Fermilab Machine Learning Group Coordinator Coordinated the monthly meetings for the inter-experimental community for knowledge transfer and collaboration on machine learning applications in particle physics.	2018-2020

RESEARCH
EXPERIENCE

Neutrino Physics

- Applied principal component analysis for reducing systematic uncertainties for NOvA's oscillation and cross section measurements.
- Calculated the effects of systematic uncertainties for NOvA's neutrino oscillation measurements.
- Analyzed neutral pion decays to test NOvA's reconstruction and energy scale techniques.
- Computed the predicted data-driven event numbers and uncertainties in the far detector.

NOvA Detector Operations

- Expert in the NOvA DAQ software and hardware. Includes providing on-call service for the experiment.
- Served as an expert on detector timing system components. Provided support for any incident involving the timing system or power outages.
- Maintained and improved NOvA's online data monitoring systems.

Reconstruction and Machine Learning

- Improved the design and training techniques for machine learning algorithms used by NOvA.
- Developed and implemented a machine learning technique for full event reconstruction.
- Measured the performance of NOvA's particle identification algorithm for neutrino events.

Cosmic Rays

- Measured the effects of atmospheric variations on the rate of multiple-muon cosmic ray air showers.
- Incorporated CORSIKA simulation software into the NOvA simulation framework.

Python Analysis Framework

- Created a python-based analysis framework using pandas for NOvA analyses.
- Accelerated the NOvA machine learning workflow using a new data format.

DUNE 35 Ton Prototype

- Modeled the effects of space charge on the DUNE 35 ton prototype detector.

TEACHING &
MENTORING

Students Supervised

- **Abhilash Dombara**, Graduate Student, Syracuse University 2019-2020
Instructed student on the components of the NOvA DAQ and timing systems to become an on-call expert. Also, instructed on analysis of the wrong sign fraction in the NOvA detector.
- **Elijah Guess**, Undergraduate Student, Indiana University 2018-2020
Taught student on the use of machine learning to compare data and simulated neutrinos.

NOvA Experiment Workshops

2017-2020

Selected to present tutorials for collaborators on detector DAQ, machine learning, and analysis methods at yearly Young NOvA Tutorial Days, yearly DAQ On-Call Training Weeks, and at two Deep Learning Workshops.

Teaching Assistant, Indiana University

2016

Teaching assistant for algebra based physics laboratory on electricity and magnetism.

ADVOCACY & OUTREACH	HEP Advocacy Trip , Washington, D.C.	2019, 2020
	Attended the annual trip organized by Fermilab's Users Executive Committee to advocate for continued Federal investment in high-energy physics with lawmakers on Capitol Hill.	
	IU Science Fest , Indiana University, Bloomington, IN	2018, 2019
	Volunteered at the annual open house hosted by IU science departments.	
	Fermilab Open House , Fermilab, Batavia, IL	2018
	Guided visitors to Fermilab through the neutrino campus during Fermilab's annual open house.	
	NOvA Neutrino Classification Webpage	
	Created a webpage for users to classify NOvA neutrino events and compare themselves to NOvA's machine learning algorithms.	
SCHOOLS	12th International Neutrino Summer School , Fermilab, Batavia, IL	August 2019
	2nd Computational and Data Science School for HEP , Princeton University, Princeton, NJ	August 2018
	US Particle Accelerator School , Lisle, IL	June 2017
INVITED PRESENTATIONS	High Energy Physics Seminar , Indiana University, Bloomington, IN	December 2019
	Seminar: <i>Measuring Neutrino Oscillation Parameters with Machine Learning on NOvA</i>	
	Reconstruction and Machine Learning in Neutrino Experiments , DESY, Hamburg, Germany	September 2019
	Talk: <i>Overview of Machine Learning on NOvA</i>	
	International Workshop on Next Generation Nucleon Decay and Neutrino Detectors , TRIUMF, Vancouver, Canada	November 2018
	Parallel Talk: <i>Machine Learning Techniques on NOvA</i>	
CONTRIBUTED PRESENTATIONS	Neutrino Physics and Machine Learning , Chicago, IL	July 2020
	Talk: <i>Full Event Reconstruction on NOvA using Instance Segmentation</i>	
	29th International Conference on Neutrino Physics and Astrophysics , Chicago, IL	June 2020
	Poster: <i>Improvements and New Applications of Machine Learning on NOvA</i>	

- April Meeting of the American Physical Society,**
Washington, D.C. April 2020
Parallel Talk: *Physics Capabilities of the NOvA Experiment*
- April Meeting of the American Physical Society,**
Denver, CO April 2019
Parallel Talk: *Systematic Uncertainties for the NOvA Oscillation Analyses*
Poster: *Neutrino Event Reconstruction with Machine Learning on NOvA*
- Fermilab New Perspectives,** Fermilab, Batavia, IL June 2018
Talk: *NOvA Reconstruction using Deep Learning*
- 28th International Conference on Neutrino Physics and Astrophysics,**
Heidelberg, Germany June 2018
Poster: *Systematic Uncertainties and Cross-Checks for the NOvA Joint $\nu_\mu + \nu_e$ Analysis*
Poster: *Neutrino Physics with Deep Learning: Techniques and Applications on NOvA*

PUBLICATIONS
LIST

Peer Reviewed Publications

1. M. A. Acero et al. Search for slow magnetic monopoles with the NOvA detector on the surface. *Phys. Rev. D*, 103(1):012007, 2021
2. Fernanda Psihas, Micah Groh, Christopher Tunnell, and Karl Warburton. A Review on Machine Learning for Neutrino Experiments. *Int. J. Mod. Phys. A*, 35(33):2043005, 2020
3. B. Abi et al. First results on ProtoDUNE-SP liquid argon time projection chamber performance from a beam test at the CERN Neutrino Platform. *JINST*, 15(12):P12004, 2020
4. B. Abi et al. Neutrino interaction classification with a convolutional neural network in the DUNE far detector. *Phys. Rev. D*, 102(9):092003, 2020
5. B. Abi et al. Long-baseline neutrino oscillation physics potential of the DUNE experiment. *Eur. Phys. J. C*, 80(10):978, 2020
6. M.A. Acero et al. Adjusting neutrino interaction models and evaluating uncertainties using NOvA near detector data. *Eur. Phys. J. C*, 80(12):1119, 2020
7. M.A. Acero et al. Supernova neutrino detection in NOvA. *JCAP*, 10:014, 2020
8. M.A. Acero et al. Search for multimessenger signals in NOvA coincident with LIGO/Virgo detections. *Phys. Rev. D*, 101(11):112006, 2020
9. M.A. Acero et al. First Measurement of Neutrino Oscillation Parameters using Neutrinos and Antineutrinos by NOvA. *Phys. Rev. Lett.*, 123(15):151803, 2019
10. F. Psihas, E. Niner, M. Groh, R. Murphy, A. Aurisano, A. Himmel, K. Lang, M.D. Messier, A. Radovic, and A. Sousa. Context-Enriched Identification of Particles with a Convolutional Network for Neutrino Events. *Phys. Rev. D*, 100(7):073005, 2019

11. M.A. Acero et al. Observation of seasonal variation of atmospheric multiple-muon events in the NOvA Near Detector. *Phys. Rev. D*, 99(12):122004, 2019
12. M.A. Acero et al. Measurement of neutrino-induced neutral-current coherent π^0 production in the NOvA near detector. *Phys. Rev. D*, 102(1):012004, 2020
13. M.A. Acero et al. New constraints on oscillation parameters from ν_e appearance and ν_μ disappearance in the NOvA experiment. *Phys. Rev. D*, 98:032012, 2018
14. P. Adamson et al. Search for active-sterile neutrino mixing using neutral-current interactions in NOvA. *Phys. Rev. D*, 96(7):072006, 2017

Submitted for Publication

15. B. Abi et al. Prospects for Beyond the Standard Model Physics Searches at the Deep Underground Neutrino Experiment. 8 2020
16. B. Abi et al. Supernova Neutrino Burst Detection with the Deep Underground Neutrino Experiment. 8 2020

Conference Proceedings

17. Reddy Pratap Gandrajula and Micah Groh. Systematic Uncertainties and Cross-Checks for the NOvA Joint $\nu_\mu + \nu_e$ Analysis. In *28th International Conference on Neutrino Physics and Astrophysics*, 8 2018

Design Reports

18. Babak Abi et al. Volume IV. The DUNE far detector single-phase technology. *JINST*, 15(08):T08010, 2020
19. Babak Abi et al. Volume III. DUNE far detector technical coordination. *JINST*, 15(08):T08009, 2020
20. Babak Abi et al. Deep Underground Neutrino Experiment (DUNE), Far Detector Technical Design Report, Volume II DUNE Physics. 2 2020
21. Babak Abi et al. Volume I. Introduction to DUNE. *JINST*, 15(08):T08008, 2020
22. B. Abi et al. The DUNE Far Detector Interim Design Report, Volume 3: Dual-Phase Module. 7 2018
23. B. Abi et al. The DUNE Far Detector Interim Design Report, Volume 2: Single-Phase Module. 7 2018
24. B. Abi et al. The DUNE Far Detector Interim Design Report Volume 1: Physics, Technology and Strategies. 7 2018
25. B. Abi et al. The Single-Phase ProtoDUNE Technical Design Report. 6 2017
26. R. Acciarri et al. Long-Baseline Neutrino Facility (LBNF) and Deep Underground Neutrino Experiment (DUNE): Conceptual Design Report, Volume 4 The DUNE Detectors at LBNF. 1 2016

27. R. Acciarri et al. Long-Baseline Neutrino Facility (LBNF) and Deep Underground Neutrino Experiment (DUNE): Conceptual Design Report, Volume 2: The Physics Program for DUNE at LBNF. 12 2015
28. R. Acciarri et al. Long-Baseline Neutrino Facility (LBNF) and Deep Underground Neutrino Experiment (DUNE): Conceptual Design Report, Volume 1: The LBNF and DUNE Projects. 1 2016