

Constraints on neutrino oscillation parameters with the NOvA experiment

by

Erika Catano-Mur

A dissertation submitted to the graduate faculty
in partial fulfillment of the requirements for the degree of
DOCTOR OF PHILOSOPHY

Major: High Energy Physics

Program of Study Committee:
Mayly Sanchez, Major Professor
Soeren A. Prell
Peter J. Sherman
Amanda Weinstein
Kerry L. Whisnant

The student author, whose presentation of the scholarship herein was approved by the program of study committee, is solely responsible for the content of this dissertation. The Graduate College will ensure this dissertation is globally accessible and will not permit alterations after a degree is conferred.

Iowa State University

Ames, Iowa

2018

Copyright © Erika Catano-Mur, 2018. All rights reserved.

TABLE OF CONTENTS

	Page
ABSTRACT	v
CHAPTER 1. INTRODUCTION	1
1.1 Preface	1
1.2 Discovery of the neutrino	2
1.3 Neutrinos in the Standard Model	5
1.4 Massive neutrinos	7
1.5 Neutrino mixing and neutrino oscillations	8
1.5.1 Two-flavor oscillations	11
1.5.2 Three-flavor oscillations	13
1.6 Current knowledge and open questions	15
CHAPTER 2. THE NOvA EXPERIMENT	17
2.1 Experimental design and physics goals	18
2.2 The NuMI neutrino beam and NOvA’s off-axis design	18
2.2.1 Off-axis design	21
2.3 The NOvA detectors	23
2.3.1 Far Detector (FD)	25
2.3.2 Near Detector (ND)	26
2.3.3 Data acquisition	27
2.4 Simulation	28
2.5 Data analyzed	30
CHAPTER 3. OSCILLATION ANALYSIS FRAMEWORK	32
3.1 Analysis strategy	32
3.1.1 Decomposition and extrapolation	34
3.2 Event topology and reconstruction goals	37
3.3 Event reconstruction	40
3.4 Particle identification	44
3.4.1 Convolutional visual network (CVN)	44
3.4.2 Reconstructed Muon Identifier (ReMId)	45
3.4.3 Other particle identifiers	46
3.5 Energy reconstruction	50
3.5.1 Calorimetric energy and calibration	50
3.5.2 Muon neutrino energy	54
3.5.3 Electron neutrino energy	54
3.6 Sources of systematic uncertainty	55
3.6.1 Flux systematics	56

3.6.2	Cross section systematics	56
3.6.3	Detector response systematics	57
3.6.4	Calibration and energy response systematics	57
3.6.5	Other systematics	58
CHAPTER 4. ν_μ DISAPPEARANCE ANALYSIS		59
4.1	Event selection	60
4.1.1	FD selection performance	62
4.1.2	ND selection	62
4.2	Analysis binning	64
4.3	Constraints from ND Data	66
4.4	Cosmic background prediction	71
4.5	Far detector prediction and systematics	72
4.5.1	Systematic uncertainties	73
CHAPTER 5. ν_e APPEARANCE ANALYSIS		75
5.1	Event selection	76
5.1.1	Core sample	77
5.1.2	Peripheral sample	78
5.1.3	FD selection performance	81
5.1.4	ND selection	81
5.2	Analysis binning	83
5.3	Constraints from ND Data	84
5.3.1	Signal prediction	85
5.3.2	Background prediction	86
5.3.3	Constrained FD components and extrapolation systematics	96
5.4	Cosmic background prediction	98
5.5	Final ν_e prediction and systematics	99
5.5.1	Systematic uncertainties	100
CHAPTER 6. RESULTS OF THE JOINT $\nu_e + \nu_\mu$ ANALYSIS		103
6.1	Fitting and determination of significance	103
6.1.1	Treatment of oscillation parameters and construction of contours	104
6.1.2	Treatment of systematic uncertainties	105
6.1.3	Confidence intervals and Feldman-Cousins corrections	106
6.2	Physics sensitivity	108
6.2.1	Systematic effects on the measurement of oscillation parameters	110
6.3	Far detector data	112
6.4	Joint fit results	118
CHAPTER 7. PROJECTED MEASUREMENTS WITH NOvA		123
7.1	Antineutrino beam running	124
7.2	Run plan optimization	125
7.3	Physics reach of NOvA	128
7.3.1	Determination of mass hierarchy	129

7.3.2	Determination of CP violation	130
7.3.3	Rejection of maximal mixing and octant determination	132
7.4	Possible improvements	135
7.4.1	Changes in beam configuration	135
7.4.2	Analysis improvements	135
7.4.3	Joint analysis with T2K	135
7.5	Conclusion	137
BIBLIOGRAPHY		138
APPENDIX OSCILLATION PARAMETERS		149
A.1	2016 best fit results	149
A.2	2017 best fit results	150

ABSTRACT

NOvA is a long-baseline neutrino oscillation experiment, which consists of two finely-segmented liquid-scintillator detectors operating 14 mrad off-axis from the NuMI muon neutrino beam. With an 810 km baseline, the combined measurements of muon neutrino disappearance and electron neutrino appearance allow the determination of neutrino oscillation unknowns, namely the mass hierarchy, the octant of the largest neutrino mixing angle, and the CP violating phase.

In this dissertation I present the joint analysis of $\nu_\mu \rightarrow \nu_\mu$ and $\nu_\mu \rightarrow \nu_e$ oscillation data with an exposure of 8.85×10^{20} protons on target on the 14 kton detector. It includes the estimation of neutrino energy distributions in the far detector using near detector data constraints, details of the implementation of the simultaneous fit of ν_μ and ν_e samples, and the effect of systematic uncertainties on the measurement of oscillation parameters. I also discuss NOvA's projected sensitivity to determine the mass hierarchy and discover CP violation in future analyses with increased exposure and the addition of antineutrino datasets.

With 66 electron neutrino and 126 muon neutrino candidates, the best fit to the data corresponds to the normal mass hierarchy, $\Delta m_{32}^2 = 2.44 \times 10^{-3} eV^2/c^4$, $\sin^2 \theta_{23} = 0.56$, and $\delta_{CP} = 1.21\pi$, with the allowed regions $\Delta m_{32}^2 \in [2.37, 2.52] \times 10^{-3} eV^2/c^4$, $\sin^2 \theta_{23} \in [0.43, 0.51] \cup [0.52, 0.60]$ and $\delta_{CP} \in [0, 0.12\pi] \cup [0.91\pi, 2\pi]$ at the 68.3% C.L. Our data disfavor maximal mixing by 0.8σ , $\delta_{CP} = \pi/2$ in the inverted hierarchy at higher than 3σ , and the entire inverted mass hierarchy at the 95% confidence level. By increasing the size of the samples and adding anti-neutrino data, NOvA can potentially resolve the neutrino mass hierarchy at 4σ C.L. or higher for some combinations of oscillation parameters. Improvements to the sensitivity to the determination of the mass hierarchy and CP violation in the neutrino sector can be achieved through gains in the NuMI beam intensity, the inclusion of external data, and refinements in the analysis methodology.

CHAPTER 1. INTRODUCTION

1.1 Preface

The data presented in this dissertation was obtained with the NOvA experiment, operated by the Fermi National Accelerator Laboratory (Fermilab) in the US. As of 2018, the NOvA collaboration was made up of more than 240 scientists and engineers from 50 institutions in six countries. The lead collaborator at Iowa State University is Professor Mayly Sanchez, Ph.D., who supervised this work.

Chapter 1 presents an overview of the physics of neutrino oscillations. It was written based on a variety of sources outside of the collaboration.

Chapter 2 is a description of the NOvA experiment and chapter 3 describes the framework for the analysis. These two chapters were written from a combination of published and internal documents. The majority of figures come from an official plot database and have been created and approved by the NOvA collaboration for public use. Other figures include citations where appropriate. Individual contributions have been credited in the text.

Chapters 4, 5 and 6 present the description and results of a combined analysis of ν_e appearance and ν_μ disappearance data. These results were published by the NOvA collaboration in Phys. Rev. D98, 032012 (2018). I was one of the three members of the committee that wrote the manuscript, alongside Dr. Gary Feldman and Dr. Jeremy Wolcott. My direct contributions include: the method for the estimation of the beam electron neutrino background, the evaluation of the effects of systematic uncertainties in the ν_e and ν_μ predictions, and the implementation of the simultaneous fit of ν_e and ν_μ data (including sensitivities and results). I created figures 4.8-4.12, 5.9-5.12, 5.16, 5.18-5.20, 6.1-6.7, and 6.10-6.13 and tables 4.1, 5.1-5.7, 6.1-6.3, with occasional input from other members of the ν_e and ν_μ analysis groups. Note that these methods were first used for

the previous publication Phys. Rev. Lett. 118, 231801 (2017), and continue to be used in 3-flavor neutrino oscillation analyses with NOvA data.

Chapter 7 presents the expected physics sensitivity of the NOvA experiment and is based on my own work.

1.2 Discovery of the neutrino

At the beginning of the 20th century, only three particles were known: the proton, the electron and the photon, having positive, negative and neutral charge, respectively. These corresponded to the three types of radioactive decay by-products named by Rutherford: α , β , γ [1]. There were a few experimental hints that this picture was incomplete, in particular with the observations of β decay.

In 1914, J. Chadwick conducted studies of the energy spectra in β decay [2]. Originally, this decay was thought to be the spontaneous emission of an electron from an atom, resulting in a different nucleus ($N \rightarrow N' + e$). As a two-body process connected with the quantized energy levels of the atom, the expectation for the final state was that the distribution of the electrons' energy would be discrete, but the measured spectrum turned out to be continuous (Figure 1.1). As it was understood, the decay would violate the principle of conservation of energy.




Figure 1.1: Expected and measured energy spectrum of the electron in β decay. From [3].

In 1930, W. Pauli proposed that a very light, neutral particle was produced in the decay together with the electron ($N \rightarrow N' + e + \nu$); the undetectable energy carried away by this particle would

allow for the continuous spectrum that was observed. Pauli intended to call it the *neutron* [4], but this name was coined for a different, heavier neutral particle discovered by Chadwick in 1932 [5]. In 1934, E. Fermi [6] published his formulation of the theory of β decay and renamed the mysterious particle from β decay *neutrino* (Italian for *little neutral one*). Figure 1.2 shows Fermi's proposed 4-fermion coupling involving the neutron, proton, electron and the neutrino.




Figure 1.2: Fermi 4-fermion coupling. From [3].

Fermi's theory also implies that neutrinos should scatter off matter through the inverse process, $\bar{\nu} + p^+ \rightarrow n + e^+$. H. Bethe and R. Peierls [7] estimated the cross-section of this process to be $\sigma_{\bar{\nu}} \leq 10^{-44} \text{cm}^2$ for an energy $E_{\bar{\nu}} \simeq 2 \text{MeV}$. This means that a neutrino would have a mean free path of thousands of light years while traveling in water. Their conclusion was that this interaction would be impossible to observe [3].

B. Pontecorvo was among the first to realize that this observation was feasible [8]. In 1946, he proposed that with a neutrino flux of $\sim 10^{11} \nu/\text{cm}^2/\text{s}$, which is typical near nuclear reactors, a detector with a mass of $\sim 1 \text{ton}$ could detect a few neutrino events per day. F. Reines and C. Cowan were first to succeed using a liquid scintillator detector near a nuclear reactor and looking for coincidences of positron annihilation and delayed neutron capture. After an early failed experiment in 1953 [9], they were successful in 1956 after realizing that the experiment had to be moved underground to reduce cosmic background [10]. The new experiment also had better segmentation and a larger volume. Their experiment, sketched in Figure 1.3, constituted the discovery of the electron antineutrino.




Figure 1.3: The Savannah River experiment by Reines and Cowan. Left: Diagram of the detection technique. An antineutrino from the reactor interacts in the detector producing a positron and a neutron. The positron annihilation creates photons that are detected in the scintillator. The delayed capture of the neutron after it slows down also generates scintillation light. Right: experimental setup. Tanks A and B contain water and dissolved CdCl_2 . Tanks I, II and III are filled with scintillators and instrumented with PMTs. From [11].

The view of particle physics grew with the discovery of lepton flavors. Cosmic ray muons had been discovered, but were initially thought to be similar to pions and not heavier analogs of electrons [12]. A neutrino that had to be different from the one in β decay was proposed to be the product of pion decay $\pi^- \rightarrow \mu^- + \bar{\nu}_\mu$ [13]. The typical energies involved in this process are higher, and in Fermi's theory this implied that the cross section is also higher and therefore it should be easier to detect.

The observation of the muon neutrino was accomplished by L. Lederman, M. Schwartz and S. Steinberger (LSS) in 1962 at Brookhaven National Lab [14]. Their experiment also involved the creation of the first neutrino beam. The basic design, shown in Figure 1.4 is one still used today: a boosted proton hits a target, producing pions and other hadrons; the latter decay into neutrinos and other particles, which are stopped with a thick shield. A neutrino detector is located behind the shield, and the neutrinos are observed from the appearance of muons in the detector.

The observation of the third neutrino type, the tau neutrino, happened much later. The tau was discovered in 1975 at Stanford Linear Accelerator Laboratory [15]. With the addition of a third charged lepton, the theory predicted that an associated third neutrino should also exist. The tau




FIG. 1. Plan view of AGS neutrino experiment.

Figure 1.4: Diagram of the LSS accelerator neutrino experiment. The neutrino beam is produced from pion decay and detected by a 10-ton aluminum spark chamber. Other particles are stopped by a thick iron wall. From [14].

neutrino was observed directly for the first time in 2000 in the DONUT experiment at Fermilab [16].

1.3 Neutrinos in the Standard Model

The Standard Model (SM) of particle physics is the theory of the strong, electromagnetic and weak interactions. It is a gauge theory based on the group $SU(3)_C \times SU(2)_L \times U(1)_Y$, where all fundamental particles are arranged as irreducible representations of the group. The fermion content of the SM is presented in Table 1.1.

The strong interaction is a $SU(3)$ gauge theory, and it involves the quarks of different colors (three colors $i = 1, 2, 3$) and flavors (six flavors u, d, c, s, t, b) and is mediated by eight gluons. Color states are confined, i.e. only combinations of quarks can exist as free particles. The known color singlets are known as hadrons, and they can be further classified as mesons (made of a quark-antiquark pair) or baryons (three quarks). The strong nuclear force between protons and neutrons is a manifestation of the underlying strong force among the quarks [17].

The electroweak interaction corresponds to the symmetry group $SU(2)_L \times U(1)_Y$, and is spontaneously broken down to the electromagnetic interaction $U(1)_Q$ at around 100 GeV. Under $SU(2)_L$, left-handed quarks and leptons are doublets and right-handed fermions are singlets; this forbids

Table 1.1: Irreducible fermion representations in the Standard Model. The notation for the representations is $(d_{SU(3)}, d_{SU(2)})_Y$. $i = 1, 2, 3$ denotes color.

Representation	Left-handed		Right-handed		
	Leptons $(\mathbf{1}, \mathbf{2})_{-\frac{1}{2}}$	Quarks $(\mathbf{3}, \mathbf{2})_{+\frac{1}{6}}$	Leptons $(\mathbf{1}, \mathbf{1})_{-1}$	Quarks $(\mathbf{3}, \mathbf{1})_{+\frac{2}{3}}$ $(\mathbf{3}, \mathbf{1})_{-\frac{1}{3}}$	
First generation	$\begin{pmatrix} \nu_e \\ e \end{pmatrix}_L$	$\begin{pmatrix} u^i \\ d^i \end{pmatrix}_L$	e_R	u^i_R	d^i_R
Second generation	$\begin{pmatrix} \nu_\mu \\ \mu \end{pmatrix}_L$	$\begin{pmatrix} c^i \\ s^i \end{pmatrix}_L$	μ_R	c^i_R	s^i_R
Third generation	$\begin{pmatrix} \nu_\tau \\ \tau \end{pmatrix}_L$	$\begin{pmatrix} t^i \\ b^i \end{pmatrix}_L$	τ_R	t^i_R	b^i_R

lepton mass terms before the electroweak symmetry breaking. After symmetry breaking, three gauge bosons (W^\pm, Z) acquire mass and become the effective mediators of the weak force, and one boson (the photon) remains massless. The Higgs mechanism also gives masses to quarks and charged leptons. There is no right-handed neutrino and thus neutrinos are massless in the SM.

Neutrinos interact in the SM via charged and neutral currents [3]:

$$\mathcal{L}_{SM} - \frac{g}{\sqrt{2}} \sum_{\alpha} \bar{\nu}_{\alpha} \gamma_{\mu} P_L l_{\alpha} W_{\mu}^{+} - \frac{g}{2 \cos \theta_W} \sum_{\alpha} \bar{\nu}_{\alpha} \gamma_{\mu} P_L \nu_{\alpha} Z_{\mu}^0 + h.c. \quad (1.1)$$

Note that the weak currents only mix the left-handed fields. The charged current (CC) term is mediated by the W^\pm bosons, and involves the charged lepton and the neutrino from the doublets in Table 1.1. In other words, the flavor of the neutrino can be identified from the flavor of the associated charged lepton in a CC interaction. The neutral current interaction (NC) is mediated by the neutral Z boson, and only involves neutrinos. With these interactions, the β decay and Fermi's 4-fermion coupling from Figure 1.2 is now understood as exchange of a W boson, shown in Figure 1.5.




Figure 1.5: β decay in the Standard Model. From [3].

1.4 Massive neutrinos

From a theoretical point of view, there is not a fundamental reason why neutrinos have to be massless in the SM. Unlike, e.g., the photon, where a zero mass is dictated by the $U(1)_Q$ gauge symmetry, the absence of a right-handed neutrino ν_R and a Dirac mass term are not mandatory, but rather a choice [18]. While all the data available at the time the SM was proposed was consistent with a massless, neutral, left-handed neutrino, none of the experiments had enough sensitivity to prohibit very small neutrino masses. Another theoretical motivation that makes massive neutrinos attractive is that they are a natural consequence of many Beyond the Standard Model (BSM) theories as a low-energy manifestation of a high-energy scale.

There are two kinds of experiments that can probe neutrino masses: kinematic and exclusive tests [18]. Kinematic tests, also known as direct tests, involve processes that are allowed even if neutrinos are massless, and depend on the experimental sensitivity to small changes in the final state energy distribution. Examples of this kind of experiment involve nuclear β decay, pion decay, and tau decay. Cosmological and astrophysical data can also set an upper limit on the sum of the neutrino masses [19]. Exclusive tests involve processes that are forbidden if neutrinos are massless. These might explicitly depend on neutrino mixing, such as neutrino oscillations and neutrino decays, or not, such as a neutrino magnetic moment or neutrino-less double-beta decay [18].

To date, none of the kinematic searches of neutrino mass have provided definitive evidence; some of the current limits that they have set are [20]:

- $m_{\bar{\nu}_e} < 2 \text{ eV}$, from ${}^3\text{H}\beta$ decay;
- $m_{\nu_\mu} < 0.17 \text{ MeV}$, from pion decay;
- $m_{\nu_\tau} < 18.2 \text{ MeV}$, from two tau decay channels in LEP.

On the other hand, exclusive tests have provided unmistakable evidence that neutrinos have masses [19]. Several experiments with solar, atmospheric, reactor and accelerator neutrinos have shown the existence of neutrino oscillations, with the implication that there are massive neutrinos and that they mix.

1.5 Neutrino mixing and neutrino oscillations

Neutrino mixing means that the left-handed neutrino fields $\nu_{\alpha L}$ that enter the CC interaction in (1.1) are linear combinations of three or more mass eigenstates:

$$|\nu_\alpha\rangle = \sum_k U_{\alpha k}^* |\nu_k\rangle \quad (\alpha = e, \mu, \tau), \quad (1.2)$$

where ν_k is the left-handed component of a field with mass m_k , and U is the leptonic mixing matrix, commonly known as the PMNS (Pontecorvo-Maki-Nakagawa-Sakata) matrix [21, 22]. In the SM and (1.1), neutrinos only couple to the charged lepton in the same flavor doublet. This implies that neutrino flavor states are orthogonal, and thus the PMNS matrix is unitary:

$$\begin{aligned} \langle \nu_\alpha | \nu_\beta \rangle &= \langle \sum_i U_{\alpha i} | \sum_j U_{\beta j}^* \nu_j \rangle \\ \delta_{\alpha\beta} &= \sum_{i,j} U_{\alpha i} U_{\beta j}^* \langle \nu_i | \nu_j \rangle \\ \delta_{\alpha\beta} &= \sum_i U_{\alpha i} U_{\beta i}^*, \end{aligned} \quad (1.3)$$

where we assumed that the mass eigenstates are orthogonal (m_i, m_j are not degenerate).




Figure 1.6: Diagram showing neutrino flavor change (oscillation) in vacuum. *Amp* denotes an amplitude. From [23]

Current oscillation data can be explained with three flavor eigenstates, and measurements on the invisible decay width of the Z boson limit the number of light active neutrinos that couple to the Z to three. The available neutrino data [20] supports the existence of three light massive neutrinos: ν_1, ν_2, ν_3 with $m_1 \neq m_2 \neq m_3$ and $m_{1,2,3} \lesssim 1\text{eV}$. The number of mass eigenstates can be greater than three if e.g. there are sterile neutrinos that couple to the light flavor neutrinos.

Neutrino oscillations are a quantum mechanical consequence of the non-degenerate neutrino masses and lepton mixing [20]. Figure 1.6 helps illustrate this phenomenon. At the source, a neutrino ν_α and a charged lepton l_α are produced. The neutrino travels a distance L and interacts with a target, and a lepton l_β is detected. If $\alpha \neq \beta$, this means that the neutrino changed flavor to ν_β . In the bottom panel of Figure 1.6, the same process is depicted in terms of the mass eigenstates. The initial neutrino ν_α is a linear combination of the mass eigenstates ν_i that propagate and reach the detector. The interaction with the target is mediated by the weak force and involves a flavor eigenstate, so back in the flavor basis there is some probability to see each of the ν_β .

The following is a simplified derivation of the oscillation probabilities using a plane wave approximation [3]. A full calculation requires wave packet or quantum field theory formalisms [20].

Consider a neutrino of flavor α is produced at t_0 . It is a superposition of the mass eigenstates that we assume to be plane waves with spatial momentum \mathbf{p} :

$$|\nu_\alpha(t_0)\rangle = \sum_i U_{\alpha i}^* |\nu_i(\mathbf{p})\rangle. \quad (1.4)$$

The mass eigenstates are eigenstates of the free Hamiltonian:

$$\hat{H}|\nu_i(\mathbf{p})\rangle = E_i(\mathbf{p})|\nu_i(\mathbf{p})\rangle, \quad E_i(\mathbf{p})^2 = \mathbf{p}^2 + m_i^2. \quad (1.5)$$

The time evolution operator from $t_0 \rightarrow t$ is given by $e^{-i\hat{H}(t-t_0)}$. The state at time t is given by

$$|\nu_\alpha(t)\rangle = e^{-i\hat{H}(t-t_0)}|\nu_\alpha(t_0)\rangle = \sum_i U_{\alpha i}^* e^{-iE_i(\mathbf{p})(t-t_0)} |\nu_i(\mathbf{p})\rangle. \quad (1.6)$$

The probability that at time t the state is in flavor β is

$$P(\nu_\alpha \rightarrow \nu_\beta)(t) = |\langle \nu_\beta | \nu_\alpha(t) \rangle|^2 = \left| \sum_i U_{\beta i} U_{\alpha i}^* e^{-iE_i(\mathbf{p})(t-t_0)} \right|^2, \quad (1.7)$$

where we used the orthogonality relation $\langle \nu_i(\mathbf{p}) | \nu_j(\mathbf{p}) \rangle = \delta_{ij}$.

For ultra-relativistic neutrinos, we can approximate

$$E_i(\mathbf{p}) - E_j(\mathbf{p}) \simeq \frac{1}{2} \frac{m_i^2 - m_j^2}{|\mathbf{p}|} + \mathcal{O}(m^4), \quad (1.8)$$

and $L \simeq (t - t_0)$. The probability becomes

$$P(\nu_\alpha \rightarrow \nu_\beta) = \sum_{i,j} U_{\alpha i}^* U_{\beta i} U_{\alpha j} U_{\beta j}^* e^{-i \frac{\Delta m_{ij}^2 L}{2|\mathbf{p}|}}, \quad (1.9)$$

where we defined

$$\Delta m_{ij}^2 \equiv m_i^2 - m_j^2. \quad (1.10)$$

Defining $W_{\alpha\beta}^{ij} \equiv [U_{\alpha i} U_{\beta i}^* U_{\alpha j}^* U_{\beta j}]$ and using the unitarity of the mixing matrix, we can rewrite the probability in the more familiar form:

$$\begin{aligned} P(\bar{\nu}_\alpha \rightarrow \bar{\nu}_\beta) &= \delta_{\alpha\beta} - 4 \sum_{j>i} \text{Re}[W_{\alpha\beta}^{ij}] \sin^2\left(\frac{\Delta m_{ij}^2 L}{4E_\nu}\right) \\ &\quad \mp 2 \sum_{j>i} \text{Im}[W_{\alpha\beta}^{ij}] \sin\left(\frac{\Delta m_{ij}^2 L}{2E_\nu}\right), \end{aligned} \quad (1.11)$$

where the plus (minus) sign applies to neutrinos (antineutrinos), $|\vec{p}| \simeq E_\nu$. The expression (1.11) shows that non-zero, non-degenerate neutrino masses are a requirement for neutrino oscillations. It also shows that the probability “oscillates” as a function of L/E , hence the name. The unitarity of U implies

$$\sum_{\beta} P(\nu_{\alpha} \rightarrow \nu_{\beta}) = 1, \quad (1.12)$$

thus the total neutrino flux remains the same but may get redistributed among flavors.

1.5.1 Two-flavor oscillations

Now let us consider the simplest case of two family mixing, where we only need one angle and one mass squared difference. The PMNS matrix is

$$U_{\text{PMNS}} = \begin{pmatrix} \cos \theta & \sin \theta \\ -\sin \theta & \cos \theta \end{pmatrix} \quad (1.13)$$

The oscillation probability is given by

$$\begin{aligned} P(\nu_{\alpha} \rightarrow \nu_{\beta}) &= \sin^2 2\theta \sin^2 \left(1.27 \frac{\Delta m^2 (\text{eV}^2) L (\text{km})}{E_{\nu} (\text{GeV})} \right), \quad \alpha \neq \beta. \\ P(\nu_{\alpha} \rightarrow \nu_{\alpha}) &= 1 - P(\nu_{\alpha} \rightarrow \nu_{\beta}). \end{aligned} \quad (1.14)$$


The probability is the same for neutrinos and antineutrinos, since there cannot be CP violation with only two families [3]. We call $\nu_{\alpha} \rightarrow \nu_{\beta}$ an “appearance” oscillation probability as the final flavor state is different than the initial, and $\nu_{\alpha} \rightarrow \nu_{\alpha}$ the “disappearance” or “survival” probability.

The probability P in (1.14) is a sinusoidal function with an amplitude that depends on the oscillation angle, and wavelength

$$L_{\text{osc}} (\text{km}) = \pi \frac{E_{\nu} (\text{GeV})}{1.27 \Delta m^2 (\text{eV}^2)}. \quad (1.15)$$

P is *maximal* for $\sin^2 2\theta = 1$ ($\theta = \pi/4$).

Figure 1.7a shows the appearance probability as a function of L . Usually neutrino oscillation experiments occur at a fixed baseline and variable energy, as in Figure 1.7b. In this case, the



(a) Appearance probability as a function of baseline, (b) Appearance probability as a function of energy, assuming constant neutrino energy.

Figure 1.7: Neutrino oscillation probability in the two flavor approximation. From [3].

position of the first oscillation maximum (counted from right to left) gives information about the mass splitting:

$$E_{\max}(\text{GeV}) = 1.27 \frac{\Delta m^2(\text{eV}^2)L(\text{km})}{\pi/2}. \quad (1.16)$$

Neutrino oscillation experiments can optimize the available neutrino energies and baselines to be on the same order of the mass splitting, $E/L \sim \Delta m^2$. In the fast oscillation regime $E/L \ll \Delta m^2$ the measurements are smeared. In slow oscillations $E/L \gg \Delta m^2$, the effects from the mixing angle and the mass splitting cannot be separated.

All the expressions above apply for neutrinos traveling in vacuum. When traveling through matter, neutrino oscillation probabilities can be significantly modified by coherent forward scattering off electrons and nuclei [24, 25]. Writing explicitly the two flavor states, the equation of motion including the effect of matter is written as:

$$i \frac{d}{dt} \begin{pmatrix} \nu_e \\ \nu_\mu \end{pmatrix} = \frac{1}{2} \begin{pmatrix} -(\frac{\Delta m^2}{2E} \cos 2\theta - \sqrt{2}G_F N_e) & \frac{\Delta m^2}{2E} \sin 2\theta \\ \frac{\Delta m^2}{2E} \sin 2\theta & (\frac{\Delta m^2}{2E} \cos 2\theta - \sqrt{2}G_F N_e) \end{pmatrix} \begin{pmatrix} \nu_e \\ \nu_\mu \end{pmatrix}, \quad (1.17)$$

where G_F is the Fermi constant, and N_e is the electron density in matter which could change in time or space. The term $\sqrt{2}G_F N_e$ acts as an extra potential due to the difference between ν_μ and ν_e scattering amplitude off electrons, and reverses sign for antineutrinos, Note that the eigenstates of this system are not the vacuum mass eigenstates ν_1 and ν_2 . Assuming that the matter density

is constant, the solution to the equation of motion can be written as (1.14) with the substitutions

$$\begin{aligned}\sin^2 2\theta &\rightarrow \sin^2 2\theta' = \frac{\tan^2 2\theta}{(1 - \frac{N_e}{N_e^{res}})^2 + \tan^2 2\theta}, \\ \Delta m^2 &\rightarrow \Delta m'^2 = \Delta m^2 \sqrt{(1 - \frac{N_e}{N_e^{res}})^2 \cos^2 2\theta + \sin^2 2\theta},\end{aligned}\tag{1.18}$$

where

$$N_e^{res} = \frac{\Delta m^2 \cos 2\theta}{2E\sqrt{2}G_F} \approx 6.56 \times 10^6 \frac{\Delta m^2 [\text{eV}^2]}{E [\text{MeV}]} \cos 2\theta \cdot N_A [\text{cm}^{-3}]\tag{1.19}$$

is called the resonance density [25, 26]. Note that in the limit of short baselines or low matter density, the vacuum probability (1.14) remains a good approximation. Under resonance, regardless of the value of θ the probability of neutrino oscillation through matter can be greatly enhanced with respect to the probability in vacuum.

1.5.2 Three-flavor oscillations

In the case of three flavor neutrinos, the PMNS matrix can be parameterized with three angles, one CP-violating phase and two Majorana phases, as:

$$\begin{aligned}U &= \begin{pmatrix} 1 & 0 & 0 \\ 0 & c_{23} & s_{23} \\ 0 & -s_{23} & c_{23} \end{pmatrix} \begin{pmatrix} c_{13} & 0 & s_{13}e^{-i\delta_{CP}} \\ 0 & 1 & 0 \\ -s_{13}e^{i\delta_{CP}} & 0 & c_{13} \end{pmatrix} \begin{pmatrix} c_{12} & s_{12} & 0 \\ -s_{12} & c_{12} & 0 \\ 0 & 0 & 1 \end{pmatrix} \begin{pmatrix} 1 & 0 & 0 \\ 0 & e^{i\frac{\alpha_{21}}{2}} & 0 \\ 0 & 0 & e^{i\frac{\alpha_{31}}{2}} \end{pmatrix} \\ &= \begin{pmatrix} c_{12}c_{13} & s_{12}c_{13} & s_{13}e^{-i\delta_{CP}} \\ -s_{12}c_{23} - c_{12}s_{23}s_{13}e^{i\delta_{CP}} & c_{12}c_{23} - s_{12}s_{23}s_{13}e^{i\delta_{CP}} & s_{23}c_{13} \\ s_{12}s_{23} - c_{12}c_{23}s_{13}e^{i\delta_{CP}} & -c_{12}s_{23} - s_{12}c_{23}s_{13}e^{i\delta_{CP}} & c_{23}c_{13} \end{pmatrix} \times \text{diag}\left(1, e^{i\frac{\alpha_{21}}{2}}, e^{i\frac{\alpha_{31}}{2}}\right)\end{aligned}\tag{1.20}$$

where $s_{ij} \equiv \sin \theta_{ij}$ and $c_{ij} \equiv \cos \theta_{ij}$, θ_{ij} are the three mixing angles and δ_{CP} is the CP-violating phase. The Majorana phases α_{ij} do not enter the oscillation probabilities and can be safely ignored.

The neutrino oscillation probability then is determined by six parameters: the mixing angles θ_{12} , θ_{13} , θ_{23} , which define the amplitude of the oscillation probability; the differences in the squared masses of the eigenstates, $\Delta m_{21}^2 = m_2^2 - m_1^2$ and $\Delta m_{32}^2 = m_3^2 - m_2^2$, that define the oscillation frequency and the position of the oscillation maxima as a function of L/E ; and the CP-violating

phase, δ_{CP} . While it is known that $\Delta m_{21}^2 > 0$, the sign of Δm_{32}^2 or “neutrino mass ordering” [20] is still unknown.

Because the oscillations are driven by two mass-squared differences Δm_{32}^2 and Δm_{31}^2 , where $\Delta m_{31}^2 = \Delta m_{32}^2 + \Delta m_{21}^2$, the interference between two oscillation frequencies leads to terms that depend on all the mixing parameters. At leading order, the ν_μ and $\bar{\nu}_\mu$ survival probabilities in vacuum take the same form as the two-flavor approximation with the effective parameters given by [27]:

$$\begin{aligned}\sin^2 2\theta &= 4 \sin^2 \theta_{23} \cos^2 \theta_{13} (1 - \sin^2 \theta_{23} \cos^2 \theta_{13}), \\ \Delta m^2 &= \Delta m_{32}^2 + \Delta m_{21}^2 \sin^2 \theta_{12} + \Delta m_{21}^2 \cos \delta_{CP} \sin \theta_{13} \tan \theta_{23} \sin 2\theta_{12}.\end{aligned}\quad (1.21)$$

The exact symmetries of the two-flavor probability under $\theta \rightarrow \pi/2 - \theta$ and $\Delta m^2 \rightarrow -\Delta m^2$ lead to approximate degeneracies in the octant of θ_{23} and mass hierarchy in (1.21).

As before, for neutrinos traveling through matter, the propagation eigenstates are modified by the MSW effect [24, 25]. The mixing angle θ_{13} is replaced by a modified version, θ_M , given by [28]:

$$\sin^2 2\theta_M = \frac{\sin^2 2\theta_{13}}{\sin^2 2\theta_{13} + (A - \cos 2\theta_{13})^2}.\quad (1.22)$$

The size of the matter effect is determined by the parameter

$$A \equiv \pm 2\sqrt{2}G_F n_e E_\nu / \Delta m_{31}^2,\quad (1.23)$$

where G_F is the Fermi weak coupling constant, n_e is the number density of electrons and the sign of A is positive (negative) for neutrinos (antineutrinos). $\sin^2 2\theta_M$ is maximal at $A = \cos 2\theta_{13}$. This condition leads to the resonant enhancement of $\nu_\mu \leftrightarrow \nu_e$ oscillations, which can significantly alter the magnitude of ν_μ disappearance. The enhancement occurs for neutrinos in the normal hierarchy and for antineutrinos in the inverted hierarchy.

For accelerator neutrinos, the $\nu_\mu \rightarrow \nu_e$ appearance probability in matter expanded to second order in $\alpha \equiv \Delta m_{21}^2/\Delta m_{31}^2$ (≈ 0.03), is given by [29]:

$$\begin{aligned}
P(\nu_\mu \rightarrow \nu_e) &\approx \sin^2 \theta_{23} \sin^2 2\theta_{13} \frac{\sin^2 \Delta(1-A)}{(1-A)^2} \\
&+ \alpha \tilde{J} \cos(\Delta \pm \delta_{CP}) \frac{\sin \Delta A \sin \Delta(1-A)}{A(1-A)} \\
&+ \alpha^2 \cos^2 \theta_{23} \sin^2 2\theta_{12} \frac{\sin^2 \Delta A}{A^2}.
\end{aligned} \tag{1.24}$$

with

$$\tilde{J} \equiv \cos \theta_{13} \sin 2\theta_{13} \sin 2\theta_{12} \sin 2\theta_{23} \tag{1.25}$$

$$\Delta \equiv \Delta m_{31}^2 L_\nu / 4E_\nu \tag{1.26}$$

and the plus (minus) sign applies to neutrinos (antineutrinos). The first term in (1.24) is proportional to $\sin^2 \theta_{23}$ and breaks the θ_{23} octant degeneracy. Finally, the dependence on A is sensitive to the mass hierarchy and the second term in the expansion is sensitive to CP violation.

1.6 Current knowledge and open questions

Table 1.2: Best-fit values and 3σ allowed ranges of the 3-neutrino oscillation parameters, derived from a global fit of the current neutrino oscillation data. From [20, 30].

Parameter	Best fit	3σ
Δm_{21}^2 [$10^{-5} \text{eV}^2/c^4$]	7.37	6.93 - 7.96
$\Delta m_{31(23)}^2$ [$10^{-3} \text{eV}^2/c^4$]	2.56 (2.54)	2.45 - 2.69 (2.42 - 2.66)
$\sin^2 \theta_{12}$	0.297	0.250 - 0.354
$\sin^2 \theta_{23}$ (NH)	0.425	0.381 - 0.615
$\sin^2 \theta_{23}$ (IH)	0.589	0.384 - 0.636
$\sin^2 \theta_{13}$ (NH)	0.0215	0.0190 - 0.0240
$\sin^2 \theta_{13}$ (IH)	0.0216	0.0190 - 0.0242
δ_{CP}/π	1.38 (1.31)	2σ : (1.0 - 1.9) (2σ : (0.92-1.88))

All the available oscillation data allows us to determine i) Δm_{21}^2 and $\sin^2 \theta_{12}$, responsible for the solar ν_e oscillations; ii) $|\Delta m_{31}^2|$ ($|\Delta m_{32}^2|$) and $\sin^2 \theta_{23}$, responsible for the dominant oscillations

of atmospheric ν_μ and $\bar{\nu}_\mu$; and iii) $\sin^2 \theta_{13}$ responsible for the $\nu_\mu \rightarrow \nu_e$ oscillations in accelerator experiments and the $\bar{\nu}_e$ oscillations observed in reactor experiments [20]. Table 1.2 presents the best fit values and the 99.73% confidence level (CL) allowed ranges of the neutrino oscillation parameters, as well as the 95% CL allowed range of the CP violation phases δ [30, 20]. The angles and squared mass differences have been measured with good precision. The 1σ uncertainties (1/6 of the 3σ interval) for Δm_{21}^2 , $\sin^2 \theta_{12}$, $|\Delta m_{31(32)}^2|$, $\sin^2 \theta_{32}$ and $\sin^2 \theta_{13}$ are 2.3%, 5.8%, 1.6%, 9.6% and 4.0%, respectively. Maximal solar mixing ($\theta_{12} = \pi/4$) has been ruled out at more than 6σ . Maximal mixing for the atmospheric angle ($\theta_{23} = \pi/4$) has not been ruled out. The neutrino mass hierarchy, i.e. $\text{sign}\Delta m_{32}^2$ also hasn't been determined. The data show some hints on the Dirac δ_{CP} phase, $\delta_{CP} \simeq 3\pi/2$ and $\delta_{CP} \neq \pi/2$, but all values of δ_{CP} including the CP conserving $\delta_{CP} = 0, \pi$, are still allowed [20].

CHAPTER 2. THE NOvA EXPERIMENT

The NuMI Off-axis ν_e Appearance (NOvA) experiment [31] is a two-detector, long-baseline neutrino oscillation experiment that observes the NuMI muon neutrino beam from Fermilab [32]. The near detector (ND) samples the beam 1 km from the source, and the far detector (FD) observes the oscillated beam 810 km downstream, near Ash River, MN.




Figure 2.1: The NOvA experiment studies the NuMI neutrino beam from Fermilab with its two detectors. The far detector is 810 km away from the source, in Minnesota, and the near detector is on-site, only 1 km away from the target, and underground. Left: Geographical location of the detectors. Right: diagram of the experimental configuration. From [33]

As the name suggests, NOvA's physics scope is centered around oscillation measurements using the electron neutrino appearance channel. In Section 2.1 we review NOvA's main physics goals and how they informed the experimental design. Section 2.2 presents the characteristics of the neutrino beam that is used for the oscillation measurements. In Section 2.3, we explain the basic detection elements and particularities of the near and far detectors. Section 2.4 covers the beam and detector simulation of neutrino interactions in NOvA. Finally, the datasets used for the analysis in this thesis are presented in Section 2.5.

2.1 Experimental design and physics goals

As discussed in Chapter 1, the discovery of neutrino oscillations gave way to new types of experiments trying to fully characterize the various parameters. The early atmospheric and solar neutrino oscillation experiments produced information about the mass differences (Δm_{21}^2 , Δm_{32}^2) and two of the three mixing angles (θ_{12}, θ_{23}). Before the reactor experiments achieved a high precision measurement, it was clear that a non-zero θ_{13} would allow the exciting new possibility of probing CP violation in the neutrino sector and determining the neutrino mass hierarchy via $\nu_\mu \rightarrow \nu_e$ oscillations in the presence of matter effects [34].

The design of NOvA was inspired by the MINOS experiment [35]. MINOS was a first-generation long-baseline neutrino oscillation experiment measuring the NuMI beam with two detectors. The MINOS far detector was located in the Soudan mine in Minnesota, 735 km downstream of the beam target and 705 m underground. The principal channel observed in MINOS was $\nu_\mu \rightarrow \nu_\mu$, with a low sensitivity to $\nu_\mu \rightarrow \nu_e$. NOvA would make use of the same beam and an analogous two-detector design, with the following differences: [34] i) improved detection of electrons by having a higher longitudinal sampling and a highly active detector by mass (5% for MINOS vs. 70% for NOvA); ii) increased detector mass (MINOS was 5.4 kton); iii) a far detector located off the beam axis, so that the neutrino flux peaks around the same energy as the oscillation probability and the high energy NC backgrounds are reduced; iv) unlike MINOS, the NOvA near detector is not magnetized, and the far detector is not underground. We will explore these aspects of the NOvA experiment in the following sections.

2.2 The NuMI neutrino beam and NOvA's off-axis design

The source of neutrinos for NOvA is the Neutrinos at the Main Injector (NuMI) beam at Fermilab [32]. The beam is created from collisions of 120 GeV protons accelerated by the Main Injector (MI) with a graphite target. Figure 2.2 presents an overview of the main accelerators and beams at Fermilab.




Figure 2.2: An overview of the Fermilab accelerator complex. The ion source is a radio-frequency quadrupole (RQF) that feeds 35-750 keV H^- ions to the linear accelerator. A carbon foil at the end of the Linac transforms the beam into a H^+ (proton) beam. The particles are successively accelerated in the Linac (400 MeV), Booster (8 GeV), and the Main Injector (120 GeV). The Recycler is a staging area that combines batches of protons from the Booster before they transition to the MI, in order to increase the intensity. It can also send particles to the test beam facilities and the muon experiments. Protons from the Booster or the Main Injector are directed toward targets to form two neutrino beams, the low energy Booster Neutrino Beam and the high energy NuMI beam. From [36].

Pictured in Figure 2.2, the Booster is a rapid-cycling synchrotron that accelerates protons from 400 MeV to 8 GeV at a rate of 15 Hz. These protons are arranged in small packets, called “bunches”. In each cycle, the Booster delivers 1.6 μs long “batches” with a 53 MHz bunch spacing. Typically, each 1.6 μs batch contains 4.3×10^{12} protons [38]. The circumference of the Main Injector is seven times larger than the Booster, and it can deliver up to six batches accelerated to 120 GeV at a time. This pulse is called a “spill”. The Recycler sits on top of the Main Injector, and its role is to facilitate the proton injection from the Booster to the MI and perform “slip-stacking”. In this




Figure 2.3: Diagram of the NuMI beamline. The 120 GeV protons incide on a graphite target and produce a secondary hadron beam, which is then focused by two magnetic horns. The hadrons are allowed to decay into a long pipe, producing the neutrino beam. Other decay products (predominately muons) are stopped right after the decay pipe by the beam dump or the rock. The configuration of the target and horns can be changed to alter the characteristics of the beam. From [37].

process, pairs of batches are injected into the Recycler and then merged to form double-intensity batches. The Recycler is capable of slip-stacking up to twelve batches, which results in six double-intensity batches for extraction to MI [38]. With this 6+6 slip-stacking, the Main Injector can deliver $\sim 5 \times 10^{13}$ protons to the NuMI target in $10 \mu\text{s}$ spills, at an average of 1.33 s [38]. The nominal expectation for one year of running is 6×10^{20} protons on target (POT) delivered to NuMI. We will use POT as the unit of neutrino beam intensity.

Figure 2.3 is an overview of the elements of the NuMI beamline. The protons of 120 GeV momentum are extracted from the Main Injector and directed toward a 1.2 m graphite hadron production target. Two toroidal magnets called “horns” sign-select and focus the secondary mesons, and send them towards an evacuated decay pipe where they may decay to muons and neutrinos. At the end of the decay pipe, an absorber and then 240 m of un-excavated rock stop the remaining hadrons and the muons.

Most of the neutrinos in the NuMI beam result from $\pi^\pm \rightarrow \mu^\pm + \nu$. The contribution from K^\pm and K_L is smaller because the kaon production cross sections in the target are lower, and the mass differences produce a more diffuse neutrino beam than the pions. A small contamination of electron neutrinos arises from muons that subsequently decay as e.g. $\mu^+ \rightarrow \bar{\nu}_\mu + e^+ + \nu_e$, and from




Figure 2.5: Diagram of the magnetic focusing horns in the NuMI beamline. The magnetic horn system acts as a lens with a focal length proportional to the momentum of the charged particles. Changing the direction of the current reverses the magnetic field and in consequence the sign of the focused hadrons and the type of neutrinos produced. From [39].

rare leptonic and semileptonic decays of pions and kaons. This is denoted as the “intrinsic beam ν_e component”.

Depending on the direction of the current in the magnetic focusing system, the hadrons that are directed into the decay pipe will have positive or negative charges, producing a beam of muon neutrinos or antineutrinos. The configuration that leads to a ν_μ ($\bar{\nu}_\mu$) beam is called Forward (Reverse) Horn Current, or FHC (RHC) (see Figure 2.5). Hadrons that leave the target very parallel to the beam direction might not be deflected, and therefore a small fraction of particles with the opposite charge will make it to the decay pipe. The resulting neutrinos are called “wrong-sign contamination”, i.e. the fraction of $\bar{\nu}_\mu$ in a ν_μ beam and vice-versa.

2.2.1 Off-axis design


The location of the detectors in the path of the beam will determine the neutrino flux that is observed. Let us consider a neutrino resulting from $\pi \rightarrow \mu + \nu$. From two-body decay kinematics, we know that for a small angle θ between the parent pion and the neutrino in the lab frame of

reference, the energy of the neutrino E_ν is given by [31]:


$$E_\nu = \frac{2E_\nu^*}{m_\pi} \frac{E_\pi}{1 + (\gamma_\pi \theta)^2} \simeq \frac{0.43E_\pi}{1 + (\gamma_\pi \theta)^2}, \quad (2.1)$$

where $E_\nu^* = (m_\pi^2 - m_\mu^2)/2m_\pi$, m_π (m_μ) is the pion (muon) mass, E_π is the pion energy, $\gamma_\pi = E_\pi/m_\pi$. The expression for $K \rightarrow \mu + \nu$ is identical to (2.1) except that 0.43 is replaced by 0.96. For a detector of area A located at a distance z from the decay point, the flux per decay Φ is given by:

$$\Phi = \left(\frac{2\gamma_\pi}{1 + \gamma_\pi^2 \theta^2} \right) \frac{A}{4\pi z^2}. \quad (2.2)$$



(a) Neutrino energy as a function of the parent pion (top) or kaon (bottom) energy for various off-axis angles. These correspond to (2.1). The shaded bands represent the energies of interest for NOvA.



(b) Simulated neutrino event rate in the absence of oscillations at a distance of 810 km and various off-axis angles. Bottom: the $\nu_\mu \rightarrow \nu_e$ oscillation probability as a function of energy for the same distance.

Figure 2.7: Neutrino energy and event rate for different off-axis locations.

The location of the NOvA far detector was decided based on available sites in the USA in the path of the NuMI beam that would optimize the measurements in the electron neutrino appearance channel. As shown in Figure 2.6b, at 14 mrad most pion decays result in neutrinos with $E \sim 2$ GeV,

and therefore the NOvA far detector can measure a narrow band beam peaked near the maximum of the $\nu_\mu \rightarrow \nu_e$ oscillation probability. Notice in Figure 2.6a the different behavior of the neutrino energies as a function of pion vs. kaon energy at 14 mrad. While the pion curve reaches a maximum and decreases slowly in the 1-2 GeV range, the kaon curve increases monotonically. Combining this with the lower production rate of kaons in the target, we expect the peak in Figure 2.6b to be dominated by neutrinos with pion ancestors, and the small kaon ancestor component to remain small and to dominate the high energy tail.

2.3 The NOvA detectors




Figure 2.8: Diagram showing the sizes of the NOvA detectors and the main detection components. From [40].

The NOvA detectors are designed to be functionally equivalent and optimized for the observation of ν_e CC interactions. The two-detector design ensures identical signal and background efficiencies, and allows for the reduction of systematic uncertainties due to neutrino flux and cross-section modeling [31]. While the two detectors have similar material, segmentation, and orientation, the sizes are different: the far detector is 14 kton to maximize statistics at the 810 km baseline, while the near detector with only 300 ton observes a higher flux due to its proximity to the neutrino source.




Figure 2.9: A schematic of the NOvA detectors, showing the alternating plane orientations. Horizontal and vertical planes provide the side and top views, and combined offer 3D reconstruction of the particle trajectory. From [36].

Figure 2.8 presents an overview of the detectors and their components. The fundamental elements are long cells of highly reflective polyvinyl chloride (PVC) that are filled with liquid scintillator (mineral oil mixed with 5% pseudocumene). The light produced in a cell as particles travel through the detector is collected with a loop of a wavelength-shifting (WLS) optical fiber. The ends of each fiber are connected to a single pixel of an avalanche photodiode (APD), which amplifies and digitizes the light.

Groups of 32 parallel cells are sealed together to form planes, which are arranged alternating between vertical and horizontal orientations to allow for 3D event reconstruction, as shown in Figure 2.9. By mass, the detectors are composed of 63% active material. High granularity and low Z (~ 0.15 radiation lengths per layer) of the detector are two of the requisites introduced in Section 2.1.

2.3.1 Far Detector (FD)



Figure 2.10: Photograph of the NOvA far detector. From [36].

The FD cells have a cross section of 3.9×6.6 cm in cross section and 15.5 m long, where the 6.6 cm side is measured in the direction of the neutrino beam. The FD comprises 896 planes with 344,064 channels and a total mass of 14 kton.

The primary goal of the far detector is the measurement of the oscillated neutrino spectra in the 1-3 GeV energy range, which is sensitive to $\nu_\mu \rightarrow \nu_e$ oscillations. As mentioned in Section 2.2, this is the reason why the FD is located 14.6 mrad away from the central axis of the NuMI beam and 810 km away from the beam target. The FD is on the surface of the Earth, so it is exposed to a high cosmic ray flux that is partially mitigated by an overburden of 1.2 m of concrete and 15 cm of barite. NOvA uses in-situ measurements to estimate the background to the beam oscillation analysis, and separate periodic minimum-bias triggers to collect cosmic data that are used for calibration and training of algorithms.

2.3.2 Near Detector (ND)



Figure 2.11: Photographs of the NOvA near detector. Left: the front of the detector, closest to the beam target. Right: the most downstream section of the ND is the muon catcher, composed of steel plates alternating with scintillator planes, and only two-thirds of the height of the active region. From [36].

Most of the ND and FD cells are identical except for their length (3.9 m in the ND). The near detector is much smaller than the far detector (15.9 m along the beam direction), so the last 3.23 m are an additional component with higher density, the “muon catcher”, designed to range out muons. The muon catcher comprises 11 pairs of horizontal/vertical scintillator planes separated by ten steel planes, 10 cm-thick. The vertical planes in the muon catcher are only 2.6 m high. There are 214 planes in the ND for a total mass of 290 ton, 130 ton of which are liquid scintillator.

The primary goal of the ND is to characterize the neutrino beam before oscillations. The ND is centered at the same off-axis angle as the FD to maximize the similarity of the neutrino flux between the two detectors. Note that the ND is only 1 km from the neutrino source, so it receives a higher intensity beam with a larger angular spread than the FD. The ND is 100 m underground, thus the cosmic background is negligible.

2.3.3 Data acquisition




Figure 2.12: Diagram of the NOvA Data Acquisition system (DAQ) in the FD. The ND system is analogous, with a smaller number of channels.

Figure 2.12 is a diagram of the data acquisition system (DAQ) in NOvA, divided in three domains: detector, trigger, and analysis. In the detector, each cell is one channel connected to one of 32 pixels in an APD. Each APD continuously sends data to a Front End Board (FEB) that discriminates and adds a time stamp to the signals. The FEBs transmit signals that are above threshold to Data Concentrator Modules (DCM). Each DCM collects information from 64 FEBs during $50 \mu\text{s}$ windows, called “microslices”. Each microslice is a single data packet that is sent to buffer nodes until the trigger system decides if they have to be recorded or rejected, for a maximum of 16 minutes.

Even though the electronics is permanently sending data, the high cosmic ray rate in the FD (148 kHz) and the big number of channels means that the volume of data produced is higher than 1 GB/s. To keep the most relevant data, the triggering system uses either i) external signals, ii) periodic (clock) triggers, or iii) the results of real-time algorithms, or “data driven trigger” [41], and builds events of a requested length based on the time stamps of the microslices. For this

analysis, we are only interested in the NuMI beam spill signal and the independent minimum-bias cosmic trigger. Both of them build $550\mu s$ events, with the former always positioning the microslices on-time with the beam spill in the 218-228 μs time relative to the beginning of the trigger event.

A timing system [42] synchronizes and provides a time base to all the DAQ components. A series of Time Distribution Units (TDU) connect to the DCMs and ensure that all the channels of the detector are synchronized to the same local time. Additionally, the TDUs maintain the global time synchronized with an external GPS time, that will be consistent between the detectors and the accelerator at Fermilab [43].

2.4 Simulation

The experiment relies on detailed simulations for the calibration of the detectors and data analysis. The main stages of the simulation chain are sketched in Figure 2.13.




Figure 2.13: Flowchart of NOvA simulation chain. The blue squares represent collections of tools for each stage, and the green circles represent the resulting data products. From [44].

First, we use a detailed model of the beamline geometry of Section 2.2 and GEANT4 [45] to simulate the hadron production, transport, and decay that produce the neutrino beam. The output

are flux files that contain the flavor, energies and directions of neutrinos, as well as the information of the ancestors that produced them. This neutrino flux is corrected using constraints from hadron production experiments, using a suite of tools created for the NuMI beam called PPFX [46]. Table 2.1 shows simulation predictions of the beam compositions at the Near and Far Detectors in the absence of oscillations; the full predicted spectra from 0-20 GeV are given in Fig. 2.14.

Table 2.1: Predicted beam composition in the 1 to 5 GeV neutrino energy region in the absence of oscillations. As published in [47].

Component	ND (%)	FD (%)
ν_μ	93.8	94.1
$\bar{\nu}_\mu$	5.3	4.9
ν_e and $\bar{\nu}_e$	0.9	1.0




Figure 2.14: Predicted PPFX/-based composition of the NuMI beam in the FHC beam mode (neutrino-enhanced) at the ND (left) and the FD (right). Curves from top to bottom: ν_μ , $\bar{\nu}_\mu$, ν_e , $\bar{\nu}_e$. The Table 2.1 gives the fractional composition for each neutrino flavor integrated from 1-5 GeV. From [33].

Second, we simulate neutrino interactions in the detectors and their surroundings with GENIE [48, 49], using the predicted flux and detector geometry files as input. GENIE uses neutrino cross-section data and theoretical models to determine whether a neutrino interacts with a material at some vertex position, the type of interaction, and the kinematics of the outgoing particles. Similar to the beam stage, we need to add corrections to the neutrino interaction simulation to compensate

for the deficiencies of the default configuration. First, we activate a type of interactions with nucleon-nucleon pairs denoted Meson Exchange Current (MEC) using a semi-empirical model [50] that is optional in GENIE, and increase their rate by an extra 20% as suggested by fits to NOvA ND data [51]. The effect of long-range nuclear charge screening as predicted in the Random Phase Approximation (RPA) [52, 53] are included by re-weighting quasi-elastic interactions, and the rates of ν_μ CC non-resonant single pion production are reduced following [54].

Third, we simulate the propagation and interaction with the detector elements of primary and secondary particles that resulted from the neutrino interactions using GEANT4. This step generates energy depositions of particles in the detector cells.

Finally, custom NOvA algorithms evaluate the transport of the light signal in the fibers, the collection in the APDs, and the response of the FEBs to the APD signals. At the end of this process, we obtain files that mimic the raw detector data, thus allowing us to apply the same algorithms to data and MC as will be described in the next chapter.

The computing scheme followed by NOvA is common in high energy experiments [55], where the data and simulated events are organized in files that are shared among the collaboration and safely stored, and they can be accessed remotely and using grid computing clusters for analysis [56, 55]. Similarly, the software is shared and centralized, based on ROOT [57] and additional tools common to many Fermilab experiments [58, 59, 60].

2.5 Data analyzed

The NOvA far detector was constructed starting in 2012 with the first out of 14 modules. The final stages of construction and commissioning were between February and September 2014. During that time, the NuMI beam was running, and data was taken with a partial detector. The data used for this thesis were recorded between February 6, 2014 and February 20, 2017, and correspond to the equivalent of 8.85×10^{20} protons on target (POT) in the full NOvA Far Detector with the NuMI beam in neutrino mode (FHC). The data collection over time is presented in Figure 2.15.




Figure 2.15: Daily protons on target (POT) delivered by NuMI (black) and recorded by NOvA (red). The lines are traced over a 28 day average. This thesis uses the 2017 dataset, taken during February 2014 - February 2017, corresponding to 8.85×10^{20} POT-equiv. of beam in neutrino mode (FHC). The collection of antineutrino data was tested in July 2016 and it is not included in this analysis. From [33].

CHAPTER 3. OSCILLATION ANALYSIS FRAMEWORK

3.1 Analysis strategy




Figure 3.1: Flowchart of the NOvA oscillation analysis strategy. The top blocks (red) represent the steps of the ν_μ analysis, and the bottom (blue) blocks represent ν_e . Near detector data is used in both analyses to constrain the far detector predictions from the simulation. Through the comparison of FD data to the prediction Far detector predictions, we can impose constraints on the neutrino oscillation parameters. From [61].

In the previous chapter we introduced NOvA’s goals to characterize neutrino oscillations in the $\nu_\mu \rightarrow \nu_\mu$ and $\nu_\mu \rightarrow \nu_e$ channels, using Fermilab’s NuMI muon neutrino beam and the two-detector technique. We will now discuss how we use ND and FD data and the simulation to produce the measurement. This chapter covers the general strategy and tools, while the following two chapters contain specific applications to the ν_μ and ν_e analysis.

The flowchart in Figure 3.1 shows the succession of processes followed for data analysis. The main categories can be described as follows:

- **Event reconstruction:** The raw data collected in the detectors in the form of individual “hits”, are grouped together and processed to construct event candidates. Vertex positions, directions, visible energy, and other quantities can be estimated. By using the same algorithms in the simulation, where “true” information is available, we can assess their efficiency. This is discussed in Section 3.3.
- **Particle identification:** Using the topological characteristics of the activity in the cells associated with a reconstructed event, we can create discriminants that indicate whether such activity is likely to result from specific types of particle interactions. These discriminants are essential to identify analysis candidates and establish a pure sample for the estimation of oscillation parameters. In Section 3.4 we review the main algorithms used for this classification.
- **Energy reconstruction:** Oscillation probabilities depend on the energy of the interacting neutrino. Therefore, we need an energy variable reconstructed from the activity of the outgoing particles in the detector that can be a reliable estimate of the original neutrino energy. As part of this process we consider two components: i) detector calibration, ii) and energy tuning depending on the particle hypotheses; both are covered in Section 3.5.
- **Event selection:** In order to isolate analyzable samples of ν_μ or ν_e CC candidates, we require a series of selection cuts based on the event’s timing, beam conditions, reconstruction quality, containment, cosmic rejection and particle classifier score. These are optimized independently and detailed in Section 4.1 (ν_μ) and Section 5.1 (ν_e).
- **Extrapolation:** The ND data can be used to improve the predicted distributions of FD events. In NOvA, this process is referred to as a “Near to Far extrapolation”. A basic description of extrapolation is presented in Section 3.1.1, while the specific implementations for ν_μ and ν_e are in Section 4.3 and Section 5.3.

- **Cosmics:** Because the NOvA far detector is located on the surface, cosmogenic interactions are an important background. The cosmic background estimates for ν_μ and ν_e are given in Sections 4.4 and 5.4, respectively.
- **Fitting:** A comparison between the selected FD data with the predictions at given oscillation parameters allows us to construct confidence intervals. Additional considerations are necessary given the relatively low number of events observed and the physics of oscillations. The essential procedures used to interpret the data are explained in Section 6.1, while sensitivities and results will be presented in Chapters 6 and 7.
- **Systematics:** The effect of systematics is assessed by repeating the steps in Figure 3.1 using modified versions of the simulation. Details on the flux, neutrino interaction, and detector uncertainties considered are given in Section 3.6. The impact on the ν_μ and ν_e predictions will be presented in Sections 4.5 and 5.5, and the impact on the measurements in Section 6.2.1.

3.1.1 Decomposition and extrapolation

Two technical terms included in the flowchart Figure 3.1, “decomposition” and “extrapolation”, are used extensively in the context of the NOvA oscillation analyses so we will introduce them first. They refer to the strategies used to combine measurements in the near detector to construct simulated predictions in the far detector.

3.1.1.1 Decomposition

“Decomposition” is a procedure to assign selected events in the ND data to the neutrino flavor components that are present in the simulation. A clear example is the ν_e decomposition, which will be discussed in detail in Section 5.3. The ν_e candidates according to the Monte Carlo correspond to three categories: ν_e CC, ν_μ CC, NC, all in a similar, non-negligible proportion. The measured ND data are inclusive, so if there is any excess or deficit we have to try to interpret it as a mis-modeling of the ν_e CC, ν_μ CC, NC, or all the interaction types. Three ways to approach this problem are:

- Assign the discrepancies to a single MC component (“ ν_μ decomposition”, “ ν_e decomposition”, “NC decomposition”).
- Assign the discrepancies to all three components, assuming that the MC correctly predicts the relative proportions but not the absolute normalization in each analysis bin (“proportional decomposition”).
- Add independent samples that provide more information, thus reducing the degeneracy, and assign the discrepancies accordingly (“data-driven decomposition”).

The first kind of decomposition is used for the ν_μ ND selection, since the fraction of background components according to the MC is negligibly small (see Section 4.3). The second and third options are used in the ν_e analysis, and in particular a data-driven decomposition is used for the official results (details in Section 5.3).

The results of the decomposition can be interpreted as estimates of each component α from the ND data:

$$\left[\text{ND}^{\text{Data}}\right]_{\text{reco}} = \sum_{\alpha} \left[\text{ND}_{\alpha}^{\text{Pred}}\right]_{\text{reco}}. \quad (3.1)$$

If they differ from the simulated distributions, that means that the corresponding type of interaction in the FD is probably mis-modeled as well, and thus they are used as corrections.

3.1.1.2 Extrapolation

The process of using decomposed ND data to create an improved FD prediction is called “extrapolation”. In NOvA, we define two types of extrapolation: “truth” and “reco” extrapolation. The names refer to the type of variables used to connect the ND observation with the FD simulation. Recall that, by definition, *true* information is only available for the simulated events, while reconstruction is done identically for data and MC. The choice between the two types of extrapolation is made based on the characteristics of the ND sample.

- **Truth extrapolation:** the decomposed ND data histogram using bins of reconstructed energy is transformed to bins of true energy using a 2D matrix from the ND MC:

$$\left[\text{ND}_\alpha^{\text{Pred}} \right]_{\text{true}} = \sum_{\text{reco}'} \frac{\left[\text{ND}_\alpha^{\text{Pred}} \right]_{\text{reco}'} \times \left[\left[\text{ND}_\alpha^{\text{Pred}} \right] \right]_{\text{reco}', \text{true}}}{\left[\text{ND}_\alpha^{\text{MC}} \right]_{\text{reco}'}}. \quad (3.2)$$

The result can be multiplied by the far/near true energy ratio to get a “corrected” FD distribution. We can apply oscillation weights and transform back to reconstructed using a 2D true-to-reco matrix from the FD MC. The result is the corrected FD prediction in bins of reconstructed energy that can be compared to the FD data.

$$\left[\left[\text{FD}_{\alpha \rightarrow \beta}^{\text{Pred}} \right] \right]_{\text{true, reco}} = \frac{\left[\text{ND}_\alpha^{\text{Pred}} \right]_{\text{true}} \times \left[\left[\text{FD}_{\alpha \rightarrow \beta}^{\text{MC}} \right] \right]_{\text{true, reco}}}{\left[\text{ND}_\alpha^{\text{MC}} \right]_{\text{true}}}, \quad (3.3)$$

$$\left[\text{FD}_{\alpha \rightarrow \beta}^{\text{Pred}} \right]_{\text{reco}} = \sum_{\text{true}} \left[\left[\text{FD}_{\alpha \rightarrow \beta}^{\text{Pred}} \right] \right]_{\text{true, reco}} \times \left[P_{\alpha \rightarrow \beta} \right]_{\text{true}}. \quad (3.4)$$

Since we use transformation matrices between reco and true energies, this extrapolation can only be used for components that have a well-behaved energy estimation, or for appearance signals where the flavor of events selected in each detector is different (thus having different selectors and energy estimators, i.e. reco' vs. reco in the equations above). For the present analysis, we use truth extrapolation for the $\nu_\mu \rightarrow \nu_\mu$ and $\nu_\mu \rightarrow \nu_e$ CC signal (Section 4.3 and Section 5.3.1 respectively).

- **Reco extrapolation:** For components where there is no change in flavor between detectors, or where the energy estimator is not expected to perform very well, only the E_{reco} information from the ND is used. The decomposed ND data in bins of reconstructed energy is weighted by the Far/Near ratio with the same binning, to give a corrected FD prediction. We only use the FD reco-to-true matrix to apply oscillation weights where necessary.

$$\left[\left[\text{FD}_{\alpha \rightarrow \beta}^{\text{Pred}} \right] \right]_{\text{true, reco}} = \frac{\left[\text{ND}_\alpha^{\text{Pred}} \right]_{\text{reco}} \times \left[\left[\text{FD}_{\alpha \rightarrow \beta}^{\text{MC}} \right] \right]_{\text{true, reco}}}{\left[\text{ND}_\alpha^{\text{MC}} \right]_{\text{reco}}}, \quad (3.5)$$

$$\left[\text{FD}_{\alpha \rightarrow \beta}^{\text{Pred}} \right]_{\text{reco}} = \sum_{\text{true}} \left[\left[\text{FD}_{\alpha \rightarrow \beta}^{\text{Pred}} \right] \right]_{\text{true, reco}} \times \left[P_{\alpha \rightarrow \beta} \right]_{\text{true}}. \quad (3.6)$$

The reco extrapolation is used for the ν_e background prediction (Section 5.3.2).

It is worth highlighting that all the steps of the extrapolation procedure including the 2D migration matrices are re-evaluated for each systematic variation, as will be explored in latter sections.

In general, very small background components cannot be constrained using any data-driven methods, so we will use the prediction given directly by the simulation (or “not extrapolated”).

Finally, we should note that the extrapolation of ND data is most useful for small discrepancies ($\mathcal{O}(10\%)$) with the MC. Very large differences would indicate a problem in the simulation, that should be addressed early in the analysis (generating new MC, or reweighting events) and long before the unblinding of FD data.

3.2 Event topology and reconstruction goals

Recall from Chapter 2 the basic structure of the NOvA acquired data: ADC counts (generally 4 numbers), a time stamp, and the plane and cell numbers. A “trigger event” includes the full readout for a given type, and sets of trigger events are grouped together in files, further labeled by run and subrun numbers that can also be linked to the beam and detector condition databases. We are mostly interested in the NuMI trigger events, the corresponding files from the simulation, and the independent cosmic trigger data. The overall goal of the reconstruction chain is to transform these figures into an useful format, and add higher-level variables (e.g. energy deposition, vertex positions, particle classification scores) that can be used in a consistent manner in the analyses.

Figures 3.2 and 3.3 show event displays for one full readout of the NuMI trigger in the ND and FD. While the beam spill is $10\ \mu\text{s}$ long, a longer window of $550\ \mu\text{s}$ is stored for data quality and background estimation. Since the ND is close to the beam source, an important consideration is pileup. We expect around 4 or 5 neutrino interactions per NuMI trigger. The FD observes very few events, not only from the $1/R^2$ reduction of the flux at the 810 km baseline, but also from the effect of neutrino oscillations. Thus, at most one neutrino event is expected per trigger. Instead of signal overlap, the concern is the 50-70 cosmic events that are detected during the beam spill window.




Figure 3.2: $550 \mu\text{s}$ readout of one NuMI trigger event in the NOvA near detector. The neutrino beam spill window occurs in the range $218\text{--}228 \mu\text{s}$ and arrives from the left. Color represents ADC charge. The detector is underground and 1 km away from the source. Most of the activity recorded is from neutrino beam events; the beam peak can be observed in the hits vs. time histogram.

Both detectors can also present cell activity due to noise, and events that originate outside of the detector boundaries.

The two NOvA detectors use the same basic technology, therefore the topologies of neutrino interactions are very similar in both of them. Figure 3.4 presents three simulated interactions with the same neutrino energy, where the outgoing particles are a proton with 0.78 GeV momentum, and a second particle with 1.86 GeV . The outgoing muon and electron are characteristic of the ν_μ and ν_e CC interactions, while the π^0 decaying to two photons is a typical NC background to the ν_e CC signal.

Typically, muons leave long straight tracks in the detector, and when they are fully contained, the time-delayed decay (Michel) electron might be found. Electrons, on the other hand, produce shorter electromagnetic showers. In both cases, low-range hadronic activity can be identified as emanating from the vertex along with the depositions by the charged lepton. In neutral current interactions, the outgoing lepton is a neutrino and cannot be used to tag the event; furthermore, it carries energy that is not visible to the experiment. In the three-flavor oscillation with a reli-




Figure 3.3: 550 μs readout of one NuMI trigger event in the NOvA far detector. The neutrino beam spill window occurs in the range 218-228 μs and arrives from the left. Color represents ADC charge. The detector is on the surface of the earth and 810 km away from the beam source. Most of the activity recorded is from cosmic rays, covering the full time range.

able energy estimation to be used in the analysis, we consider neutral current events one of the backgrounds for both the ν_μ disappearance and ν_e appearance channels. NC events containing a charged pion will occasionally mimic the signature of ν_μ CC events, but these are generally easy to reject through particle identification algorithms (this will be shown in Chapter 4). The difference between NC and ν_e CC is generally harder to spot, except in cases like the one presented in Figure 3.4 where the NC has an identifiable gap between the vertex and the start of the shower (corresponding to the distance that photons travel before converting into electron/positron pairs that produce scintillation light).

The goal of reconstruction in NOvA is to isolate and label neutrino interactions like the ones in Figure 3.4 from the full readouts (e.g. Figure 3.3), to provide pure samples with a reliable energy estimation to be used in the analysis. The following three sections will describe the ingredients of the full reconstruction: event reconstruction, particle identification, and energy estimation.

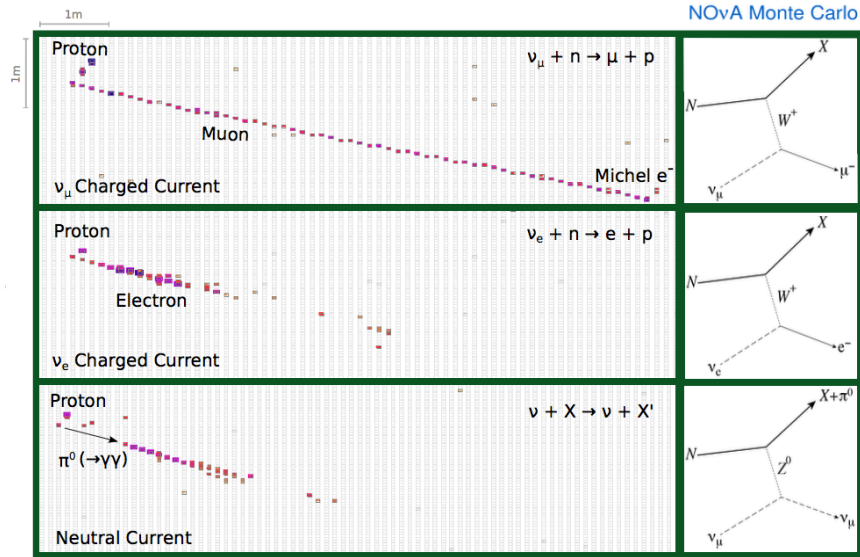


Figure 3.4: Simulated neutrino interactions in the NOvA FD: ν_μ CC (top), ν_e CC (middle), and NC (bottom). These three interactions result from a 2.15 GeV neutrino, and produce a 0.78 GeV proton plus another particle. Cells are colored by the charge deposited. The photon conversion distance in NOvA is ~ 38 cm (6 plane widths). From [62].

3.3 Event reconstruction

Figure 3.5 presents the NOvA event reconstruction chain. All the trigger events, both data and simulated, will go through most of these steps, and new “data products” are progressively added. At the end of the chain, a given event has been processed in different ways, and the associated file has collated all the information that can be needed by different users, e.g. for calibration, oscillation analysis, cross-section studies (not pictured). Figure 3.6 contains an example of such processes applied to a single event. The stages of the reconstruction are:

- **Raw hits:** The original readout from the detectors includes an identification of the pixel (plane, cell), a time stamp, and the ADC charges (2 or 4 digits). These are grouped in data files, identified by detector, run and subrun number, and trigger. This is also the type of output that is provided by the simulation, such that the reconstruction chain proceeds almost identically for data and MC (the simulation has *true* information that can be mapped to reconstructed objects).




Figure 3.5: Diagram of the NOvA event reconstruction chain. The green boxes represent processes that manipulate the information in the data and construct physically meaningful groupings and variables. The gray boxes at the bottom indicate the relevant applications of the data products. From [63]

- **CalHits:** The raw data is transformed into the *art* framework format [58] to facilitate the addition and correlation of new data products as they are added, software compatibility, and record keeping. Information about calibration, geometry, beam and detector conditions can also be accessed and incorporated. In cases where there are no analyzable data, the events are not further reconstructed. Detector “signal” and “noise” hits are indistinguishable so far.
- **Slices:** Collections of hits are clustered together to form physics event candidates. They are separated using a custom density-based clustering algorithm (Slicer4D [64] based on DBSCAN [65]) that uses the 4D distances between hits as input. Since the incident particles are traveling close to the speed of light, hits must be in light-like intervals (although limited by the time resolution). Large distances in the z direction (along the beam axis) and x or y (for hits in the same view) are penalized. Hits that are in or neighboring high-density regions are considered part of the slice, while isolated hits are labeled as noise.




Figure 3.6: An example of the reconstruction chain showing step by step algorithms applied to a NC background. From [62].

“Slices” constitute the essential reconstructed object in NOvA, and often the terms “slice” and “event” are used interchangeably.

- **Cosmic tracks:** The simplest signature of a cosmogenic event is a single vertical track that goes through the detector. The cosmic tracker [66] applies a sliding window tracking algorithm assuming these downward-going single-track topologies. This is designed to be a fast way to process single-particle events with minimal reconstruction, in particular the cosmic muons used in calibration (Section 3.5.1); the output will not be useful for e.g. beam signal events, since a single almost-straight line is a poor approximation to multi-particle depositions.
- **Kalman tracks:** Charged particles that do not produce showers in the detector, in particular muons from ν_μ CC interactions, lose energy via ionization, and form narrow tracks with occasional small changes in direction as the particle scatters [67]. NOvA uses a Kalman-filter algorithm [68] to reconstruct the particle trajectory, by iteratively collecting hits that are collinear within a few planes or only present small angular deviations. Tracks are initially reconstructed separately for each 2D view, and later matched in 3D based on the overlap in the z -direction (common to both views).
- **Vertices:** Vertex identification in general is a critical and computationally expensive reconstruction step. In NOvA, it has two components. First, two-dimensional guidelines are reconstructed via a Multi-Hough transform [69]. This pattern recognition algorithm takes pairs of pixels in each detector view and constructs a two-dimensional array, or Hough space map, using polar coordinates and a Gaussian-smearred vote [62]. Peaks above threshold in the Hough map are line candidates. To control the number and quality of the line candidates, pixels that have been associated with a dominant candidate are removed, and the process is iterated. Second, the Elastic Arms algorithm [70] uses the Multi-Hough outputs as seeds in order to find vertex candidates. A typical NOvA event has a series of particle trajectories emanating from a common point, where the neutrino interaction occurred. Starting from the intersections of Hough lines, a series of “arms”, which are vectors pointing away from

the vertex, are constructed and adjusted to better represent the trajectories. The optimum vertex is found via a minimization of an energy cost function [62]. The NOvA application of Elastic Arms includes additional terms to allow gaps between the vertex and first energy deposition that is characteristic in the NC processes [71].

- **Prongs:** A prong is a collection of hits that is associated with a single particle candidate emerging from the vertex, in the form of tracks or showers [72]. This assignment is realized with a possibilistic fuzzy-k means algorithm [73, 74]. “Possibilistic” means that the hits are not required to be 100% assigned to prongs (i.e. noise), and “fuzzy” means that a hit can belong to several prongs, and the total number of prongs is initially unknown. This algorithm uses the angular distribution of energy depositions with respect to the vertex to find peaks that would correspond to different prong *centers*. Hits are then assigned a *degree* of membership to each prong based on their distance to the center. The process is iterated avoiding duplicate prong centers, and ends when all cells belong to a prong or a maximum number of clusters has been identified. Finally, these two-dimensional prongs are matched based on their geometry and energy distribution using a Kupier metric [62].

3.4 Particle identification

3.4.1 Convolutional visual network (CVN)

The Convolutional Visual Network (CVN) [75] is a deep-learning classifier that applies a series of linear operations to extract abstract features from each event. The input to this convolutional neural network are the *slices* introduced in Section 3.3, where recorded hits have already been grouped in time and space, but no other reconstruction has been applied. CVN is based on techniques from the field of computer vision for image recognition [76, 77, 78]: the hit maps are treated like images (“pixel maps”), and each event is associated with two maps corresponding to the two detector views. These are processed separately, and only at the final layers of the network combine the information. CVN was developed using the Caffe framework [79], with an architecture

inspired by GoogleNet [77]. Figure 3.7 shows a diagrammatic view of the network architecture, and Figure 3.8 presents examples of partial outputs. The network is trained over a sample of simulated beam events and cosmic ray data [76]. From the output of the network, a multilayer perceptron [80, 81] maps these features into a normalized score that classifies events as more or less likely to be ν_e CC, ν_μ CC, ν_τ CC, NC or cosmic. Figure 3.9 compares the scores for simulated ν_e CC or ν_μ CC candidates.




Figure 3.7: Schematic view of the CVN architecture. Convolutional layers apply learned kernels to the image and the output is the same size of the input, while pooling layers reduce the size of the image. The local response normalization (LRN) layers compare and adjust the responses of adjacent kernels so that important features remain the most pronounced. From [82].

3.4.2 Reconstructed Muon Identifier (ReMId)

ReMId is a k-Nearest Neighbor (kNN) [83] classifier used to identify muon candidates among the different particles in an event [84]. It uses the Kalman track candidates introduced in Section 3.3, and scores them based on four reconstructed variables: dE/dx likelihood, scattering likelihood,

total track length, and non-hadronic plane fraction. dE/dx refers to the particle's energy deposition along its path; the expected dE/dx profile of a muon is different of the one for a pion, which can lose energy through hadronic scattering in addition to ionization. Thus the first variable compares the dE/dx distribution against the pion or muon hypotheses using only planes with minimum contamination from other activity, and returns a likelihood score. Similarly, the scattering likelihood considers the deviations in the particle trajectory from a straight line. Since the NOvA detectors are not magnetized, muons can only acquire a slight curvature through multiple soft-scattering and occasionally hard Coulomb scattering. Pions experience the same processes, plus hadronic interactions through the strong force. The third variable, track length, considers the fact that hadronic showers tend to contain multiple short tracks as opposed to the single, longer muon tracks. Finally, the non-hadronic plane fraction counts the planes that had energy outside of the track that were excluded from the dE/dx likelihood calculation. It is expected that muon tracks have almost no contamination away from the vertex, which is the opposite of the pion tracks as they are part of the hadronic shower [85].

The kNN algorithm uses a training set of simulated signal and background events, and determines the four-dimensional distance between the Kalman tracks in the candidate event and the simulated tracks. This way, a track that is *closer* to signal-like is assigned a higher score. The most muon-like of the tracks is taken as the muon candidate, and this score is used for the entire event.

Figure 3.10 shows the distributions of the four variables used in the training of ReMID and the output score in a selected sample of contained ν_μ candidates. ReMID shows excellent separation of the NC background from the ν_μ CC signal in the FD: for a ReMID score of 0.70, the efficiency of the selection is 81% and the purity is 98% [85]

3.4.3 Other particle identifiers

The following are additional particle classification algorithms that play smaller roles in the analysis:

- **CVN prong identification:** A second application of CNNs in NOvA is labeling individual prongs within an event [86]. One of the difficulties of this type of classifier is the reliance on other reconstruction algorithms for the assignment of hits (of fractions thereof) to separate prongs. The architecture of ProngCVN is analogous to the classifier from Section 3.4.1 using four views (two for the full event, two for the prong), and outputs the labels: electron, photon, muon, pion, proton [82]. In the context of this thesis, ProngCVN is only used as part of the ν_e energy estimation in Section 3.5.3, to differentiate the electromagnetic component from the rest of the deposited energy.
- **ν_μ cosmic rejection BDT:** The ν_μ analysis uses a boosted decision tree (BDT) [87, 88, 89, 90] for increased cosmic rejection [91]. There are seven reconstructed variables included in the training: angle between the lepton and the neutrino, length of the muon track, verticality of the muon, distance from the track start or stop to the top of the detector, distance of the track to the front or back of the detector, fraction of hits in the track, and the CVN cosmic score. The BDT is trained with cosmic data, and overlays of MC and cosmic data for the signal, with a pre-selection of ReMId > 0.75 . The combination of cutting on ReMId and training with the CVN cosmic score signals that the algorithm is designed to improve over these selectors, trying to remove the cosmic events that have been deemed as signal-like by the previous classifiers. The portion of the analysis that uses this BDT will be covered in Section 4.1.
- **ν_e peripheral selection BDT:** Similar to the ν_μ BDT, the ν_e analysis uses a BDT to increase the purity of a subsample selected that is prone to cosmogenic backgrounds. This BDT is trained on five variables [92]: number of hits, reconstructed distance from the end of the prong to the top of the detector, minimum distance from the end of the prong to any face of the detector, reconstructed forward momentum p_T/p , and sparseness asymmetry. More information about this selection is presented in Section 5.1.



(a) True ν_e CC event.(b) True ν_μ CC event.

Figure 3.8: Example of CVN feature extraction. Each of the two panels show the y view of a neutrino interaction (bottom left), the 256 feature maps extracted from that event by the end of the first inception module of the CVN network (bottom right), and three highlighted feature maps from the ensemble (top) which appear to have become sensitive to muons (green, top left), electromagnetic showers (blue, top center) and hadronic activity (purple, top right) respectively. From [33].





Figure 3.9: The CVN classifier outputs for simulated samples of ν_e (left) and ν_μ (right) events in the far detector. For the ν_e analysis, appeared ν_e CC are the signal, and beam ν_e are an irreducible background; both are expected to score high in this classifier. For the ν_μ analysis, the survived ν_μ CC candidates are the signal, and NC are the main backgrounds. Note that the signature long tracks that muons leave in the detector generally make them easier to separate from other types of beam events. From [33].



(a) Input variables for the kNN training: dE/dx likelihood (top left), scattering likelihood (top right), track length FD. (bottom left), plane fraction (bottom right). (b) ReMId score for a ν_μ -selected sample in the

Figure 3.10: The Reconstructed Muon Identification (ReMId) kNN identifies ν_μ CC candidates based on the differences between muons and pions in a four-dimensional space. From [85].

3.5 Energy reconstruction

The essential idea behind neutrino oscillation experiments is the comparison of observed event distributions with the expectation in the absence of oscillations. In the case of NOvA, with a fixed long-baseline of 810 km, the measurements rely on probing the dependence of the oscillation probabilities on the true energy. Therefore, one of the most important requirements for the analysis is having a sensible estimate of the energy of neutrino signal candidates.


In practice, energy estimation can be divided in three separate steps. The first step is calibration, where standard candles are used to evaluate and adjust the raw data so that its translation to energy is consistent among hits, along time, and comparable to the simulation. The second step is event reconstruction, where the calibrated hits are grouped together into event and prong candidates. Finally, the third step is the analysis-level construction of the estimator. Energy estimators are not unique and always imperfect, so we need to use knowledge of the signal/backgrounds of interest to make the choices that make more sense for the measurement. Some elements of event reconstruction have been discussed in Section 3.3. In this section, we will briefly describe the calibration of the NOvA detectors, and the energy estimators used in the ν_μ and ν_e oscillation analysis.

3.5.1 Calorimetric energy and calibration

The purpose of calibration is to ensure that the amount of energy deposited and registered in two different detector units can be compared no matter where or when the activity occurred [93], and that later it can be expressed in physically meaningful units. Cosmic ray muon data is used for the calibration of the NOvA FD, as they can provide uniform or well understood energy deposition profiles. There are three main stages in the energy calibration chain:

- **Attenuation calibration:** Recall from Chapter 2 that the NOvA detectors are constructed with narrow cells arranged in planes of alternating directions, so that the combination of the two views can provide three-dimensional reconstruction. However, each individual cell is very long (15.7 m in the FD), and signals occurring away from the readout elements get attenuated as they travel through the fiber. In other words, even if two energy deposits have

the same magnitude, the one furthest from the ADC would be erroneously assigned a lower energy. The attenuation calibration [93] addresses this effect, using cosmic ray muon tracks with clear information in three neighboring cells (e.g. two vertical, one horizontal) so that the hit positions and expected energy output can easily be determined. A fit of a parametric function to the cosmic data will then establish the corrections to be used for any hits in that cell in other data samples [93]. Figure 3.11 illustrates some of these steps. An additional correction is applied to data and simulation to correct biases due to hits that fail the readout threshold or self-shielding effects [94].



(a) Example of attenuation for a single horizontal cell in the FD. The uncorrected number of photoelectrons (PE) collected is higher for hits that occur closer to the readout. The effect of attenuation calibration would be to flatten this distribution.

(b) Comparison of the mean ratios of reconstructed to true energy as a function of length along the cell before (red) and after (blue) calibration. These profiles include all the horizontal (Y-view) cells.

Figure 3.11: The attenuation of the signal as it is transported in the wavelength-shifting fiber inside the cell creates a bias for hits closer to the electronic readout. The calibration corrects the outputs using cosmic rays as standard candles. The position of a hit in a (horizontal) cell, W , is assessed using the neighboring (vertical) cells. In the FD, W is defined in the range $-780 \text{ cm} < W < 780 \text{ cm}$, where the higher values are closer to the readout. From [33].

- **Drift calibration:** This calibration considers changes in detector response over time, including changes in the APD gain, variation of temperature, aging of the scintillator and changes in electronics [95]. These corrections use cosmic events similar to the attenuation calibration,

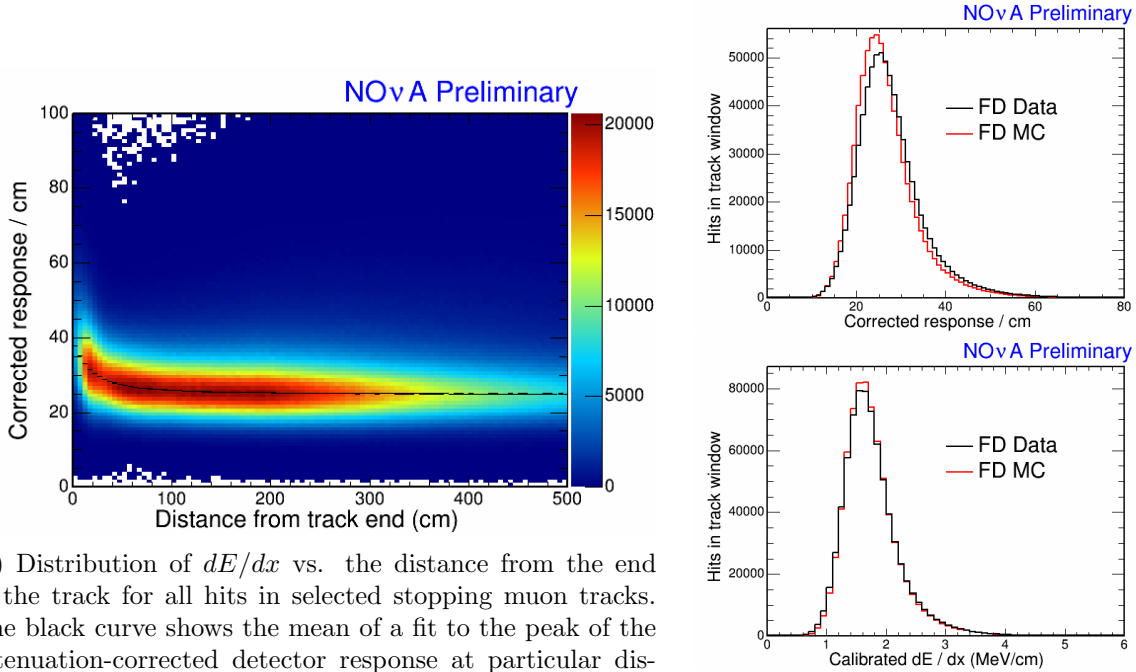
but comparing the variations in response across different periods of data taking. The drift calibrations are applied at the FEB level [96].

- **Absolute calibration:** The last stage of calibration allows us to obtain a physically meaningful energy scale (GeV) from the normalized outputs of the previous steps. The standard candles for this calibration are a subset of stopping cosmic ray muons [94]. The Bethe-Bloch equation predicts the rate of energy loss of the particle as it goes through a material [19]; by comparing the dE/dx profiles in the arbitrary energy units after relative calibration ($PE_{\text{corr}}/\text{cm}$) to the physical expectation (GeV/cm), we obtain the absolute energy associated with a hit. Figure 3.12 illustrate this process. The stopping muons are selected based on the position of the reconstructed end of track, given directly by reconstruction or from Michel tagging ¹ [92]. The selection includes hits that have been calibrated in the previous stages, and that are away from regions with high edge or threshold effects where the attenuation calibration result might be less reliable. Only the deposits in the 1-2 m window away from the end of the muon track, where the dE/dx is close to minimum ionizing, are used to set the energy scale [97].

All the procedures as described above apply only to the far detector. The near detector is underground and cannot have a significant cosmic data sample; calibration is done using beam muons, resulting from the neutrino beam interactions in the detector or the rock surrounding it. The results of the calibration are stored in a database, including all the parameters and the relevant detector coordinates [93]. These are loaded early in the reconstruction chain from Section 3.3, and used in all the steps that follow.

Despite sustained efforts to improve the simulation, reconstruction and calibration, there can still be discrepancies between energy scales in the data and MC [94]. For the current dataset, they are below 5%. These lead to the systematic uncertainties that will be described in Section 3.6, while the effect on specific parts of the analysis will be evaluated in Chapters 4, 5 and 6.

¹A short discussion about Michel electrons and an important use in the analysis will be presented in Section 5.3.2.2



(a) Distribution of dE/dx vs. the distance from the end of the track for all hits in selected stopping muon tracks. The black curve shows the mean of a fit to the peak of the attenuation-corrected detector response at particular distances from the track end. The absolute calibration uses hits in the 100-200 cm, where the detector response is almost uniform.

(b) Distribution of dE/dx for hits in the calibration window in units of $\text{PE}_{\text{corr}}/\text{cm}$ (top), and after calibration in units of MeV/cm (bottom). The calibration fixes the misalignment between the data (black) and the simulation (red).

Figure 3.12: The absolute calibration uses stopping muons selected in the cosmic data to convert dE/dx from the arbitrary units of corrected response ($\text{PE}_{\text{corr}}/\text{cm}$) to the physical units GeV/cm . From [33].

The corrected energy per hit that is the output of the calibration is denoted as “calorimetric energy”, and constitutes the fundamental energy estimator in NOvA. Since not all the particles involved in the neutrino interaction will deposit energy in the detector, and numerous processes might be mis-modeled or unaccounted, we construct additional estimators that are specific to the analyses and show better resolution for the desired signal. The 2017 estimators for ν_e and ν_μ are discussed below.

3.5.2 Muon neutrino energy

Recall from Section 3.2 the basic topology of CC neutrino interactions in the NOvA detectors: a track or shower associated with the lepton, and a hadronic component from the nuclear system. In the case of ν_μ CC, muons have a low MIP dE/dx , leaving long tracks in the detector. The muon energy can be estimated using the track length, which has an uncertainty corresponding to a single plane length. This energy estimate is more precise than calorimetric energy, which is used for the hadronic component. In order to improve the energy resolution, both estimates are obtained from spline fits as shown in Figure 3.13. The muon energy has an average resolution of 3%, while the hadronic part has a resolution of 30%. The total ν_μ CC energy estimate is given by

$$\hat{E}_\nu = \hat{E}_\mu + \hat{E}_{\text{had}}. \quad (3.7)$$




Figure 3.13: Spline fits used in reconstructing both the \hat{E}_μ (left) and \hat{E}_{had} (right). From [33].

3.5.3 Electron neutrino energy

Similar to ν_μ , the ν_e energy has two components: the electron shower and the hadronic part. Unlike the muon track, the energy deposition of the electron has a complicated shape, and the total energy cannot simply be measured from its length. The size of an electromagnetic shower grows as $\ln E$, and the brems activity towards the end is stochastic and adds to the perceived length. Therefore, both electromagnetic and hadronic components are measured calorimetrically, with the

EM label assigned to hits in prongs that are electron- or photon-like according to the CVN prong label of Section 3.4.3.

In contrast to the ν_μ case, both components of the neutrino energy have similar resolutions. In order to improve the final resolution of E_ν , these two components are not simply added, but we perform a fit to a quadratic function of the form

$$E_\nu = E_\nu = aE_{\text{EM}} + bE_{\text{had}} + cE_{\text{EM}}^2 + dE_{\text{had}}^2 \quad (3.8)$$

The training sample contains simulated events that have been reweighted to avoid biasing towards the peak energies in the flux. After the fit, the neutrino energy resolution obtained is 11%.




Figure 3.14: Left: Distribution of hadronic energy vs. electromagnetic energy for ν_e CC events. Colors indicate the average true energy per bin. The ν_e energy reconstruction results from a quadratic fit to this distribution. Right: fractional difference between true and reconstructed neutrino energy. The energy resolution is 11%. From [33].

3.6 Sources of systematic uncertainty

We evaluate the effect of systematic uncertainties by reproducing the analysis steps with different simulated samples. This involves reweighting events based on some combination of variables that are already stored in the MC files, or creating new simulated samples by introducing different parameters in some step in Section 2.4 and running reconstruction from scratch [98].

The main sources of uncertainty can be classified as: flux related, detector, cross-sections and final state interactions, calibration and energy response. Below is a brief overview of what types

of uncertainties were considered. The effects on the analysis will be presented in the following chapters.

3.6.1 Flux systematics

There are two categories of flux uncertainties: *beam transport*, which cover differences between the simulation and the working conditions of the NuMI beam, and *hadron production*, which concern the rates of production of pions and kaons from the proton collisions on the carbon target.

The beam transport uncertainties considered include the horn and target position, the horn current, the beam position on the target and the beam spot size. The effect of each uncertainty is below 5% on the flux spectrum at the near or far detector for almost all the uncertainties investigated [99].

As mentioned in Section 2.4, we used a package, PPFX, to constrain the hadron production models for the NuMI beam using external constraints. The systematic effect is assessed by generating a number of alternative weightings where the uncertainty on the fixed target data and any theoretical assumptions are allowed to float. These “alternative universe” PPFX weights are reduced into a set of uncorrelated weights via principal component analysis [100]. Four principle components are used in the fit.

3.6.2 Cross section systematics

The majority of the cross section and final state interaction (FSI) uncertainties are evaluated using reweighting factors included in GENIE [49], which can be applied to the simulated events based on the true information of the neutrino interaction.

A few additional modifications are added, informed by results from the cross section community and the NOvA ND data. First, for MEC interactions we consider uncertainties in the dependence with neutrino energy, energy-transfer, and final-state nucleon-nucleon pair composition, based on the different theoretical models available [101, 102, 103]. Second, uncertainties on the RPA model are included following the prescription in [53]. Third, the uncertainty on the value of M_A^{QE} , the

axial mass in quasi-elastic scattering, is reduced from the GENIE default value of ${}^{+25\%}_{-15\%}$ to $\pm 5\%$. The higher default value is driven by the missing nuclear effects that are now being included [104]. Fourth, based on measurements of NOvA ND data, the uncertainty applied to non-resonant pion production with three or more pions and invariant hadronic mass of $W < 3 \text{ GeV}$ was increased to 50% to match the default for 1- and 2-pion cases. Finally, we include two uncertainties on the ratio of ν_e CC and ν_μ CC cross sections: 2% to account for potential differences due to radiative corrections, and 2% to consider the possibility of second-class currents in CCQE events [107, 108].

3.6.3 Detector response systematics

The amount of light that is produced in the scintillator as particles traverse the detector involve two effects: scintillation and Cherenkov radiation. Uncertainties related to light levels and thresholds are included by altering the light level by $\pm 10\%$ with a compensating change made to the absolute calibration constants [110]. The size of this shift is taken from the measured spread in cell-by-cell brightness relative to the simulated value. This gives us an idea of how mis-modeling true energy deposition of different charged particles can bias the reconstruction and PID algorithms. In addition, an adjustment to the Cherenkov model shifts the proton response down by 2.6% while leaving the muon response unchanged [110].

3.6.4 Calibration and energy response systematics

Uncertainties in the calibrations and energy scales can affect the near and far detector differently, so we divide them in absolute (fully correlated) and relative (anticorrelated or uncorrelated). The overall energy response uncertainty is driven by the uncertainty in the calorimetric energy calibration. The absolute and relative muon energy scale uncertainties are both less than 1%. The calibration uncertainties are quantified by comparing simulated and measured data distributions of numerous channels, including the energy deposits of cosmic and beam muons, the energy distribution of Michel electrons which result from the decay of stopped muons, the invariant mass spectrum of neutral pion decays into photons, and the proton energy scales in quasi-elastic-like

events in the ND. The uncertainty used corresponds to the case with the largest discrepancies, namely the proton energy scale, at 5%. The size of the sample in the FD is not enough to evaluate the relative uncertainty, therefore the relative uncertainty is also set at 5% [94].

3.6.5 Other systematics

Normalization systematics: We quantify the uncertainty arising from potential imperfections in the simulation of beam-induced pileup in the ND by overlaying a single extra simulated event onto samples of both simulated and data events. We then examine the selection efficiency of this extra event and assign the 3.5% difference between the data and simulation samples as a conservative uncertainty on the normalization of the ND rate. These are added in quadrature with much smaller uncertainties in the detector mass and the total beam exposure to yield an overall normalization systematic.

ν_e extrapolation systematics: These are analysis-level uncertainties that arise from differences between the acceptances of the ND ν_μ selection criteria and the FD ν_e sample into which the ND corrections are extrapolated in the ν_e analysis. More details are presented in Section 5.3.3.


CHAPTER 4. ν_μ DISAPPEARANCE ANALYSIS

Figure 4.1: A ν_μ CC candidate from FD data, with the characteristic long muon track and small hadronic component. In this event display the neutrino beam arrives from the left, the color scale represents the ADC charge in a cell, and the two panels are the top and side views of the FD zoomed-in. Taken from [33].

NOvA can measure the atmospheric parameters Δm_{32}^2 and $\sin^2 \theta_{23}$ with the $\nu_\mu \rightarrow \nu_\mu$ channel. The analysis is based on the identification of ν_μ CC contained events in both detectors. The combination of an intense muon neutrino beam with the clear signature of muon tracks in the NOvA detectors allows us to have samples with good statistics and high purity.

The energy of a CC interaction can be reconstructed based on the depositions by the charged lepton and the visible energy of the hadronic system. In the case of ν_μ , the energy of the muon can be estimated to a high accuracy using the characteristics of the track (Section 3.5.2). Therefore,

the ν_μ CC energy resolution is limited by the amount of energy in the event that is not carried by the muon. For the current analysis, the ν_μ samples are split by the fraction of hadronic energy, which in turn separates the high-quality events with good energy resolution from those with lower resolution and backgrounds.

Section 4.1 describes the criteria for selecting events in each detector. Section 4.2 details the construction of histograms for the ν_μ analysis, including the variable binning of the reconstructed neutrino energy and the division in four subsets depending on the fraction of hadronic energy. In Section 4.3, the procedure to use ND data to correct the ν_μ signal prediction in the FD is explained. Section 4.4 presents the estimation of the cosmic background. The final predictions and expected effects of the oscillation parameters and systematics are examined in Section 4.5.

4.1 Event selection

In Chapter 3, we reviewed the basic event topologies that are pertinent to the analysis, in addition to the reconstruction and particle identification algorithms. While ν_μ CC events are visually identifiable via the long muon track, the oscillation requires a pure sample with reliable energy estimation. Therefore, selection criteria must combat cosmic backgrounds, neutral current events, and possible ν_μ CC beam events with poor reconstruction (included uncontained events). Figure 4.2 shows the successive stages of the ν_μ selection in the FD. These steps are:

- **Basic quality:** A minimal selection must ensure that the events are properly reconstructed and that a muon track candidate can be identified. To that end, we require at least 20 hits in the slice, four contiguous planes with activity, and the presence of a reconstructed track with a non-zero ReMID score.
- **Containment:** We want to eliminate interactions that occur outside of the detector (cosmics, rock) and keep only the beam candidates where all the energy is deposited inside the detector. Rejected events have activity too close to the edges of the detector (60 cm to the top, 12 cm to the bottom or west, 16 cm to the east, 18 cm to front or back), or their projected distance from a track to the edge is less than 6 cells.




Figure 4.2: Flowchart of the selection of ν_μ candidates in the FD. In the last step, the sample is further divided in four sets, depending on the fraction of hadronic energy in the event. The same sequence is applied on MC, NuMI data and cosmic trigger data. From [33].

- PID:** A combination of two identification algorithms from Section 3.4 are used to get most ν_μ CC-like events. The ν_μ output of the CVN classifier, CVNm, with a cut value of 0.5 was used for the first time in this analysis. Additionally, we also require scores of the ReMID classifier higher than 0.5 to keep events with a clear muon track and reject some cosmogenic events.
- Cosmic rejection:** To reduce the cosmic background even further, we remove any remaining events where the angle between the Kalman track and the beam direction is large ($\cos \theta > 0.5$) or the activity in the detector is low (require at least 400 hits in the FD). Finally, we require a score of the ν_μ cosmic BDT higher than 0.5.

Table 4.1: Predicted composition of the ν_μ CC candidate sample in the FD, in event counts, at various stages in the selection process. The oscillation parameters used in the prediction are the best fit values from Section A.2. As published in [47].

Selection	$\nu_\mu \rightarrow \nu_\mu$ CC	NC	ν_e CC	ν_τ CC	$\nu_e \rightarrow \nu_\mu$ CC	Cosmic
No selection	963.7	612.1	126.6	9.6	0.6	4.91×10^7
Containment	160.8	219.9	61.5	2.4	0.3	1.95×10^4
CVN	132.1	3.0	0.3	0.4	0.2	26.4
Cosmic BDT	126.1	2.5	0.3	0.4	0.2	5.8

4.1.1 FD selection performance

The efficiency of the basic quality and containment selection with respect to all ν_μ CC with true interaction vertices within a fiducial volume in the FD is 41.3%. Most of the events rejected have activity escaping the detector, energies higher than 5 GeV, or too few hits so that the muon track energy cannot be reconstructed. An example of a NOvA analysis that recovers the uncontained FD events is [111]. Out of the remaining events, the PID and cosmic rejection cuts preserve around 78% of the ν_μ CC events.

As seen in Table 4.1, the combination of selection cuts effectively reduces the number of cosmic backgrounds expected by 7 orders of magnitude. It is also evident that the PID algorithms are especially important to reject beam backgrounds: the dominant NC background, in particular, is reduced by 2 orders of magnitude. The final sample of ν_μ selected events is 93% pure ν_μ CC, over an expected background of 6 cosmic and 3 beam events. The analysis binning in Section 4.2 below will aid in further separation of signal and backgrounds.

4.1.2 ND selection

There are two differences between FD and ND selections in the ν_μ analysis. First, cosmic rejection cuts are not applied to the ND, since the rate of cosmogenic background is very low 100 m underground. Second, the containment cuts are modified to account for the smaller detector size, and the presence of the muon catcher. The cut-flow in the ND is as follows:

- **Basic quality:** Same requirements as the FD.

- **Containment:** All showers must be fully contained in the detector, with start and stop positions in the ranges $-180 \text{ cm} \leq x, y \leq 180 \text{ cm}$ and $20 \text{ cm} \leq z \leq 1525 \text{ cm}$. Similarly, all reconstructed tracks must start in the active region, and only the primary muon track can range out in the muon catcher. Since the latter is only $\frac{2}{3}$ of the height of the PVC planes, tracks that are estimated to have crossed at the top $\frac{1}{3}$ (“air gap”) are also rejected.
- **PID:** The same combination of $\text{CVN} > 0.5$ and $\text{ReMID} > 0.5$ is used.

Figures 4.3 and 4.4 show some comparisons between selected ND data and the simulation. Figure 4.3 displays the CVN scores for ND events that pass basic quality and containment cuts, and shows good agreement between both distributions, and a clear separation of the simulated background at lower scores. Figure 4.4 uses the full selection, in bins of reconstructed muon, hadronic and neutrino energies. As explained in Section 3.5.2, the muon neutrino energy estimate is constructed from the muon energy (related to the length of the muon track) and the hadronic component (related to the visible energy outside of the track).




Figure 4.3: Comparison between the ν_μ pre-selected distributions in ND data and MC, in bins of the CVN classifier. Higher scores indicate that events look more ν_μ -like. In the ND, the backgrounds to ν_μ CC are almost exclusively NC, which tend to have a very low CVN score. The analysis only considers events with $\text{CVN} > 0.5$. From [33].




Figure 4.4: Comparison between the ν_μ -selected distributions in ND data and MC, in bins of reconstructed muon energy (top left), hadronic energy (bottom left) and neutrino energy (right). The neutrino energy estimator is constructed as a function of the two components on the left. From [33].

4.2 Analysis binning

Recall from Section 3.5.2 how the ν_μ CC energy estimation requires two components: the muon energy and the visible energy of the hadronic system. The estimation of the muon energy in the FD uses the reconstructed track, and the resolution for the selected ν_μ CC signal is 3.5%. In contrast, the hadronic energy resolution is $\sim 40\%$, thus becoming the limiting factor for the E_ν estimate. The ν_μ CC energy resolution over the whole FD sample is 9.1%. At the ND, the lower active fraction of the muon catcher degrades the accuracy of the muon energy estimation, bringing the overall ν_μ CC energy estimation to 11.8%. The distributions comparing the muon, hadronic and neutrino energies in the ND for data and MC are in Figure 4.4.

The ν_μ CC interactions in which most of the energy goes to a contained muon not only have better resolution, but often are easier to classify as muon neutrinos. The amount of hadronic

activity is therefore closely related to the quality of the events, and can be exploited to increase the sensitivity of the analysis [112]. We use the hadronic energy fraction, $\frac{E_{\text{had}}}{E_\nu}$, as a function of reconstructed neutrino energy, to separate the selected events into sub-samples that we will occasionally refer to as “resolution bins”.




Figure 4.5: Distribution of selected ν_μ CC events in the FD, in bins of hadronic energy fraction and reconstructed neutrino energy. No oscillation weights are applied. The lines correspond to the 25%, 50% and 75% quartiles for each energy bin. These boundaries are used to create four sub-samples, ranging from low to high E_{had} fraction. From [33].

The ranges of fraction of hadronic energy as a function of reconstructed neutrino energy used to construct the four sub-samples are presented in Figure 4.5. While these boundaries start as quartiles in a strict sense, dividing the un-oscillated ν_μ CC population in sets with 25% of the events, after adding the backgrounds and applying oscillation weights, the relative proportions will vary. Nonetheless, we will keep using the term “quartile” to refer to them.

Figure 4.6 shows a comparison of the FD predictions for two choices of oscillation parameters, broken down by quartiles and signal/background components. It can be seen that, as a consequence of the resolution binning, most of the backgrounds are isolated in the fourth quartile, together with the ν_μ CC events that have more energy going into the hadronic system. The average energy resolution in the FD for each quartile is 6.2%, 8.2%, 10.3% and 12.4%, respectively.




Figure 4.6: The ν_μ -selected events are separated in quartiles of hadronic-energy fraction, labeled 1 through 4 and ranging from best to worst energy resolution. Quartile 1 will have higher sensitivity to ν_μ disappearance, opposite to quartile 4 where most background events are concentrated. From [33].

We use non-constant bin widths for the reconstructed neutrino energy, which is the variable used for the oscillation measurement. As shown in the figures below, the bins are narrower between 1 and 2 GeV, where the oscillation dip is observed, and wider at high energy, where fewer events are expected and the effect of oscillations is milder. The bin edges are located at (0, 0.75, 1, 1.1, 1.2, 1.3, 1.4, 1.5, 1.6, 1.7, 1.8, 1.9, 2, 2.25, 2.5, 2.75, 3, 3.5, 4, 5) GeV, for a total of 19 analysis bins per quartile [97].

4.3 Constraints from ND Data

The criteria defined in Section 4.1 select highly pure ν_μ CC samples in both detectors. According to the simulation, the ND events are 99.3% ν_μ CC and 0.7% NC. Therefore, we can only use the

ν_μ -selected data in the ND to constrain the $\nu_\mu \rightarrow \nu_\mu$ signal in the FD. The basic steps of the ν_μ signal extrapolation are sketched in Figure 4.7.




Figure 4.7: Diagram illustrating the steps for the ν_μ signal extrapolation. ν_μ -selected events in the ND are used to correct the ν_μ signal prediction in the FD by translating any discrepancies with the MC to bins of true energy. This information is combined with the $\nu_\mu \rightarrow \nu_\mu$ oscillation probability and the FD simulation to get the corrected signal in ν_μ analysis bins, to be later added with the background prediction and compared with the FD data. The same process is repeated for each of the quartiles. From [33].

The ν_μ ND selection is applied to both data and simulation. Any Data/MC discrepancy is assumed to be due to the ν_μ (or $\bar{\nu}_\mu$) charged current component, and is split according to their proportions in the simulation:

$$\nu_\mu^{\text{CC,data}} = (\text{data} - \text{not } \nu_\mu \text{ CC}) \left(\frac{\nu_\mu}{\nu_\mu + \bar{\nu}_\mu} \right). \quad (4.1)$$

The estimated $\nu_\mu^{\text{CC,data}}$ component is transformed from reconstructed to true energy bins using a 2D migration matrix from the simulation. The ratio of the true energy distribution in the FD to that in the ND is used to scale the result of the previous step, to get a corrected, un-oscillated FD prediction. By applying the energy-dependent oscillation weights for some set of parameters, and using the FD true-to-reco migration matrix, we finally obtain an oscillated FD prediction in bins of reconstructed energy that can be compared to the FD data.

As mentioned earlier, discrepancies between the ND data and MC can be due to mis-modeling of the neutrino flux, cross sections, detector effects, etc. The main idea behind the extrapolation

is to use an in-situ measurement to correct the central values of the simulation. That way, our observations in the FD are more likely to be related to neutrino oscillations, while deficiencies in the simulation become a next-order effect.

With the introduction of quartiles in Section 4.2 come different choices in how to use the ND data: i) use the same hadronic energy fraction vs energy cuts in both detectors; ii) use the ND data as a single sample, reused four times with the four FD MC predictions; iii) use different quartiles in each detector. Studies with fake ND data show that i) is the most robust approach, returning predictions that more closely resemble the FD expectation in systematically-shifted universes [112, 113].

Figure 4.8 presents the ND distributions in bins of reconstructed energy, with the quartile cuts applied. While the ND selection alone shows a 1.3% excess of data over the MC, the quartiles have offsets of +12%, -13%, -13% and +4%, respectively. These differences are almost exclusively a normalization effect, as can be seen by comparing Figs. 4.8 and 4.9. This implies that the shapes of the distributions are rather well-modeled in the MC, and only rates of events selected in each sample require an adjustment.




Figure 4.8: Comparison of selected ν_μ CC candidates (black dots) in the ND data to the simulation (red histograms) in the hadronic energy fraction quartiles, where the prediction is absolutely normalized to the data by exposure. The expected background contributions (gray) are smaller in the quartiles with better resolution. The shaded band represents the quadrature sum of all systematic uncertainties. These distributions are the input to the extrapolation procedure. As published in [47].




Figure 4.9: The same as Figure 4.8, but with the MC histograms normalized to the number of events in data. By removing the effect of normalization, it can be seen that the shapes of the distributions from data and MC have a good agreement.

4.4 Cosmic background prediction

Recall that the NOvA FD is located on the surface, and a number of cosmogenic events will inevitably contaminate the beam data sample. Combined with the known timing of the beam spill, the selection cuts from Section 4.1 progressively reduce the expected background by 7 orders of magnitude. However, the $\nu_\mu \rightarrow \nu_\mu$ measurement relies on the energy distribution of the signal events, especially around the oscillation dip (1–2 GeV). In consequence, we also need to obtain an expected distribution of the cosmic background.

To obtain an estimate of the cosmic backgrounds with identical detector conditions to the data, we can use the 420 μs around the blinded beam spill window from the FD NuMI-trigger dataset. There is a low number of events before scaling that is insufficient to provide a reliable distribution that can be used as an input to get an estimate in each quartile [91]. For this reason, we use a fit to cosmic trigger data to provide the shape of the prediction. This dataset has the advantage of larger statistics, with the downside that the detector conditions might not be a perfect match [112, 91]. The later is accounted for by averaging over periods, and using the counts from the NuMI sideband to fix the normalization. The results of this process are illustrated in Figure 4.10.




Figure 4.10: Estimation of the cosmic background for the ν_μ analysis. The cosmic trigger data (black) are used to estimate the shape of the distribution, while the out-of-time spill data (blue) set the normalization. Given the relatively low number of events in the samples, especially if they are split by quartiles (four plots to the right), a smooth function from the fit to data (red) is used. From [33].

4.5 Far detector prediction and systematics

The final ν_μ FD expectation combines the constrained signal prediction (Section 4.3), the cosmic background estimate (Section 4.4), and simulated beam backgrounds (NC, $\nu_e \rightarrow \nu_\mu$ CC, $\nu_e \rightarrow \nu_{e,\mu}$ CC, $\nu_{e,\mu} \rightarrow \nu_\tau$ CC). A breakdown of these components in the four quartiles for a given set of oscillation parameters is given in Table 4.2 and plotted in Figure 4.11.




Figure 4.11: Predicted distribution of ν_μ candidates in the FD, including the constraints from ND data and the cosmic estimate, in bins of reconstructed energy and divided by the hadronic energy fraction quartiles. The prediction is oscillated to the parameters from Section A.2.

Both the event counts and their distribution in energy bins, depend on the choice of oscillation parameters. In Figure 4.12 (left) we compare the prediction from the table above to one in the absence of oscillations (for simplicity, all quartiles have been added together). If all muon neutrinos survived, we would expect 763 events in the FD. Similarly, Figure 4.12 (right) compares that

Table 4.2: Predicted composition of the ν_μ CC samples in the FD, including the constraints from ND and cosmic data. These correspond to the integrals of Figure 4.11.

Sample	Total	$\nu_\mu \rightarrow \nu_\mu$ CC	NC	CC background	Cosmic
Quartile 1	34.69	34.35	0.08	0.08	0.18
Quartile 2	25.74	25.35	0.14	0.10	0.15
Quartile 3	27.75	26.80	0.33	0.13	0.50
Quartile 4	41.02	33.69	1.95	0.38	5.00
All quartiles	129.20	120.19	2.50	0.69	5.82

prediction (with $\sin^2 \theta_{23} = 0.558$) with the prediction at maximal mixing ($\sin^2 \theta_{23} = 0.5$) and the 2016 result ($\sin^2 \theta_{23} = 0.69$, from Section A.1).

4.5.1 Systematic uncertainties

The effects from different types of systematic uncertainties on the event counts are summarized in Table 4.3. These correspond to the error bands in Figure 4.11. Note that they depend on the oscillation parameters, and that systematics that shift events from one bin to the next (either energy or resolution bins) might have a significant effect on the measurement while not necessarily changing the total event counts. An approach to assess the impact of systematics directly on the $\sin^2 \theta_{23}$ and Δm_{32}^2 measurement is introduced in Section 6.2.

Table 4.3: Effect of 1σ variations of the systematic uncertainties on the total ν_μ predictions. Simulated data were used and oscillated with the parameters from Section A.2.

Source of uncertainty	Quartile 1	Quartile 2	Quartile 3	Quartile 4
	(%)	(%)	(%)	(%)
Calibration	5.1	7.1	11	9.1
Cross sections and FSI	3.1	10	9.9	4.3
Muon energy scale	2.8	2	1.3	1.1
Normalization	3.5	3.5	3.5	3.5
Detector response	1	1.3	2	1.9
Neutrino flux	0.39	0.43	0.54	0.88
Total systematic uncertainty	7.7	13	15	11
Statistical uncertainty	18	18	18	17




Figure 4.12: Expected distribution of ν_μ CC candidates in the FD for different assumptions of the oscillation parameters, summing over all quartiles. Left: In the absence of ν_μ disappearance ($P(\nu_\mu \rightarrow \nu_\mu) = 1$), we would observe six times more events, and no dip before the flux peak. Right: Predictions for three realistic sets of oscillation parameters. A very marked dip between 1-2 GeV would indicate that the data are compatible with maximal mixing ($\sin^2 \theta_{23} = 0.5$). The introduction of quartiles and variable bins of reconstructed energy in this analysis is meant to increase the sensitivity of the measurement.


CHAPTER 5. ν_e APPEARANCE ANALYSIS

Figure 5.1: A ν_e CC candidate from FD data, with the characteristic electromagnetic shower and smaller hadronic component. In this event display the neutrino beam arrives from the left, the color scale represents the ADC charge in a cell, and the two panels are the top and side views of the FD zoomed-in. The CVNe score of this event is 0.993, and its reconstructed energy is 2.36 GeV. Taken from [33].

The comparison between the NOvA ν_e appearance data with the expectation under oscillations allows us to probe δ_{CP} , $\sin^2 \theta_{23}$, and the neutrino mass hierarchy. In order to investigate $\nu_\mu \rightarrow \nu_e$ oscillations, we require a pure sample of ν_e CC events in the FD. The ν_e analysis uses events selected by the CVNe classifier (Section 3.4.1), and a series of data-driven methods to improve the prediction from the simulation. In contrast with the ν_μ analysis described in the previous chapter, in the ν_e analysis the nature of the events selected in each detector is different, and there is a higher incidence of beam-related backgrounds. To correct the appearance signal in the FD, we use

ν_μ -selected events in the ND. The ν_e candidates in the ND constrain the NC, ν_μ CC and intrinsic ν_e CC backgrounds.

The selection criteria applied to reconstructed ν_e events are listed in Section 5.1. Section 5.2 describes the variables and the binning that will be used for the oscillation measurement. Section 5.3 is concerned with the constraints to the signal and background predictions using the ND data, while Section 5.4 covers the expected background from cosmogenic activity using FD data. The resulting prediction, and additional effects that can affect the oscillation measurement, are detailed in Section 5.5.

5.1 Event selection

In Chapter 3 we described the prong-based event and energy reconstruction algorithms that are more appropriate for ν_e CC events (compared to the ν_μ CC events that rely on the muon track). We also introduced CVNe, the output of the CVN classifier for the electron neutrino hypothesis. These reconstructed variables are available for all data and MC events, but only a well-defined subset will be acceptable for the appearance measurement. The selection criteria are tuned to maximize the figure of merit (FOM)

$$FOM = \sqrt{\sum_i^{\text{bins}} \frac{s_i^2}{s_i + b_i + c_i}}, \quad (5.1)$$

where s_i, b_i, c_i represent the expected number of ν_e appearance signal, beam background and cosmic background events in each bin i of reconstructed energy and PID. Two samples with independent optimizations are used in the analysis: *core* and *peripheral*; as their names indicate, the essential difference between them is the distance of the events to the edges of the detector. Figure 5.2 provides an overview of the selection steps for both samples, which will be described in detail in the subsections below. The cuts are tuned using the MC and cosmic-trigger datasets [114], while the final cosmic background estimate and the beam measurement use the NuMI trigger dataset (with two appropriate timing selections, see Section 5.4).




Figure 5.2: Flowchart of the selection of ν_e candidates into the two FD samples used in the analysis. Events that fail the core preselection or cosmic rejection cuts are considered for the peripheral sample. The same sequence is applied on MC, NuMI data and cosmic trigger data. From [33].

5.1.1 Core sample

The core sample comprises fully-contained events that are compatible with ν_e CC signal topologies. The selection criteria are analogous to those used in the ν_μ analysis (Chapter 4) and in published versions of the ν_e analysis [115, 76], in contrast with the new peripheral sample that will be introduced below. The selection steps charted in the left side of Figure 5.2 proceed as follows:

- **Basic selection:** These cuts remove obvious cosmic background (vertical tracks and through-going events), reconstruction failures (absent vertex or prongs, or too many hits per plane), APD saturation effects, and some events from a partial detector (reconstruction is unreliable when less than 4 out of 14 consecutive diblocks are taking good data).

- **Containment and energy (“preselection”)**: We remove events that have hits too close to the six walls of the detector (63 cm to the top, 18 cm to the front or back, 12 cm to the other walls), indicating that they are likely cosmogenic, rock backgrounds, or uncontained (energy estimation is not reliable). Further, we only keep events with 1–4 GeV reconstructed energies, 30–150 hits, and main prong lengths between 100 and 500 cm.
- **Cosmic rejection**: Cosmic backgrounds that are contained but not aligned with the beam direction are removed by cutting on the event transverse momentum fraction, p_T/p . A fraction of the remaining events correspond to cosmic photon showers entering the back of the detector where the overburden is thinner [116, 117]. These can be reduced using their *sparseness asymmetry*, defined as the difference in the number of planes without hits in the first and last eight planes of the event, divided by their sum. The basic idea behind this is that a shower in the beam direction is expected to have activity concentrated upstream, and more disperse downstream. Finally, a *nearest slice* cut looks for cosmogenic-like activity close in space or time to the selected event, to reduce cases where a single background event is split into two or more pieces sliced separately, and one of them is mistaken as a signal [114].
- **CVN**: The ν_e classifier CVNe, described in Section 3.4.1, achieves most of the background rejection, for both cosmic and beam interactions. The FOM defined in (5.1) is optimized with $\text{CVNe} \geq 0.75$.

Table 5.1 presents the predicted number of events at each of these selection stages, for the signal and background components. The intrinsic beam ν_e background is noticeably reduced by the containment and energy cuts, while the NC and other CC components require the topology-based CVNe cuts. After the full selection, the level of cosmic rejection reached is $\mathcal{O}(10^7)$. The distributions of energy and CVNe for the core sample are presented in Figure 5.3.

5.1.2 Peripheral sample

Comparing the different cut stages in Table 5.1, it is clear that the energy and containment step is the least efficient, eliminating a third of the potential signal events. The “peripheral” selection

Table 5.1: Predicted number of ν_e CC signal and background events at the FD, for each stage of the core selection. These predictions use the oscillation parameters listed in Section A.2. As published in [47].

Selection	$\nu_\mu \rightarrow \nu_e$ CC	Beam ν_e CC	NC	ν_μ, ν_τ CC	Cosmic
No selection	77.9	48.7	612.1	973.8	4.91×10^7
Containment/energy cut	52.3	8.0	121.4	49.3	2.05×10^4
Pre-CVN cosmic rejection	51.3	7.9	114.3	47.0	1.58×10^4
CVN	41.4	6.0	5.3	1.3	2.0




Figure 5.3: Distributions of reconstructed energy (left) and CVNe (right) for ν_e candidates in the FD core sample, divided by interaction type and neutrino flavor. From [33].

[117] was introduced in this analysis as an attempt to recover some of those events. The two steps to the right of Figure 5.2 are:

- **Basic quality and energy (“preselection”)**: The events that are considered for the peripheral sample must pass the same basic quality cuts from the core selection, but not the containment or cosmic rejection requirements. To achieve good purity, only events in the 1-4.5 GeV energy range are included.
- **BDT and CVN**: In this preselected sample, the high PID region, $\text{CVN}_e > 0.95$, has a significant signal component of 5.45 signal events, with low beam backgrounds (1.34) and

relatively low cosmic events (20.0). To increase the cosmic rejection, a Boosted Decision Tree (BDT) classifier was introduced, using some of the variables from the core cosmic rejection as discriminators [114]. Events with $CVN > 0.99$ or $CVNE > 0.95$ and $BDT > 0.53$ are preserved.

Figure 5.4 shows the two-dimensional BDT-CVN space for the preselected peripheral events. A high-purity sample is obtained by cutting on both variables simultaneously. Table 5.2 presents the predicted composition of this sample.




Figure 5.4: Distribution of BDT vs CVNe for events preselected for the ν_e peripheral sample. The black boxes represent the number of signal events, and the shaded color represents the purity in each bin. Bins with lower purity are dominated by cosmic backgrounds, and are concentrated to the left of the red line; events in that region are rejected from the analysis. As published in [47].

Table 5.2: Predicted number of ν_e CC signal and background events at the FD, at each stage of the peripheral selection. These predictions use the oscillation parameters listed in Section A.2. As published in [47].

Selection	$\nu_\mu \rightarrow \nu_e$ CC	Beam ν_e CC	NC	ν_μ, ν_τ CC	Cosmic
Basic quality/energy cut	20.4	6.6	199.9	160.9	2.79×10^6
CVN + BDT	5.9	1.0	0.2	0.1	2.2

5.1.3 FD selection performance

As indicated earlier, before the addition of the peripheral sample, only 67% of the ν_e CC signal events would have been considered for the analysis before cosmic rejection and particle identification. The new combined preselection keeps 93% of the candidates. Out of this population, 67.4% of the events remain after full selection. While the efficiency of the peripheral CVN + BDT cut is only 30%, compared to 82% for the core CVN + CosRej cut, the overall gains from the optimizations performed for this analysis represent a 17.4% effective gain in exposure compared to our previous results [76].

From Tables 5.1 and 5.2, it is clear that CVNe is a powerful tool for rejecting both cosmic and beam backgrounds. It reduces the estimated cosmogenic component by 4 (5) orders of magnitude in the core (peripheral) sample. The level of beam background rejection with CVN is also high: 92% for the core selection (loose CVNe > 0.75 cut) core and 99.7% for the peripheral (tight CVNe > 0.95 cut).

Finally, it is worth noting that most backgrounds that remain after the selection are expected to present electromagnetic activity, which is the signature for ν_e signal events. This is illustrated in Figure 5.5. The beam ν_e CC events are, of course, true ν_e interactions that produce an electron, and are only labeled as backgrounds since the desired signal is the oscillated $\nu_\mu \rightarrow \nu_e$. The misidentified ν_μ CC and NC background events very often involve π^0 s that subsequently produce the electromagnetic showers via $\pi^0 \rightarrow \gamma\gamma$.

5.1.4 ND selection

The ND selection is designed to resemble the FD core selection from Section 5.1.1. The most relevant characteristics of the ND that need to be considered are: i) the detector size, which affects the containment criteria, and ii) its location (closer to the source and underground), which affects the shape of the flux, creates intensity effects, and makes the cosmic component negligible. The selection cuts proceed as follows:




Figure 5.5: The FD background, decomposed into ν_μ CC, ν_e CC, NC and cosmic. The colors represent the type of true electromagnetic activity: a final state electron, a final state π^0 with kinetic energy higher than 0.5 GeV, or none. The cosmic estimate is also shown in gray. From [33].

- Basic quality:** The same criteria for reconstruction quality are applied to both detectors. The data from a partial ND is never used.
- Containment:** We require vertices within the fiducial volume $-100 < x < 160$ cm, $-160 < y < 100$ cm, $100 < z < 900$ cm, and shower start and stop points within $-170 < x, y < 170$ cm and $100 < z < 1225$ cm. The first requirement keeps neutrino interactions that occurred inside the detector, not the surrounding rock, while the second one keeps only events where all activity is expected to be inside the detector, thus maintaining a reliable energy reconstruction.
- Preselection:** The ND preselection is slightly looser compared to the FD, keeping events with 1-4.5 GeV reconstructed energies and 20-200 hits. The extra high energy bin is now used for the extrapolation of the peripheral sample; while the number of hits range is consistent with the dimensions of the ND.
- p_T/p :** While the rock overburden of the ND almost eliminates the cosmic background, the p_T/p cut for cosmic rejection in the FD is also applied. This cut sculpts the kinematics of

selected events in the FD, so we want to guarantee that the ND constraints are compatible with the FD. Furthermore, as the ND is closer to the beam source, the angular distribution of events in each detector is expected to be very different without this cut.

- **CVN:** we use the same CVNe cuts in both detectors.

The distributions of reconstructed energy and CVNe for the ND data and MC are compared in Figure 5.6. The behavior of the ν_e component stands out in both distributions, presenting a high energy tail and a concentration of events at high CVN scores. These are expected features of the classification, as beam ν_e 's constitute an irreducible background to the appearance signal, especially around the 2 GeV peak of the ν_μ flux.

The distributions in Figure 5.6 also show good data/MC agreement. The use of the ND data for further checks of the ν_e selection and to constrain the FD prediction will be discussed in Section 5.3.




Figure 5.6: Comparison between the ν_e -selected distributions in ND data and MC, in bins of reconstructed energy (left) and CVNe (right). The dotted lines correspond to the raw simulation, while the solid lines correspond to corrected distributions after applying the techniques of Section 5.3.2. From [33].

5.2 Analysis binning

In the previous section we discussed some key characteristics of the distributions in reconstructed neutrino energy and CVNe for signal and backgrounds. Specifically, cosmic backgrounds

are expected to concentrate in the low-CVN core region and the peripheral; beam ν_e CC dominate at high CVN and high energy; and NC and ν_μ CC are rather flat in both variables. For this reason, both variables will be used to break down the FD spectrum for the oscillation measurement. The different signal/background composition of the analysis bins means different behavior under oscillations, thus increasing the statistical power of the binned fit [114].

In terms of PID, the core sample is divided into three regions, “low”, “mid” and “high”, defined by the ranges $0.75 < \text{Low CVN} < 0.87 < \text{Mid CVN} < 0.95 < \text{High CVN}$. The peripheral sample is not further divided by PID, and for simplicity is depicted as a fourth bin.

Based on the number of events expected in the FD, the energy axis is simply divided in bins of 500 MeV. For the core sample, this means 6 analyzable bins per PID region. While the peripheral selection includes events in the 1–4.5 GeV range, they are close to the edges of the detector and might be uncontained. For this reason, all peripheral events are summed into a single bin. In total, the ν_e appearance measurement will use $6 \times 3 + 1 = 19$ analysis bins.

Almost identical binning is used for ν_e candidates in the ND, with two differences: the energy range is 1–4.5 GeV; and there are only 3 PID regions, as there is not an equivalent of the peripheral sample in the ND. For the signal prediction, we use ν_μ CC candidates in the ND with the same selection and binning as the ν_μ analysis (see Sections 4.1.2 and 4.2).

5.3 Constraints from ND Data

Measurements in the ND are used to identify necessary corrections to the simulated FD predictions. Unlike the ν_μ disappearance analysis, where we search for the same flavor neutrino in both detectors, the ν_e appearance signal is associated with ν_μ in the ND and ν_e in the FD. The ν_e ND selection serves to probe the performance of the ν_e classifier, and to constrain three beam-related backgrounds to the $\nu_\mu \rightarrow \nu_e$ CC measurement. In this section, we discuss the signal and background extrapolation, with an emphasis on the data-driven techniques to separate the ν_e ND data into the ν_e CC, ν_μ CC and NC background components.

5.3.1 Signal prediction




Figure 5.7: Reconstructed energy distributions for ν_μ -selected events in the ND, for data and MC. This sample is used to constrain the predicted $\nu_\mu \rightarrow \nu_e$ signal in the FD. The scale of the background is very small compared to the ν_μ CC component, making it difficult to differentiate from the horizontal axis.

To predict the ν_e appearance signal we follow the same steps as in the ν_μ signal prediction from Section 4.3. The ν_μ -selected events in the ND data (Figure 5.7) are extrapolated to the FD via the reco-to-true migration matrix, Far/Near ratio, oscillation probability, and true-to-reco migration matrix. Figure 5.8 illustrates this process, which can be immediately compared to the ν_μ case Figure 4.7. The features that pertain to the ν_e analysis are: i) the ν_μ events are not split into the four resolution quartiles, ii) the $\nu_\mu \rightarrow \nu_e$ oscillation probabilities are orders of magnitude smaller than the $\nu_\mu \rightarrow \nu_\mu$ survival, and iii) while we use ν_μ selection and binning in the ND, the FD migration matrix has ν_e selection and analysis bins.

It should be noted that while there is an overall 1.3% excess of events in data over MC in the ν_μ ND sample, it is not straightforward to quantify its effect on the ν_e signal. The corrected signal prediction is affected by how the excess is distributed as a function of energy, as well as the precise values of the oscillation parameters. This will be explored further in Section 5.5.




Figure 5.8: Diagram illustrating the steps for the ν_e signal extrapolation. ν_μ -selected events in the ND are used to correct the ν_e signal prediction by translating any discrepancies with the MC to bins of true energy. This information is combined with the $\nu_\mu \rightarrow \nu_e$ oscillation probability and the FD simulation to get the corrected signal in ν_e analysis bins, to be later added with the background prediction and compared with the FD data. From [33].

5.3.2 Background prediction

Three types of beam-related ν_e backgrounds in the FD are estimated using the ND: neutral currents, ν_μ charged currents, and the intrinsic ν_e component in the NuMI beam (beam ν_e CC). Since each one propagates differently to the FD, we use a combination of data-driven techniques to correct their relative proportions from the simulation. Figure 5.9 is a comparison between the number of data and MC ν_e -selected events in the ND (Section 5.1.4) as a function of CVN and reconstructed energy (Section 5.2). The total number of events in data is 11.5% higher than the simulation, with different proportions depending on the analysis bin. This discrepancy means that the FD background prediction needs to be increased, but the results depend on how it is *decomposed*, i.e. divided among the three types.

The simplest way to handle the data/MC disagreement is to scale all components by the same factor, matching the total amount from the selected data bin-by-bin. This “proportional” decomposition was used for the first published results [115], and is useful whenever there are no other handles to better estimate the proportion, or as a first approximation while the analysis is in de-




Figure 5.9: Selected ν_e events in the ND, in bins of CVN and reconstructed energy. Any data/MC discrepancy must be split among the three components, which are extrapolated separately to the FD.

velopment. An error can be assessed by assigning 100% of the discrepancy to a single component [115, 114].

For the current analysis, the “combo” decomposition relies on two independent methods, each best suited to adjust one of the CC components; because the corrected total matches the data by design, the system is closed and all three components are determined. Details of the beam ν_e estimation using ν_μ CC samples are presented in Section 5.3.2.1, while the ν_μ CC and NC background estimates and the final result are in Section 5.3.2.2 and Section 5.3.2.3 respectively. Unlike the signal extrapolation, where a migration matrix translates the ND data to bins of true energy, we only use the reconstructed energy bins to constrain these three components in the FD. As introduced in Section 3.1.1, this is more appropriate for backgrounds, where energy estimation is not expected to be reliable.

5.3.2.1 Beam ν_e estimation

ν_μ CC samples in the ND are used to estimate the intrinsic ν_e CC component, by connecting them via the shared hadron ancestors that decay into either flavor. As explained in Section 2.2, the

neutrino beam results from the decay of hadrons, mostly pions and kaons, that were produced in the NuMI target and focused in the direction of the detectors. Figure 5.10 shows the true energy spectra of simulated ν_e and ν_μ CC events with true interaction vertex inside the ND. The muon neutrinos that dominate around the 2 GeV peak mainly result from the decay $\pi^+ \rightarrow \nu_\mu + \mu^+$. A few anti-muons that subsequently decay as $\mu^+ \rightarrow \bar{\nu}_\mu + e^+ + \nu_e$ give rise to the intrinsic ν_e component at low energies. At higher energies, the majority of ν_μ and ν_e originate in kaon decays. We will use ν_μ data events selected in the ND to obtain corrections to the pion and kaon yields, and apply them to the simulated beam ν_e component. Note that this estimate does not rely on the ν_e selection, and as such is independent of the measured excess in data seen in Figure 5.9.




Figure 5.10: True energy spectrum of muon (left) and electron (right) neutrino events with interaction vertex in the Near Detector according to the simulation, split by their ancestor at the target.

A special selection was developed in order to obtain two orthogonal sets of ν_μ CC data dominated by events with either pion or kaon ancestors [118, 119, 120]. This selection is a simplified version of the ν_μ selection used in the signal extrapolation. The shared criteria include: basic quality requirements; reconstructed vertex inside a fiducial volume; and a high score in the ReMID ν_μ classifier. An additional containment requirement on the muon track candidate splits the sample in two; the ν_μ contained and uncontained data and MC spectra are presented in Figure 5.11. Corrections to the pion (kaon) yield will use events in the (un)contained sample and the low-(high-) energy region.




Figure 5.11: ν_μ selected data and MC events, in the contained (left) and uncontained (right) samples. Contained events in the 0-5 GeV are used to calculate pion weights, while 1-3 GeV contained and 4.5-10 GeV uncontained events are used for the kaon scale.

There is an important distinction between the two samples affecting how the corrections are calculated. The energy estimator is designed to perform well for contained ν_μ events, and the resolution for uncontained or high-energy events is expected to be very poor. The kaon corrections result from events with both characteristics and therefore cannot rely too heavily on the reconstructed neutrino energy. On the other hand, the sample used for the pion corrections closely matches the one used for the signal extrapolation, and a more elaborate approach can be followed.

As seen in Figure 5.10b, the high energy tail of the selected sample is dominated by ν_μ CC from kaons. The kaon scale, S_K , is calculated from

$$S_K = \frac{\sum_j \left(N_{\nu_\mu}^{\text{data}} - S_\pi \times N_{\nu_\mu}^{\pi \text{CC}} - N_{\nu_\mu}^{\text{bkg}} - N_{\nu_\mu}^{\text{other}} \right)}{\sum_j \left(N_{\nu_\mu}^K \right)}, \quad (5.2)$$

where the sum is over uncontained events with ν_μ energies in the range 4.5–10 GeV. In (5.2), S_π is a pion normalization scale similarly obtained from contained events in the 0.75–3 GeV energy range. After a couple of iterations, a final $S_K - 1 = 3.84\%$ is obtained and applied to all simulated ν_e CC events with a kaon ancestor.

To calculate the pion corrections, we will use the reconstructed energy in data and MC from Figure 5.11a, as well as kinematic information of the ancestors for both ν_μ and ν_e from the simula-

tion. Figure 5.12 shows the distribution of transverse and forward momenta of the pion ancestors for the selected samples. It can be seen that the phase space of pions that generate beam ν_e 's overlaps with that of the ν_μ CC.




Figure 5.12: Selected ν_μ CC (left) and ν_e CC (right) events from a pion ancestor as a function of the pion transverse and forward momentum after it leaves the target. Pion weights obtained from the ν_μ samples are applied to ν_e events according to the ancestor p_T and p_z , and collapsed back to the ν_e analysis bins.

The calculation of pion weights starting from the ν_μ contained samples proceeds as follows. First, any excess of data over the sum of the simulated background and the contributions from *other* ancestors is split among pions and kaons according to their proportions in the simulation. The histogram with background-subtracted data is then divided by the simulated ν_μ CC from π to construct a set of weights w_{ν_μ} as a function of reconstructed energy E_{ν_μ} , as

$$w_{\nu_\mu}(E_{\nu_\mu}) = \left(\frac{N_{\nu_\mu CC}^\pi}{N_{\nu_\mu CC}^\pi + N_{\nu_\mu CC}^K} \right) \left(\frac{N_{\nu_\mu}^{\text{data}} - N_{\nu_\mu}^{\text{bkgd}} - N_{\nu_\mu CC}^{\text{other}}}{N_{\nu_\mu CC}^\pi} \right), \quad (5.3)$$

where N^X is the number of events in the sample X in a bin centered around the energy E_{ν_μ} . Note that the factor of $N_{\nu_\mu CC}^\pi$ cancels in (5.3).

Second, the weights are propagated to the level of the pion ancestors, in momentum (p_T, p_z) space,

$$w(p_T^\pi, p_z^\pi) = \frac{\sum_k w_{\nu_\mu}(E_{\nu_\mu}) \times E_{\nu_\mu}(k)}{\sum_k E_{\nu_\mu}(k)}, \quad (5.4)$$

where the sum is over all the energies that contribute to a (p_T, p_z) bin. Third, the (p_T, p_z) weights are applied to all the ν_e CC from π^+ events selected in the MC. Finally, a sum over (p_T, p_z) is done on these weighted ν_e events in each analysis bin, in order to give a corrected ν_e from π^+ spectrum.

The final beam ν_e estimate is the sum of the corrected pion and kaon components, and the uncorrected MC for other ancestors, with the added restriction that the total prediction on any given bin cannot exceed the data count in that bin. The results are summarized in Table 5.3. The combination of the 3.84% increase in the kaon component with the small pion corrections represents an overall increase of 1.13% in the total event count, and 0.95% in the high CVN bin.

Table 5.3: Differences between the beam ν_e estimation and nominal MC in the ND, for different ancestors and the three ranges of the CVNe selector.

0-5 GeV	Low CVN	Mid CVN	High CVN	Total
All	+1.45%	+1.33%	+0.95%	+1.13%
Pion	+0.05%	-0.03%	-0.04%	-0.02%
Kaon	+3.84%	+3.84%	+3.84%	+3.84%

Uncertainties in the procedure can be assessed by using systematically-shifted MC samples instead of ND data, and comparing the results with the *truth*. Using the $\pm 1\sigma$ systematically shifted MC variants, we find that the kaon scale is $3.84 \pm 0.95\%$, and the average pion scale is $0.02 \pm 0.94\%$. The relative small size of the corrections also reflect improvements in the simulation, especially through tuning. For example, the previous results published in [76] used a 17% kaon scale [121], that is now reduced to less than 4% after the flux improvements introduced for this analysis mentioned in Section 2.4.

Given that this estimation method relies on characteristics of the neutrino flux, it does not impose constraints on the ν_μ CC or NC backgrounds which are the result of misidentification. Any remaining discrepancies between ν_e -selected data and MC need to be assigned to those components either by dividing proportionally or using another handle from data.

5.3.2.2 ν_μ CC and NC background corrections

Time-delayed electrons from muon decay can be used to tag events and assign the ND data counts to one of the three ν_e background components. A true ν_μ CC interaction necessarily contains a muon, which decays to produce a Michel electron (ME) via $\mu^- \rightarrow e^- + \nu_\mu + \bar{\nu}_e$ with a high branching ratio [114]. Generally, the ν_μ CC events that are selected have the muon hidden in the shower associated with the hadronic recoil, but often the time-delayed ME can be identified. While the hadronic system can also produce MEs via a charged pion that decays to a muon, ν_μ CC interactions on average have one more ME than beam ν_e CC and NC, thus the number of MEs (N_{ME}) in an event can in principle be used to separate them. Figure 5.13 shows a simulated ν_μ CC background that passes ν_e selection; the muon has a short track that is hard to find, while an energetic π^0 produces the dominant electromagnetic shower.





Figure 5.13: Example of an event display for a ν_μ CC background with a Michel electron. From [92].

NOvA has a dedicated set of algorithms to identify Michel-like activity in data and simulation [92, 122]. There are two types of ME: those produced at the end of a clearly defined-muon track, or those following a short or hidden muon; ν_μ CC backgrounds to the ν_e analysis are of the second kind. Since the first type of ME are easier to find, they can be used to benchmark an identification algorithm, and be used for other tasks such as calibration [122]. Michel candidates are constructed with clusters of hits that are close in space but delayed in time from the parent physics slice. Four variables (calorimetric energy, number of hits, minimum distance to parent slice, time delay) are used as inputs to a log-likelihood PID (“MID”) to differentiate ME backgrounds. The distributions



(a) The variables used to identify a Michel candidate cluster are: calorimetric energy (top left), number of hits (top right), time delay (bottom left) and distance to the parent slice (bottom right).

Figure 5.14: Input variables and output score of the Michel electron classifier. The panels compare the distributions of these variables for ND data (black) and simulation (red), split in true Michel events (blue) and background (green). Events with $\text{MID} > 0$ are likely to have a true Michel and are used in the decomposition. From [92].

in data and MC of the four variables and the MID score are shown in Figure 5.14. Only clusters with $\text{MID} > 0$ contribute to the N_{ME} counts for the decomposition.

Figure 5.15 shows the distribution of the number of Michel electrons in data and simulation in ν_e -selected ND events. Note that the excess in data seen here is the same as in Figure 5.9 but with a different binning. The “Michel decomposition” method creates N_{ME} distributions for each bin of reconstructed energy and CVN, and rescales the ν_μ CC and NC components to match the data. The ν_e CC and NC are degenerate under this approach, so to reduce the number of degrees of freedom the beam ν_e estimate from Section 5.3.2.1 is used as an input. In cases where the statistics are limited, especially when the fraction of ν_μ CC is lower than 20%, this approach is not as robust against systematic uncertainties [122]. In those bins we revert to the simpler proportional decomposition described at the beginning of the section.




Figure 5.15: Distribution of the number of Michel electron candidates in ND Data and MC. On average, ν_μ CC backgrounds have more Michel candidates than beam ν_e and NC. Similar distributions are constructed for each analysis bin. We perform a fit of the ν_μ CC and NC scales to data, using the beam ν_e estimation results as fixed input. From [33].

5.3.2.3 Decomposition results

Figure 5.16 shows the final decomposition of the ND data into the three background components, using the beam ν_e estimate and the fit to number of Michel electrons. The corresponding number of events in each bin of CVN, before and after the corrections, is presented in Table 5.4.

Table 5.4: Breakdown of the number of ν_e -selected events in the ND, split by component and CVN bin, before and after the decomposition procedure. These numbers correspond to integrals in Figure 5.16.

(a) ND Data and uncorrected MC					(b) Corrected MC				
	LowCVN	MidCVN	HighCVN	Total		LowCVN	MidCVN	HighCVN	Total
Data	12213	9173	10862	32248	Total MC	12213	9171	10862	32246
Total MC	10750	8477	9690	28917	Beam ν_e	2135	3178	7272	12585
Beam ν_e	2053	3153	7022	12228	NC	5345	2915	1992	10252
NC	4639	2760	1366	8764	ν_μ CC	4732	3079	1598	9409
ν_μ CC	4058	2564	1303	7924					




Figure 5.16: Results of the decomposition of ν_e ND data. The uncorrected MC components (dashed) are scaled so that the total matches the data, bin by bin, using the beam ν_e and Michel estimation methods described above. Published in [47].

While the decomposition creates by design perfect agreement between data and MC in the analysis bins as shown above, applying the weights to simulated events also improves agreement in other variables [114, 123]. This impact was illustrated in Figure 5.6. The data-driven methods only slightly alter the proportions of each component out of the total: ν_μ CC from 27 to 29%; NC from 30 to 32%; beam ν_e from 42 to 39%. Compared to the simpler proportional decomposition, the total background count after extrapolation yields roughly the same number of events. However, it is

Table 5.5: Comparison between the proportions of ν_μ CC, ν_e CC and NC selected events in the ND MC, with(out) the corrections from the decomposition methods. From [114].

	ν_μ CC	NC	Beam ν_e CC	Total
Data	–	–	–	32248
Combo	9409 (29%)	10252 (32%)	12585 (39%)	32246
Proportional	8917 (28%)	9823 (31%)	13446 (42%)	32186
Uncorrected	7924 (27%)	8764 (30%)	12228 (42%)	28917

important to note that the main advantage of the combo decomposition is a reduction in systematic uncertainties, compared to the simulation-only case [114, 124].

5.3.3 Constrained FD components and extrapolation systematics

Table 5.6 summarizes the effect of the extrapolation of ND data for the ν_e signal and background predictions in the FD. Overall, all the background components are scaled up, with the higher corrections applied to the NC component. In contrast, the small excess of ν_μ data is translated into a -1.5% correction to the signal expectation at the chosen combination of oscillation parameters. The final predictions will be further discussed below in Section 5.5.

Table 5.6: Comparison of the ν_e FD signal and background predictions, before and after extrapolating the decomposed ND data. Predictions are oscillated using the parameters in Section A.2.


	ν_μ CC	NC	Beam ν_e CC	Total bkg.	Signal
No extrapolation	1.13	5.52	7.07	14.1	47.4
After extrapolation	1.26	6.40	7.27	15.3	46.7
Correction	+11%	+15.9%	+2.8%	+8.5%	-1.5 %

The details of the ν_e extrapolation call for additional checks on potentially unmodeled Far/Near behavior that can skew the prediction [114]. Two acceptance systematics need to be added to the list from Section 3.6: an uncertainty covering ν_μ ND vs ν_e FD kinematics for the signal, and a 1.2% normalization effect for the ν_e backgrounds. There is no equivalent to this uncertainty in the ν_μ analysis that would require adding new correlations to the fit.

The signal prediction uses ν_μ selected events in the ND to correct the expected ν_e distribution in the FD. Each have their own event topologies, energy estimators and selections, which can sculpt the properties of the events used in the analysis. While most effects are taken into account with the extrapolation and other systematics, a residual effect arises from the differences in kinematics. The effect on the prediction is evaluated by reweighting the ND events, so that the distributions of a given variable matches that of FD events. These variables are: the squared four-momentum transfer Q^2 , the fraction of transverse momentum p_T/p , and the angle with respect to the beam

direction $\cos \theta_{\text{NuMI}}$. The biggest effect was found to be in Q^2 ; the differences between this and the nominal prediction are taken as the uncertainty [125]. A comparison of the Q^2 distributions for ν_μ in the ND and ν_e in the FD, as well as a comparison of the three reweighted predictions to the nominal, are presented in Figure 5.17.

Another source of Far/Near discrepancies is the detector size. The ND is smaller and closer to the neutrino source, and interactions in certain regions of the detector could be affected differently by the selection criteria (in particular, containment). These effects are studied by using only one half of the near detector for the extrapolation, divided along one of its three dimensions (top vs bottom, east vs west, front vs back), or according to the distance from the center (inside vs outside). The largest deviation with respect to the nominal prediction was used to set the $\pm 1.2\%$ scale systematic [125].



(a) Q^2 distributions of ν_μ events selected in the ND and ν_e events in the FD. The ND events are reweighted according to their Q^2 to match the FD distribution.

(b) Comparison of the extrapolated ν_e signal prediction, with(out) reweighting the ν_μ events in the ND, in bins of reconstructed neutrino energy. Three variables, Q^2 , p_T/p and $\cos \theta$ are used to construct the weights; the largest differences to the nominal prediction are used for the systematic uncertainty.

Figure 5.17: The ν_μ selection in the ND and the ν_e selection in the FD can sculpt the characteristics of the samples differently and affect the $\nu_\mu \rightarrow \nu_e$ signal extrapolation. An acceptance systematic is constructed to account for this effect [33].

5.4 Cosmic background prediction

As it was pointed out in previous sections, the FD is located on the surface of the earth, making cosmic rejection a fundamental part of event selection. The cosmic rejection cuts from Section 5.1 are tuned using the minimum-bias cosmic trigger dataset, a sample independent of the beam data to be used for the measurement. To obtain an estimate of the cosmic background, we apply the tuned cuts on the NuMI-trigger FD dataset, with an additional timing cut. Out of the $500 \mu\text{s}$ recorded, we ignore $12 \mu\text{s}$ around the beam time (where the blinded signal lives) and $25 \mu\text{s}$ at the beginning and end of the stream for data quality. Using this dataset guarantees that the detector conditions of the estimate are a perfect match to the data, which can't be achieved with any other set, and that the characteristics of the reconstructed events correspond to actual backgrounds in the detector, which cannot be achieved with the simulation. In consequence, the cosmic background estimate is virtually free from systematic errors.




Figure 5.18: Distribution of ν_e -selected events in the out-of-time (cosmic) sideband in the FD. To obtain the estimated background to the beam signal, this histogram is scaled down to match the NuMI beam livetime.

Since the recorded time of this sample is larger than the beam spill, the histograms containing these events are scaled down by livetime, from 12886 to 434 seconds. The uncertainty introduced

by this scaling is 8.3%, and is included in the fit. The final distribution of the cosmic background estimate using the analysis bins is shown in Figure 5.18.

5.5 Final ν_e prediction and systematics

Figure 5.19 and Table 5.7 present the final ν_e prediction at given oscillation parameters, including the cosmic background estimate from the timing sideband, the extrapolation results ($\nu_\mu \rightarrow \nu_e$ CC signal; $\nu_\mu \rightarrow \nu_\mu$ CC, $\nu_e \rightarrow \nu_e$ CC and NC backgrounds), and the small components taken directly from the simulation ($\nu_e \rightarrow \nu_\mu$ CC, $\nu_{e,\mu} \rightarrow \nu_\tau$ CC). A small correction of 0.39 events from “rock” interactions that occur outside of the detector has been added to the beam components [124].




Figure 5.19: Predicted distribution of ν_e -selected events in the FD, including signal and background extrapolation and the cosmic background estimate. From [33].

Choosing a different set of oscillation parameters would produce different predictions above, with the largest changes in the $\nu_\mu \rightarrow \nu_e$ appearance channel. As an example, the predicted spectra for four sets of oscillation parameters are shown in Figure 5.20a. In Figure 5.20b we present the total number of events expected for a range of oscillation parameters. All in all, the total prediction

Table 5.7: Number of ν_e -selected events predicted in each category using the extrapolation and cosmic estimate. Oscillation parameters are as described in Section A.2. These are the integrals of Figure 5.19.

	LowCVN	MidCVN	HighCVN	Peripheral	Total
ν_e CC Signal	4.01	6.91	30	5.85	46.7
Beam ν_e CC	0.857	1.51	3.96	0.939	7.27
NC	3.24	1.85	1.09	0.223	6.4
ν_μ CC	0.576	0.333	0.194	0.161	1.26
ν_τ CC	0.0932	0.111	0.151	0.0286	0.384
Beam backgrounds	4.77	3.81	5.39	1.35	15.3
Cosmic background	1.6	0.612	0.136	2.55	4.9
Total	10.4	11.3	35.5	9.76	66.9

can range between 34 (IH $\delta_{CP} = \pi/2, \sin^2 \theta_{23} = 0.4$) and 75 (NH, $\delta_{CP} = 3\pi/2, \sin^2 \theta_{23} = 0.6$ total events over approximately 20 backgrounds (5 cosmogenic and 15 beam events)).



5.5.1 Systematic uncertainties

Recall that we evaluate systematic uncertainties by creating new extrapolated predictions from reweighted simulated events, or from alternative MC samples. Thanks to the two-detector technique, the ND measurements greatly reduce the size of the uncertainty, and only residual effects that can mimic the effect of oscillations on the predictions need to be included in the fit.

Table 5.8 summarizes the changes in the total signal and background predictions from these systematic variations. For the fit to data, the extensive list of systematics from Section 3.6 and the full spectral information are used.

Table 5.8: Effect of 1σ variations of the systematic uncertainties on the total ν_e signal and background predictions. Simulated data were used and oscillated with the parameters in Section A.2. As published in [47].

Source of uncertainty	ν_e signal (%)	Total beam background (%)
Cross sections and FSI	7.7	8.6
Normalization	3.5	3.4
Calibration	3.2	4.3
Detector response	0.67	2.8
Neutrino flux	0.63	0.43
ν_e extrapolation	0.36	1.2
Total systematic uncertainty	9.2	11
Statistical uncertainty	15	22
Total uncertainty	18	25

(a) Predicted spectra of ν_e candidates for four choices of oscillation parameters.(b) Total number of predicted ν_e candidates as a function of δ_{CP} . The color bands correspond to normal (blue, top) and inverted (red, bottom) hierarchies, and the range $0.43 \leq \sin^2\theta_{23} \leq 0.60$.Figure 5.20: Effect of neutrino oscillations in the predicted number of ν_e candidates.

CHAPTER 6. RESULTS OF THE JOINT $\nu_e + \nu_\mu$ ANALYSIS

In chapters 4 and 5, we detailed the procedures to select ν_μ and ν_e samples in the NOvA FD, and to construct predictions for ν_μ disappearance and ν_e appearance based on the simulation with constraints from ND and cosmic data. We also showed how the expected number of events and the shape of the predicted spectra change under different assumptions of the oscillation parameters. The comparison between the measured (unblinded) FD data spectra with the predictions allows us to make inferences about the underlying neutrino oscillation physics. At NOvA's 810 km baseline, the fit to ν_μ data is mostly sensitive to $\sin^2 \theta_{23}$ and Δm_{32}^2 , while the fit to ν_e data adds to the sensitivity to δ_{CP} , the θ_{23} octant and the neutrino mass hierarchy. In Section 6.1, we elaborate on the techniques for fitting the binned data, including the treatment of oscillation parameters, systematics, and Feldman-Cousins corrections. In Section 6.2, we use fake FD data to assess the sensitivity of the analysis. Finally, the distributions measured in the FD are revealed in Section 6.3, while the implications for neutrino oscillations are discussed in Section 6.4.

6.1 Fitting and determination of significance

In order to infer neutrino oscillation parameters from the FD data, NOvA uses a frequentist statistical approach [19]. Given a binned prediction for a vector of oscillation parameters $\vec{\theta}$, we compute the Poisson log-likelihood function

$$-2 \ln \lambda(\vec{\theta}) = 2 \sum_{i=1}^{\text{bins}} \left[\nu_i(\vec{\theta}) - n_i + n_i \ln \frac{n_i}{\nu_i(\vec{\theta})} \right], \quad (6.1)$$

where ν_i is the expected number of events in bin i and n_i is the observation. The sum runs over all the analysis bins (in our case, one ν_e and four ν_μ binned samples) [126]. We will render this quantity to χ^2 via

$$\chi^2 \sim -2 \ln \lambda(\vec{\theta}), \quad (6.2)$$

The vector of parameters that minimizes (6.1) is the best-fit result $\hat{\theta} \equiv \vec{\theta}_{\text{best}}$,

$$\chi_{\text{best}}^2 \equiv \chi^2(\hat{\theta}) = \min_{\vec{\theta}'} \left(\chi^2(\vec{\theta}') \right), \quad (6.3)$$

and we can define the test statistic $\Delta\chi^2$ relative to the minimum χ^2 as

$$\Delta\chi^2(\vec{\theta}) = \chi^2(\vec{\theta}) - \chi^2(\hat{\theta}). \quad (6.4)$$

With this definition, $\Delta\chi^2$ is non-negative and ordered so that if $\Delta\chi^2(\vec{\theta}_A) < \Delta\chi^2(\vec{\theta}_B)$, the combination of parameters $\vec{\theta}_A$ is a better fit to the data than $\vec{\theta}_B$. Strictly, it is the asymptotic distribution of $-2 \ln \lambda(\vec{\theta})$ that follows a χ^2 distribution, given some regularity conditions (Wilk's theorem) [127]. We will use the ‘‘Gaussian approximation’’ $-2 \ln \lambda(\vec{\theta}) = \chi^2(\vec{\theta})$ for sensitivities and first-order assessments, and the unified approach of Feldman and Cousins [128] to construct confidence intervals for the final results.

6.1.1 Treatment of oscillation parameters and construction of contours

The underlying physics model in our analysis is three-flavor neutrino oscillations including matter effects, using the full formulae introduced in Section 1.5. The number of degrees of freedom is reduced by holding the following parameters fixed:

$$L = 810 \text{ km}, \quad (6.5)$$

$$\rho = 2.84 \text{ g/cm}^3, \quad (6.6)$$

$$\Delta m_{21}^2 = 7.53 \times 10^{-5} \text{ eV}^2, \quad (6.7)$$

$$\sin^2 2\theta_{12} = 0.851. \quad (6.8)$$

The solar oscillation parameters are the 2017 PDG values [19]; ρ is the density of the earth estimated with the CRUST 2.0 model [129], using the average depth underground between the two detectors (9.38 km) [130]. Small variations in any of these parameters have negligible effect in the results.

While NOvA's ν_e and ν_μ oscillation channels are sensitive to θ_{13} , reactor experiments have already measured this parameter to a precision higher than current long-baseline experiments can

achieve. Thus θ_{13} is allowed to vary in the fit, treated as a constrained nuisance parameter. We use the average result [19]

$$\sin^2 2\theta_{13} = 0.082 \pm 0.004, \quad (6.9)$$

and add a Gaussian penalty term to the test statistic

$$\chi'^2 = \chi^2 + \frac{(x - \mu)^2}{\sigma^2}, \quad (6.10)$$

where we assumed the form $x = \mu \pm \sigma$. Using this “reactor constraint” guarantees that our results for the other parameters are compatible with current measurements. Alternatively, fit results without such constraint can be used as a cross check when directly compared to different experiments, but the lower precision would make NOvA measurements less interesting.

The remaining parameters (δ_{CP} , Δm_{32}^2 and $\sin^2 \theta_{23}$) are free in the fit. While the neutrino mass hierarchy is not an additional parameter per se ($\sim \text{sign}(\Delta m_{32}^2)$), the combination of the symmetry in the first-order terms of the oscillation probabilities and the precision of the experiment create two disjoint regions centered around $|\Delta m_{32}^2| \sim 2.5 \times 10^{-3} \text{eV}^2/c^4$. Therefore, we display the results for normal and inverted hierarchy separately, and use “fitting over the mass hierarchy” to signify the use of both positive and negative seeds for Δm_{32}^2 in the minimization algorithms.

Finally, in order to present the results, we construct 1- or 2-dimensional graphs based on the $\Delta\chi^2$ distributions, where the remaining parameters are treated as nuisance (“profiled”). To first order, this can be understood as taking projections of a multi-dimensional $\Delta\chi^2(\theta_{23}, \Delta m_{32}^2, \delta_{CP}, \theta_{13}, \dots)$, using the minimum value over all the other parameters that are not the axes of interest. The alternative method, “marginalization”, would use a weighted average instead of the minimum, but requires prior assumptions on the nuisance parameters.

6.1.2 Treatment of systematic uncertainties

Systematic uncertainties are incorporated in the likelihood function as nuisance parameters [131]; each one of them is associated with an additional free term in the fit, which requires an additional measurement or external constraint. As mentioned earlier, the near-to-far extrapolation

procedure aids in this function: predictions are created from the systematically-shifted MC samples and the ND data, effectively correcting the central values. Any leftover discrepancies between the nominal and the shifted predictions determine the allowed ranges for the systematic pulls. Penalty terms are added to the χ^2 to keep their pulls in the 1σ ranges

$$\chi^2(\vec{\theta}) = \min_{\vec{s}} \left(\chi^2(\vec{\theta}, \vec{s}) + \sum_i^{\text{sys}} \frac{s_i^2}{\sigma_i^2} \right) \quad (6.11)$$

where s_i are the values of the individual systematic shifts, and σ_i are the 1-sigma ranges used to assign the penalty term for each systematic [114]. This method assumes that all sources of uncertainty are independent of each other. Alternative approaches could use a covariance matrix method, or a simultaneous two-detector fit.

Recall from Section 3.6 that 28 sources of systematic uncertainty are included in this analysis, and for presentation they are divided under the categories of calibration, cross sections and final state interactions, neutrino flux, detector response, and others. Given that the ν_e and ν_μ predictions are created following the same procedure, the correlations among samples are maintained by construction.

6.1.3 Confidence intervals and Feldman-Cousins corrections

While the best fit result $\hat{\theta}$ provides the point estimates of the oscillation parameters given the data, we also want to express the statistical precision of the measurement. In the frequentist approach, “confidence intervals” are regions of parameter space that include, by construction, their true values, with a probability higher than a certain level (“coverage probability”) [19].

The Neyman construction [132] is the basic procedure to create frequentist intervals, appropriate for multi-dimensional problems with nuisance parameters [131]. We use the test statistic based on the likelihood ratio defined earlier in (6.4) as the ordering rule of the construction, which is known in the field as the Unified Approach of Feldman and Cousins [128]. The intervals are constructed as

$$C_{1-\alpha} = \left\{ \vec{\theta} \mid \chi_{\text{data}}^2(\vec{\theta}) - \chi_{\text{data}}^2(\hat{\theta}) \leq k_{1-\alpha}(\vec{\theta}) \right\}, \quad (6.12)$$

with $C_{1-\alpha}$ the set of $\vec{\theta}$ values for which $\hat{\theta}$ falls in the region of sampling probability $1 - \alpha$ where the log-likelihood ratio in favor is better than the average outside the region, and $k_{1-\alpha}(\vec{\theta})$ is a constant fixed by the coverage requirement at the given $\vec{\theta}$.

In the Gaussian approximation, a $(1 - \alpha)$ CL interval can be simply constructed by taking the set of values of θ such that

$$\Delta\chi^2(\vec{\theta}) \leq \chi_{d,1-\alpha}^2 \quad (6.13)$$

with $\chi_{d,1-\alpha}^2$ the $(1 - \alpha)$ quantile of a χ^2 distribution with d degrees of freedom, and is independent of $\vec{\theta}$. For example, all parameters with a $\Delta\chi^2 < 2.71$ (4.61) are allowed at 90% C.L. for a 1D (2D) limit. Frequently, these levels are displayed as a $N\sigma$ significance, mapping back to two-sided probabilities in the tails of a 1-dimensional Gaussian distribution (e.g. $1 - \alpha = 68.27\% \rightarrow 1\sigma$). In high energy physics, 3σ is deemed as evidence against a null hypothesis, and 5σ is considered enough for discovery [131].

Neutrino oscillation measurements tend to move away from the Gaussian approximation because the sizes of the samples can be small, factors such as $\sin^2 \theta_{23}$ are near physical boundaries, and δ_{CP} is a cyclical parameter to which current experiments including NOvA have very little sensitivity. Therefore, the Gaussian approximation might not suffice, and the $k_{1-\alpha}(\vec{\theta})$ limits from (6.12) need to be computed to obtain the physics results.

In order to construct the 1- or 2-dimensional confidence intervals, we create an ensemble of mock experiments ($\mathcal{O}(10^4)$) at each point of parameter space and compute $\Delta\chi^2 = \chi_{\text{true}}^2 - \chi_{\text{best}}^2$, where χ_{true}^2 and χ_{best}^2 are the best fits to the mock data keeping the known parameters fixed or allowing them to float, respectively. In other words, we are probing how often the statistical fluctuations in the mock data with some hypothesis of oscillation parameters look more compatible with some *wrong* combination of parameters.

For this process, systematic and oscillation parameters that are not the axes of the desired graphs are profiled over. The resulting $\Delta\chi^2$ distributions are used to obtain critical values for the 68%, 95%, etc., C.L. intervals for each set of $\vec{\theta}$, and can be compared to the measured $\Delta\chi_{\text{data}}^2(\vec{\theta})$ to get the significance of the result at the given point. We often refer to these results as having


“Feldman-Cousins corrections”, in the sense that we can quickly use the Gaussian critical values (6.13) as an initial approximation, and later evaluate deviations using the mock experiments for the final results, since they demand a considerable amount of computational resources.

One of the difficulties of this method is to ensure proper coverage in the presence of nuisance parameters, since their true values are unknown. For this analysis we use the best fit parameters of the data in the creation of the mock experiments. This gives coverage for the most likely systematic shifts given the data measured and the hypothesis of oscillation parameters [114, 92].

6.2 Physics sensitivity

In order to understand the potential of the NOvA oscillation measurements before unblinding the data, we can study the $\Delta\chi^2$ distributions using fake data with some known oscillation parameters. We use the Asimov prediction [133], which means that the fake data histograms are not forced to have an integer number of events per bin (unlike real data), but do have the statistical errors expected for the exposure of 8.85×10^{20} POT-equiv. In addition, these sensitivities are created following the procedure from Section 6.1 in the Gaussian approximation. They use the ν_e and the four ν_μ extrapolated predictions, and the reactor constraint $\sin^2 2\theta_{13} = 0.082 \pm 0.004$, and simultaneously fit δ_{CP} , $\sin^2 \theta_{23}$, Δm_{32}^2 and $\sin^2 2\theta_{13}$. The contours in Figures 6.1 and 6.2 correspond to $\sqrt{\Delta\chi^2} = 2.30, 6.18, 11.83$ levels for $\{\delta_{CP}, \sin^2 \theta_{23}\}$, and $\{\sin^2 \theta_{23}, \Delta m_{32}^2\}$, profiling over the two variables that are not represented in the axes. The fake data parameters used in Figure 6.1 are the 2016 best fit results [76] listed in Section A.1, and Figure 6.2 uses the parameters in Section A.2. While the contours have different shapes depending on the choice of parameters, it should be noted that all values of δ_{CP} , both hierarchies and both octants would still be allowed according to data. Occasionally certain combinations, such as maximal mixing or $\delta_{CP} = \pi/2$, might be rejected at higher than 3σ C.L.

In Fig. 6.3 we use a similar approach to obtain the potential to determine the neutrino mass hierarchy, CP violation, non-maximal mixing and the θ_{23} octant, for a wider set of true parameters.



(a) 1, 2 and 3 σ allowed regions in $\sin^2 \theta_{23}$ vs δ_{CP} (b) 1, 2 and 3 σ allowed regions in Δm_{32}^2 vs $\sin^2 \theta_{23}$ space, and normal (top) and inverted (bottom) hierarchy. Profiling over $\sin^2 2\theta_{13}$ and Δm_{32}^2 (b) 1, 2 and 3 σ allowed regions in Δm_{32}^2 vs $\sin^2 \theta_{23}$ space, and normal (top) and inverted (bottom) hierarchy. Profiling over $\sin^2 2\theta_{13}$ and δ_{CP} .

Figure 6.1: Sensitivities for the 2017 joint $\nu_e + \nu_\mu$ analysis, assuming the true oscillation parameters from Section A.1 (NH, $\delta_{CP} = 1.48\pi$, $\sin^2 \theta_{23} = 0.404$, $\Delta m_{32}^2 = 2.7 \times 10^{-3} eV^2$).

They are calculated as follows:


$$\text{Hierarchy: } \Delta\chi_{\text{best} | \text{wrong hierarchy}}^2 - \Delta\chi_{\text{best}}^2,$$

$$\text{CP violation: } \Delta\chi_{\text{best} | \delta_{CP}=0, \pi}^2 - \Delta\chi_{\text{best}}^2,$$

$$\text{Maximal mixing: } \Delta\chi_{\text{best} | \sin^2 \theta_{23}=0.5}^2 - \Delta\chi_{\text{best}}^2,$$

$$\text{Octant: } \Delta\chi_{\text{best} | \text{wrong octant}}^2 - \Delta\chi_{\text{best}}^2.$$

These graphs showcase the combinations of δ_{CP} , mass hierarchy and octant that have good potential for the determination of the mass hierarchy or octant, and the low sensitivity of the experiment to CP violation at the current level of statistics. The relative flat shape of the rejection



(a) 1, 2 and 3 σ allowed regions in $\sin^2 \theta_{23}$ vs δ_{CP} (b) 1, 2 and 3 σ allowed regions in Δm_{32}^2 vs $\sin^2 \theta_{23}$ space, and normal (top) and inverted (bottom) hierarchy. Profiling over $\sin^2 2\theta_{13}$ and Δm_{32}^2 space, and normal (top) and inverted (bottom) hierarchy. Profiling over $\sin^2 2\theta_{13}$ and δ_{CP} .

Figure 6.2: Sensitivities for the 2017 joint $\nu_e + \nu_\mu$ analysis, assuming the true oscillation parameters from Section A.2 (NH, $\delta_{CP} = 1.21\pi$, $\sin^2 \theta_{23} = 0.56$, $\Delta m_{32}^2 = 2.44 \times 10^{-3} eV^2$).

of maximal mixing as a function of δ_{CP} indicates how the $\nu_\mu \rightarrow \nu_\mu$ channel has higher sensitivity to $\sin^2 \theta_{23}$ than $\nu_\mu \rightarrow \nu_e$.

6.2.1 Systematic effects on the measurement of oscillation parameters

We estimate the impact of systematic uncertainties on the measurement of oscillation parameters by running a fit using predictions created with fake ND data. This approach cleanly separates the effect of extrapolating real ND data from the effect of the systematic [112, 97].

The relative size of the impacts are quantified using the increase in the one-dimensional 68% C.L. interval, relative to the size of the interval when only statistical uncertainty is included in the




Figure 6.3: Potential rejection of the wrong hierarchy, CP conservation, maximal mixing and the wrong octant, as functions of δ_{CP} , assuming true normal (blue) or inverted (red) hierarchy, and θ_{23} in the lower (solid) or upper (dashed) octant. For the lower (upper) octant, we used true values of $\sin^2 \theta_{23} = 0.404$ (0.623), and constrained the $\sin^2 \theta_{23}$ fit result to be in the same octant. Other true oscillation parameters: $\Delta m_{32}^2 = \pm 2.5 \times 10^{-3} eV^2$, $\theta_{13} = 0.022$.

fit [134, 47]. Simulated data were used and oscillated with the same parameters as in Table 5.8. Given the asymmetry of the $\sin^2 \theta_{23}$ interval with respect to its best fit value, only the change in the upper edge is included. The total systematic uncertainty is calculated by adding the individual components in quadrature.

Table 6.1: Sources of uncertainty and their estimated average impact on the oscillation parameters in the joint fit. As published in [47].

Source of uncertainty	Uncertainty in $\sin^2 \theta_{23}$ ($\times 10^{-3}$)		Uncertainty in Δm_{32}^2 ($\times 10^{-6} \text{eV}^2/c^4$)		Uncertainty in δ_{CP}
Calibration	+	7.3	+	27 / - 27	$\pm 0.05\pi$
Cross sections and FSI	+	6.9	+	14 / - 19	$\pm 0.08\pi$
Muon energy scale	+	2.4	+	8.5 / - 12	$\pm 0.01\pi$
Normalization	+	4.4	+	7.3 / - 12	$\pm 0.05\pi$
Detector response	+	0.8	+	6.2 / - 7.7	$\pm 0.01\pi$
Neutrino flux	+	1.1	+	4.0 / - 4.4	$\pm 0.01\pi$
ν_e extrapolation	+	0.1	+	0.2 / - 0.7	$\pm 0.01\pi$
Total systematic uncertainty	+	12	+	33 / - 38	$\pm 0.12\pi$
Statistical uncertainty	+	38	+	75 / - 84	$\pm 0.66\pi$
Total uncertainty	+	40	+	82 / - 92	$\pm 0.67\pi$

6.3 Far detector data

Using the ν_μ CC selections from chapter 4 we observe 126 candidates in the FD. In the absence of oscillations, we would have expected 720_{-47}^{+67} (syst.) total events, including a predicted background of 3.4 beam and 5.8 cosmic events. The observed energy spectra split by quantile are shown in Figure 6.4, and added together in Figure 6.5. The combined distribution highlights the sharp dip around 1.6 GeV and a good agreement between data and MC. Recall that the values of $\sin^2 2\theta_{23}$ and Δm_{32}^2 essentially correspond to the depth and location of the dip, respectively. While the bin below 3 GeV presents a rather large fluctuation, its location away from the dip means that we expect the effect on the measurement to be small. Additionally, these measurements are sensitive to the energy resolution, and we expect the quartile with best energy resolutions and lower backgrounds to have the highest impact.

Similarly, using the selection from Chapter 5, we observe 66 ν_e CC candidates in the FD over an expected background of 20.3 ± 2.0 (syst.) events (7.3 beam ν_e CC, 6.4 NC, 1.3 ν_μ CC, 0.4 ν_τ CC, 4.9 cosmic). Figure 6.6 shows the distribution of these events in the ν_e analysis bins (as a function of




Figure 6.4: Comparison of the reconstructed energy spectra of selected ν_μ CC candidates in FD data (black dots) and best-fit prediction (purple). The sample is split into four reconstructed hadronic energy fraction quartiles labeled 1 through 4, where 1 (4) has the best (worst) energy resolution. The majority of the total background (gray, upper) including the cosmogenic subcomponent (blue, lower) lies in the fourth quartile. As published in [47].

the reconstructed neutrino energy for the three CVN classifier bins and for the peripheral sample), overlaid with the best fit prediction.

In Figure 6.7, we compare the observed number of ν_e candidates to the predictions from Figure 5.20b. While the binned log-likelihood fit uses more information, this figure gives some context to the observation, hinting at the normal mass hierarchy and δ_{CP} values away from $\pi/2$ as the most compatible with the data.

Finally, Figures 6.8 and 6.9 give additional context to the accumulation of ν_e data in NOvA over time. Unexpected (a)symmetries could indicate problems in the data acquisition. In Figure 6.8, the out-of-time events seem randomly distributed on both sides (before/after) the beam window,




Figure 6.5: Data from Fig. 6.4 summed over the four quartiles.

and across time. These are the same events used to construct the background estimated from Figure 5.18. Similarly, both 6.8 and 6.9 show a normal increase in the number of signal candidates as a function of the exposure.




Figure 6.6: Comparison of the neutrino energy spectra of selected ν_e CC candidates in the FD data (black dots) with the best fit prediction (purple lines) in the three CVN classifier bins and the peripheral sample. The total expected background (gray, upper) and the cosmic component of it (blue, lower) are shown as shaded areas. The events in the peripheral bin have energies between 1 and 4.5 GeV. As published in [47].




Figure 6.7: Total number of ν_e CC candidate events observed in the FD (gray) compared to the prediction (color) as a function of δ_{CP} . The color lines correspond to the best fit values of $\sin^2 \theta_{23}$ and Δm_{32}^2 from Table 6.2, with the upper two curves (blue) representing two octants in the normal mass hierarchy ($\Delta m_{32}^2 > 0$) and the lower curve (red) the inverted hierarchy ($\Delta m_{32}^2 < 0$). The color bands correspond to $0.43 \leq \sin^2 \theta_{23} \leq 0.60$. All other parameters are held fixed at the best-fit values. As published in [47].




Figure 6.8: ν_e -selected events in the beam time window (black) and out-of-time (red). The gray and cyan bands represent accelerator shutdown periods and antineutrino (RHC) mode, respectively. The NuMI timing window is outlined with blue. The second shifted timing window is a result of GPS initialization issue, removed with a firmware patch in October 2014. From [33].




Figure 6.9: Accumulation of the 66 ν_e candidates over time, overlaid with the accumulated exposure (POT weighted by the fiducial mass). A Kolmogorov-Smirnov test was applied on the shapes of distributions. From [33].

6.4 Joint fit results

The values of the oscillation parameters at the best fit point in the two mass hierarchies and octants are given in Table 6.2. The first line is the best fit overall, corresponding to the normal hierarchy and upper $\sin^2 \theta_{23}$ octant, and used for the predictions in Figures 6.4 and 6.6. The other two points in the table are the normal hierarchy, lower octant result, which is slightly less significant, and the best fit in the inverted hierarchy. As hinted earlier in Figure 6.7, the predicted number of events in the inverted hierarchy, lower octant, is much lower than the measured data, and thus highly disfavored. The χ^2 for the overall best fit is 84.6 for 72 degrees of freedom. Figure 6.10 shows the systematic shifts contributing to the $\Delta\chi^2$, as explained in Section 6.1.2. The biggest pulls come from calibration and neutrino interaction systematics.

Table 6.2: Best fit values for the 2017 NOvA oscillation analysis. As published in [47].

Hierarchy/Octant	$\delta_{CP} (\pi)$	$\sin^2 \theta_{23}$	$\Delta m_{32}^2 (10^{-3} \text{eV}^2/c^4)$	$\Delta\chi^2$
Normal/Upper	1.21	0.56	2.44	0.00
Normal/Lower	1.46	0.47	2.45	0.13
Inverted/Upper	1.46	0.56	-2.51	2.54




Figure 6.10: Systematic pulls (in units of $N\sigma$) on the best fit point for each of the uncertainties included in the fit.

All of the contours and significance levels that follow are constructed following the unified approach of Feldman and Cousins [128], profiling over unspecified physics parameters and systematic uncertainties.




Figure 6.11: Regions of parameter space consistent with the ν_e appearance and the ν_μ disappearance data at various levels of significance. The top panels corresponds to normal mass hierarchy and the bottom panels to inverted hierarchy. The color intensity indicates the confidence level at which particular parameter combinations are allowed. As published in [47].

Figures 6.11a and 6.10b show the 1, 2, and 3 σ two-dimensional contours for $\sin^2 \theta_{23}$ vs. δ_{CP} and Δm_{32}^2 vs. $\sin^2 \theta_{23}$, respectively. Additionally, Figure 6.12 compares the 90% C.L. Δm_{32}^2 vs. $\sin^2 \theta_{23}$ contours in the normal hierarchy with the corresponding results of T2K [108], MINOS [135], IceCube [136], and Super-Kamiokande [137]. At the current levels of sensitivity, all these results are compatible with each other and compatible with maximal mixing.

The contours in Figure 6.11 can be compared to the sensitivities from Figure 6.2, keeping in mind that differences arise from the Feldman-Cousins corrections and the use of real data vs.




Figure 6.12: Comparison of measured 90% confidence level contours for Δm_{32}^2 vs. $\sin^2 \theta_{23}$ for this result (black line; best-fit value, black point), T2K [108] (green dashed), MINOS [135] (red dashed), IceCube [136] (blue dotted), and Super-Kamiokande [137] (purple dash-dotted). As published in [47].


the Asimov predictions. Both figures share the same general features, only differing by a slightly higher rejection of $\delta_{CP} = \pi/2$, and the appearance of a 1σ island in the inverted hierarchy Δm_{32}^2 vs. $\sin^2 \theta_{23}$ contour that resulted from the corrections.

Table 6.3: 1σ confidence intervals for physics parameters in the normal mass hierarchy. As published in [47].

Parameter (units)	1σ interval(s)
Δm_{32}^2 ($10^{-3}\text{eV}^2/c^4$)	[2.37, 2.52]
$\sin^2 \theta_{23}$	[0.43, 0.51] and [0.52, 0.60]
δ_{CP} (π)	[0, 0.12] and [0.91, 2]


Figures 6.13a, 6.13b and 6.13c show the significance results for δ_{CP} , $\sin^2 \theta_{23}$ and Δm_{32}^2 , for the two mass hierarchies. Table 6.3 shows the corresponding 1σ confidence intervals. Note that the minima in the inverted hierarchy correspond to different significances even though the $\Delta\chi^2$ are identical. This is a consequence of the Feldman-Cousins corrections, since the parameters that are profiled can lead to a different translation from $\Delta\chi^2$ to significance (unlike the Gaussian approximation, where the mapping would be one-to-one). The near-degeneracy between the two

octants in the normal mass hierarchy mentioned earlier can be readily observed in figures 6.13a and 6.13b, in contrast with the inverted hierarchy where only one local minimum is present. Our data disfavor maximal mixing by 0.8σ , $\delta_{CP} = \pi/2$ in the inverted hierarchy at higher than 3σ , and the entire inverted mass hierarchy at the 95% confidence level [47].



(a) Significance at which each value of δ_{CP} is disfavored in the normal (blue, lower) or inverted (red, upper) mass hierarchy. The normal mass hierarchy is divided into upper (solid) and lower (dashed) θ_{23} octants corresponding to the near degeneracy in $\sin^2 \theta_{23}$.

(b) Significance at which each value of $\sin^2 \theta_{23}$ is disfavored in the normal (blue, lower) or inverted (red, upper) mass hierarchy. The vertical dotted line indicates the point of maximal mixing.



(c) Significance at which each value of $|\Delta m_{32}^2|$ is disfavored in the normal (blue, lower) or inverted (red, upper) mass hierarchy.

Figure 6.13: One-dimensional significances for δ_{CP} , $\sin^2 \theta_{23}$ and Δm_{32}^2 . As published in [47].

CHAPTER 7. PROJECTED MEASUREMENTS WITH NOvA




Figure 7.1: $\bar{\nu}_\mu$ (left) and $\bar{\nu}_e$ (right) antineutrino candidates in the FD. Since the NOvA detectors are not magnetized, the signatures of ν and $\bar{\nu}$ CC candidates are almost identical. From [33].

In the previous chapters, we reviewed the analysis techniques and results of the joint $\nu_e + \nu_\mu$ oscillation analysis, with a dataset corresponding to an exposure of 8.85×10^{20} POT-eq. with the NuMI beam in neutrino mode. Improvements in the measurement of oscillation parameters in NOvA will arise primarily from increased statistics and the combination of neutrino and antineutrino data.

We are particularly interested in the ν_e appearance channel, since the differences in the oscillation probabilities of neutrinos and antineutrinos as a function of δ_{CP} and the neutrino mass hierarchy create measurable effects. In Section 7.1 we describe these effects, and discuss the implications for the NOvA run plan in Section 7.2. We explore the potential for determination of the neutrino mass hierarchy, CP violation, maximal mixing and octant in the upcoming years in Section 7.3. Further improvements to the experiment or the analysis are briefly stated in Section 7.4. Finally, Section 7.5 summarizes the current and future contributions of the NOvA experiment to the global knowledge of neutrino oscillations.

7.1 Antineutrino beam running

Recall from Section 2.2 that the NuMI beam has two main configurations depending on the direction of the horn current, FHC and RHC, that give a neutrino beam dominated by muon neutrinos or antineutrinos, respectively. Since the antineutrino cross sections are smaller than their neutrino counterparts, an antineutrino run will naturally have fewer events and a higher content of wrong sign neutrinos for the same exposure in both detectors.

Additionally, the functional forms of $\nu_\mu \rightarrow \nu_e$ oscillation probabilities including matter effects from Section 1.5 show opposing enhancement/deficit effects for neutrinos vs. antineutrinos depending on the values of δ_{CP} and the neutrino mass hierarchy. This means that a combination of parameters that is favorable in FHC such as $\delta_{CP} = 3\pi/2$ in the normal hierarchy, would give a low event count in RHC with a high incidence of wrong-sign background.




Figure 7.2: Expected total number of ν_e candidates in the NOvA FD for an exposure of 9×10^{20} POT in neutrino (left) or antineutrino (right) mode, as a function of δ_{CP} . The color bands correspond to normal (blue) or inverted (red) hierarchy, and the range $0.4 \leq \sin^2 \theta_{23} \leq 0.6$, and the solid lines represent maximal mixing $\sin^2 \theta_{23} = 0.5$.

Figure 7.2 shows the expected number of ν_e selected events for an identical exposure of 9×10^{20} POT in neutrino or antineutrino mode, for a range of oscillation parameters. Notice that the scales of the two graphs differ, with maximum values of $\mathcal{O}(70)$ vs. $\mathcal{O}(30)$ events in each mode. The maxima are expected for opposite hierarchies and CP violating values of δ_{CP} . Only the behavior with respect to $\sin^2 \theta_{23}$ is preserved in both cases, meaning that the upper octant will always

produce a higher event count than the lower octant for a given hierarchy and δ_{CP} . These event numbers are represented in two dimensions in Figure 7.3. Points that are further to the right (top) have higher number of neutrinos (antineutrinos), with the extreme cases corresponding to $\delta_{CP} = \pi/2, 3\pi/2$.




Figure 7.3: Expected number of ν_e candidates in neutrino and antineutrino mode for given combinations of oscillation parameters, for a total exposure of 18×10^{20} POT split equally between the two modes. Select curves from Figure 7.2 can be combined in two dimensions to form ellipses, each one representing the full δ_{CP} range for a set of $\sin^2 \theta_{23}$ and Δm_{32}^2 . Modified from [138, 139].

7.2 Run plan optimization

The considerations in the previous section mean that, depending on the true values of the oscillation parameters, there can be an optimal balance between neutrino and antineutrino running that would lead to better measurements in less time. To illustrate this point, consider the alternative versions of Figure 7.3 in Figure 7.4.




Figure 7.4: Comparison of measurement scenarios overlaid on the two-dimensional ν_e and $\bar{\nu}_e$ event count expectations. The bands represent uncertainties in a ν_e event count in FHC (vertical band) or RHC (horizontal band), and the stars are possible true values. These cartoons are for illustration only and do not reflect complete analysis results. Modified from [138, 139].

Starting from a FHC-only measurement like the vertical yellow band in Figure 7.4a where both octants are allowed, adding RHC data in the form of the horizontal bands can lead to disambiguation while adding FHC data could not. In Figure 7.4b, if the true oscillation parameters are those that lead to the highest possible count in neutrino mode, adding FHC data can lead to determination of the mass hierarchy, $\sin^2 \theta_{23}$ octant and possibly the rejection of CP conservation faster than any alternative. In a case like Figure 7.4c, it is unlikely that we would ever be able to choose normal over inverted hierarchy, therefore other measurements could be prioritized. If the data are highly compatible with $\sin^2 \theta_{23} = 0.5$ (the ellipses collapse toward the center), the determination of the octant and rejection of maximal mixing are out of the question, and only the determination of hierarchy and CP violation would be relevant for the run plan.

It is worth pointing out that we are not trying to differentiate among these curves presented in Figure 7.4, but an infinite family of them covering the allowed ranges of $\sin^2 \theta_{23}$, Δm_{23}^2 and $\sin^2 2\theta_{13}$. The ν_μ data helps constraining the first two parameters with increasing precision, but the dependency of the ν_μ survival probability on $\sin^2 2\theta_{23}$ instead of $\sin^2 \theta_{23}$ implies that it cannot

lift the degeneracy. Similarly, it is not possible to determine the neutrino mass hierarchy from ν_μ data alone.

Another important observation is that given the *separation* between δ_{CP} conserving values $(0, \pi)$, the lack of knowledge of the other oscillation parameters creates a wide range of points that are compatible with CP conservation, making a high-confidence determination of CP violation difficult to achieve with the NOvA experiment during its expected lifetime. We should also point out that there are many combinations of oscillation parameters where definitive (3 or 5σ) measurements of the neutrino mass hierarchy, octant or maximal mixing, might not be achieved by the end of the NOvA experiment, and we would only be able to set limits with increasing precision. For that reason, other considerations for a run plan include the physics goals of other experiments that use the NuMI beam [140, 141], and non-oscillation (or non-standard oscillations) measurements with NOvA.

Depending on the oscillation scenarios, we can try to find run plans that have the optimal FHC/RHC proportions to achieve a very good measurement, but it is not possible to do it for all the values of interest without making very strong assumptions [142, 143]. Given that at our current precision both hierarchies, octants and all values of δ_{CP} are still allowed, we will only explore possible milestones with an even split of the beam between neutrino and antineutrino mode after 2019. Table 7.1 breaks down the expected accumulated POT as a function of year. The POT accumulated follows the design conditions of the NuMI beam, plus the assumptions of 40 weeks of 98% uptime per year, and the planned ramp up of the beam power from the current 700 kW to a proposed 1000 kW by 2023 [141]. It is expected that NuMI will go off-line after 2024 to allow the modifications to the accelerator complex required for the future LBNF/DUNE neutrino beamline [144].

Table 7.1: Assumptions for the beam exposure for projected sensitivities. The 2017-18 values correspond to the analyzed datasets. Other values are the projected end-of-year accumulated POT, with an even split between neutrino (FHC) and antineutrino (RHC) modes. Assuming beam power ramping up to 1000 kW (in three 100 kW increments), beam running 40 weeks per year after 2019, and a combined beam and detector uptime of 98%. From [145].

Year	10^{20} POT FHC	10^{20} POT RHC
2017	9.5	0
2018	9.5	7.0
2019	12.6	12.6
2020	16.7	16.7
2021	21.2	21.2
2022	25.8	25.8
2023	30.9	30.9
2024	35.9	35.9

7.3 Physics reach of NOvA

In this section, we present the sensitivities of the NOvA oscillation analysis to the determination of the neutrino mass hierarchy, CP violation, the θ_{23} octant, and the rejection of maximal mixing. These are based on a joint $\nu_e + \nu_\mu$ analysis of neutrino and antineutrino data, using the Asimov predictions as fake data to assess the average sensitivity [143]. The methodology is identical to the one presented in Sections 6.1 and 6.2, with the addition of RHC predictions and variable POT assumptions.

While we cover a wide range of oscillation parameters, statistical fluctuations and appropriate corrections to the significances using the Unified Approach would modify these results.

The definitions of the quantities of interest are the ones presented in Section 6.2:

$$\begin{aligned}
 \text{Hierarchy:} & \quad \Delta\chi_{\text{best} \mid \text{wrong hierarchy}}^2 - \Delta\chi_{\text{best}}^2, \\
 \text{CP violation:} & \quad \Delta\chi_{\text{best} \mid \delta_{CP}=0,\pi}^2 - \Delta\chi_{\text{best}}^2, \\
 \text{Maximal mixing:} & \quad \Delta\chi_{\text{best} \mid \sin^2 \theta_{23}=0.5}^2 - \Delta\chi_{\text{best}}^2, \\
 \text{Octant:} & \quad \Delta\chi_{\text{best} \mid \text{wrong octant}}^2 - \Delta\chi_{\text{best}}^2.
 \end{aligned}$$

As stated earlier, even if some of the highlighted milestones are not achieved, we expect that the precision of the $\sin^2 \theta_{23}$ and Δm_{32}^2 estimates will increase over time, and larger sections of

parameter space will be excluded. These results are difficult to quantify and present generically and thus will not be included in this document.

7.3.1 Determination of mass hierarchy

The NOvA experiment has the potential for the determination of the neutrino mass hierarchy. This would have a very high impact in the field of neutrino physics, and in particular the design and expected performance of future long-baseline experiments [144, 146] and direct neutrino mass measurements [20].




Figure 7.5: Projected significance of resolution of the neutrino mass hierarchy with NOvA as a function of year, assuming true normal (left) or inverted (right) hierarchy. The color bands correspond to different true values of δ_{CP} , and a range of $\sin^2 \theta_{23} = 0.4 - 0.6$. Joint fit combines the electron neutrino appearance and muon neutrino disappearance channels, and global reactor constraint $\sin^2 \theta_{13} = 0.082 \pm 0.004$. Assuming an equal fraction of beam exposure in neutrino and antineutrino mode after 2019, and a total of 72×10^{20} POT by 2024 (details in table 7.1). From [145].

In Figure 7.5 we present the projected significance for hierarchy determination as a function of year for a range of oscillation parameters. The most interesting combinations are true NH $\delta_{CP} = 3\pi/2$ and true IH $\delta_{CP} = \pi/2$: for some values of $\sin^2 \theta_{23}$ we could achieve a 3σ result by 2020, and 5σ or higher by 2024.

In addition to these extreme cases, Figure 7.6 shows that we can expect to know the mass hierarchy at 95% C.L. for more than 40% of the true values of δ_{CP} in the $[0, 2\pi)$ range, independently of $\sin^2 \theta_{23}$ and the assumed true hierarchy.

Finally, Figure 7.7 provides a snapshot of the results by 2024. While there are a few special values of δ_{CP} as the ones mentioned above where we can determine the neutrino mass hierarchy at higher than 4σ , it is still safe to expect a $> 3\sigma$ measurement for $\sim 50\%$ of the cases.




Figure 7.6: Fraction of δ_{CP} values in the range $[0, 2\pi)$ for which the neutrino mass hierarchy could be resolved by NOvA at a 95% C.L. or higher, as a function of year. The color lines correspond to different true values of $\sin^2 \theta_{23}$. Joint fit combines the electron neutrino appearance and muon neutrino disappearance channels, and global reactor constraint $\sin^2 \theta_{13} = 0.082 \pm 0.004$. Assuming an equal fraction of beam exposure in neutrino and antineutrino mode after 2019, and a total of 72×10^{20} POT by 2024 (details in table 7.1). From [145].

7.3.2 Determination of CP violation

We mentioned in Section 7.2 how the determination of CP violation is a difficult measurement with NOvA. Given the wide range of neutrino oscillation parameters that are allowed, and the ambiguities that systematic uncertainties and finite statistics introduce, the simultaneous rejection of the two CP conserving values of $\delta_{CP} = 0, \pi$ is very unlikely with the exposure that NOvA is expected to accumulate by 2024.




Figure 7.7: Fraction of δ_{CP} values in the range $[0, 2\pi)$ for which the neutrino mass hierarchy could be resolved by NOvA at a given sensitivity σ by the year 2024. The lines correspond to different true values of $\sin^2 \theta_{23}$ and true NH (blue) or IH (red). Joint fit combines the electron neutrino appearance and muon neutrino disappearance channels, and global reactor constraint $\sin^2 \theta_{13} = 0.082 \pm 0.004$. Assuming an equal fraction of beam exposure in neutrino and antineutrino mode after 2019, and a total of 72×10^{20} POT by 2024 (details in table 7.1). From [145].

In Figure 7.8 we present the evolution of the projected significance with time, assuming the extreme CP-violating cases of $\delta_{CP} = \pi/2, 3\pi/2$, and the range $0.4 \leq \sin^2 \theta_{23} \leq 0.6$. For all combinations, we would expect to remain under the 3σ limit.

Figure 7.9 shows that for 30-40% of δ_{CP} values we could establish CP violation at 95% C.L. by 2024. Figure 7.10 presents the end-of-life expectations in more detail. An interesting feature in this graph is the symmetry between the two mass hierarchies. While this could be construed as a consequence of the 50/50 split between neutrino and antineutrino running, turns out that similar levels of confidence can be achieved with other combinations of POT [142]. The run plan cannot be optimized to significantly improve this measurement, and other possibilities at the analysis level including the addition of external data need to be explored.




Figure 7.8: Projected significance of rejection of CP conservation with NOvA as a function of year, assuming true normal (left) or inverted (right) mass hierarchy. The color bands correspond to different true values of δ_{CP} , and a range of $\sin^2 \theta_{23} = 0.4 - 0.6$. Joint fit combines the electron neutrino appearance and muon neutrino disappearance channels, and global reactor constraint $\sin^2 \theta_{13} = 0.082 \pm 0.004$. Assuming an equal fraction of beam exposure in neutrino and antineutrino mode after 2019, and a total of 72×10^{20} POT by 2024 (details in table 7.1). From [145].

7.3.3 Rejection of maximal mixing and octant determination

The determination of the θ_{23} octant and rejection of maximal mixing are simpler results that essentially depend on how far away the true value of $\sin^2 \theta_{23}$ is from 0.5, and therefore related to the precision of the ν_μ disappearance measurement. As mentioned before, the ν_e data serves as the means to break the octant degeneracy. How early this can occur depends on the true mass hierarchy and the true octant: recall the balance between neutrino/antineutrino events for the hierarchy/octant combinations presented in Figure 7.3. The impact of the true value of δ_{CP} is smaller, and almost negligible for rejecting maximal mixing.

Figures 7.11 and 7.12 show the projected significances for rejection of maximal mixing and octant determination, respectively. In this case, the color bands cover the δ_{CP} range $[0, 2\pi)$. Note that these bands would collapse to 0 as $\sin^2 2\theta_{23}$ approaches the physical limit of 1 corresponding to maximal mixing. For the chosen values of $\sin^2 \theta_{23}$, we can expect the rejection of maximal mixing to reach 3σ by 2020 in all cases, followed by the octant determination after that, depending on the ν_e observation.




Figure 7.9: Fraction of δ_{CP} values in the range $[0, 2\pi)$ for which CP conservation could be rejected by NOvA at a 95% C.L. or higher, as a function of year. The color lines correspond to different true values of $\sin^2 \theta_{23}$. Joint fit combines the electron neutrino appearance and muon neutrino disappearance channels, and global reactor constraint $\sin^2 \theta_{13} = 0.082 \pm 0.004$. Assuming an equal fraction of beam exposure in neutrino and antineutrino mode after 2019, and a total of 72×10^{20} POT by 2024 (details in table 7.1). From [145].




Figure 7.10: Fraction of δ_{CP} values in the range $[0, 2\pi)$ for which CP conservation could be rejected by NOvA at a given sensitivity σ by the year 2024. The lines correspond to different true values of $\sin^2 \theta_{23}$ and true NH (blue) or IH (red). Joint fit combines the electron neutrino appearance and muon neutrino disappearance channels, and global reactor constraint $\sin^2 \theta_{13} = 0.082 \pm 0.004$. Assuming an equal fraction of beam exposure in neutrino and antineutrino mode after 2019, and a total of 72×10^{20} POT by 2024 (details in table 7.1). From [145].




Figure 7.11: Projected significance of rejection of maximal mixing with NOvA as a function of year. The color bands correspond to different true values of $\sin^2 \theta_{23}$ and all possible values of δ_{CP} . Joint fit combines the electron neutrino appearance and muon neutrino disappearance channels, and global reactor constraint $\sin^2 \theta_{13} = 0.082 \pm 0.004$. Assuming an equal fraction of beam exposure in neutrino and antineutrino mode after 2019, and a total of 72×10^{20} POT by 2024 (details in table 7.1). From [145].




Figure 7.12: Projected significance of octant determination with NOvA as a function of year. The color bands correspond to different true values of $\sin^2 \theta_{23}$ and all possible values of δ_{CP} . Joint fit combines the electron neutrino appearance and muon neutrino disappearance channels, and global reactor constraint $\sin^2 \theta_{13} = 0.082 \pm 0.004$. Assuming an equal fraction of beam exposure in neutrino and antineutrino mode after 2019, and a total of 72×10^{20} POT by 2024 (details in table 7.1). From [145].

7.4 Possible improvements

7.4.1 Changes in beam configuration

The exposure assumptions from Table 7.1 include an expected ramp up of the NuMI beam power from 700 to 1000 kW in the next five years, and a total of 40 weeks of beam in a year compared to the 34 weeks originally proposed. These changes already represent an increase in exposure of 50% compared to the NOvA proposal [34].

There are preliminary studies [147] that indicate that a change in the geometry of the beam target elements can lead to an additional 17% effective exposure. Higher benefits could be observed in the low energy NuMI beam configuration [148, 149]. Since NOvA is an off-axis experiment, changes from the medium- to the low-energy beam configurations are not as evident as they would be on-axis [148, 149]. Nonetheless, some of these options might be worth exploring if the target modifications prove feasible, or in case the experiment runs longer than the planned 2024 shutdown.

7.4.2 Analysis improvements

Some analysis gains can be achieved by controlling the systematic uncertainties, especially as the statistical errors get smaller over time. Some of the dominant systematics affecting the results in Chapter 6 involve the neutrino energy, in particular calibration, and neutrino interaction models. Other options to improve the analysis involve reducing neutral current and wrong sign backgrounds, and increasing the number of selected events. In practice, they can involve studies with test beam data [150], adding new samples in the far detector, or using constraints from cross section measurements with the near detector.

7.4.3 Joint analysis with T2K

Along with improvements using the NOvA detectors, another way to improve our knowledge of neutrino oscillations is combining external data to constrain a measurement. For example, there have been two long-baseline experiments concurrent with NOvA: MINOS and T2K. Together, they probe various baselines, neutrino energies, and detector technologies. MINOS had a similar two-

detector setup and shared the NuMI beam with NOvA but at a different angle; the analysis of the combined data would result in a better understanding of the flux and additional information for a range of neutrino energies. Meanwhile, the T2K experiment has a radically different design and a higher sensitivity to reject CP conservation while NOvA is more sensitive to the neutrino mass hierarchy. Therefore, a joint fit with these complementary datasets should impose stronger limits on oscillation parameters. In both examples, the combination can be challenging as experiments use different simulation and analysis techniques, but it will lead to measurements of neutrino interactions and oscillation results with higher sensitivities.




Figure 7.13: Probabilities of electron neutrino and antineutrino appearance in the NOvA and T2K experiments. The NOvA baseline and average neutrino energies increase the separation between normal and inverted hierarchy, while the higher overlap in T2K could potentially allow a measurement of CP non-conservation.

7.5 Conclusion

The joint analysis of electron neutrino appearance and muon neutrino disappearance data in the NOvA experiment can set constraints on the three-flavor oscillation parameters $|\Delta m_{32}^2|$, θ_{23} , δ_{CP} and the neutrino mass hierarchy. In this thesis, I have presented the results of the analysis of a dataset corresponding to an exposure of 8.85×10^{20} protons on target, with the NuMI beam running in neutrino-enhanced mode.

The $\nu_\mu \rightarrow \nu_\mu$ channel is sensitive to the values $\sin^2 \theta_{23}$ and Δm_{32}^2 , and as the size of the datasets increase, the precision on both measurements will improve, potentially allowing NOvA to have the most precise measurement of these parameters in the short term. The current results of the joint analysis in the normal mass hierarchy have the best fit values $\Delta m_{32}^2 = 2.44 \times 10^{-3} eV^2/c^4$, $\sin^2 \theta_{23} = 0.56$, with the allowed regions $\Delta m_{32}^2 \in [2.37, 2.52] \times 10^{-3} eV^2/c^4$, $\sin^2 \theta_{23} \in [0.43, 0.51] \cup [0.52, 0.60]$ at the 68.3% C.L.

The sensitivity to δ_{CP} and the neutrino mass hierarchy comes from the $\nu_\mu \rightarrow \nu_e$ channel. The best fit value obtained in this analysis is $\delta_{CP} = 1.21\pi$, with the allowed range $\delta_{CP} \in [0, 0.12\pi] \cup [0.91\pi, 2\pi]$ in the normal mass hierarchy. In the inverted hierarchy, values around $\delta_{CP} = 3\pi/2$ are disfavored at higher than 3σ independently of the value of $\sin^2 \theta_{32}$. The inverted mass hierarchy is disfavored at the 95% confidence level for all choices of the other oscillation parameters.

Even though NOvA was designed to observe electron neutrinos in the FD, a combination of factors including the low neutrino cross sections, the small probabilities of appearance and the lack of knowledge of other oscillation parameters, create difficulties for discovery. While it is unlikely that NOvA will confirm CP violation in the lepton sector, there is a range of values of δ_{CP} that are favorable for the determination of neutrino the mass hierarchy at a high confidence level.

Finally, there are several ways to improve the measurements in NOvA, including: i) optimization of the run plan based on increased knowledge of the true oscillation parameters; ii) improvements in the NuMI beam design and performance; iii) reduction of systematic uncertainties and better energy estimation and selection via complementary measurements with the NOvA detectors; and iv) novel combinations of NOvA and other experiments' data.

BIBLIOGRAPHY

- [1] T. J. Trefft, “Rutherford on the alpha-beta-gamma classification of radioactive rays,” *Isis* **67**, 61–75 (1976).
- [2] J. Chadwick, “The intensity distribution in the magnetic spectrum of beta particles from radium (b+ c),” *Verh. Phys. Gesell.* **16**, 383–391 (1914).
- [3] P. Hernandez, “Neutrino Physics,” in *Proceedings, 8th CERN Latin-American School of High-Energy Physics (CLASHEP2015): Ibarra, Ecuador, March 05-17, 2015* (2016) pp. 85–142.
- [4] W. Pauli, “On the Earlier and more recent history of the neutrino,” *Camb. Monogr. Part. Phys. Nucl. Phys. Cosmol.* **14**, 1–22 (2000).
- [5] J. Chadwick, “Possible Existence of a Neutron,” *Nature* **129**, 312 (1932).
- [6] E. Fermi, “Versuch einer theorie der beta-strahlen,” *Zeitschrift für Physik* **88** (1934).
- [7] H. Bethe and R. Peierls, “The ‘neutrino’,” *Nature* **133**, 532 (1934).
- [8] B. Pontecorvo, “Inverse beta process,” *Solar Neutrinos: An Overview*, *Camb. Monogr. Part. Phys. Nucl. Phys. Cosmol.* **1**, 25–31 (1991), [97(1946)].
- [9] F. Reines and C. L. Cowan, “Detection of the free neutrino,” *Phys. Rev.* **92**, 830–831 (1953).
- [10] C. L. Cowan, F. Reines, F. B. Harrison, H. W. Kruse, and A. D. McGuire, “Detection of the free neutrino: A Confirmation,” *Science* **124**, 103–104 (1956).
- [11] F. Reines, “The neutrino: From poltergeist to particle,” *Rev. Mod. Phys.* **68**, 317–327 (1996).
- [12] H. Yukawa, “On the Interaction of Elementary Particles I,” *Proc. Phys. Math. Soc. Jap.* **17**, 48–57 (1935), [Prog. Theor. Phys. Suppl.1,1(1935)].
- [13] B. Pontecorvo, “Electron and Muon Neutrinos,” *Sov. Phys. JETP* **10**, 1236–1240 (1960), [*Zh. Eksp. Teor. Fiz.*37,1751(1959)].
- [14] G. Danby, J.-M. Gaillard, K. Goulianos, L. M. Lederman, N. Mistry, M. Schwartz, and J. Steinberger, “Observation of high-energy neutrino reactions and the existence of two kinds of neutrinos,” *Phys. Rev. Lett.* **9**, 36–44 (1962).
- [15] M. L. Perl *et al.*, “Evidence for Anomalous Lepton Production in $e^+ - e^-$ Annihilation,” *Phys. Rev. Lett.* **35**, 1489–1492 (1975), [193(1975)].

- [16] K. Kodama *et al.* (DONUT), “Observation of tau neutrino interactions,” *Phys. Lett.* **B504**, 218–224 (2001).
- [17] U. Sarkar, *Particle and astroparticle physics*, Series in high energy physics, cosmology, and gravitation (Taylor and Francis, New York, USA, 2008).
- [18] R. N. Mohapatra and P. B. Pal, “Massive neutrinos in physics and astrophysics. Second edition,” *World Sci. Lect. Notes Phys.* **60**, 1–397 (1998), [World Sci. Lect. Notes Phys.72,1(2004)].
- [19] C. Patrignani *et al.* (Particle Data Group), “Review of Particle Physics,” *Chin. Phys.* **C40**, 100001 (2016).
- [20] M. Tanabashi *et al.* (Particle Data Group), “Review of Particle Physics,” *Phys. Rev.* **D98**, 030001 (2018).
- [21] Z. Maki, M. Nakagawa, and S. Sakata, “Remarks on the unified model of elementary particles,” *Prog. Theor. Phys.* **28**, 870–880 (1962), [,34(1962)].
- [22] B. Pontecorvo, “Neutrino Experiments and the Problem of Conservation of Leptonic Charge,” *Sov. Phys. JETP* **26**, 984–988 (1968), [Zh. Eksp. Teor. Fiz.53,1717(1967)].
- [23] B. Kayser, “Neutrino physics,” *Proceedings, 32nd SLAC Summer Institute on Particle Physics: Cosmic Connections (SSI 2004): Menlo Park, California, August 2-13, 2004*, eConf **C040802**, L004 (2004).
- [24] L. Wolfenstein, “Neutrino Oscillations in Matter,” *Phys. Rev.* **D17**, 2369–2374 (1978), [,294(1977)].
- [25] S. P. Mikheyev and A. Yu. Smirnov, “Resonance Amplification of Oscillations in Matter and Spectroscopy of Solar Neutrinos,” *Sov. J. Nucl. Phys.* **42**, 913–917 (1985), [,305(1986)].
- [26] V. D. Barger, K. Whisnant, S. Pakvasa, and R. J. N. Phillips, “Matter Effects on Three-Neutrino Oscillations,” *Phys. Rev.* **D22**, 2718 (1980), [,300(1980)].
- [27] H. Nunokawa, S. J. Parke, and R. Zukanovich Funchal, “Another possible way to determine the neutrino mass hierarchy,” *Phys. Rev.* **D72**, 013009 (2005).
- [28] C. Giunti, C. W. Kim, and M. Monteno, “Atmospheric neutrino oscillations with three neutrinos and a mass hierarchy,” *Nucl. Phys.* **B521**, 3–36 (1998).
- [29] A. Cervera, A. Donini, M. B. Gavela, J. J. Gomez Cadenas, P. Hernandez, O. Mena, and S. Rigolin, “Golden measurements at a neutrino factory,” *Nucl. Phys.* **B579**, 17–55 (2000), [Erratum: Nucl. Phys.B593,731(2001)].
- [30] F. Capozzi, E. Di Valentino, E. Lisi, A. Marrone, A. Melchiorri, and A. Palazzo, “Global constraints on absolute neutrino masses and their ordering,” *Phys. Rev.* **D95**, 096014 (2017).

- [31] D. S. Ayres *et al.* (NOvA), *The NOvA Technical Design Report*, Tech. Rep. FERMILAB-DESIGN-2007-01 (Fermilab, 2007).
- [32] P. Adamson *et al.*, “The NuMI Neutrino Beam,” *Nucl. Instrum. Meth.* **A806**, 279–306 (2016).
- [33] P. Adamson *et al.* (NOvA), “NOvA Official Plot Database,” NOvA Internal Documents (2014–2018).
- [34] D. S. Ayres *et al.* (NOvA), “NOvA: Proposal to Build a 30 Kiloton Off-Axis Detector to Study $\nu_\mu \rightarrow \nu_e$ Oscillations in the NuMI Beamline,” (2004).
- [35] E. Ables *et al.* (MINOS), *P-875: A Long baseline neutrino oscillation experiment at Fermilab*, Tech. Rep. FERMILAB-PROPOSAL-0875, NUMI-L-79 (Fermilab, 1995).
- [36] “Fermilab Creative Services,” (2018).
- [37] J. Evans (MINOS), “The MINOS experiment: results and prospects,” *Adv. High Energy Phys.* **2013**, 182537 (2013).
- [38] Accelerator Division - Operations Department, *Concepts Rookie Book*, Tech. Rep. (Fermilab, 2013).
- [39] K. Abe *et al.* (T2K), “Measurements of neutrino oscillation in appearance and disappearance channels by the T2K experiment with 6.6×10^{20} protons on target,” *Phys. Rev.* **D91**, 072010 (2015).
- [40] R. P. Gandrajula and M. Groh (NOvA), “Systematic Uncertainties and Cross-Checks for the NOvA Joint $\nu_\mu + \nu_e$ Analysis,” in *28th International Conference on Neutrino Physics and Astrophysics (Neutrino 2018) Heidelberg, Germany, June 4-9, 2018* (2018).
- [41] A. Norman *et al.*, “Performance of the NOvA Data Acquisition and Trigger Systems for the full 14 kT Far Detector,” *Proceedings, 21st International Conference on Computing in High Energy and Nuclear Physics (CHEP 2015): Okinawa, Japan, April 13-17, 2015*, *J. Phys. Conf. Ser.* **664**, 082041 (2015).
- [42] A. Norman, R. Kwarcianny, G. Deuerling, and N. Wilcer (NOvA), “The NOvA Timing System: A System for Synchronizing a Long Baseline Neutrino Experiment,” *Proceedings, 19th International Conference on Computing in High Energy and Nuclear Physics (CHEP 2012): New York, USA, May 21-25, 2012*, *J. Phys. Conf. Ser.* **396**, 012034 (2012).
- [43] A. Norman, E. Niner, and A. Habig, “Timing in the NOvA detectors with atomic clock based time transfers between Fermilab, the Soudan mine and the NOvA Far detector,” *Proceedings, 21st International Conference on Computing in High Energy and Nuclear Physics (CHEP 2015): Okinawa, Japan, April 13-17, 2015*, *J. Phys. Conf. Ser.* **664**, 082040 (2015).
- [44] R. Hatcher, “Simulation tools in the neutrino experiments,” NOvA Internal Document 14441 (2013).

- [45] S. Agostinelli *et al.* (GEANT4), “GEANT4: A Simulation toolkit,” *Nucl. Instrum. Meth.* **A506**, 250–303 (2003).
- [46] L. Aliaga *et al.* (MINERvA), “Neutrino Flux Predictions for the NuMI Beam,” *Phys. Rev.* **D94**, 092005 (2016), [Addendum: *Phys. Rev.*D95,no.3,039903(2017)].
- [47] M. A. Acero *et al.* (NOvA), “New constraints on oscillation parameters from ν_e appearance and ν_μ disappearance in the NOvA experiment,” *Phys. Rev.* **D98**, 032012 (2018).
- [48] C. Andreopoulos *et al.*, “The GENIE Neutrino Monte Carlo Generator,” *Nucl. Instrum. Meth.* **A614**, 87–104 (2010).
- [49] C. Andreopoulos, C. Barry, S. Dytman, H. Gallagher, T. Golan, R. Hatcher, G. Perdue, and J. Yarba, “The GENIE Neutrino Monte Carlo Generator: Physics and User Manual,” (2015).
- [50] T. Katori, “Meson Exchange Current (MEC) Models in Neutrino Interaction Generators,” *Proceedings, 8th International Workshop on Neutrino-Nucleus Interactions in the Few GeV Region (NuInt 12): Rio de Janeiro, Brazil, October 22-27, 2012*, *AIP Conf. Proc.* **1663**, 030001 (2015).
- [51] J. Wolcott (NOvA), “Recent Cross Section Work From NOvA,” in *18th International Workshop on Neutrino Factories and Future Neutrino Facilities Search (NuFact16) Quy Nhon, Vietnam, August 21-27, 2016* (2016).
- [52] J. Nieves, J. E. Amaro, and M. Valverde, “Inclusive quasi-elastic neutrino reactions,” *Phys. Rev.* **C70**, 055503 (2004), [Erratum: *Phys. Rev.*C72,019902(2005)].
- [53] R. Gran, *Model Uncertainties for Valencia RPA Effect for MINERvA*, Tech. Rep. FERMILAB-FN-1030-ND (Fermilab, 2017).
- [54] P. Rodrigues, C. Wilkinson, and K. McFarland, “Constraining the GENIE model of neutrino-induced single pion production using reanalyzed bubble chamber data,” *Eur. Phys. J.* **C76**, 474 (2016).
- [55] A. Habig, A. Norman, and C. Group, “Recent Evolution of the Offline Computing Model of the NOvA Experiment,” *Proceedings, 21st International Conference on Computing in High Energy and Nuclear Physics (CHEP 2015): Okinawa, Japan, April 13-17, 2015*, *J. Phys. Conf. Ser.* **664**, 032011 (2015).
- [56] B. Rebel, M. C. Sanchez, and S. Wolbers, “Working Group Report: Computing for the Intensity Frontier,” in *Proceedings, 2013 Community Summer Study on the Future of U.S. Particle Physics: Snowmass on the Mississippi (CSS2013): Minneapolis, MN, USA, July 29-August 6, 2013* (2013).
- [57] R. Brun and F. Rademakers, “ROOT: An object oriented data analysis framework,” *New computing techniques in physics research V. Proceedings, 5th International Workshop, AI-*

- HENP '96, Lausanne, Switzerland, September 2-6, 1996*, *Nucl. Instrum. Meth.* **A389**, 81–86 (1997).
- [58] C. Green, J. Kowalkowski, M. Paterno, M. Fischler, L. Garren, and Q. Lu, “The Art Framework,” *Proceedings, 19th International Conference on Computing in High Energy and Nuclear Physics (CHEP 2012): New York, USA, May 21-25, 2012*, *J. Phys. Conf. Ser.* **396**, 022020 (2012).
- [59] R. A. Illingworth, “A Data Handling System for Modern and Future Fermilab Experiments,” *Proceedings, 20th International Conference on Computing in High Energy and Nuclear Physics (CHEP 2013): Amsterdam, The Netherlands, October 14-18, 2013*, *J. Phys. Conf. Ser.* **513**, 032045 (2014).
- [60] C. Group, S. Fuess, O. Gutsche, M. Kirby, R. Kutschke, A. Lyon, A. Norman, G. Perdue, and E. Sexton-Kennedy, “Fermilab Computing at the Intensity Frontier,” *Proceedings, 21st International Conference on Computing in High Energy and Nuclear Physics (CHEP 2015): Okinawa, Japan, April 13-17, 2015*, *J. Phys. Conf. Ser.* **664**, 032012 (2015).
- [61] A. Himmel, “New Oscillation Measurements from NOvA (CERN Seminar),” NOvA Internal Document 26151 (2018).
- [62] M. Baird, J. Bian, M. Messier, E. Niner, D. Rocco, and K. Sachdev, “Event Reconstruction Techniques in NOvA,” *Proceedings, 21st International Conference on Computing in High Energy and Nuclear Physics (CHEP 2015): Okinawa, Japan, April 13-17, 2015*, *J. Phys. Conf. Ser.* **664**, 072035 (2015).
- [63] M. Baird, “Reconstruction tutorial,” NOvA Internal Document 20171 (2017).
- [64] M. Baird, “Technote for slicer4d,” NOvA Internal Document 9195 (2013).
- [65] M. Ester, H.-P. Kriegel, J. Sander, X. Xu, *et al.*, “A density-based algorithm for discovering clusters in large spatial databases with noise.” in *Kdd 96*, 34 (1996) pp. 226–231.
- [66] B. Rebel, “A Window Tracking Algorithm for Cosmic Ray Muons,” NOvA Internal Document 15977 (2016).
- [67] N. Raddatz, “Kalman Track Technical Note,” NOvA Internal Document 13545 (2015).
- [68] R. E. Kalman, “A new approach to linear filtering and prediction problems,” *Journal of basic Engineering* **82**, 35–45 (1960).
- [69] P. V. Hough, “Method and means for recognizing complex patterns,” (1962), US Patent 3,069,654.
- [70] M. Ohlsson, “Extensions and explorations of the elastic arms algorithm,” *Comput. Phys. Commun.* **77**, 19–32 (1993).

- [71] M. Baird, “Vertex reconstruction based on elastic arms,” NOvA Internal Document 7530 (2012).
- [72] E. Niner, “Vertex Clustering with Possibilistic Fuzzy-K Means Algorithm,” NOvA Internal Document 7648 (2012).
- [73] R. Krishnapuram and J. M. Keller, “A possibilistic approach to clustering,” *IEEE transactions on fuzzy systems* **1**, 98–110 (1993).
- [74] M.-S. Yang and K.-L. Wu, “Unsupervised possibilistic clustering,” *Pattern Recognition* **39**, 5–21 (2006).
- [75] A. Aurisano, A. Radovic, D. Rocco, A. Himmel, M. D. Messier, E. Niner, G. Pawloski, F. Psihas, A. Sousa, and P. Vahle, “A Convolutional Neural Network Neutrino Event Classifier,” *JINST* **11**, P09001 (2016).
- [76] P. Adamson *et al.* (NOvA), “Constraints on Oscillation Parameters from ν_e Appearance and ν_μ Disappearance in NOvA,” *Phys. Rev. Lett.* **118**, 231801 (2017).
- [77] C. Szegedy, W. Liu, Y. Jia, P. Sermanet, S. Reed, D. Anguelov, D. Erhan, V. Vanhoucke, and A. Rabinovich, “Going deeper with convolutions,” in *Proceedings of the IEEE conference on computer vision and pattern recognition* (2015) pp. 1–9.
- [78] D. E. Rumelhart, G. E. Hinton, and R. J. Williams, “Learning representations by back-propagating errors,” *Nature* **323**, 533 (1986).
- [79] Y. Jia, E. Shelhamer, J. Donahue, S. Karayev, J. Long, R. Girshick, S. Guadarrama, and T. Darrell, “Caffe: Convolutional architecture for fast feature embedding,” in *Proceedings of the 22nd ACM international conference on Multimedia* (ACM, 2014) pp. 675–678.
- [80] F. Rosenblatt, *Principles of neurodynamics. perceptrons and the theory of brain mechanisms*, Tech. Rep. (Cornell Aeronautical Lab Inc. Buffalo NY, 1961).
- [81] R. Reed and R. J. Marks II, *Neural smithing: supervised learning in feedforward artificial neural networks* (MIT Press, 1999).
- [82] B. Behera, G. Davies, and F. Psihas (NOvA), “Event Reconstruction in the NOvA Experiment,” in *Proceedings, Meeting of the APS Division of Particles and Fields (DPF 2017): Fermilab, Batavia, Illinois, USA, July 31 - August 4, 2017* (2017).
- [83] T. Cover and P. Hart, “Nearest neighbor pattern classification,” *IEEE transactions on information theory* **13**, 21–27 (1967).
- [84] N. Raddatz, “ReMId Technical Note,” NOvA Internal Document 11206 (2014).
- [85] N. J. Raddatz, *Measurement of Muon Neutrino Disappearance with Non-Fiducial Interactions in the NOvA Experiment*, Ph.D. thesis, Minnesota U. (2016).

- [86] F. Psihas, *Measurement of Long Baseline Neutrino Oscillations and Improvements from Deep Learning*, [Ph.D. thesis](#), Indiana U. (2018).
- [87] C. G. Broyden, “The convergence of a class of double-rank minimization algorithms 1. general considerations,” *IMA Journal of Applied Mathematics* **6**, 76–90 (1970).
- [88] R. Fletcher, “A new approach to variable metric algorithms,” *The computer journal* **13**, 317–322 (1970).
- [89] D. Goldfarb, “A family of variable-metric methods derived by variational means,” *Mathematics of computation* **24**, 23–26 (1970).
- [90] D. F. Shanno, “Conditioning of quasi-newton methods for function minimization,” *Mathematics of computation* **24**, 647–656 (1970).
- [91] K. Bays, “NOvA cosmic rejection package and algorithms technical note,” NOvA Internal Document 11205 (2017).
- [92] D. Pershey, *A Measurement of ν_e Appearance and ν_μ Disappearance Neutrino Oscillations with the NOvA Experiment*, [Ph.D. thesis](#), California Institute of Technology (2018).
- [93] C. Backhouse *et al.*, “The Attenuation and Threshold Calibration of the NOvA Detectors,” NOvA Internal Document 13579 (2017).
- [94] R. Nichol, B. Zamorano, *et al.*, “Executive summary of calibration for the third analysis,” NOvA Internal Document 23558 (2017).
- [95] P. Singh *et al.*, “Calibration of the NOvA Detectors,” NOvA Internal Document 15786 (2016).
- [96] B. Behera *et al.*, “Drift calibration,” NOvA Internal Document 16984 (2017).
- [97] L. Vinton, *Measurement of Muon Neutrino Disappearance with the NOvA Experiment*, [Ph.D. thesis](#), Sussex U. (2018).
- [98] E. Catano-Mur *et al.*, “Systematic Uncertainties for the ν_e Appearance and ν_μ Disappearance 2017 Analysis,” NOvA Internal Document 24029 (2017).
- [99] L. Aliaga, G. Brunetti, L. Cremonesi, K. Maan, and A. Radovich, “NuMI flux systematic uncertainties for the NOvA third analyses,” NOvA Internal Document 17608 (2017).
- [100] N. Nayak, “Hadron Production Systematics for the NOvA Oscillation Analysis,” NOvA Internal Document 22352 (2017).
- [101] R. Gran, J. Nieves, F. Sanchez, and M. J. Vicente Vacas, “Neutrino-nucleus quasi-elastic and 2p2h interactions up to 10 GeV,” [Phys. Rev. D](#) **88**, 113007 (2013).

- [102] M. Martini, M. Ericson, G. Chanfray, and J. Marteau, “A Unified approach for nucleon knock-out, coherent and incoherent pion production in neutrino interactions with nuclei,” *Phys. Rev.* **C80**, 065501 (2009).
- [103] G. Megias, J. Amaro, M. Barbaro, J. Caballero, T. Donnelly, and I. Ruiz Simo, “Charged-current neutrino-nucleus reactions within the superscaling meson-exchange current approach,” *Phys. Rev.* **D94**, 093004 (2016).
- [104] T. Katori and M. Martini, “Neutrinonucleus cross sections for oscillation experiments,” *J. Phys.* **G45**, 013001 (2018).
- [105] A. Bodek, S. Avvakumov, R. Bradford, and H. S. Budd, “Vector and Axial Nucleon Form Factors:A Duality Constrained Parameterization,” *Eur. Phys. J.* **C53**, 349–354 (2008).
- [106] A. S. Meyer, M. Betancourt, R. Gran, and R. J. Hill, “Deuterium target data for precision neutrino-nucleus cross sections,” *Phys. Rev.* **D93**, 113015 (2016).
- [107] M. Day and K. S. McFarland, “Differences in Quasi-Elastic Cross-Sections of Muon and Electron Neutrinos,” *Phys. Rev.* **D86**, 053003 (2012).
- [108] K. Abe *et al.* (T2K), “Measurement of neutrino and antineutrino oscillations by the T2K experiment including a new additional sample of ν_e interactions at the far detector,” *Phys. Rev.* **D96**, 092006 (2017).
- [109] P. Adamson *et al.* (NOvA), “Measurement of the neutrino mixing angle θ_{23} in NOvA,” *Phys. Rev. Lett.* **118**, 151802 (2017).
- [110] A. Aurisano, “Tech note: 2017 light model,” NOvA Internal Document 23228 (2017).
- [111] J. A. Sepulveda-Quiroz, *Measurement of the Kaon Production Normalization in the NuMI Target Using Uncontained Charged-Current Muon Neutrino Interactions in the NOvA Far Detector*, *Ph.D. thesis*, Iowa State U. (2018).
- [112] J. Hartnell and A. Radovic, “Summary of the 2017 ν_μ disappearance analysis,” NOvA Internal Document 22562 (2017).
- [113] M. Baird and L. Vinton, “Extrapolation technote for the ν_μ third analysis,” NOvA Internal Document 23390 (2017).
- [114] C. Backhouse and A. Himmel, “Overview of the 2017 NOvA ν_e Appearance Analysis,” NOvA Internal Document 22475 (2017).
- [115] P. Adamson *et al.* (NOvA), “First measurement of electron neutrino appearance in NOvA,” *Phys. Rev. Lett.* **116**, 151806 (2016).
- [116] K. Sachdev, “Cosmic Rejection and ν_e Selection Cuts in the Far Detector,” NOvA Internal Document 15387 (2016).

- [117] D. Pershey, L. Kolupaeva, and N. Nayak, “Event selection in the 2017 ν_e analysis,” NOvA Internal Document 22445 (2017).
- [118] I. Anghel, Z. Djurcic, and M. Sanchez, “Development of Beam ν_e Estimation Method in NOvA Oscillation Analysis,” NOvA Internal Document 15442 (2016).
- [119] K. K. Maan *et al.*, “K+ normalization and the ν_e flux using ν_μ CC ND Data,” NOvA Internal Document 15379 (2016).
- [120] E. Catano-Mur, “BEN (Beam Electron Neutrino) Decomposition Technote 2017,” NOvA Internal Document 23174 (2017).
- [121] E. Catano-Mur, “BEN (Beam Electron Neutrino) Decomposition Technote,” NOvA Internal Document 15392 (2016).
- [122] D. Pershey, “MEFinder Tech Note,” NOvA Internal Document 14789 (2015).
- [123] C. Backhouse and N. Nayak, “ND data/MC comparisons for the 2017 ν_e analysis,” NOvA Internal Document 23237 (2017).
- [124] D. Pershey, “Extrapolation for the ν_e Appearance Analysis - Technote,” NOvA Internal Document 23108 (2017).
- [125] T. Nosek, “ND Subcomponents / Kinematics Systematics for 2017 ν_e Appearance Analysis,” NOvA Internal Document 22478 (2017).
- [126] E. Catano-Mur and D. Pershey, “Technical Note for Fitting, Feldman-Cousins, and Sensitivities for the 2017 Joint Analysis,” NOvA Internal Document 23126 (2017).
- [127] S. S. Wilks, “The large-sample distribution of the likelihood ratio for testing composite hypotheses,” *Ann. Math. Statist.* **9**, 60–62 (1938).
- [128] G. J. Feldman and R. D. Cousins, “A Unified approach to the classical statistical analysis of small signals,” *Phys. Rev.* **D57**, 3873–3889 (1998).
- [129] C. Bassin, G. Laske, and G. Masters, “The current limits of resolution for surface wave tomography in north america,” *Eos* **81** (2000).
- [130] A. Radovic, “Oscillation Parameters for First NOvA Analyses,” NOvA Internal Document 13640 (2015).
- [131] O. Behnke, K. Kröninger, G. Schott, and T. Schörner-Sadenius, eds., *Data analysis in high energy physics* (Wiley-VCH, Weinheim, Germany, 2013).
- [132] J. Neyman, “Outline of a theory of statistical estimation based on the classical theory of probability,” *Philosophical Transactions of the Royal Society of London. Series A, Mathematical and Physical Sciences* **236**, 333–380 (1937).

- [133] G. Cowan, K. Cranmer, E. Gross, and O. Vitells, “Asymptotic formulae for likelihood-based tests of new physics,” *Eur. Phys. J.* **C71**, 1554 (2011), [Erratum: *Eur. Phys. J.*C73,2501(2013)].
- [134] E. Catano-Mur, “Joint-fit systematics tables for oscillation parameters,” NOvA Internal Document 26589 (2018).
- [135] P. Adamson *et al.* (MINOS), “Combined analysis of ν_μ disappearance and $\nu_\mu \rightarrow \nu_e$ appearance in MINOS using accelerator and atmospheric neutrinos,” *Phys. Rev. Lett.* **112**, 191801 (2014).
- [136] M. G. Aartsen *et al.* (IceCube), “Measurement of Atmospheric Neutrino Oscillations at 656 GeV with IceCube DeepCore,” *Phys. Rev. Lett.* **120**, 071801 (2018).
- [137] K. Abe *et al.* (Super-Kamiokande), “Atmospheric neutrino oscillation analysis with external constraints in Super-Kamiokande I-IV,” *Phys. Rev.* **D97**, 072001 (2018).
- [138] E. Catano-Mur, “ ν_e FHC and RHC event counts,” NOvA Internal Document 17935 (2017).
- [139] E. Catano-Mur and N. Nayak, “2018 Bi-Event Plots – Blessing Package,” NOvA Internal Document 29663 (2018).
- [140] D. Drakoulakos *et al.* (MINERvA), “Proposal to perform a high-statistics neutrino scattering experiment using a fine-grained detector in the NuMI beam,” (2004).
- [141] P. Vahle *et al.*, “NOvA update for the PAC,” NOvA Internal Document 30924 (2018).
- [142] E. Catano-Mur, “Future sensitivities and FHC/RHC POT optimization,” NOvA Internal Document 16134 (2016).
- [143] E. Catano-Mur and M. Sanchez, “NOvA Sensitivity Projections Summary,” NOvA Internal Document 16111 (2016).
- [144] R. Acciarri *et al.* (DUNE), *Long-Baseline Neutrino Facility (LBNF) and Deep Underground Neutrino Experiment (DUNE)*, Tech. Rep. FERMILAB-DESIGN-2016-01 (2016).
- [145] E. Catano-Mur, “Projected significances for the joint $\nu_e + \nu_\mu$ analysis,” NOvA Internal Document 30027 (2018).
- [146] K. Abe *et al.*, “Letter of Intent: The Hyper-Kamiokande Experiment — Detector Design and Physics Potential —,” (2011).
- [147] T. Jyoti, “Study of a new target design with an additional horn for numi beam,” arXiv preprint arXiv:1710.03699 (2017).
- [148] D. Kalra, “Increased Neutrino yield with the New NOvA target design: Simulation Study,” NOvA Internal Document 22328 (2017).

- [149] G. Brunetti, J. Cooper, D. Kalra, and J. Tripathi, “Off-Axis Beam Studies for NOvA,” NOvA Internal Document 16233 (2016).
- [150] A. Sousa, R. Nichol, K. Lang, and J. Nelson, “Test beam task force report,” NOvA Internal Document 15750 (2016).

APPENDIX. OSCILLATION PARAMETERS

Unless explicitly stated, we assume that the following parameters are constant and these values presented are used everywhere:

$$L = 810 \text{ k}, \tag{A.1}$$

$$\rho = 2.84 \text{ g/cm}^3, \tag{A.2}$$

$$\Delta m_{21}^2 = 7.53 \times 10^{-5} \text{ eV}^2/c^4, \tag{A.3}$$

$$\sin^2 2\theta_{12} = 0.851. \tag{A.4}$$

The solar oscillation parameters are the 2017 PDG values [19]; ρ is the density of the earth estimated with the CRUST 2.0 model [129], using the average depth underground between the two detectors (9.38 km) [130].

We used the central value from reactor experiments [19] whenever θ_{13} is constant.

$$\sin^2 2\theta_{13} = 0.082 \pm 0.004, \tag{A.5}$$

A.1 2016 best fit results

The 2016 NOvA best fit results from [76]:

$$\delta_{CP} = 1.48\pi, \tag{A.6}$$

$$\sin^2 \theta_{23} = 0.404, \tag{A.7}$$

$$\Delta m_{32}^2 = 2.7 \times 10^{-3} \text{ eV}^2/c^4. \tag{A.8}$$

A.2 2017 best fit results

The 2017 NOvA best fit results, from this thesis and published in [47]:

$$\delta_{CP} = 1.21\pi, \tag{A.9}$$

$$\sin^2 \theta_{23} = 0.56, \tag{A.10}$$

$$\Delta m_{32}^2 = 2.44 \times 10^{-3} \text{eV}^2/c^4. \tag{A.11}$$