Measuring θ_{13} via Muon Neutrino to Electron Neutrino Oscillations in the MINOS Experiment

Ruth Toner
Pembroke College, University of Cambridge

This dissertation is submitted to the University of Cambridge for the degree of Doctor of Philosophy

Abstract

One of the primary goals in neutrino physics at the present moment is to make a measurement of the neutrino oscillation parameter θ_{13}. This parameter, in addition to being unknown, could potentially allow for the introduction of CP violation into the lepton sector. The MINOS long-baseline neutrino oscillation experiment has the ability to make a measurement of this parameter, by looking for the oscillation of muon neutrinos to electron neutrinos between a Near and Far Detector over a distance of 735 km. This thesis discusses the development of an analysis framework to search for this oscillation mode. Two major improvements to pre-existing analysis techniques have been implemented by the author. First, a novel particle ID technique based on strip topology, known as the Library Event Matching (LEM) method, is optimized for use in MINOS. Second, a multiple bin likelihood method is developed to fit the data. These two improvements, when combined, increase MINOS' sensitivity to $\sin^2(2\theta_{13})$ by 27% over previous analyses. This thesis sees a small excess over background in the Far Detector. A Frequentist interpretation of the data rules out $\theta_{13} = 0$ at 91%. A Bayesian interpretation of the data is also presented, placing the most stringent upper boundary on the oscillation parameter to date, at $\sin^2(2\theta_{13}) < 0.09(0.015)$ for the Normal (Inverted) Hierarchy and $\delta_{CP} = 0$.

Declaration

This dissertation is the result of my own work, except where explicit reference is made to the work of others, and has not been submitted for another qualification to this or any other university. This dissertation does not exceed the word limit for the respective Degree Committee.

Ruth Toner

Acknowledgements

This thesis would not have been possible without the help and support of many people. Thank you first to my advisor, Mark Thomson, for inviting me to Cambridge in the first place, and for four years of excellent insight and guidance. Despite your busy schedule, you taught me both how to do a physics analysis and how to interpret data correctly. We are still waiting for your final statistics lecture.

Thank you also to the other students and postdocs I worked with on the Library Event Matching particle ID. Pedro Ochoa introduced me to the MINOS ν_e analysis and LEM and took the time to teach me throughout my first two years. Ryan Patterson acted as a second advisor during much of the work in this thesis. He answered literally hundreds of my emails, and was there day and night to discuss new ideas or analysis problems. Mhair Orchanian was an excellent fellow traveler during this analysis, providing good ideas and humour even during the darkest moments.

Thank you to all the members of the MINOS ν_e working group over the years: Lisa Whitehead, Greg Pawloski, Jeff Nelson, João Coelho, Patricia Vahle, Steve Cavanaugh, and many others. You have been a wonderful group of people to work with, and I am proud of the physics we have done together. A special thank you to Nick West for his tireless help in battling the UK Grid. Thank you also to the MINOS spokespeople past and present, Rob Plunkett, Stan Wojcicki, and Jenny Thomas, for supporting my work and making me feel welcome on the collaboration.

Thank you to my fellow Cambridge Minoans. Andy Blake taught me the basics of MINOS from my first week in the UK, and answered many of my questions over the years. Jessica Mitchell, you have been an excellent office-mate, traveling companion, and above all friend.

Thank you a second time to Mark, Ryan, Lisa, and Andy, for reading drafts of this thesis and providing many helpful suggestions.

Thank you to my family. Emily, you are growing up to be awesome, and I am looking forward to living near you again. Mom and Dad, you have made me the person I am today and have always been there no matter what. I could not have asked for two more loving and supportive parents.

Simon, you have stood by me every step of the way these past four years. You have kept me sane. I would never have made it to this point without you, and I can't wait to start our life together.

Finally, this thesis is dedicated to Betty Johnson and to the memory of Douglas Bushnell. This voyage would have never even started without your love and support, and for that I am eternally grateful.

Contents

“Yeah, well, you know, that’s just, like, your opinion, man.”

— The Dude, *The Big Lebowski*

Chapter 1

Introduction

Some of the most exciting particle physics discoveries of the last few decades have come from the study of neutrinos. These particles are present everywhere and arrive from many different sources: the sun, cosmic rays, nuclear reactors, decays inside the earth, supernovae, particle accelerators, and even the Big Bang. Despite this abundance of neutrinos, little was known about the physics governing their behavior until the past few decades. This is in large part due to the fact that neutrinos only interact with matter via the weak force, and are therefore very difficult to detect.

The discovery of neutrino oscillations opened up a whole new avenue of scientific inquiry. The fact that oscillations occur in nature at all means that neutrinos must have non-zero mass, something not predicted in the current version of the Standard Model. A decade's worth of research by various experiments has made a good start on quantifying these oscillations, and increasingly precise measurements have been made of the parameters involved.

One major neutrino oscillation parameter remains unmeasured at this point: the mixing angle θ_{13}. Recent experimental results from T2K (see Appendix A) have hinted that the value of this term is non-zero. Most measurements to date, however, have simply set upper limits, suggesting that the value is small. Beyond simple curiosity, there are compelling physics motivations for measuring this last mixing angle. In particular, a non-zero value for θ_{13} would allow for the possible incorporation of CP violation in the lepton sector. Consequentially, measuring θ_{13} is one of the top priorities of the neutrino physics community at the moment, driving a new generation of oscillation experiments.

This thesis will attempt to make a measurement of θ_{13} using the MINOS long baseline neutrino oscillation experiment. MINOS has the ability to measure θ_{13} by observing the oscillation of muon neutrinos to electron neutrinos between its Near and Far Detectors. This oscillation mode will only occur if θ_{13} is non-zero. MINOS, however, was not designed to look for electron neutrino interactions. A combination of poor detector resolution and a high degree of background contamination will necessitate the use of various sophisticated analysis techniques.

This thesis will begin with a description of the history of neutrinos and neutrino oscillations, as well as a basic overview of the related physics concepts. A description of the MINOS experiment itself will follow, including the neutrino beamline, the two detectors, and the preparation of data for analysis. The main body of the document will focus on the search for $\nu_\mu \rightarrow \nu_e$ oscillations. It will describe the development of a particle ID scheme, the prediction of backgrounds, and the fitting of the results. The document will conclude with an interpretation of the Far Detector data.

The work in this thesis uses techniques developed for previous ν_e analyses, while also attempting to improve on these methods wherever possible. Two major improvements over previous MINOS $\nu_e \rightarrow \nu_\mu$ searches will be highlighted. First, a particle ID technique known as the Library Event Matching (LEM) method was extensively studied and optimized by the author. The work in this thesis has both increased the performance of the particle ID and led to a better understanding of its behavior, so that LEM can be used effectively and reliably to search for a ν_e appearance signal. Second, unlike previous versions of this analysis, which were counting experiments, this thesis will employ a multiple bin likelihood fit to the shape of the LEM and energy distributions in order to measure θ_{13}.

The work in this thesis was conducted in parallel with the author's work on the 2011 MINOS ν_e analysis. The results from this published analysis are described in Appendix A. While much of the work performed for this thesis was also used in that analysis, the author has also further expanded or optimized many of the techniques used.

Chapter 2

Neutrinos: Theory and Background

2.1 A Brief History of Neutrino Physics

The important role played by the neutrino in high energy physics today is perhaps ironic given its origins as a “desperate remedy.” The particle was first proposed as an explanation for anomalous results observed in 1914 by Chadwick [1], in the energy spectrum of β decay. Both α and γ radiation produced a discrete spectrum of energy, equal to the energy lost between the initial and final state of the atomic nucleus. The spectrum for β decay was found to be continuous, with energy varying between zero and a maximum value. Only a single electron was observed as the decay product, so it was unclear where the remaining energy of the interaction had gone. Later calorimetric experiments by Ellis and Woolsey in 1927 [2] confirmed that this continuous energy spectrum was a real feature of β decay, rather than an experimental effect. This mounting experimental evidence was worrying enough that Bohr even considered the potential violation of energy conservation as an explanation.

In a letter in 1930, Wolfgang Pauli proposed a less drastic solution: the existence of a particle (which he termed a “neutron”) with spin 1/2 and a light mass, which would only interact via the weak force [3]. Turning β decay into a three body decay would solve the energy conservation problem. Giving the particle these specific physics properties would explain the lack of experimental detection. Enrico Fermi soon hypothesized a mechanism for β decay: $n \rightarrow p + e^- + \nu$. The β decay spectrum required the mass of the neutrino to be very small [4].

The neutrino would not be directly detected until nearly 20 years later, in 1954, with an experiment by Reines and Cowan [5]. A detector was placed near a nuclear reactor at the Hanford Site, in Washington. Electron antineutrinos from the reactor would interact with nucleons in the detector via the inverse beta decay reaction ($\bar{\nu}_e + p \rightarrow n + e^+$).

The photons emitted in neutron capture were then used to detect the interaction. A 2σ excess of interactions was observed when the reactor was turned on, versus when it was turned off. This was the first observation of a neutrino signal, and of the $\bar{\nu}_e$. Two other flavours of neutrino have since been discovered. Muon neutrinos were observed in 1962 in an experiment at the Brookhaven AGS facility by Lederman, Schwartz, and Steinberger [6]. Tau neutrinos were finally observed in 2001 in the DONUT experiment [7] at Fermilab.

The phenomenon now known as neutrino oscillation was detected for the first time in the Homestake experiment [8], although this result was not definitively identified as such until some time later. This experiment, which was started by Ray Davis in 1965, was designed to detect low energy (MeV scale) neutrinos produced by nuclear reactions in the sun. Davis used a radio-chemical technique to detect neutrino interactions. The detector consisted of a tank of Cl_2C_4 situated in the Homestake Mine in South Dakota. Electron neutrinos entering the detector tank interacted with the Chlorine nuclei to produce an unstable isotope of argon: $\nu_e + ^{37}\text{Cl} \rightarrow e^- + ^{37}\text{Ar}$. The argon atoms were then counted to determine the rate of neutrino interactions. When this rate was compared to the predictions from Bahcall's Standard Solar Model [9], a surprising discrepancy was seen: the SSM predicted 8.1 ± 1.2 SNU[1], but the Homestake experiment only observed 2.56 ± 0.25 SNU, less than half the predicted flux. There were three possible explanations for this discrepancy, which became known as the "Solar Neutrino Problem": the Solar Standard Model was wrong, the experiment was wrong, or something was happening to the neutrinos in between.

Numerous sources of strong experimental evidence have since confirmed the Standard Solar Model. Much of this proof comes from helioseismology: acoustic waves which propagate through the sun are detected via Doppler shifted light and used to characterize solar structure and dynamics. As an example of this evidence, the SSM prediction of the velocity of these sound waves as a function of radial distance agrees with the data to within 0.5% [10]. Certainly, no disagreement between the SSM prediction and data has been found on a scale to explain the Solar Neutrino Problem.

Following Homestake, numerous other experiments gathered additional evidence for the solar neutrino discrepancy, suggesting that the Homestake experiment itself was not at fault. The Kamiokande-II water Cerenkov detector in Japan detected solar neutrinos through the elastic scattering channel $\nu_e + e^- \rightarrow \nu_e + e^-$ and once again observed a deficit, with 0.45±0.02 SNU observed versus 1.0±0.2 SNU expected [11]. The GALLEX

---

[1] SNU $=10^{-26}$ neutrino interactions / target atom / s, a unit designed to accommodate the very low event rates in these experiments.

[12] and SAGE [13] experiments used a radio-chemical technique employing Gallium, which allowed the two experiments to explore a much lower energy range of the solar neutrino flux. Once again, a large deficit was found. As evidence for the soundness of the Standard Solar Model mounted, and more experiments with different methodologies showed evidence for there being a real deficit in the solar neutrino flux, explanations for the Solar Neutrino Problem began to focus on the third possible source of discrepancy: the neutrinos themselves. In particular, it was proposed that neutrinos could "oscillate" between their weak eigenstates as they traveled between the Sun and Earth. This possibility of this type of oscillation had first been suggested by Pontecorvo, in 1957 [14]. As will be described later in this chapter, neutrinos were finally shown to transition between separate weak eigenstates in the SNO experiment in 2001, opening up a whole new era of neutrino physics.

2.2 Neutrinos in the Standard Model

The Standard Model of particle physics provides a complete description of the electromagnetic, weak, and strong force interactions, tying them together with the associated force-carrying gauge bosons. The model also incorporates all of the known fundamental fermions, both quarks and leptons, in three generations of matter. Table 2.1 lists the properties and forces associated with the gauge bosons, while Table 2.2 summarizes the basic characteristics of the fermions. Quantum Chromodynamics (QCD) describes the strong force as a $SU(3)$ gauge group, while the electromagnetic and weak forces are described together under electroweak unification with a gauge symmetry of $SU(2)_L \otimes U(1)_Y$. Particles in the quark sector carry colour charge and can therefore participate in the strong force, in addition to electromagnetic and weak interactions. Electron, muon, and tau leptons carry electromagnetic charge, and can interact both electromagnetically and weakly. The three neutrinos ν_e, ν_μ, and ν_τ comprise the final three particles of the lepton section. They have no charge, and therefore only interact via the weak force.

2.2.1 Neutrinos and the Weak Force

Fermi made the first attempt to describe weak force interactions by constructing a "four point" model with a vector current very similar to that of the electromagnetic interaction. There were several problems with this model. Firstly, when the cross-

Boson	Charge	Spin	Mass (GeV)	Interaction
γ	0	1	0	Electromagnetic
$W^{\pm}$	±1	1	80	Weak
Z^0	1	1	91	Weak
Gluon (8)	0	1	0	Strong
Higgs	0	0	(Unknown)	-

Table 2.1: Summary of the basic force-carrying bosons in the Standard Model, with their charges, spins, and measured masses [15].

Lepton	Mass (MeV)	Charge	Quark	Mass (MeV)	Charge
e	0.511	-1	u	1.7-3.1	$\frac{2}{3}$
ν_e	~ 0	0	d	4.1-5.7	-$\frac{1}{3}$
μ	105.7	-1	c	1,290	$\frac{2}{3}$
ν_μ	~ 0	0	s	100	-$\frac{1}{3}$
τ	1,776.8	-1	t	172,900	$\frac{2}{3}$
ν_τ	~ 0	0	b	4,200	-$\frac{1}{3}$

Table 2.2: A summary of the charge and masses for Standard Model particles in the lepton sector (left) and the quark sector (right). Values are taken from [15]. The u, d, and s quark masses are current masses. The b and c masses are the running masses. The t mass is from direct observation.

section for a “Charged Current” (CC) interaction ($\nu_e + n \to p + e^-$) was calculated, it was found to increase linearly with energy to infinity. Secondly, it was necessary to provide a mechanism by which parity violation could be introduced into the weak interaction.

The first problem, of the divergent cross-sections for the weak interaction, can be solved by the introduction of a massive force propagator in the form $1/(M_W^2 - q^2)$ to the weak Charged Current interaction, to reduce the coupling strength of the force. The electroweak unification theory of Glashow [16], Salam [17], and Weinberg [18] proposed the existence of this massive $W^{\pm}$ boson, which governs Charged Current interactions. The theory also theorized the existence of a second, neutral boson, the Z^0 boson, which was responsible for a form of weak scattering known as a “Neutral Current” (NC) interaction. In the early 1980s, both bosons were observed in the UA1 and UA2 experiments [19] [20] at the CERN $p\bar{p}$ collider.

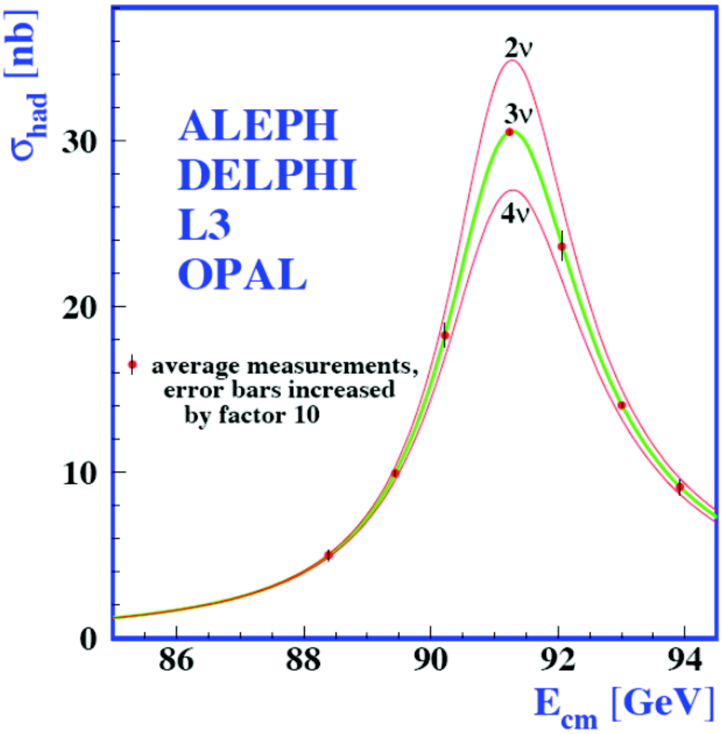

Figure 2.1: LEP study [21] of the decay of the Z^0 boson, showing the cross-section for hadron production around the resonance. A curve showing the predicted values for two, three, and four neutrino species is given.

Studies of the Z^0 boson have also allowed for a determination of the number of neutrino generations which couple to the weak force. This is done by examining the decay width of the Z^0 boson to determine the number of decay modes. From this measurement (see Figure 2.1), it was determined that only three neutrinos couple with the Z^0 boson [21]. This does not rule out the possibility of an additional neutrino more massive than the Z^0 boson.

The need to include parity violation in the weak force interaction came about from experimental observations. A parity transformation performs a spatial inversion through the origin. It is described by the operator $\hat{P}$, which, for a Dirac spinor $\psi(t, x, y, z)$, is the γ^0 matrix:

$$\psi'(t', x', y', z') = \hat{P}\psi(t, x, y, z) = \gamma^0\psi(t, x, y, z) = \psi(t, -x, -y, -z) \tag{2.1}$$

It can be shown that interactions in both QED and QCD conserve parity. The original formulation of the weak interaction, with its vector field current, assumed the conservation of parity as well. In 1956, an experiment by C.S. Wu [22] showed this assumption to be wrong: the weak force in fact violates parity. Wu constructed an experiment in which ^{60}Co emitted β decay electrons. This emission occurred while the atoms were in a magnetic field, causing their spin vectors to align. If parity is conserved, these β electrons should have been emitted symmetrically with respect to the nuclear spin. Instead, Wu observed that the β electrons were emitted preferentially in the direction opposite to the spin. This asymmetry was evidence of parity violation.

Parity violation can be built into the Charged Current weak interaction by choosing an alternate form for the interaction vertex. In addition to the vector interaction $(\bar{\psi}\gamma^\mu\phi)$, there can also be an axial vector component: $\bar{\psi}\gamma^\mu\gamma^5\phi$ (where $\gamma^5 = i\gamma^0\gamma^1\gamma^2\gamma^3$). Experiment has shown that the actual form of the weak interaction is **V-A**, with a current of the form:

$$j^\mu \propto \bar{u}(\gamma^\mu - \gamma^\mu\gamma^5)u = \bar{u}\gamma^\mu(1-\gamma^5)u \tag{2.2}$$

While the separate vector and axial vector parts of the current each individually conserve parity, the combination of them does not.

The **V-A** vertex also has important consequences for neutrino physics. Equation 2.2 contains the left-handed chiral projection operator, $\frac{1}{2}(1-\gamma^5)$:

$$\bar{\psi}\gamma^\mu\frac{1}{2}(1-\gamma^5)\phi = \bar{\psi}\gamma^\mu\phi_L \tag{2.3}$$

Since $\bar{\psi}_R\gamma^\mu\phi_L = 0$, only the left-handed chiral components of a particle spinor (or the right-handed components of an anti-particle spinor) can participate in the weak interaction:

$$\bar{\psi}\gamma^\mu\frac{1}{2}(1-\gamma^5)\phi = \bar{\psi}\gamma^\mu\phi_L = (\bar{\psi}_L + \bar{\psi}_R)\gamma^\mu\phi_L = \bar{\psi}_L\gamma^\mu\phi_L \tag{2.4}$$

Neutrinos are ultra-relativistic particles, so that their helicity is equal to their chirality. Therefore, only left-handed neutrinos and right-handed antineutrinos can participate in the Charged Current weak interaction. In neutrino experiments, the neutrinos are created via the weak Charged Current interaction, and are therefore all left-handed.

In addition to not conserving parity, the weak interaction also does not conserve charge conjugation symmetry. The charge conjugation operator $\hat{C}$ transforms a particle into its antiparticle. Applying this symmetry to an interaction which includes a left-handed neutrino, such as $\pi^+ \to \mu^+ + (\nu_\mu)^L$, results in a left-handed antineutrino, which cannot be created in a CC interaction: $\pi^- \to \mu^- + (\overline{\nu}_\mu)^L$. The combination of these

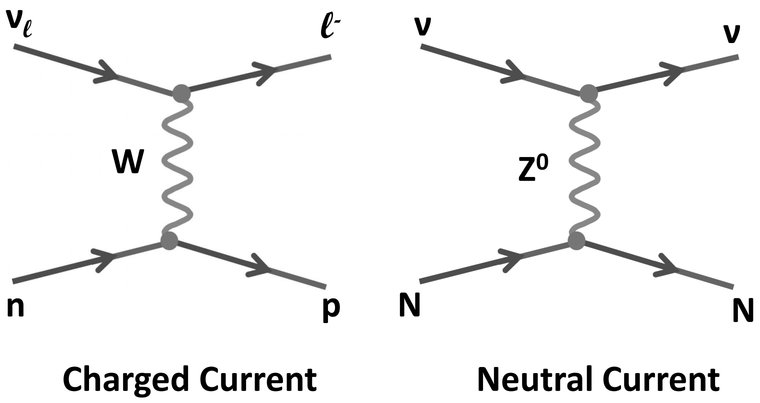

Figure 2.2: Feynman diagrams for the weak Charged Current (left) and Neutral Current (interactions). These particular examples are of a neutrino interacting with a nucleon.

two symmetries (CP) can, however, be conserved, as CP symmetry turns the left left-handed neutrino in the above interaction to a right handed antineutrino. CP violation has already been observed in the quark sector. As will be explained later, it may exist in the lepton sector as well.

2.2.2 Interaction Types

Neutrinos are detected either via the Charged Current or Neutral Current weak interaction. Examples of these two interactions are shown in Figure 2.2 for the case of a neutrino interacting with a nucleon. Firstly, in the Charged Current interaction, an electron neutrino, muon neutrino, or tau neutrino is converted to (respectively) an electron, muon, or tau via the emission of a W boson. A neutrino interacting with a nucleon will produce a charged lepton of the same flavour, which can then be detected by experiment. Secondly, the Neutral Current interaction is mediated by the Z^0 boson. The Z^0 boson couples fermion to fermion, without any transfer of charge, but with the ability to couple to right handed chiral states. This proceeds as a scattering interaction, in which the flavour of the original lepton is not detected.

In additional to these broad categories of Neutral Current and Charged Current events, there are several different subcategories of interaction which can be observed in experiment. Feynman diagrams for three such processes (Deep Inelastic Scattering,

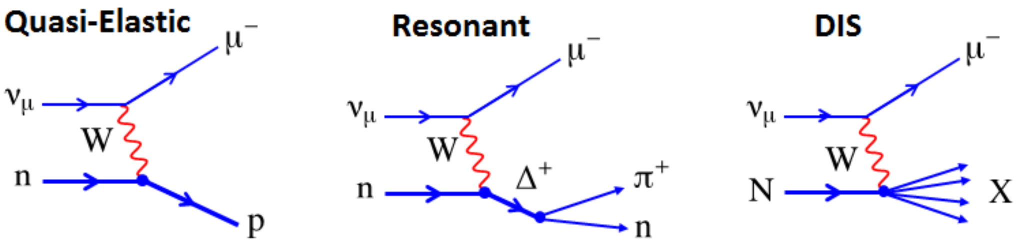

Figure 2.3: Feynman diagrams for three forms of Charged Current interaction: Quasielastic, Resonance Production, and Deep Inelastic Scattering. Image from [23].

Quasielastic, and Resonance Production interactions) are shown in Figure 2.3. The cross-sections for the ν_μ CC interaction, $\nu_\mu + N \to \mu^- + X$ are shown in Figure 2.4 as a function of energy. Quasielastic interactions (QE) proceed as $\nu_\ell + N \to \ell^- + p$. These interactions dominate the neutrino cross-section at low energy. QE interactions are often the signal in a neutrino experiment, as the event topology is very simple and thus it is easy to determine the original neutrino flavour. Deep Inelastic Scattering (DIS) interactions dominate the cross-section at higher energies. In this case, a virtual W or Z^0 boson probes the structure of the nucleon. The nucleon is broken up, resulting in a large number of daughter particles in a hadronic shower, which can obscure the lepton. Resonance interactions (RES) occur at an intermediate energy. In the case shown, a Δ resonance is created: $\nu_\ell + N \to \Delta^+ + \ell$. The Δ then decays, producing various daughter particles such as pions. Finally, so-called Coherent Pion Production (COH) can occur when the neutrino interacts coherently with the whole nucleus. This can produce a single pion: a π^+ for CC interactions and a π^0 for NC interactions. This process is rare and the cross-section poorly understood, but it can constitute an important background for many neutrino experiments.

2.2.3 Neutrino Mass

The current form of the Standard Model assumes a massless neutrino. The phenomenon of neutrino oscillation, however, strongly suggests that neutrinos have mass. The mass of the three neutrinos is known from experiment to be extremely small. The current best upper limits [15] on the neutrino masses from decay experiments are $m_{\nu_e} < 2$ eV, $m_{\nu_\mu} < 0.19$ MeV, and $m_{\nu_\tau} < 18.2$ MeV. Studies of the Cosmic Microwave background from the most recent WMAP results, combined with distance information from supernovae and Baryon Acoustic Oscillations (BAO), place an upper limit on the total sum of neutrino

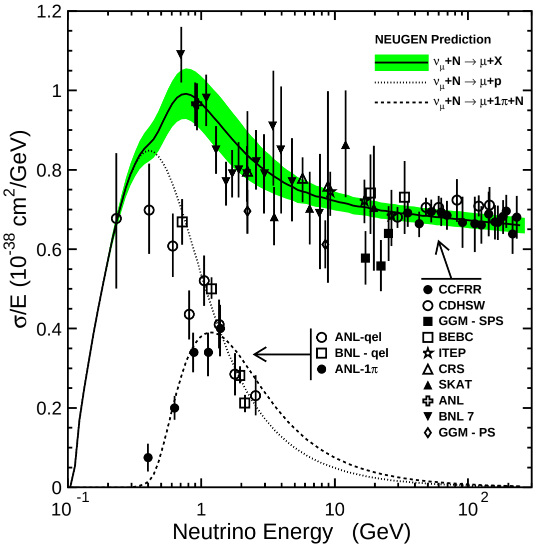

Figure 2.4: NEUGEN-v3 [74] cross-section for ν_μ Charged Current interactions as a function of neutrino energy. Cross-sections shown include QE, RES, and the inclusive cross-section (which includes DIS). The shaded band shows the uncertainty on the cross-section. Cross-sections are for an isoscalar target. Data from past experiments (compiled by [24]) are shown as well. Image from [53].

masses, at $\sum m_\nu < 0.67$ eV (at 95% CL) [28].

The so-called "Seesaw" Mechanism provides a means of introducing a non-zero, but still very small neutrino mass into the current theoretical framework. For the fermions in the Standard Model, one can write a Dirac mass term:

$$\mathcal{L}_D = -m_D(\overline{\nu}_L \nu_R + \overline{\nu}_R \nu_L) \tag{2.5}$$

All quarks and charged leptons are Dirac particles, with mass arising from the Higgs mechanism and Yukawa couplings. If one assumes the existence of a right-handed neutrino field, this method can be used to describe the nonzero mass arising from a Dirac neutrino:

$$\mathcal{L}_D = -f_\nu \phi \overline{\nu}_L \nu_R = -f_\nu \langle \phi \rangle \overline{\nu}_L \nu_R = -f_\nu V \overline{\nu}_L \nu_R \tag{2.6}$$

where f_ν is a Yukawa coupling constant, ϕ is the neutral Higgs field, and V is the vacuum expectation value, 246 GeV [15]. To produce a neutrino mass of the expected size, the coupling constant must be on the order of 10^{-13}, as compared to that of the charged leptons, 10^{-4}. While this is not impossible, there is no reason that f_ν should be so comparatively small. Some additional mechanism must be found to explain the small observed mass.

In addition to the above Dirac mass term, one can also write down a Majorana [25] mass term, constructed solely out of the left-handed field:

$$\mathcal{L}_{M_L} = -\frac{1}{2} m_M^L (\overline{\nu}_L \nu_L^c + \overline{\nu}_L^c \nu_L) \tag{2.7}$$

Here, ν is a neutrino field (either left handed L or right handed R) and ν^c is the charged-conjugated field: $\nu^c = C\overline{\nu}^T = C\gamma_0\nu^*$. Unfortunately, Equation 2.7 violates gauge invariance, requiring that $\mathcal{L}_{M_L} = 0$. However, there are no such restrictions on a right-handed Majorana mass term, $\mathcal{L}_{M_R}$, once one assumes the existence of a right-handed neutrino field:

$$\mathcal{L}_{M_R} = -\frac{1}{2} m_M^R (\overline{\nu}_R \nu_R^c + \overline{\nu}_R^c \nu_R) \tag{2.8}$$

For a particle to be a Majorana particle, it must be identical to its own antiparticle, a feature which is impossible if the particle possesses charge. Neutrinos are the only fermions in the Standard model which are electrically neutral and therefore the only fermions which could be Majorana particles. Notably, the Majorana mass term in 2.8 also violates lepton number, a feature beyond the current Standard Model.

Various forms of Grand Unified Theory, however, do permit both the violation of lepton number. Under these GUTs, right handed neutrinos can exist, and neutrinos can therefore be (massive) Majorana particles. Gell-Mann, Ramond, and Slansky [26], as well as Yanigada [27], therefore proposed the See-Saw Mechanism to account for the small observed neutrino mass. Under this mechanism, the mass terms $\mathcal{L}_D$ and $\mathcal{L}_{M_R}$ can be combined together to form a general mass term in matrix form:

$$\mathcal{L} = -\frac{1}{2} \begin{pmatrix} \overline{\nu}_L & \overline{\nu}_R^c \end{pmatrix} \begin{pmatrix} 0 & m_D \\ m_D & m_M^R \end{pmatrix} \begin{pmatrix} \nu_L^c \\ \nu_R \end{pmatrix} + h.c. \tag{2.9}$$

The matrix in the above expression is known as the “seesaw matrix.” Diagonalizing this matrix results in two eigenvalues:

$$m\pm = \frac{1}{2} m_M^R \pm \frac{1}{2}\sqrt{(m_M^R)^2 + 4m_D^2} \tag{2.10}$$

If $m_M^R \gg m_D$, the two corresponding eigenstates are approximately:

$$\nu_+ \approx (\nu_R + \nu_R^c) + \frac{m_D}{m_M^R}(\nu_L + \nu_L^c) \tag{2.11}$$

$$\nu_- \approx (\nu_L - \nu_L^c) + \frac{m_D}{m_M^R}(\nu_R + \nu_R^c) \tag{2.12}$$

From this $m_M^R \gg m_D$ assumption, there is first one "heavy" state $N = \nu_+$, which is dominated by the sterile right handed term, with mass $m_N \approx m_M^R$. Second, there is a "light" state $\nu = \nu_-$ with mass $m_\nu \approx \frac{m_D^2}{m_M^R}$. With a Dirac mass m_D of the same approximate size as the charged leptons (MeV scale), and a right handed Majorana mass M_M^R just below the GUT scale (10^{15} eV), the light mass m_ν will be on the order of 1 meV. This is the neutrino observed in experiment. The heavy right handed partner N has never been observed, but could have potentially been created in the early universe and then decayed during cooling. This model is easily adapted to three neutrino generations, with many different theoretical variations.

Support for the above mechanism can come from a variety of different experimental studies. Neutrino oscillations, as will be explained, provide concrete evidence for a non-zero neutrino mass, as well as a measurement of the size of the difference between those masses (squared). Other experiments are currently attempting to make a direct measurement of neutrino mass using the spectrum of β decay electrons. In the presence of a non-zero neutrino mass, the slope of this spectrum will change, and the maximum allowed decay energy will shift. Using this method for tritium decay, the Mainz experiment has placed the current best limit on the ν_e mass, at < 2.3 eV [29].

Finally, experiments searching for neutrino-less double beta decay can both measure the neutrino mass, and determine if neutrinos are a Majorana particle. In regular two-neutrino double beta decay, the decay proceeds as $(Z, A) \rightarrow (Z + 2, A) + 2e^- + 2\overline{\nu}_e$. This is a rare process, with a very long half life. The existence of a massive neutrino, however, also allows for "neutrino-less" double beta decay: $(Z, A) \rightarrow (Z + 2, A) + 2e^-$. This process, which violates lepton number, occurs when a neutron emits a right-handed antineutrino, which is then re-absorbed by another neutron as a left-handed neutrino. This can only occur if a neutrino is a Majorana particle, i.e. if it both has mass and is its own anti-particle. In addition to being an important physics result in its own right, observing this interaction would allow for a direct measurement of neutrino mass, as the rate of this interaction is a function of the effective mass[2]. The controversial

---

[2] Using the mixing matrix format introduced in the oscillation section, the effective neutrino mass is defined as $m_{ee} \equiv \sum_i U_{ei}^2 m_i$

2004 Heidelberg-Moscow result [30] claims a 4σ discovery of this process, with neutrino effective mass < 0.35 eV. A new generation of experiments will soon test this claim.

2.3 Neutrino Oscillations

Neutrinos are produced and interact with matter in three separate weak eigenstates: ν_e ν_μ and ν_τ. However, they also propagate through space and time via three separate mass eigenstates: ν_1, ν_2, and ν_3. If the masses of these eigenstates are non-zero, and if the mass eigenstates are not equivalent to the three weak eigenstates, the phenomenon of neutrino oscillation can occur.

2.3.1 Oscillation with Two Neutrinos

The basic physics of neutrino oscillation can be described with the simple case of only two neutrinos, which have weak eigenstates $|\nu_\alpha\rangle$ and $|\nu_\beta\rangle$ and mass eigenstates $|\nu_1\rangle$ and $|\nu_2\rangle$. A neutrino produced in weak eigenstate $|\nu_\alpha\rangle$ is also a superposition of the two mass eigenstates:

$$|\Psi(0)\rangle = |\nu_\alpha\rangle = U_{\alpha 1}|\nu_1\rangle + U_{\alpha 2}|\nu_2\rangle \tag{2.13}$$

where $U_{\alpha 1}$ and $U_{\alpha 2}$ are constants. These mass eigenstates, in turn, propagate as free-particle plane waves:

$$|\nu_1(t)\rangle = |\nu_1\rangle e^{i\vec{p}_1\cdot\vec{x}-iE_1t},\ |\nu_2(t)\rangle = |\nu_2\rangle e^{i\vec{p}_2\cdot\vec{x}-iE_2t} \tag{2.14}$$

These expressions can be further simplified by defining a phase $\phi_1 = iE_1t - \vec{p}_1\cdot\vec{x}$:

$$|\nu_1(t)\rangle = |\nu_1\rangle e^{-i\phi_1},\ |\nu_2(t)\rangle = |\nu_2\rangle e^{-i\phi_2} \tag{2.15}$$

The two weak eigenstates can be described in terms of the mass eigenstates with a 2×2 unitary matrix, parametrized by a "mixing" angle, θ:

$$\begin{pmatrix}\nu_\alpha\\ \nu_\beta\end{pmatrix} = \begin{pmatrix}\cos\theta & \sin\theta\\ -\sin\theta & \cos\theta\end{pmatrix}\begin{pmatrix}\nu_1\\ \nu_2\end{pmatrix} \tag{2.16}$$

One can now write an expression for the wave function of $|\Psi(0)\rangle = |\nu_\alpha\rangle$, after it has propagated distance L:

$$|\Psi(L)\rangle = \cos\theta|\nu_1\rangle e^{-i\phi_1} + \sin\theta|\nu_2\rangle e^{-i\phi_2} \tag{2.17}$$

Converting back to the weak eigenstate basis yields the following expression:

$$|\Psi(L)\rangle = (\cos^2\theta e^{-i\phi_1} + \sin^2\theta e^{-i\phi_2})|\nu_\alpha\rangle + (\cos\theta\sin\theta)(e^{-i\phi_2} - e^{-i\phi_1})|\nu_\beta\rangle \tag{2.18}$$

If $\phi_1 = \phi_2$, the neutrino will continue to propagate in its original weak eigenstate, i.e. $|\Psi(L)\rangle = |\nu_\alpha\rangle$. If $\phi_1 \neq \phi_2$, the neutrino will propagate as a superposition of *two* weak eigenstates. There is now a non-zero probability of observing the neutrino in the ν_β eigenstate:

$$P(\nu_\alpha \to \nu_\beta) = |\langle\nu_\beta|\Psi(L)\rangle|^2 = \sin^2(2\theta)\sin^2\left(\frac{\phi_2 - \phi_1}{2}\right) \tag{2.19}$$

The phase difference $\phi_2 - \phi_1$ can be rigorously derived with a proper wave packet treatment. However, the same final result can be obtained by assuming $p_1 = p_2 = p$, $t \approx L$ (in natural units), and $p \gg m_{1,2}$:

$$\begin{aligned}
\phi_2 - \phi_1 &= (E_2 t - pL) - (E_1 t - pL) = (E_2 - E_1)t \\
&= \left[p\left(1 + \frac{m_2^2}{p^2}\right)^{\frac{1}{2}} - p\left(1 + \frac{m_1^2}{p^2}\right)^{\frac{1}{2}}\right] t \\
&\approx \frac{m_2^2 - m_1^2}{2p} L = \frac{\Delta m_{21}^2}{2p} L \approx \frac{\Delta m_{21}^2}{2E} L
\end{aligned} \tag{2.20}$$

Substituting this into Equation 2.19 yields:

$$P(\nu_\alpha \to \nu_\beta) = \sin^2(2\theta)\sin^2\left(\frac{\Delta m_{21}^2}{4E} L\right) \tag{2.21}$$

Thus, if the two mass eigenstates have different masses and are distinct from the weak eigenstates (via the above neutrino mixing matrix), there is a non-zero probability that a neutrino produced in weak eigenstate $|\nu_\alpha\rangle$ will be observed as $|\nu_\beta\rangle$ after some distance of propagation. This phenomenon is known as neutrino oscillation.

2.3.2 Oscillation with Three Neutrinos

While this shows the basic mechanism of neutrino oscillation, a full picture must account for the existence of three types of neutrino: weak eigenstates $|\nu_e\rangle$, $|\nu_\mu\rangle$, and $|\nu_\tau\rangle$, and mass eigenstates $|\nu_1\rangle$, $|\nu_2\rangle$, and $|\nu_3\rangle$. This section will describe the physics of a neutrino traveling in a vacuum; the case of a neutrino traveling through matter will be discussed later. The relationship between the mass and weak eigenstates is again parametrized by a unitary matrix:

$$\begin{pmatrix} \nu_e \\ \nu_\mu \\ \nu_\tau \end{pmatrix} = \begin{pmatrix} U_{e1} & U_{e2} & U_{e3} \\ U_{\mu 1} & U_{\mu 2} & U_{\mu 3} \\ U_{\tau 1} & U_{\tau 2} & U_{\tau 3} \end{pmatrix} \begin{pmatrix} \nu_1 \\ \nu_2 \\ \nu_3 \end{pmatrix} \tag{2.22}$$

so that a given weak eigenstate is written as follows:

$$|\nu_\alpha\rangle = \sum_{i=1}^{3} U_{\alpha i}|\nu_i\rangle \tag{2.23}$$

Using methods similar to before, one can derive the probability of one weak eigenstate ν_α oscillating to ν_β:

$$\begin{aligned} P(\nu_\alpha \to \nu_\beta) &= |\langle \nu_\beta | \Psi(L)\rangle|^2 = |U_{\alpha 1}U^*_{\beta 1}e^{-i\phi_1} + U_{\alpha 2}U^*_{\beta 2}e^{-i\phi_2} + U_{\alpha 3}U^*_{\beta 3}e^{-i\phi_3}|^2 \\ &= \sum_{ij} U_{\beta i}U^*_{\alpha i}U^*_{\beta j}U_{\alpha j}e^{-i(\phi_i - \phi_j)} = \sum_{ij} U_{\beta i}U^*_{\alpha i}U^*_{\beta j}U_{\alpha j}e^{-i\frac{\Delta m^2_{ij}}{2E}L} \end{aligned} \tag{2.24}$$

This equation can be further simplified by exploiting a few of its mathematical features. First, the unitarity of the 3×3 mixing matrix means that $\sum_{i=1}^{3} U_{\beta i}U^*_{\alpha i} = \delta_{\alpha\beta}$. Additionally, term ij is equal to the complex conjugate of term ji, so that only the real components for i>j need be summed. The probability can be rewritten as:

$$\begin{aligned} P(\nu_\alpha \to \nu_\beta) &= \delta_{\alpha\beta} + 2\sum_{i>j} \Re[U_{\beta i}U^*_{\alpha i}U^*_{\beta j}U_{\alpha j}(e^{-i\frac{\Delta m^2_{ij}}{2E}L} - 1)] \\ &= \delta_{\alpha\beta} - 4\sum_{i>j} \Re[U_{\beta i}U^*_{\alpha i}U^*_{\beta j}U_{\alpha j}] \sin^2\left(\frac{\Delta m^2_{ij}}{4E}L\right) \\ &\quad +2\sum_{i>j} \Im[U_{\beta i}U^*_{\alpha i}U^*_{\beta j}U_{\alpha j}] \sin^2\left(\frac{\Delta m^2_{ij}}{2E}L\right) \end{aligned} \tag{2.25}$$

In addition to describing three-neutrino oscillations, this expression has important physics implications. In particular, it allows for the possibility of neutrino oscillation exhibiting CP violation. If CP symmetry is conserved, then $P(\nu_\alpha \rightarrow \nu_\beta) = P(\overline{\nu}_\alpha \rightarrow \overline{\nu}_\beta)$. Similarly, if oscillations conserve time reversal symmetry (T), then: $P(\nu_\alpha \rightarrow \nu_\beta) = P(\nu_\beta \rightarrow \nu_\alpha)$. Neutrino oscillations are expected to conserve the combined CPT symmetry: $P(\nu_\alpha \rightarrow \nu_\beta) = P(\overline{\nu}_\beta \rightarrow \overline{\nu}_\alpha)$. Therefore, either both CP and T are conserved, or both are violated. Time reversal symmetry can be broken if any of the elements of the 3×3 mixing matrix are complex[3], so that for the probability in Equation 2.25, $P(\nu_\alpha \rightarrow \nu_\beta) \neq P(\nu_\beta \rightarrow \nu_\alpha)$. In order to preserve CPT, neutrinos must then also exhibit CP violation: $P(\nu_\alpha \rightarrow \nu_\beta) \neq P(\overline{\nu}_\alpha \rightarrow \overline{\nu}_\beta)$.

CP violation has already been observed in the quark sector, via the mechanism of the CKM matrix. It is natural to expect to observe it in the lepton sector as well. In fact, several theories in particle physics depend on this possibility. In the theory of leptogenesis [31], neutrinos are Majorana particles, with a heavy right handed neutrino partner. This partner existed in the earlier universe and then decayed to left handed neutrinos or right handed antineutrinos, plus Higgs bosons (which then further decay into heavy quarks). If CP violation occurred in this decay process, favoring one decay mode over another (even to a small degree), it could help to explain the observed matter/antimatter asymmetry of today's universe. Observing CP violation in the oscillation of today's light neutrinos would first step in studying this asymmetry in the lepton sector.

2.3.3 The PMNS Mixing Matrix

The 3×3 mixing matrix in Equation 2.22 is formally known as the Pontecorvo-Maki-Nakagawa-Sakata (PMNS) mixing matrix [32] [33]. The PMNS mixing matrix is parametrized with three separate rotation or "mixing" angles, θ_{12}, θ_{23}, θ_{13}, and a complex phase δ. The matrix can be written in the following convenient parametrization, with $c_{ij} \equiv \cos\theta_{ij}$

---

[3] Note that the mixing matrix can have an overall complex phase without causing time reversal asymmetry. Rather, the asymmetry is brought about, for instance, by an element having a **different** complex phase than the other elements.

and $s_{ij} \equiv \sin\theta_{ij}$:

$$U = \begin{pmatrix} c_{12}c_{13} & s_{12}c_{13} & s_{13}e^{-i\delta} \\ -s_{12}c_{23} - c_{12}s_{13}s_{23}e^{i\delta} & c_{12}c_{23} - s_{12}s_{13}s_{23}e^{i\delta} & c_{13}s_{23} \\ s_{12}s_{23} - c_{12}s_{13}c_{23}e^{i\delta} & -c_{12}s_{23} - s_{12}s_{13}c_{23}e^{i\delta} & c_{13}c_{23} \end{pmatrix} \tag{2.26}$$

$$U = \begin{pmatrix} 1 & 0 & 0 \\ 0 & c_{23} & s_{23} \\ 0 & -s_{23} & c_{23} \end{pmatrix} \times \begin{pmatrix} c_{13} & 0 & s_{13}e^{-i\delta} \\ 0 & 1 & 0 \\ -s_{13}e^{i\delta} & 0 & c_{13} \end{pmatrix} \times \begin{pmatrix} c_{12} & s_{12} & 0 \\ -s_{12} & c_{12} & 0 \\ 0 & 0 & 1 \end{pmatrix} \tag{2.27}$$

$$U = U_{23} \times U_{13} \times U_{12} \tag{2.28}$$

Each of these “sub-matrices” affects different types of neutrino oscillation, depending on the length L and energy E of oscillation. Solar neutrino oscillations are observed in neutrinos either from the sun ($L \sim 10^8$ km) or reactors ($E \sim 1$ MeV), and therefore have a larger L/E. These oscillations are mostly dominated by the U_{12} part of the matrix and the mass splitting term $\Delta m^2_{12} \approx 8 \times 10^{-5}\text{eV}^2$. So-called atmospheric neutrino oscillations are instead dominated by the U_{23} part of the matrix, with $\Delta m^2_{32} \approx 3 \times 10^{-3}\text{eV}^2$. These oscillations occur for neutrinos produced either in cosmic rays or accelerator beams, with a relatively short baseline ($L < 10^4$ km) and large energy (~ 1 GeV), for a smaller value of L/E. The middle matrix, U_{13}, is the least well-explored by experiment. It is also extremely important, as it couples to a complex phase factor δ, which could potentially introduce CP violation into neutrino oscillation.

2.3.4 Solar Neutrino Oscillations

As described in Section 2.1, the first evidence for neutrino oscillations came from the study of solar neutrinos. Another solar neutrino experiment, SNO [34], was the first experiment to confirm that neutrinos can change flavour. The first generation of solar neutrino oscillation experiments (Homestake, Super-Kamiokande, GALLEX, SAGE) were only able to detect neutrinos via the Charged Current interaction. These experiments operated at low energies, in the range of <30 MeV, well below the energy threshold needed to produce a muon or tau in a CC interaction. Oscillation was therefore observed as a deficit in the expected ν_e CC interaction rate, rather than the appearance of a second flavour. Neutral Current interactions have no such threshold problem. Unlike the earlier solar experiments, SNO was able to observe these NC interactions in addition to the CC mode. This allowed for a measurement of the overall neutrino flux, regardless of flavour.

SNO consisted of a 1 kiloton tank of heavy water (D_2O instead of H_2O) situated 2 km underground in a salt mine in Sudbury, Ontario. SNO was a water Cerenkov detector, and was capable of observing three main channels of neutrino interaction:

- **Charged Current:** $\nu_e + d(n,p) \rightarrow p + p + e^-$. Because of energy threshold requirements, only ν_e undergoes this interaction. This interaction provides a measurement of the ν_e flux: Rate(CC) $\propto \phi(\nu_e)$.
- **Elastic Scattering:** $\nu + e^- \rightarrow \nu + e^-$. All three flavours of neutrino undergo this interaction via Neutral Current interactions. However, ν_e can also interact this way via a Charged Current interaction: Rate(ES) $\propto \phi(\nu_e) + 0.15(\phi(\nu_\mu) + \phi(\nu_\tau))$.
- **Neutral Current:** $\nu + d(n,p) \rightarrow n + p + \nu$. All three flavours undergo this interaction, so the rate is proportional to the overall flux: Rate(NC) $\propto \phi(\nu_e) + \phi(\nu_\mu) + \phi(\nu_\tau)$.

When the rates of all three interaction types were measured, the total neutrino flux was found to be consistent with the Solar Standard Model, with the flux of the remaining two neutrinos accounting for the deficit observed in the ν_e flux. This is clear evidence for neutrino flavour transitions.

Other experiments have confirmed the SNO result. The KamLAND [35] experiment investigates neutrino oscillation at the solar oscillation scale by studying electron antineutrinos from reactors throughout Japan. The $\overline{\nu}_e$ have energies between approximately 1 and 10 GeV. The effective baseline L_0 of the KamLAND experiment, i.e., the flux-weighted average, is 180 km. The antineutrinos oscillate rapidly as they propagate, and are detected when they interact in the KamLAND detector via inverse beta decay in liquid scintillator. The antineutrino survival probability measured in KamLAND as a function of L_0/E is consistent with a specific hypothesis of neutrino oscillation.

Both solar neutrino experiments and KamLAND place limits on the solar oscillation parameters θ_{12} and Δm^2_{12}. These limits are shown in Figure 2.5. When the results from both solar neutrino experiments and KamLAND are combined, the following best fits [15] for the solar oscillation parameters are found:

$$\Delta m^2_{12} = (7.59 \pm 0.21) \times 10^{-5}\mathrm{eV}^2, \;\; \sin^2(2\theta_{12}) = 0.861^{+0.026}_{-0.022} \tag{2.29}$$

The discussion of neutrino oscillation so far has assumed oscillation in a vacuum. A full understanding of the solar results requires a consideration of the effects of matter on neutrino oscillation - a phenomenon known as the Mikheyev-Smirnov-Wolfenstein (MSW) effect [46] [47]. Electron neutrinos produced in the core of the Sun must first

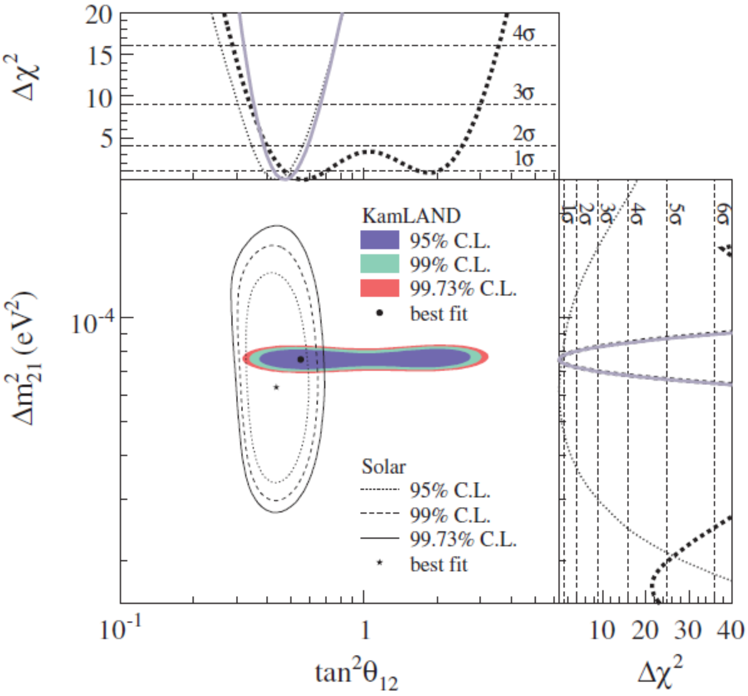

Figure 2.5: Allowed values for solar neutrino oscillation parameters (θ_{12} and Δm^2_{12}. Limits from KamLAND (colour) and solar neutrino experiments (black and white) are shown [35].

pass through the Sun's dense outer layers, interacting weakly with matter along the way. Due to these interactions, the energy E of the neutrino state becomes $E+V$, where V is an interaction potential. Some of this potential will come from Neutral Current interactions:

$$V_m^{NC} = \frac{\mp 1}{\sqrt{2}} G_F n_n \tag{2.30}$$

where G_F is the Fermi Constant and n_n is the density number of neutrons. All three neutrino flavours undergo this scattering with the same strength. V_m^{NC} therefore does not affect the phase of the weak eigenstates, and there is no impact on oscillations. Some of the potential also comes from Charged Current interactions with electrons:

$$V_m^{CC} = \pm\sqrt{2} G_F n_e \tag{2.31}$$

where n_e is the density number of electrons. This interaction will only occur for ν_e, as the energy threshold is too high for ν_μ and ν_τ. Because only one of the three weak states is affected, V_m^{CC} introduces a phase difference in the weak eigenstates, which *will* have an effect on neutrino oscillations.

The impact of matter effects can be demonstrated using a two-neutrino case, where only one weak state $|\nu_\alpha\rangle$ is affected by potential V_α. In a vacuum, the Hamiltonian operator H for a neutrino can be written as follows in terms of H_0 (removing any terms proportional to the identity matrix) and the PMNS mixing matrix U:

$$U = \begin{pmatrix} \cos\theta & \sin\theta \\ -\sin\theta & \cos\theta \end{pmatrix}, \quad H_0 = \frac{1}{2E}\begin{pmatrix} m_1^2 & 0 \\ 0 & m_2^2 \end{pmatrix} = \begin{pmatrix} 0 & 0 \\ 0 & \frac{\Delta m_{21}^2}{2E} \end{pmatrix} \tag{2.32}$$

$$H = U H_0 U^\dagger \tag{2.33}$$

The presence of potential V_α results in a modified Hamiltonian, H':

$$H' = H + V = U H_0 U^\dagger + \begin{pmatrix} V_\alpha & 0 \\ 0 & 0 \end{pmatrix} = \frac{\Delta m_{21}^2}{2E}\begin{pmatrix} \sin^2\theta + \frac{2E}{\Delta m_{21}^2} V_\alpha & -\sin\theta\cos\theta \\ -\sin\theta\cos\theta & \cos^2\theta \end{pmatrix} \tag{2.34}$$

The new Hamiltonian H' can now be diagonalized, to form a matrix with a new "effective" mixing angle θ_m, with eigenvalues λ_1 and λ_2 (where any terms proportional to the identity matrix are again removed):

$$H' = \frac{\Delta m_{21}^2}{2E} U(\theta_m) \begin{pmatrix} \lambda_1 & 0 \\ 0 & \lambda_2 \end{pmatrix} U^\dagger(\theta_m) = \frac{\Delta m_{21}^2}{2E} U(\theta_m) \begin{pmatrix} 0 & 0 \\ 0 & \lambda_2 - \lambda_1 \end{pmatrix} U^\dagger(\theta_m) \tag{2.35}$$

This now resembles the vacuum oscillation expression. Solving for the eigenvalues above yields both the effective mixing angle θ_m and the effective mass splitting term Δm_m^2:

$$\sin 2\theta_m = \frac{\sin 2\theta}{\sqrt{\sin^2 2\theta + (\cos 2\theta - \frac{2E}{\Delta m_{21}^2} V_\alpha)^2}}, \tag{2.36}$$

$$\Delta m_m^2 = \Delta m_{21}^2 \sqrt{\sin^2 2\theta + \left(\cos 2\theta - \frac{2E}{\Delta m_{21}^2} V_\alpha \right)^2} \tag{2.37}$$

which can be inserted into a new two-neutrino oscillation equation:

$$P_m(\nu_\alpha \rightarrow \nu_\beta) = \sin^2(2\theta_m) \sin^2 \left(\frac{\Delta m_m^2}{4E} L \right) \tag{2.38}$$

These equations have a few important features. First, the size of the new mixing angle is now dependent on the sign of Δm_{21}^2. Second, as $V_\alpha \rightarrow 0$, $\theta_m \rightarrow \theta$ and $\Delta m_m^2 \rightarrow \Delta m_{21}^2$, and the expression for vacuum oscillation is obtained. However, as $V_\alpha \rightarrow \infty$, $\sin 2\theta_m \rightarrow 0$, and there is no oscillation. Finally, for some values of potential V_α, $\sin 2\theta_m \rightarrow 1$, and oscillation will be maximal, regardless of the vacuum oscillation properties. These values are known as MSW "resonances."

The MSW effect has a large impact on the oscillation of neutrinos from the Sun. The potential $V_\alpha = V_m^{CC}$ (from Equation 2.31) is proportional to the solar electron number density n_e, which varies with solar radius. The oscillation of neutrinos is initially suppressed in the dense core of the Sun, where V_m^{CC} is very high. Lower energy solar neutrinos (< 2 MeV), in fact, will effectively not oscillate at all inside the Sun, instead oscillating as they propagate through the vacuum between the Sun and Earth. Higher energy neutrinos (> 5 MeV) will initially not oscillate, but as the value of n_e drops with solar radius, the conditions will briefly be met for resonance and maximal oscillation. These neutrinos will consequentially undergo matter-enhanced $\nu_\mu \rightarrow \nu_e$ oscillations before they leave the Sun. This transition between low energy vacuum oscillations and higher energy MSW oscillations, and the resulting characteristic energy dependence of oscillations, was the final detail needed to explain the experimental results seen in different solar neutrino experiments.

2.3.5 Atmospheric Neutrino Oscillations

Cosmic rays interacting in the earth's upper atmosphere produce large showers of particles. Atmospheric neutrinos are produced in several decay chains in these showers. Pions decay to form muon neutrinos and muons. These muons subsequently decay and produce electron neutrinos. Contributions to this neutrino flux at higher energies can come from kaons as well. These atmospheric neutrinos are produced everywhere in the earth's atmosphere and can easily travel through the Earth, so that a detector positioned on the Earth's surface will receive neutrinos arriving from all directions. These neutrinos will have traveled anywhere between 15 km and 13,000 km and are more energetic than solar neutrinos, with an average energy on the order of 1 GeV. In models of atmospheric neutrino production, there is a large ($\sim$ 20%) uncertainty on the overall neutrino flux. However, the uncertainty on the relative ν_μ and ν_e flux is small (<2%), with an expected ν_μ to ν_e ratio of 2 to 1 [36].

In 1988, the Kamiokande experiment [37] attempted to make a measurement of the relative fluxes of ν_μ and ν_e. The electron neutrino flux was found to match experimental predictions. However, the muon neutrino flux was $\sim$60% lower than expected. This deficit was soon observed by other experiments (IMB [38], MACRO [39], and Soudan-2 [40]) and became known as the "Atmospheric Neutrino Anomaly." In 1998, the upgraded Super-Kamiokande experiment confirmed that the deficit was caused by neutrino oscillation [41]. Super-Kamiokande's water Cerenkov detector allowed for the reconstruction of the direction from which the incoming neutrino had entered. Interaction rates for ν_μ and ν_e events could therefore be plotted as a function of zenith angle, as shown in Figure 2.6. The ν_e flux agreed well with predictions, but the ν_μ flux deficit was found to be heavily dependent on the zenith angle - essentially, on the distance that the neutrino had traveled since its production. This observed ν_μ disappearance was consistent with the oscillation of $\nu_\mu \rightarrow \nu_\tau$.

The MINOS experiment, the focus of this thesis, has both confirmed these results and tightened the limits on the U_{23} atmospheric mixing parameters. The MINOS experiment uses an accelerator-produced neutrino beam to study neutrino oscillations at the Δm^2_{32} scale. In this case, L is fixed, at 735 km, with an average energy E of 3 GeV (in a wideband spectrum). MINOS is designed to detect mainly ν_μ Charged Current interactions. The experiment measures $\sin^2(2\theta_{23})$ and Δm^2_{32} by looking for muon neutrino disappearance. The two-neutrino approximation for ν_μ survival can be written as

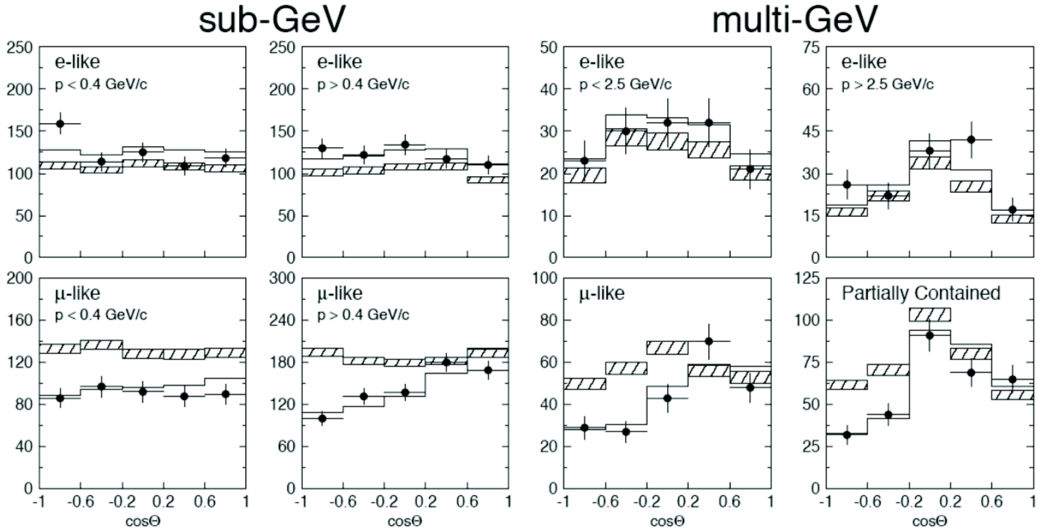

Figure 2.6: Results for the Super-Kamiokande atmospheric neutrino oscillation study [37]. Number of events observed (black data markers) versus number of events predicted in the absence of oscillations (hatched region). The solid black line gives the best fit (assuming ν_μ disappearance oscillations). The events are plotted in groups of ν_e-like and ν_μ-like events, in different energy regimes, as a function of zenith angle.

follows:

$$P(\nu_\mu \to \nu_\mu) \approx 1 - \sin^2(2\theta_{23}) \sin^2\left(\frac{1.27\Delta m^2_{32} L}{E}\right) \tag{2.39}$$

Oscillation is measured by comparing the unoscillated ν_μ CC rate at a Near Detector (at Fermilab) to the oscillated rate at a Far Detector (in a mine in Soudan, Minnesota). The MINOS results [42], shown in Figure 2.7, are consistent with a hypothesis of $\nu_\mu \to \nu_\tau$ oscillation. The results have also placed the tightest limits to date on Δm^2_{32} and θ_{23}. The oscillation appears to be maximal, with $\sin^2(2\theta_{23})$ consistent with 1.0 (i.e., $\theta_{23} = 45°$):

$$\Delta m^2_{32} = (2.32^{+0.12}_{-0.08}) \times 10^{-3}\ \mathrm{eV}^2, \ \ \sin^2(2\theta_{23}) > 0.90 \ \ (90\%\ \mathrm{CL}). \tag{2.40}$$

2.4 The experimental search for θ_{13}

As described above, there is now ample evidence for the oscillation of neutrinos. Most of the focus on the U_{12} and U_{23} matrices is now on making precision measurements of the parameters involved. In the case of the U_{13} matrix, and θ_{13}, much work remains

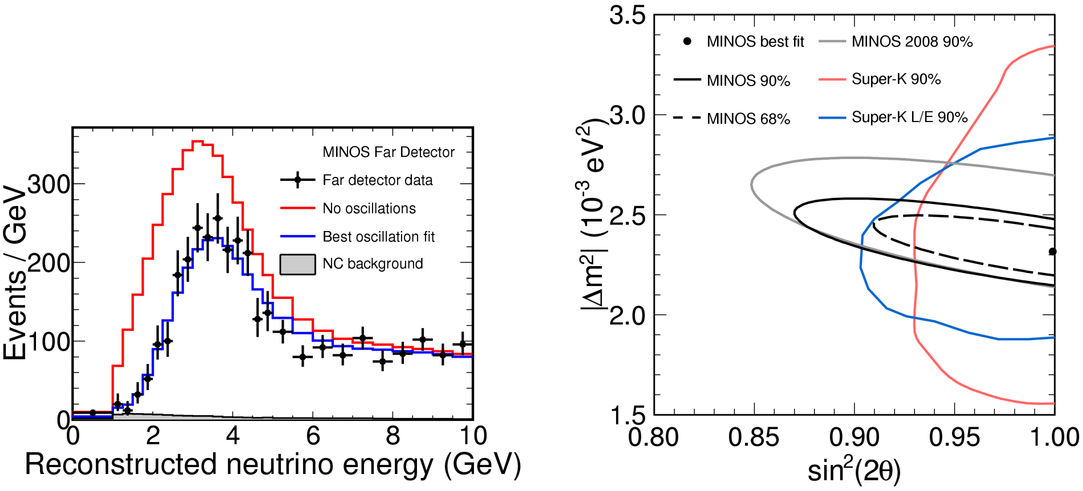

Figure 2.7: The latest (2010) results from the MINOS experiment for ν_μ disappearance oscillations [42]. The left hand plot shows the unoscillated prediction versus data, with a best fit assuming $\nu_\mu \rightarrow \nu_\tau$ oscillations. The right hand plot shows the resulting limits on the oscillation parameter values Δm^2_{32} and θ_{23}, with constraints given from Super-Kamiokande as well.

to be done[4]. Recent results from the T2K [145] experiment (described in Appendix A) have hinted at a non-zero value for this parameter. However, most measurements to date have yielded only upper limits, and it is clear that θ_{13} will be small compared to the other two angles, even if it is non-zero. However, even a small non-zero θ_{13} can potentially introduce CP violation via the term δ in Equation 2.26. Measuring θ_{13} has therefore become one of the most important goals in neutrino physics today.

2.4.1 Current Knowledge: Reactor Experiments

The best previous limits on the value of θ_{13} have come from the CHOOZ reactor experiment [43] in France. CHOOZ attempted to measure θ_{13} by looking for the disappearance of reactor electron antineutrinos. The two-neutrino approximation for $\overline{\nu}_e$ oscillation survival is written as follows:

$$P(\overline{\nu}_e \rightarrow \overline{\nu}_e) \approx 1 - \sin^2(2\theta_{13}) \sin^2 \left(\frac{1.27 \Delta m^2_{32} L}{E} \right) \tag{2.41}$$

---

[4] The third mass difference term, Δm^2_{31}, is approximately equal to Δm^2_{32}

The CHOOZ detector was located 1 km from a nuclear reactor which provided O(MeV) electron antineutrinos. The detector volume itself was liquid scintillator, doped with Gadolinium. The reactor $\overline{\nu}_e$ interacted with the protons in the detector via inverse beta decay: $\overline{\nu}_e + p \rightarrow e^+ + n$. PMTs detected photons from positron annihilation, followed by a delayed photon signal from neutron capture on the Gadolinium. CHOOZ found a $\overline{\nu}_e$ rate which was consistent with the unoscillated prediction. Nonetheless, the experiment placed an upper limit[5] [15] on the value of θ_{13}.

$$\sin^2(2\theta_{13}) < 0.16 \;\; (90\% \text{ CL}) \tag{2.42}$$

Until new MINOS and T2K results appeared in 2011, this was the best limit on θ_{13}. It will therefore often be quoted as such throughout this thesis.

In addition to electron antineutrino disappearance, θ_{13} can also be measured via a second oscillation mode: $\nu_\mu \rightarrow \nu_e$. The approximate two-neutrino probability for this oscillation is:

$$P(\nu_\mu \;\rightarrow \nu_e) \approx \sin^2(2\theta_{13}) \sin^2(\theta_{23}) \sin^2\left(\frac{1.27\Delta m^2_{32} L}{E}\right) \tag{2.43}$$

MINOS, T2K, and other accelerator experiments can therefore attempt to measure θ_{13} by looking for the appearance of electron neutrinos after some distance of oscillation. The search for this mode of oscillation in MINOS will be the focus of this thesis.

2.5 Summary

Neutrinos occupy a somewhat unusual place in the Standard Model. Chargeless, colourless, and only interacting via the weak force, they are also very difficult to observe. Until recently, neutrinos were also presumed massless. Oscillation, in addition to explaining observed deficits in solar and atmospheric neutrino flux, provides evidence that they do indeed have mass. The as-yet-unmeasured mixing angle θ_{13}, if it is non-zero, may also provide an avenue for introducing CP violation into the lepton sector. This thesis will describe an attempt to measure θ_{13}, by searching for the oscillation of $\nu_\mu \rightarrow \nu_e$ in the MINOS experiment.

---

[5] Note: this value is after including the current MINOS best limit [42] on Δm^2_{32}, and recent reevaluations of reactor flux [44] [45].

Chapter 3

The MINOS Experiment

The Main Injector Neutrino Oscillation Search (MINOS) experiment was built with the goal of studying neutrino oscillations at the atmospheric Δm^2_{32} scale. A source of mainly muon neutrinos is provided by the NuMI (Neutrinos at the Main Injector) beamline at Fermilab. Two MINOS detectors observe the neutrinos from this beamline. A first Near Detector (ND), at Fermilab, measures neutrinos prior to oscillation. The Far Detector (FD), located at an on-axis beamline distance of 735 km, then measures the neutrinos after they have traveled and potentially oscillated. The detectors themselves are magnetized tracking calorimeters which have been optimized for the study of muon neutrino Charged Current (CC) events. Using these two operationally equivalent detectors permits a much more precise understanding of the behavior of the beamline, as systematic effects such as flux mismodeling and cross-section uncertainties largely cancel out in any physics analysis. A detailed description of the MINOS detectors and their construction can be found in [52], and of the NuMI beamline in [53].

3.1 The NuMI beamline

3.1.1 Producing a Neutrino Beam

The NuMI beamline is part of the larger Fermilab Accelerator complex. A map of the components of this beamline is shown in Figure 3.1. The beam begins when the Booster sends protons to the Main Injector, where the protons are accelerated from 8 GeV to 120 GeV. The protons are accelerated in six batches per cycle: five or six are sent to the NuMI Target Hall, with any remaining batches going to the Tevatron. During the first year of running, the MI typically provided 2.1×10^{13} protons-on-target (POT) every 2.2−2.4 seconds. This intensity has since been improved to 2.5×10^{13} POT.

In the Target Hall, the protons from the Main Injector collide with a target to produce a shower of hadrons, most of which are pions and kaons. The target itself is a 6.4 mm $\times$ 15 mm$\times$ 940 mm rectangular rod made of graphite. It consists of 47 longitudinal segments called fins and is protected upstream by a collimating baffle. The particles produced in the target collision are focused by two sets of magnetic horns (Figure 3.2) [54]. The two magnetic horns are separated by 10 m and each contain a parabolic inner conductor, to which a pulsed current is provided in time with beam spills. This current produces a magnetic field, which focuses or separates particles based on their charge and momentum to form a parallel beamline. The greater the momentum of a charged particle, the longer the focal length of the focusing. The toroidal field has been measured to fall off in a manner consistent with 1/r to within 1%. The current in the horn is known to within $\pm 0.5\%$, with less than a 0.2% variation during data taking [53]. NuMI is designed to focus pions, which produce neutrinos in the desired oscillation energy range. Normally, the beam is run with the horn current in "forward" mode, to focus π^+ and K^+ and produce a neutrino beam. The beam can also be run with the horn current in a "reverse" mode to focus π^- and K^- and produce an antineutrino beam. Only neutrino data is used in this thesis.

After the pions and kaons are focused, they pass into a 675 m long by 2 m wide decay pipe, which is sealed with two aluminium windows. During the first two years of running, the interior of the tube was a 0.5 Torr vacuum. Helium gas was later pumped into the tube for safety reasons. The length of the tube was chosen to be of the order of the decay length of a 10 GeV pion. As both the pions and kaons pass through this tube, they decay to produce neutrinos and other secondary particles. The beam is mainly composed of muon neutrinos from the decay $\pi^+ \rightarrow \mu^+ + \nu_\mu$. At the end of this decay pipe, the remaining hadrons are screened out with an absorber of cooled aluminium, followed by a layer of steel and concrete and then 250 m of rock. The beam is angled downwards at $-3°$ with respect to horizontal at the Near Detector. It then passes through the Earth's surface, and enters the Far Detector at an angle of $+3°$.

The beam intensity and quality are monitored in several different locations. The intensity of the initial proton beam is determined to within $\pm 1\%$ using a series of toroidal beam current transformers, while the size and position of the beam are also determined to be stable to within ± 50 μm by a series of beam position monitors [55]. Downstream from the decay pipe, three muon monitors in the rock measure the alignment of the NuMI beamline. From these monitors, the angle of the beam is known to within 20 μrad. The intensity of the neutrino beam is relatively stable, with charge per spill per protons on target (POT) varying by $\pm 2\%$ in the first two years of running [53].

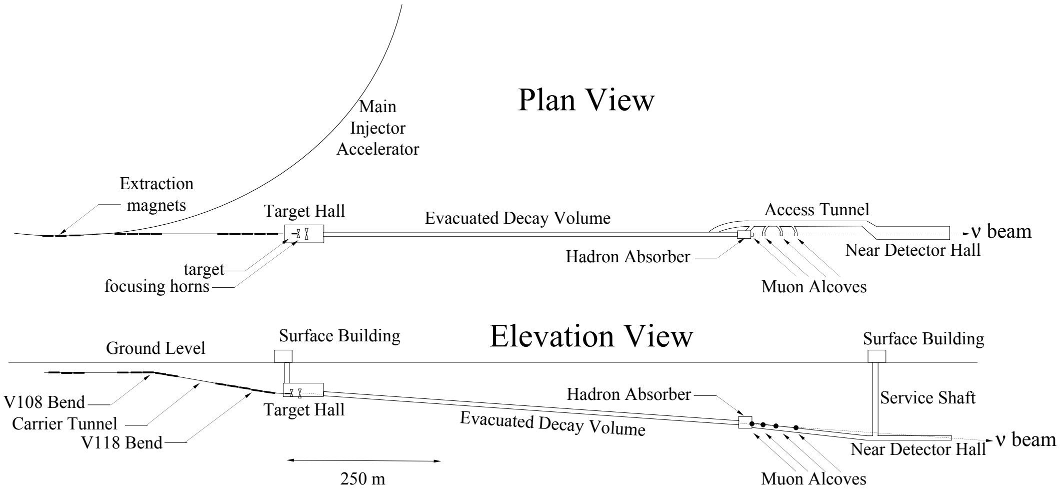

Figure 3.1: Plan of the components of the NuMI beamline from a top view (top) and from a side view (bottom). Image from [53].

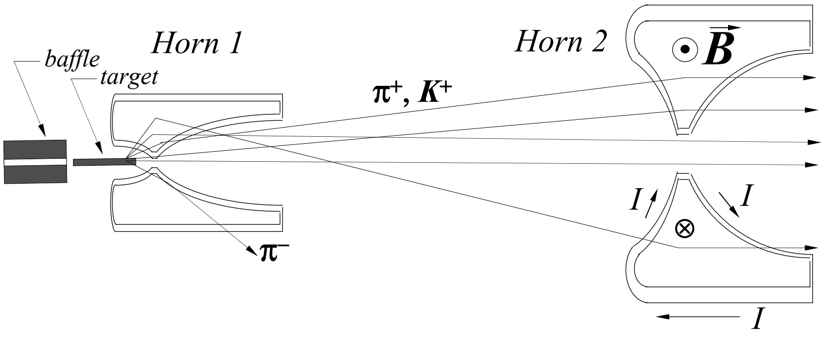

Figure 3.2: Schematic of the two NuMI magnetic focusing horns. The configuration is shown in standard mode, with the target and baffle upstream of the two horns (which are separated by 10 m). The vertical scale here is shown to four times that of the horizontal scale. Image from [53].

3.1.2 Characterizing the NuMI Beamline

The components of the NuMI beamline were designed to produce a beam consisting of nearly all muon neutrinos, in an energy range favorable for oscillations. NuMI is a wideband beam, with a large neutrino flux and and a relatively broad spectrum of energy. The energy range can be adjusted by altering both the position of the target and the current in the magnetic focusing horns. The target is mounted on a 2.5 m long rail-drive system. When the target is at its furthest position downstream (position = 250 cm), it is 65 cm inside the first horn. The maximum possible horn current is 200 kA, which produces a 3 T toroidal field. For most data-taking, the beam is run in Low Energy mode (LE10/185kA). In this configuration, the target is placed 10 cm from its default position inside the first focusing horn. The magnetic focusing current is run at 185 kA. This produces a beam with a peak in intensity at 3.3 GeV (with an RMS spread of 1.1 GeV). While the beam can be run in other modes, this configuration is the main setting of NuMI, with an energy optimized for the detection of atmospheric-scale muon neutrino oscillations. A plot of the shape and relative flux of several NuMI configurations can be seen in Figure 3.3. The total protons on target (POT) exposure of the beam to date (2005-2011) for the various modes is presented in Figure 3.4.

The branching ratios for π^+ and K^+ decay are very favorable for the production of muon neutrinos, with the decay $\pi^+ \rightarrow \mu^+ + \nu_\mu$ occurring 99.99% of the time [15]. The vast majority of the neutrinos in the beam originate from pion decay. In the Low Energy mode, 87% of muon neutrinos were produced by the decay of π^+s, with the remaining 13% coming from K^+s. There is also a small contamination of electron neutrino events in the beamline. Because of the large fraction of pions produced in the target, most of these electron neutrinos are produced through secondary muon decay, with a smaller fraction from K^+ decay. The decay of K_Ls also provides a tiny remaining fraction of both the ν_μ and ν_e flux, as well as a small anti-neutrino component. Some negative particles are still focused by the horn, leading to additional antineutrinos. For the standard Low Energy mode, the beam contains 92.9% ν_μ, 5.8% $\overline{\nu}_\mu$ and 1.3% $\nu_e + \overline{\nu}_e$.

The Far Detector data analyzed in this document were collected in a series of runs between May 2005 and July 2010. The total exposure for standard neutrino beam running during this time was 8.2×10^{20} POT. For analysis purposes, this data is often split up into three discrete chunks: Run 1 (1.2×10^{20} POT), Run 2 (1.9×10^{20} POT), and Run 3 (actually runs 3, 4, 5, and 6, for a total 5.0×10^{20} POT). These divisions were made based on the beam conditions during each time period. For example, helium was added to the decay pipe for Run 3 and later.

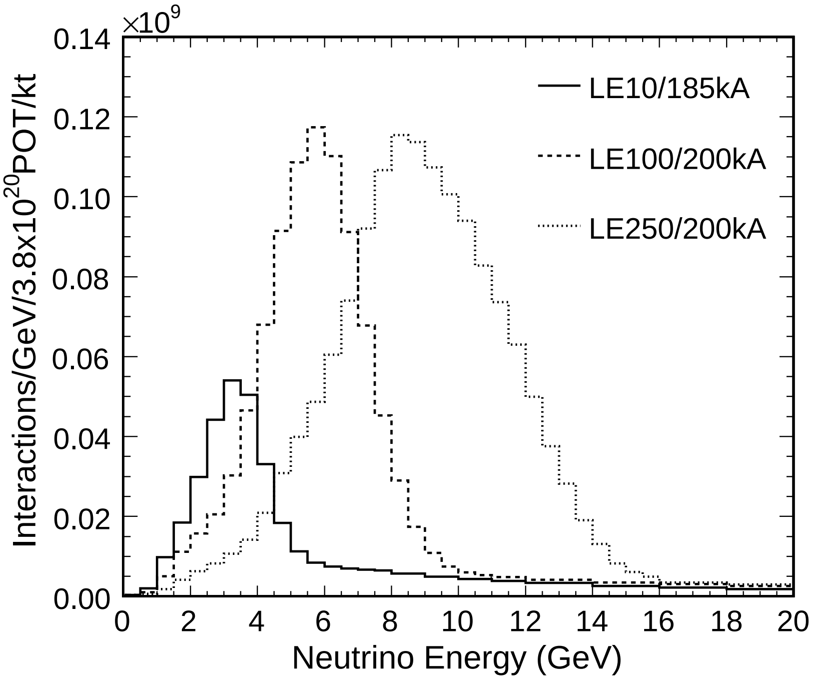

Figure 3.3: Projected rate of Near Detector ν_μ Charged Current interaction resulting from three separate NuMI configurations (from [53]): low energy (LE10/185kA), high energy (LE250/200kA), and medium energy mode (LE100/200kA). These rates are calculated from Monte Carlo.

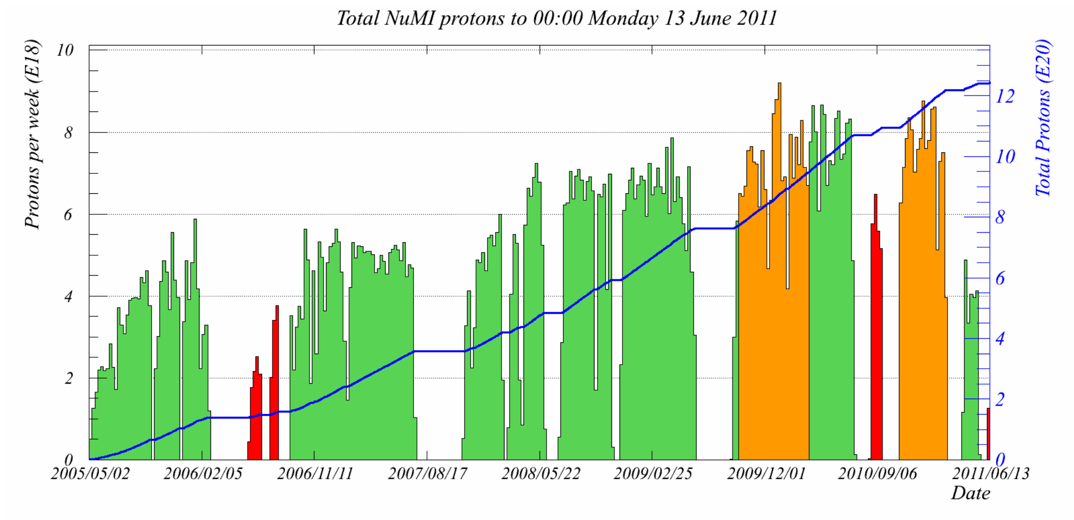

Figure 3.4: Total exposure of the NuMI beamline from May 2005 to June 2011. The green area indicates beam run in standard LE010/185kA neutrino mode. The orange indicates beam run in LE010/185kA reverse horn current antineutrino mode. The red indicates alternate configurations, such as the high energy (LE250/200kA) mode. The blue line shows the total integrated exposure. Image from [56].

3.2 The MINOS Detectors

The Near and Far MINOS detectors are as similar as possible, to allow for an accurate comparison of their observed neutrino spectra. Both detectors have been primarily designed for measurements of muon neutrino Charged Current (ν_μ CC) events. A description of the two detectors, and their basic operation, follows.

3.2.1 Basic components: Steel, Scintillator System, and Magnetic Fields

Both MINOS Detectors consist of alternating planes of steel and solid scintillator. The steel planes form the target mass (95% of total mass) of the experiment. Neutrinos from the NuMI beamline interact with nuclei in these planes to produce secondary particle tracks and showers. In both detectors, the steel planes are 2.54 cm thick. The steel itself is low-carbon, low-radioactivity steel.

Between each steel plane there is a second layer constructed from strips of extruded polystyrene scintillator. These scintillator planes actively detect the particles produced by these neutrino interactions and allow for calorimetry and event reconstruction. These

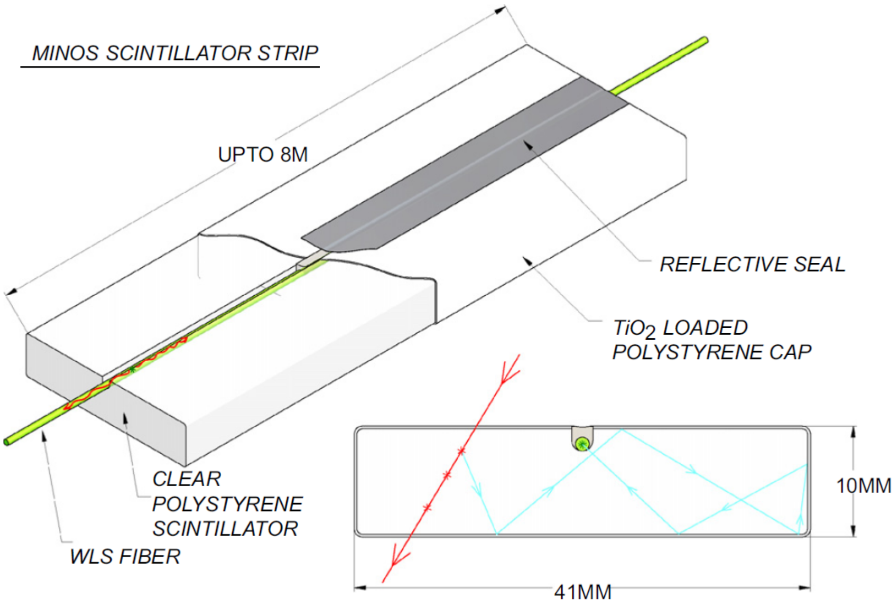

Figure 3.5: Schematic of a scintillator strip and wavelength shifting fibre in the MINOS detector. The top part of the image shows the basic setup of the individual strips, with the WLS fibre inserted into a groove on the top surface of the scintillator strip, and sealed in with a reflective seal. The bottom image shows how light is collected. Energy from an incident ionizing particle is absorbed and re-emitted as light within the scintillator strip. This light is then repeatedly reflected off the side of the strip and absorbed by the WLS, which then carries it to the end of the strip, to the readout system. Image from [52].

strips are 4.1 cm wide, 1.0 cm thick, and have a maximum length of 8.0 m. The scintillator strips themselves are covered in a reflective titanium dioxide layer which both aids light yield and prevents damage.

A normally incident minimum ionizing particle (MIP) passing through the strip will produce scintillator light. This is collected using wavelength shifting (WLS) fibres, which are polystyrene doped with Y11, with an outer cladding of acrylic and polyfluor. These fibres absorb light at a peak of 420 nm and then re-emit it at 470 nm, to prevent self-absorption. Each scintillator strip contains a single WLS fibre, which sits on the wide surface at the top of each strip. A cutaway drawing of the WLS/scintillator system is shown in Figure 3.5.

In both detectors, the strips and WLS fibres are grouped into independent modules, consisting of a layer of parallel scintillator strips which are glued together and held in place between two layers of aluminium. This aluminium provides both structural support and light shielding. At the end of each module, the WLS are bundled together into a bulk optical cable and connected to clear plastic cables, which transport the light

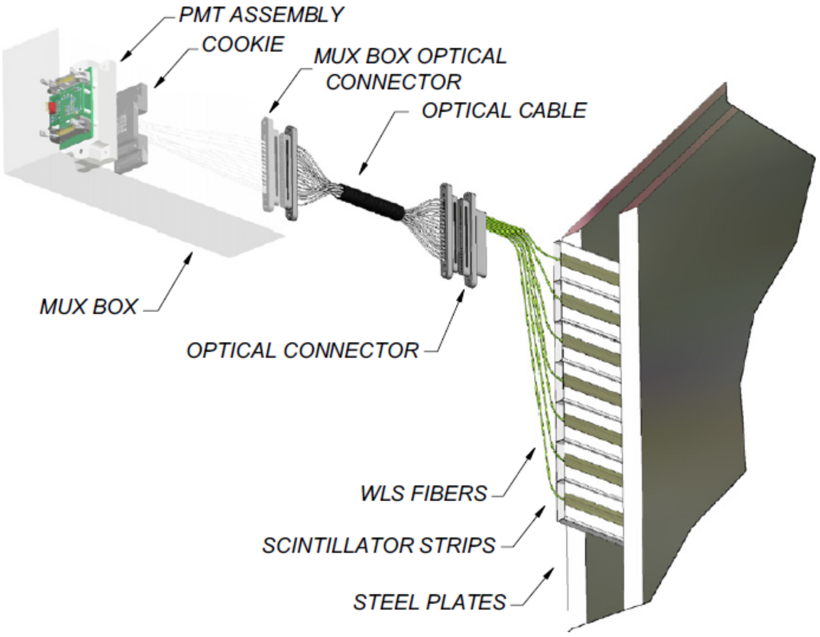

Figure 3.6: Diagram showing the readout system by which a light signal in the MINOS detector is recorded. The original signal is picked up in the scintillator strips in the plane shown on the far right. This light is then carried by the wavelength shifting fibres out of the module to a clear plastic cable. The clear plastic cable carries the signal to the PMT enclosure and the surface of the PMT itself. Image from [52].

to photomultiplier tubes (PMTs). In each detector, the PMTs are housed in light-tight boxes which shield them from light, magnetic fields, and electronics noise. The Near Detector employs 64-anode (M64) Hamamamtsu PMTs, with a single PMT in each enclosure. The Far Detector employs 16 anode (M16) Hamamamtsu PMTs, with three to an enclosure. The PMT signal is digitized, and the hits are recorded by the data acquisition system in units of ADCs. A schematic of the full scintillator and readout system can be seen in Figure 3.6.

From these readouts, one can construct a full three dimensional picture of an event in MINOS. The plane number gives information about the event in the longitudinal z direction, parallel to the beamline. The readout of the individual scintillator strips gives information about the event in the transverse direction. All the strips in a single plane are oriented in the same direction. The strip orientations in alternating planes, however, are orthogonal to one another, at $+45°$ (the u direction) and $-45°$ (the v direction). This provides a stereoscopic picture of the event as it travels through the detector and allows for a full three-dimensional reconstruction. The u and v coordinates can then be translated back into x (horizontal) and y (vertical) coordinates.

Both MINOS detectors are magnetized. A toroidal magnetic field is produced by the current in a magnetic coil threaded through a "coil hole" in the center of each detector module. The magnetic field curves the path of muon tracks from ν_μ CC interactions, both better containing the tracks inside the detector, and allowing for the determination of muon momentum and charge sign from track curvature. The strength of the magnetic fields is similar between the two detectors, with an average field of 1.28 T in the Near Detector, and 1.42 T in the Far Detector.

3.2.2 The Near Detector

Interactions in the Near Detector are used to characterize the neutrino flux before oscillation has occurred. The Near Detector sits 1040 m from the NuMI target, in a cavern on the Fermilab site at a depth of 100 m, with 225 meters water equivalent (mwe) of rock on top. A schematic of the Near Detector can be seen in Figure 3.7. Due to its proximity to the start of the beamline, there is a high rate of neutrino interaction in the Near Detector. During each 10 μs beam spill, an average of 16 neutrino interactions typically occur in the detector. Data acquisition is triggered in both a "spill-gate mode" and a "dynode mode." The spill-gate mode is used for standard data taking. The beam timing triggers a 13 μs window during which hits are continuously digitized, beginning 1.5 μs before the delivery of the beam spill. The dynode mode triggers digitization on a PMT activation threshold, and is used for taking out-of-spill cosmic ray data. Because of the high event rate at the Near Detector, this trigger is designed to operate without dead time.

Due to this high neutrino flux, the Near Detector requires a smaller fiducial volume than the Far Detector. The Near Detector contains 282 planes and has a total mass of 980 metric tonnes. The steel planes have a 6.2 m wide by 3.8 m high "squashed octagon" shape. The magnetic coil is offset from the center of the plane to create a 2 meter diameter fiducial area through which the beam passes. The Near Detector is separated into both a Calorimeter and a Spectrometer region. The first 120 upstream planes of the Near Detector comprise the Calorimeter region, designed to make precise measurements of neutrino interaction topology and energy. Planes in this region are all active (i.e., are all read out by PMTs). This region contains two types of plane: partially instrumented planes, which only span the fiducial region, and fully instrumented planes which also span the coil hole. Plans of these two plane types are shown in Figure 3.8. The Calorimeter planes are arranged in a 10-plane pattern, with each fully instrumented plane followed by

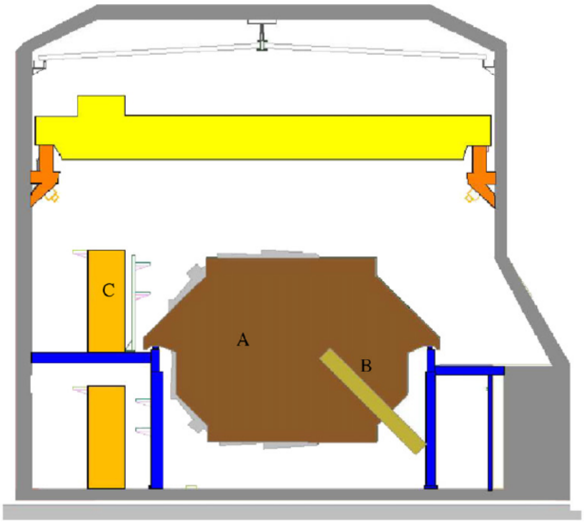

Figure 3.7: Schematic of the Near Detector setup. A shows the (front) steel plane of the Near Detector. B shows the magnetic coil, leading into the detector. C shows an electronics rack on an elevated walkway. Images from [52].

four partially instrumented planes. During each beam spill, approximately half of the ND neutrino events will occur in the Calorimeter region and will therefore be reconstructed. The remaining Near Detector planes 121-282 comprise the Spectrometer region, which is responsible for tracking downstream muons. In this section, only every fifth plane is instrumented, with alternating u and v fully-instrumented views, and no partially instrumented planes in between.

The Near Detector scintillator system modules differ from the Far in that only one side of the scintillator strip is read out. This results in a lower overall light level in the detector (i.e., a smaller amount of light is detected from the same deposited charge). Strips in the upstream Calorimeter region are all read out individually. Strips in the downstream Spectrometer region are read out with anode pads linked in parallel in groups of four, to reduce the number of readout electronics (the four strips are chosen to be approximately a meter apart, to allow for easy discrimination). The PMTs are each operated at a voltage of approximately 800 V to reach a required minimum average gain of 0.8×10^6.

3.2.3 The Far Detector

The MINOS Far Detector measures the rate of neutrino interaction after the neutrinos have propagated and hopefully oscillated. A location was chosen 753.3 km northwest of the NuMI target, in an iron mine in Soudan, Minnesota. The detector is located in a

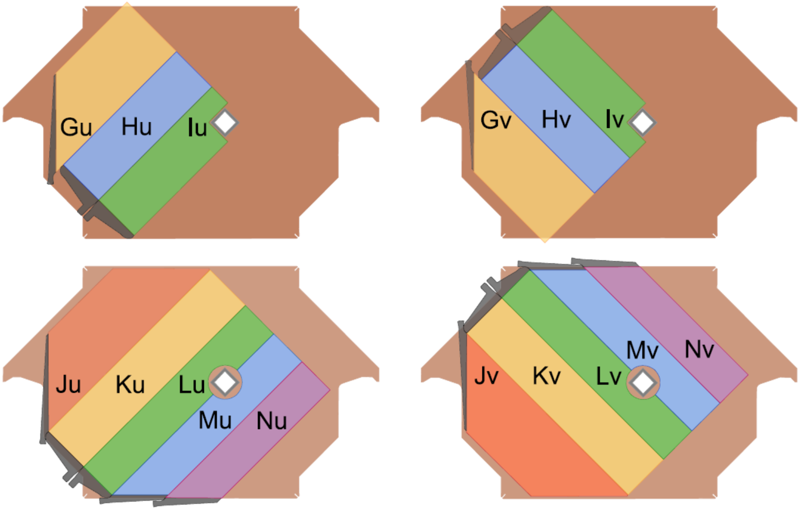

Figure 3.8: Configuration of the Near Detector scintillator strip modules. The views on the left are oriented in the u direction, and the views on the right in the v direction. The top two images show two "partially instrumented" planes. The bottom two images show two "fully instrumented" planes. The beam is aimed in the center of the region to the left of the coil hole, which has full coverage by scintillator modules. Image from [52].

cavern 705 m underground (with a rock overburden of 2070 mwe). A schematic of the detector is shown in Figure 3.9. The Far Detector is significantly larger than the Near Detector, with a total mass of 5400 metric tonnes. The planes are octagonal, with a diameter of 8 meters. The detector was constructed in two separate "supermodules," separated by a 1.15 m gap. The first supermodule is 14.78 m long with 249 planes, and the second supermodule is 14.10 m long with 237 planes. Each of the supermodules has a separate magnet coil, which is threaded directly through the center of the plane. Unlike the Near Detector, the Far Detector scintillator strips fully cover the surface of each plane, with 192 scintillator strips to a plane. The layout of the modules is shown in Figure 3.10. As in the Near Detector, the strip orientation is rotated 90° from plane to plane.

The underground location of the Far Detector was chosen to block most surface radiation, including the cosmic ray muons which still dominate detector activity at an interaction rate of 0.8 Hz. Due to the Far Detector's distance from the target, the rate of beam neutrino interactions/kg is a factor of 10^6 smaller than in the Near Detector (several neutrino interactions per day, rather than several per spill).

The Far Detector strips are read out from both ends. The signal is routed from the end of the module to 484 light-tight boxes which each contain three PMTs. Eight clear

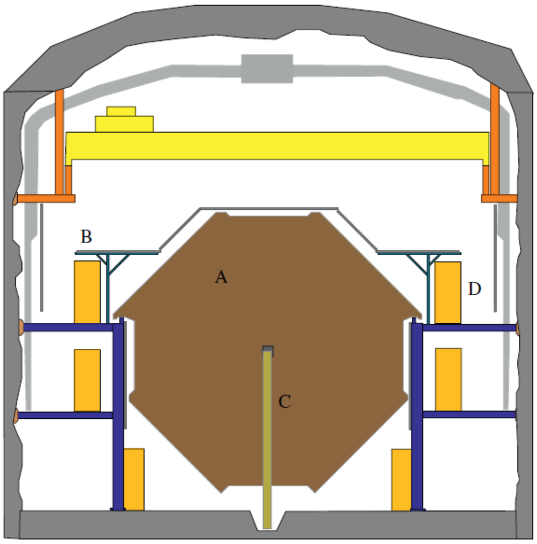

Figure 3.9: Schematic of the Far Detector setup. A shows the (front) steel plane of the Far Detector. B shows the muon veto shield. C shows the magnetic coil, leading into the detector. D shows an electronics rack on an elevated walkway. Images from [52].

fibres are attached to a single PMT pixel, where the signal is optically summed. The placement of these fibres is such that the readout strips are separated by at least 1 m. A "demultiplexing" algorithm [57] is used to reconstruct the original pattern of strip hits in the detector. This pattern helps to reduce the effects of PMT crosstalk in the reconstruction.

The recording of FD Data is triggered by both spill time and PMT activation. A GPS timestamp is sent from the Near Detector the Far Detector to indicate that a beam spill has occurred. Data are then recorded from a several second buffer in a 100 μs window. In order to reduce electronics dead time, hits are only digitized if two or more PMTs in the same vicinity reach the 0.25 PE digitization threshold at the same time.

3.2.4 Calibration Detector

A Calibration detector [60] (often referred to as the CalDet) was also built for MINOS. This detector was exposed to a test beam of muons, hadrons, electrons, at the CERN PS facility. The calibration detector was smaller than both the Near and Far Detectors, with a total of 60 1 m × 1 m planes of steel and scintillator, and no magnetic field. The detector read out events using electronics from both the Near and Far detectors. The data from this detector are used to determine the absolute calibration of the MINOS

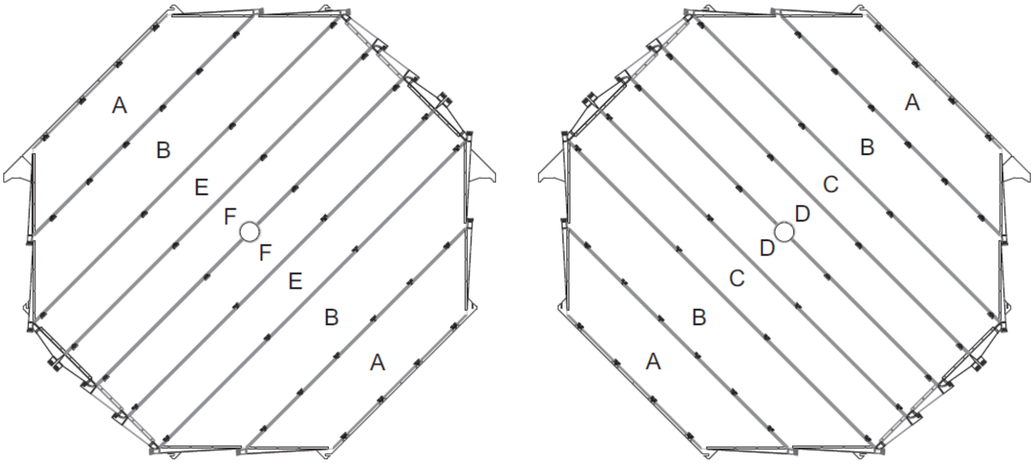

Figure 3.10: Diagram of the configuration of the separate Far Detector scintillator strip modules. The left image is oriented in the "U" direction, and the right image in the "V" direction. WLS fibres are read out by the PMT MUX box at the end of each module. Image from [52].

detector response.

3.3 Reconstruction

To identify neutrino interactions in MINOS, one must first glean basic information about event topology and energy. This is the purpose of the MINOS event reconstruction algorithm. Prior to reconstruction, event information is stored as digitized strip hits (ADCs) measured in individual electronics channels. These data include information about strip location, pulse height, and hit timing.

The cluster of data found in a beam spill is known as a snarl. In the Far Detector, where the neutrino flux is much lower, a snarl contains one or zero neutrino events. In the Near Detector, a single snarl will contain multiple events. The first task of the reconstruction is to break down a snarl into different candidate events, called slices. This is done by grouping the snarl hits both by physical location and by the timing of their energy deposition. Once the hits have been grouped into potential events, the reconstruction attempts to find particle tracks and showers in each interaction. ν_μ Charged Current interactions are characterized by a long particle track, produced by a muon. At the neutrino interaction vertex, the ν_μ CC event will usually also have a cluster of hits from a hadronic shower. Neutral Current and ν_e CC interactions typically produce only showers, although these events can sometimes contain smaller particle

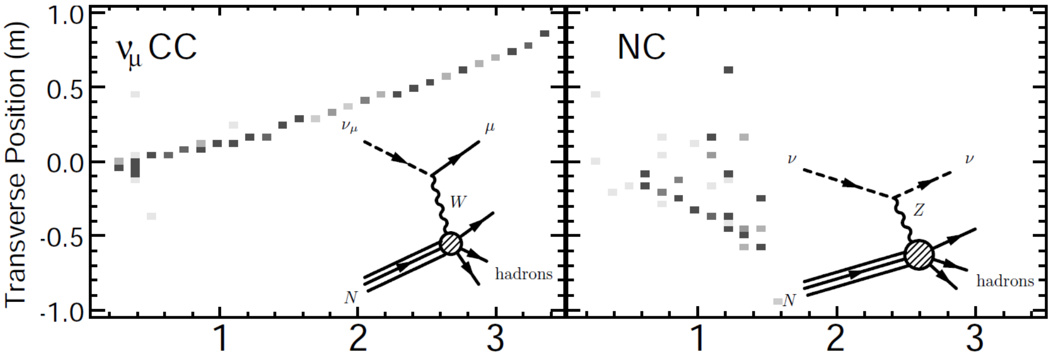

Figure 3.11: Side-view of two types of event in the MINOS Detector: ν_μ charged current (left) and Neutral Current (right). Darker pixels indicate a larger pulse height in a hit. The x axis is parallel to the beam line, and the y axis is transverse position. Image from [53].

tracks. Figure 3.11 shows examples of these event topologies.

A pattern recognition algorithm is applied to identify hits which make up the segments of a track. These segments are then chained together into larger track, which is fitted using a Kalman Filter algorithm. The momentum of the resulting track is determined in two separate ways. For muons which stop in the detector, track range is used to determine the momentum. If the muon track exits the detector, the curvature of the muon in the detector's magnetic field is used to determine the momentum. The range method has a 5% resolution for momentum, and the curvature method a 10% resolution [53]. The neutrino interaction vertex is identified from the earliest plane of the track.

After any tracks have been identified, the remaining hits are sorted into showers, based on their relative positions in space and time. The shower energy is determined based on the summed pulse-height of the hits (determined from the calibration, which will be described in the next section). Some strip hits will contain both a shower and a track. In this case, the estimated pulse height of the track is subtracted. In the longitudinal direction, the shower vertex is identified as the first plane in the series of contiguous planes in an event. In the transverse direction, the shower vertex is calculated from the energy-weighted mean position of the event strips.

3.4 Calibration

Searching for ν_e appearance requires a precise understanding of shower energy in MINOS. The calibration determines this shower energy by first calculating the detector

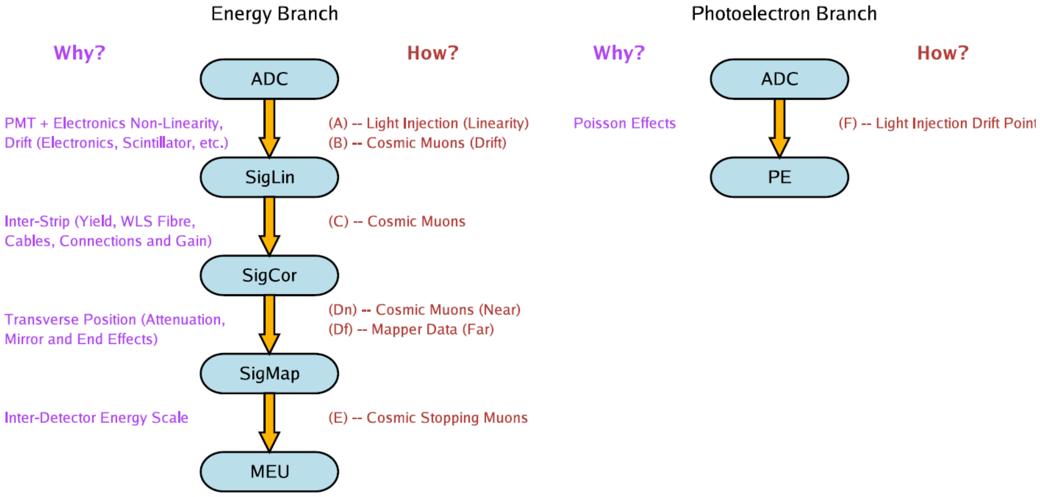

Figure 3.12: Flow chart of the two separate branches of the detector calibration, with (left to right) the rationale, the variable, and the method. The energy deposition branch is shown on the left, and the photoelectron production branch on the right. Image from [61].

response and then performing a relative and absolute calibration to reconstruct the original deposited energy. For the detector response, two separate chains of calibration are employed in MINOS: an “Energy” chain which converts the raw digitized signal (in ADCs) into real deposited energy, and a “Photoelectron” chain which turns the signal into photoelectron units. The individual steps in each calibration chain are shown in Figure 3.12, along with the associated MINOS variables.

The Energy Chain is the primary chain used for the calibration of MINOS. This chain starts with a raw ADC pulse height which can be described as a quantity $Q_{raw}(i, t, x, s, d)$, dependent on readout channel i, interaction time t, strip s, interaction position x, and detector d (ND or FD). The goal of the calibration is to turn $Q_{raw}(i, t, x, d)$ into Q_{corr}, the actual energy deposited in the detector, via several stages of corrections:

$$Q_{corr} = Q_{raw}(i, t, x, s, d) \times D(d, t) \times L(d, s, Q_{raw}) \times S(d, s, t) \times A(d, s, x) \times M(d) \tag{3.1}$$

The first three steps, the drift correction $D(d, t)$, linearity correction $L(d, s, Q_{raw})$, and strip-to-strip correction $S(d, s, t)$ are carried out prior to event reconstruction. The attenuation correction $A(d, s, x)$ is carried out following the reconstruction. Finally, cosmic ray muons which stop in the detector are used to set the inter-detector relative energy scale, and data from the Calibration Detector are used to normalize the detector response. This relative and absolute calibration are step $M(d)$.

Several sets of tools are employed in the calibration in Equation 3.1. First, all three detectors - Near, Far, and Calibration - contain a Light Injection (LI) system [63] [64] to assess the response of the detector readout. This system consists of a series of pulser boxes, each containing 20 or 40 UV LEDs which send a light signal to the WLS at the edge of the detector. This light imitates a signal produced in the scintillator, but has a known intensity which is monitored by PIN photodiodes. This system can be both used to monitor the response of the detector channels over time and to characterize the linearity of the response.

Second, through-going cosmic ray muons provide a good way of monitoring the response of the scintillator system. These muons are taken from out-of-spill cosmic rays. In the Near Detector, these muons have a typical average energy of 55 GeV with a rate of approximately 10 Hz. In the Far Detector, they have a mean energy of 200 GeV with a rate of approximately 0.8 Hz. A small portion of cosmic ray muons also stop in the Near and Far Detectors. This "stopping muon" sample will be used to calibrate the inter-detector energy scale.

Third, data from a Module Mapper are also used throughout the calibration. During the construction of the detectors, a 5 mCi ^{137}Cs γ-ray source enclosed in a lead pig was precision-guided along the length of each scintillator module to create a standard ionizing signal. This ionizing signal was then measured with optical fiber cables and used to map the scintillator response throughout the module.

3.4.1 Drift Correction

The first stage in energy calibration is the drift correction $D(d,t)$, which accounts for changes in detector response over time. This correction can be made using either the Light Injection system or through-going cosmic rays. The LI system accomplishes this by pulsing approximately 300 (for the Far Detector) or 1000 (for the Near Detector) times per hour, for an average pulse of about 50 PE per PMT pixel. From this response, one can determine the average ADC count per photoelectron in each channel (the gain). This gain data is monitored over time and is used throughout the reconstruction. Environmental conditions in the MINOS caverns are relatively stable, causing the detectors' response to vary by less than 4% over the course of a given year [52].

While this LI technique measures the response of the readout electronics and PMTs, it does not monitor the response of the scintillator and wavelength shifting fibres. To account for this, cosmic ray muons are used to make the drift correction instead. The overall flux of these muons, and the median energy deposited in the detector, should remain constant with time. The distribution of energy per plane is plotted daily, and the median of the distribution is recorded. The drift of this median with respect to a standard time t_0 is used to make the drift correction:

$$D(d,t) = \frac{\text{Median Response}(d,t_0)}{\text{Median Response}(d,t)} \tag{3.2}$$

Because of the aging of the scintillator and other MINOS detector elements, the response of the detector has slowly dropped over time. This drift is typically no more than 2% per year [52].

3.4.2 Linearity Correction

The Light Injection system can also be used to characterize the linearity of the readout response. Both the PMTs and the Far Detector electronics exhibit nonlinearity on the order of 5-10% when the light signal is over 100 photoelectrons [65] [66]. PMT nonlinearity is measured and then corrected for by determining readout response versus true illumination. The LI system LED box pulses 1000 times over a broad range of light levels. This LED light signal is monitored by a pair of PIN diodes, one with a high gain, and one with a low gain. The resulting PMT response is then parametrized as a function of light level. The calibration applies this linearity correction to determine the true illumination produced by the signal.

3.4.3 Strip-to-Strip Correction

There are numerous effects which cause variations in the detector response on a strip-to-strip basis. These variations can be caused by differences in PMT response, WLS fibre efficiency, attenuation in the clear plastic cables, etc. These disparate effects are grouped under a single correction for variations between individual strips. The correction is calculated with through-going cosmic rays, which are selected to be normally incident through the center of the strip. The deposited energy must also be corrected for the photoelectron production threshold and for hits which clip a strip corner. The mean response of an individual strip with respect to the detector average is then used as a correction for that strip:

$$D(d,t) = \frac{\text{Mean Detector Response}(d,t)}{\text{Mean Strip Response}(s,d,t)} \tag{3.3}$$

The Near Detector strips have been found to have variations from the mean on the order of 29%, with the calibration stable to 2.1%, while the Far Detector strips vary on the order of 30% and are stable to within 4.8% [52].

3.4.4 Attenuation Correction

Following reconstruction, a correction must next be made based on the original location of the signal along the scintillator strip, to account for the attenuation of light as it travels along the WLS fibre. This is done using scintillator response data from the Module Mapper. The mapper response was fit using the following equation:

$$A(x) = A_1 e^{-x/L_1} + A_2 e^{-x/L_2} \tag{3.4}$$

where x is the distance along the strip, L_1 and L_2 are two attenuation constants, and A_1 and A_2 are two normalizing constants. $A(x)$ is then used to correct for attenuation. This method has been cross-checked using cosmic ray muons, and the two methods have been found to be consistent.

3.4.5 Relative and Absolute Calibration

Following attenuation corrections, the calibration is nearly complete. Two final steps must now be accomplished: the calibration of the relative energy scale of the two detectors, and the calibration of the absolute energy scale. These two steps comprise stage $M(d)$ in Equation 3.1.

While Sections 3.4.1 to 3.4.4 make the response of the detector uniform in terms of interaction time and location, the relative response of the two detectors must still be determined. The two detectors are designed to be functionally equivalent, but there are still several features, such as fiducial volume size, readout design, electronics, etc., which cause differences in response. This calibration of the inter-detector energy scale is carried out using stopping muons, which comprise 0.3% of the cosmic ray events in the Far Detector. The Bethe-Bloch relationship in Figure 3.13 can be used to find the total ionizing energy deposited in the MINOS detector, or stopping power, as a function of muon momentum. Muon momentum is determined in MINOS from track range, with curvature in the magnetic field as a cross-check. Only the portion of the track in the 0.5 to 1.1 GeV range is used, as dE/dx varies relatively slowly in this region [67]. The energy deposition is then used to set the inter-detector energy scale, in units of MIPS (also known as MEUs, or Muon Energy Units). The error on this relative energy scale

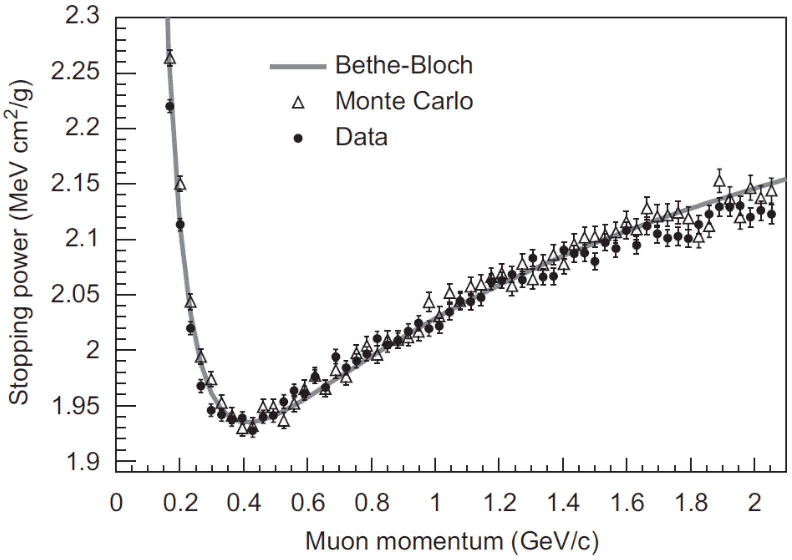

Figure 3.13: Muon stopping power, used for the relative calibration of the MINOS detectors. The Bethe-Bloch prediction is shown in gray. The triangles indicate MINOS GEANT3 Monte Carlo, and the circles indicate Far Detector data (these plots have been normalized to the Bethe-Bloch curve). Image from [52].

is approximately 2% [52].

Finally, this relative energy (in MEU) must be turned into an absolute energy (in GeV). This is done using data from the Calibration detector. As described earlier, the Calibration detector was exposed to a CERN beam of electrons, muons, and hadrons at several steps of fixed momentum from 200 MeV up to 10 GeV. Topology, time of flight, and measurements from fixed Cerenkov detectors were used to distinguish between particle types. The data were then plotted as a function of calorimeter response. The response was found to be well-described by the GEANT/GCALOR [72] simulation used in the Near and Far Detectors. The data are used to set the normalization of the detector response in GeV, as shown in Figure 3.14. Stopping muons were used to verify the accuracy of this approach to within 3%. For electrons, there was found to be a better than 2% agreement between the Monte Carlo and calorimeter data [68], while the simulations of hadronic showers caused by pions and protons agreed with data to within 6% [69]. The energy resolution was found to be $21.4\%/\sqrt{E} \oplus 4\%$ for electromagnetic showers, and $56\%/\sqrt{E} \oplus 2\%$ for hadronic showers (where E is energy in GeV) [52].

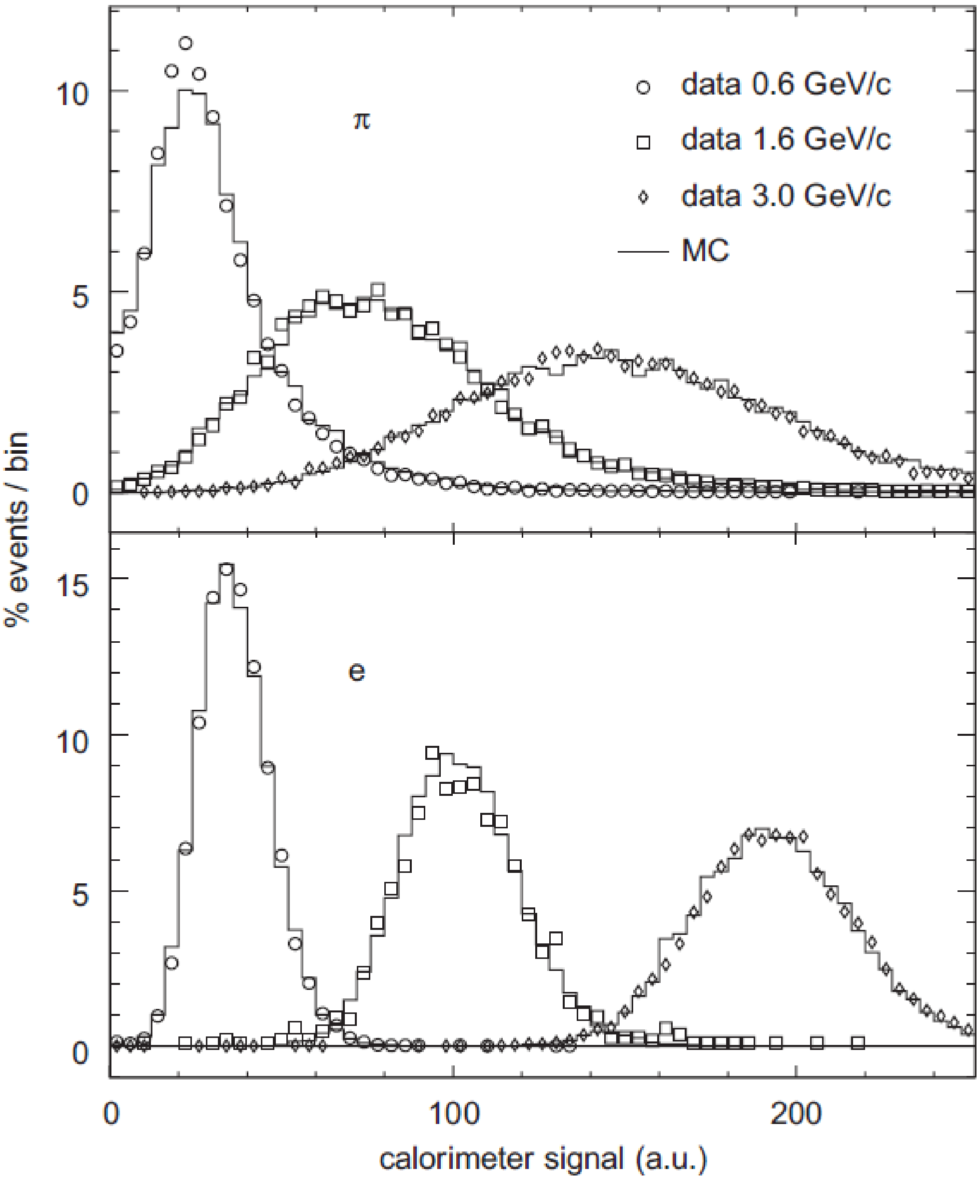

Figure 3.14: Response of the MINOS Calibration Detector to pions and electrons at three fixed momenta from a CERN test beam. These data are compared to the relevant MINOS Monte Carlo simulation. Image from [52].

3.4.6 Photoelectron Branch

The second calibration chain in Figure 3.12, the photoelectron branch, does not play an immediate role in the energy calibration of MINOS. Instead, gain information from the LI system is used, along with photoelectron statistics, to turn ADCs directly into units of photoelectrons in individual strips. These gains are calculated as follows. As described in Section 3.4.1, an average of $n = 50$ PEs are provided per LI pulse to measure the readout response. If the resulting ADC response distribution in a channel has mean μ and RMS width σ, then one can calculate the gain G as follows [50]:

$$G = \frac{\mu}{n} = \frac{\mu}{(\mu^2/\sigma^2)} = \frac{\sigma^2}{\mu} \tag{3.5}$$

Test stand measurements [62] indicate that the PMTs do not have a perfect 1 PE resolution, but instead a fractional width w. This width can be used to apply a correction to the gain measurement, as follows:

$$G = \frac{\sigma^2}{\mu}(1 + w^2) \tag{3.6}$$

The PE counts calculated from these gains are used in the reconstruction, in physics analyses, and in any other task requiring the use of raw photoelectrons [61].

3.5 Monte Carlo Generation

Understanding the behavior of MINOS also requires the use of simulated Monte Carlo data. The MINOS Monte Carlo is produced in three steps: the simulation of the NuMI beamline, the simulation of neutrino interactions, and the simulation of the detector response.

3.5.1 Simulation of the Beam Flux

In order to model the NuMI beamline, it is necessary to simulate the following: hadron production at the target, the propagation of those hadrons through the horn and decay pipe, and the propagation of the neutrinos to the Near and Far Detectors. The collision of 120 GeV protons on the hadron target is modeled by FLUKA05 [70]. These hadrons then travel forward into the magnetic focusing horn fields. A combination of GEANT3 [71] and FLUKA, referred to as GNuMI, simulates the decay chains and hadronic interactions which produce the neutrinos in the beamline. The magnetic focusing horns are assumed

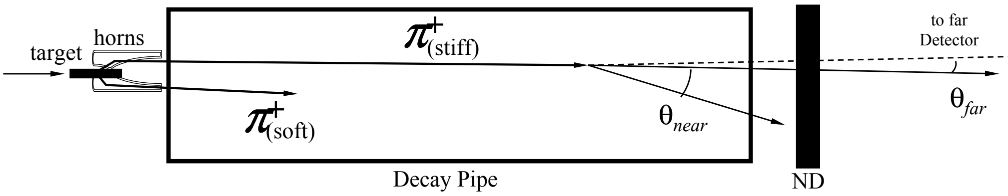

Figure 3.15: Diagram showing the different angular acceptances for neutrino production in the Near and Far Detectors. While the Near Detector accepts a large solid decay angle, the Far Detector only accepts very forward going neutrinos, with a small angle and a position further down the decay pipe. Image from [53].

to have a uniform current inside their conductors, producing a 1/r toroidal magnetic field.

The final stage of beam simulation is the propagation of the beam neutrinos to both the Near and the Far Detector. This stage assigns both an energy and a trajectory to each of the neutrinos. The energy is derived from decay kinematics [53]. For a two body decay $\pi/K \rightarrow \mu + \nu$, the neutrino energy is given by:

$$E_\nu = \frac{(1 - \frac{m_\mu}{M^2})E}{1 + \gamma^2 \tan^2 \theta_\nu} \tag{3.7}$$

where M is parent hadron mass, E is parent hadron energy, m_μ is muon mass, θ_ν is decay angle in the lab frame, and $\gamma = E/M$ is the Lorentz boost. The probability of a meson decay occurring at a given solid angle is given by the following expression:

$$\frac{dP}{d\Omega} \approx \frac{1}{4\pi} \frac{4\gamma^2(1 + \tan^2 \theta_\nu)^{3/2}}{(1 + \gamma^2 \tan^2 \theta_\nu)^2} \tag{3.8}$$

As is illustrated in Figure 3.15, the kinematics of neutrino production in the NuMI beamline result in the Near and Far Detectors each accepting different angles of hadronic decay, and therefore different neutrino energy spectra. The Near Detector can accept a much wider solid angle of decay, including higher energy pions which decay further down the pipe. The Far detector only accepts very forward-going neutrinos with a small angle of decay, resulting in a different solid angle acceptance and spectrum of neutrino energy. The beam flux drops with distance as a function of $1/\mathrm{R}^2$, resulting in an overall neutrino flux 10^6 smaller in the Far Detector. The resulting ratio of the unoscillated Far Detector neutrino spectrum to the Near Detector spectrum can be seen in Figure 3.16.

There are several sources of uncertainty in the simulation of the beam flux. The

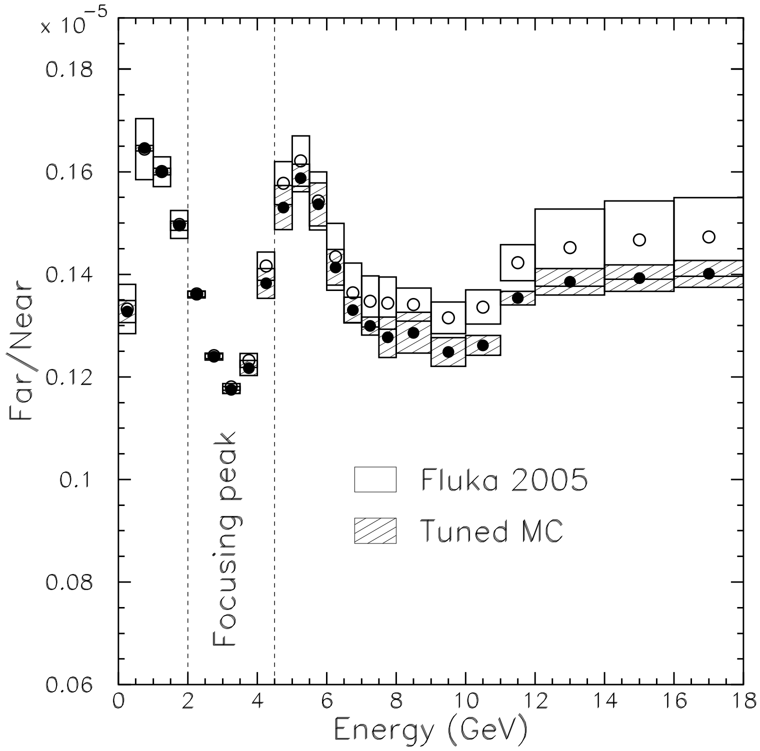

Figure 3.16: Monte Carlo ratio of the predicted unoscillated ν_μ spectrum in the Far Detector to the spectrum in the Near Detector. Boxes indicate uncertainties. Both the 2005 FLUKA Monte Carlo and the SKZP-tuned version are shown. Image from [53].

biggest uncertainties come from the production of hadrons at the target. Other uncertainties arise from physics interactions in the beamline, and from the magnetic horn focusing (due to current, alignment, etc.). To correct for these uncertainties, the Near Detector ν_μ CC data is used to "tune" the Monte Carlo. The horn and target arrangement determine the Near Detector acceptance as a function of pion longitudinal and transverse momentum (p_z, p_T). A parametrization of $\mathrm{d}^2N/\mathrm{d}p_z\mathrm{d}p_T$ is fit to the data for several beam configurations, and the results of the fit are then used to tune the Monte Carlo. This so-called SKZP tuning [73] is applied to both detectors. Figure 3.16 compares the Far/Near ratio of the standard MC to that of the SKZP-tuned MC as a function of energy.

3.5.2 Simulation of Neutrino cross sections and Hadronic Final States

Neutrino interactions in MINOS are simulated using NEUGEN (Neutrino Event Generator) - v3 [74]. NEUGEN is responsible for modeling the cross sections for each interaction type. It also simulates hadronization for those events containing showers. Finally, it models the intranuclear rescattering of hadrons as they exit the iron nuclei.

The cross sections of the various types of Neutral and Charged Current interactions are all calculated differently. The cross section for quasi-elastic (QE) scattering events, which dominate neutrinos with low energies < 1 GeV, is simulated using weak form factors: a dipole axial form factor, and a pseudoscalar weak form factor (which is negligible except for ν_τ scattering). In the case of resonance interactions (RES), a production model based on the Rein and Seghal model [75] is used, which ignores lepton mass. The resonance decay is assumed to occur isotropically in the reference frame of the hadron. The cross section model for deep inelastic scattering (DIS) events is based on parton model structure functions, taken from PDFLIB [76]. Finally, coherent pion production (COH) is again calculated using the Rein-Seghal model [77].

The NEUGEN event generator chooses the interaction type based on these relative cross sections. The hadronic final state is then modeled, using the AGKY model [105]. AGKY, developed in large part by members of MINOS, represents an attempt to improve on previous simulations of hadronization for neutrino interactions. For events with high invariant mass ($W > 3.0$ GeV), PYTHIA/JETSET [79] is used for shower modeling. For events with low invariant mass ($W < 2.3$ GeV), a phenomenological model based on KNO scaling [80] is used, which has been tuned to external bubble chamber data. The data are taken from hydrogen, deuterium, freon, and neon targets, from the BEBC [106]

[107] [108] [109], SKAT [110], and the Fermilab 15' [111] bubble chamber experiments. The intermediate region 2.3 GeV < W < 3.0 GeV uses both models and gradually scales between the full use of KNO (at W=2.3 GeV) and the full use of JETSET/PYTHIA (at W=3.0 GeV). AGKY is responsible for all aspects of hadronization, including final shower particle multiplicity and transverse momentum selection.

The final step of the NEUGEN simulation is to simulate the scattering of the hadronic shower particles as they leave the nucleus. This is handled by the INTRANUKE simulation [78], which is also tuned using bubble chamber data (exiting data from neutrino interactions in neon at BEBC [112] and measurements of pion interactions in deuterium at the ANL 12' bubble chamber) and takes into account effects such as the absorption and exchange of charge and the scattering of pions.

3.5.3 Simulation of the Detectors

Once NEUGEN and INTRANUKE have produced neutrino interactions and their final hadronic states, these events must be turned into simulated data in the MINOS detector. This is done in several stages. First, a GEANT3 simulation models the resulting particle transport, and the energy deposition in the scintillator strips. This simulation contains a detailed model of the MINOS detectors and the surrounding cavern rock. A package known as PhotonTransport [81] simulates scintillation, and propagates this signal through the detector to the PMT photocathode. This process makes use of the calibration constants derived from the correction process in Figure 3.12. The DetSim package then simulates the detector readout and digitization to produce an output PE signal. The final output of DetSim is a Monte Carlo event which resembles the raw MINOS data. The event signal is stored in units of ADCs, which can be calibrated and reconstructed in the same manner as the data events.

One important element in the simulation is the modeling of cross-talk, i.e. the leaking of charge between two separate pixels on a photomultiplier tube. This can happen in two separate ways. Charge cross-talk results either when electrons leak from one dynode chain to another, or when charge is induced by coupled anodes. Optical cross-talk occurs during the initial photoelectron production, when photons at the end of a readout fibre reflect or refract to a second pixel, or when photoelectrons produced near the edge of a pixel accidentally leak into another dynode chain. Cross-talk is simulated in both the Near and Far Detectors as part of the official Monte Carlo. Both sets of PMTs exhibit up to a 4% level of cross-talk for any signal [52]. This is simulated using the separate readout maps of the two detectors, the earlier gains measured by the LI system, and photoelectron statistics. As will be discussed later, this part of the simulation is known to be poorly modeled. Because cross-talk hits will normally be incorporated into showers, rather than tracks, this effect will need to be corrected for in any search for ν_e CC events. The general MINOS reconstruction corrects for this mismodeling by disregarding any strip hits which have less than a 2 PE pulse height.

Chapter 4

The ν_e Appearance Analysis

MINOS searches for two major modes of oscillation. The first of these is the disappearance of muon neutrinos. The survival probability for ν_μ is as follows:

$$P(\nu_\mu \rightarrow \nu_\mu) \approx 1 - \sin^2(2\theta_{23}) \sin^2 \left(\frac{1.27 \Delta m^2_{32} L}{E} \right) \tag{4.1}$$

This mode allows for a measurement of Δm^2_{32} and θ_{23}. The disappearance of ν_μ is interpreted as $\nu_\mu \rightarrow \nu_\tau$ oscillation, although the oscillation energy is typically not high enough to produce ν_τ CC events.

MINOS can also look for a second oscillation mode, $\nu_\mu \rightarrow \nu_e$:

$$P(\nu_\mu \ \rightarrow \nu_e) \approx \sin^2(2\theta_{13}) \sin^2(\theta_{23}) \sin^2 \left(\frac{1.27 \Delta m^2_{23} L}{E} \right) \tag{4.2}$$

This oscillation will occur very infrequently if the value of θ_{13} is zero[1]. This chapter will describe the challenges of a ν_e appearance analysis. It will also introduce the basic tools required for the work in this thesis. These tools were developed during the two previous MINOS analyses in 2009 [85] and in 2010 [86].

4.1 MINOS and the Search for θ_{13}

A number of upcoming accelerator and reactor neutrino experiments have been designed with the primary purpose of making a measurement of θ_{13}. As a running experiment, MINOS has the ability to make the first measurement, or at least constrain the current best limits. Figure 4.1 shows spectrum-averaged oscillation probability for $\nu_\mu \rightarrow \nu_e$ as

---

[1] In the full oscillation equations, there is some small amount of $\nu_\mu \rightarrow \nu_e$ oscillation due to Δm^2_{21}-scale solar neutrino oscillations

a function of distance, with the MINOS Far Detector indicated. The actual predicted oscillation in MINOS is much more complicated than the two-neutrino approximation in Equation 4.2. Much of this complexity comes from the inclusion of matter effects on neutrino oscillation, as the NuMI beamline passes through the Earth. When the full-three neutrino case is considered, and these matter effects are taken into account, the oscillation probability is approximately[2]:

$$\begin{aligned}P(\nu_\mu \rightarrow \nu_e) \approx\; & \sin^2(2\theta_{13})\sin^2\theta_{23}\frac{\sin^2(A-1)\Delta}{(A-1)^2} \\ & +2\alpha\sin\theta_{13}\cos\delta\sin 2\theta_{12}\sin 2\theta_{23}\frac{\sin A\Delta}{A}\frac{\sin(A-1)\Delta}{A-1}\cos\Delta \\ & -2\alpha\sin\theta_{13}\sin\delta\sin 2\theta_{12}\sin 2\theta_{23}\frac{\sin A\Delta}{A}\frac{\sin(A-1)\Delta}{A-1}\sin\Delta \end{aligned} \tag{4.3}$$

where the following substitutions have been made:

$$\Delta = \frac{\Delta m^2_{32}L}{4E}, \quad A = \frac{G_F n_e L}{\sqrt{2}\Delta}, \quad \alpha = \frac{\Delta m^2_{21}}{\Delta m^2_{32}} \tag{4.4}$$

The oscillation probability now depends not only on θ_{13}, but on the CP violation phase[3] δ and the sign of the mass splitting term Δm^2_{32} as well. Figure 4.3 shows $\nu_\mu \rightarrow \nu_e$ oscillation probability as a function of neutrino energy, for a single non-zero θ_{13} and for various choices of these parameters. The choice of sign for Δm^2_{32} determines the so-called Mass Hierarchy - the ordering of the neutrino mass states. This hierarchy can be either Normal or Inverted (see Figure 4.2). Matter effects will cause oscillations in the Inverted Hierarchy case to be suppressed. MINOS will therefore not make a direct measurement of θ_{13}, but a measurement which is also dependent on δ and mass hierarchy choice.

Prior to summer 2011, there were some hints of a non-zero θ_{13}. As described earlier, CHOOZ set an upper limit of $\sin^2 2\theta_{13} < 0.16$ at 90% C.L [43] [42] [45]. Further constraints on θ_{13} also come from solar and KamLAND data, as well as other long-baseline and reactor experiments. A 2009 global fit [87] combined these results together and found $\theta_{13} > 0$ at a level of approximately 2σ, with a best fit of $\sin^2 2\theta_{13} = 0.08$. A 2010 MINOS $\nu_\mu \rightarrow \nu_e$ search also placed new limits on θ_{13} [86]. This search predicted 49.1±9.7 events and saw 54 events, for a non-significant 0.7 σ excess. This resulted in limits of $\sin^2 2\theta_{13} < 0.12$ for the Normal Hierachy, and $\sin^2 2\theta_{13} < 0.20$ for the Inverted Hierarchy, both assuming $\delta = 0$. New results by T2K and MINOS were also announced towards

---

[2] The full derivation and oscillation equation with higher order terms can be found elsewhere [48] [49] [50].

[3] δ is referred to as δ_{CP} throughout this document.

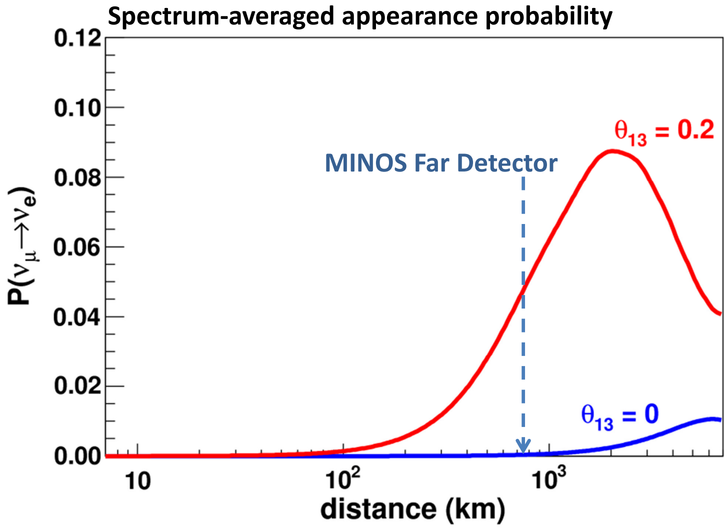

Figure 4.1: $\nu_\mu \rightarrow \nu_e$ oscillation probability for both a large and zero value of θ_{13}. The oscillation probability is averaged over the 1-6 GeV range for a MINOS-style wide band beam. The non-zero oscillation probability for $\theta_{13} = 0$ is due to second order effects from the solar terms θ_{12} and Δm^2_{21}. Plot produced by Ryan Patterson [83].

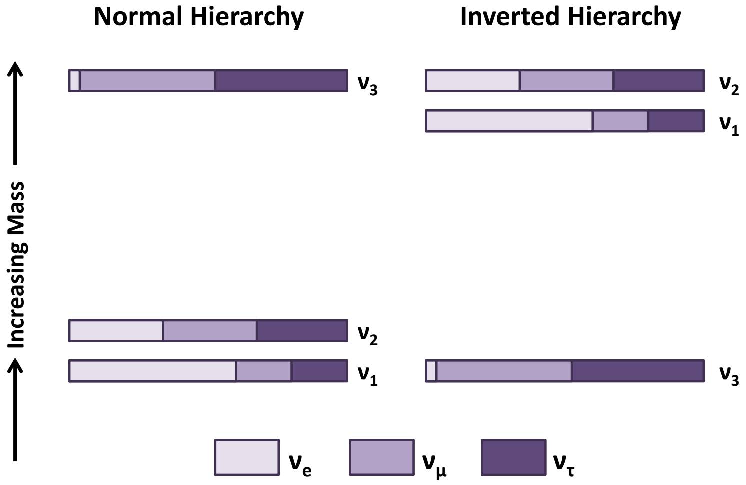

Figure 4.2: Illustration of the mass hierarchy question for neutrinos. The mass eigenstates (shown with the relative mixings of weak eigenstates) are stacked from lowest mass at the bottom to highest mass at the top. The absolute values of the mass differences between these states (squared), Δm^2_{21} and Δm^2_{32}, are known, but the actual ordering is still unknown. The sign of Δm^2_{21} is known from solar neutrino experiments. The left stack is referred to as the Normal mass hierarchy, while the right stack is the Inverted mass hierarchy.

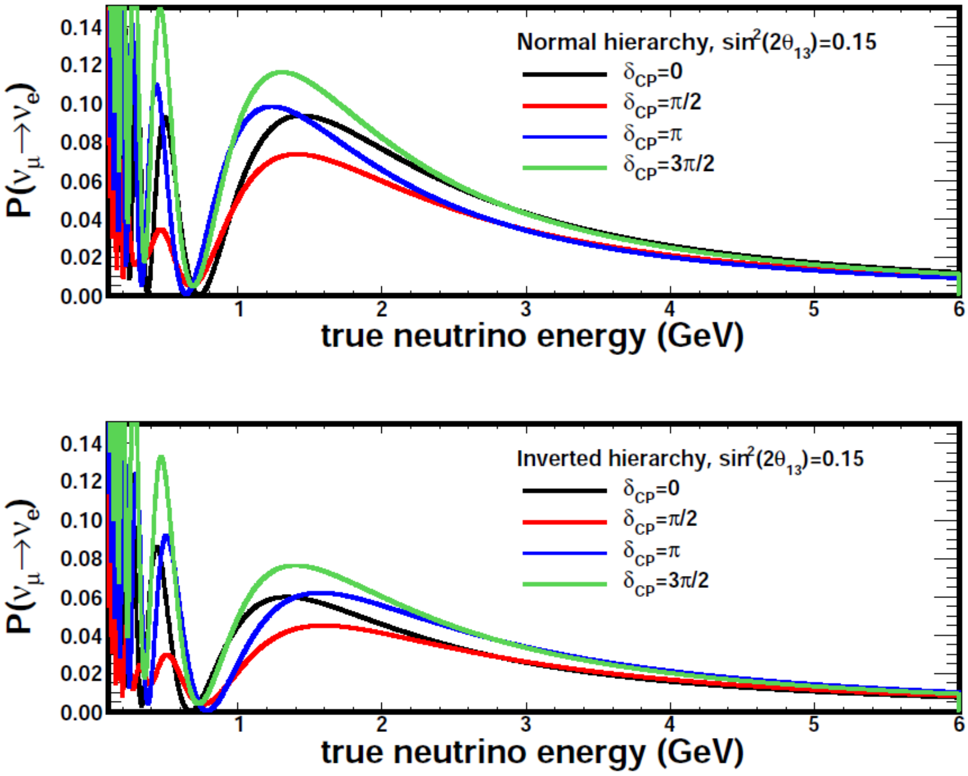

Figure 4.3: The probability for $\nu_\mu \to \nu_e$ oscillation, shown as a function of true neutrino energy assuming the MINOS oscillation length (735 km). The oscillation probability is shown for both the Normal (top) and Inverted (bottom) hierarchies, for various values of the CP violation phase term δ. A MINOS measurement of θ_{13} will depend on both of these factors (mass hierarchy sign and δ_{CP}). Plots produced by Pedro Ochoa [50].

the end of the writing of this thesis, and will be described in Appendix A.

4.2 Analysis Considerations

The signal in this analysis consists of electron neutrino Charged Current (ν_e CC) events which appear via oscillation in the MINOS Far Detector. There are several major challenges to overcome in observing this ν_e appearance signal. The first of these is the relatively small size of θ_{13} itself, and the corresponding low probability of $\nu_\mu \rightarrow \nu_e$ oscillation. Even with a signal at the size of the CHOOZ limit, 2% of events in the MINOS Far Detector at most are expected to be ν_e CC signal events. Second, MINOS has been designed to search for ν_μ CC interactions, not ν_e CC interactions. Consequentially, the detector resolution for signal event topology is poor. Finally, partially as a consequence of this poor resolution, there are numerous sources of background contamination, particularly from Neutral Current events, which can be very difficult to separate from signal. All three of these features necessitate a sophisticated analysis for observing ν_e appearance at the Far Detector.

4.2.1 The signal

An example of a simulated signal ν_e CC event can be seen in the first image in Figure 4.4. The signature of a ν_e CC interaction is the production of an electron. Because of its small mass, this electron is easily deflected by atomic nuclei as it travels forward and passes through matter, producing bremsstrahlung radiation photons. These photons then produce more e^+/e^- pairs, which are similarly deflected. This process results in a compact energy deposition pattern known as an electromagnetic (EM) shower. The mean longitudinal energy profile of the shower, $\mathrm{d}E/\mathrm{d}t$, can be approximately described by the following function [15]:

$$\frac{\mathrm{d}E}{\mathrm{d}t} = E_0 b \frac{(bt)^{a-1} e^{-bt}}{\Gamma(a)} \tag{4.5}$$

where t is the distance measured in units of radiation length (dependent upon material), E_0 is the original electron energy, a is a variable describing shower rise, b describes shower fall, and Γ is the gamma function. Table 4.1 shows the average size of an EM shower in MINOS compared to the relevant detector parameters. Rather than a distinct event topology, an electromagnetic shower in MINOS looks like a compact clustering of strips with very little detail.

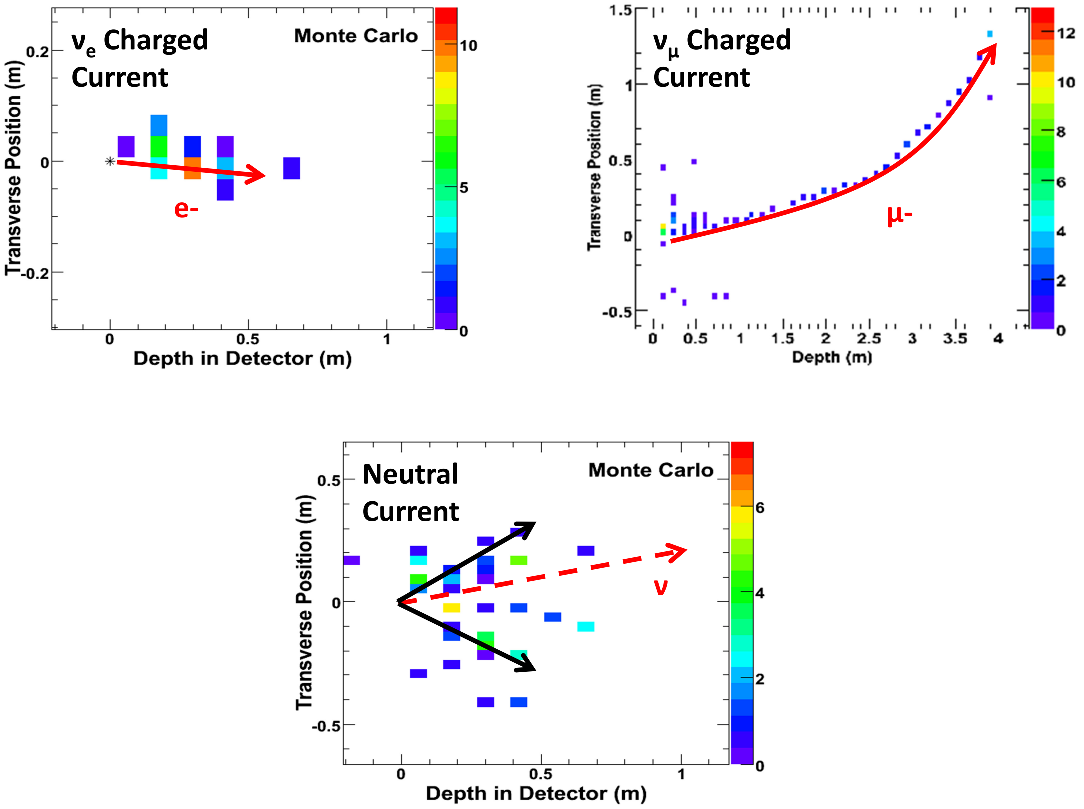

Figure 4.4: Transverse views (taken in one plane view only) of three types of event in the MINOS detector: ν_e Charged Current (top left), ν_μ Charged Current (top right), and Neutral Current (bottom). Events are Far Detector Monte Carlo. Hit pulse-heights are measured in MIPs, and have been cleaned of hits < 2 PE to remove cross-talk. The lines showing the presence of leptons and hadrons are purely illustrative.

Average MINOS EM Shower	Relevant Detector Parameter
Effective radiation length = 4.06 cm	Plane separation = 5.95 cm
Molière radius = 3.7 cm	Strip width = 4.12 cm

Table 4.1: Average size of an electromagnetic shower in MINOS, versus the relevant detector parameter. Molière radius is here defined as the radius of a cylinder situated along the shower axis which contains 90% of the energy in an electromagnetic shower. Numbers from [50].

In addition to this electromagnetic shower component, many ν_e CC events also contain a large hadronic shower component. The fraction of total neutrino energy going to the hadronic shower in a ν_e CC event is quantified by the variable y (often known as the inelasticity):

$$y = \frac{E_\nu - E_\ell}{E_\nu} = 1 - \frac{E_\ell}{E_\nu} \tag{4.6}$$

where E_ν is the incoming neutrino energy, and E_ℓ is the outgoing lepton (in this case, electron) energy. The smaller the value of y, the more electromagnetic-like an event is, and the more easily it can be distinguished from background events such as Neutral Current. While QE events tend to have a value of y peaking close to 0.0, DIS events have a cross-section which is nearly flat as a function of y. The inelasticity of RES and COH events can usually be found somewhere in between.

4.2.2 Background events

The following sources of background complicate the search for signal ν_e Charged Current events:

Neutral Current (NC): Neutral Current (NC) interactions normally appear in the MINOS detector as hadronic showers, consisting of large number of particles at a relatively low transverse momentum. NC events often appear more "diffuse" than ν_e CC electromagnetic showers. An example of a simulated NC event can be seen in Figure 4.4. Two major factors make make NC events the most important background for a ν_e search. First, identifying a NC event becomes more difficult when the hadronic shower contains a π^0. The mean lifetime of a π^0 is 8.4×10^{-17} seconds, after which it will usually decay to two photons (branching ratio=98.8%) or to a photon and e^+/e^- pair (branching ratio=1.2%) [15]. Some fraction of NC events will therefore have an electromagnetic-like component, making them difficult to separate from the signal. Second, only some of the incoming neutrino energy for a NC interaction is transferred to the hadronic shower, while the remainder is carried, undetected, by the outgoing neutrino. The reconstructed energy of a NC event is consequentially only a fraction of the true neutrino energy. NC events from higher energy neutrinos can therefore be reconstructed with an energy in the $\nu_\mu \rightarrow \nu_e$ oscillation range, contaminating the region of analysis interest. In fact, about half of the NC events in the $1 - 8$ GeV region come from neutrinos with true energy > 7 GeV [50].

ν_μ Charged Current (CC): Muon neutrino CC events (see Figure 4.4) are normally very easy to recognize. Because the muon produced in the CC interaction has a mass approximately 200 times larger than that of an electron, the muon will produce a long ionizing track in the detector. A ν_μ CC event with a high inelasticity y can be harder to recognize, however. In these events, a large fraction of the neutrino energy is transferred to the hadronic shower, resulting in a shorter track which can become embedded in the shower. Without the tagging muon, these ν_μ CC hadronic shower events can be difficult to distinguish from ν_e CC EM shower events.

Beam ν_e CC: Another form of background contamination comes from the $\nu_e/\overline{\nu}_e$ component of the beam. These "beam" electron neutrinos are produced primarily from secondary muon decay:

$$\pi^+ \rightarrow \mu^+ + \nu_\mu \tag{4.7}$$

$$\mu^+ \rightarrow e^+ + \nu_e + \bar{\nu}_\mu \tag{4.8}$$

These events are topologically identical to ν_e CC appearance events, and must therefore be included in an estimate of the total background contamination.

ν_τ CC: Producing a ν_τ CC interaction requires a neutrino energy >3.5 GeV. The peak for $\nu_\mu \rightarrow \nu_\tau$ oscillation is smaller than this threshold, making these interactions rare and as-yet-unobserved in MINOS. When a ν_τ CC event does occur, however, the τ particle will decay typically into into an electron or a hadronic system. This poses a small source of potential background.

4.3 Overview of the Analysis

The basic goal of this thesis is to measure θ_{13} by observing the statistically significant appearance of ν_e Charged Current events at the MINOS Far Detector. To do this, one must be able to predict the amount of background contamination present in the absence of $\nu_\mu \rightarrow \nu_e$ oscillations. An excess of events on top of this background prediction is interpreted to be a signal. The basic steps of the analysis are as follows.

1. **Event Selection.** A selection is applied to reject background events and select signal ν_e CC events in the MINOS Far Detector. Some basic cuts can be applied to eliminate obvious background contamination, but a statistically significant measurement of electron neutrino appearance requires the use of a sophisticated Particle ID (PID) algorithm. This thesis will employ a novel PID known as the Library Event Matching, or LEM, method. In the LEM PID, event strip hits are compared to those of events in a massive Monte Carlo library of signal and background events. Out of this comparison, a list of best matches is compiled and used to form a PID discriminant.

2. **Near Detector Decomposition:** The event selection in **1.** is applied to the Near Detector Data. This ND Data selection must then be broken down (or "decomposed") into its different constituent background types. Because of discrepancies between Data and MC, the ND MC alone cannot carry out this decomposition. Instead, an algorithm known as the HOOHE method is used, which exploits the relative background compositions of different NuMI beamline configurations.

3. **Extrapolation:** The ND Data selection is then extrapolated to form a Far Detector prediction, using Far/Near detector ratios taken from the Monte Carlo. Because of ν_μ CC oscillation, detector effects, and beam kinematics, the individual Near Detector background components (NC, ν_μ CC, and beam ν_e CC) must each be extrapolated separately to the Far Detector.

4. **Analysis Checks:** All potential sources of systematic error are identified and their impact assessed. The analysis framework is checked using two Far Detector analysis sidebands, which consist of FD Data separate from the ν_e analysis region.

5. **Unblinding the Far Data:** The full FD Data set is examined. Up until this point, the FD Data is kept blinded in the analysis region, to prevent bias. A multiple bin shape fit is used to assess the likelihood of a non-zero signal having been observed. Several possible interpretations of the result will be given.

4.4 $\boldsymbol{\nu_e}$ Appearance Analysis: Initial Steps

4.4.1 Preparing the Data and Monte Carlo for Analysis

Several changes must be made to the ND and FD Data and MC before they can be analyzed. The first of these changes is putting the events in the correct energy scale. The main MINOS calibration determines shower energy based on a hadronic energy scale. For this analysis, the hadronic energy scale must be converted to an electromagnetic energy scale. This conversion is calculated using signal and beam ν_e CC MC events in the Far Detector, and beam ν_e CC MC events in the Near Detector. Quasi-elastic events are used to set the EM energy scale, because a greater fraction of their energy is located on average in an electromagnetic shower. The true neutrino energies of these QE events (in GeV) are plotted versus their deposited energy in Minimum Ionizing Particles (MIPs), and a linear fit is performed between 1 and 8 GeV. In the Far Detector, a factor of MIP/GeV = 23.56 ± 0.03 is found, with an offset of -2.5 MIP. In the Near Detector, a similar factor of MIP/GeV = 23.36 ± 0.6 is found, with an offset of -2.8 MIP [88]. These factors are used to convert the reconstructed energy to the EM scale.

For Monte Carlo events, several weights are applied on an event by event basis. The SKZP weights are applied to both the Near and Far Detectors to correct for beam mismodeling. In the Far Detector, additional weights must also be applied to simulate the effects of oscillation. Three separate types of Monte Carlo file are used to simulate the Far Detector. In the so-called **beam** files, the beam has not oscillated at all between the Near and Far detector. In the $\boldsymbol{\nu_e}$ files, all of the muon neutrinos have been oscillated to electron neutrinos. In the $\boldsymbol{\nu_\tau}$ files, all of the muon neutrinos have instead been oscillated to tau neutrinos. The neutrino events in each file type are reweighted to produce the correct predicted neutrino spectrum for a given set of oscillation parameters. The standard oscillation parameters used in this analysis are listed in Table 4.2. Unless otherwise stated, the value of θ_{13} is set at the current CHOOZ limit.

Finally, both Monte Carlo samples and the Near Detector Data are reweighted to match a standard exposure. The FD Monte Carlo is normalized to the same exposure as the Far Detector data, 8.2×10^{20} protons on target (POT). The ND Data and Monte Carlo are both normalized to a standard exposure of 1×10^{19} POT.

Oscillation Parameter	Value
θ_{23}	$\pi/4$
θ_{12}	0.60
θ_{13}	0.21
$\|\Delta m^2_{32}\|$	2.32×10^{-3} eV^2
$\|\Delta m^2_{21}\|$	7.59×10^{-5} eV^2
δ_{CP}	0.0

Table 4.2: Oscillation parameters [15] [89] used in the oscillation of Far Detector Data. Angles are given in radians. θ_{13} is set at the CHOOZ limit.

4.4.2 Basic Cuts: Data Quality, Cosmic Rays, and Fiducial

Before selecting for ν_e Charged Current events, a round of cuts is applied to all Data and MC events to ensure event quality. These cuts remove any events which are either potentially poorly reconstructed or not part of a beam spill. A set of standard MINOS data quality cuts removes events which were recorded when the detector was not taking data properly (e.g., when the DAQ was malfunctioning). Additional cuts also ensure that the magnetic coil current was on and being run in forward horn current neutrino mode, and that the event was not a Light Injection event.

Several cuts are applied only to the Far Detector. A timing cut constrains events to be in a 14 μs window around the beam spill. While the vast majority of spills in the Far Detector will only have a single reconstructed event, occasionally there may be two or more. When this is the case, only the event with the highest pulse height is accepted, as the other event is probably a reconstruction artifact.

A fiducial volume cut is applied to both the Data and Monte Carlo, to ensure that events were far enough from the detector edge to be correctly reconstructed. If z is longitudinal distance in the detector, and r is radius from the beam center (which is offset by $x = 1.49$ m, $y = 0.14$ m in the Near Detector), the following events are retained:

- **Near Detector**: 1.0 m $< z <$ 4.0 m, $r <$ 0.80 m
- **Far Detector**: 0.5 m $< z <$ 14.3 m (supermodule 1), 16.3 m $< z <$ 28.0 m (supermodule 2), 0.5 m $< r <$ 3.7 m

These regions are shown as part of the larger detector ensemble in Figure 4.5.

Finally, the Far Detector sample is cleaned of cosmic ray interactions. Cosmic rays typically produce events which have a large vertical component or a large angle with

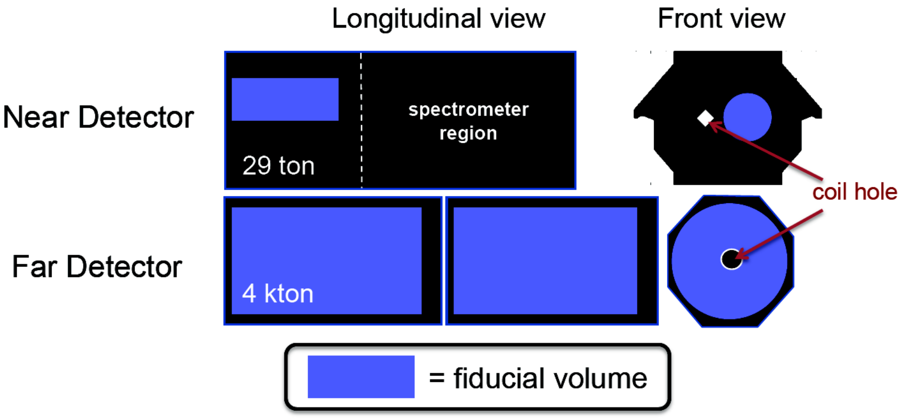

Figure 4.5: The fiducial volume regions in both the Near and Far Detectors. Images provided by Pedro Ochoa [50].

respect to the beam. If an event has a shower, the event is cut if its slope relative to the beam path (z) is >10. If an event has a track, the event is also cut if the track's vertical component is larger than 10 m, and if the angle with respect to vertical is less than 35°.

4.4.3 Preselection Cuts

After Data Quality, Fiducial, and Cosmic Ray cuts have been applied, only ~2% of the Far Detector MC sample consists of ν_e Charged Current signal events for a CHOOZ-sized signal. This ratio can be improved by applying a series of basic Preselection cuts which remove the most obvious background events.

The first Preselection cut is placed on reconstructed energy. Below 1 GeV, the accuracy of the reconstruction breaks down, and Neutral Current events dominate the sample. Above 8 GeV, essentially no $\nu_\mu \rightarrow \nu_e$ oscillation is expected, and almost all remaining ν_e Charged Current events will come from beam ν_e CC events. Only the intermediate region is therefore accepted (this selection can be seen in Figure 4.6):

- 1 GeV < Reco. E < 8 GeV.

A second set of Preselection cuts removes events which are almost certainly ν_μ CC events, due to the presence of a long muon-like track. A first cut selects events based on the number of "track planes," (track end plane) - (track beginning plane). A second cut selects events based on the number of "track-like planes," the number of planes which contain only a track hit with no shower hits. These two cuts, which are shown in Figure 4.7, select only events with the following short track-like topology:

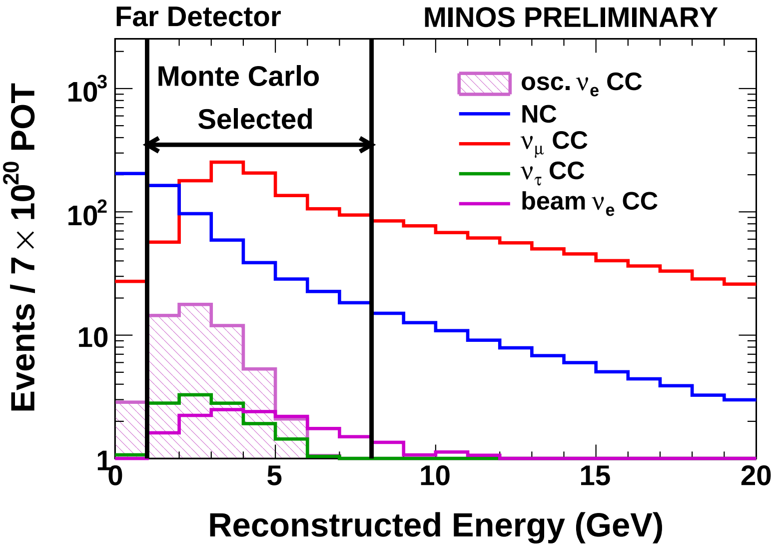

Figure 4.6: Reconstructed energy Preselection cut in MINOS. This plots is for Far Detector Monte Carlo, with the accepted region demarcated by black lines. This plot was produced for the 7.0 × 10^{20} POT ν_e appearance analysis [86] in MINOS.

- Track planes < 25
- Track-like planes < 16

The final two cuts remove events which were likely Neutral Current interactions. Electromagnetic showers in ν_e CC events tend to be much more dense than hadronic shower in NC events. Many NC events can therefore be cut by requiring that an event have at least four contiguous planes. Each of these planes must have at least 0.5 MIPs of deposited energy. While both NC and ν_μ CC events can contain a reconstructed track with no shower, almost all ν_e Charged Current events contain a shower. Therefore, all events are required to have at least one reconstructed shower. The following two selections are made:

- Contiguous Planes (min. 0.5 MIPS) > 4
- Number of Showers > 0

The total number of Far Detector Monte Carlo events selected after Fiducial and Data Cuts, and then after Preselection cuts, are listed in Table 4.3. The Preselection cuts are 76.8% efficient for signal, while removing 82.6% of the total background. A

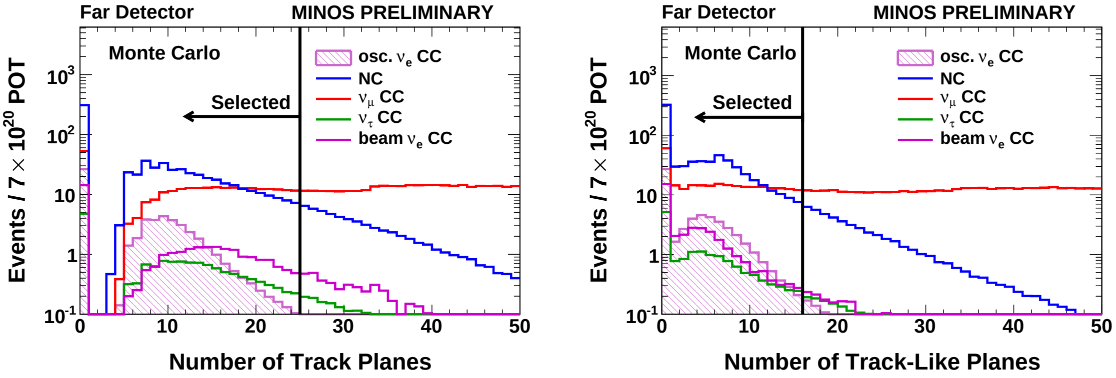

Figure 4.7: Track Preselection cuts in MINOS. These plots are for Far Detector Monte Carlo, with the accepted region demarcated by black lines. The top plot shows the cut on the total track length. The bottom plot shows the cut on the number of "track like" planes (planes containing only a track, with no shower hits). These plots were produced for the 7.0×10^{20} POT ν_e appearance analysis [86] in MINOS.

CHOOZ-sized signal comprises 10% of the Preselection sample, a large gain in sample purity over the Fiducial sample. More than 95% of the events which are cut by the Preselection will also be cut later by the particle ID algorithm. Applying the above Preselection cuts is still valuable, however, as it reduces analysis and processsing time and permits the PID to focus on more difficult cases of particle discrimination.

Selection	osc. ν_e CC	Total bg.	NC	ν_μ CC	ν_τ CC	Beam ν_e
Fiducial vol.	70.9	3173.8	873.6	2238.1	19.5	42.6
Preselection	54.4	553.0	338.1	189.7	10.5	14.74
efficiency (vs. Fid.)	76.8%	17.4%	38.7%	8.5%	54.0%	34.6%

Table 4.3: Number of Far Detector MC events selected by the Fiducial (including Data Quality) and then Preselection cuts. The corresponding selection efficiency of the Preselection with respect to the Fiducial sample is also given. A CHOOZ-sized signal is used.

4.4.4 Particle Identification Example: ANN11

Numerous particle identification methods have been explored over the course of the ν_e appearance analysis in MINOS. In both the 2009 [85] and 2010 [86] published analyses, an algorithm called ANN11 was used. The ANN11 PID will mentioned throughout this thesis, mostly as a point of reference for comparing the sensitivity gains of the LEM particle ID. A more detailed description of this PID can be found elsewhere [90].

ANN stands for artificial neural net, and the 11 refers to the eleven separate input variables of the PID. Each of these variables quantifies some aspect of the shape and energy profile of a candidate event. These variables include:

- Shower rise (a from Equation 4.5)
- Shower fall (b from Equation 4.5)
- Shower Molière radius (i.e., the radius of 90% energy containment)
- RMS of the transverse energy loss profile
- Magnitude of the longitudinal energy projection vectors (weighted by pulse height)
- The sum of the distance between hits in a Minimal Spanning Tree, formed from hits with pulse heights higher than the event mean
- Fraction of energy lost in 2, 4, and 6 planes
- Fraction of energy lost in 8 strips
- Fraction of the total energy contained within a 1.5 strip distance from the main shower axis

The output of the ANN11 PID is a number between 0.0 and 1.0, where 0.0 indicates high probability for a background event, and 1.0 indicates high probability for a signal event. A cut of ANN11>0.7 was found to have an optimal performance in separating signal from background while still keeping a high signal efficiency.

4.5 Some Analysis Tools

4.5.1 FOM and superFOM

In developing a particle ID for the ν_e analysis, some means of assessing relative PID performance must be found. Carrying out the full MINOS analysis for each new iteration of a PID discriminant would be prohibitively time-intensive. The sensitivity of a given discriminant to signal is therefore assessed initially using its "Figure of Merit" (FOM). A given particle ID selects N_{sig} signal ν_e CC events and N_{bkg} background events. The FOM measures how significant that signal selection is in comparison to statistical fluctuations in the background:

$$\text{FOM} = \frac{N_{\text{sig}}}{\sqrt{N_{\text{bkg}}}}. \tag{4.9}$$

In reality, the fluctuations on the background will have a contribution from systematic error as well. For this reason, the "super-FOM" is calculated. If f is the fractional systematic error on the background, the super-FOM is defined as follows:

$$\text{super-FOM} = \frac{N_{\text{signal}}}{\sqrt{\sigma^2_{\text{stat}} + \sigma^2_{\text{syst}}}} = \frac{N_{\text{signal}}}{\sqrt{N_{\text{bg}} + (\text{f} \times N_{\text{bg}})^2}} \tag{4.10}$$

At this point, the systematic error on the PID selection is typically unknown. As an approximation, a systematic error of 5%, the previously measured value for ANN11 [86], can be assumed. As will be shown later, this value is actually quite close to the real LEM systematic error.

4.5.2 The PORP Method

The above FOM method requires a prediction of the background N_{bkg} and signal N_{sig} selected by the PID. This prediction can be estimated with the Far Detector MC. The prediction is more accurately made, however, by first applying the PID selection to the Near Detector data, and then extrapolating this ND selection component by component to the Far Detector using Far/Near ratios from MC:

$$\text{FD}^{\text{pred}}_{\alpha,j} = \text{ND}^{\text{data}}_{\alpha,j} \frac{\text{FD}^{\text{MC}}_{\alpha,j}}{\text{ND}^{\text{MC}}_{\alpha,j}}, \tag{4.11}$$

where $\text{FD}^{\text{pred}}_{\alpha,j}$ is the Far Detector prediction for background component α (NC, ν_μ CC, or beam ν_e CC) and PID or Reconstructed Energy bin j, and $\text{ND}^{\text{data}}_{\alpha,j}$, $\text{FD}^{\text{MC}}_{\alpha,j}$, and $\text{ND}^{\text{MC}}_{\alpha,j}$ are the same but for ND Data, FD MC, and ND MC, respectively.

Performing this extrapolation requires knowledge of the fraction of each background type in the Near Detector. In the full analysis, the HOOHE method is used to perform this decomposition. This would again be time-consuming to perform for each iteration of the selection. As a first order approximation, however, one can use a simplified method called PORP[4] [100] to perform the background decomposition. In this method, the

---

[4] Named for its creators, Pedro Ochoa and Ryan Patterson.

assumption is made that the hadronic showers in NC and ν_μ CC events are mismodeled to an equal degree, while beam ν_e CC events have been modeled correctly. This is a reasonable assumption, as the background is dominated by Neutral Current events, with a smaller sample of NC-like ν_μ CC events. In comparison to the large error on the hadronic backgrounds, the error on the beam ν_e CC sample is comparatively small, and EM showers are known to be well-modeled from studies of the Calibration detector [99].

The PORP decomposition of the Near Detector data is performed by scaling the NC and ν_μ CC Near Detector Monte Carlo events by the same amount, and leaving the beam ν_e CC events unchanged from Monte Carlo:

$$\begin{aligned}
\mathrm{ND}^{\mathrm{data}}_{\mathrm{beam}\ \nu_e,j} &\simeq \mathrm{ND}^{\mathrm{MC}}_{\mathrm{beam}\ \nu_e,j} && (4.12)\\
\mathrm{ND}^{\mathrm{data}}_{\mathrm{NC},j} &\simeq \mathrm{ND}^{\mathrm{MC}}_{\mathrm{NC},j}\frac{\mathrm{ND}^{\mathrm{data}}_{\mathrm{all},j}-\mathrm{ND}^{\mathrm{MC}}_{\mathrm{beam}\ \nu_e,j}}{\mathrm{ND}^{\mathrm{MC}}_{\mathrm{all},j}-\mathrm{ND}^{\mathrm{MC}}_{\mathrm{beam}\ \nu_e,j}} && (4.13)\\
\mathrm{ND}^{\mathrm{data}}_{\nu_\mu\ \mathrm{CC},j} &\simeq \mathrm{ND}^{\mathrm{MC}}_{\nu_\mu\ \mathrm{CC},j}\frac{\mathrm{ND}^{\mathrm{data}}_{\mathrm{all},j}-\mathrm{ND}^{\mathrm{MC}}_{\mathrm{beam}\ \nu_e,j}}{\mathrm{ND}^{\mathrm{MC}}_{\mathrm{all},j}-\mathrm{ND}^{\mathrm{MC}}_{\mathrm{beam}\ \nu_e,j}} && (4.14)
\end{aligned}$$

This PORP method provides an excellent approximation of the final HOOHE background prediction. As an example, one can compare the background predictions for the optimized single-bin LEM particle ID cut, LEM > 0.70. The PORP method predicts a total of 46.6 events (33.9 NC events, 7.0 ν_μ CC events, and 5.7 beam ν_e CC events) for LEM>0.7. The HOOHE method predicts a total of 47.4 events (35.1 NC events, 7.1 ν_μ CC events, and 5.2 beam ν_e CC events). Given the simplicity of this algorithm, this is an excellent degree of agreement (within 2%) with the far more complicated HOOHE method.

Chapter 5

The Library Event Matching (LEM) Particle ID

5.1 The Principle Behind LEM

In order to detect the transition of ν_μ CC events to ν_e CC events, one must be able to distinguish ν_e CC electromagnetic showers from the various forms of background present in the MINOS detectors. The Preselection described in the previous section goes part of the way towards removing obvious background. However, a more sophisticated algorithm is still needed to identify ν_e CC events in the Far Detector. The aforementioned ANN11 analysis is one example of how to perform this discrimination. ANN11 and all other previous MINOS ν_e analysis particle IDs relied on the use of reconstructed quantities such as energy profile and event shape to distinguish between neutrino interaction types. This process of constructing higher-order variables results in a loss of information.

The Library Event Matching (LEM) particle ID solves this problem by, in principle, using all the event information available to identify ν_e CC events. Rather than using post-reconstruction variables, the LEM PID uses the information from the event strip hits themselves. While LEM is novel, the concept is also very simple. Each candidate input event is compared to a large library of Monte Carlo ν_e CC and NC events. Out of this comparison, a list of N best matches is compiled. Figure 5.1 shows an example of a good match and a bad match, for a candidate input event. The information from these N best matches is then fed into a discriminant, which produces a single number describing how signal-like or background-like the original input event was. This final number is the value of the LEM particle ID. The flow chart in Figure 5.2 shows this process graphically.

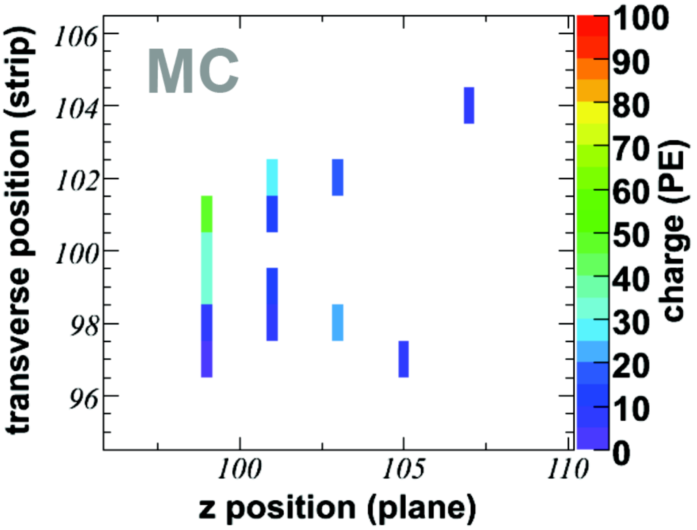

(a) Input Event

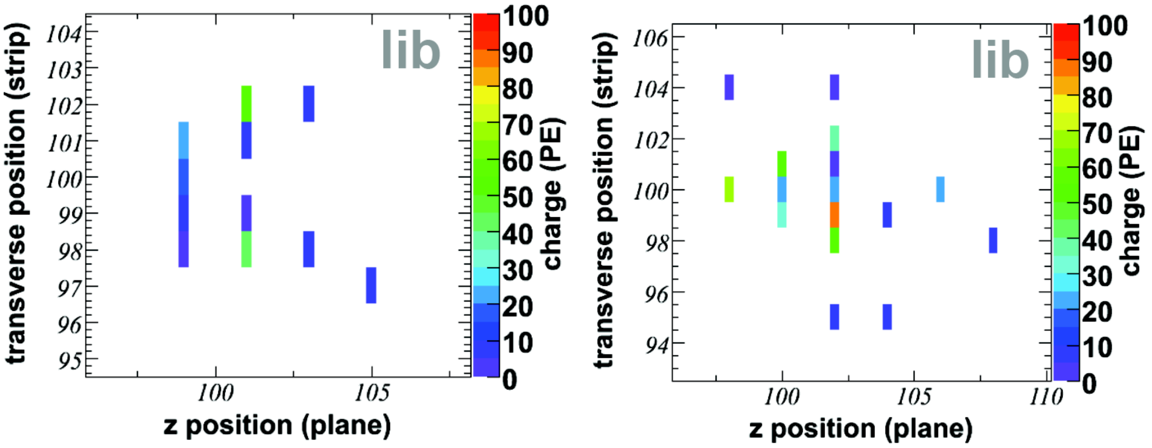

(b) A Good Library Match **(c)** A Bad Library Match

Figure 5.1: An example of the matching process. The event in (a), which one wants to identify as signal or background, is compared to the events in the LEM library. Some of these events will be good matches (b), while most will be bad matches (c). The candidate event is Far Detector Monte Carlo. Source of images: [92].

The downside of the Library Event Matching code is that it is computationally intensive. However, a number of methods have been found to decrease processing time drastically. Additionally, when this PID is incorporated into the full analysis chain (see Chapter 9), there is a 14% improvement in sensitivity to $\sin^2 2\theta_{13}$, an increase in senstivity which is equivalent to adding 30% more data.

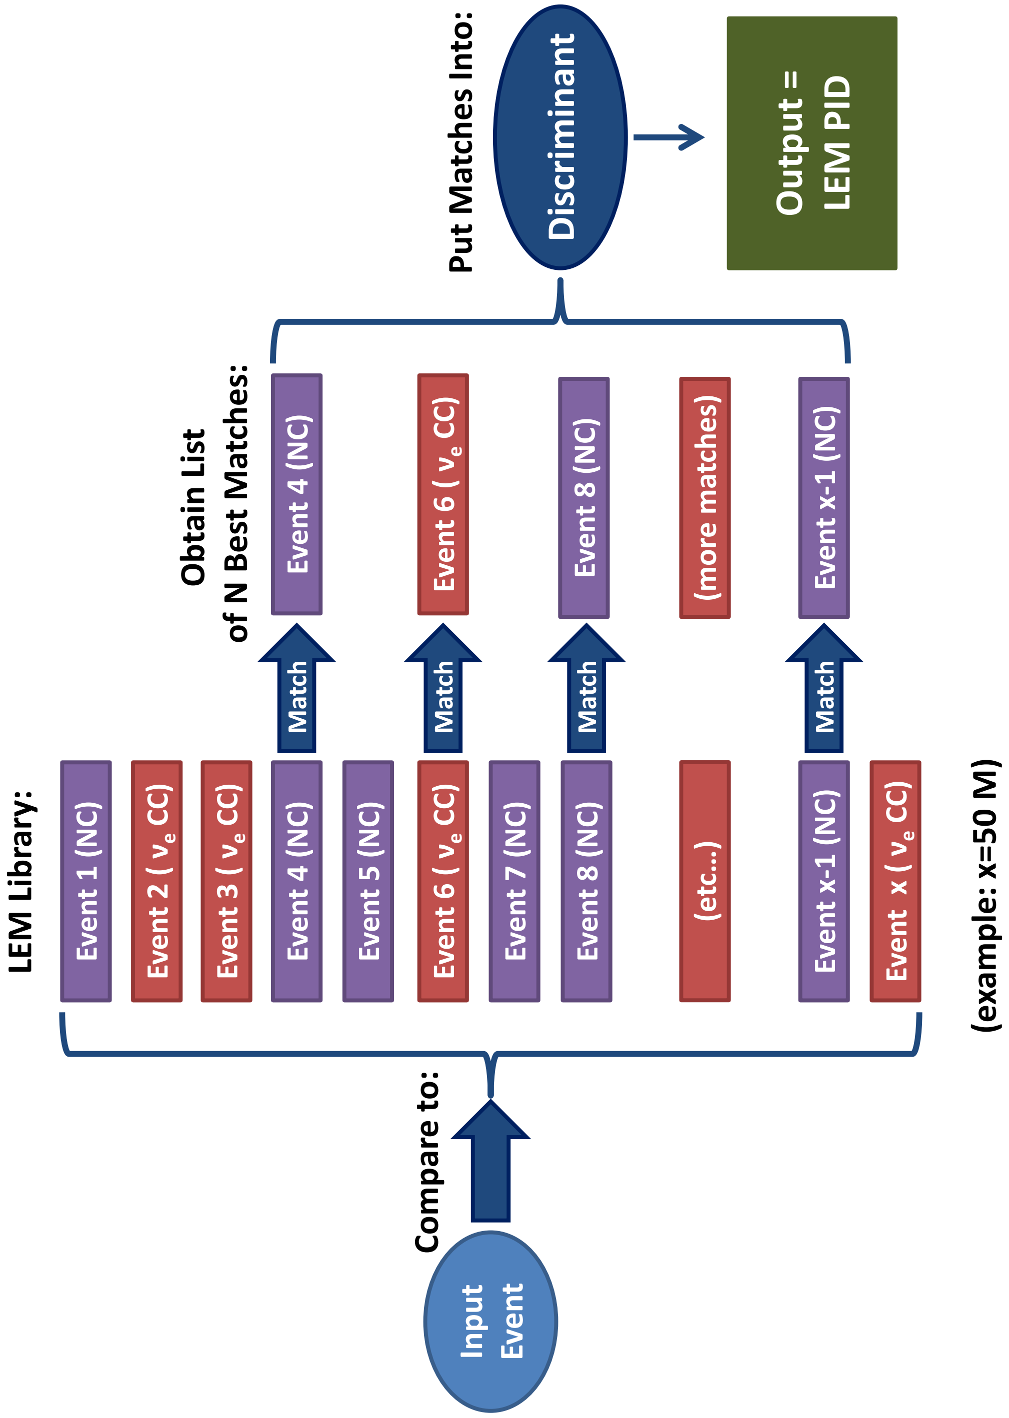

Figure 5.2: A flowchart showing the means by which LEM accomplishes particle identification. An input candidate event (Data or MC) is compared with a large library of Monte Carlo signal (ν_e CC) and background (NC) events. Out of this comparison, N best matches are determined. These N best matches are then put into a discriminant. The output of this discriminant is a single number, the LEM particle ID, which states how likely the input event is to be signal or background.

5.1.1 Context of Thesis Work

The development of the LEM PID is one of the primary focuses of this thesis. The basic LEM algorithm was developed in 2006-2008 [91] [92]. A first version of the PID[1] was used as a cross-check in the 2009 MINOS ν_e appearance analysis [85]. By the time the work in this thesis began, the basic elements of LEM were in place, including an early version of the MC library, the matching algorithm, and the PID discriminant variables. The author's work in this thesis has been to turn this early version into a fully optimized PID. This includes, most significantly, the development of a new multivariate discriminant to interpret matching data. Work has also been done to develop the pre-matching steps, including further development of event compactification and matching precuts. This thesis has also made major advances in quantifying and assessing the behavior of LEM, including studies of matching quality, positional bias, etc. The end product of this work is a fully optimized particle ID, with a well-understood behavior, which can reliably and effectively be used for a $\nu_\mu \rightarrow \nu_e$ search.

5.2 The LEM Library

The LEM Library consists of approximately 50 million simulated Far Detector MC events. These events were produced in the Low Energy beam configuration. The raw Monte Carlo was then reconstructed using the most recent version of the MINOS software. This entire process was carried out at the Rutherford Appleton Laboratory on the LCG Grid, and on the Fermilab Computing Grid, over the course of approximately six months.

Of these 50 million events, approximately 20 million were ν_e CC events, and approximately 30 million were Neutral Current events. The relative composition of the ν_e CC and NC events was chosen in a study performed for the first version of the LEM PID [92]. Rejection sampling methods were used to calculate the sensitivity of a simple version of the LEM algorithm with different proportions of the two event types. The ideal library composition was found to have a fraction of 0.3-0.4 ν_e CC events. ν_μ CC interactions were not included in the library, as the vast majority of ν_μ CC events are removed by the track cuts of the Preselection. Those ν_μ CC events which remain (with short embedded tracks) look very similar to NC hadronic showers [93]. Because LEM will be applied to both detectors, and then extrapolated to form a FD prediction, any mis-identification of ν_μ CC events will be an issue of performance, not accuracy. While the addition of

---

[1] Developed largely by Pedro Ochoa [50].

ν_μ CC library events could have potentially further improved background rejection, this decision would have increased the processing time prohibitively. The final choice of 50 million NC and ν_e CC library events almost doubled the previous library size.

Not all of the events in the LEM library are of interest to this analysis. Therefore, a series of preselection cuts are applied to the overall library to only retain those events relevant to the ν_e CC-like sample. The following events are kept:

- Events are within a ND and FD Fiducial region;
- $5 < N_{\text{planes}} < 30$;
- $7 < N_{\text{strips}} < 80$;
- $N^V_{\text{strips}} > 2$;
- $N^U_{strips} > 2$;
- $50 < Q_{tot} < 3000$ PE

where N_{planes} is the number of planes traversed by an event, N_{strips} is the number of strips hit by an event, and N^U_{strips} and N^V_{strips} are the number of strips hit in the U and V views respectively. Q_{tot} is the total summed pulse-height of an event in PEs. These preselection cuts both speed up processing time and help maintain match quality. Following the cuts, approximately 24 million reconstructed ν_e CC and NC events remain, which constitute the LEM library.

5.3 Preparing Events for Matching

Before matching occurs, several stages of cuts and formatting must be applied to ensure that both the library events and the input events can be compared on the same footing in the two MINOS detectors.

5.3.1 Strip Charge Units

When performing the LEM matching comparison, one must choose a unit of charge in which to store event strips. The LEM matching algorithm assumes Poisson-based photon statistics. In terms of quantities in the MINOS calibration chain, ADCs which have been converted into PEs (via the photoelectron chain, see Section 3.4) are a natural unit in which to compare library and input event hits. Converting ADCs into PEs requires knowledge of the PMT gain calibration. The PMT gains are only known to within ± 5%, with a random channel-to-channel variation of 7% [94]. When PEs were used for an earlier version of LEM, the resulting topological uncertainty resulted in an 8.5% systematic effect on the background prediction [95].

In choosing the strip charge units for a new analysis, several alternative approaches were considered [96]. The method ultimately chosen relies instead on a quantity taken from the Energy chain, the SigCor.[2] At this point in the calibration chain, an event's hits have been corrected for linearity and drift, as well as strip-to-strip response differences. Using SigCors avoids the gain uncertainties in the ADC→PE conversion which lead to a high systematic error.

In order to use a Poisson statistics-based matching metric, it is necessary to scale the SigCor charges to the size of the earlier PE-based pulse height. To determine this scaling, the total number of SigCors and PEs were calculated separately for the Near and Far Monte Carlo. The ratio of (Total SigCors)/(Total PEs) was then then used as a scaling factor to convert the SigCors into effective PEs [97].

5.3.2 Compacting Events

Before they are matched to one another, input and library events must undergo a so-called Event Compacting process. The events are reduced in size, given a standard position in the detector, and generally prepared to be compared on the same footing. The compacting steps are shown graphically in Figure 5.4, and are as follows:

1. A 3 PE cut is applied to the strips in the event. Almost all hits below 3 PE are poorly modeled low pulse height hits or hits caused by cross-talk. The reconstruction already applies a 2 PE cut to mitigate the effects of this mismodeling. The stricter LEM 3 PE cut further reduces detector-specific differences in topology, as cross-talk is known to affect the two detectors differently [82].
2. The charge-weighted (centroid) mean of the event is calculated. This is done in such a way as to avoid bias from stray high energy hits at the edge of an event which may accidentally have been included during reconstruction. In both the u and v views, the strips are first looped over to determine the ± 4 strip window with the highest summed charge. The centroid is then calculated using the strips within this window, with the outliers having been disregarded.
3. Attenuation corrections are applied to each of the strips in the event, using the centroid mean as the reference position. These corrections get rid of position and detector-dependent differences in light level, and will be described in detail in the next section.
4. The event is given a new position, so that its centroid mean position is at generic position strip 100, plane 100. This allows the events to be compared in a position-

---

[2] See Figure 3.12 for the location of SigCors on the Energy Calibration chain.

independent manner.

5. The strips are regrouped around the centroid according to the pattern in Figure 5.3. While the information at the center of the event is largely preserved, strips are increasingly consolidated as distance from the centroid increases.

The purpose of this compacting is twofold. First, it has a practical purpose in reducing the size of events at matching time. This both decreases the memory requirement, and greatly speeds up the matching process. Second, the process also improves the quality of the matches themselves, by making the procedure less sensitive to random spatial fluctuations. Most of the topological information in a ν_e Charged Current event tends to be confined to the compact electromagnetic shower at the core of the event, unlike the more diffuse Neutral Current and ν_μ CC hadronic showers. Therefore, hits at the outer edges of an event can be grouped together to de-emphasize their contribution in matching, with no reduction in performance. Studies have found that the particle ID's sensitivity to ν_e CC detection decreases by $\sim 7\%$ when this strip grouping is not used [92]. While an infinite library would be able to match to every single pattern, the strip grouping simplifies event toplogy and permits the use of a smaller library to sample all event types.

The centroid mean calculation in Step 2 was previously very sensitive to high pulse-height outlier hits. These stray high pulse-height hits are often found at the edge of ND Data input events, and result from imperfect separation of the numerous interactions in each snarl. An older version of the algorithm occasionally located the centroid mean in a position outside the main body of the event, causing the strips to be incorrectly grouped. This failure occurred in about 2% of all cases. The new method, used in step 2, ensures that the centroid mean is found in the main body of the event, where the largest number of strips are located. The updated version of LEM also truncates the strip grouping pattern after 16 strips, to ensure that these outlying high pulse-height hits are not incorporated into the event at matching time.

5.3.3 Charge Attenuation Corrections

By using raw strip hit information instead of reconstructed tracks and showers, the LEM PID makes itself largely immune to changes or uncertainties in the reconstruction code. This also means that the SigCor variable used to store LEM strip hits does not have the standard calibration attenuation corrections. Figure 5.5 shows the amount of light produced by muons from ν_μ CC events (with tracks spanning > 30 planes) passing through different parts of a scintillator strip, in the two detectors. In the Far Detector,

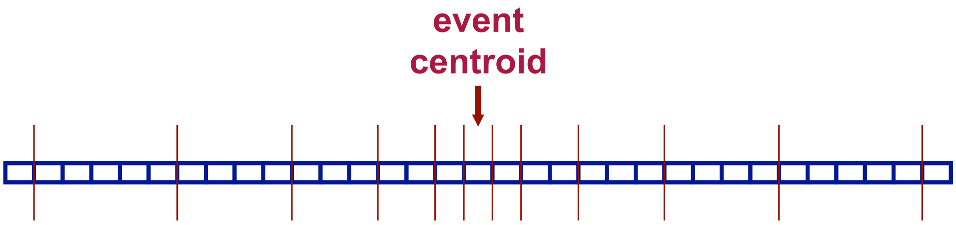

Figure 5.3: The pattern used for compacting event strips, as shown for the strips in an individual plane in an event. Strips at the center of the event are left largely untouched, in order to preserve the strip information in the core of the event. As distance from the centroid mean increases, however, strips are compacted together in increasing numbers, so that the contribution of outlying hits to the matching is greatly reduced.

the total light yield of a muon is 15% lower if it passes through the center of a strip rather than an edge. Additionally, the light level in the Near Detector is approximately 50% lower than in the Far Detector, due to the single-ended readout of strips.

Consequentially, the light level of individual events will be dependent on position. All of the LEM library events are Far Detector MC events. One could in theory attempt to match input events to library events in the same region in either of the two detectors. Producing a library large enough to match both topology and position, however, while still maintaining match quality, would be impossible. With the current LEM library, these light level differences can cause an input event to match to a different set of LEM library events depending on its position and detector, resulting in reduced performance and potential systematic bias. LEM input and library events must therefore be "calibrated" with a set of LEM attenuation corrections.

The process by which these attenuation corrections were applied is described in detail in reference [92] and is summarized briefly here. A phenomenological fit to the data in Figure 5.5 found that light attenuation in MINOS can be described by the following expression:

$$a_w(d_w) \simeq \frac{2}{3}e^{-\frac{-d_w}{7.05}} + \frac{1}{3}e^{-\frac{-d_w}{1.05}} \tag{5.1}$$

where d_w is distance (in meters) traveled along the WLS fibre. Light traveling through the clear plastic cable also attenuates in a similar way (although with a single, longer attenuation distance). In the Far Detector Data, raw light Q_{raw} is read out at both the east and west ends of the strip: $Q_{raw} = Q^E_{raw} + Q^W_{raw}$. Both Q^E_{raw} and Q^W_{raw} are corrected for attenuation. The corrected sum, $Q = Q^W + Q^E$ is then reweighted by the average

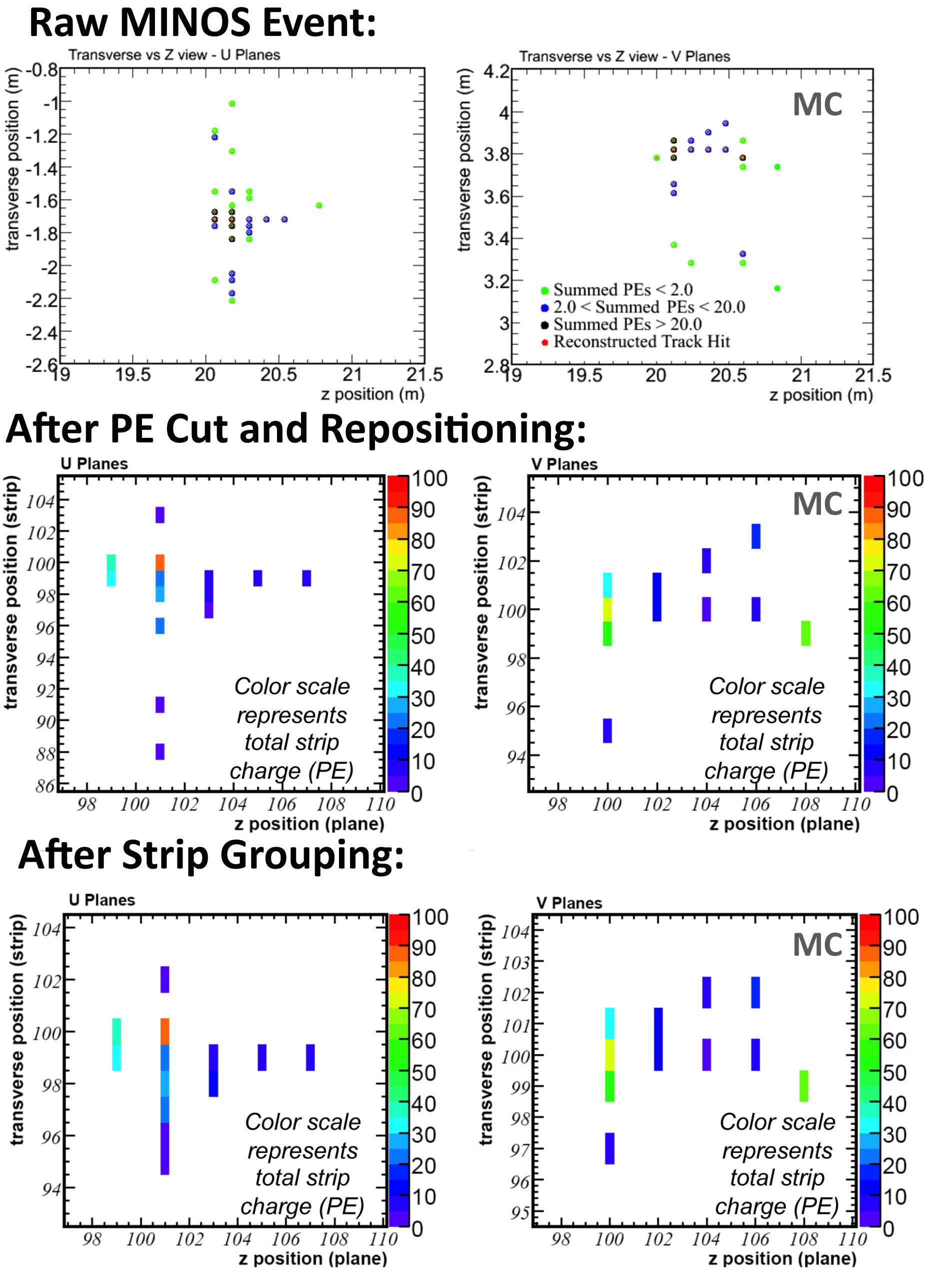

Figure 5.4: The event compacting process, for a single MINOS MC event. The top image shows the raw event, prior to any compacting. The second image shows the event following the 3 PE cut and attenuation corrections, with the even centroid having being relocated to plane 100, strip 100. The bottom image shows the fully compacted event, after strip grouping has been applied as well. Source of images: [92].

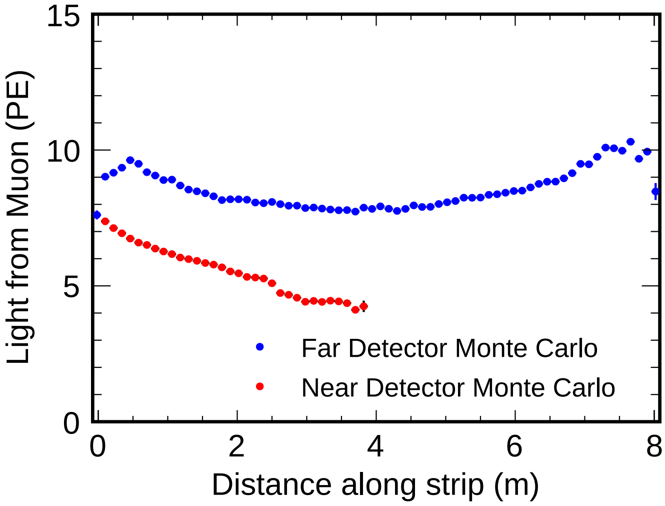

Figure 5.5: Light levels for charge deposited by long beam muons in the Near and Far Detectors. Distance along strip refers to the muon's location in the detector, i.e. distance from the West readout. The charge is collected at both strip ends for the Far Detector (blue), and at only the West end for the Near Detector (red). Data provided by Pedro Ochoa [92].

FD attenuation factor $\langle A^{FD} \rangle = \langle \frac{Q_{raw}}{Q} \rangle$ to match the real average charge deposition. A similar, single-sided correction is applied to the Near Detector. This ND correction must also take into account the reflectivity R of the non-readout strip side. The attenuation-corrected Near Detector light is then scaled to the same average light level as the Far Detector. The effects of this calibration on both the Near and Far Detectors can be seen in Figure 5.6. The light levels are now essentially flat in the Fiducial regions of both detectors, and events can be compared on the same footing, without bias.

5.4 Matching LEM Events

5.4.1 Deciding which events to match

Although the library preselection cuts and compacting procedure reduce some of the time required for matching, comparing each input event to the full library would still be computationally intensive. To further optimize the matching, a relaxed version of the standard ν_e analysis Preselection is also applied to all input events. These cuts remove events outside of the ν_e oscillation energy range, and events which have a large number of reconstructed track planes and are therefore likely to be ν_μ CC events. The following

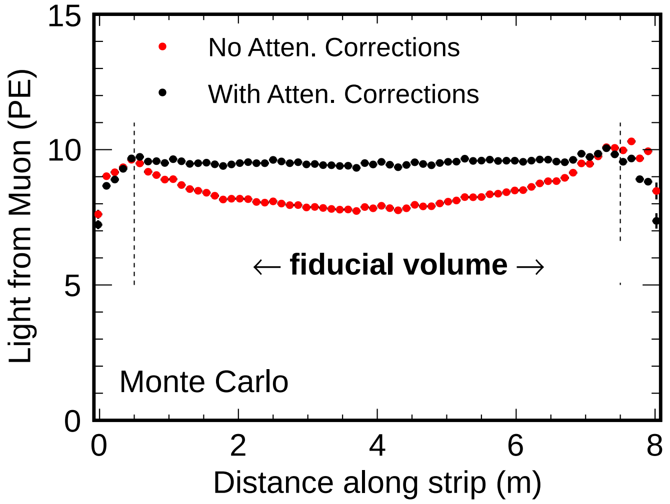

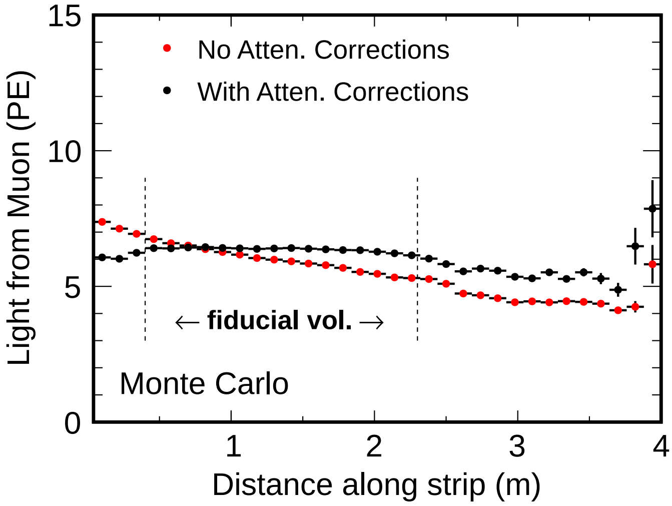

Figure 5.6: Light levels in the Far (top) and Near (bottom) detectors, before and after attenuation corrections. Distance along strip refers to the muon's location in the detector, i.e. distance from the West readout. The charge is collected at both strip ends for the Far Detector (top), and at only the West end for the Near Detector (bottom). The red curve shows the light yield prior to attenuation corrections, and the black curve following attenuation corrections. The analysis Fiducial region is also indicated. Data provided by Pedro Ochoa [92].

events are retained:

- 0.5 GeV < $\mathrm{E}_{\mathrm{reco}}$ < 12 GeV
- $N^{\mathrm{track}}_{\mathrm{planes}} < 30$

During the matching process, a series of "compatibility" cuts are also applied to prevent the comparison of very dissimilar input and library events. Events A and B will not be compared if they have a drastically different number of planes (a large difference in the number of planes N^{A}_{planes} and N^{B}_{planes}), more than a 20% difference in their total compactified and attenuation-corrected charges (Q^{A}_{total} and Q^{B}_{total}), or more than a 20% difference in the number of compacted strips (N^{A}_{strips} and N^{B}_{strips}):

- $|N^{A}_{\mathrm{planes}} - N^{B}_{\mathrm{planes}}| \leq x$ where $x = 4$ if $N^{avg}_{\mathrm{planes}} = \frac{N^{A}_{\mathrm{planes}}+N^{B}_{\mathrm{planes}}+1}{2} > 20$, $x = 3$ if $N^{avg}_{\mathrm{planes}} > 15$ and $x = 2$ otherwise,
- $2|\frac{N^{A}_{\mathrm{strips}}-N^{B}_{\mathrm{strips}}}{N^{A}_{\mathrm{strips}}+N^{B}_{\mathrm{strips}}}| \leq 0.2$,
- $2|\frac{Q^{A}_{total}-Q^{B}_{total}}{Q^{A}_{total}+Q^{B}_{total}}| \leq 0.2$,

Any input event which either fails the initial preselection or fails the compatibility checks for all library events is assigned a LEM PID value of -1. Following the standard ν_e Preselection cuts, only $\sim 0.2\%$ of input events have LEM=-1 [98].

The Matching Algorithm

All input events which pass these cuts undergo LEM matching. The matching procedure quantifies how likely two separate patterns of (effective) photoelectrons were to have been produced by the same energy deposition. The detector response can be described using Poisson statistics. A given true energy deposition has some probability of producing n photoelectrons in strip i. This probability is described as $P(n^i;\lambda)$, the probability of seeing n^i photoelectrons given a detector response described by a Poisson distribution with a mean of λ photoelectrons (these probabilities are calculated prior to the matching and stored for run time). The likelihood $S^i(n^i_A, n^i_B)$ that two separate photoelectron pulse heights n^i_A and n^i_B were produced by the same energy deposition is found by taking the product of the two Poisson probabilities and integrating over all possible values of λ:

$$S^i(n^i_A, n^i_B) = \int_0^\infty P(n^i_A;\lambda)P(n^i_B;\lambda)d\lambda \qquad (5.2)$$

Moving from strips to full events, the probability of two collections of strips A and B being produced by the same energy deposition is found by taking the product of S^i over all strips, or more conveniently, the sum of $-\ln(S^i)$:

$$-\mathcal{L} = -\sum_{i=1}^{\mathrm{N}^{\mathrm{strips}}} \ln\left(\int_0^{\infty} P(n_A^i;\lambda)P(n_B^i;\lambda)d\lambda\right) \tag{5.3}$$

The log likelihood $-\mathcal{L}$ will vary based on the size or number of strips in an event. For convenience, the self-likelihood of the event is subtracted out from the matches as an arbitrary constant:

$$-\Delta\mathcal{L} = -(\mathcal{L}^{\mathrm{lib}} - \mathcal{L}^{\mathrm{self}}) = \mathcal{L}^{\mathrm{self}} - \mathcal{L}^{\mathrm{lib}} \tag{5.4}$$

where the self-likelihood $\mathcal{L}^{\mathrm{self}}$ is defined as follows (if event A is the input candidate event):

$$-\mathcal{L}^{\mathrm{self}} = -\sum_{i=1}^{\mathrm{N}^{\mathrm{strips}}} \ln\left(\int_0^{\infty} P(n_A^i;\lambda)P(n_A^i;\lambda)d\lambda\right) \tag{5.5}$$

A smaller value of $-\Delta\mathcal{L}$ corresponds to a better match.

This matching procedure runs over the complete set of Preselected library events. For every input event, $-\Delta\mathcal{L}$ is calculated three times. First, the matching is carried out with both compacted events' centroids centered at Plane 100, Strip 100. Then, the matching is carried out twice more, first with the input event's position shifted by $+1$ plane, and then by -1 plane. The smallest value of $-\Delta\mathcal{L}$ is chosen as the likelihood for that particular library event. It has been found that approximately two thirds of input events have their best match without the ±1 plane shift [92]. Like the strip grouping procedure, this process reduces the sensitivity of the matching to random spatial variations and increases the effective statistics of the library. Because strip compactification has already been applied to reduce the effects of these fluctuations in the transverse direction, this shifting procedure is not necessary for the strips.

As the matching proceeds, a running list is kept of the 200 library events with the lowest value of $-\Delta\mathcal{L}$, i.e. the best matches. In order to reduce processing time, the matching process in Equation 5.3 is terminated once an event reaches a value of $-\Delta\mathcal{L}$ which is too large to be one of the 200 best matches.

5.4.2 Match Quality

Several studies were conducted to characterize the behavior and accuracy of the LEM matching algorithm for this thesis. One such study is simply to look at the magnitude of $-\Delta\mathcal{L}$ for different types of events. Figure 5.7 shows the value of $-\Delta\mathcal{L}$ for the 50 best matches for Preselected input events, both for NC matches and ν_e CC matches. This plot has been broken down into individual input event types. FD MC ν_e input events match more frequently to ν_e CC library events, while NC input events match ν_e CC library events both less frequently and with a higher average value of $-\Delta\mathcal{L}$.

Another good diagnostic of the matching procedure is how well the reconstructed energy of an input event corresponds to the true energy of its ν_e CC best matches[3]. A discrepancy between match and input event energy would indicate some inconsistency or bias in the matching process, or some problem with the library sample. As can be seen in Figure 5.8, the energy of FD MC signal and background input events corresponds well with the ν_e CC match true energy. This relationship is at least in part a result of the LEM matching compatibility cuts, which ensure the similarity of events being compared. Notably, the ND MC background events also match to similar energies of ν_e CC library events. This agreement between energies is a good cross-check of the matching process. A similar study examining the relationship between input and matching events in terms of their hadronic properties will be presented later (see Section 6.2).

5.5 Designing a LEM discriminant

The final output of the matching is a list of 200 matches, with basic information such as whether those matched library events were ν_e CC or NC, their true energy, their inelasticity y, etc. The next stage in the LEM particle ID algorithm is to use this information to form an optimal PID discriminant.

5.5.1 Constructing variables from the LEM Matches

The first step in this process is to decide how many matches will be retained, and to construct variables from those matches. Studies done for the first version of LEM found that using 50 matches produced an optimal PID discriminant. Less than 50 matches contained too little information, and match quality degraded above 50. From these 50

---

[3] This relationship does not hold for Neutral Current matches, because the visible energy of a NC event is only a fraction of its total energy.

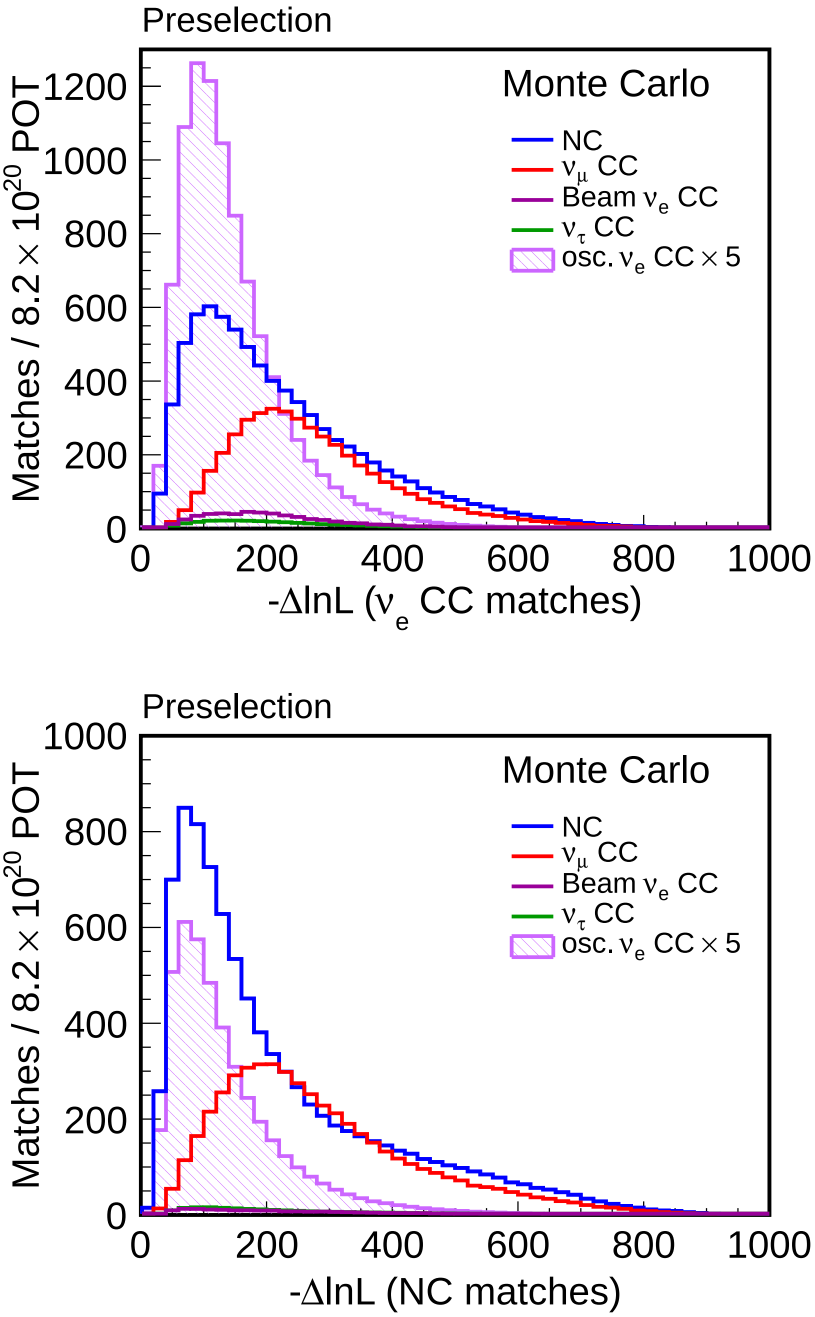

Figure 5.7: Distribution of $-\Delta\mathcal{L}$ for the 50 best LEM library matches. The matches are taken from input events passing the standard Preselection. The above two plots are broken into ν_e CC library matches (top) and NC library matches (bottom). This plot is for Far Detector MC only, and θ_{13} is at the CHOOZ limit.

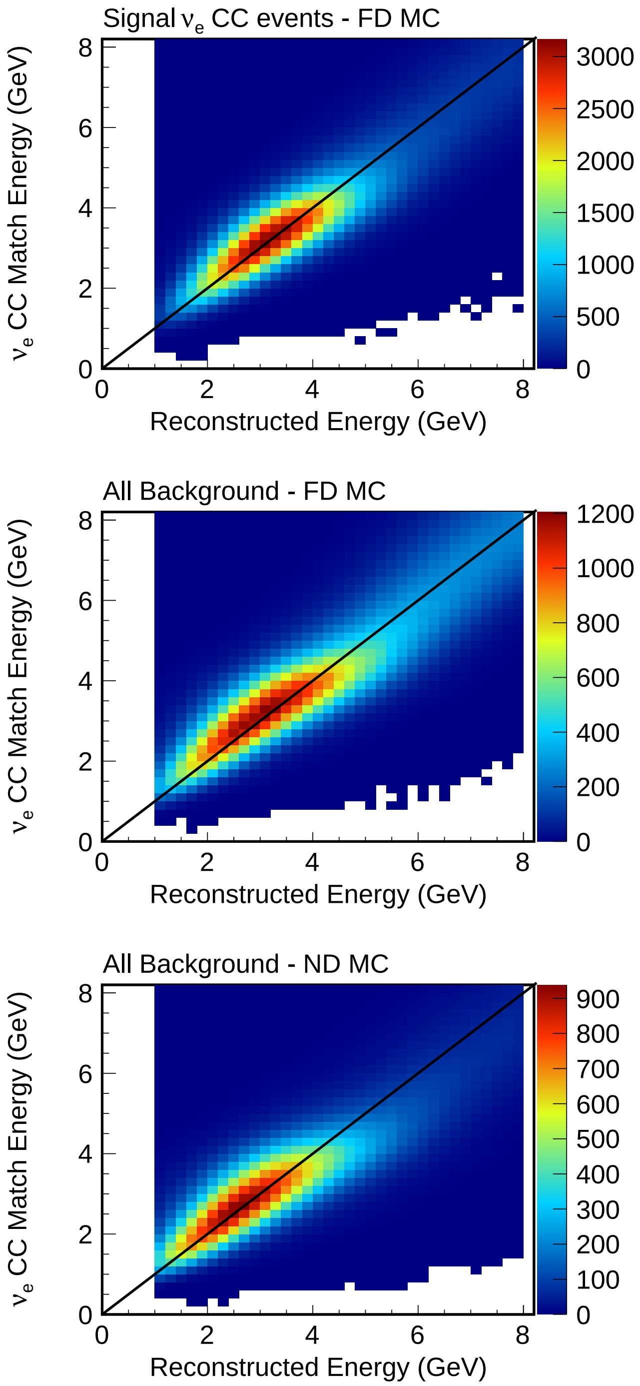

Figure 5.8: Comparison of reconstructed input event energy versus true LEM event match energy, for those of the 50 best matches which are ν_e CC matches, for Far Detector MC events at the Preselection level. The top plot shows this relationship for FD signal ν_e CC input events. The middle plot shows the same for FD MC background input events, and the bottom plot for ND MC background events. The black line represents x=y.

best matches, three separate quantities were then constructed to answer three separate questions:

1. How many of the 50 best matches were signal (ν_e CC) events?
2. How electromagnetic-like were the ν_e CC match showers?
3. How well did the charge distribution of the input and the ν_e CC matches overlap?

These characteristics are described with the following variables:

f_{50} - The fraction of matches which are ν_e CC: This is the simplest of the three particle ID variables, quantifying how often the input event matched to signal. If a larger fraction of the 50 best matches were ν_e CC signal events, instead of NC events, the input event was more likely to be signal. A higher value f_{50} therefore corresponds to a more signal-like event.

y_{50} - The mean y of the ν_e CC Matches: This variable quantifies how electromagnetic-like an event was, by calculating the mean inelasticity (hadronic y) of its ν_e CC best matches. As a reminder, y is defined as follows:

$$y = \frac{E_\nu - E_\ell}{E_\nu} = 1 - \frac{E_\ell}{E_\nu} \tag{5.6}$$

where E_ν is incoming neutrino energy and E_ℓ is the amount of energy transferred to the lepton (in this case, the electromagnetic shower). The rest of the energy is transferred to the hadronic system. The signature of a ν_e CC event is an EM shower, so that library events with a low value of y will look most like signal. ν_e CC library events with a higher value of y look more similar to background. A signal event will therefore tend to match to low-y ν_e CC library events more often, resulting in a y_{50} closer to 0. Hadronic shower background events can also have a small electromagnetic-like component from π^0 decay, but will tend to match to higher-y events, resulting in a larger value of y_{50}.

q_{50} - The average fractional matched charge for ν_e CC matches This final variable quantifies the quality of the matching between the input event and the ν_e CC library events, by looking at Q_{frac}, the fraction of the strip charges which overlap between the input event and a ν_e CC library event:

$$Q_{frac} = \frac{Q_{\mathrm{matched}}}{Q_{\mathrm{matched}} + Q_{\mathrm{unmatched}}}, \tag{5.7}$$

> The PID variable q_{50} consists of the average value of Q_{frac} for the ν_e CC best matches. If the ν_e CC matching was good on average (i.e., if q_{50} is high), the event is more likely to be signal.

The Preselection-level distributions for these LEM particle ID variables for specific event types are shown in in Figure 5.9. In each case the expected signal is shown at the CHOOZ limit ($\sin^2 2\theta_{13} = 0.16$).

Finally, because the signal and background distributions each have distinct energy distributions (see Figure 5.10), reconstructed energy (in GeV, in the EM scale) is also incorporated as a fourth PID variable.

5.5.2 Choosing a LEM discriminant

One of the major tasks undertaken in this thesis is to optimize the performance of LEM, by combining these 3 PID variables and reconstructed energy into a multivariate discriminant. Numerous techniques were studied for this analysis, including a simple set of cuts and a likelihood-based method employing two-dimensional probability distribution functions. However, the choice of method is constrained by two factors. First, as is shown in Figure 5.11, the three LEM variables are highly correlated, both for ν_e CC signal events and for background events. These correlations arise due to the fact that the three LEM PID variables are not fully independent. An event which is very electromagnetic-like ($y_{50} \rightarrow 0.0$), for instance, will consequentially mostly match to ν_e CC events ($f_{50} \rightarrow 1.0$). This necessitates the use of a discriminant which can adequately take correlations into account.

Second, the choice of discriminant is also limited by the amount of available Far Detector Monte Carlo statistics. Less than 300,000 Far Detector Monte Carlo events are available for both signal and background. While this is an adequate number of events to produce a good discriminant, it does not allow for, say, the population of a four-dimensional probability distribution function.

5.5.3 Neural Networks

A particle ID for LEM must therefore be able to handle highly correlated variables and a relatively limited body of statistics. An artificial neural network discriminant meets both of these requirements. Broadly speaking, an artificial neural network consists of a set of artificial neurons which mimic those found in the brain. Each neuron reacts to stimuli in some characteristic way and in turn communicates with other neurons via a

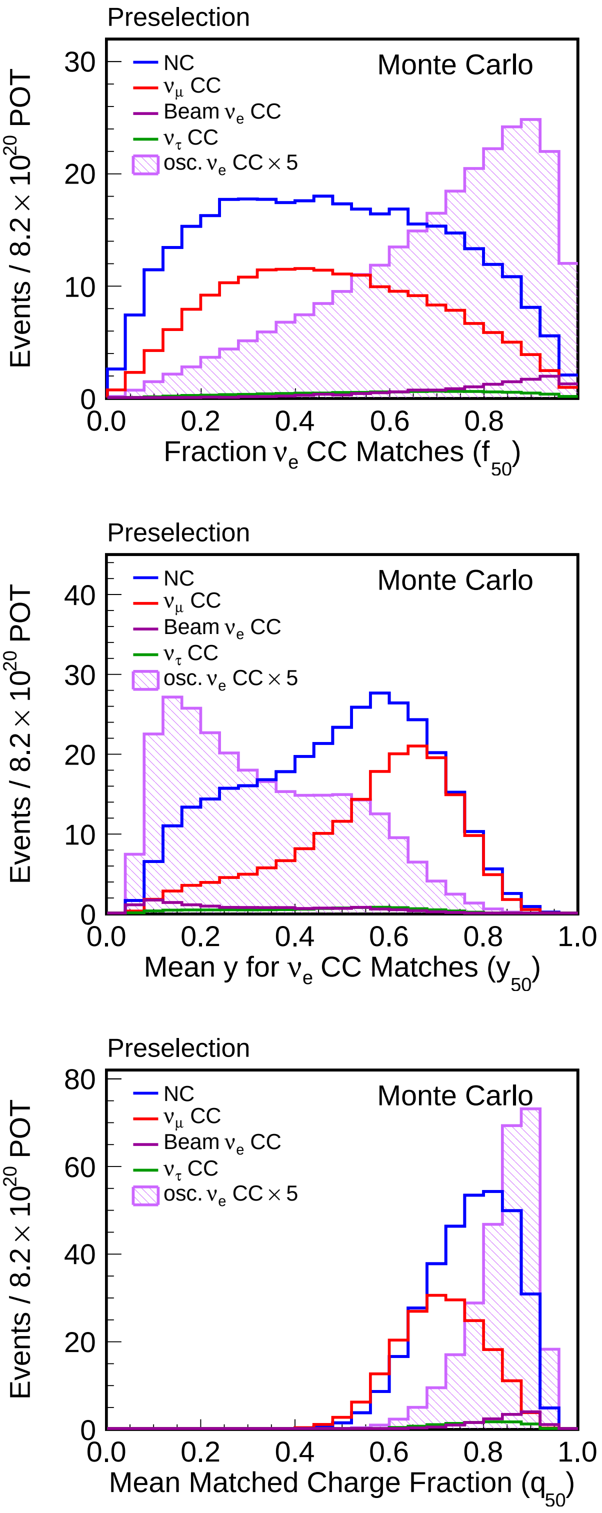

Figure 5.9: The LEM variables for Preselected Far Detector Monte Carlo events in the MINOS detector: f_{50} (top), y_{50} (middle), and q_{50} (bottom). Signal is at the CHOOZ limit, and has been scaled × 5 for ease of viewing.

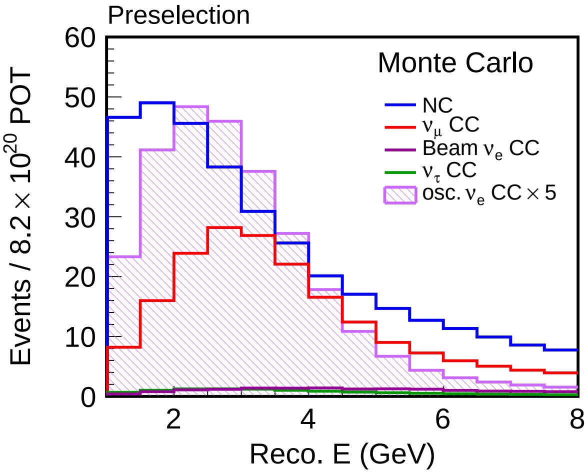

Figure 5.10: Reconstructed energy (in GeV) Preselected Far Detector Monte Carlo events in the MINOS detector. This variable is the fourth input to the LEM PID. Signal is at the CHOOZ limit, and has been scaled × 5 for ease of viewing.

network of weighted connections. Using the response of this network of neurons, the system can be taught to perform a specific task, e.g. to discriminate between a ν_e CC event and a Neutral Current event. A multidimensional space of input variables can then be mapped to a single one-dimensional space, the discriminant output variable.

Every neural network has a characteristic architecture, or an arrangement of neurons and a set of weighted connections. Each neuron is given an activation $a = \sum_i w_i$, where w_i is the ith weighted input to the neuron. The neuron has a characteristic response to the input, $y = f(a)$. The network is also given a so-called learning rule, which determines how the system evolves over time. Using these components [102], a neural network can be "trained" to perform various tasks, such as classifying events into two categories. The type of neural net discussed in this document is a supervised network, in which the network is trained using examples of input events (e.g., FD MC events with values of f_{50}, y_{50}, q_{50}, and energy) which have a known target classification (signal or background).

A Multilayer Perceptron (MLP) style of neural net was employed as a discriminant for the LEM PID. This discriminant was trained using the TMVA [101] multivariate analysis package. In a MLP neural net, the neurons are organized into layers: an input layer, one or more central "hidden" layers, and an output layer (in this case, a single neuron). For the neuron activation function $f(a)$, one can choose between a variety of different sigmoid functions [103].

The neural net is trained using the learning rule known as back-propagation, which

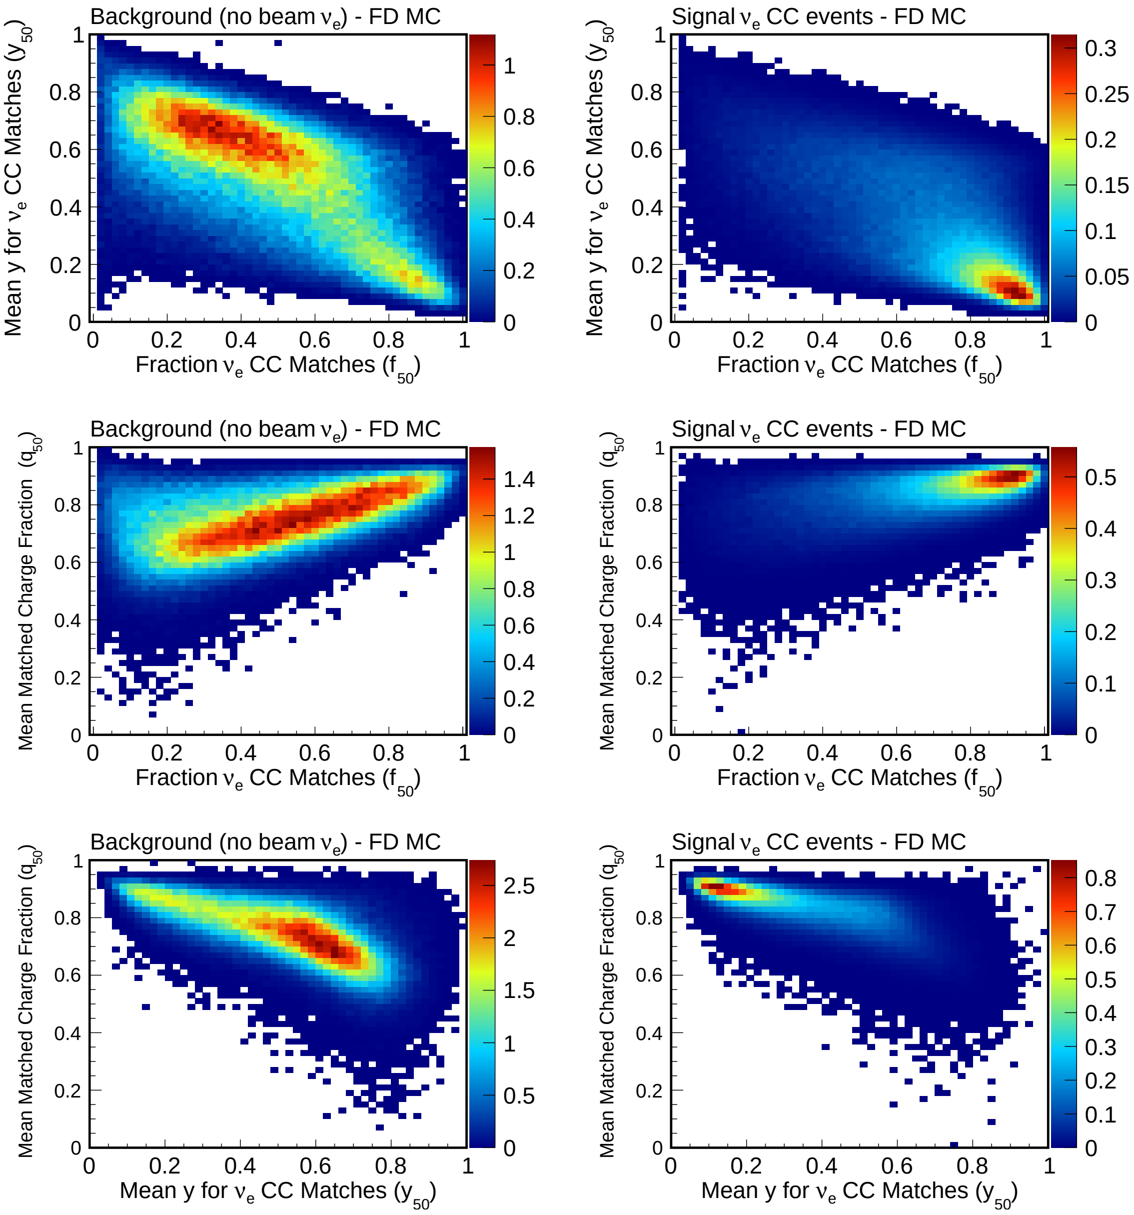

Figure 5.11: Correlations between the three LEM PID variables for background events (left hand images: NC, ν_μ CC, and ν_τ CC, with no beam ν_e events) and signal ν_e CC events (right hand). Events are Preselected Far Detector MC.

works as follows. The output of a neural network with a single hidden layer and single output node can be written as:

$$y_{ANN} = \sum_{j=1}^{m} y_j^{(2)} w_{j1}^{(2)} = \sum_{j=1}^{m} f(\sum_{i=1}^{n} x_i w_{ij}^{(1)}) w_{j1}^{(2)}, \tag{5.8}$$

where n is the number of input neurons, m is the number of hidden neurons, x_i is training input i, y_j is the output of neuron j, $w_{ij}^{(i)}$ is the weight between input i and hidden neuron j, and $w_{j1}^{(2)}$ is the weight between hidden neuron j and the single output neuron. N training events are provided to the neural network. Each event a has input $\mathbf{x}_a = (x_1, \ldots, x_n)$. $y_{ANN,a}$ is calculated each time and compared to the target value $\hat{y}_a \epsilon \{1, 0\}$. From this output, a so-called error function is calculated, which states how close the output was to the target:

$$E_a(\mathbf{x}_a|\mathbf{w}) = \frac{1}{2}(y_{ANN,a} - \hat{y}_a)^2, \tag{5.9}$$

In order to train the neural net, this error function must be minimized. This is typically done by adjusting the weights in the direction in weight space where the value of E is falling the steepest (i.e., where the value of $-\nabla_w E$ is the largest). TMVA employs the “online learning” method, where the weights are adjusted after every single training event. The sample of training events is iterated through (potentially hundreds of times) until the error function has found a minimum, and the neural net is considered trained.

5.5.4 The LEM Neural Network

The LEM Multilayer Perceptron has three layers of neurons. First, there is a layer of four input neurons, one each for f_{50}, y_{50}, q_{50}, and reconstructed energy. After the input layer, there is one hidden layer with 9 nodes. Finally, there is a single output neuron, which in this case produces a number between 0.0 and 1.0 quantifying how likely an input event was to have been signal (1.0) or background (0.0). The architecture of the system was therefore 4 : 9 : 1. For the neurons themselves, a neuron response function of tanh(a) was chosen. A maximum number of 600 cycles was set for the training process.

The TMVA neural net was trained with a sample of FD MC events divided into signal (ν_e CC) and background (ν_μ CC, NC, and ν_τ CC). Signal and background events were weighted in the correct proportions for oscillation assuming the values in Table 4.2 (i.e., CHOOZ-sized signal), with other standard MC weights applied as well. The two samples were also normalized so that they contained roughly the same number of signal and background events.

Several variations on the above 4 : 9 : 1 neural net were also considered. A three-variable neural net omitting reconstructed energy was also attempted, and found to have a 15% reduced performance. A fifteen-variable neural net which incorporated both the ANN11 and LEM variables had comparable performance, to within 2% of the four-variable LEM, suggesting there was little additional sensitivity to be gained by combining the two PIDs. The 4 : 9 : 1 neural net with inputs f_{50}, y_{50}, q_{50}, and reconstructed energy was therefore chosen. The resulting distribution of the LEM PID for various event types can be seen in Figure 5.12.

5.6 Assessing the Performance of LEM

A simple PID cut will here be used to assess the relative performance of LEM, using a FOM calculated from a PORP prediction. In the final $\nu_\mu \to \nu_e$ search, a multiple bin shape fit will be used. The simple PID cut will be used to assess the basic PID performance and behavior.

5.6.1 LEM Performance with PORP and approximate error

Figure 5.13 shows the performance of a range of cuts for the LEM particle ID, with a CHOOZ-size signal. Background predictions are calculated using the PORP method. The three separate performance measures are as follows:

- **Signal/Background:** The sample purity, which rolls off around 0.95 due to very limited statistics in the high PID region.
- **Figure of Merit:** FOM, with statistical error only (Equation 4.9).
- **Figure of Merit:** super-FOM, with a 5% systematic error in addition to statistical error (Equation 4.10).

A maximum super-FOM of 4.03 is obtained with a cut of LEM > 0.76. At this point, 21.8 signal events and 27.4 background events are selected, giving a sample purity of 44.3%. Figure 5.14 shows the relative FOM and super-FOMs for LEM and ANN11 (with the same systematic error). ANN11 achieves a maximum super-FOM of 3.21 at ANN11>0.70, with a purity of 31.2%. Comparing the FOMs of the PIDs shows that LEM is capable of producing up to a 25% improvement in sensitivity over the previous PID technique, ANN11.

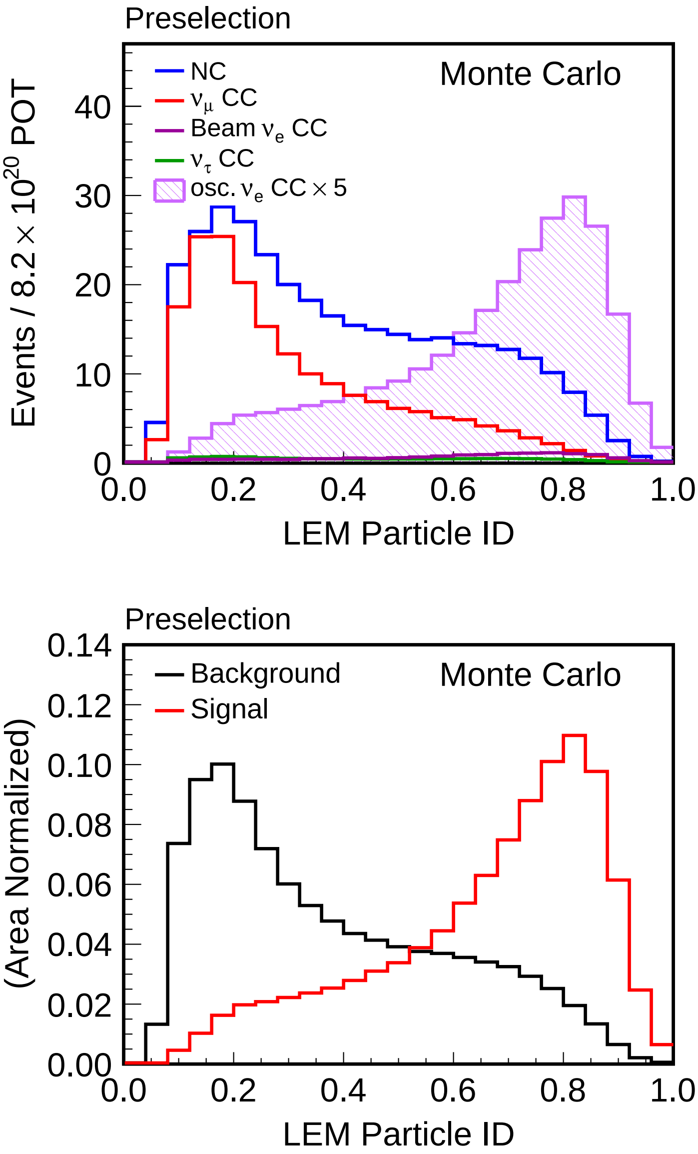

Figure 5.12: The LEM particle ID distribution for Preselected Far Detector Monte Carlo events. The ν_e CC events have been oscillated at the level of the CHOOZ limit. The top plot shows the particle ID distribution for separate event types. Signal has been scaled × 5 for ease of viewing. The bottom plot shows the signal distribution versus the total background distribution, with both distributions area normalized.

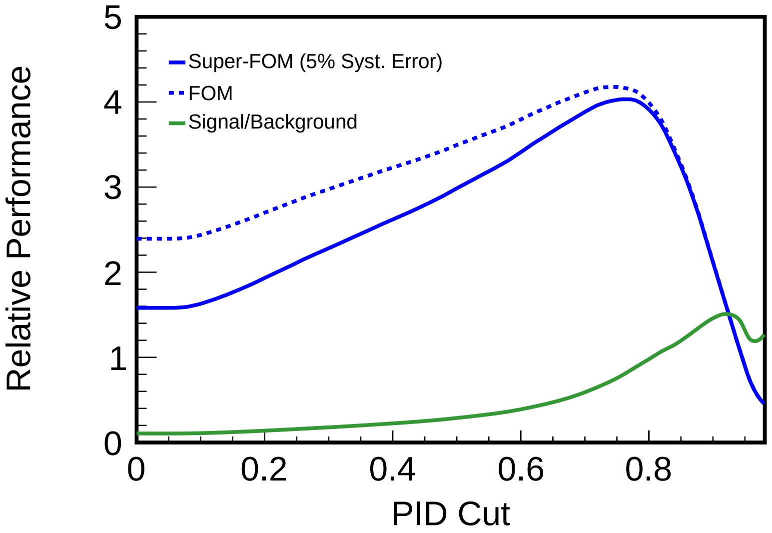

Figure 5.13: Relative performance of the LEM PID as a function of cut. Performance is shown as Signal/Background, Figure of Merit (FOM), and super-FOM with a 5% systematic error. All the quantities have PORP data-based corrections applied.

The efficiency of the LEM>0.76 cut for various background types is listed in Table 5.1. This cut has a 30.8% selection efficiency for Signal events relative to the Fiducial sample. This selection efficiency rises to 40.1% relative to the Preselection Sample. Either with or without the PORP correction, this cut rejects approximately 99% of all background events in the Fiducial sample. For the PORP-corrected case, 95% of the background in the Preselection sample is rejected. LEM therefore manages both to eliminate the vast majority of the background in the MINOS detector, while also preserving just under half of a potential signal sample.

5.6.2 LEM $>$ 0.70: Official Cut Performance and Efficiency

The numbers discussed in the previous section assume an approximate systematic error of 5% and a PORP decomposition. In reality, the systematic error on the LEM selection will both not be exactly 5% and will also vary with the choice of LEM cut. When the full systematic error is calculated, the optimal cut ends up being at LEM $>$ 0.70, with a 4.8% systematic error on the background. With PORP corrections, this cut selects 28.7 signal and 48.7 background events, for a purity of 37.1% and a super-FOM of 3.90.

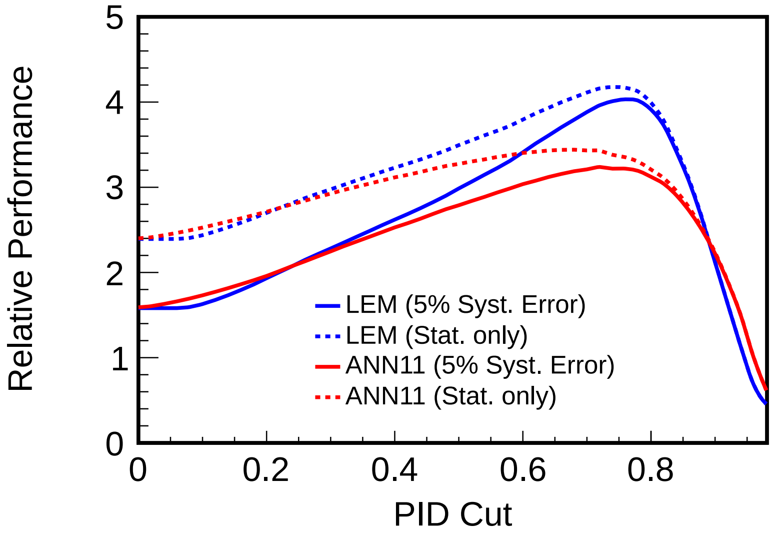

Figure 5.14: Relative performance of the LEM PID versus the ANN11 PID as a function of cut. Performance is shown for both PIDs as a function of FOM and super-FOM (with a 5% systematic error). All the quantities have PORP data-based corrections applied

Selection	osc. ν_e CC	Total bg.	NC	ν_μ CC	ν_τ CC	Beam ν_e
Fiducial vol.	70.9	3173.8	873.6	2238.1	19.5	42.6
Preselection	54.4	553.0	338.1	189.7	10.5	14.74
LEM (No PORP) $>$ 0.76	21.8	36.7	26.3	4.8	1.2	4.0
efficiency (vs. Fid.)	30.8%	1.1%	3.0%	0.2%	6.6%	9.4%
efficiency (vs. Presel.)	40.1%	6.6%	7.9%	2.5%	12.2%	27.2%
LEM (PORP) $>$ 0.76	21.8	27.4	18.8	3.3	1.3	4.0
efficiency (vs. Fid.)	30.8%	0.9%	2.2%	0.2%	6.6%	9.4%
efficiency (vs. Presel.)	40.1%	5.0%	5.6%	1.7%	12.2%	27.2%

Table 5.1: Number of Far Detector raw MC events selected by a LEM $>$ 0.76 cut, and the corresponding selection efficiencies relative to the Fiducial and ν_e Preselection selections. The LEM selection is shown both with and without PORP corrections.

This constitutes a 22% improvment in sensitivity over ANN11. For a PORP-corrected sample, LEM>0.70 has a 40.5% efficiency versus the Fiducial sample, and a background rejection of 98.5% (see Table 5.2). This cut will be assumed to be the optimum single-bin cut for all future discussions of the LEM particle ID in this document, unless otherwise stated.

Selection	osc. ν_e CC	Total bg.	NC	ν_μ CC	ν_τ CC	Beam ν_e
Fiducial vol.	70.9	3173.8	873.6	2238.1	19.5	42.6
Preselection	54.4	553.0	338.1	189.7	10.5	14.74
LEM (No PORP) > 0.70	28.7	61.7	44.5	9.5	2.0	5.74
efficiency (vs. Fid.)	40.5%	1.9%	5.1%	0.4%	10.3%	13.5%
efficiency (vs. Presel.)	52.7%	11.2%	13.2%	5.0%	19.1%	38.9%
LEM (PORP) > 0.70	28.7	48.7	33.9	7.0	2.0	5.74
efficiency (vs. Fid.)	40.5%	1.5%	3.9%	0.3%	10.3%	13.5%
efficiency (vs. Presel.)	52.7%	8.8%	10.0%	3.7%	19.1%	38.9%

Table 5.2: Number of Far Detector raw MC events selected with the official analysis LEM > 0.70 cut, and the corresponding selection efficiencies versus the Fiducial and Standard ν_e Preselection selections. The LEM selection is shown both with and without PORP corrections.

5.6.3 Events Selected by LEM

Certain types of neutrino interaction tend to look more ν_e-like, and are therefore more likely to be selected by the LEM PID. Figure 5.15 shows the LEM PID broken down by interaction type (DIS, QE, RES, COH) for both signal and background, while Table 5.3 lists the percentage of each type for both the overall Preselection sample, and for the LEM>0.70 selection. For signal events, the LEM PID preferentially selects QE events, as well as many RES events. DIS events overwhelmingly dominate the background sample, with small contributions from the other event types (including COH). Table 5.4 breaks down the selections by how many π^0s were present in the interaction. At the preselection level, 61.9% of background events contained at least one π^0. Following the LEM>0.70 cut, this percentage rises to 84.9%. This is not surprising, as LEM is expected to select hadronic shower events which have a larger EM component from π^0 decay, and therefore look more ν_e CC-like.

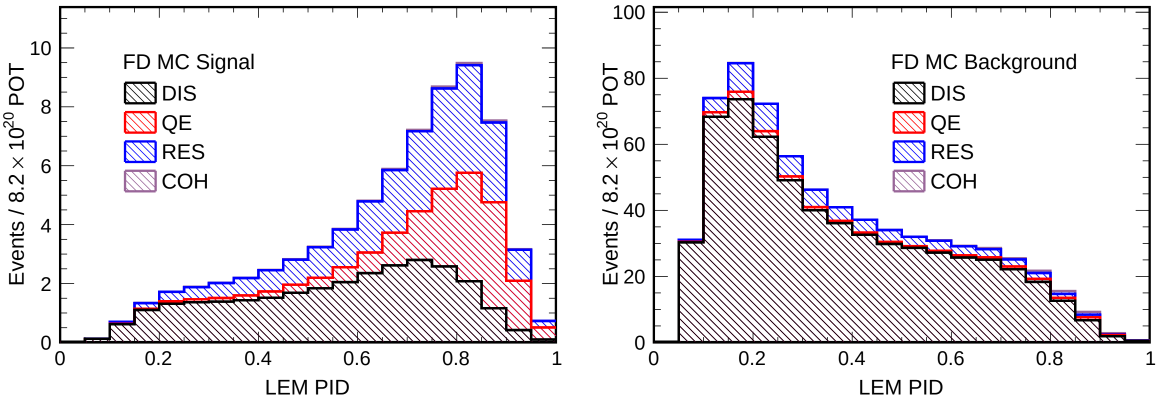

Figure 5.15: The LEM PID distribution (FD MC) for signal (left) and background (right) broken down into a stack of different interaction types: Deep Inelastic Scattering (DIS), Quasi-elastic (QE), Resonance (RES), and Coherent Pion Production (COH).

Event Type	DIS	QE	RES	COH
Preselection				
Signal	40.8%	24.8%	33.7%	0.6%
Background	87.8%	2.4%	9.2%	0.6 %
LEM $>$ 0.70				
Signal	24.8%	37.1%	33.3%	0.8%
Background	82.3%	5.3%	8.2%	4.2 %

Table 5.3: Interaction types selected in the FD MC sample, both for Preselection and LEM>0.70.

5.6.4 Cross-Checks of LEM

It is valuable to examine the efficiency of the LEM cut with respect to position, energy, and other variables, to check for any remaining biases in either of the two detectors. Once again, this is more a concern for efficiency than for accuracy, as most systematic biases will cancel out to first order in the extrapolation. Figures 5.16 and 5.17 show the efficiency of LEM relative to the Preselected sample as a function of position in the detector. This is shown both in the plane perpendicular to the beam (u vs v) and along the direction of the beam (z). In each case, this is shown for ND Data and MC, and for FD MC. For reasons that will be discussed in Chapter 6, the discrepancy between the data and Monte Carlo efficiency in the Near Detector is expected. There are some small

Event Type	0 π^0	1 π^0	>1 π^0
Preselection			
Signal	71.7%	23.2%	5.1%
Background	38.1%	40.6%	21.3%
LEM > 0.70			
Signal	78.4%	18.4%	3.2%
Background	15.2%	54.4%	30.5%

Table 5.4: Number of π^0s present in events in the FD MC sample, both for Preselection and LEM>0.70.

spatial variations from the mean. As can be seen in the case of the Near Detector, these spatial variations appear to be well-modeled between Data and Monte Carlo.

One can also assess the light level corrections in a more direct way. Light level can be approximated by looking at the ratio between the event's pulse height in two scales, MIPs and PEs. MIPs describe the response of the scintillator itself. PEs[4] describe the actual amount of light which is eventually read out at the PMTs. The ratio of (total event PEs)/(total event MIPs) is therefore a good approximation of an event's mean light level, i.e. how much readout light was produced on average by an event's deposited charge. This ratio is not well-defined for all topologies, with very high or low values of PE/MIP being dominated by low energy events. However, in the central regions where the light level is well defined (and all event energies are well-represented), the efficiency is flat, as shown in Figure 5.18. In both the spatial and light level efficiency plots, the efficiencies are also quite close for both the Near and Far Detectors. These plots therefore suggest that the attenuation corrections discussed in Section 5.3.3 have been largely successful in making the LEM selection independent of detector position effects.

Finally, one can also consider the efficiency of the LEM detector with respect to reconstructed energy. The particle ID has been trained to look for a signal with a peak at approximately 2 GeV. LEM will therefore have a lower efficiency for events with an energy higher or lower than this peak. Figure 5.19 shows the efficiency (with respect to Preselection) of LEM in terms of reconstructed energy, showing that this is in fact true.

---

[4] These are actual PEs, from the ADC→PE conversion, rather than effective PEs.

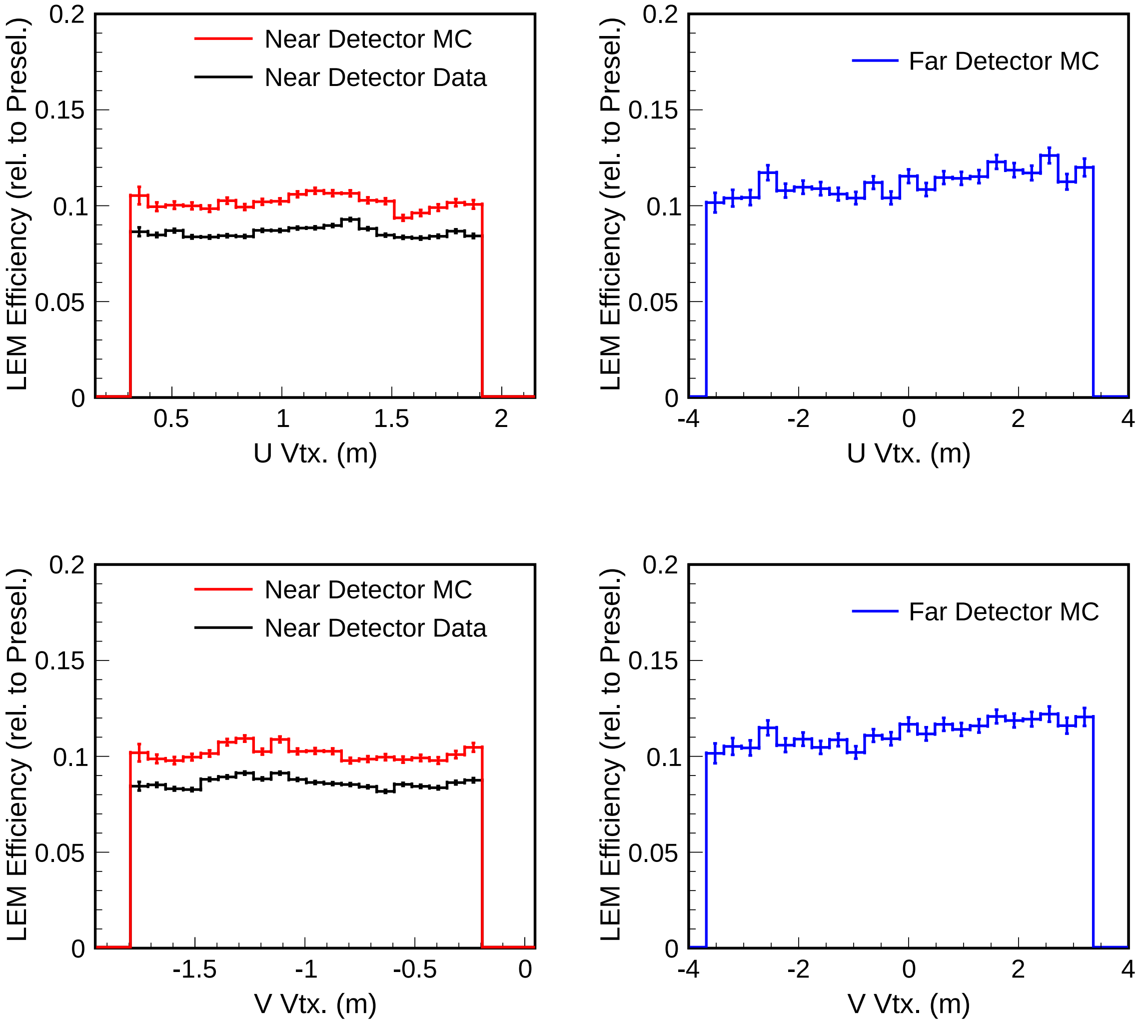

Figure 5.16: The efficiency of the LEM>0.70 cut as a function of U (top) and V (bottom) direction vertex in the MINOS detectors, for ND Data and MC (left) and FD MC (right). Efficiency is with respect to the standard Preselection Sample. The FD MC sample contains only background events (NC, ν_μ CC, Beam ν_e CC). No PORP or other data-based corrections are applied.

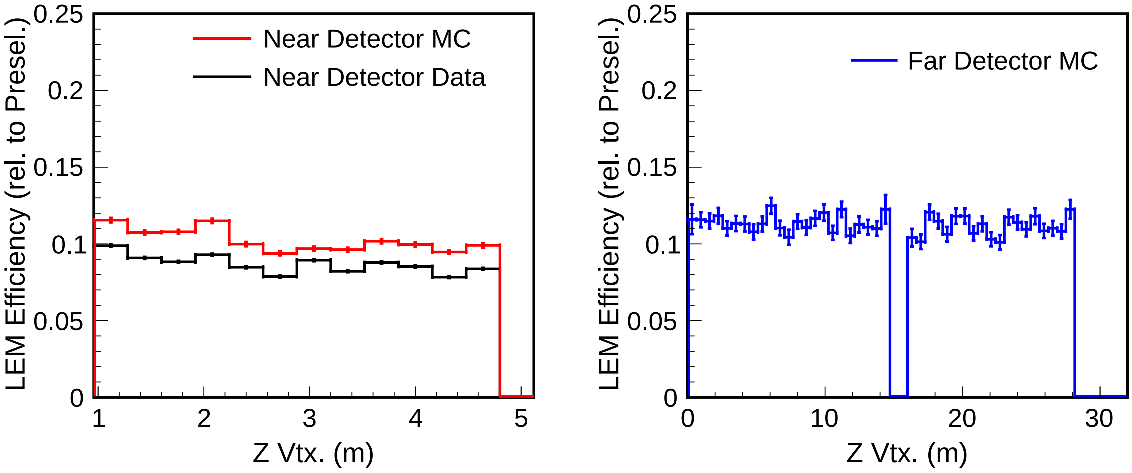

Figure 5.17: The efficiency of the LEM>0.70 cut as a function of Z vertex in the MINOS detectors, for ND Data and MC (left) and FD MC (right). Efficiency is with respect to the standard Preselection Sample. The FD MC sample contains only background events (NC, ν_μ CC, Beam ν_e CC). No PORP or other data-based corrections are applied.

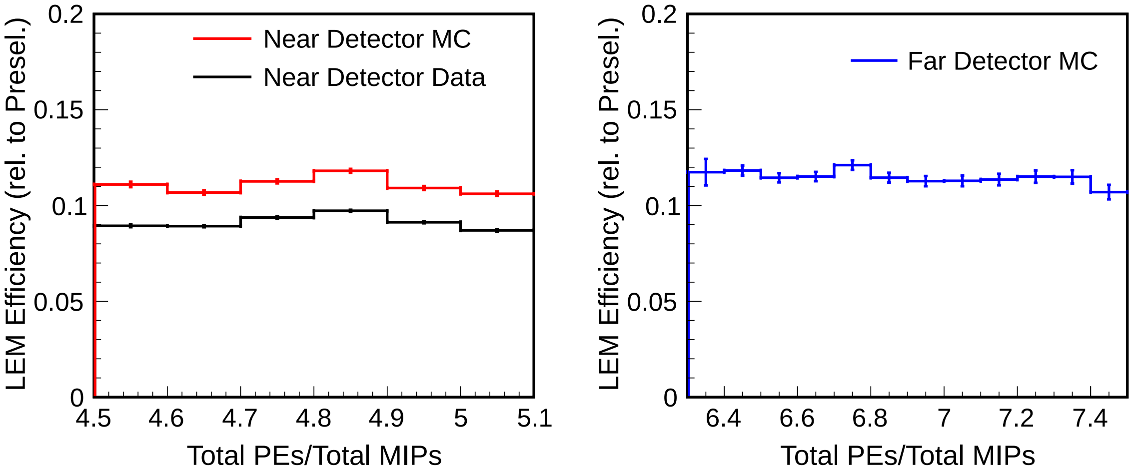

Figure 5.18: The efficiency of the LEM>0.70 relative to Preselection as a function of light level. Light level is here defined the ratio (pulse height in PEs)/(pulse height in MIPs). Only the 'well-defined" region is shown - outside of this region, low energy events dominate. ND Data and MC are shown on the right, and FD MC on the left.

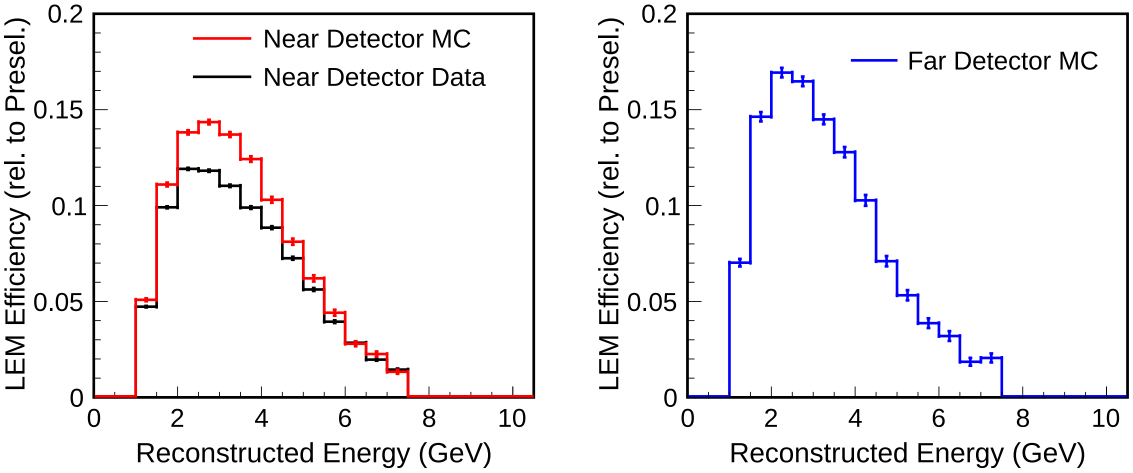

Figure 5.19: The efficiency of the LEM>0.70 cut as a function of Reconstructed Energy in the MINOS detectors. Efficiency is with respect to the standard Preselection Sample. The efficiencies for Near Detector Data and Monte Carlo are shown on the top, and for Far Detector Monte Carlo on the bottom. The Far Detector sample contains only background events (NC, ν_μ CC, Beam ν_e CC). No PORP or other data-based corrections are applied.

5.7 Summary

The Library Event Matching (LEM) PID improves signal-background discrimination in MINOS by in principle using all available strip hit information. A basic form of the PID existed prior to the work in this thesis. The new version of LEM has been optimized by the author in several ways, most significantly in the introduction of a new multivariate discriminant. The new version of the PID is potentially up to 22% more sensitive to signal than the previous ν_e analysis PID, ANN11. The behavior of the matching and efficiency has now been thoroughly checked, as well. No unexpected biases or systematic errors have been found.

Chapter 6

The Near Detector Selection

6.1 Introduction

The previous section focused almost exclusively on signal/background discrimination in the MINOS Far Detector. The Near Detector plays an equally important role in this analysis. In order to make a prediction of signal and background at the Far Detector, the LEM selection is first applied to the unoscillated Near Detector data. This selection is then extrapolated to the Far Detector, to make a data-corrected prediction of the selected FD background in the absence of oscillations.

Because of ν_μ CC oscillations and other effects, the individual ND Data background components (ν_μ CC, NC, and beam ν_e CC) must be extrapolated separately to the Far Detector. One therefore needs to make an accurate measurement of the relative fraction of each background type. In theory, one could use the Near Detector Monte Carlo to determine these relative ND background rates. However, there is large discrepancy between the ND Data and MC for the LEM particle ID selection. Figure 6.1 shows the LEM distribution and the three PID variables for ND Data and MC. In all cases, the shape of the distribution is different between Data and MC. The degree of mismodeling is particularly strong in the high-PID region. Figure 6.2 also shows reconstructed energy for the LEM>0.70 selection. In this signal-like region, there is a 15%-25% excess in the Monte Carlo.

This chapter will focus on this discrepancy in two separate ways. First, it will describe the author's attempts to determine the potential sources of the large Data/MC discrepancy. Because the LEM method is highly dependent on the details of the Monte Carlo modeling, it is not surprising that there is a large discrepancy between Data and MC. Prior to the work in this thesis, the ND Data/MC discrepancy had been observed, but the sources of the mismodeling were not well understood. Notably, as long as any

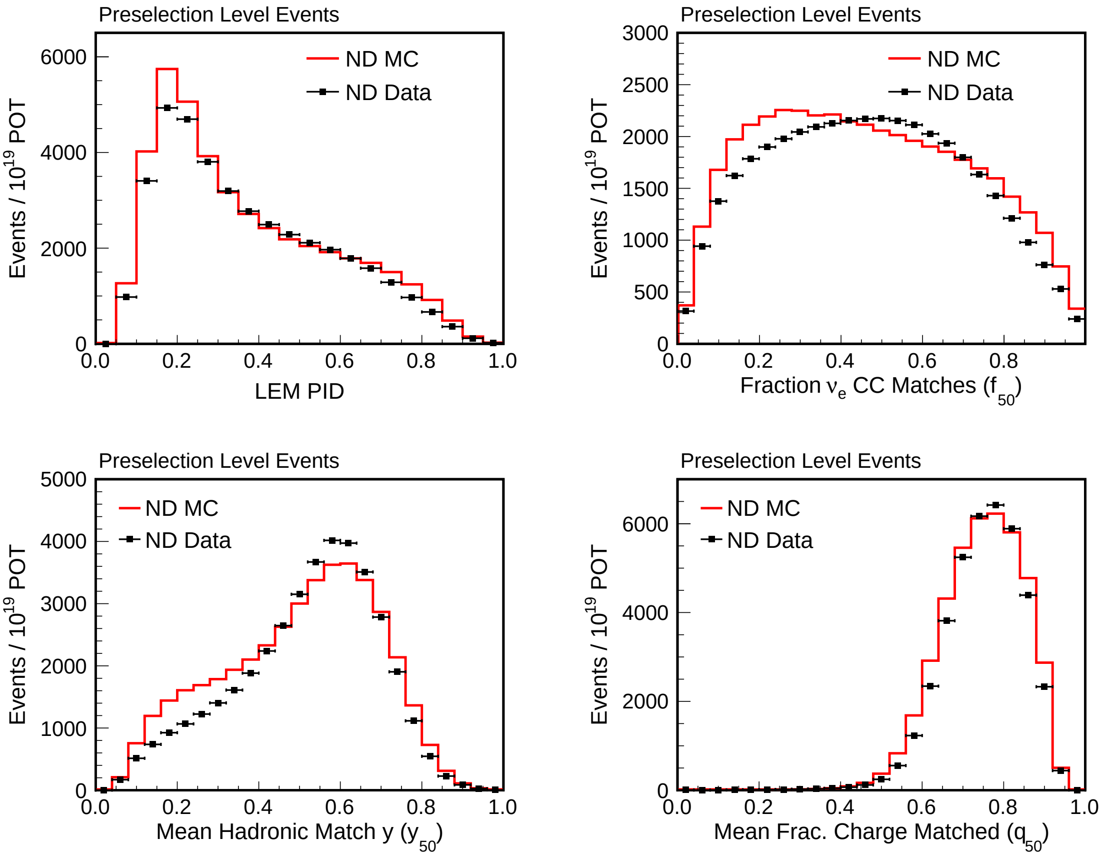

Figure 6.1: The LEM Particle ID distribution, and the three LEM PID input variables, in the Near Detector. Events passing the Preselection cuts are shown. Both Data and MC are scaled to 10^{19} POT.

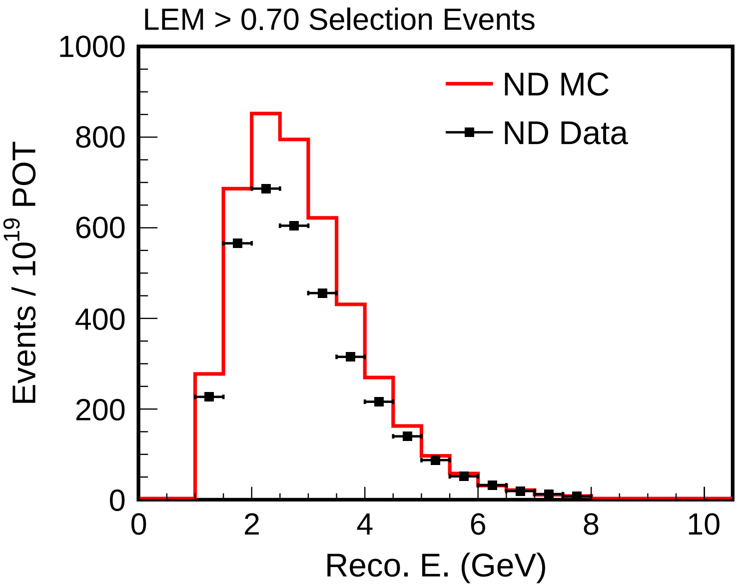

Figure 6.2: Reconstructed energy (GeV) in the Near Detector Data. Events passing both the LEM>0.70 cut and the Preselection are shown. Both Data and MC are scaled to 10^{19} POT.

Far/Near effects are adequately modeled, this discrepancy will not affect the final result.

Second, this chapter will briefly describe the means by which the Near Detector can be decomposed into separate background components. This method, the Horn On Horn Off High Energy (HOOHE) beam decomposition method, was developed for the MINOS ν_e appearance analysis over the course of several analyses. A brief summary of this method will be given by the author.

6.2 Exploring the Source of the Data/MC Discrepancy

6.2.1 Hadronic Shower Modeling in MINOS

In Figure 6.1, there is a large difference in shape between Data and MC for the variable y_{50}, the average true y of an event's ν_e CC matches. This suggests that the mismodeling of hadronic showers could be the source of the large Data/MC difference in the LEM selection. Hadronic showers at a higher invariant masses are well modeled by the PYTHIA/JETSET model [79]. This modeling breaks down for events with a lower

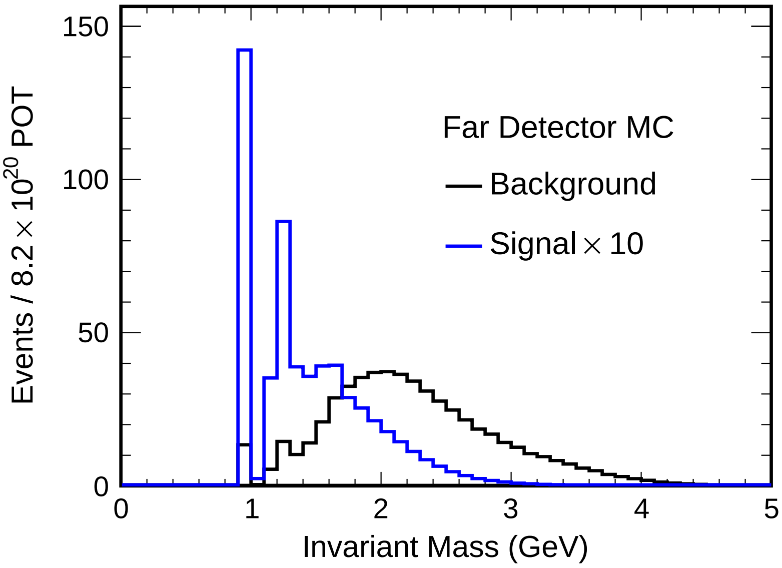

Figure 6.3: The true invariant mass W of signal (blue) and background (black) events in the FD Monte Carlo. All events have passed the Preselection. Signal (CHOOZ-size) has been scaled × 10 for ease of viewing.

invariant mass closer to the pion production threshold. For this reason, a phenomenological model, AGKY [105], was developed to better model the type of hadronic showers present in MINOS. This model was largely tuned with and checked against data from bubble chamber experiments (BEBC [106] [107] [108] [109], SKAT [110], and the Fermilab 15' [111] bubble chamber). All events with invariant mass $W > 3.0$ GeV are still described using PYTHIA/JETSET. All events with $W < 2.3$ GeV are described with a KNO-based model. These regions are linked by a transition region, between 2.3 and 3.0 GeV, in which fraction of events described by PYTHIA/JETSET rises from 0% to 100%, and the fraction of events described by KNO drops from 100% to 0%. As shown in Figure 6.3, most interactions of interest to the ν_e analysis are found in the KNO and transition regions.

The KNO-based portion of AGKY carries out two major tasks: first, it chooses the multiplicity of the particles which comprise the hadronic shower, and second, it selects the 4-momenta of these particles. The first task (populating the hadronic shower) involves several steps. First, the average charged hadron multiplicity is chosen: $\langle \mathrm{n_{ch}} \rangle = a_{ch} + b_{ch} W^2$, where a_{ch} and b_{ch} are constants derived from bubble chamber experiment data. Next, the average total hadron multiplicity is calculated: $\langle \mathrm{n_{tot}} \rangle = 1.5 \langle \mathrm{n_{ch}} \rangle$. Finally, the KNO scaling law [80] is used to turn this average into an actual hadron multiplicity:

$$\langle n \rangle \times P(n) = f(n/\langle n \rangle) \tag{6.1}$$

where $P(n)$ is the probability of obtaining n hadrons, and f(x) is the Levy function (with a constant c derived from experiment):

$$f(z,c) = \frac{2e^{-c}c^{cz+1}}{\Gamma(cz+1)} \tag{6.2}$$

The kinematics and charge of the system determine what ratio of each particle type populates the shower. A single baryon is also produced in the hadronic shower, based on quark models.

Next, the momenta of the resulting particles must be described. First, a 4-vector $P_N^* = (E_N^*,\mathbf{p}_N^*)$ is chosen for the baryon, which dominates the system. This choice is based both on the transverse momentum of the nucleon, p_T^2, and on parton distributions functions (with respect to x_F) calculated from earlier experiments [113] [114]. Momenta are next chosen for the hadrons that populate the remainder of the system. These hadron momenta are generated using a phase space decay in center of mass. During this decay, the transverse momentum of the shower is constrained with a "squeezing" factor A, applied as $e^{-A\,\mathrm{p}_T}$. The system is then boosted back in to the lab frame.

6.2.2 y_{50} as a probe of hadronization

In Section 5.4.2, it was shown that the reconstructed energy of an input event corresponds closely with the true energy of its ν_e CC best LEM matches. A similar relationship could also be expected between the true inelasticity y of a signal event and the average y of its ν_e CC matches - i.e., the LEM variable y_{50}. Indeed, Figure 6.4 shows a strong correlation between these variables. This is unsurprising, as electromagnetic showers in ν_e CC signal events are known to be relatively well-modeled in the MINOS detector [99].

A similar study can be performed for the hadronic showers in NC input events as well. Many hadronic shower events also have a electromagnetic component, caused by a π^0 decaying to two photons. For these MC events, one can construct the variable y_{NC}, the fraction of the NC shower which is electromagnetic. This is the fraction of the NC visible energy E^{vis} carried by π^0s:

$$y_{NC} = 1 - \frac{E^{\pi^0}}{E^{vis}} \tag{6.3}$$

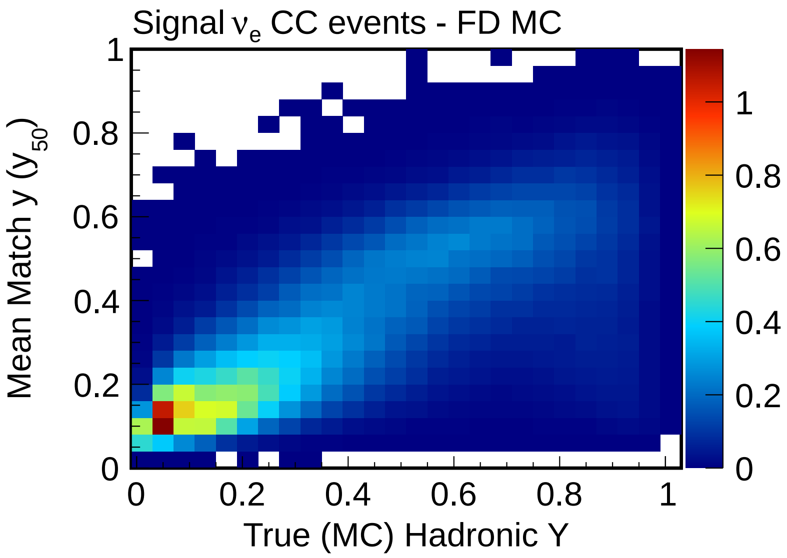

Figure 6.4: The relationship between true hadronic y and y_{50} for ν_e CC input events in the Far Detector MC.

Figure 6.5 shows the relationship between y_{NC} and y_{50} (for NC input events). There is a clear correlation between these two variables. The quantity y_{50} distribution, and the LEM PID as a whole, therefore probe the EM fraction of a hadronic shower and the hadronic modeling of a MC event in general. Mismodeling in AGKY would be reflected in a resulting discrepancy between Data and MC for hadronic shower background events selected by the LEM PID.

6.2.3 Sources of Mismodeling in AGKY

A number of approximations and ambiguities in AGKY could be responsible for the ND Data/MC discrepancy in LEM. During the evaluation of systematic error for the first ν_e analysis, six separate possible sources of error in AGKY were identified [115]. In order to determine their effect on the ν_e analysis, these errors were evaluated using a series of probability distribution function histograms which reweighted existing Monte Carlo as a function of invariant mass W, summed transverse momentum $\sum p_T$ and hadronic shower EM fraction. The histograms themselves were 2 dimensional histograms of p_T vs EM fraction, in 0.1 GeV slices of W^2 [95]. These histograms were normalized to keep the total event count constant. The histograms modeled the following uncertainties:

Baryon x_f Selection (Model 1) This histogram reweights the MC so that the selection of baryon momentum results in showers which more closely match reality.

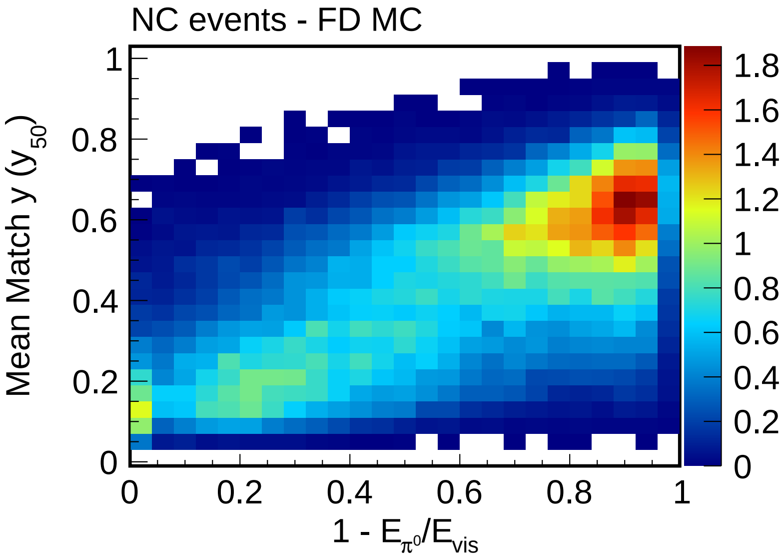

Figure 6.5: The relationship between the EM fraction of a NC hadronic shower ($1 - E_{\pi^0}/E_{vis}$) and y_{50} (the mean matched y of ν_e CC matches) for NC input events in the Far Detector MC.

The baryon momentum, as described before, is chosen based on parton models. This four-vector is most likely to be found in a backward direction in the center of mass frame, due to the location of the di-quark. This results in most of the pions and other hadrons moving in the forward direction in a jet-like distribution. Comparisons with experimental data found that AGKY overestimated the multiplicity of charged hadrons with a forward momentum and underestimated those with a backward momentum. Model 1 changes this by generating the baryon momentum isotropically in the center of mass using a phase-space decay [116]. This is believed to more closely match reality, and some feature resembling this isotropic decay may be included in future versions of GENIE [117].

π^0 Selection (Model 2±) This reweighting model accounts for the uncertainty in the probability of π^0 production in the shower, a feature which could have a large impact on the hadronic shower EM fraction, and the resulting LEM PID value. AGKY uses an experimentally derived selection probability of 30%. There is some uncertainty on this number, due to both the difficulty of detecting π^0s in a hydrogen bubble chamber environment, and to the presence of intranuclear rescattering in experiments utilizing the heavier freon or neon. Comparisons [115] of the predicted AGKY π^0 multiplicity and π^0 dispersion with experiment (SKAT freon measurements [110] and BEBC neon [107] and hydrogen [108] measurements) indicated a ±20% uncertainty on the π^0 selection probability. These two reweighting histograms reset the probability from 30% to 21% and 39%.

Charged-Neutral Multiplicity Correlation (Model 3) This reweighting model considers the degree of correlation between the multiplicity of charged and neutral hadronic particles. The current AGKY model selects the charged particle multiplicity first, and from that the overall shower multiplicity. This introduces correlation between the charged and neutral multiplicities. This correlation is in disagreement with the data, which suggests that the two multiplicities are largely independent when W is on the order of a few GeV. This histogram reweights the MC to simulate independent multiplicities.

Implementation Ambiguities (Model 4) This reweighting model considers several differences in the hadron 4-vector selection process (e.g., the procedure taken during p_T squeezing) which were discovered between the NEUGEN v3.5.0 and GENIE v2.0.0 versions of AGKY. This model simulates a GENIE-only model.

$\mathbf{p_T}$ Squeezing (Model 5) This reweight model simulates uncertainty in the p_T rejection factor (the squeezing factor A described in 6.2.1) which is used to select the transverse momentum of a shower. The parameter A is changed from 3.5 to 1.5, resulting in broader showers.

Isotropic 2-body Decay (Model 6) AGKY models 2-body decays isotropically in the center of mass. This reweighting model tests the impact of this assumption by performing the decay orthogonally to the direction of momentum transfer.

Each of these reweighting factors was applied to the ND MC to investigate the effect of potential mismodeling on Monte Carlo. Figure 6.6 shows the effect of each reweight on the LEM>0.70 reconstructed energy distribution. The majority of models studied changed the distribution by between 0% and 5%. The uncertainty in the π^0 multiplicity was expected to have a major impact on the LEM selection, due to the important role played by the $\pi^0 \rightarrow \gamma\gamma$ decay, but could not fully account for the $15-25\%$ discrepancy between ND Data and MC. One of the models, Baryon x_f Selection (Model 1), however, did have the required effect on the high PID region. As can be seen in Figure 6.6, reweighting for Model 1 results in the best agreement between ND MC and data.

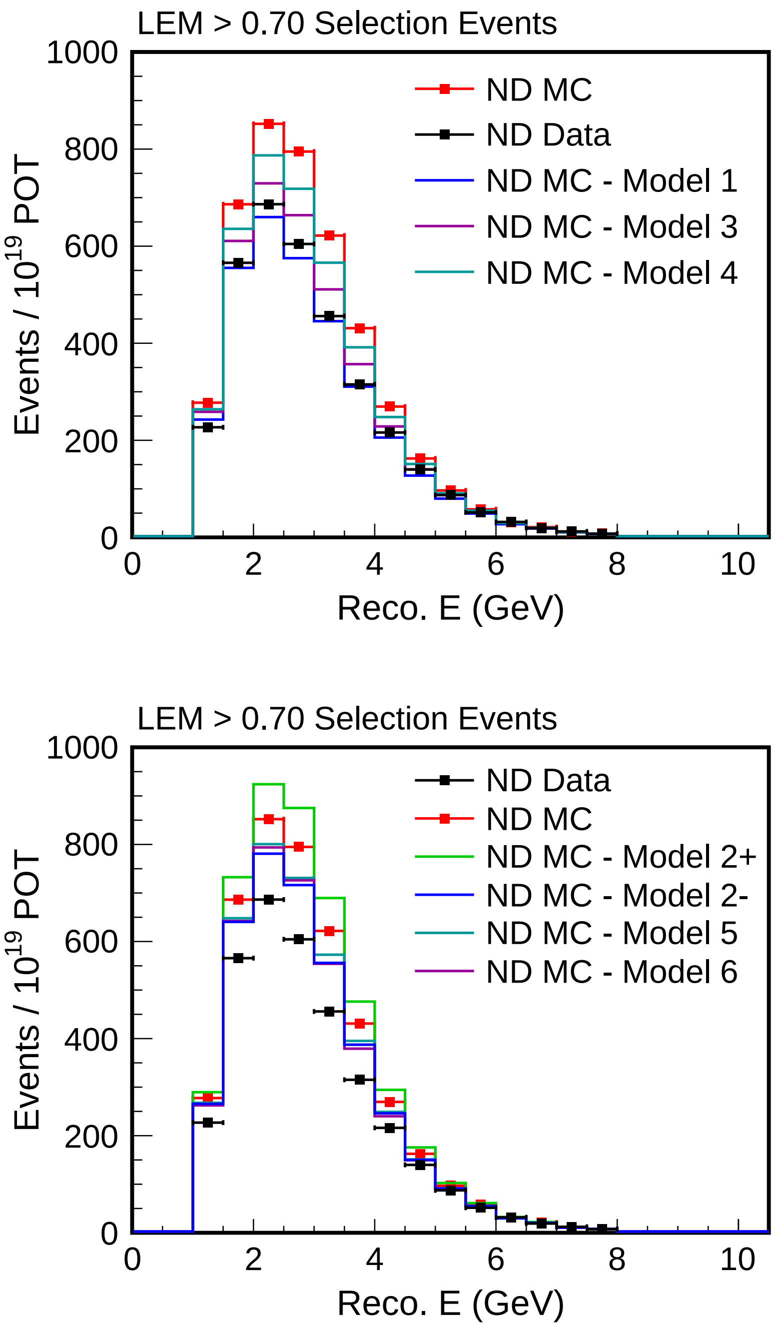

Figure 6.6: The reconstructed energy distribution for the LEM>0.70 selection in the Near Detector, for Data, Monte Carlo, and Monte Carlo reweighted to various hadronic models. The models are as follows (with the aspect of AGKY they modify): Model 1 (baryon $\mathbf{x}_f$ selection), Model 2± (π^0 selection probability), Model 3 (Charged-Neutral particle multiplicity correlation), Model 4 (implementation ambiguities), Model 5 (p_T squeezing), Model 6 (isotropic 2-body decays). Both samples are scaled to 10^{19} POT.

6.2.4 Reweighting Events for Model 1

The effect of the Model 1 reweighting on the LEM PID distribution is shown in Figure 6.7. In the high-PID region, the ND Data/MC discrepancy is drastically reduced. The reason for this reduced discrepancy can be seen in the reweighted y_{50} distribution, in 6.8. The shape of the reweighted y_{50} MC curve, the best probe of the hadronic modeling, now much more closely matches the data. The LEM>0.70 reconstructed energy distribution, in Figure 6.9, now has a discrepancy on the order of 5% rather than 20%. The effect of the baryon 4-vector modeling can also be studied a step further back in the LEM PID, as well, by using Model 1 reweighted MC events to retrain the LEM neural net itself. In this case, a similar level of Data/MC agreement is obtained.

The results of these reweighting studies strongly suggest that hadronic mismodeling is the primary cause of the large Data/MC disagreement found in LEM. In particular, the LEM PID is very sensitive to the electromagnetic content of hadronic showers, via the PID variable y_{50}. When the Monte Carlo is reweighted to simulate a model which better describes the EM content (as indicated by the shape of the y_{50} distribution), the ND data and MC agree much more closely.

6.2.5 Impact of Hadronic Modeling on super-FOM

Although uncertainties in hadronic modeling have a large effect on the shape and size of the LEM distribution for Monte Carlo, the overall effect on the prediction and sensitivity is minimal. The MINOS two-detector setup is designed to cancel out systematic errors and uncertainties of this sort to first order. Figure 6.10 shows super-FOM for a single-bin cut for Standard LEM (regular MC), Model 1 Reweighted LEM (MC events reweighted for Model 1, no retraining), and Model 1 Retrained LEM (both input and PID training events reweighted for Model 1), assuming a systematic error of 5%. Reweighting is now applied to both the ND and FD MC. As described earlier, Standard LEM has a maximum super-FOM of 4.03 at LEM>0.76, with 21.8 signal events and 27.4 background events selected. Reweighted LEM (no retraining) has a maximum super-FOM of 3.97 at LEM>0.76, with 21.4 signal events and 27.1 background events selected. The Retrained LEM has a maximum super-FOM of 4.07, with 21.1 signal and 25.4 background events selected. The behavior of the Standard and Reweighted LEM are essentially the same, while the Retrained LEM has slightly improved sensitivity. This suggests that while the Data/MC discrepancy in LEM is caused by real physics, the presence of poorly modeled hadronic showers in LEM does not appear to have resulted in any appreciable reduction in performance. Any remaining second order effects from mismodeling will incorporated

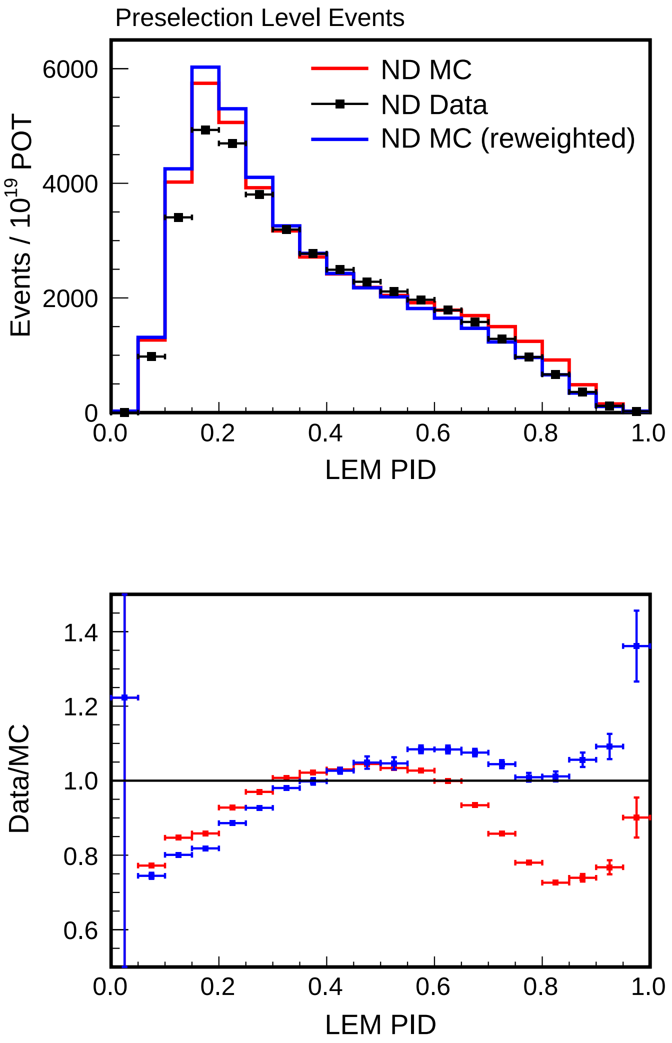

Figure 6.7: The LEM PID distribution (top) in the Near Detector, for Monte Carlo and Data events passing the standard ν_e preselection. The Monte Carlo is shown both raw (red) and reweighted to the Model 1 baryon selection model (blue). All samples are scaled to 10^{19} POT. The ratio of the Data / MC is shown on the bottom.

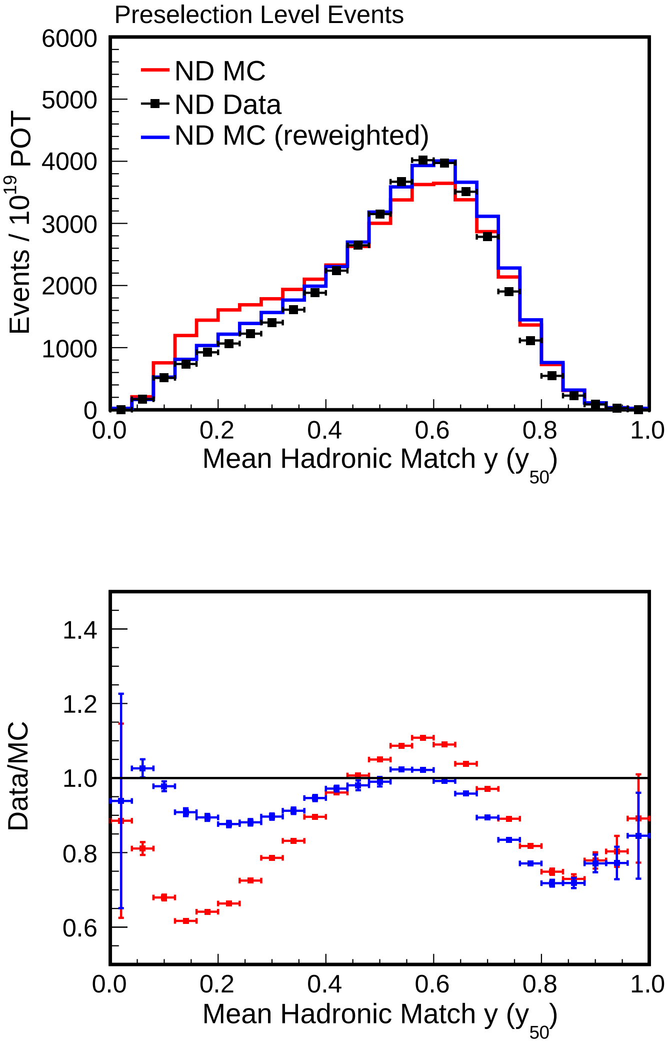

Figure 6.8: Mean Match y (y_{50}) in the Near Detector, for Monte Carlo and Data events passing the standard ν_e preselection. The Monte Carlo is shown both raw (red) and reweighted to the Model 1 baryon selection model (blue). All samples are scaled to 10^{19} POT. The ratio of the Data / MC is shown on the bottom.

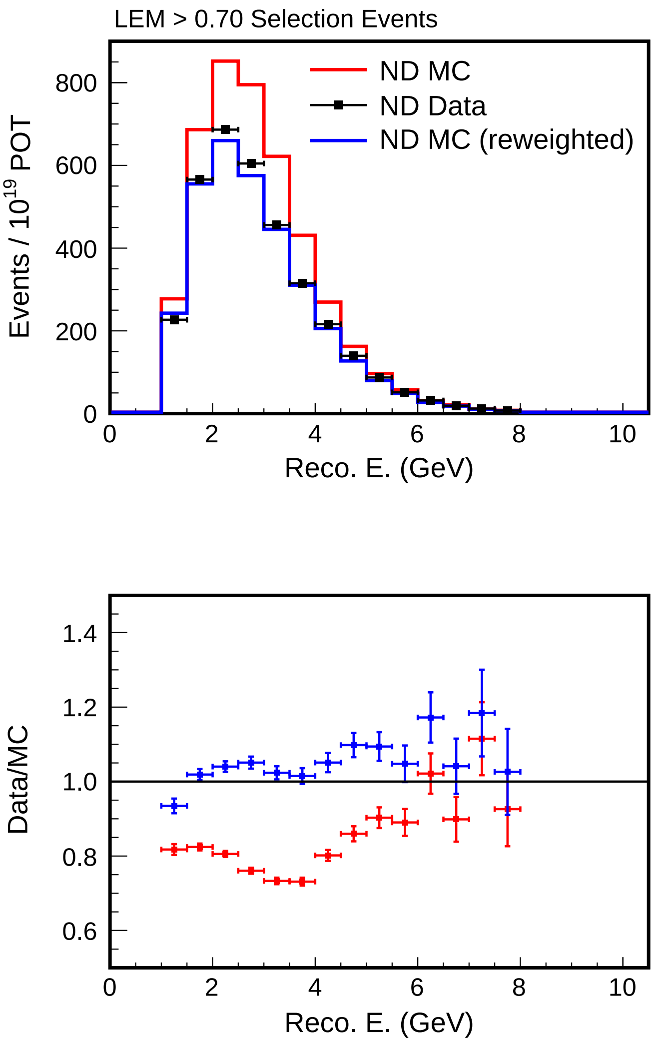

Figure 6.9: Energy spectrum of Near Detector events passing LEM PID > 0.70. The Monte Carlo is shown both raw (red) and reweighted to the Model 1 baryon selection model (blue). All samples are scaled to 10^{19} POT. The ratio of the Data / MC is shown on the bottom.

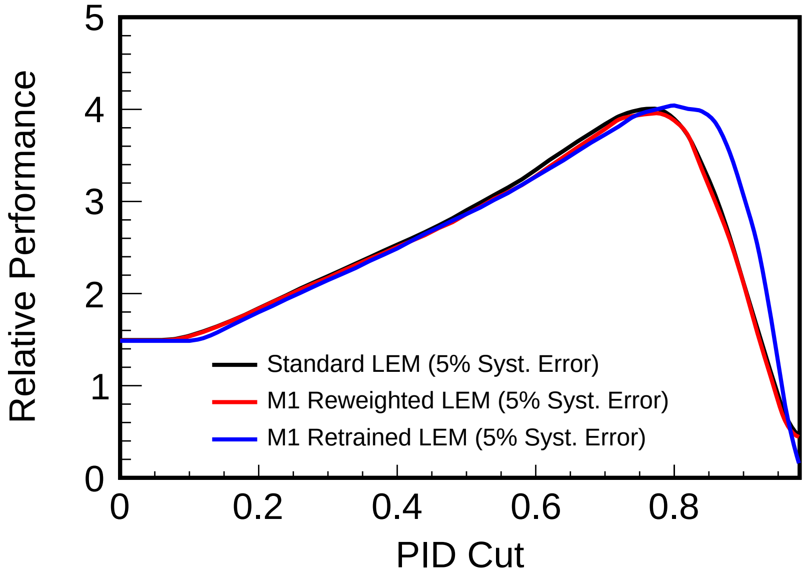

Figure 6.10: Super-FOM with 5% systematic error as a function of cut for the Standard, M1 Reweighted, and M1 Retrained LEM PIDs. All quantities have PORP-based corrections and a CHOOZ-size signal.

into the analysis as systematic error.

6.3 The HOOHE Method

The large discrepancy between Data and Monte Carlo in the LEM Selection means that one cannot rely on Monte Carlo alone to decompose the Near Detector data into individual background types. A more sophisticated decomposition technique must therefore be used. The Near Detector decomposition procedure described here is referred to as the "Horn On, Horn Off, High Energy" (HOOHE) method. A summary of the process is given below, but full details can be found in reference [118].

6.3.1 Three Separate Beam Configurations

The Near Detector decomposition procedure makes use of the fact that, while the NuMI beamline tends to run primarily in Low Energy (LE10/185kA) mode, it is also possible to change the configuration of the magnetic focusing horns and graphite target to produce different beam configurations. These beam configurations not only have different overall neutrino fluxes and energy spectra, but also different relative rates for each event type

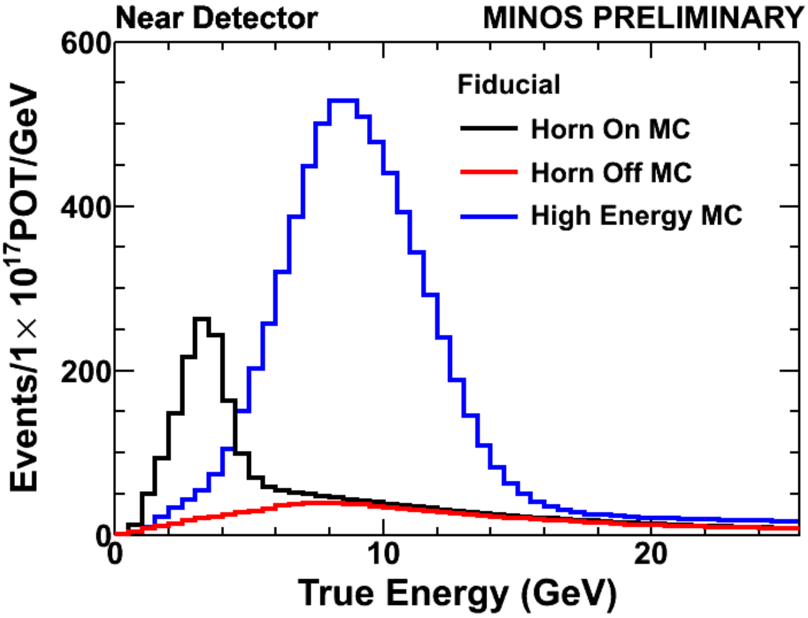

Figure 6.11: Expected true energy spectra for the three beam configurations used in the HOOHE method: Low Energy ("Horn On"), Horn Off, and High Energy [119].

(ν_μ CC, beam ν_e CC, NC). This algorithm will use MINOS data recorded in both a High Energy mode (LE250/200kA) and a Horn Off mode. The expected true energy spectra for these three beam types are shown in Figure 6.11. Using the ND Data and Monte Carlo selections for each configuration, a decomposition of the Low Energy beam configuration can be performed.

In the standard LE10/185kA configuration, the magnetic focusing horns are given a current of 185 kA, and the graphite target is placed slightly upstream (10 cm) from the first of the two magnetic focusing horns. This produces a beam with a peak at 3.1 GeV (with an RMS of 1.1 GeV) [53]. When the LEM>0.70 selection is applied to Low Energy Monte Carlo, a distribution with the composition seen in in the top image in Figure 6.12 is obtained.

For the Horn Off configuration, the magnetic focusing horns are turned off. When the hadrons are no longer focused, the low energy peak present in the standard beam configuration disappears, and a larger part of the spectrum comes from high energy neutrino events (a peak of 7.4 GeV, with a broad RMS of 4.1 GeV). With the loss of the low energy peak, the LEM particle ID selects almost no ν_μ CC events at high energy, causing the relative composition of Neutral Current events to be drastically increased (see the middle image in Figure 6.12).

Finally, in the High Energy LE250/200kA configuration, the horn current is increased to 200kA and the target is moved to a position of 250 cm, upstream from the first magnetic focusing horn. In this configuration, the horn focuses higher energy hadrons, causing a distribution peaked at 8.6 GeV (and 2.7 GeV RMS). As was the case with the Horn Off configuration, the presence of higher energy hadrons causes NC events to dominate the sample.

6.3.2 Making a Linear System

From the above three beam configurations, a system of three separate linear equations can be constructed:

$$\mathrm{N}^{\mathrm{LE\ Data}}_{\mathrm{i,Total}} = \mathrm{N}^{\mathrm{LE\ Data}}_{\mathrm{i,NC}} + \mathrm{N}^{\mathrm{LE\ Data}}_{\mathrm{i,CC}} + \mathrm{N}^{\mathrm{LE\ Data}}_{i,\nu_e} \tag{6.4}$$

$$\mathrm{N}^{\mathrm{HO\ Data}}_{\mathrm{i,Total}} = \left[\frac{\mathrm{N}^{\mathrm{HO\ MC}}_{\mathrm{i,NC}}}{\mathrm{N}^{\mathrm{LE\ MC}}_{\mathrm{i,NC}}}\right]\mathrm{N}^{\mathrm{LE\ Data}}_{\mathrm{i,NC}} + \left[\frac{\mathrm{N}^{\mathrm{HO\ MC}}_{\mathrm{i,CC}}}{\mathrm{N}^{\mathrm{LE\ MC}}_{\mathrm{i,CC}}}\right]\mathrm{N}^{\mathrm{LE\ Data}}_{\mathrm{i,CC}} + \left[\frac{\mathrm{N}^{\mathrm{HO\ MC}}_{i,\nu_e}}{\mathrm{N}^{\mathrm{LE\ MC}}_{i,\nu_e}}\right]\mathrm{N}^{\mathrm{LE\ Data}}_{i,\nu_e} \tag{6.5}$$

$$\mathrm{N}^{\mathrm{HE\ Data}}_{\mathrm{i,Total}} = \left[\frac{\mathrm{N}^{\mathrm{HE\ MC}}_{\mathrm{i,NC}}}{\mathrm{N}^{\mathrm{LE\ MC}}_{\mathrm{i,NC}}}\right]\mathrm{N}^{\mathrm{LE\ Data}}_{\mathrm{i,NC}} + \left[\frac{\mathrm{N}^{\mathrm{HE\ MC}}_{\mathrm{i,CC}}}{\mathrm{N}^{\mathrm{LE\ MC}}_{\mathrm{i,CC}}}\right]\mathrm{N}^{\mathrm{LE\ Data}}_{\mathrm{i,CC}} + \left[\frac{\mathrm{N}^{\mathrm{HE\ MC}}_{i,\nu_e}}{\mathrm{N}^{\mathrm{LE\ MC}}_{i,\nu_e}}\right]\mathrm{N}^{\mathrm{LE\ Data}}_{i,\nu_e} \tag{6.6}$$

where $\mathrm{N}^{\mathrm{BC\ Data(MC)}}_{\alpha}$ is the number of ND Data (MC) events in analysis bin i for configuration BC (LE, HO, HE) and for background component α (NC, beam ν_e CC, ν_μ CC). Bin i can be, for instance, a bin of PID and/or energy.

First, $\mathrm{N}^{\mathrm{LE\ Data}}_{\mathrm{i,Total}}$, $\mathrm{N}^{\mathrm{HO\ Data}}_{\mathrm{i,Total}}$, and $\mathrm{N}^{\mathrm{HE\ Data}}_{\mathrm{i,Total}}$ are the total number of ND Data events counted in each configuration. Second, the relative frequencies of each background type between configurations, $\left[\frac{\mathrm{N}^{\mathrm{HO\ MC}}_{i,\alpha}}{\mathrm{N}^{\mathrm{LE\ MC}}_{i,\alpha}}\right]$ and $\left[\frac{\mathrm{N}^{\mathrm{HE\ MC}}_{i,\alpha}}{\mathrm{N}^{\mathrm{LE\ MC}}_{\alpha}}\right]$, are known with good accuracy from Monte Carlo. Much as is the case when comparing the Near and Far Detector MC, systematic errors cancel out to first order when taking the ratio of background rates between two beam configurations. Third and finally, the individual LE data background components $\mathrm{N}^{\mathrm{LE\ Data}}_{i,NC}$, $\mathrm{N}^{\mathrm{LE\ Data}}_{i,CC}$, and $\mathrm{N}^{\mathrm{LE\ Data}}_{i,\nu_e}$ are the unknowns of the system. With three equations and three unknowns, the system can in theory be solved to find the individual Near Detector Background components.

In reality, the process is somewhat more complicated. Both the Horn Off and High Energy configurations are dominated by Neutral Current events, causing Equations 6.5

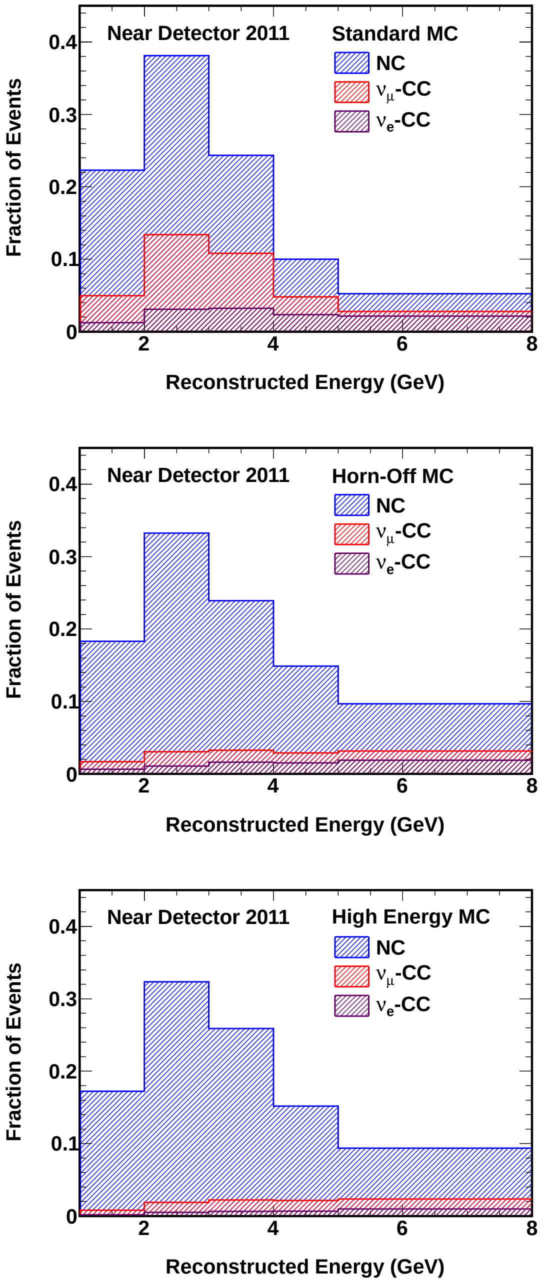

Figure 6.12: Relative compositions [119] of background events for the three separate beam configurations used in the Near Detector decomposition, for a LEM>0.7 selection. Horn On (standard beam) is shown on top, followed by Horn Off (middle) and High Energy (bottom). Distributions are area-normalized.

and 6.6 to be nearly linearly dependent [118]. The uncertainty on the solutions to the system will therefore be quite high. While the original Near Detector Monte Carlo predictions are not accurate, they can be used as an additional constraint, so that the uncertainty on the decomposition will be no greater than the original uncertainty on the Monte Carlo. With this additional requirement, the system is now over-constrained. Rather than obtaining an exact solution for $\mathrm{N}_{i,NC}^{\text{LE Data}}$, $\mathrm{N}_{i,CC}^{\text{LE Data}}$, and $\mathrm{N}_{i,\nu_e}^{\text{LE Data}}$, a χ^2 minimization is fit. Statistical, flux, and systematic errors (see Chapter 8) on the ND samples are incorporated into this fit in the form of a covariance matrix [118].

6.3.3 Results - the Near Detector Decomposition

The decomposition is performed with the ND Data split into bins of PID and Energy. For the energy binning, a grouping of five bins with edges [1 GeV, 2 GeV, 3 GeV, 4 GeV, 5 GeV, 8 GeV] was chosen. The choice of LEM binning will be discussed in Section 9.3, but the decomposition of the LEM>0.7 selection will here be considered as an example (i.e., one PID bin). In addition to bins of PID and energy, the decomposition is also broken down into separate analysis runs (Runs 1, 2, and 3) and background types (ν_μ CC, NC, beam ν_e CC) when the decomposition χ^2 fit is performed. The χ^2 fit also requires the solution to not have a negative solution in any bin.

The resulting decomposition of the LEM>0.7 sample is shown in Figure 6.13. The fit has a χ^2 of 20.2, for 44 degrees of freedom. The breakdown of the sample into NC, ν_μ CC, and beam ν_e CC samples can be found in Table 6.1. The NC and ν_μ CC components have a correlation of -0.172.

Event Type	% of ND Data Sample	Error on Sample.
NC	61.5	± 1.1
ν_μ CC	23.5	± 1.1
beam ν_e CC	15.0	± 1.4

Table 6.1: Percentage breakdown of the individual background components in the LEM>0.70 sample, with errors. These numbers are shown graphically in Figure 6.13

.

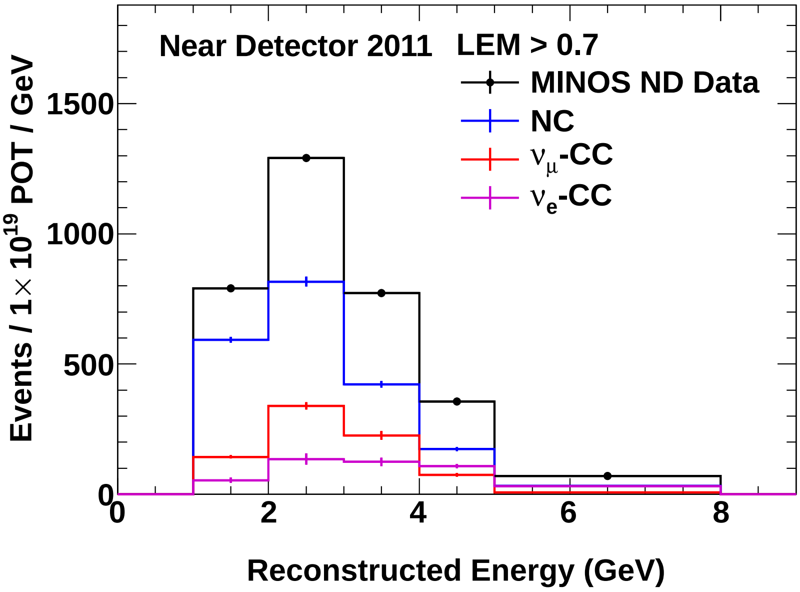

Figure 6.13: The results of the example of a Near Detector Decomposition for LEM>0.70 sample, with five energy bins. The total curve (in black) is derived from the Near Detector data. The individual background components have been derived from the decomposition. The plot is normalized to an exposure of 10^{19} POTs. Plot produced by Joao Coehlo [119].

6.4 Summary

A firm understanding of the Near Detector Data is vital for the $\nu_\mu \to \nu_e$ search presented in this thesis. In the next chapter, the ND Data selection is extrapolated to the Far Detector to make a background prediction. Because the ν_μ CC component oscillates, while the NC component does not, it is vital to break down the ND data into its constituent parts. The HOOHE method accomplishes this decomposition.

Studying the Near Detector Data and Monte Carlo also allow for a more sophisticated understanding of the LEM PID, however. The LEM PID was previously known to have a large discrepancy between the ND Data and MC in the signal-like region (LEM>0.70). The work in this chapter has revealed that the LEM variable y_{50} is a good probe of the EM content of hadronic shower events. The Data/MC discrepancy is therefore mostly caused by hadronic shower mismodeling. This type of systematic error will largely cancel out to first order in the Far/Near extrapolation, having little effect on the final prediction. With the ND data now well understood, the analysis can proceed to a prediction of the Far Detector background.

Chapter 7

Predicting the Far Detector Data

Now that the ND Data has been decomposed, this chapter will discuss how a prediction is made for the five different expected FD interaction types: Neutral Current, ν_μ CC, beam ν_e CC, ν_τ CC, and a ν_e CC signal. Most of the techniques in this chapter were developed for previous analyses, with the exception of an assessment of the rock muon background, which has been carried out for this thesis.

7.1 The NC, ν_μ CC, and beam ν_e CC Background Predictions

The first three interaction types (NC, ν_μ CC, and beam ν_e CC) are the most straightforward to predict. The rates of these three background types are now known in the Near Detector from the HOOHE decomposition. Using Far/Near ratios from Monte Carlo, these backgrounds can each be separately extrapolated to form a Far Detector prediction:

$$\mathrm{F}_i^\alpha = \frac{\mathrm{f}_i^\alpha}{\mathrm{n}_i^\alpha}\mathrm{N}_i^\alpha \tag{7.1}$$

where F_i^α is the FD prediction for background type α in extrapolation bin i (a bin of PID and/or energy), and f_i^α, n_i^α, and N_i^α are respectively the FD MC, ND MC, and ND Data samples.

Several effects make it necessary to extrapolate these background components separately. Muon neutrinos oscillate between the Near and Far Detectors, causing a ν_μ CC-specific FD deficit. The Neutral Current rate will undergo no such oscillation. Additionally, the energy spectrum for beam ν_e CC events is different between the Near

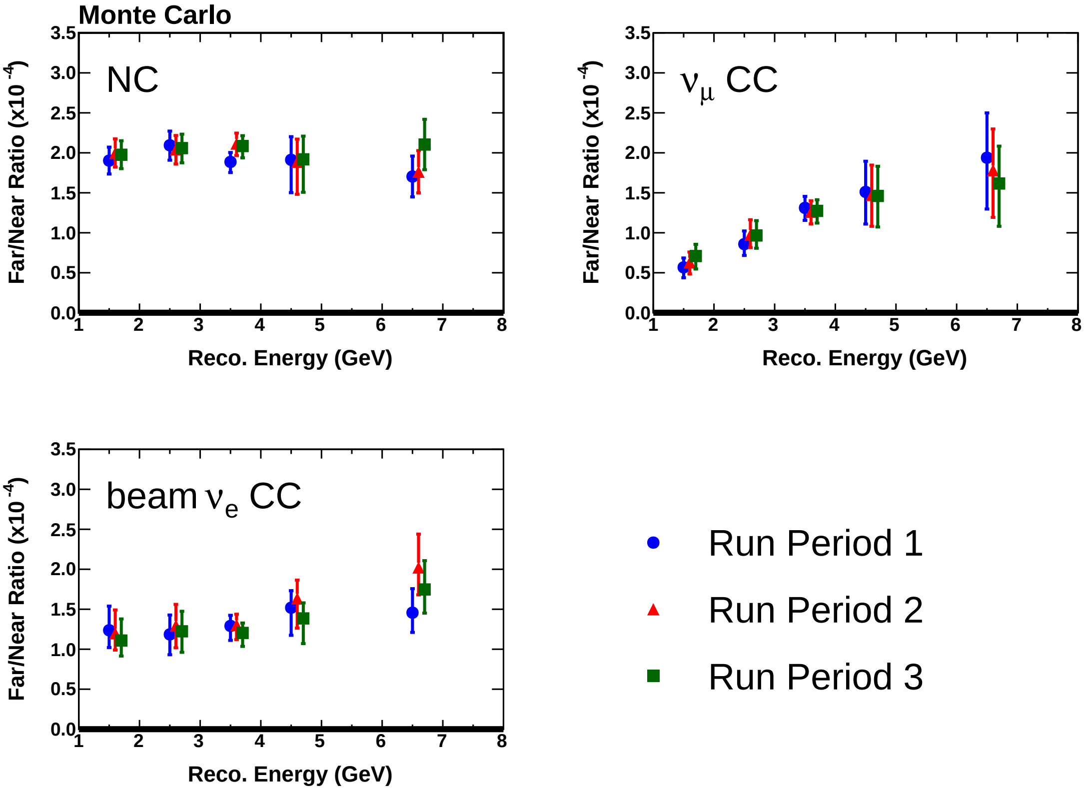

Figure 7.1: Examples of Far/Near MC ratios used in the extrapolation of the Near Detector backgrounds (Neutral Current, ν_μ CC, and beam ν_e CC) to the Far Detector. This plot shows extrapolation to the Far Detector in 5 bins of Reconstructed Energy (GeV), for the LEM>0.7 signal-enriched region. These values are shown for each of the three analysis runs (1, 2, and 3456). The error bars indicate the total systematic error on the extrapolation; the contributions to this error will be discussed in Chapter 8.

and Far Detectors. Beam ν_es are produced by secondary muon decays, which will occur further along the beam pipe than pion and kaon decays. This gives the Near Detector a different angular acceptance than for ν_μ CC events. The Far/Near ratios for all three components are therefore different.

Several other effects also contribute to an overall Far/Near difference. The factor of 10^6 difference in neutrino flux between the two detectors has already been described (see Section 3.5.1). Far/Near differences also arise from physical differences between the two detectors, in fiducial volume and the readout system. All of these differences are folded into the Monte Carlo Far/Near ratio in equation 7.1. The Far/Near ratio as a function of reconstructed energy for each background is shown in Figure 7.1.

7.2 Predicting ν_τ CC and signal ν_e CC appearance

While NC, ν_μ CC, and beam ν_e CC events can be predicted with relative ease from the Near Detector, there are two other event types which only exist in the Far Detector: ν_τ CC and signal ν_e CC. Both of these event types are the product of $\nu_\mu \rightarrow \nu_\ell$ oscillation between the Near and the Far Detector ($\ell = e$ or τ). The flux of ν_ℓ therefore depends on the total flux of ν_μ CC events in the ND fiducial volume, rather than the Preselected HOOHE rates. These ν_μ CC events are selected in a manner similar to the methods used in the ν_μ disappearance analysis [53], with basic cuts on muon track topology, event position, and a particle ID cut. This ν_μ CC selection is applied to the ND Data and to the FD and ND MC. The ND Data ν_μ CC fiducial selection is then extrapolated to the Far Detector as in 7.1 to make a prediction of the reconstructed ν_μ CC spectrum in the Far Detector.

This reconstructed ν_μ energy is then turned into a ν_e or ν_τ appearance signal in a process with three major steps:

1. **The ν_μ reconstructed spectrum is transformed into a ν_μ true spectrum**. This is done by correcting the FD ν_μ CC prediction with a two-dimensional matrix $\mathrm{RT}^{\nu_\mu \mathrm{CC}}$, which transforms reconstructed to true energy. The spectrum is also corrected for the purity P and the efficiency E of the ν_μ CC selection:

$$\mathrm{F}_j^{\nu_\mu \mathrm{CC}} = \left[\sum_k \left[\frac{\mathrm{f}_k^{\nu_\mu \mathrm{CC}}}{\mathrm{n}_k^{\nu_\mu \mathrm{CC}}} \mathrm{N}_k^{\nu_\mu \mathrm{CC}} \right] \mathrm{RT}_{kj}^{\nu_\mu \mathrm{CC}} \right] \frac{\mathrm{P}_j}{\mathrm{E}_j} \tag{7.2}$$

   where j is here a bin of true energy, k a bin of reconstructed energy, and $\mathrm{f}_k^{\nu_\mu \mathrm{CC}}$, $\mathrm{n}_k^{\nu_\mu \mathrm{CC}}$, and $\mathrm{N}_k^{\nu_\mu \mathrm{CC}}$ are the respective CC selections for FD MC, ND MC, and ND data.

2. **The ν_μ true spectrum is oscillated to make a true ν_e or ν_τ spectrum.** The true ν_μ spectrum is reweighted for the probability of $\mathrm{P}(\nu_\mu \rightarrow \nu_\ell)$ oscillation. This is transformed into true ν_ℓ CC energy by scaling to the correct cross-section, i.e. σ^ℓ instead of σ^μ:

$$F_j^{\nu_\ell\, CC} = \mathrm{F}_j^{\nu_\mu \mathrm{CC}} \mathrm{P}_j^{\nu_\mu \rightarrow \nu_\ell} \frac{\sigma_j^\ell}{\sigma_j^\mu} \tag{7.3}$$

3. **The true ν_e or ν_τ spectrum is turned into a reconstructed spectrum.** A second transfer matrix TR^{ν_ℓ} is applied to convert true ν_ℓ energy back into reconstructed energy. The final signal is corrected for the efficiency ϵ of LEM (in

PID bin i) as a function of reconstructed energy:

$$\mathrm{F}_{ik}^{\nu_\ell \mathrm{CC}} = \sum_j F_j^{\nu_\ell\ \mathrm{CC}} \mathrm{TR}_{kj}^{\nu_\ell} \epsilon_{ik}^{\nu_\ell} \tag{7.4}$$

The output of these steps is $\mathrm{F}_{ik}^{\nu_\ell \mathrm{CC}}$, a prediction of the number of ν_e or ν_τ events selected in reconstructed energy bin k and LEM bin i.

7.3 Corrections to the Extrapolation

7.3.1 Corrections to the Signal Prediction: the MRE Correction

An additional correction must be applied to the ν_e CC appearance prediction. The mismodeling of the hadronic shower component in MC ν_e CC events could cause the efficiency term ϵ in Equation 7.4 to be incorrect. A selection of ν_μ CC-derived hadronic showers with an injected EM shower are therefore used to make a data-based correction to correct for this mismodeling. The hadronic showers are taken from known ν_μ CC events possessing an identifiable muon track. The muon is removed from the shower, and its momentum and vertex used to produce an electron instead. A MC electromagnetic shower is generated from this electron and merged with the “muon-removed” hadronic shower. The new event is passed through the reconstruction and analysis chain again. Finally, a series of cuts are applied to these events, ensuring that the original event was a good-quality ν_μ CC event. This new sample is known as the Muon Removed, Electron Added, or MRE sample [121].

The MRE sample is produced for both ND Data and Monte Carlo hadronic showers, and the selection efficiency is then calculated for different value of the LEM PID. The ratio of the ND Data and ND MC efficiencies, $\epsilon_{Data}/\epsilon_{MC}$ is then used as a correction factor for ϵ^{ν_e} in Equation 7.4. As an example, the LEM>0.70 PID cut has an efficiency of 35.5% for the ND Data MRE sample, and an efficiency of 36.5% for the ND MC MRE sample. The MRE correction factor is then C_{MRE}=0.97. Table 7.1 shows the size of the MRE correction factor [122] for a variety of ranges of the LEM PID.

It is interesting to compare the effects of the MRE correction to the effects of the Model 1 hadronic reweighting studied in the previous chapter. To do this, the ratio (M1 Reweighted)/(Standard) is calculated for standard FD MC ν_e CC signal events for a given range of PID. This ratio is shown in comparison to the MRE correction in Table

LEM Bin	MRE correction	Signal: (Model 1)/Standard
>0.70	0.97	0.98
0-0.2	0.93	1.06
0.2-0.3	0.95	1.07
0.3-0.4	0.96	1.04
0.4-0.5	0.98	1.02
0.5-0.6	0.97	1.01
0.6-0.7	0.97	0.98
0.7-0.8	0.97	0.98
>0.80	0.99	0.98

Table 7.1: Size of MRE correction factor for various ranges in the LEM particle ID. The effect of the Model 1 reweighting on ν_e CC FD MC events is also given as a point of comparison, in the final column.

7.1. In the high PID region, where Model 1 is known to have the largest effect, the effects of the MRE correction are comparable to the effects of the Model 1 reweight. This suggests that the MRE correction reflects many of the same hadronic modeling issues that were discussed in the previous chapter. The similarity between the MRE correction and Model 1 breaks down for lower PID values, however, suggesting that other hadronic effects cause the discrepancy for this range of the PID.

7.3.2 Other backgrounds: Cosmic Rays

The cosmic ray cuts described as part of the Preselection should in theory remove all background events from cosmic rays. The effectiveness of these cuts can be measured by looking at the number of events selected in the NuMI anti-coincidence spills, i.e. in the windows out of time with the beam [123]. Out of the total 43 million anti-coincidence spills analyzed, 258 events are selected at Preselection level, with 72 having LEM>0.70. When the cosmic ray cuts are then applied, the Preselection selects 4 events, with none passing LEM>0.7. These numbers are then normalized to the total data-taking period of the in-time beam spills. After this normalization, it was determined that cosmic rays add <0.5 events to the Preselection sample, and <0.3 events to the LEM>0.70 signal-enriched region. This contamination can be disregarded.

7.3.3 Other backgrounds: Rock Interactions

The NuMI beam spreads out substantially as it travels, so that in both the Far and Near Detectors, neutrinos from the NuMI beamline can interact in the rock surrounding the cavern. Daughter particles from these rock interactions, such as muons and hadrons, continue to travel forward, potentially into the MINOS detectors, where they can be reconstructed as neutrino interactions. These so-called rock events are typically found either at the front or the outer radius of the detector, and tend to have low energies (because they are partially-reconstructed events). Rock muons are not be accounted for in the Far/Near background extrapolation. Because the Near Detector fiducial volume uses only a fraction of the total detector volume, there is a self-shielding effect as the non-fiducial steel stops many of the particles. The Far Detector fiducial volume, however, extends nearly to the edge of the steel planes, providing far less shielding and leaving the detector open to background from rock events. A correction must be applied to the FD prediction to account for this.

The rock event background was assessed using special Far Detector Monte Carlo. In this simulation [124], the cavern is modeled as a plain rectangular box, surrounded by Ely greenstone rock. Adding more detailed features to the cavern shape has a very small effect on the simulation. As shown in Figure 7.2, most of the Preselected rock events occur near the edge of the fiducial volume. Table 7.2 lists the total number of rock events versus the total number of detector background events (raw FD MC) for various selection levels. Figure 7.3 shows the distribution of the rock muon events as a function of reconstructed energy and LEM PID value. The rock events are also shown as a fractional excess on the standard detector FD MC background sample.

As might be expected, most of the rock events are concentrated at low energy. The events also tend to have low PID values. Very few rock events are selected, and these events form a negligible fraction of the LEM>0.70 selection. The fractional excess in the low energy region, however, was considered large enough to add a correction to the extrapolation. The fractional rock muon background is therefore calculated for each background type and used as a correction factor in the final prediction.

7.4 Predictions for the Far Detector

Using the tools outlined in this chapter, a prediction can now be made of the Far Detector background and signal. Unlike the earlier LEM>0.70 studies, this extrapolation is now performed in bins of both PID and energy, to predict the full shape of the distribution.

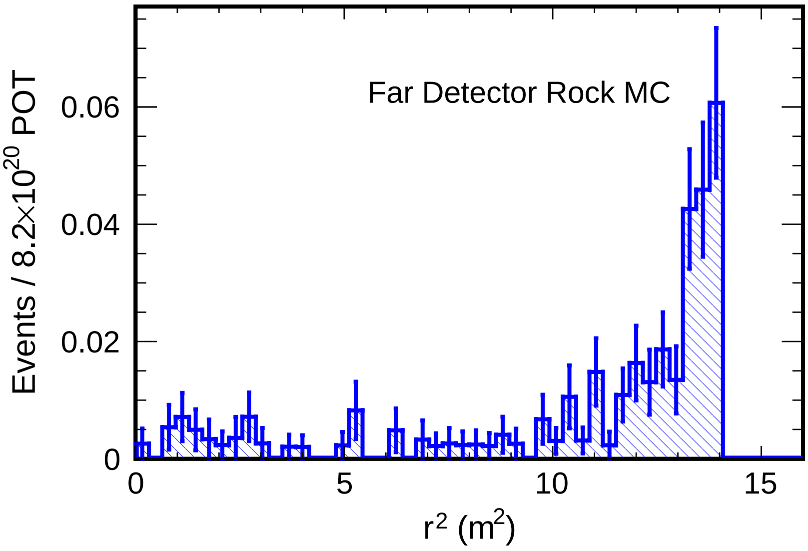

Figure 7.2: Expected rock muon background (Far Detector Monte Carlo) as a function of radius squared for all preselected events.

Selection	Standard FD MC	Rock Muon FD MC	% Contamination
Fiducial	3173.3	93.0	2.9%
Preselection	553.0	2.5	0.5%
LEM > 0.70	61.7	0.1	0.0%

Table 7.2: Number of FD MC rock muon events selected in the Far Detector, versus number of standard FD MC background events. No PORP weights are applied. Numbers are for an exposure of 8.2 $\times 10^{20}$ POTs.

The reconstructed energy is extrapolated in the HOOHE decomposition binning. The LEM PID is extrapolated in 8 bins, with edges {0.0, 0.2, 0.3, 0.4, 0.5 ,0.6, 0.7, 0.8, 1.0}. This binning will be justified in Chapter 9. The extrapolation is performed separately in the three data-taking runs, and then summed to form a FD prediction. Figure 7.4 shows the resulting FD prediction as a projection of the full LEM PID distribution. Figure 7.5 shows the prediction as a projection of the reconstructed energy distribution. The prediction in each PID bin can be found in Table 7.3 for both a zero and non-zero θ_{13}.

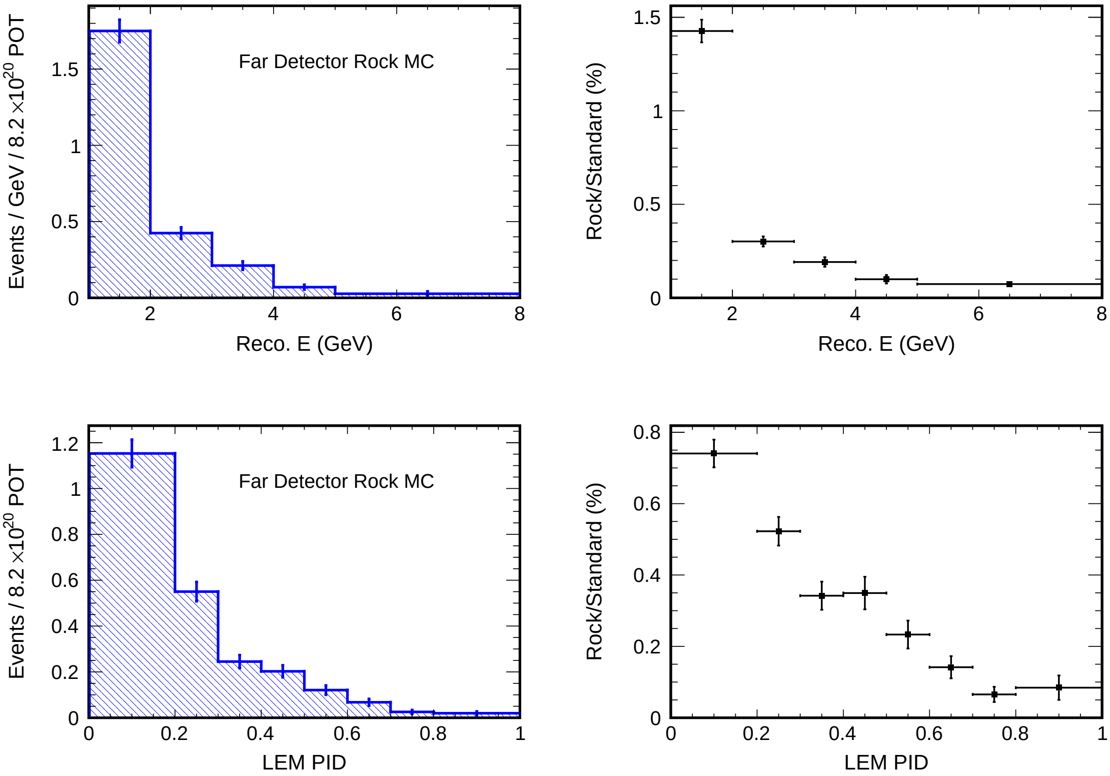

Figure 7.3: Expected rock muon background (Far Detector Monte Carlo), in bins of reconstructed energy (top) and LEM (bottom). The left plots show the total number of MC rock events selected for a 8.2 $\times 10^{20}$ POT exposure. The right plots show these rock event distributions as a fractional excess on top of the total FD MC background sample.

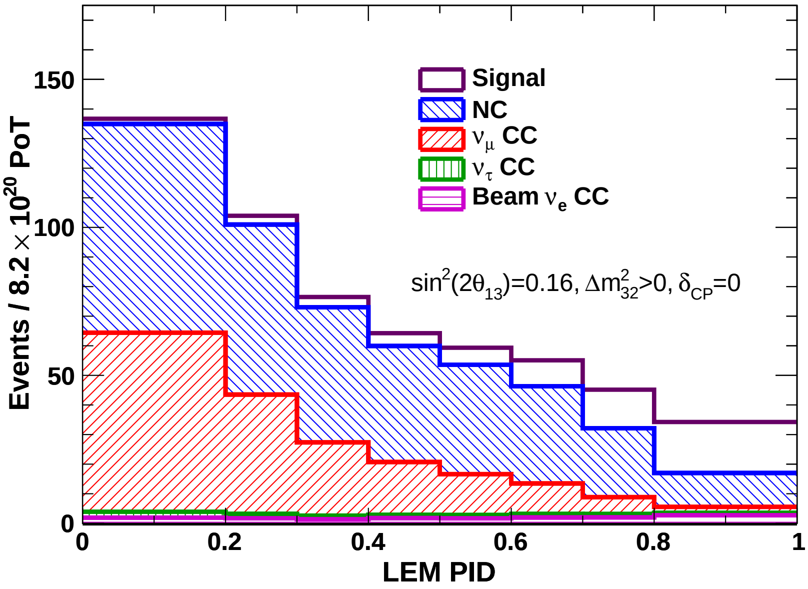

Figure 7.4: PID distribution prediction for an 8 LEM × 5 E bin extrapolation. Signal is CHOOZ-sized ($\sin^2 2\theta_{13} = 0.16$). This prediction is for an exposure of 8.2 $\times 10^{20}$ POT.

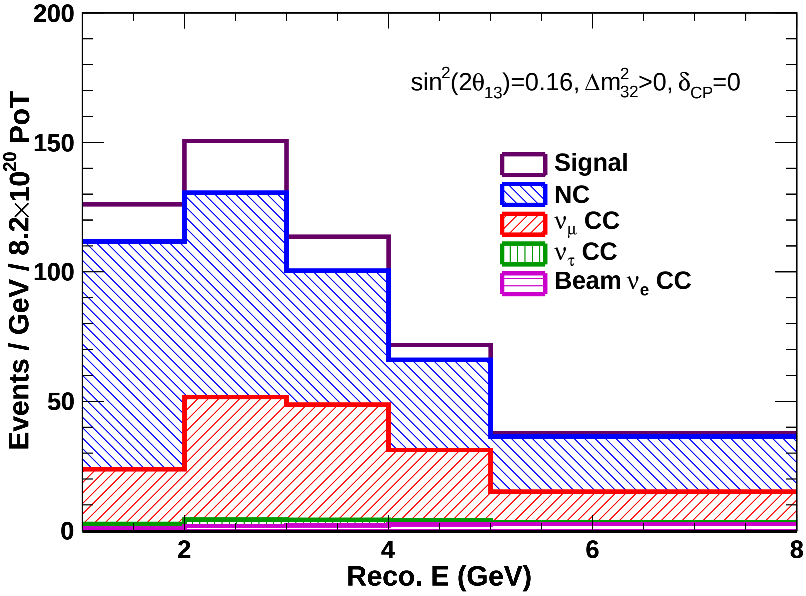

Figure 7.5: Reconstructed energy distribution prediction for an 8 LEM × 5 E bin extrapolation. This plot is equivalent to the Preselection reconstructed energy distribution. Signal is CHOOZ-sized ($\sin^2 2\theta_{13} = 0.16$). This prediction is for an exposure of 8.2 $\times 10^{20}$ POT.

	LEM Particle ID Bin Lower Edge								
Event Type	0.0	0.2	0.3	0.4	0.5	0.6	0.7	0.8	>0.70
$\sin^2 2\theta_{13} = 0.0$									
NC	72.32	58.12	46.27	39.97	37.39	33.59	23.80	11.31	35.11
ν_μ CC	60.92	40.78	25.01	18.03	13.92	9.92	5.44	1.64	7.09
beam ν_e CC	1.12	1.58	1.21	1.50	1.55	2.12	2.07	3.46	5.53
ν_τ CC	2.24	1.66	1.31	1.22	1.25	1.33	1.27	0.92	2.19
Total Bkgd.	136.60	102.14	73.79	60.72	54.11	46.97	32.59	17.33	49.92
Signal	0.01	0.02	0.02	0.03	0.04	0.06	0.08	0.11	0.19
$\sin^2 2\theta_{13} = 0.16$									
NC	72.32	58.12	46.27	39.97	37.39	33.59	23.80	11.31	35.11
ν_μ CC	60.77	40.64	24.92	17.96	13.86	9.88	5.42	1.64	7.05
beam ν_e CC	1.10	1.53	1.15	1.43	1.46	2.02	1.99	3.22	5.21
ν_τ CC	2.05	1.52	1.20	1.12	1.14	1.22	1.17	0.84	2.00
Total Bkgd.	136.24	101.81	73.54	60.48	53.86	46.71	32.37	17.01	49.38
Signal	1.65	2.63	3.11	4.05	5.50	8.42	12.54	17.11	29.65

Table 7.3: Far Detector predictions in individual PID bins, for both an 8 bin full-shape LEM, and a single LEM>0.7 cut. Predictions are given for all event types, for both no signal ($\sin^2 2\theta_{13} = 0.0$) and a CHOOZ-size signal ($\sin^2 2\theta_{13} = 0.16$)

Chapter 8

Systematic error

The previous chapter discussed the means by which a Near Detector selection is turned into a Far Detector prediction. During this process, most potential errors in physics modeling, calibration, etc. cancel out to first order in the Far/Near MC ratio. This does not mean that the measurement is free from systematic error. Some uncertainties affect both detectors and will have a second-order impact on the final background prediction. Other systematic errors may predominantly affect a single detector, and therefore have a direct effect on the prediction. Finally, some systematic errors will specifically affect the prediction of ν_e and ν_τ appearance. The systematic effects described in this chapter were identified and evaluated for the 2011 ν_e analysis. A more detailed description of each error, and several which were ultimately not included in this analysis, can be found in [125].

8.1 Simulating Systematic Error

The vast majority of systematic errors were simulated by shifting the relevant parameter by one standard deviation. This shift was simulated either by reweighting existing MC, or producing new reconstructed MC events. Either way, these shifted MC samples are used to produce a new FD prediction. The fractional change in the prediction due to the systematic uncertainty is then calculated:

$$\mathrm{R}_{bkgd} = \frac{\mathrm{P}^{+1\sigma}_{\mathrm{NC}} + \mathrm{P}^{+1\sigma}_{\nu_\mu\,\mathrm{CC}} + \mathrm{P}^{+1\sigma}_{\mathrm{beam}\,\nu_e\,\mathrm{CC}}}{\mathrm{P}^{0}_{\mathrm{NC}} + \mathrm{P}^{0}_{\nu_\mu\,\mathrm{CC}} + \mathrm{P}^{0}_{\mathrm{beam}\,\nu_e\,\mathrm{CC}}} \tag{8.1}$$

$$\mathrm{R}_{\nu_\tau} = \frac{\mathrm{P}^{+1\sigma}_{\nu_\tau\,\mathrm{CC}}}{\mathrm{P}^{0}_{\nu_\tau\,\mathrm{CC}}} \tag{8.2}$$

$$\mathrm{R}_{\nu_e} = \frac{\mathrm{P}^{+1\sigma}_{\nu_\tau\,\mathrm{CC}}}{\mathrm{P}^{0}_{\nu_e\,\mathrm{CC}}} \tag{8.3}$$

where $\mathrm{P}^{+1\sigma}_{\alpha}$ is the prediction for for background type α (NC + ν_μ CC + beam ν_e CC background, ν_τ CC, or signal ν_e CC) shifted by one standard deviation for a given systematic error, and P^{0}_{α} is the nominal, unshifted prediction, and R_α is the resulting fractional shift in the prediction.

In the case of the NC + ν_μ CC + ν_e CC background error, the systematic shifts affect the Far/Near MC extrapolation ratio. In the case of the ν_e and ν_τ CC errors, there are two types of effects. First, the same uncertainties which affect the Far/Near ratio also affect the selection efficiencies ϵ^{ν_e} and ϵ^{ν_τ} (in Equation 7.4). Second, the underlying ν_μ CC spectrum determining the ν_τ and ν_e CC flux is also affected by several systematic errors specific to the ν_μ CC sample.

In this chapter, each uncertainty will be characterized by the fractional shifts R_{bkgd}, R_{ν_e}, and R_{ν_τ} it produces in the LEM>0.7 prediction. For the largest systematic errors, plots will be shown of these fractional shift as a function of PID value. The impact of the systematic errors on the full PID and energy distribution will be further quantified in the discussion of the multiple bin likelihood fit framework in Chapter 9.

8.2 Direct Far/Near Uncertainties

The largest systematic errors are those which directly impact the Far/Near ratio, and therefore are less likely to be canceled out in the course of the extrapolation. These include uncertainties in the relative normalization of the two detectors, the calibration, and the modeling of cross-talk.

8.2.1 Normalization

Several sources of uncertainty affect the normalization of the prediction. These include uncertainties on the exact geometry and mass of the detectors, the total counted POTs, and the live time of the Near and Far Detectors. A comprehensive study of these factors [126] found a 0.32% uncertainty on the overall normalization due to the total measured POT exposure, a 0.2% uncertainty from the steel thickness, and a 0.2% uncertainty from the scintillator thickness.

Another source of normalization error comes from uncertainty in the relative number of events accepted by the Near and Far fiducial volume cuts. This uncertainty is evaluated by comparing the Far/Near ratio between Data and MC for muon-removed shower events (MRCC events, which are essentially MRE events without an injected electron). From this study, there was found to be a 2.0% uncertainty in the fiducial volume selection. The combined effects of POT counting, detector components, and fiducial volume uncertainties result in a 2.0% normalization systematic error on the background. Additional uncertainty in the selection of ν_μ CC fiducial events raises this uncertainty to 2.4% for ν_e and ν_τ appearance events.

8.2.2 Cross-talk mismodeling

Much of the discriminating power of the LEM PID comes from its direct use of the event's strip hits. This gives rise to another potential source of systematic error, resulting from cross-talk mismodeling. Studies [128] have found a large Data/MC disagreement for detector hits below 1 PE. Most hits below 2 PE are caused by PMT cross-talk, suggesting that this aspect of the simulation may be imperfect. Because Near and Far Detector cross-talk behave according to two different PMT readout maps, this mismodeling could pose a significant source of Far/Near systematic error. In order to mitigate these effects, the reconstruction already disregards any hits below 2 PE, and the LEM PID does not use any hits below 3 PE, but there could still be residual effects.

Cross-talk is produced in two separate forms: electrical and optical. The effects of electrical cross-talk were found to be small enough to disregard. The effects of optical cross-talk are assessed by comparing the prediction from Monte Carlo using the current cross-talk map, to that from Monte Carlo produced with a more accurately simulated map. The effect of this change on the LEM>0.70 prediction is on the order of 2%. As will be described in Section 8.6, much of this 2% shift comes from statistical fluctuations in the shifted sample. A linear parametrization will be needed to better approximate this systematic error.

8.2.3 Calibration

Calibration uncertainties constitute one of the largest possible sources of systematic error in the analysis. The calibration process not only determines the energy of the final event, but also can affect its topology and reconstruction. Resulting uncertainties in light level can change the LEM match list, and the final PID value. Calibration systematic errors were mostly evaluated using specially generated Monte Carlo samples, which were reconstructed with a given calibration parameter shifted by one standard deviation. The following sources of uncertainty were investigated:

Relative ND and FD Energy Scale The largest single calibration error comes from uncertainty in the relative energy scale of the Near and Far Detectors. This systematic error arises from several uncertainties in the inter-detector energy calibration, including uncertainties on the muon stopping power, drift, detector aging, and random noise. This uncertainty is evaluated separately for the two detectors. The Near Detector MC energy scale (the factor which converts the pulse height to MIPs) is shifted by its uncertainty, 1.9%. The Far Detector MC energy scale is shifted by 0.9%. Both uncertainties have a large effect on the extrapolated background prediction. For the LEM >0.70 sample, the ND energy scale uncertainty leads to a −3.8% systematic effect on the total background prediction. The FD energy scale uncertainty leads to a 1.6% systematic effect. The fractional shift on the background prediction as a function of PID can be seen in Figure 8.1. Because these systematic errors only affect the extrapolated energy distribution, they were not evaluated for the ν_τ or ν_e CC appearance predictions.

ND and FD Gains Another possible source of Far/Near systematic error comes from the PMT gain calibration. There is a 5% overall uncertainty on PMT gain measurements, and a 7% random variation in gains between PMTs [94]. While these gains are not used in the calculation of event energy, they are used to convert ADCs to PEs for each hit. Therefore, the uncertainty can have an impact on various analysis tasks such as the reconstruction's 2 PE cut. The effect of this uncertainty is evaluated by producing a new MC sample in which the average PMT gain is shifted by 5%. The gains of the individual PMTs are also given a 7% random Gaussian variation. This is done separately for the Near Detector and Far Detector. The Near Detector Gain uncertainty has a systematic effect of -0.5% on the total predicted background, and the Far Detector Gain systematic error a 0.7% effect. The FD gain error also has a -1.0% effect on the ν_τ CC prediction, and a 0.2% effect on the ν_e CC signal prediction.

Linearity Correction When the PMT linearity is measured, there is an associated uncertainty on the measurement. In order to propagate these uncertainties to the prediction, these uncertainties are used set the width of a Gaussian distribution, which is then used to randomly vary the linearity correction in a new MC sample. This uncertainty has a <1% effect on the LEM>0.70 FD prediction.

Attenuation Correction Following the attenuation corrections, detector response may still vary by approximately 1% along the length of the scintillator strip. This variation occurs in data, but not in the Monte Carlo (which is perfectly corrected). To assess the impact of these attenuation uncertainties, Monte Carlo events are given the same attentuation factor as the data and then recalibrated. This uncertainty has a <1% impact on the LEM>0.70 prediction.

Strip-to-Strip Correction Some variations in response between strips on the order of 0.5% from the mean are still expected even after the strip-to-strip corrections. This uncertainty is simulated by fluctuating the strip response within a Gaussian distribution of width 0.5%. This variation has a <1% effect on the LEM>0.70 prediction.

Absolute Energy Scale There is a 5.7% uncertainty on the absolute energy scale in both detectors [127]. To approximate this uncertainty, a new Monte Carlo sample was reconstructed with this scale shifted by one standard deviation. This uncertainty produces a relatively large uncertainty on the LEM>0.7 background prediction, on the order of 1%. The resulting fractional shift on the prediction, as a function of PID, is shown in Figure 8.2.

Hadronic Energy Scale There is a potential uncertainty on the relative hadronic and electromagnetic energy scales. Data from the Calibration Detector suggests that the error on this relative scale is on the order of ±5%. This uncertainty is simulated by both reweighting events based on their hadronic energy, and by shifting the total energy by 5%. The reweighting determines the effect on the selection efficiency, while the energy shift determines the effect on the reconstructed energy. The uncertainty has a total effect of <1% on the background prediction.

The total systematic uncertainty due to calibration, and the contributions of each of the individual uncertainties, are summarized in Table 8.1. Calibration uncertainties have the largest impact on the LEM>0.70 prediction of any systematic error type, with an error on the order of 4% for the background and signal prediction, and 10% for the ν_τ CC prediction.

8.3 Uncertainties affecting both detectors

The previous section outlined the most serious potential sources of systematic error, in which the uncertainty directly impacts the Far/Near ratio in the extrapolation. Other

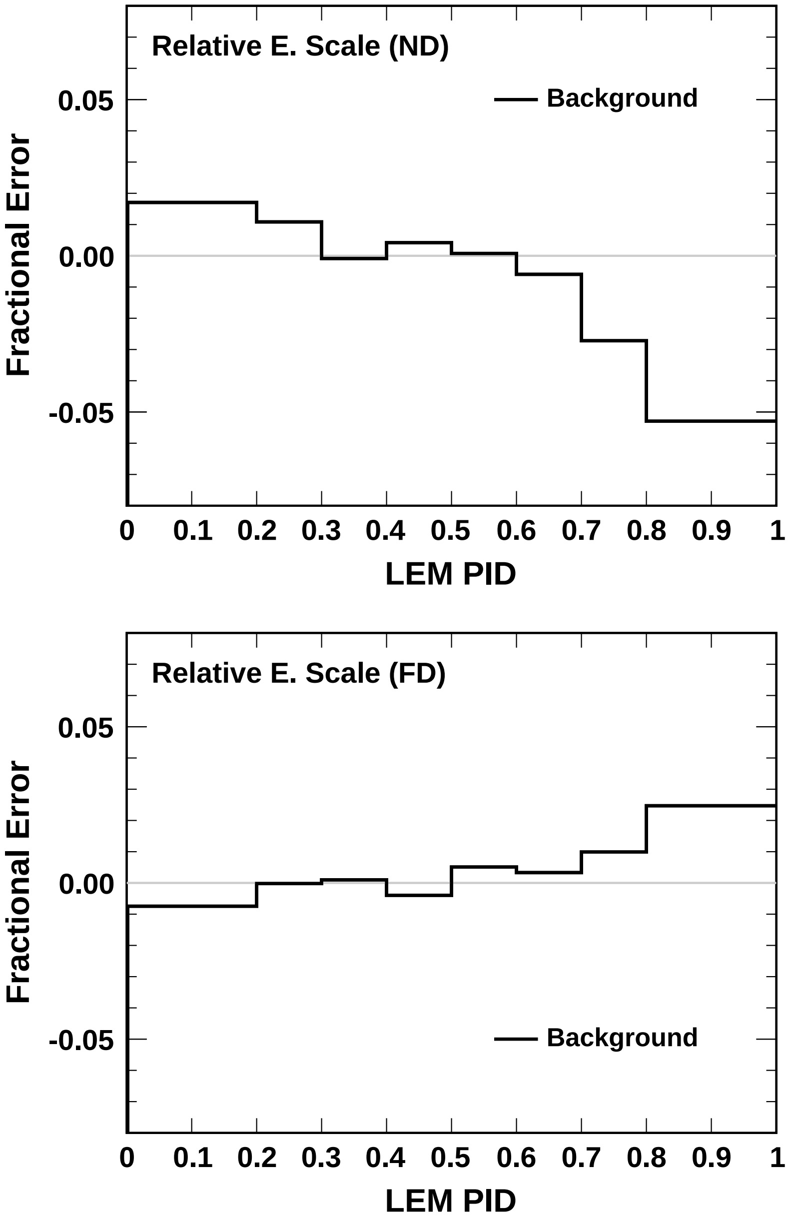

Figure 8.1: Fractional shift in the background prediction (including ν_τ CC) due to uncertainties in the relative energy scale, as a function of particle ID (LEM). The top plot shows the effects of a one standard deviation uncertainty in the Near Detector energy scale, and the bottom plot the effects of a one standard deviation uncertainty in the Far Detector energy scale.

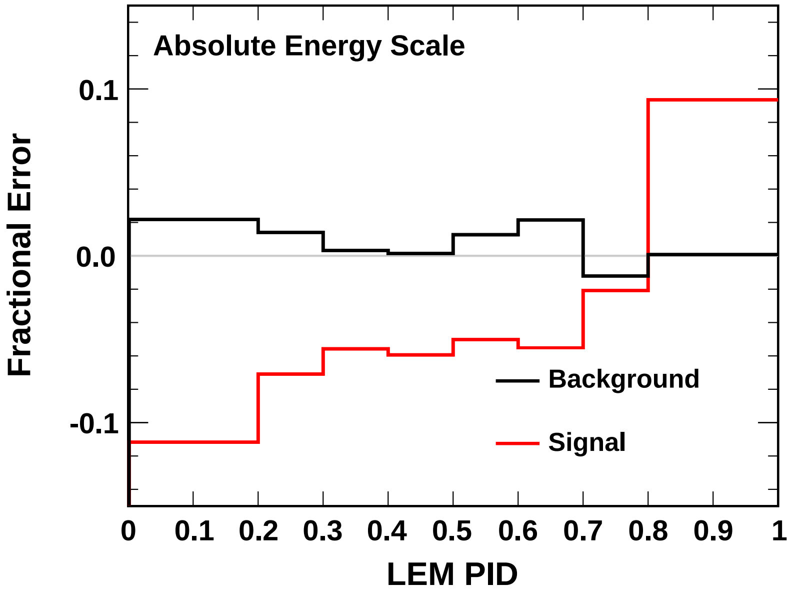

Figure 8.2: Fractional shift in the background (including ν_τ CC) (black) and signal (red) predictions due to uncertainty in the absolute energy scale, as a function of particle ID (LEM). The plot shows the impact on the prediction when the absolute energy scale in the Near and Far MC is shifted by one standard deviation.

Calibration Errors			
Syst.	F/N Bkg.	ν_e CC	ν_τ CC
Attenuation	-0.11%	0.05%	-0.02%
Gain (F $\pm1\sigma$)	0.65%	0.21%	-1.03%
Gain (N $\pm1\sigma$)	-0.46%	–	–
Linearity	-0.21%	0.00%	-0.01%
Rel E-scale (F $\pm1\sigma$)	1.60%	–	–
Rel E-scale (N $\pm1\sigma$)	-3.81%	–	–
Strip-To-Strip	-0.12%	0.05%	-0.12%
Abs. Energy Scale	-1.21%	4.41%	9.99%
EM vs. Had. Scale	0.28%	-0.48%	-0.82%
Total Error	4.11%	4.44%	10.71%

Table 8.1: The % shift in the LEM>0.70 prediction, for F/N backgrounds (NC, ν_μ CC, and Beam ν_e CC) and the appearance predictions (signal ν_e CC and ν_τ CC.), due to various calibration systematic uncertainties. Numbers are for 1 σ of uncertainty. The systematic uncertainties are added in quadrature to form the total uncertainty. Error values denoted by "–" indicate that the systematic error was not evaluated for that event type. The FD and ND Relative Energy Scale errors and the ND Gain error were not relevant for the prediction of ν_τ and ν_e appearance.

systematic errors have an impact on both detectors and thus largely cancel out in the Far/Near extrapolation. The remaining second-order effects are evaluated below.

8.3.1 Beam Model Systematic Error

The MINOS MC beam simulation has several potential sources of systematic error. These are as follows:

SKZP Weights The ND ν_μ CC flux simulated by FLUKA disagrees with the ND Data. The SKZP weights described in Section 3.5.1 are applied to correct this discrepancy. There are several sources of uncertainty on the correction, including uncertainties in hadronic production, target position, current, and other physical parameters in the beam and target. The Monte Carlo is reweighted as a function of energy to simulate the impact of shifting the SKZP weights by one standard deviation of uncertainty. This reweighting has a <1% effect on the background prediction.

FLUGG vs. FLUKA While FLUKA is the default beam simulation, later files were also produced with FLUGG [134]. FLUGG has a beam ν_e CC flux 20% larger than FLUKA in the Near Detector and 10% larger in the Far Detector. To evaluate the impact of this uncertainty in beam ν_e flux, the standard (FLUKA) Monte Carlo is reweighted by the ratio of the FLUGG ν_e/ν_μ flux ratio to the FLUKA ν_e/ν_μ flux ratio. This systematic uncertainty, which primarily affects the beam ν_e prediction, has a less than 1% effect on the total background prediction.

Target Degradation Finally, the gradual degradation of the NuMI target was found to have a time-dependent effect on the beam flux. Simulations found that this degradation was consistent with the loss of two fins in the target, due to radiation damage. The Monte Carlo was reweighted to simulate this damage and assess the impact on the prediction. The overall effect on the LEM>0.70 background prediction was found to be less than 1%.

The contributions from beam systematic error are summarized in Table 8.2; the total impact is on the order of 1%. All three of these uncertainities have an additional effect on the ν_μ CC Flux used to calculate ν_τ and ν_e CC appearance, as will be discussed in Section 8.4.2.

Beam Model Errors	
Syst.	F/N Bkg.
FLUGG	−0.69%
SKZP	0.27%
Target Degradation	−0.10%
Total Error	0.75%

Table 8.2: The % shift in the LEM>0.70 prediction, for extrapolated F/N backgrounds (NC, ν_μ CC, and Beam ν_e CC), due to beam modeling systematic uncertainties. Numbers are for 1 σ of uncertainty. The systematic uncertainties are added in quadrature to form the total uncertainty.

8.3.2 Physics Modeling: Cross Sections

Other systematic uncertainties can also arise from uncertainties in the cross sections for neutrino interactions. While there are many different possible sources of cross section mismodeling in MINOS, most of these will cancel out in the Far/Near extrapolation.

Only a few cross section parameters were therefore identified and studied as possible sources of Far/Near systematic error in the ν_e analysis [95]. These parameters include uncertainties on the axial-vector mass term M_A for quasi-elastic and resonance interactions. The uncertainty on both M_A^{QE} and M_A^{RES} is taken to be ±15%. There is also a 50% uncertainty on the KNO parameters determining the DIS cross section. These uncertainties are simulated by reweighting events based on the ratio of a modified cross section (shifted within one standard deviation of uncertainty) to the nominal cross section. As expected, these effects largely cancel in the F/N extrapolation, and lead to a $< 0.5\%$ uncertainty on the extrapolated background prediction. These systematic errors are summarized in table 8.3. As will be discussed in Section 8.4.2, these uncertainities also have an effect on the ν_μ CC Flux used to calculate ν_e and ν_τ CC appearance.

Cross Section Errors	
Syst.	F/N Bkg.
KNO	0.18%
M_A^{QE}	-0.04%
M_A^{RES}	-0.14%
Total Error	0.22%

Table 8.3: The % shift in the LEM>0.70 prediction, for F/N backgrounds (NC, ν_μ CC, and Beam ν_e CC), due to various cross section systematic uncertainties. Numbers are for 1 σ of uncertainty. The systematic uncertainties are added in quadrature to form the total uncertainty. The effects of cross-section uncertainties on the ν_e and ν_τ appearance predictions will be discussed separately in Sections 8.4.1 and 8.4.2.

8.3.3 Physics Modeling: the AGKY model and Hadronic Showers

The impact of hadronic shower modeling on the LEM PID was discussed in some detail in the context of describing the Data/MC discrepancy in the Near Detector LEM distribution (see Section 6.2.3). These uncertainties in AGKY are modeled using a series of weights, which are now applied to both the ND and FD MC to assess the impact on the prediction. The models are here described again briefly, with an additional model assessing the impact of uncertainty in hadron multiplicity:

Baryon x_f selection: The MC is reweighted so that the selection of the baryon momentum produces showers which more closely match reality.

π^0 selection: The probability of π^0 production is shifted by +20% and −20%.

Multiplicity Correlation: The MC is reweighted to simulate a completely independent charged and neutral hadron multiplicity.

Implementation Ambiguities: A GENIE-only version of the AGKY model is simulated.

p_T Squeezing: The transverse momentum rejection factor is changed from −3.5 to −1.5, resulting in broader showers.

Isotropic 2 body decay: Two body decays are performed orthogonally to the direction of momentum transfer, rather than isotropically in the center of mass.

Hadron Multiplicity: The uncertainty on the total hadron multiplicity is evaluated using a slightly different type of reweighting histogram than the first six systematic errors. As described earlier, the charged hadron multiplicity is parametrized as a function of invariant mass W with the expression $\langle \mathrm{n_{ch}} \rangle = a + b \, \log W^2$. This is then converted to total multiplicity using $\langle \mathrm{n_{tot}} \rangle = 1.5 \langle \mathrm{n_{ch}} \rangle$. While the parameters a and b are derived from bubble chamber experiments, the resulting predictions do not always match experimental data. To simulate this uncertainty, a and b are varied over the range permitted by experiment. The events are then reweighted separately for each interaction type (νp, νn, $\bar{\nu} p$, or $\bar{\nu} n$).

The resulting fractional changes in the FD prediction as a function of PID for each of these systematic errors are shown in Figure 8.3. Assessing the impact of these systematic errors is very important, as uncertainties in AGKY have a large effect on the separate ND and FD Monte Carlo samples. The Baryon x_f (Model 1) uncertainty alone has a 20% systematic effect on ND MC for LEM>0.70. This systematic uncertainity is reduced by the Far/Near extrapolation to a secondary effect on the order of 1% for the background prediction. The error on the signal prediction is 3%. These systematic uncertainties are summarized in Table 8.4. Note that most of the AGKY models were not extended to study the impact on ν_τ CC hadronization. These uncertainties would have little effect on the final ν_τ CC systematic error, as the error on ν_τ appearance is completely dominated by cross section uncertainties (Section 8.4.1).

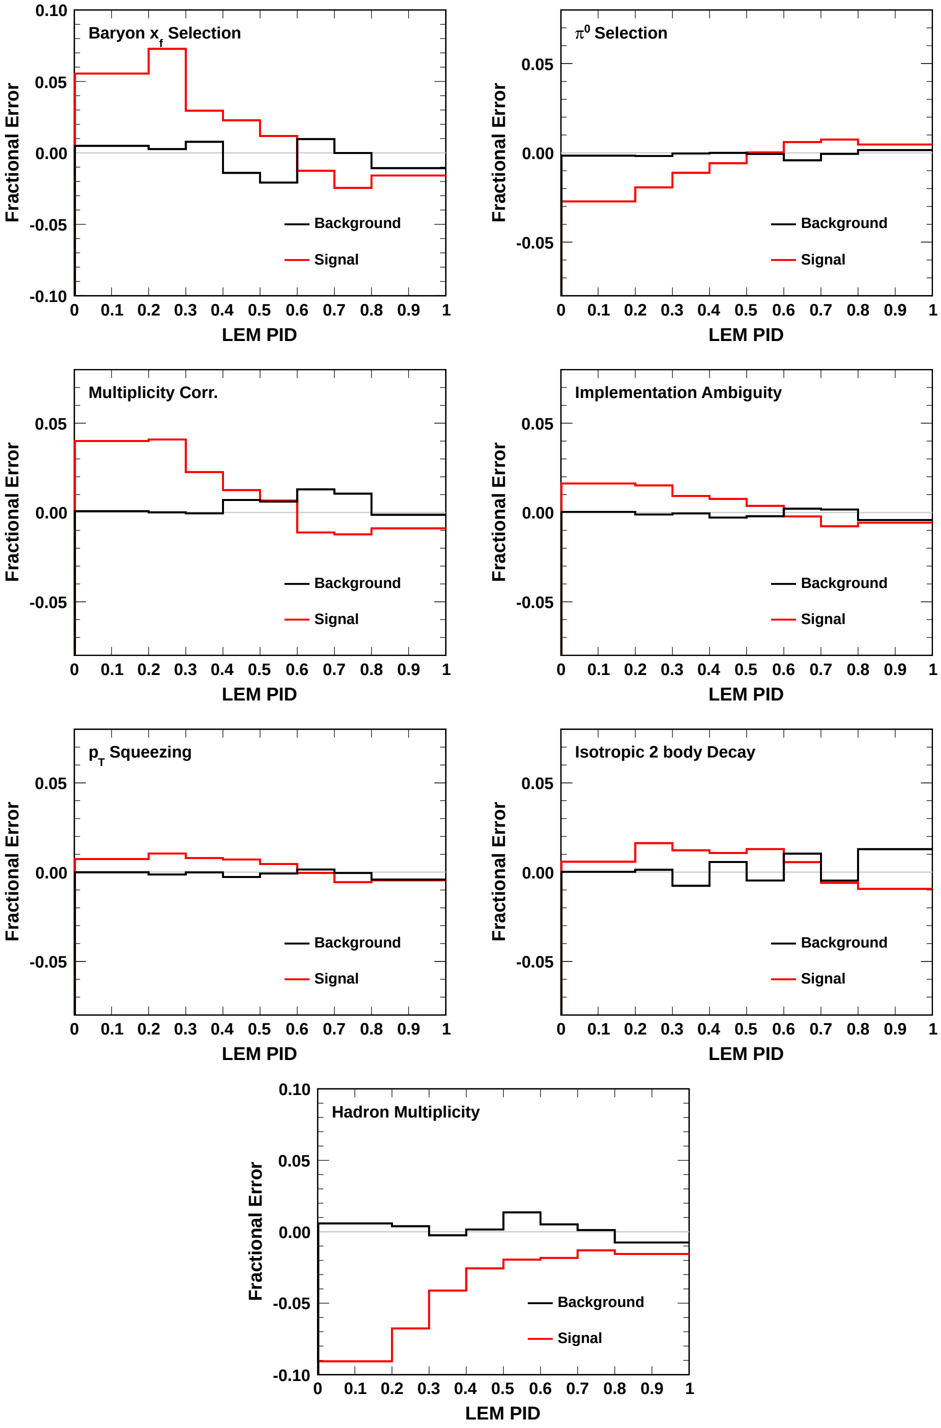

Figure 8.3: Fractional shifts in the background (including ν_τ CC) (black) and signal (red) predictions due to uncertainties in various AGKY hadronic modeling parameters, as a function of particle ID (LEM). These plots show the effects of shifting the Monte Carlo by one standard deviation for each uncertainty.

AGKY Errors				
Model	Syst.	F/N Bkg.	ν_e CC	ν_τ CC
1	Baryon x_f selection	0.40%	−1.96%	–
2+	π^0 selection (+1 σ)	0.02%	0.59%	–
2-	π^0 selection (-1 σ)	−0.02%	−0.59	–
3	Multiplicity correl.	0.66%	−1.03%	–
4	Implementation Ambiguities	−0.09%	−0.66%	–
5	p_T Squeezing	−0.23%	−0.50%	–
6	Isotropic 2 body decay	0.57%	−0.28%	–
7	Hadron Multiplicity	−0.17%	−1.44%	−0.76%
	Total Error	1.14%	2.84%	−0.76%

Table 8.4: The % shift in the LEM>0.70 prediction, for F/N backgrounds (NC, ν_μ CC, and Beam ν_e CC) and appearance predictions (ν_e CC and ν_τ CC), due to various systematic uncertainties in the AGKY hadronization model. Numbers are for 1 σ of uncertainty. The systematic uncertainties are added in quadrature to form the total uncertainty. Error values denoted by "–" indicate that the systematic error was not evaluated for that event type. ν_τ CC errors were only evaluated for Hadron Multiplicity.

8.3.4 Physics Modeling: Intranuclear Re-scattering

The final source of uncertainty in the MINOS Monte Carlo comes from uncertainties in intranuclear re-scattering, simulated by the INTRANUKE program [78]. Ten possible INTRANUKE systematic errors arising from uncertainties on cross sections or branching ratios were identified for this analysis [130]. These uncertainties were simulated by reweighting the Monte Carlo by one standard deviation for each parameter. Table 8.5 presents each of the separate INTRANUKE systematic errors, which together have a <1% effect on the prediction for LEM>0.7. A bug prevented the evaluation of these systematic errors for ν_τ CC events. As with the AGKY errors, any INTRANUKE errors would be dwarfed by the much larger ν_τ CC cross section error (Section 8.4.1).

8.4 Uncertainties on ν_e and ν_τ CC Appearance

Most of the systematic errors mentioned above affect not only the Far/Near extrapolated background, but also the ν_e and ν_τ appearance prediction. For these uncertainties, the systematic error impacts the selection efficiency for ν_e and ν_τ CC events in the Far De-

INTRANUKE Errors			
Model	Syst.	F/N Bkg.	ν_e CC
INUKE0	q-exchange	-0.03%	-0.15%
INUKE1	π elastic	-0.07%	-0.08%
INUKE2	π inel.	0.46%	0.58%
INUKE3	π abs.	-0.09%	0.58%
INUKE4	$\pi \rightarrow 2\pi$	-0.02%	-0.09%
INUKE5	N-knockout	0.13%	-0.15%
INUKE6	$N \rightarrow (2)\pi$	-0.03%	2.37%
INUKE7	form. time	-0.42%	-0.59%
INUKE8	π xsec	0.13%	0.81%
INUKE9	N xsec	-0.09%	-1.46%
	Total Error	0.66%	3.05%

Table 8.5: The % shift in the LEM>0.70 prediction, for F/N backgrounds (NC, ν_μ CC, and Beam ν_e CC) and ν_e CC appearance, due to various systematic uncertainties in the INTRANUKE hadronic intranuclear re-scattering simulation. Numbers are for 1 σ of uncertainty. The systematic uncertainties are added in quadrature to form the total uncertainty. A bug prevented the evaluation of the ν_τ CC uncertainty, but the INTRANUKE contribution is negligible when compared to the total ν_τ CC systematic error

tector. In addition to these errors, there are several additional systematic uncertainties which exclusively affect the ν_e and ν_τ samples.

8.4.1 ν_τ Cross section

There is a large degree of uncertainty on the ν_τ cross section in the MINOS MC simulation. This uncertainty is 10% for deep inelastic scattering events, and rises to 50% for quasielastic and resonance events [131] [133]. This results in an average cross section uncertainty of approximately 45% for the ν_τ CC events considered in this analysis. The large systematic error on the QE and RES interactions comes from uncertainties in the cross section's pseudoscalar form factor, which is otherwise negligible for ν_e and ν_μ CC events. The ν_τ pseudoscalar form factor is well known for low values of Q^2, but largely unmeasured at the level of the MINOS experiment [132]. In addition to these cross section uncertainties, the DIS error also includes uncertainties on other properties of the ν_τ simulation such as the impact of τ polarization [133]. The overall uncertainty on the

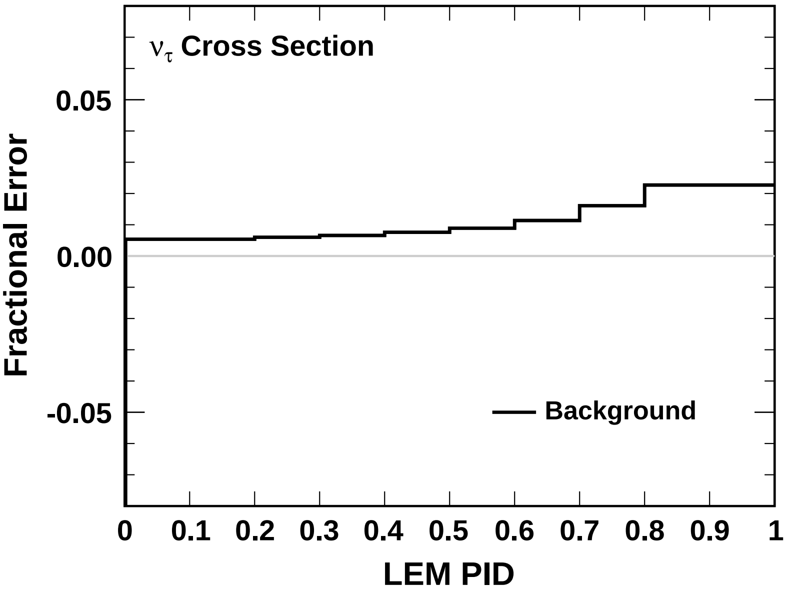

Figure 8.4: Fractional shift in the background prediction (including ν_τ CC) due to uncertainty in the ν_τ cross section as a function of particle ID (LEM). The fractional error is for a $+1\ \sigma$ shift on the cross section.

ν_τ CC cross section dominates the overall error on the ν_τ CC prediction. It is also large enough to have a >1% effect on the overall background uncertainty. The fractional shift on the total background LEM distribution due to the ν_τ CC cross section uncertainty is shown in Figure 8.4.

8.4.2 ν_μ CC Spectrum Uncertainties

Several systematic errors also affect the underlying ν_μ CC spectrum that determines the flux of the ν_τ and ν_e CC events. These errors impact the Far/Near extrapolation ratio for the ν_μ CC Fiducial samples. To simulate these errors, the FD and ND MC ν_μ CC Fiducial samples are shifted for the systematic error under consideration, and the resulting effect on the FD prediction is calculated. The relevant systematic errors were chosen from systematic error studies done for the ν_μ disappearance analysis [53]. These errors include a 50% error on the number of NC events selected in the ν_μ CC sample, and a 10% error on the ν_μ CC energy scale. The beam simulation errors in Section 8.3.1 and the cross section errors in Section 8.3.2 can also affect the ν_μ CC flux. The effects of these ν_μ CC errors are summarized in Table 8.6.

ν_μ CC Spectrum Errors		
Syst.	ν_e CC	ν_τ CC
ν_μ FLUGG Reweight	-0.01%	-0.01%
ν_μ KNO	0.35%	-0.25%
ν_μ M_A^{QE}	-0.11%	-0.37%
ν_μ M_A^{RES}	0.25%	-0.64%
ν_μ SKZP	-0.52%	0.08%
ν_μ Target Deg	0.01%	-0.30%
CC Energy Shift	3.90%	-0.92%
NC sample	-0.55%	-0.59%
Total	4.01%	1.86%

Table 8.6: The % shift in the LEM>0.70 prediction, for ν_e CC and ν_τ CC appearance, due to various systematic uncertainties on the ν_μ CC fiducial sample. Numbers are for 1 σ of uncertainty. The systematic uncertainties are added in quadrature to form the total uncertainty.

8.4.3 CalDet and MRE

Data from the Calibration Detector (CalDet) and MRE samples are used to evaluate, respectively, the uncertainty on the modeling of electromagnetic showers and the uncertainty on the hadronic components of ν_e CC events. For the EM shower uncertainty, the LEM PID is applied to data and Monte Carlo from electrons interacting in the Calibration Detector. A fitted version of this Data/MC ratio is then used to reweight the Far Detector ν_e CC distribution as a function of true electron momentum. This correction has a 1.4% effect on the predicted ν_e CC distribution. The MRE sample is used to place a similar systematic uncertainty on the hadronic component of selected ν_e CC events by shifting the ν_e selection efficiency within the statistical uncertainty on the Near Detector Data/MC MRE correction [122]. This has a 1-2% effect on the ν_e CC distribution.

The MRE and CalDet errors were both incorporated into the analysis in the form of a covariance matrix rather than a simple fractional error. This was done to account for the fact that the MRE and CalDet systematic errors were essentially statistical errors, and hence had 0% bin-to-bin correlations. For $i = j$, the elements of these matrices were $M_{ij}^{MRE} = \Delta_i^{MRE}\Delta_j^{MRE}$ and $M_{ij}^{CalDet} = \Delta_i^{CalDet}\Delta_j^{CalDet}$, where $\Delta_i^{CalDet(MRE)}$ is the shift in bin i of the FD prediction due to a 1σ CalDet or MRE error shift. For $i \neq j$, $M_{ij}^{MRE} = 0$ and $M_{ij}^{CalDet} = 0$. Both covariance matrices can then be added to the larger covariance error matrix which is used in the fit described in the next chapter.

When considered as part of a multi-bin fit, both errors had <0.1% effect on the analysis sensitivity.

8.5 HOOHE errors

In order to evaluate the total error resulting from the HOOHE decomposition, all of the Near Detector systematic errors were also evaluated for the Horn Off and High Energy samples. These systematic errors were evaluated previously, for the covariance matrix used to perform the HOOHE χ^2 fit. For each extrapolation bin i (a bin of PID and energy), the corresponding covariance matrix element for HOOHE uncertainty is calculated as follows:

$$V_{ij}^{HOOHE} = \sum_k \sum_m (P_{ik}\sigma_{ik}^{HOOHE})(P_{jm}\sigma_{jm}^{HOOHE}) \tag{8.4}$$

In this covariance matrix, the full decomposition binning (energy × PID × Run × Background type) is summed over, rather than just the extrapolation binning. Therefore, P_{ik} is the prediction in extrapolation bin i for run and background type bin k (these two variables are subsumed into a single subscript). σ_{ik}^{HOOHE} is the corresponding systematic uncertainty in that decomposition bin; this error is taken from the earlier HOOHE χ^2 fit covariance matrix. This HOOHE systematic error covariance matrix is then incorporated into the fit in the same manner as the CalDet and MRE errors. The HOOHE systematic was found to have an effect of $< 0.1\%$ on the sensitivity.

8.6 Linear Fits to Systematic Error

Most of the systematic errors discussed in this chapter were either estimated by reweighting existing Monte Carlo, or by reconstructing new Monte Carlo. In both cases, a lack of statistics in the simulation can pose a problem. This is primarily an issue for the systematic errors requiring a new reconstructed MC sample, which is often smaller in size than the nominal MC sample. To combat this lack of statistics, several of the systematic errors were approximated with a linear fit for the analysis in this thesis.

Each systematic error was plotted as the fractional shift it produces in the FD signal and background prediction, as a function of LEM PID (see Figures 8.1, 8.3, etc. for an example). From examining these distributions, several systematic errors were found to have a high degree of bin to bin fluctuation, particularly around the central 0% line. These systematic uncertainties were the strip-to-strip, optical cross-talk, and isotropic 2 body decay systematic errors. In each case, the fractional uncertainty was parametrized with a linear fit as a function of PID value. All three systematic errors chosen were essentially flat as a function of energy, so no additional linear fit along the energy axis was required. In each case, the signal (ν_e CC) and background (NC + ν_μ CC + beam ν_e CC + ν_τ CC) distributions were parametrized separately. It is necesary to separate the signal systematic error, as its impact on the total prediction is dependent on θ_{13} and δ_{CP}. Figure 8.5 shows the raw fractional uncertainty for each systematic error, along with the linear fit. For each of these systematic errors, the final systematic error is taken from this linear fit.

8.7 Total Systematic Error

The error band in Figure 8.6 shows the total systematic error on the ND MC distribution, in comparison with ND and Monte Carlo distribution, as a function of LEM PID. Figure 8.7 shows the same, but for the LEM>0.70 reconstructed energy distribution. The systematic error band amply covers the observed ND Data/MC discrepancy in nearly all cases. When the Far/Near extrapolation is performed, however, most of these errors are reduced to secondary effects, and total error drops drastically. Figures 8.8 and 8.9 again show background predictions for the LEM and reconstructed energy distributions, with much-reduced systematic error bands. Finally, Table 8.7 provides a summary of the total systematic error on the signal-enriched LEM>0.7 region. There is a total 4.8% systematic uncertainty on the extrapolated NC, ν_μ CC, and beam ν_e CC prediction, a 46.6% uncertainty on the ν_τ CC prediction, and a 7.7% uncertainty on the signal ν_e prediction.

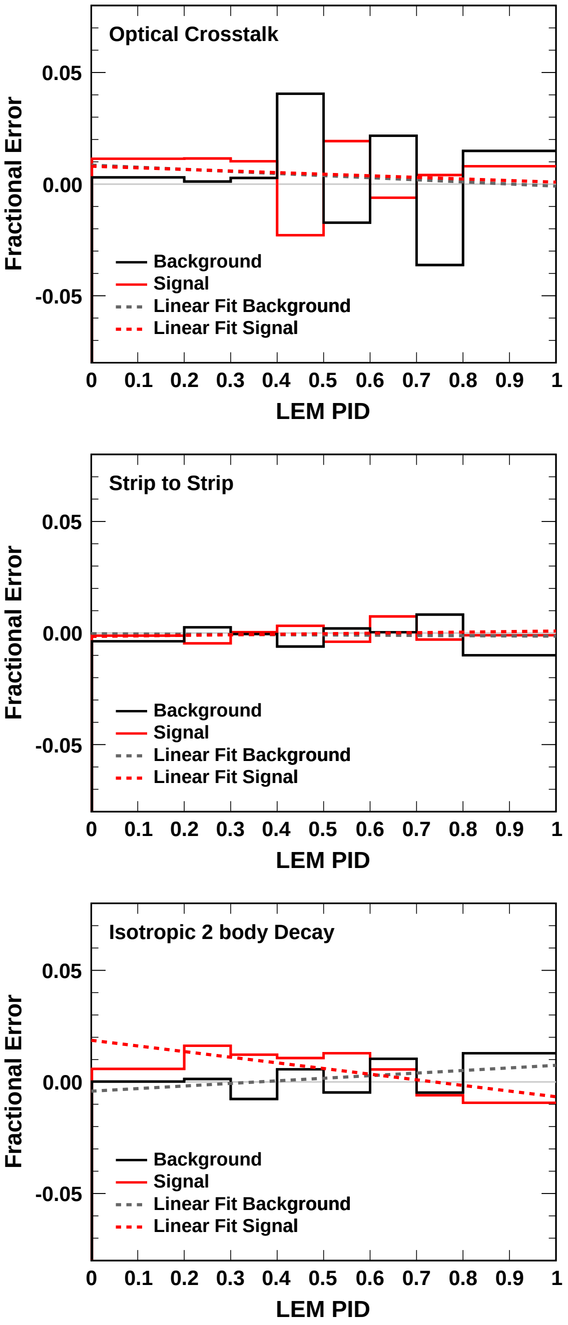

Figure 8.5: The three systematic errors shown above had a large degree of statistical fluctuation in their fractional errors, which required a linear parametrization. These systematic errors were the uncertainties on optical cross-talk (top), strip-to-strip variations (middle), and AGKY isotropic 2 body decays (bottom). The raw fractional shifts on the background prediction (black, including ν_τ CC) and signal prediction (red) are shown as a function of particle ID (LEM), along with the linear fits for each distribution (the dashed lines).

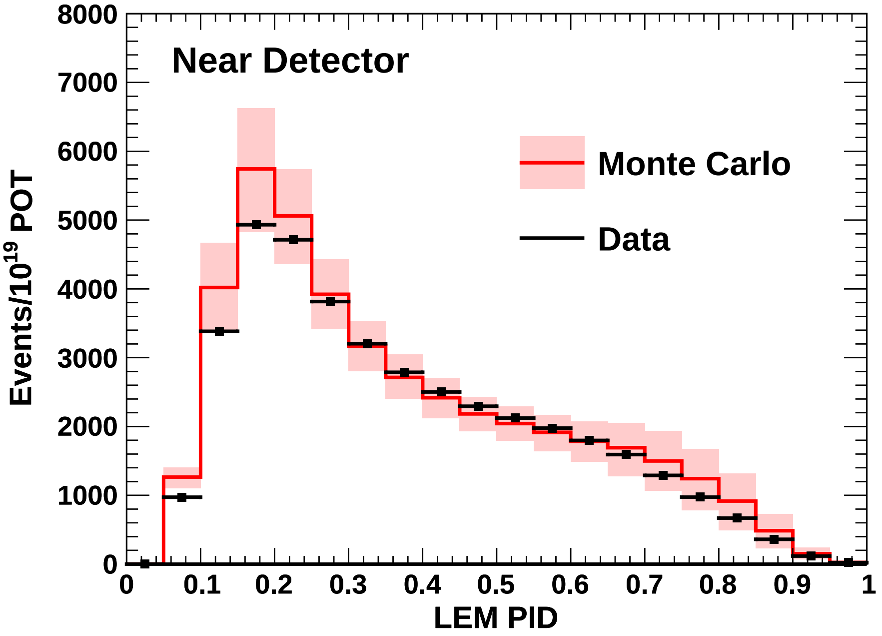

Figure 8.6: The LEM PID distribution for Data (black line) versus Monte Carlo (red line). The systematic error on the Monte Carlo is shown as a red band. This plot is normalized to 10^{19} protons on target.

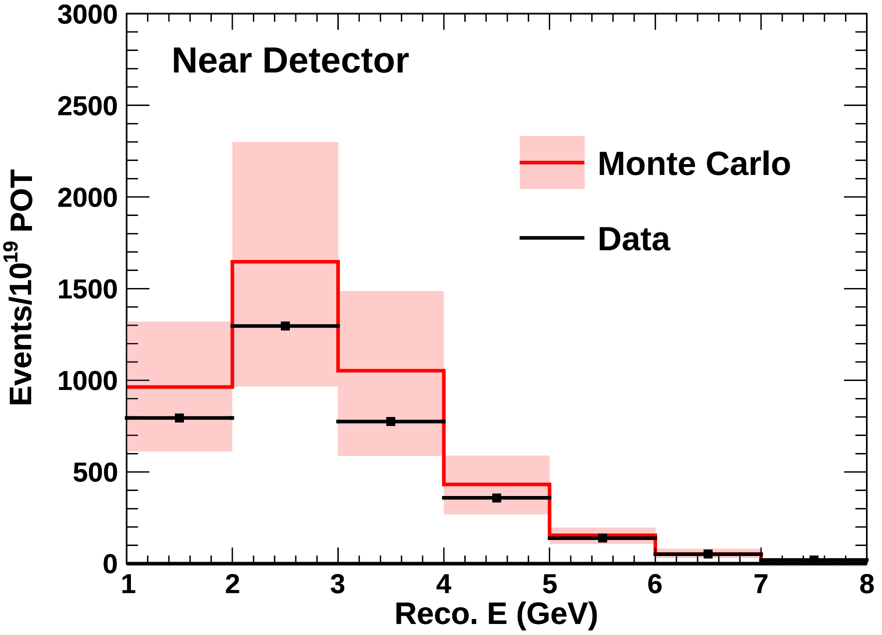

Figure 8.7: The reconstructed energy (LEM>0.70) distribution for Data (black line) versus Monte Carlo (red line). The systematic error on the Monte Carlo is shown as a red band. This plot is normalized to 10^{19} protons on target.

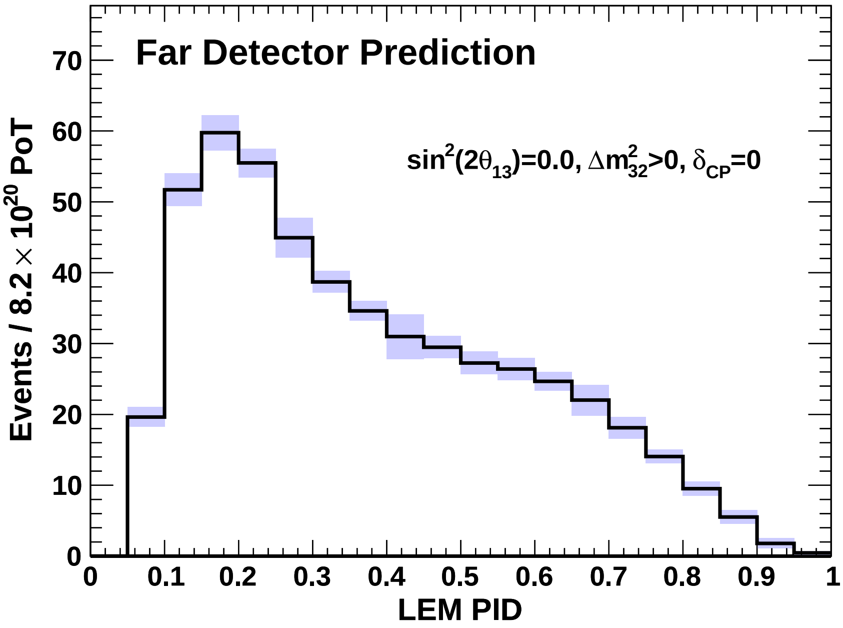

Figure 8.8: The FD prediction for the LEM PID distribution. The systematic error on the prediction is shown as a blue band. This plot is normalized to 8.2×10^{20} protons on target.

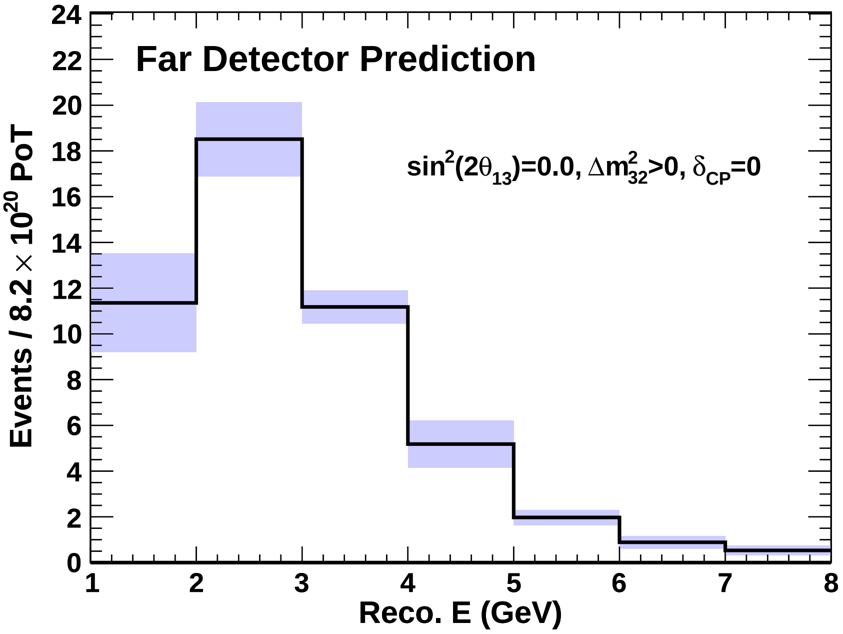

Figure 8.9: The FD prediction for the reconstructed energy (LEM>0.70) distribution. The systematic error on the prediction is shown as a blue band. This plot is normalized to 8.2×10^{20} protons on target.

Summary of Systematic Errors			
Syst.	F/N Bkg.	ν_e CC	ν_τ CC
Normalization	2.00%	2.40%	2.40%
Optical Cross-talk	0.05%	0.20%	0.05%
Calibration	4.11%	4.44%	10.71%
AGKY Model	1.14%	2.84%	−0.76%
INUKE Model	0.66%	3.05%	–
Cross Sections	0.22%	–	45.20%
Beam Model	0.75%	–	–
ν_μ CC Spectrum	–	4.01%	1.86%
Total Error	4.82%	7.68%	46.56%

Table 8.7: Total effect of systematic uncertainties on the components of the LEM >0.70 prediction, divided into in various categories of systematic error. The uncertainties are added in quadrature to calculate the total systematic error. Error values denoted by "–" indicate that the systematic error was not evaluated for that event type. The INTRANUKE uncertainties were not evaluated at all for ν_τ CC events, due to a technical problem. The ν_e CC cross section uncertainties, and the ν_e and ν_τ CC beam model uncertainties are evaluated via the underlying ν_μ CC Spectrum uncertainties.

Chapter 9

Fitting the Far Detector Data

Because the value of θ_{13} is expected to be small, and the amount of background contamination in the selected sample is significant, it is important to extract as much information as possible out of the Far Detector Data. The first two published versions of the MINOS ν_e appearance analysis were counting experiments, in which a single total number of observed events passing a PID cut (ANN11 > 0.70) was compared to an expected number of background events. As will be shown in this chapter, fitting in multiple bins (of PID and energy) increases sensitivity to $\sin^2(2\theta_{13})$ by 14%. This is a gain in sensitivity comparable to that obtained by using the LEM particle ID.

The fit described in this chapter was developed by the author for the 2011 MINOS ν_e published analysis and expanded upon for this thesis. This chapter will describe both the methods used in the fit, and the choice of optimal binning. At the end of this chapter, the projected sensitivity of this analysis to $\sin^2(2\theta_{13})$ is described.

9.1 Fitting Methodology

The number of both signal and background events involved in this analysis is small. This means that it is not possible to use a standard χ^2 fit, minimizing:

$$\chi^2 = \sum_i \frac{(n_i - \mu_i)^2}{\mu_i} \tag{9.1}$$

Instead, a likelihood method based on Poisson statistics is used. The negative of the logarithm of the likelihood is defined as:

$$-\ln L = -\sum_i \ln\left(\frac{\mu_i^{n_i} e^{-\mu_i}}{n_i!}\right) = \sum_i \ln(n_i!) + \mu_i - n_i \ln \mu_i \tag{9.2}$$

where μ_i is the expected number of events and n_i is the observed number of events in bin i. This can re-written as a likelihood ratio by subtracting out the constant self-likelihood $\ln L(\{n_i\};\{n_i\})$. The expression can then be written as follows:

$$\begin{aligned}
-\ln L \to -\ln L &= (-\ln L(\{n_i\};\{\mu_i\})) - (-\ln L(\{n_i\};\{n_i\})) \\
&= \sum_i (\ln(n_i!) + \mu_i - n_i \ln \mu_i) - \sum_i (\ln(n_i!) + n_i - n_i \ln n_i) \\
&= \sum_{i=0}^{nbins} \mu_i - n_i + n_i \ln(n_i) - n_i \ln(\mu_i) \qquad (9.3)
\end{aligned}$$

This form of the likelihood has the advantage that, in the high-statistics limit, $-2\ln L$ will approximate to a χ^2 distribution:

$$-2\ln L \approx \sum_i \frac{(n_i - \mu_i)^2}{\mu_i} = \chi^2 \qquad (9.4)$$

A fit using $-2\ln L$ from Equation 9.3 will therefore be used in this analysis.

9.1.1 Some notes about methodology

The data distribution $\{n_i\}$ is constant. In this thesis, $\{\mu_i\}$ is the background prediction, with a signal component dependent on $\sin^2(2\theta_{13})$ and δ_{CP}. For the fit, $-2\ln L$ is calculated as a function of $\sin^2(2\theta_{13})$ (for a given value of δ_{CP}). From this range of values, the minimum $-2\ln L_{min}$ is found. The value of $\sin^2(2\theta_{13})$ at this minimum is the best fit. This minimum is then subtracted to form a surface of $-2\Delta\ln L = (-2\ln L) - (-2\ln L_{min})$ as a function of $\sin^2(2\theta_{13})$. It can be shown that in the Gaussian limit for an estimate of a single parameter, 90% of the probability density function for that parameter lies within the area demarcated by $-2\Delta\ln L \leq 2.71$ [135]. This analysis has little sensitivity to δ_{CP}. In this chapter, for the purposes of determining the binning which optimizes the sensitivity to $\sin^2 2\theta_{13}$, the fit will be performed in one dimension, with δ_{CP} assumed equal to 0. This chapter will also consider only the sensitivity for the Normal Hierarchy.

This chapter will assess the performance of the fit using two separate variables:

Projected Sensitivity This variable assesses the physics reach of a fit. The observed distribution $\{n_i\}$ is chosen to be the prediction for $\theta_{13} = 0$ (i.e., the background-only prediction). The resulting upper 90% CL for $\sin^2(2\theta_{13})$ is then calculated. This upper 90% CL is the projected sensitivity. A smaller projected sensitivity is desirable, as it indicates that a fit is capable of placing a tighter limit on $\sin^2(2\theta_{13})$.

Null Exclusion This variable assesses the ability to reject the null hypothesis of $\theta_{13} = 0$. $\{n_i\}$ is chosen to be the prediction for background plus a non-zero signal with $\sin^2(2\theta_{13}) = 0.10$. The fit is performed, and the value of $-2\Delta \ln L$ at $\theta_{13} = 0$ is calculated. This value of $-2\Delta \ln L$ is the null exclusion. A larger exclusion indicates that a given fit can better distinguish a non-zero ν_e appearance signal from $\theta_{13} = 0$.

9.2 Incorporation of Systematic Error

Two methods were explored for introducing systematic errors into the likelihood fit. The first method uses a single nuisance parameter for each separate systematic error:

$$-2\ln L = \sum_{i=0}^{nbins} [\mu'_i - n_i + n_i \ln(n_i) - n_i \ln(\mu'_i)] + \sum_k \frac{f_k^2}{2} \tag{9.5}$$

$$\mu'_i = b_i + s_i + \sum_k f_k(b_i\sigma_{ik}^b + s_i\sigma_{ik}^s) \tag{9.6}$$

where b_i and s_i are the respective background and signal predictions in fitting bin i, σ_{ik}^b and σ_{ik}^s are the fractional errors on background and signal due to systematic error k, and f_k is a nuisance parameter scaling systematic error k. The values of the f_k parameters are determined in the MINUIT [136] minimization of Equation 9.5. This method will be referred to throughout this chapter as the *nuisance parameter* method. There are both advantages and disadvantages to this method. The first disadvantage is the processing time. As described in the previous chapter, over forty separate systematic errors are considered in the analysis. Each time the likelihood is calculated, Equation 9.5 must be minimized for each of the systematic error nuisance parameters f_k. Second, this method also assumes 100% bin-to-bin correlations for each systematic error. This is not an unreasonable assumption, given that most of the systematic errors are approximated by shifting the whole Monte Carlo distribution by one standard deviation. However, there are a few systematic errors for which the bin-to-bin correlations are not 100%. These include the HOOHE, MRE, and CalDet systematic uncertainties. Another method must be devised for incorporating these systematic errors.

A second method of introducing systematic error addresses both of these concerns. Although the fit statistics are Poisson, one can assume that the systematic errors under consideration behave according to Gaussian statistics. One can therefore calculate a systematic error covariance matrix V_{ij}:

$$V_{ij} = \sum_k \Delta_i^k \Delta_j^k \tag{9.7}$$

where Δ_i^k is a change in the prediction in bin i due to a $+1\sigma$ shift in systematic error k. This matrix includes contributions from each of the fractional systematic errors in Chapter 8, as well as contributions from the HOOHE, MRE, and CalDet covariance matrices. This covariance matrix is then employed in a modified version of the likelihood fit:

$$-2\ln L = \sum_{i=0}^{nbins} [\mu_i' - n_i + n_i \ln(n_i) - n_i \ln(\mu_i')] + \frac{1}{2}\mathbf{f}^T V^{-1}\mathbf{f} \tag{9.8}$$

$$\mu_i' = \mu_i + f_i \tag{9.9}$$

Equation 9.8 is minimized with respect to the nuisance parameters f_i (in a vector $\mathbf{f}$) which now correspond to fitting bin i. This method will be referred to throughout this chapter as the *bin-by-bin* method. The bin-by-bin method carries the advantage of only having as many nuisance parameters as there are bins. If few bins are used, the fitting time is reduced, and additional systematic errors can be added with little impact on speed. Additionally, the HOOHE, MRE, and CalDet errors can be incorporated directly into the covariance matrix. There are several possible downsides to using this method. First, when too few systematic errors are included in the sum in Equation 9.7, the covariance matrix V can be nearly singular and inversion is subject to numerical issues. Second, when a large number of bins are used (e.g., if energy binning is incorporated into the fit), the method loses its speed advantage. Matrix inversion difficulties also occur more frequently as the size of the matrix is increased.

When a large number of systematic errors are used with a smaller number of fitting bins, the methods in Equations 9.8 and 9.5 produce essentially the same result. The needs of an individual binning scheme will therefore drive the choice of method.

9.3 Choosing an appropriate binning

In the binning optimization studies which follow, a PORP decomposition is used to make the background and signal predictions. These studies also ignore systematic errors on the ν_μ CC fiducial flux, as the inclusion of these errors is time-consuming. CalDet, HOOHE, and MRE systematic errors are also disregarded, as these covariance matrices must be re-calculated for any new binning. As will be shown in Section 9.4, not including these errors has little impact on the final sensitivity. Due to the instability of the bin-by-bin covariance matrix method (Equation 9.8), the nuisance parameter method (Equation 9.5) is used to include systematic error.

9.3.1 An optimized single-bin LEM cut

While the final result in this thesis will be a multiple-bin likelihood fit, it is desirable to use a single-bin fit as a cross-check to the main analysis. This is the counting experiment style of fit used in the two previous MINOS published analyses. For this counting experiment, cuts were studied in increments of 0.05 between 0.55 and 0.90. Table 9.1 shows the resulting projected sensitivity and null exclusion for the various LEM cuts. A cut at LEM>0.7 is found to give the best projected sensitivity (the lowest upper 90% Confidence Level). This cut is a different from the earlier optimum LEM>0.76 cut, due to the inclusion of the true LEM systematic error.

Binning	Projected Sensitivity	Null Exclusion
LEM>0.55	0.0743	4.63
LEM>0.60	0.0708	5.07
LEM>0.65	0.0684	5.44
LEM>0.70	0.0659	5.81
LEM>0.75	0.0667	5.89
LEM>0.80	0.0728	5.22
LEM>0.85	0.0983	3.38
LEM>0.90	0.1875	2.23

Table 9.1: Sensitivity for various single-bin LEM cuts. The first column lists the LEM PID cut. The second column shows the 90% CL projected sensitivity for $\sin^2 2\theta_{13}$. The third column shows the expected exclusion of the null hypothesis ($\theta_{13} = 0$) in terms of $-2\Delta \ln L$. Systematic errors are incorporated using the nuisance parameter method (Equation 9.5), and all predictions are calculated using PORP, for an exposure of 8.2×10^{20} POT.

9.3.2 Fitting in bins of Particle ID

Prior to the work in this thesis, studies had suggested [138] that the most gain in projected sensitivity came from fitting in multiple bins of PID, rather than bins of energy.

Therefore, the PID binning will be considered first. There are a number of different ways to approach this PID binning. First, one can perform a true shape fit, incorporating the full PID range from 0 to 1.0. This incorporates all of the available information in the PID distribution, allowing for the differentiation of a shape-dependent signal and the normalization. A potential downside of this approach is the relative lack of systematic error information for the lower PID range. The systematic errors chosen for this analysis concentrated on uncertainties in the ν_e CC-like region, rather than the background-like sample. The ND Data/MC discrepancy on the low end of the PID distribution is less well understood. A second option is therefore to confine the PID binning to the signal-like LEM>0.50 region.

Too fine a PID binning may simply increase processing time with no accompanying gain in sensitivity. With this in mind, several PID binnings are considered:

- LEM>0.7
- 2 bin full-shape fit with bin edges: {0, 0.7, 1.0}
- 3 bin full shape fit: {0, 0.35, 0.65, 1.0}
- 5 bin full shape fit: {0.0, 0.2, 0.4, 0.6, 0.8, 1.0}
- 8 bin full shape fit: {0.0, 0.2, 0.3, 0.4, 0.5, 0.6, 0.7, 0.8, 1.0}
- 10 bin full shape fit: {0.0, 0.1, 0.2, 0.3, 0.4, 0.5, 0.6, 0.7, 0.8, 0.9, 1.0}
- High PID 3 bin fit: {0.6, 0.7, 0.8, 1.0}

The eight bin full shape fit uses a larger first and last bin to account for lack of statistics in the LEM<0.1 and LEM>0.9 regions. The three bins confined to the high PID region, {0.6, 0.7, 0.8, 1.0}, were chosen from an exhaustive scan of the LEM>0.50 region, which considered combinations of two, three, and four bins [139] [140].

Table 9.2 lists the resulting projected sensitivity and null exclusion for each of the binnings. The high PID three bin fit has a better sensitivity than any of the binnings, with the exception of the eight and ten bin cases. The eight bin fit gives an additional 2% improvement in sensitivity over this binning. However, there is essentially no difference between the full 10 bin fit and the 8 bin fit.

The full PID eight bin fit is therefore chosen as the PID binning for this analysis. The high PID three bin fit, which was used in the published version of the analysis (see Appendix A), will be used as a cross-check, along with the LEM>0.7 single-bin cut.

Binning	Projected Sensitivity	Null Exclusion
0.7, 1.0	0.0659	5.81
0, 0.7, 1.0	0.0634	6.22
Full PID, 3 bins	0.0660	5.75
High PID 3 bins: {0.6, 0.7, 0.8, 1.0}	0.0586	7.19
Full PID, 5 bins	0.0590	7.08
Full PID, 8 bins	0.0574	7.40
Full PID, 10 bins	0.0572	7.41

Table 9.2: Sensitivity for various LEM-only binnings. The first column lists the binning. The second column shows the 90% CL projected sensitivity for $\sin^2 2\theta_{13}$. The third column shows the expected $-2\Delta \ln L$ exclusion of the null θ_{13}=0 hypothesis for $\sin^2(2\theta_{13}) = 0.1$. Systematic errors are incorporated via the nuisance parameter method (Equation 9.5), and all predictions are calculated using PORP, for an exposure of 8.2×10^{20} POT.

9.3.3 Fitting in Energy

Fitting in multiple bins of LEM provides a significant improvement in sensitivity. However, additional information may be gained from fitting in bins of energy as well. Table 9.3 shows the results of combining the three LEM binnings (single cut, three bin, and eight bin) with various reconstructed energy binnings. The decomposition binning, with edges {1, 2, 3, 4, 5, 8} (in GeV), gives both the best projected sensitivity reach and the best null exclusion. The total improvement from introducing this binning is on the order of $1 - 2\%$ - small, compared to the 14% improvement gained by optimizing the LEM binning. The reason for this small improvement can be seen in Figure 9.1, which shows Signal/Background as a function of LEM binning and energy binning. While S/B varies strongly with LEM binning, the relationship is much flatter with energy.

The above results suggest that fitting in both energy and PID provides the greatest sensitivity. A combined PID and energy fit uses all available shape information and is capable of accounting for non-standard PID and energy distributions in the data. The decision was therefore made to fit in 40 separate bins: eight bins of LEM PID with edges {0.0, 0.2, 0.3, 0.4, 0.5, 0.6, 0.7, 0.8, 1.0}, and five bins of reconstructed energy with edges (in GeV) {1, 2, 3, 4, 5, 8}. Once this binning was chosen, a full HOOHE decomposition and formal extrapolation was performed for the 8 LEM × 5 E binning. The remainder of the studies used in this chapter use this full extrapolation.

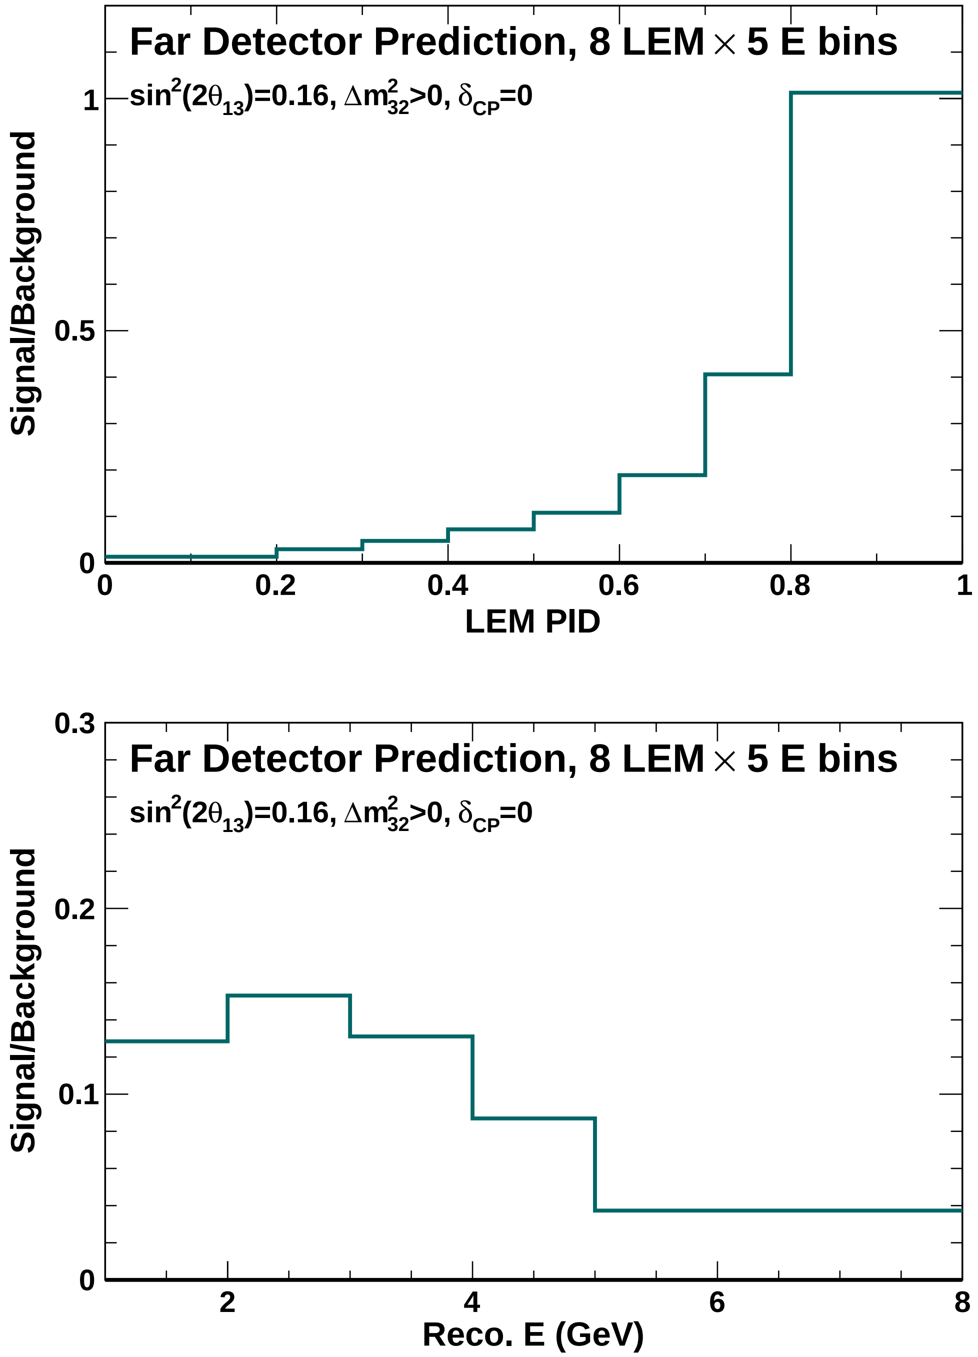

Figure 9.1: Signal to Background ratio as a function of LEM binning (top) and Reconstructed Energy binning (bottom). The predictions for signal and background are made as projections of an 8 LEM × 5 E bin extrapolation, with a CHOOZ-size signal.

LEM Bins	E Bins	Projected Sensitivity	Null Exclusion
LEM>0.70	1 8	0.0659	5.81
LEM>0.70	1 5 8	0.0652	5.92
LEM>0.70	1 3 5 8	0.0653	5.92
LEM>0.70	1 2 3 4 5 8	0.0651	5.96
LEM 3 bin	1 8	0.0585	7.19
LEM 3 bin	1 5 8	0.0580	7.31
LEM 3 bin	1 3 5 8	0.0583	7.26
LEM 3 bin	1 2 3 4 5 8	0.0580	7.32
LEM 8 bin	1 8	0.0574	7.40
LEM 8 bin	1 5 8	0.0569	7.54
LEM 8 bin	1 3 5 8	0.0571	7.49
LEM 8 bin	1 2 3 4 5 8	0.0568	7.56

Table 9.3: Sensitivity for various combinations of LEM and Energy binning. The first column lists the PID binning, and the second column the bin edges for energy binning (in GeV). The third column shows the 90% CL projected sensitivity for $\sin^2 2\theta_{13}$. The fourth column shows the $-2\Delta \ln L$ exclusion of the null hypothesis (θ_{13}=0). Systematic errors are incorporated with the nuisance parameter method (Equation 9.5), and all predictions are calculated using PORP, for an exposure of 8.2×10^{20} POT.

9.4 Systematic Error

In the previous chapter, 49 separate sources of systematic error were identified, and the fractional error on the LEM>0.7 signal-enhanced region determined. For a multiple bin fit, the shape information of the individual systematic errors also becomes important. To assess the impact of systematic error k on the 8×5 bin fit, the projected sensitivity and null exclusion are calculated when k is the only systematic error in the fit. These values are then compared to the sensitivity of a statistics-only fit, and the fractional change in the projected sensitivity and in the null exclusion due to systematic error k are calculated. The average of these two values is taken to characterize the loss in sensitivity incurred by that particular systematic error.

This sensitivity loss was determined for 46 of the total 49 systematic errors, using the nuisance parameter method (the bin-by-bin method is unreliable when only one or two systematic errors are present). The inversion of the HOOHE, CalDet and MRE covariance matrices failed for the 8×5 binning, so these matrices were again ignored. Indeed, computational inaccuracies in the inversion of the near-singular 40×40 covariance matrix caused the bin-by-bin method to fail in general for this binning.

The systematic errors were sorted into nine different groups of five systematics each, in descending order of their impact on the sensitivity. These groups are listed in Table 9.4. The five systematic errors having the largest impact on the 8×5 bin analysis were found to be: Normalization, Near and Far Relative Energy Scales, Absolute Energy Scale and ν_τ cross-section. These sets of error were added cumulatively to the likelihood fit, and the total loss in sensitivity assessed after each step. When all of the errors are included in the fit, there is a loss in projected sensitivity of approximately 10%, from $\sin^2(2\theta_{13}) = 0.0528$ to 0.0584. Figure 9.2 shows the fraction of this total loss in sensitivity achieved after adding each group of five systematic errors. After including the 15 largest systematic errors, 98% of the total systematic error effect has been accounted for. After twenty five errors have been added, this fraction is greater than 99%.

Of the 46 total non-covariance matrix systematic errors, therefore, less than half make up 99% of the total systematic error effect. However, all 46 errors were retained for the fit, as the processing time remained manageable. This study, however, provided validation of the decision to neglect the MRE, CalDet, and HOOHE errors. Because of technical difficulties in inverting the covariance matrix for the 8×5 binning, the 3×5 binning was used instead to evaluate the impact of these errors. This is a valid comparison: an identical study by the author [141] assessed sensitivity losses in the 3×5 binning and found a ranking nearly identical to that in Table 9.4. The impact of the HOOHE, CalDet, and MRE errors on the sensitivity of the 3×5 fit is shown in Table 9.5. In all cases, the errors have a negligible effect on sensitivity, smaller than most of systematic errors in 9.4. They can therefore be safely disregarded in the LEM 8×5 binning.

An alternate method of assessing the impact of systematic error on the analysis is to look at the change in the fit result when a modified version of the prediction is used as the observed distribution $\{n_i\}$. The nominal observed 8×5 bin distribution is chosen to be the background prediction with signal at $\sin^2(2\theta_{13})$=0.10. The best fit, upper 90% CL, and null exclusion are calculated for this distribution. Next, this observed distribution is reweighted with a $+1\sigma$ systematic error shift on the signal and background, and the new best fit, upper 90% CL and exclusion at zero are calculated. These values are then compared to the nominal fit. This study was performed for the five largest systematic errors (Group 1 in Table 9.4), with the results listed in Table 9.6. The best fits of the shifted distributions all remain within 10% of the nominal value. The upper 90% CL and null exclusion show a similar degree of variation.

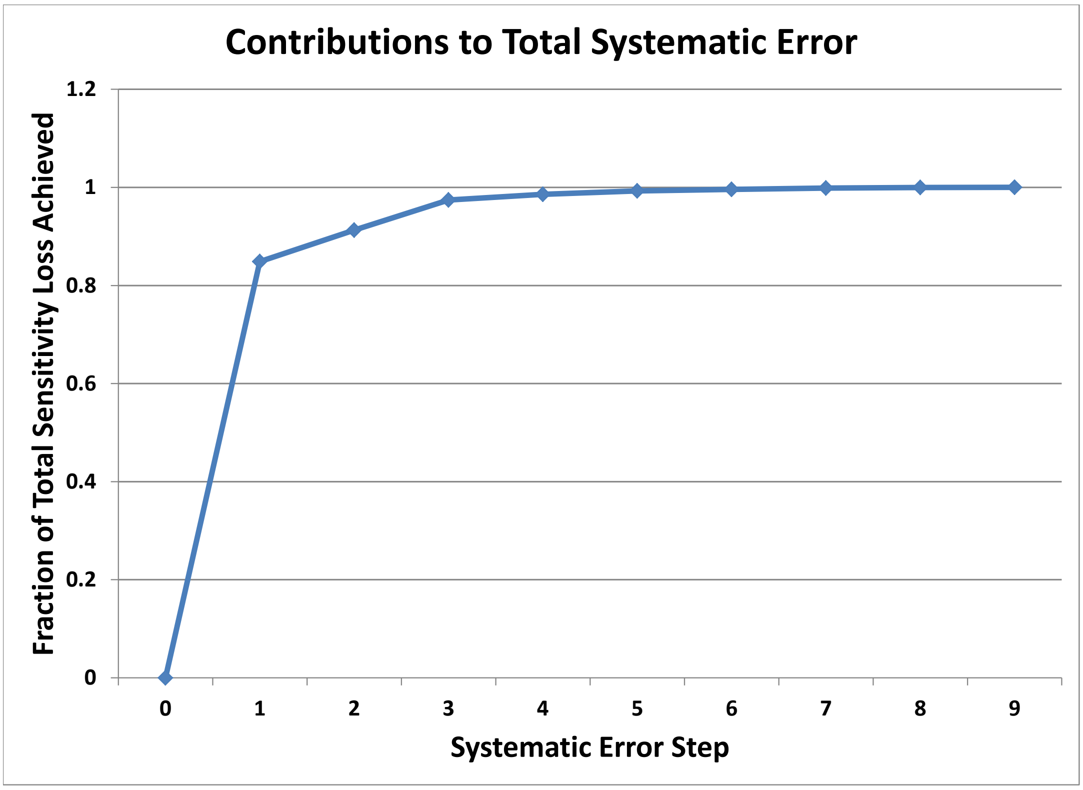

Figure 9.2: Total loss in sensitivity achieved after adding various systematic errors. The x axis refers to each of the groups in Table 9.4. The groups are added cumulatively to the fit, from 0 (statistics only) to 9 (all systematics included). The y axis shows what fraction of the total loss in sensitivity has been achieved after each step.

9.5 Projected Sensitivity

The final result of this analysis will be obtained with a fit to the data, in 8 PID bins covering the full PID distribution and 5 reconstructed energy bins. The signal and background predictions in each of the 40 fit bins are shown in Figure 9.3, and the resulting signal/background ratio in Figure 9.4. The projected sensitivity of the 8 LEM × 5 E bin fit is now determined using the full HOOHE extrapolation and the full set of systematic errors. The high-PID 3 LEM × 5 E bin fit and a LEM>0.7 counting experiment fit (with only one energy bin) will be used as cross-checks. Table 9.7 lists the projected sensitivity for the Normal Hierarchy for these three binnings. For the sake of comparison, the sensitivity of ANN11>0.7 is also shown.

Figure 9.5 shows the full 2D projected sensitivity contours for each of these fits. This plot shows how the sensitivity of this analysis has been improved relative to the ANN11 single bin fit. Employing a multiple bin shape fit provides an improvement in sensitivity comparable to that of using the LEM PID rather than ANN11. The gain from going from ANN11 to LEM is approximately 14%, while the gain from LEM>0.7 to 8 LEM × 5 E bins is also approximately 14%. The analysis improvements in this thesis have increased the MINOS sensitivity to $\sin^2(2\theta_{13})$ by a total of 27% - a gain in sensitivity equivalent to a 54% increase in data. As shown in Figure 9.6, the new 8×5 fit reaches a sensitivity well below the CHOOZ limit, for both the Normal and Inverted Hierarchy.

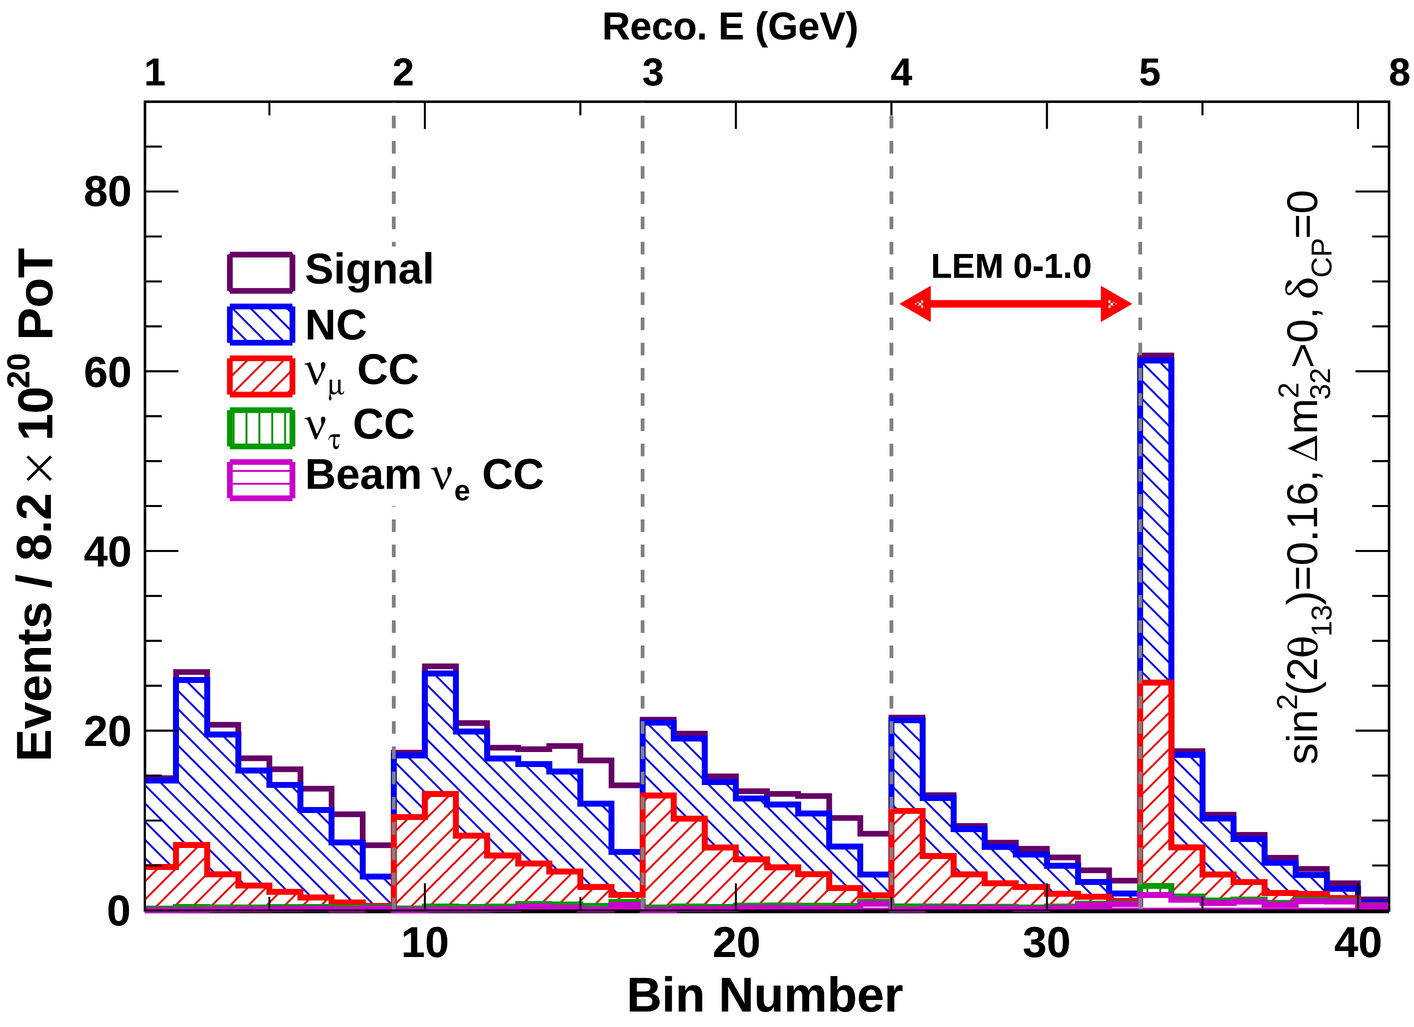

Figure 9.3: The FD prediction in each extrapolation bin for the 8 LEM × 5 E bin fit. In this plot, Bin Number = (LEM PID bin) + 5 × (Reconstructed Energy bin). The space between each dashed line indicates a single energy bin. A signal corresponding to the CHOOZ limit is presented as well. The prediction is for an exposure of 8.2×10^{20} protons on target.

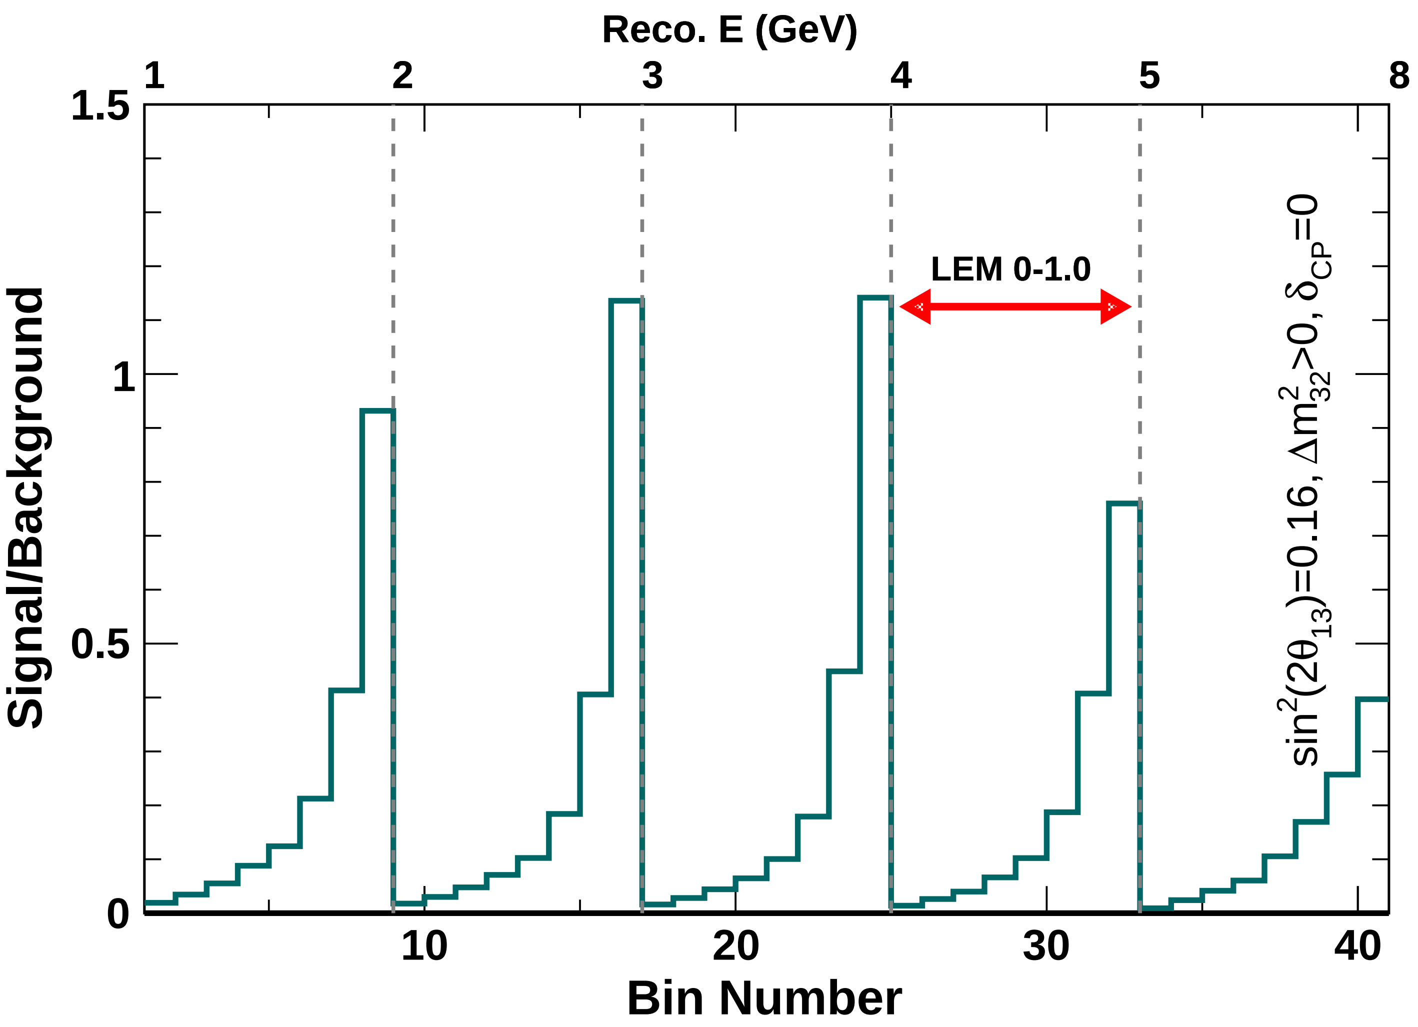

Figure 9.4: Predicted Signal/Background ratio in each extrapolation bin for the 8 LEM × 5 E bin fit. In this plot, Bin Number = (LEM PID bin) + 5 × (Reconstructed Energy bin). The space between each dashed line indicates a single energy bin. The signal prediction corresponds to the CHOOZ limit.

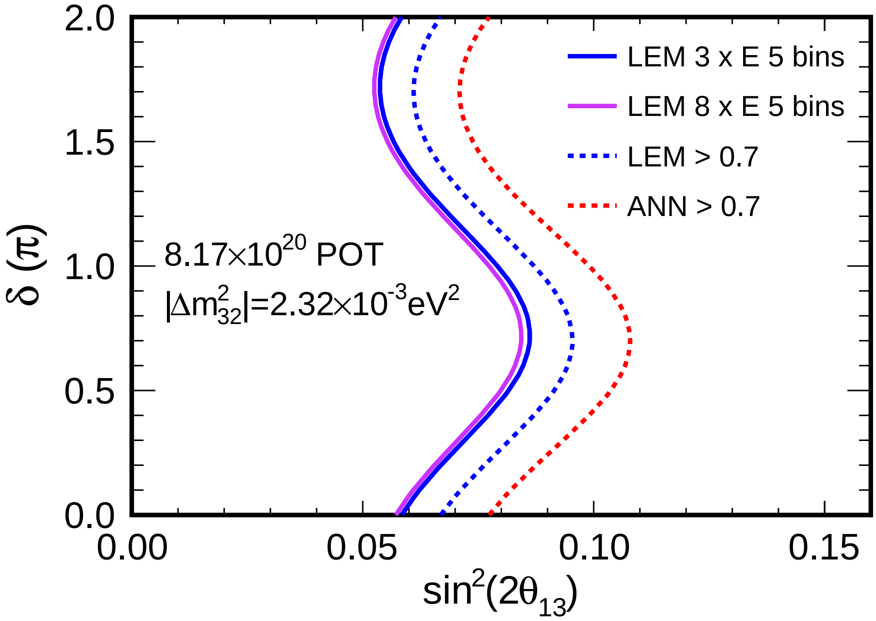

Figure 9.5: Comparative projected sensitivity for an 8 × 5 fit, a 3 × 5 fit, a LEM>0.70 fit, and an ANN11>0.70 fit. This plot assumes the Normal Hierarchy ($\Delta m_{32} > 0$) and an exposure of 8.2 $\times 10^{20}$ POT.

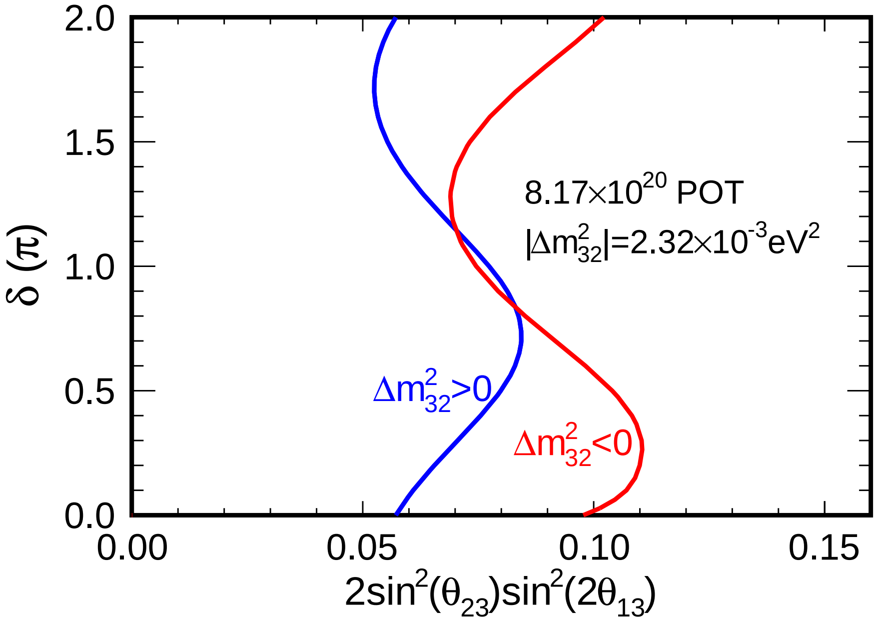

Figure 9.6: Projected sensitivity for an 8 × 5 fit for both the Normal and Inverted Hierarchy. This plot assumes an exposure of 8.2 $\times 10^{20}$ POT.

% of Total Sensitivity Loss	Systematic Errors in Group
Group 1: 84.88%	Background Normalization, Relative E. Scale (ND), ν_τ Cross-section, Absolute Energy Scale, Relative E. Scale (FD)
Group 2: 6.42%	INUKE 2, AGKY (Multiplicity Correlations), INUKE 7, Gain (FD), FLUGG
Group 3: 6.12%	AGKY (Isotropic 2 Body Decay), CC Energy Shift INUKE 5, Optical Crosstalk, INUKE 9
Group 4: 1.15%	ν_e Normalization, AGKY (π^0 Selection -1σ) AGKY (π^0 Selection $+1\sigma$), Attenuation, Gain (ND)
Group 5: 0.72%	INUKE 8, INUKE 1, Target Degradation, Strip to Strip, AGKY (Baryon x_f selection)
Group 6: 0.30%	INUKE 3, INUKE 0, ν_τ Normalization, EM vs. Had. Energy Scale, Hadron Multiplicity
Group 7: 0.25%	M_A^{RES}, INUKE 6, CCSigNC, KNO, ν_μ CC SKZP
Group 8: 0.13%	AGKY (p_T Squeezing), INUKE 4, SKZP, AGKY (Implementation Ambiguities), ν_μ CC KNO
Group 9: 0.02%	ν_μ CC M_A^{QE}, ν_μ CC M_A^{RES}, Linearity, ν_μ CC Target Degradation, M_A^{QE}, ν_μ CC FLUGG

Table 9.4: Systematic errors, sorted into various groups based on their impact on the sensitivity of the LEM 8×5 bin fit. Group 1 has the largest impact, while Group 9 has the smallest input. The number in the left hand column indicates what percentage of the total loss in sensitivity is caused by the systematics in that group.

	Systematic Errors			
Performance Measure	Statistics only	HOOHE	CalDet + HOOHE	MRE + HOOHE
Projected Sensitivity:	0.05440	0.05442	0.05442	0.05442
Null Exclusion	8.348	8.343	8.343	8.343

Table 9.5: Impact of the three covariance matrix errors (HOOHE, CalDet, and MRE) on the 3 LEM × 5 E bin fit.

Systematic Error	Best Fit	Upper 90% CL	Null Exclusion
Nominal	0.100	0.180	7.47
Normalization	0.104	0.188	8.36
Absolute Energy Scale	0.107	0.187	8.60
Rel. Energy Scale (FD)	0.103	0.183	7.84
Rel. Energy Scale (ND)	0.093	0.172	6.66
ν_τ Cross-section	0.106	0.187	8.21

Table 9.6: Best fit, upper 90% CL, and exclusion at zero for 8 LEM × 5 E bin fit, both for a nominal observed distribution with $\sin^2(2\theta_{13}) = 0.10$ and for distrubutions shifted by $+1\sigma$. The upper 90% CL and best fit are values of $\sin^2 2\theta_{13}$. The null exclusion is a value of $-2\Delta \ln L$.

Binning	Projected Sensitivity
ANN11>0.70	0.0792
LEM>0.70	0.0683
3 LEM × 5 E bins	0.0594
8 LEM × 5 E bins	0.0584

Table 9.7: Projected sensitivities, i.e. the upper 90% CL for a background-only observation, for four separate binnings. The fits are performed using the full set of systematic errors (minos MRE, CalDet, and HOOHE).

Chapter 10

The Far Detector Data

Nearly all of the studies outlined in the previous chapters have used ND Monte Carlo, FD Monte Carlo, and ND Data. Prior to this point, the Far Detector Data in the ν_e analysis region has been left blinded, to avoid possible analysis bias. With the analysis procedure finalized, the Far Detector Data can finally be examined. This chapter describes the unblinding procedure and the interpretation of the FD Data using the analysis framework developed in this thesis. The first two sections describe two final sideband cross-checks which were carried out before the data in the ν_e analysis region were studied.

10.1 Sideband 1: The Anti-PID sideband

The Anti-PID sideband tests the ability of the analysis (the PID, decomposition, and extrapolation) to predict the FD Data in the lower end of the PID region. Even for $\sin^2(2\theta_{13}) = 0.16$, the Signal/Background ratio in all bins in this region is less than approximately 0.1. A cut of LEM<0.50 has a FOM of $N_{sig}/\sqrt{N_{bkg}} = 11.4/\sqrt{373.3} = 0.57$. In comparison, the signal-like LEM>0.70 region has a FOM of $30.3/\sqrt{49.1} = 4.3$. This sample is therefore safe to use as a sideband to the main analysis.

The data in the LEM<0.50 sideband are compared with the prediction for both a no-signal case, $\sin^2(2\theta_{13}) = 0$, and with a signal at $\sin^2(2\theta_{13}) = 0.16$. Comparisons of the Anti-PID prediction and data are shown in Figure 10.1. 377 events in total are observed, versus a prediction of both 373.3 events for $\sin^2(2\theta_{13}) = 0$, and 384.4 events for $\sin^2(2\theta_{13}) = 0.16$.

The agreement between the Anti-PID data and prediction distributions is then assessed. The log-likelihood of the data distribution compared to the prediction is first calculated. To assess the significance of this log-likelihood, 10,000 toy Monte Carlo experiments are then generated by applying Poisson fluctuations to the predicted distribution. The log-likelihood of the data is then compared to that obtained from the toy experiments. For experiments generated from the $\sin^2(2\theta_{13}) = 0$ prediction, 86% of experiments have a better likelihood than the data. For experiments generated from the $\sin^2(2\theta_{13}) = 0.16$ prediction, 87% of experiments have a better likelihood than the data. This is an acceptable degree of variation between the Anti-PID FD data and prediction, which could have easily been produced by statistical fluctuation.

10.2 Sideband 2: The MRCC Sideband

While the Anti-PID sideband adequately tests the decomposition and extrapolation, it does not test the analysis for hadronic showers in the high PID signal-like region, where there is known to be a greater than 20% discrepancy between ND Data and MC. The Muon Removed Charge Current (MRCC) sample tests the analysis for these types of events. The MRCC sample is very similar to the MRE sample. A muon removal algorithm [142] is used to remove the tracks from ν_μ CC events. This leaves behind a sample of hadronic showers. Because the sample is designed to contain events which were originally ν_μ CC events, there is no chance of accidentally observing a ν_e appearance signal in the MRCC FD Data. The sideband can therefore be used to check that the Data/MC discrepancy observed in the Near Detector is also present in the Far Detector. If this is not the case, it indicates that there is some unidentified difference between the two detectors which must be accounted for in the analysis.

For this sideband, the ND Data MRCC distribution is extrapolated to form a FD MRCC prediction, which is then compared to the MRCC FD Data. This side band is then evaluated in two ways. First, the statistical goodness-of-fit between the MRCC Data and Prediction is calculated. Second, the multiple bin likelihood fit is applied to the MRCC data, to determine the extent to which the MRCC FD Data mimics a signal. Prior to studying this sideband, the acceptable level of agreement between the prediction and data was defined to be 2σ.

10.2.1 MRCC Prediction vs. Data

A series of cuts is applied to the muon-removed ND and FD Data and MC to produce the MRCC sample. These cuts are listed in Table 10.1. An initial round of MRCC Preselection cuts are first applied. These cuts ensure data quality. They also include a series of cuts taken from the ν_μ CC disappearance analysis which ensure that the original event was a golden ν_μ CC event with a definite muon track. An MRCC fiducial

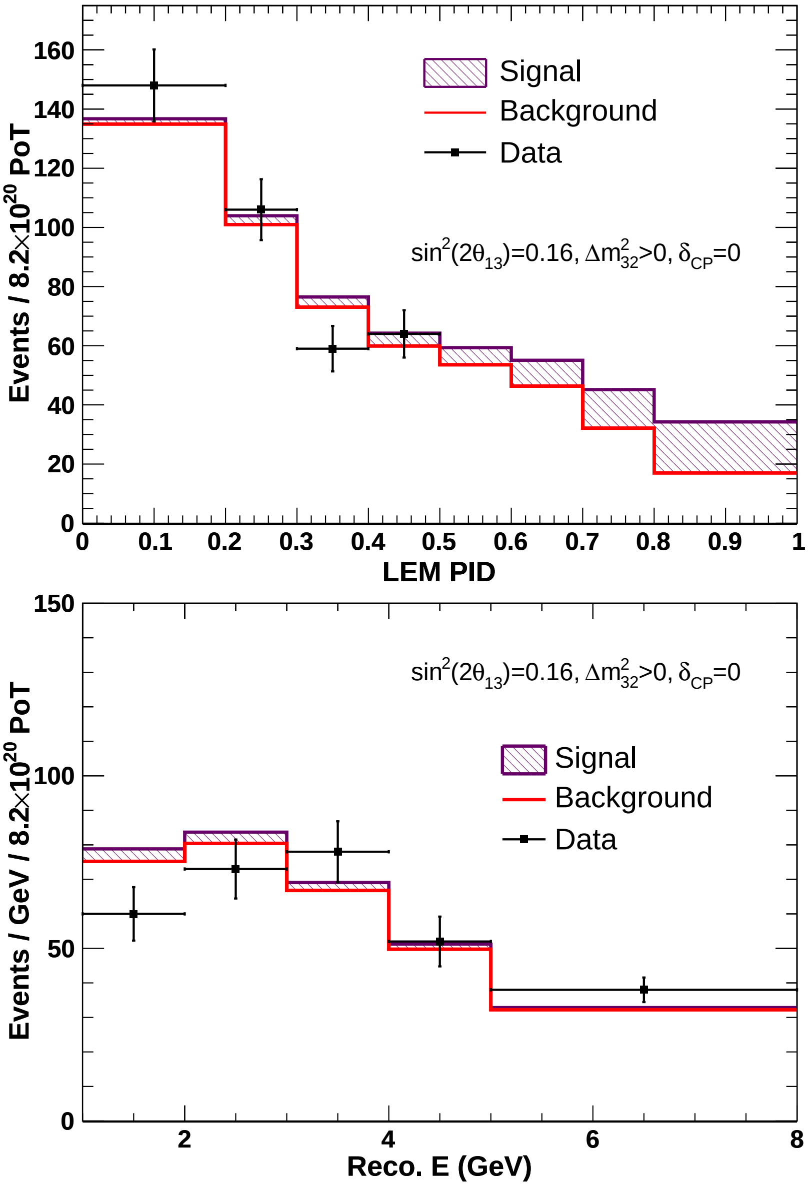

Figure 10.1: The Anti-PID sideband. The top plot shows the LEM PID distribution, with only the Far Detector data bins in the Anti-PID region (LEM<0.50) opened. The bottom plot shows the reconstructed energy prediction and data for LEM<0.5. The expected signal for $\sin^2(2\theta_{13}) = 0.16$ is also shown.

cut (slightly more relaxed than the ν_e analysis version) is also applied to the event vertex. The reconstructed energy of the events after muon removal is constrained to be within the standard ν_e analysis Preselection range of 1.0 to 8.0 GeV. Following this MRCC Preselection, the standard ν_e Fiducial and Preselection cuts are applied. Finally, the PID selection is applied, and the events are sorted into the 8 LEM × 5 E binning.

MRCC Preselection	
Data Quality	L010185N running Event successfully muon removed Beam quality cuts (data only) Detector quality cuts (data only) Largest event in snarl (FD only)
Event Energy	1.0 GeV $<$ Reco. Energy $<$ 8.0 GeV Sorted into binning $\{1.0, 2.0, 3.0, 4.0, 5.0, 8.0\}$ (GeV)
MRCC Fiducial	Original event within MRCC fiducial volume
MRCC Preselection	Has a track Track pass fit Original roCC PID > 0.3
Standard Preselection	
ν_e Fiducial	MRCC event within ν_e fiducial volume
Cosmic Cuts	Event slope < 10 Angle of any track > 0.6 Distance between start and end of any track < 2 m
ν_e Preselection	Track-like planes < 16 Track Planes < 25 Contiguous Planes > 4 Number of Showers > 0
PID Selection	
ν_e PID	LEM PID, with binning: $\{0.0, 0.2, 0.3, 0.4, 0.5, 0.6, 0.7, 0.8, 1.0\}$

Table 10.1: The series of cuts used in the MRCC sideband. The MRCC Preselection cuts are used to make the MRCC sample itself. The Standard Preselection and PID cuts are identical to what is used in the standard analysis.

Because the MRCC ND Data consist almost entirely of ν_μ CC events, there is no need for a decomposition. Following selection cuts, the ND Data are extrapolated to the Far Detector as follows:

$$\mathrm{P}_{ij} = \mathrm{N}^{\mathrm{LEM},ij}_{\mathrm{Data}} \times \frac{\mathrm{F}^{\mathrm{LEM},ij}_{\mathrm{MC}}}{\mathrm{N}^{\mathrm{LEM},ij}_{\mathrm{MC}}} \times \frac{\mathrm{MRCC}^{\mathrm{Data}}}{\mathrm{MRCC}^{\mathrm{Pred}}} \tag{10.1}$$

where P_{ij} is the MRCC FD prediction for LEM bin i and Energy bin j, $\mathrm{N}^{\mathrm{LEM,ij}}_{\mathrm{Data}}$ is the ND Data MRCC Selection, and $\mathrm{N}^{\mathrm{LEM,ij}}_{\mathrm{MC}}$ and $\mathrm{F}^{\mathrm{LEM,ij}}_{\mathrm{MC}}$ are the ND and FD MC selections, respectively. $\mathrm{MRCC}^{\mathrm{Data}}$ is the MRCC Preselection for the FD Data, and $\mathrm{MRCC}^{\mathrm{Pred}}$ is the predicted MRCC Preselection, defined as follows:

$$\mathrm{MRCC}^{\mathrm{Pred}} = \sum_j \mathrm{MRCC}^{j}_{\mathrm{ND\ Data}} \times \frac{\mathrm{MRCC}^{j}_{\mathrm{FD\ MC}}}{\mathrm{MRCC}^{j}_{\mathrm{ND\ MC}}} \tag{10.2}$$

where $\mathrm{MRCC}^{j}_{\mathrm{ND\ Data}}$ is the MRCC Preselected ND Data selection in energy bin j, and $\mathrm{MRCC}^{j}_{\mathrm{ND(FD)\ MC}}$ the same but for the ND (FD) Monte Carlo.

The MRCC Preselection correction in Equation 10.2 ensures that the correct number of events are predicted at the level of the MRCC Preselection. This is necessary, due to the presence of a known Far Detector excess in the MRCC Preselection sample discovered [143] in Run 3. This thesis is only concerned with the behavior of ν_e CC-like events, rather than that of ν_μ CC-like events. Therefore, the above correction is made to mitigate the effects of this discrepancy.

The extrapolation in Equation 10.1 is performed to make an MRCC prediction in the standard 8×5 binning. This extrapolation is performed separately for the three data runs and summed. This prediction is then compared to the unblinded MRCC FD Data. The resulting LEM, Energy, and 40 bin distributions are shown in Figures 10.2, 10.3, and 10.4. The goodness of fit was assessed using toy MC experiments, in the same manner as the Anti-PID sideband. For the full 8 × 5 distribution, 80% of experiments had a better absolute likelihood, an acceptable level of agreement. As a cross-check, 88% of experiments had a better absolute likelihood for the 3×5 binning. If a cut of LEM>0.70 is applied, 50.7 events are predicted and 64 observed, a 1.87 σ discrepancy. There is therefore an acceptable level of agreement between the prediction and data for the MRCC distribution. Note also that the significance of the above numbers is somewhat exagerrated, as none of these tests considered the systematic error on the MRCC sample.

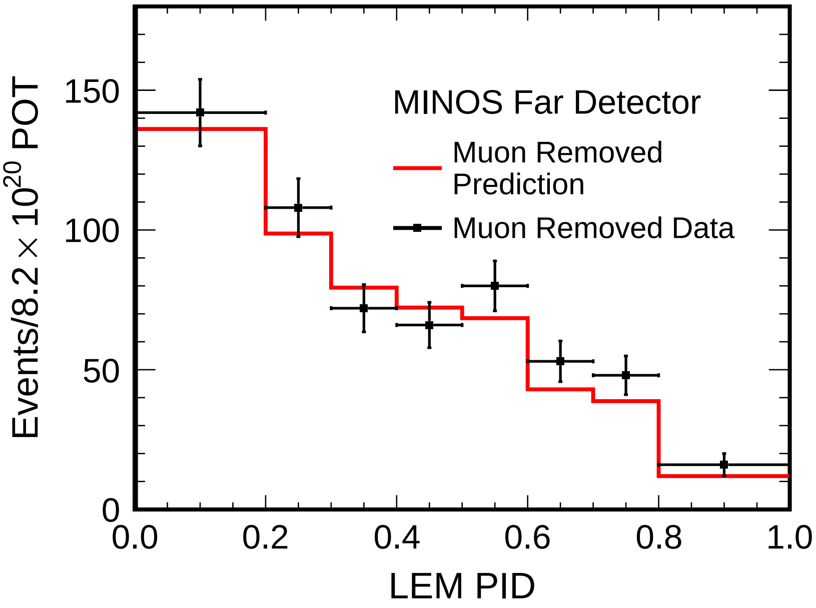

Figure 10.2: MRCC Far Detector Prediction (red) versus Data (black) for the Full PID distribution for the LEM 8×5 E bin extrapolation.

10.2.2 MRCC: Multiple Bin Fit Studies

A second way to evaluate the MRCC sideband is to test how well the FD Data events mimic a signal when a multiple bin fit is conducted. The MRCC sideband contains no signal by design. The MRCC FD Data is instead fitted to the MRCC prediction plus an injected signal-like term which varies with θ_{13}: $f(\theta_{13})$ = MRCC Prediction + Signal(θ_{13}). Signal(θ_{13}) is simply the standard ν_e signal prediction. When the MRCC FD Data is fitted to $f(\theta_{13})$, the best fit should be consistent with the null hypothesis of $\theta_{13} = 0$.

The significance of this fit is assessed with a toy MC study, similar to that used in the Feldman-Cousin method [144]. Systematic errors are not considered in this analysis in either fitting or generating toy experiments, as the relationship between the Standard and (ν_μ CC based) MRCC systematic errors is not well understood. The fit itself is a modified version of the likelihood fitting code described in the previous section. The steps for this analysis are as follows:

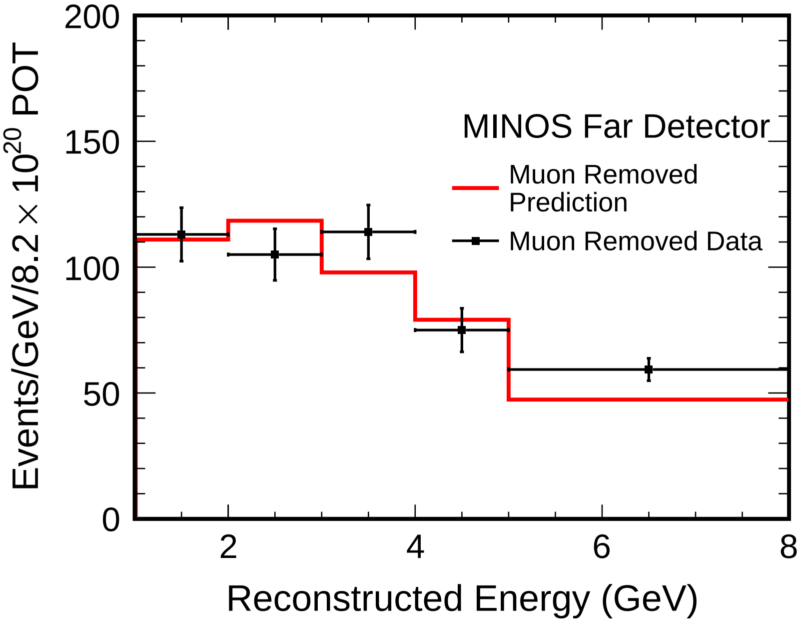

Figure 10.3: MRCC Far Detector Prediction (red) versus Data (black) for Reconstructed Energy.

1. The likelihood ratio of the data is calculated as a function of θ_{13} using:

$$-2\ln L^{Data} = \sum_{i=0}^{nbins} \mu_i - n_i + n_i \ln(n_i) - n_i \ln(\mu_i) \tag{10.3}$$

   The best fit θ_{13}^{best} is located from the minimum likelihood ratio $-2\ln L_{min}^{Data}$.

2. The null exclusion, $-2\Delta\ln L^{Data}$ at $\theta_{13} = 0$ (i.e., for the MRCC Prediction), is calculated:

$$-2\Delta\ln L^{Data} = -2\ln L_0^{Data} - (-2\ln L_{min}^{Data}) \tag{10.4}$$

   where $-2\ln L_0^{Data}$ is the likelihood at $\theta_{13} = 0$.

3. The MRCC prediction is fluctuated using Poisson statistics to produce 10,000 pseudo-experiments.

4. Steps 1 and 2 are repeated with the pseudo-experiments in place of the Data, so that $-2\Delta\ln L$ and best fit θ_{13}^{best} are found for each of the experiments.

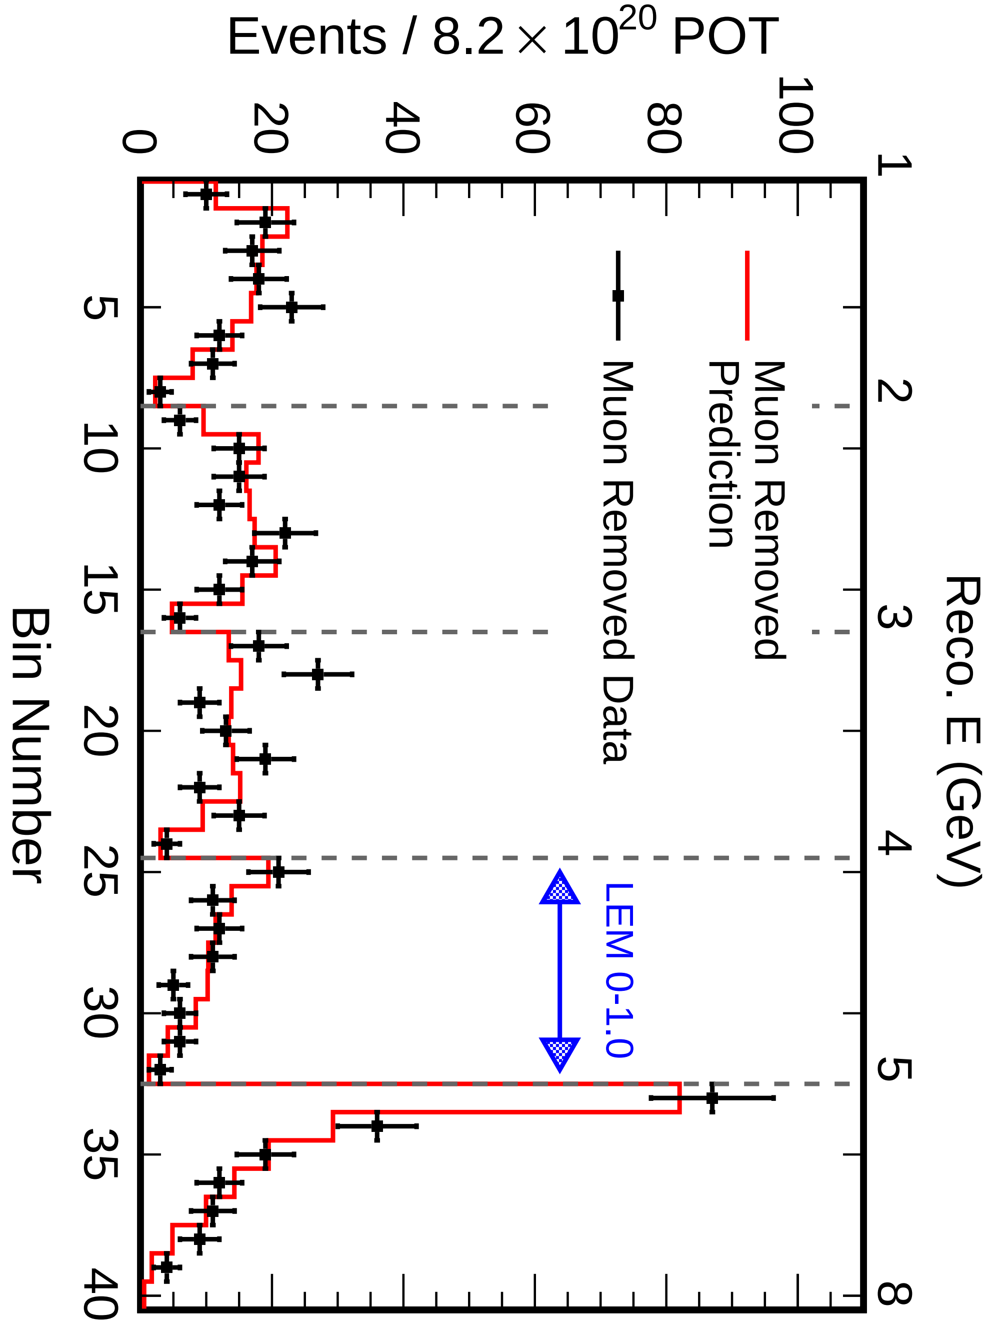

Figure 10.4: MRCC Far Detector Prediction (red) versus Data (black) for the LEM 8×5 E binning.

5. The fraction $f_{\mathcal{L}}$ of experiments for which $-2\Delta \ln L \leq -2\Delta \ln L^{Data}$ is calculated. The fraction of experiments $f_{\theta_{13}}$ for which $\theta_{13}^{best} \leq \theta_{13}^{best,Data}$ is also found.

The value $f_{\mathcal{L}}$ is the Confidence Level of a non-zero "signal." It can be interpreted as the probability that the fitted excess is real, rather than a fluctuation. The value $f_{\theta_{13}}$ is the probability that a random fluctuation in the data would be fit to a smaller signal. Notably, both of these quantities are single-sided distributions, so that the cutoff for 2σ is 97.7%. The values of $f_{\mathcal{L}}$ and $f_{\theta_{13}}$ for both an 8×5 and 3×5 bin fit to the MRCC are listed in Table 10.2. The two variables $f_{\mathcal{L}}$ and $f_{\theta_{13}}$ give similar results. For the full shape distribution, $f_{\mathcal{L}} = 92.2\%$ permits the inclusion of $\theta_{13} = 0$ at 1.51σ. This is an acceptable level of agreement with the prediction.

Quantity	LEM 3×5 bin	LEM Full PID
Best Fit $\sin^2 2\theta_{13}^{\text{best}}$	0.030	0.032
$-2\Delta \ln L$ at $\theta_{13} = 0$	1.85	2.11
$f_{\mathcal{L}}$	92.2%	93.5%
Significance	1.42 σ	1.51 σ
$f_{\theta_{13}}$	92.2%	93.7%
Significance	1.42 σ	1.53 σ

Table 10.2: Results of the toy MC study carried out for the MRCC sideband. $f_{\mathcal{L}}$ is the fraction of pseudo-experiments with a better fit to the prediction than the Data. $f_{\theta_{13}}$ is the fraction of pseudo-experiments which have a smaller signal-like excess. Not that this study did not include systematic error.

While the excess is below 2σ, it is also still somewhat high, with $f_{\mathcal{L}}$ and $f_{\theta_{13}}$ being greater than 90%. The significance of this excess is reduced somewhat by the fact that all of the above MRCC studies were statistics-only, and did not consider systematic error. The systematic uncertainty on the MRCC sample is not well understood, but will be different from that affecting the standard sample. Some of this systematic uncertainty will result, for example, from the muon removal algorithm and from the cuts used in the creation of the MRCC sample itself. An exhaustive series of tests were also conducted on the MRCC sample, to attempt to discover the source of the excess. These tests looked for pathologies in numerous variable distributions, for position biases, for problems in the muon removal process, etc. In the end, it was concluded that the excess was a statistical fluctuation.

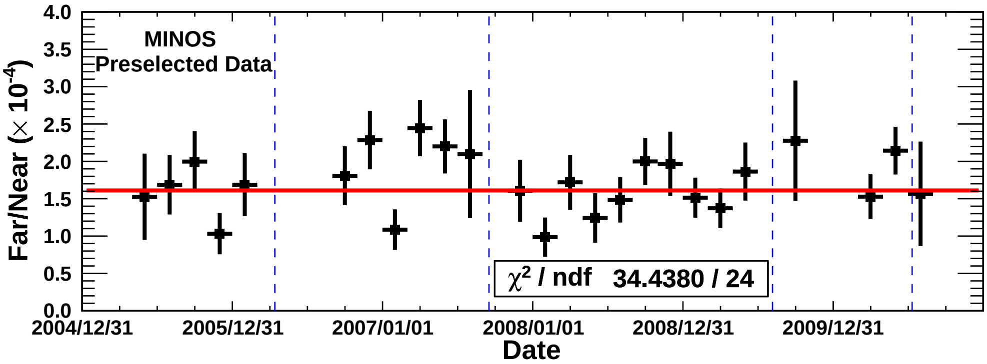

Figure 10.5: Far/Near ratio of Preselected MINOS data as a function of time. The dashed lines separate data into (in order) running periods 1, 2, 3, 4+5, and 6. Plot provided by Mhair Orchanian.

10.3 Opening the Data

Once the two sidebands have been checked, the analysis in this thesis is considered to be finalized, and the FD Data over the full PID range can be unblinded. Low-level distributions showing event timing, the physical location of the events, etc. are presented first, followed by the distributions directly linked to the analysis (energy and PID variables). The data are then interpreted in terms of the physical parameters.

10.3.1 Variable Distributions

Several checks are made of the FD Data to ensure basic data quality. Figure 10.5 shows the Far/Near ratio for Preselected data as a function of time, ensuring that the event rate has been steady. The data agrees with a flat mean to within statistical fluctuations, with a χ^2/ndf of 34.4/24.

Figure 10.6 shows the time of the Preselected Far Detector data events relative to the nearest beam spill. The Preselection requires this relative time to be -2×10^{-6} s $< \Delta t < 12 \times 10^{-6}$ s. All of the Preselected events are well contained in this window, and there are no peaks towards the end of the distribution which might indicate pathological behavior or contamination by non-beam events.

Figures 10.7 and 10.8 show the position of event vertices in the detector. Both in the transverse and longitudinal directions, there is no indication of any statistically significant excesses towards the edges of the detector.

Figure 10.9 compares the FD Data and prediction for the Preselected reconstructed energy distribution in 1 GeV bins. This distribution initially gave some cause for concern,

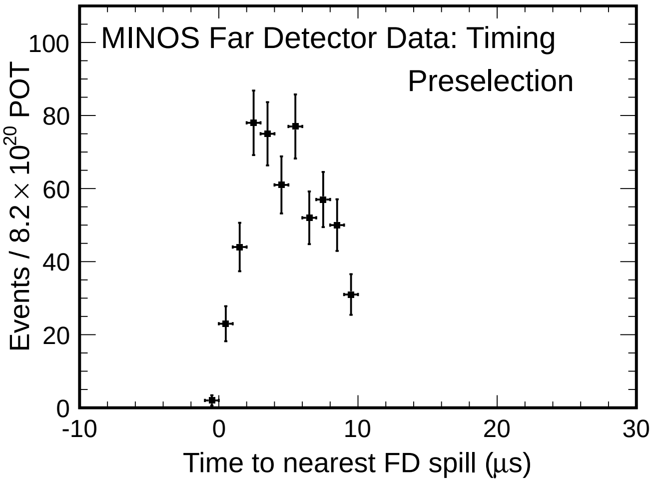

Figure 10.6: Time to nearest spill for Preselected Far Detector data events.

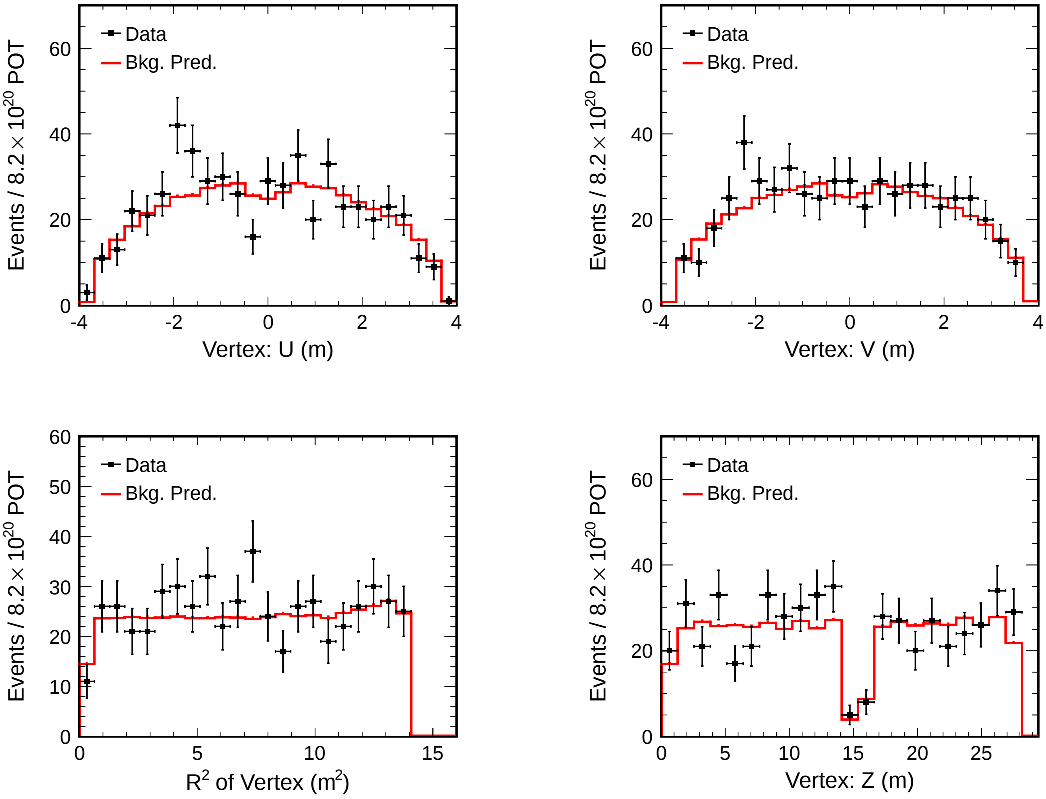

Figure 10.7: Event position vertices for Preselected Far Detector data vs. prediction. Shown are u vertex (top left), v vertex (top right), vertex radius squared (bottom left) and z vertex (bottom right).

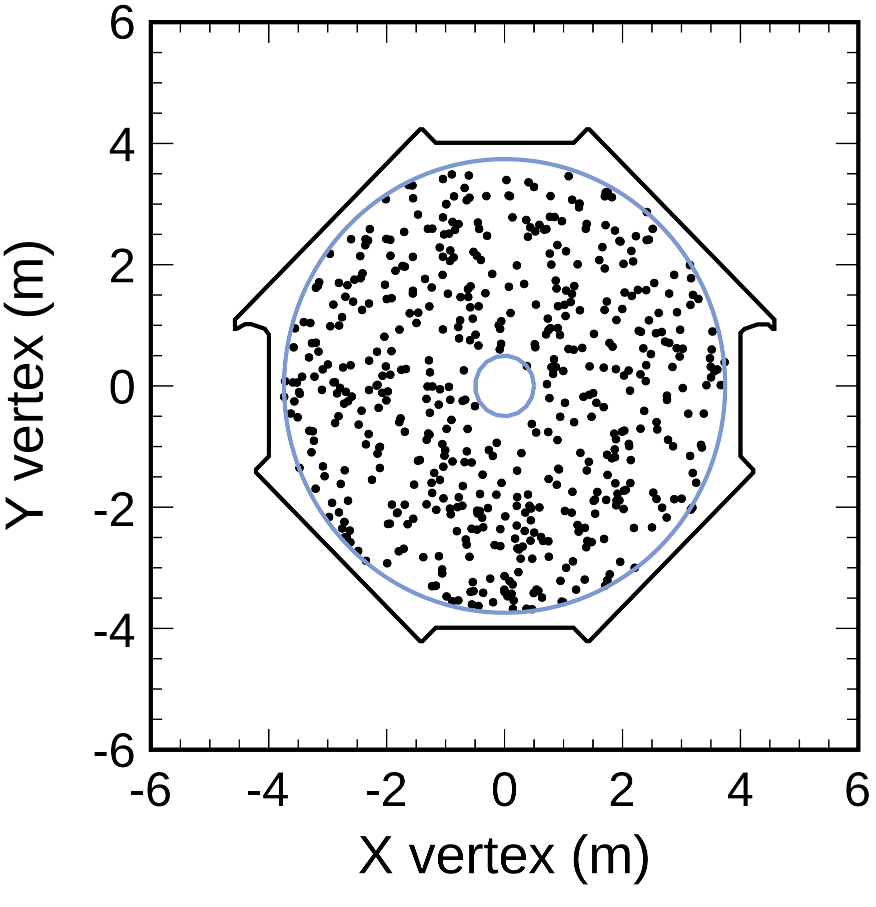

Figure 10.8: The xy positions of the Preselected FD Data. The blue lines indicate the fiducial region, and the black lines the physical edge of the detector.

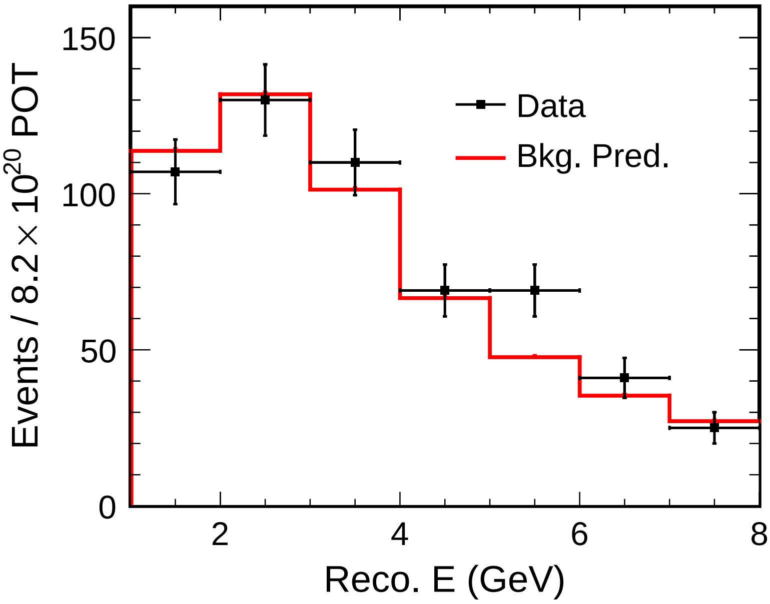

Figure 10.9: Event reconstructed energy for Preselected FD Data vs. prediction. For the standard decomposition, extrapolation, and fit, the final three energy bins (5-8 GeV) are merged into a signal bin.

due to the presence of an excess in the $5-8$ GeV fitting bin. While 109.5 events are expected in this range, 135 are observed, with most of the excess in the $5-6$ GeV range. The events in this bin were thoroughly studied, looking for any sign of anomalous behavior. Figure 10.10 shows one such study, looking at the vertex position of the events. All of the events appeared normal. A statistics-only Pearson's χ^2 of the distribution in Figure 10.9 yields a χ^2/ndf of 12.0/7, corresponding to a probability of p=0.10. The excess in the $5-6$ GeV range was therefore concluded to be a statistical fluctuation. The impact of this excess on the final fit is relatively small, due to the fit being conducted in bins of both PID and energy.

Figure 10.11 shows the prediction and FD Data distributions for the three LEM input variables. There is a good level of agreement between the FD Data and prediction for these distributions. While there is no immediate sign of a signal, there is also no sign of pathologies.

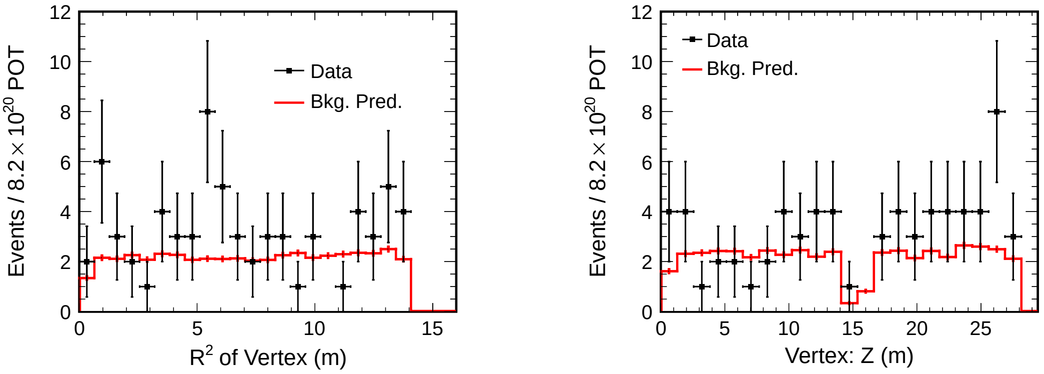

Figure 10.10: Positions of Preselected event vertices in the $5-6$ GeV region, as radius squared (left), and longitudinal position (right).

10.3.2 Fitting the Data

A fit to the FD Data is now performed in 8 LEM and 5 energy bins, using the methods described in Chapter 9. One final source of uncertainty must also be considered in the fit: oscillation uncertainties. These are uncertainties in the other oscillation parameters which are not being measured in this experiment, as summarized in Table 10.3. These uncertainties are incorporated into the fit in a manner similar to that described in Chapter 8, by shifting the FD MC by one standard deviation of the uncertainty on each parameter and then calculating the resulting fractional shift on the prediction.

Parameter	Best Fit Value	Uncertainty
$\|\Delta m^2_{32}\|$	2.32×10^{-3} eV^2	$(+0.12, -0.08) \times 10^{-3}$ eV^2
$\|\Delta m^2_{21}\|$	7.59×10^{-5} eV^2	$(+0.19, -0.21) \times 10^{-5}$ eV^2
$\sin^2(2\theta_{23})$	1.0	-0.06
$\sin^2(2\theta_{12})$	0.861	+0.03,-0.02

Table 10.3: Uncertainties on the oscillation parameters used in this analysis [42] [89].

The results of the 8×5 fit (and the cross-checks) are given in Table 10.4. For the 8×5 case, a best fit of $\sin^2(2\theta_{13}) = 0.042$ is obtained, with an upper 90% CL of $\sin^2(2\theta_{13}) = 0.115$. Figure 10.12 shows the prediction, data, and best fit signal in each of the 40 fit bins. Figure 10.13 shows the energy and PID distributions. Similar results are obtained for the LEM 3×5 bin fit and LEM>0.7 counting experiment. For the

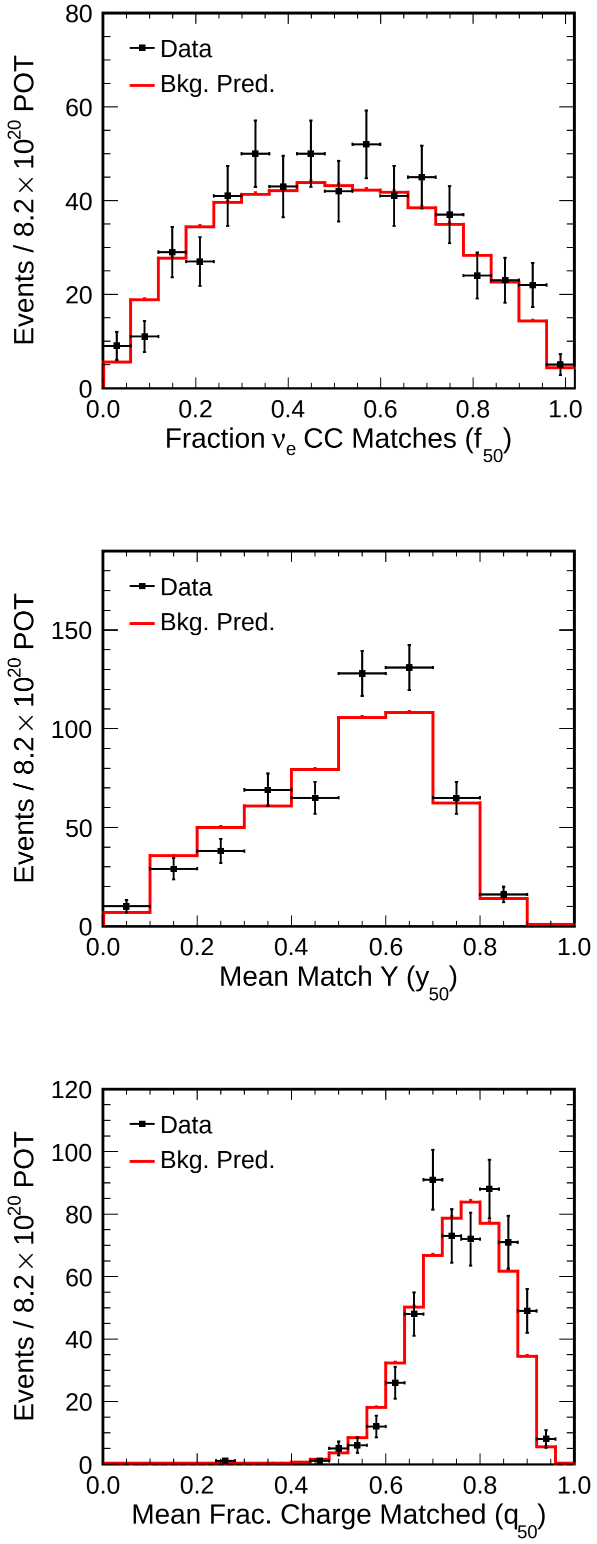

Figure 10.11: Far Detector data vs. prediction for the three LEM particle ID input variables: f_{50} (top), y_{50} (middle) and q_{50} (bottom). All events have passed the ν_e Preselection.

LEM>0.7 counting experiment, $49.9 \pm 7.1(\text{stat}) \pm 2.4(\text{syst})$ events are expected and 62 are observed, which corresponds to an excess with a statistical significance of 1.3σ. For comparison, the results are also shown for the previous ν_e appearance particle ID cut, ANN11>0.7; these results are consistent with the LEM fit.

Binning	Best Fit $\sin^2(2\theta_{13})$	Upper 90% CL	$-2\Delta \ln L$ at $\theta_{13} = 0$
8 LEM × 5 E	0.042	0.115	1.54
3 LEM × 5 E	0.037	0.122	1.20
LEM > 0.7	0.064	0.170	2.47
ANN11 > 0.7	0.040	0.147	0.83

Table 10.4: Results of various LEM fits to the Far Detector data. The numbers presented here assume δ_{CP}=0 and the Normal Hierarchy. Statistical, systematic, and oscillation parameter uncertainties are all included in the fit.

10.4 Interpreting the Results

The final result of the work in this thesis is a likelihood surface of $-2\Delta \ln L$ at each point in the parameter space.[1] The log likelihood surface for the 8×5 fit for both the Normal and Inverted Hierarchy can be seen in Figure 10.14. The lines in this image mark where $-2\Delta \ln L = 1.0$ and 2.71.

This thesis presents two different interpretations of this result, one Bayesian and the other Frequentist. The Frequentist approach answers the question of how often the same experimental results would be obtained, if the experiment were repeated many times. For this result, the exclusion of $\sin^2(2\theta_{13}) = 0$ is calculated using the Feldman-Cousins method [144]. The Bayesian approach gives a probabilistic interpretation of the data, in terms of θ_{13} and δ_{CP} for an assumed prior probability of θ_{13}. A credible interval is calculated for this result based on a prior flat in θ_{13} and δ_{CP}.

10.4.1 Interpretation 1: Feldman-Cousins Exclusion at Zero

The results of the fit in Table 10.4 become more difficult to interpret near $\sin^2(2\theta_{13}) = 0$. Below this value, any solutions are non-physical. Downward fluctuations, however, can

---

[1] Note that this log likelihood surface can easily be transformed into a raw likelihood surface by converting it to $e^{-Z/2}$, where $Z = -2\Delta \ln L$.

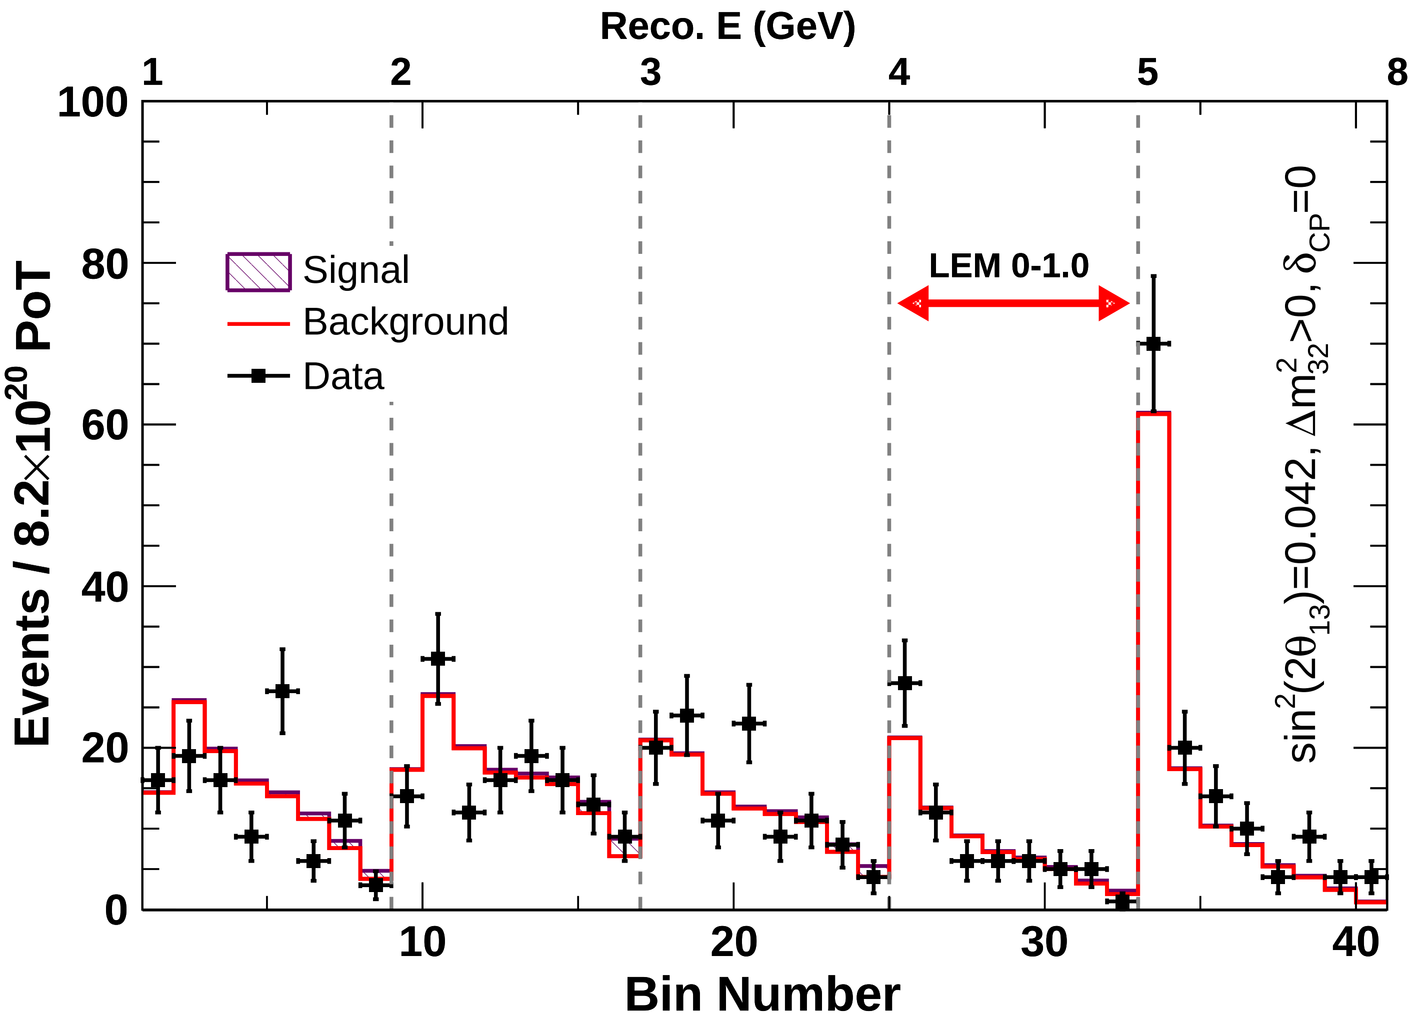

Figure 10.12: The FD Data compared to the prediction in each of the LEM 8 × 5 bins. The best-fit signal of $\sin^2(2\theta_{13}) = 0.042$ is also shown.

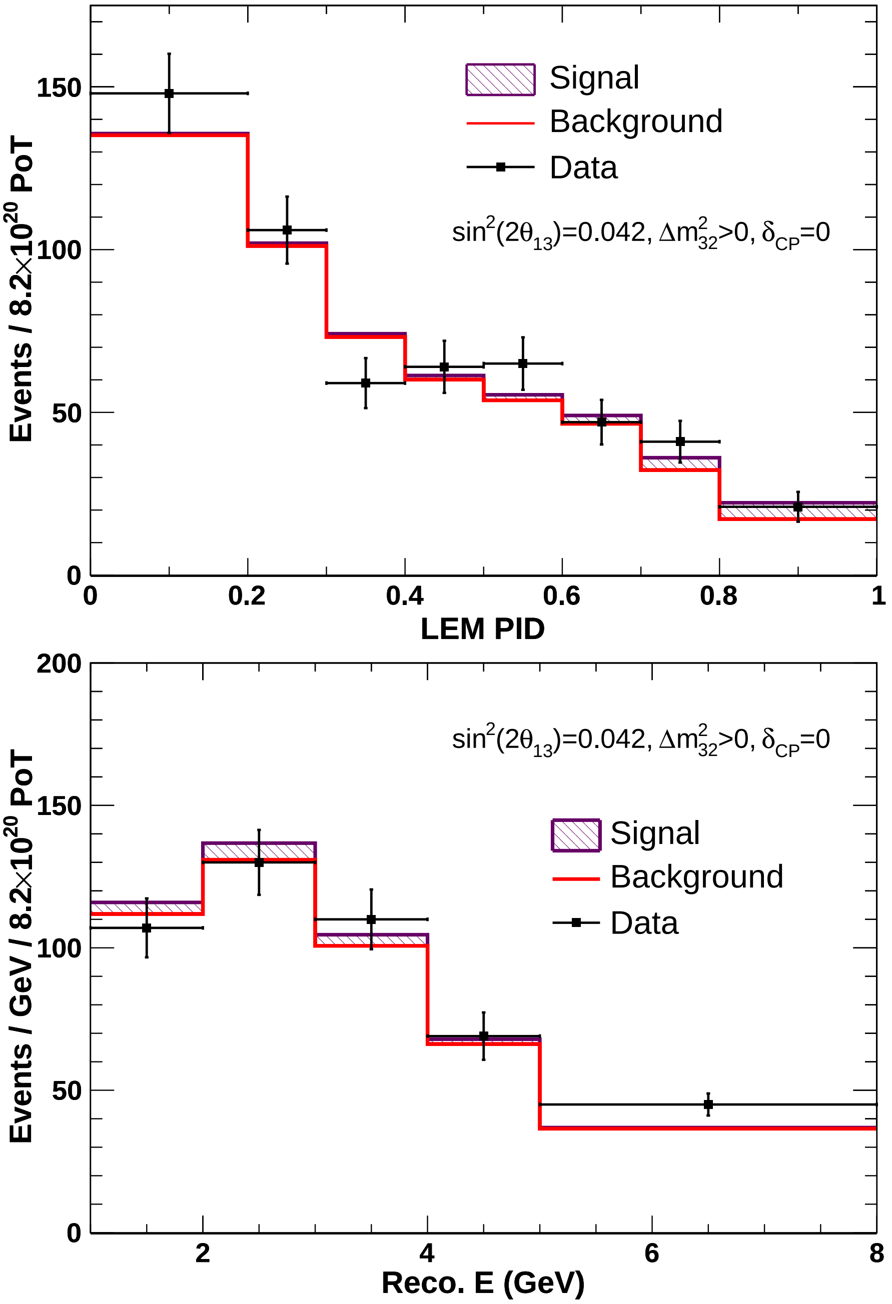

Figure 10.13: The LEM PID (top plot) and reconstructed energy (bottom plot) FD Data distributions, compared to predicted background and the best fit signal $\sin^2(2\theta_{13}) = 0.042$.

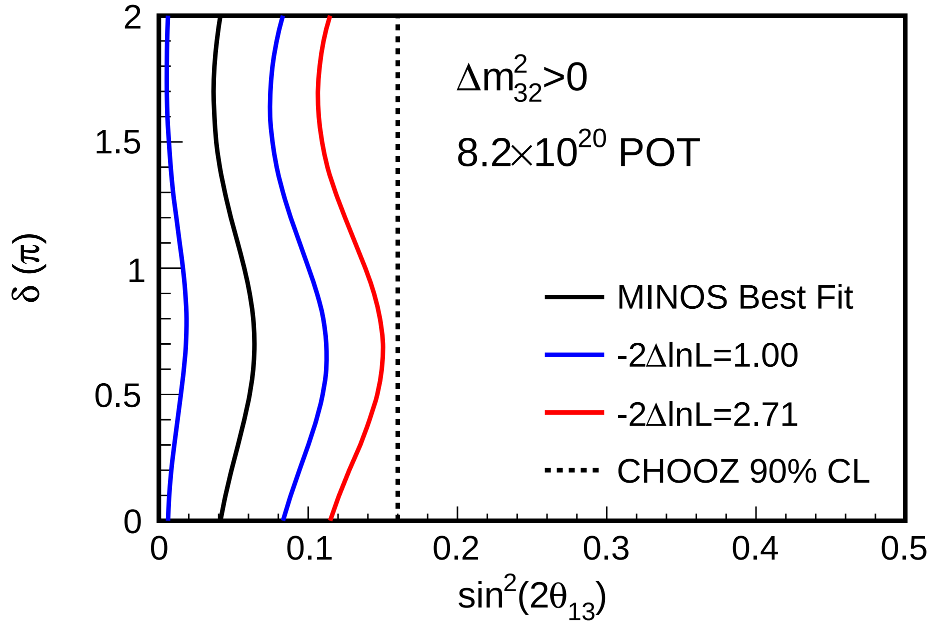

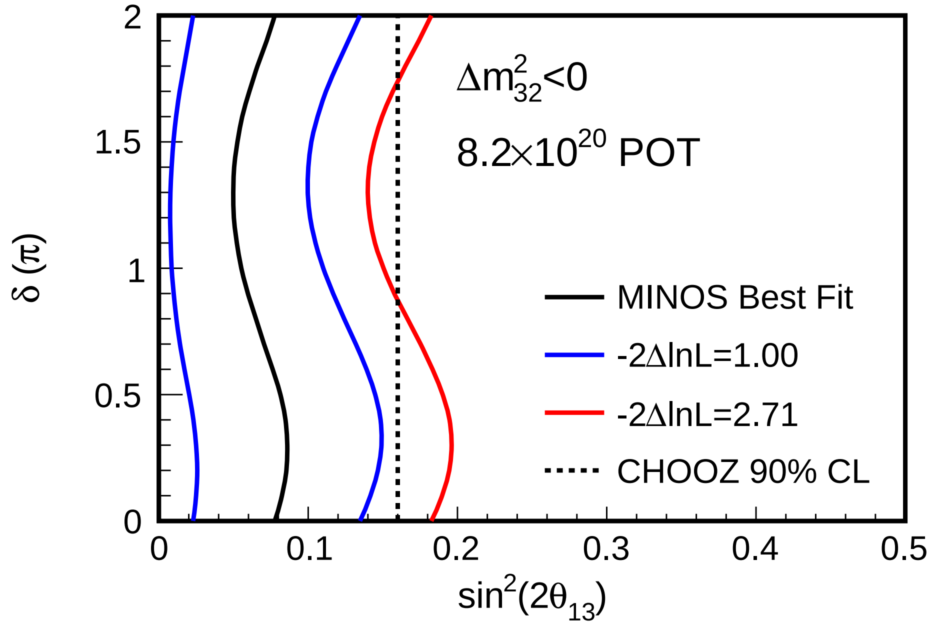

Figure 10.14: The log likelihood surface ($-2\Delta\ln L$) obtained from the experimental result. The best fit and values where $-2\Delta\ln L = 1.0$ and 2.71 are shown for the 8 LEM × 5 E bins fit. The 90% CL CHOOZ limit is indicated as well. Both the Normal (top) and Inverted (bottom) Hierarchy are shown.

result in a solution which may be beyond this boundary. For cases such as this, where the amount of data is not large, non-Gaussian behavior can occur. It is then appropriate to use a Feldman-Cousins approach to interpret the results and to produce Confidence Level contours in the physical region. In this thesis, the Feldman-Cousins method is only used to assess the significance of the exclusion of $\sin^2(2\theta_{13}) = 0$. The Feldman-Cousins approach is very similar to that used for the MRCC sideband toy MC study (Section 10.2.2). In this case, however, the fake experiments are generated with systematic error and oscillation uncertainty taken into account as well.

The first step in the FC process is already finished: the log likelihood of the data is calculated as a function of θ_{13}, and $-2\Delta \ln L^{data}$ is calculated at $\theta_{13} = 0$. These are the numbers listed in Table 10.4.

Next, 10,000 MC fake data pseudo-experiments are generated [120]. This is done in the following steps, using Gaussian and Poisson random number generators:

1. The predicted signal and background distributions (for $\theta_{13} = 0$) are fluctuated within their respective total Gaussian systematic errors σ_{sig} and σ_{bkg}.

2. The fluctuated signal and background distributions are added together, and the total distribution is fluctuated within the Gaussian oscillation uncertainty σ_{osc}. σ_{osc} is calculated by fluctuating the predicted distribution within the uncertainties outlined in Table 10.3 to produce a modified prediction. This is done 2000 times, and the fractional change calculated each time, to produce a probability distribution function which provides σ_{osc}.

3. Poisson statistical fluctuations are applied to the total distribution from 2. to form a pseudo-experiment.

Next, $-2\Delta \ln L^{pseudo}$ is calculated at $\theta_{13} = 0$ for each of the 10,000 pseudo-experiments. A histogram containing the values of $-2\Delta \ln L^{pseudo}$ is plotted and then integrated to find the fraction of experiments α for which $-2\Delta \ln L^{pseudo} < -2\Delta \ln L^{data}$. This value α is the confidence level of the result at that point in parameter space, in this case for $\sin^2(2\theta_{13}) = 0$.

When this procedure is carried out for the full 8×5 binning, $\sin^2(2\theta_{13}) = 0$ is found to be excluded at the 91% CL. Note that this limit applies for all values of δ_{CP} where $\theta_{13} = 0$. The exclusion for the various cross-checks are listed in Table 10.5.

Binning	$-2\Delta \ln L$	Feldman-Cousins CL at $\sin^2(2\theta_{13}) = 0$
8 LEM × 5 E	1.54	91%
3 LEM × 5 E	1.20	88%
LEM > 0.7	2.47	94%

Table 10.5: Feldman-Cousins exclusions of $\sin^2(2\theta_{13}) = 0$, for various binnings. The value of $-2\Delta \ln L$ at $\theta_{13} = 0$ is also listed.

10.4.2 Interpretation 2: Bayesian Credible Interval

The Frequentist approach asks how often a certain result would occur, if the experiment were repeated many times. The Bayesian interpretation attempts instead to quantify the resulting state of knowledge about the value of θ_{13}. Through Bayes' Theorem, one can calculate the *posterior probability* density function $P(x; data)$, the probability density function for a parameter x given the data:

$$P(x; data) = \frac{P(data; x)P(x)}{P(data)} \tag{10.5}$$

where $P(data; x)$ is the probability of the data as a function of x, $P(x)$ is the prior probability of x, and $P(data)$ is the prior probability of the data.

The first element in this theorem, $P(data; x)$, is simply the likelihood surface in Figure 10.14. The other two terms in Equation 10.5 are not determined by experiment. The prior probability of the data $P(data)$ simply acts as a normalization term, as it is independent of the fit parameters. The choice of $P(x)$, the prior probability for the θ_{13} and δ_{CP}, is subjective, because it involves making an assumption about what pre-existing information about the parameters to include. Nonetheless, there are reasonable assumptions that can be made. The most obvious choice is to use a flat prior, where $P(x)$ =constant. What is meant by "flat" is again subjective. ν_e appearance oscillation is detected by measuring the value of $\sin^2(2\theta_{13})$, so a prior flat as a function of $\sin^2(2\theta_{13})$ could be chosen. The underlying natural parameter being measured, however, is θ_{13}, the actual eigenstate rotation angle. A prior flat in θ_{13} and δ_{CP} is therefore the more natural choice.

This flat prior can easily be constructed by filling the likelihood surface with respect to θ_{13}, rather than $\sin^2(2\theta_{13})$. The log likelihood surface is then converted to probability and normalized to 1.0 to produce the probability density function $P(x; data)$, the outcome of Bayes' Theorem. From this function, one can then construct a Bayesian

Credible Interval to determine the parameter range in which there is a fractional probability α of the true value of θ_{13} being found. This is not a uniquely defined region, so likelihood ordering is used to determine the interval: the region is chosen so that it both contains a fraction α of the total probability, and so that the probability is higher in the included region than in the excluded region.

Both the 68% and 90% intervals for this result are shown in Figure 10.15 as a function of θ_{13}. Figure 10.16 translates these intervals back to $\sin^2(2\theta_{13})$, the experimentally measured parameter. The upper 90% limit is substantially closer to zero than the values estimated by the fit in Table 10.4. At $\delta_{CP} = 0$, the upper 90% limit is $\sin^2(2\theta_{13}) = 0.09(0.15)$ for the Normal (Inverted) Hierarchy. This lower limit results from the flat prior: when all values of θ_{13} are equally probable, a lower value of $\sin^2(2\theta_{13})$ becomes more probable. The probability drops more quickly for larger values of $\sin^2(2\theta_{13})$, causing a tighter upper 90% CL limit. This new limit is well below that of CHOOZ. Additionally, $\sin^2(2\theta_{13}) = 0$ is excluded at a confidence interval value of 86%, an exclusion not dissimilar from the above Feldman-Cousins result.

The need to choose a prior for θ_{13} makes this Bayesian Credible Interval a subjective interpretation of the result. For this reason, the true experimental result for the work in this thesis is the likelihood surface in Figure 10.14. This likelihood surface represents the sum total of knowledge gained about $\sin^2(2\theta_{13})$ from the data.

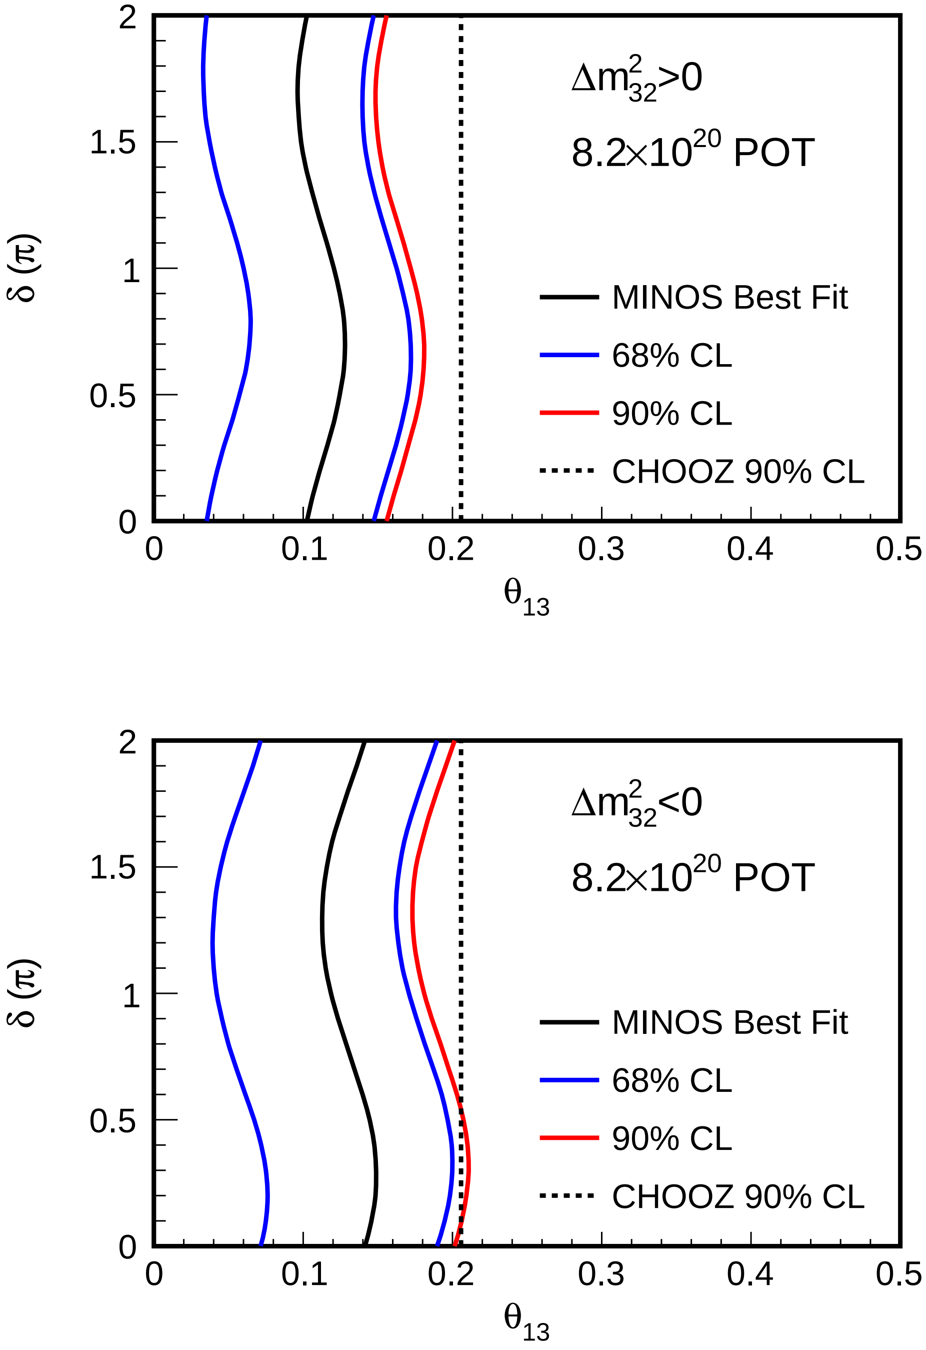

Figure 10.15: Bayesian Credible Interval for the experimental result, assuming a prior probability which is flat in θ_{13}. In this case, the result has been kept as a function of θ_{13}. The best fit, 68% CL, and 90% CL are shown; the fit was performed with 8 LEM × 5 E bins, for the Normal (top) and Inverted (bottom) Hierarchy.

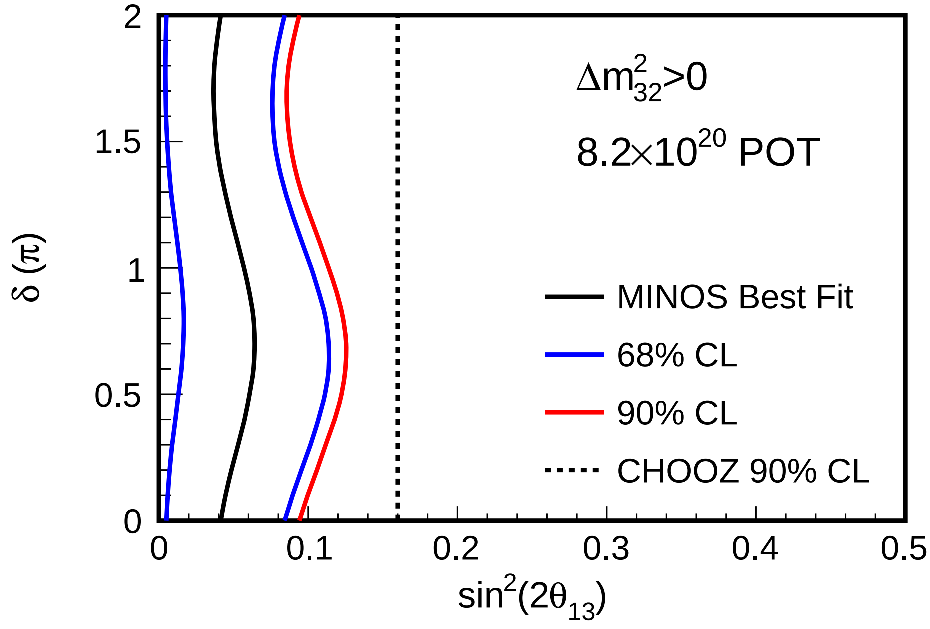

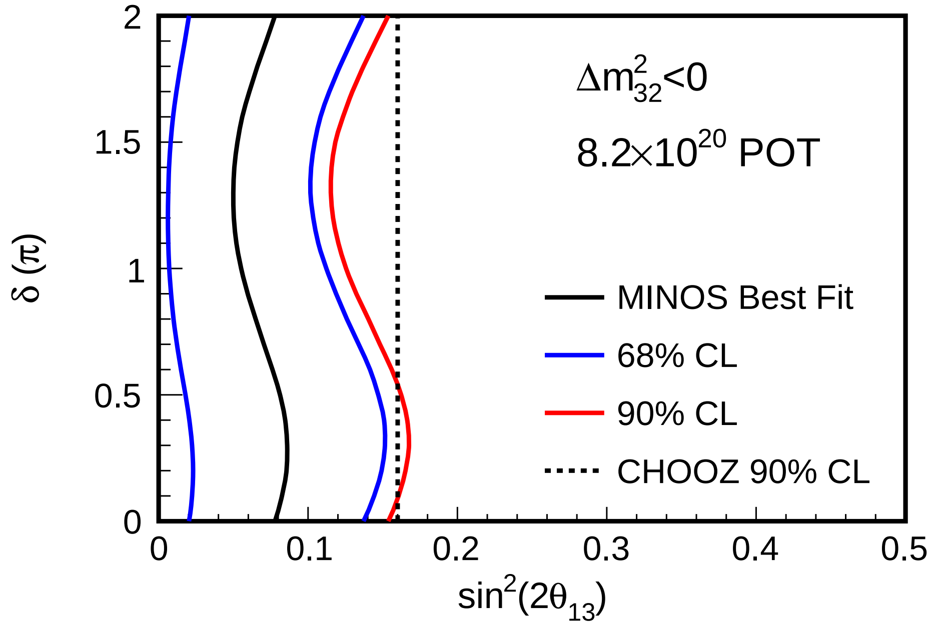

Figure 10.16: Bayesian Credible Interval for the experimental result, assuming a prior probability which is flat in θ_{13}. In this case, the result has been plotted versus $\sin^2(2\theta_{13})$, with the limits translated from Figure 10.15. The best fit, 68% CL, and 90% CL are shown; the fit was performed with 8 LEM × 5 E bins, for the Normal (top) and Inverted (bottom) Hierarchy.

Chapter 11

Conclusions: the Future of Neutrino Physics

11.1 Searching for $\nu_\mu \rightarrow \nu_e$ oscillation in MINOS

Several new experiments will come online in the next ten years which have been designed with the primary goal of making a measurement of θ_{13}. MINOS, with its large amount of existing neutrino data, can investigate new regions of the $\sin^2 2\theta_{13}$ parameter space, by looking for $\nu_\mu \rightarrow \nu_e$ oscillations between its Near and Far Detectors. As this thesis has shown, this is not an easy measurement to make. The expected appearance signal is very small, with background contamination and poor MINOS resolution for ν_e CC events posing major challenges for detection.

The work in this thesis greatly improves MINOS' sensitivity to $\sin^2(2\theta_{13})$. The basic analysis framework developed in two previous published MINOS analyses [85] [86] has been improved in two major ways. First, the author has implemented and optimized a novel particle ID scheme, the Library Event Matching method, which uses information from the event strip hits themselves. This new PID gives a 14% improvement in sensitivity to $\sin^2(2\theta_{13})$. Second, the work in this thesis has transformed the analysis from a counting experiment to a multiple bin likelihood fit. An additional 14% gain in sensitivity is gained by fitting the final data in bins of both PID and reconstructed energy. This thesis fully develops the framework for this fit, and for the interpretation of the final result.

In the end, no statistically significant proof of $\nu_\mu \rightarrow \nu_e$ oscillation was found. However, there are hints of a possible excess. A Feldman-Cousins analysis of the final result excludes $\sin^2 2\theta_{13} = 0$ at the 91% CL. The result also places the tightest limits on

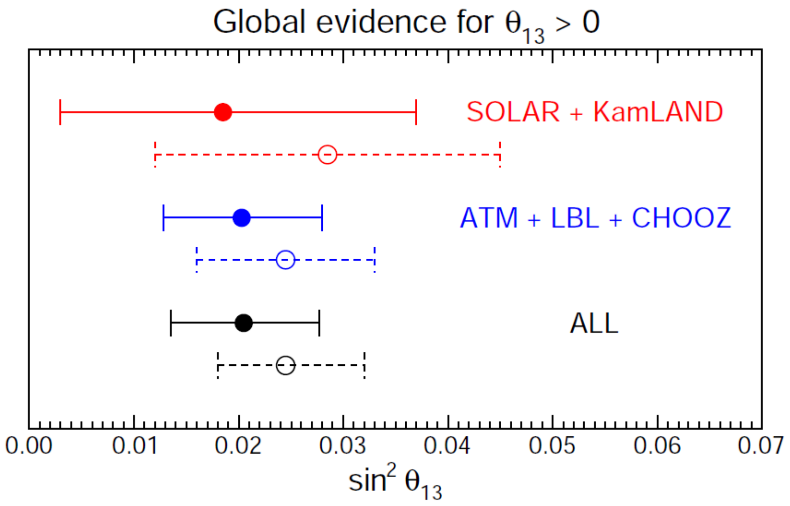

Figure 11.1: Results of a 2011 global fit [146] for for the value of θ_{13}. Fit assumes three-neutrino oscillation. Solid lines use old reactor neutrino fluxes, while dashed lines use the new fluxes [44].

$\sin^2(2\theta_{13})$ to date, for a reasonable choice of Bayesian prior. A prior flat in θ_{13}, for instance, places the upper 90% limit of a credible interval at $\theta_{13} = 0.16$, or $\sin^2 2\theta_{13} = 0.09$ (for the Normal Hierarchy).

11.2 Future Searches for θ_{13}

In the summer of 2011, evidence began to emerge for a non-zero value of θ_{13}. Some of this evidence came from the MINOS experiment, both from the work in this thesis, and from the published result described in [51]. Further proof came from the T2K experiment, which saw a small signal with a statistical significance of 2.5σ. Both the T2K result and the 2011 official MINOS result are briefly described in Appendix A. Figure 11.1 shows the results of a global fit [146] for θ_{13}, taking into account both the new measurements and previous constraints from long baseline, atmospheric, and solar results. This fit finds evidence for $\theta_{13} > 0$ at more than 3σ, with a best fit of $\sin^2(2\theta_{13}) = 0.098$.[1]

This evidence for a relatively large, non-zero value of θ_{13} gives extra urgency to the next generation of experiments. First, several new inverse β decay reactor experiments will attempt to make precise, direct measurements of the value of θ_{13}. Double CHOOZ [147], which adds a new detector to CHOOZ, hopes to reach a projected sensitivity of $\sin^2 2\theta_{13} < 0.03$ after three years of data taking. The Daya Bay experiment [148] in

---

[1] This best fit assumes the newest measurements of flux for reactor experiments [44]. A best fit of $\sin^2(2\theta_{13}) = 0.082$ is found if the older versions of the fluxes are used.

China aims for a sensitivity of $\sin^2 2\theta_{13} < 0.01$. A third reactor experiment in South Korea, RENO [149], will have a comparable sensitivity. In addition to making very precise measurements of θ_{13}, these experiments will also allow a comparison of ν_e and $\bar{\nu}_e$ oscillation.

The T2K experiment will continue to gather data and refine its measurement. The NOvA experiment [150] will also conduct a search for $\nu_\mu \rightarrow \nu_e$ oscillations. NOvA will use an upgraded version of the Fermilab NuMI beamline, with a Near Detector at Fermilab, and a Far Detector sitting at an off-axis angle of 14.6 mrad in Ash River, Minnesota, 810 km away. NOvA's sensitivity to θ_{13} is comparable to that of T2K. The longer baseline of NOvA causes matter effects to have a much larger impact on neutrino oscillation. This could allow NOvA to make a determination of the sign of Δm^2_{32} (i.e., to determine the mass hierarchy). NOvA could also place significant constraints on δ_{CP}.

The next decade will be an extremely exciting one for neutrino physics, in which many fundamental questions may be answered. The results of this thesis are a good first step, but only a prelude to the discoveries which will potentially come soon.

Appendix A

Other Results: MINOS and T2K

A.1 The T2K Experiment

In June of 2011, the T2K experiment [145] announced that they had observed a statistically significant ν_e appearance signal result. The basic concept of the T2K experiment, and a brief summary of the results, are discussed here. Like MINOS, T2K (Tokai-2-Kamioka) searches for $\nu_\mu \rightarrow \nu_e$ oscillation by comparing the flux of neutrinos at both a Near and a Far Detector.

The NuMI beam used by MINOS is a wideband neutrino beam. A large neutrino flux is produced, at a relatively wide spread of energies, requiring a very precise simulation of the beam kinematics and decays. T2K differs from MINOS in using an off-axis neutrino beam. An off-axis beam produces a much narrower neutrino spectrum. As explained earlier, the energy of a neutrino produced in a two-body decay is dependent both on decay angle and on the parent energy:

$$E_\nu = \frac{(1 - \frac{m_\mu}{M^2})E}{1 + \gamma^2 \tan^2 \theta_\nu} \tag{A.1}$$

Figure A.1 shows the energy of neutrinos produced from pions decaying at a given momentum, for various angles off a beamline. As the angle off the beamline increases, the neutrino beam becomes nearly monochromatic. This monochromatic beam can be centered near the $\nu_\mu \rightarrow \nu_e$ oscillation energy peak, and is also relatively free of difficult-to-model kaon decays. An off-axis beam is therefore very effective for making precise measurements of θ_{13}. T2K uses a beam which is 2.5° off-axis, resulting a neutrino energy peak of ~ 0.6 GeV. At the beginning of the muon neutrino beam (280 m downstream), the beam is monitored both by an on-axis detector (INGRID) which characterizes the beam profile, direction, and intensity, as well as an off-axis Near Detector which provides

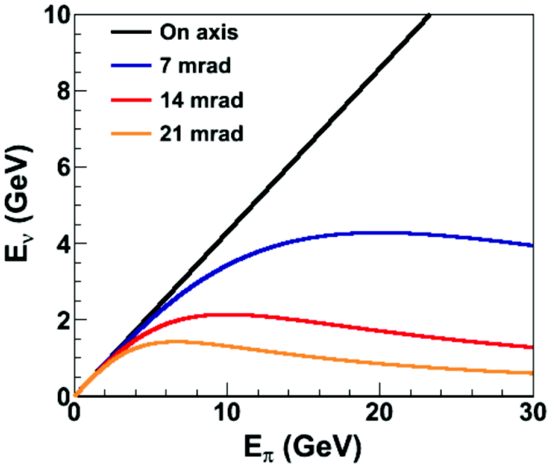

Figure A.1: The concept behind an off-axis beam. Neutrino energy is shown as a function of decay pion parent energy for various angles off the main beam axis. A larger angle results in a more monochromatic beam. Figure is for a NOvA-style beam, and was produced by Patricia Vahle [150].

comparisons with the Far Detector. The Far Detector is the Super-Kamiokande detector, at a distance of 295 km. Signal ν_e appearance events are identified at the Far Detector via electrons from quasi-elastic Charged Current interactions. The largest sources of background are the intrinsic beam ν_e component and misidentified π^0s.

Operation of T2K was shut down after the March 2011 earthquake, but in June 2011 results were presented from the first 1.43×10^{20} protons on target. Following all analysis cuts, T2K expected to see a total of 1.5 ± 0.3 background events in the Far Detector. They observed 6, corresponding to a 2.5σ excess. There is a 0.7% chance that this result could have been caused by a statistical fluctuation. The resulting Feldman-Cousins limits on $\sin^2 2\theta_{13}$ are shown in Figure A.2. The Normal Hierarchy best fit for the T2K result was found to be $\sin^2 2\theta_{13} = 0.11$, with a lower 90% CL limit of 0.03, and an upper 90% CL limit of 0.28.

A.2 The MINOS 2011 ν_e Analysis

The 2011 MINOS ν_e analysis, published in [51], proceeded along lines very similar to the work in this thesis. Many of the techniques used were developed concurrently by the author for this thesis, including the updated LEM PID and the multiple bin fit. There were several major differences, which will be briefly described here. For the fit, the bin-by-bin covariance matrix method was used to incorporate systematic error. The

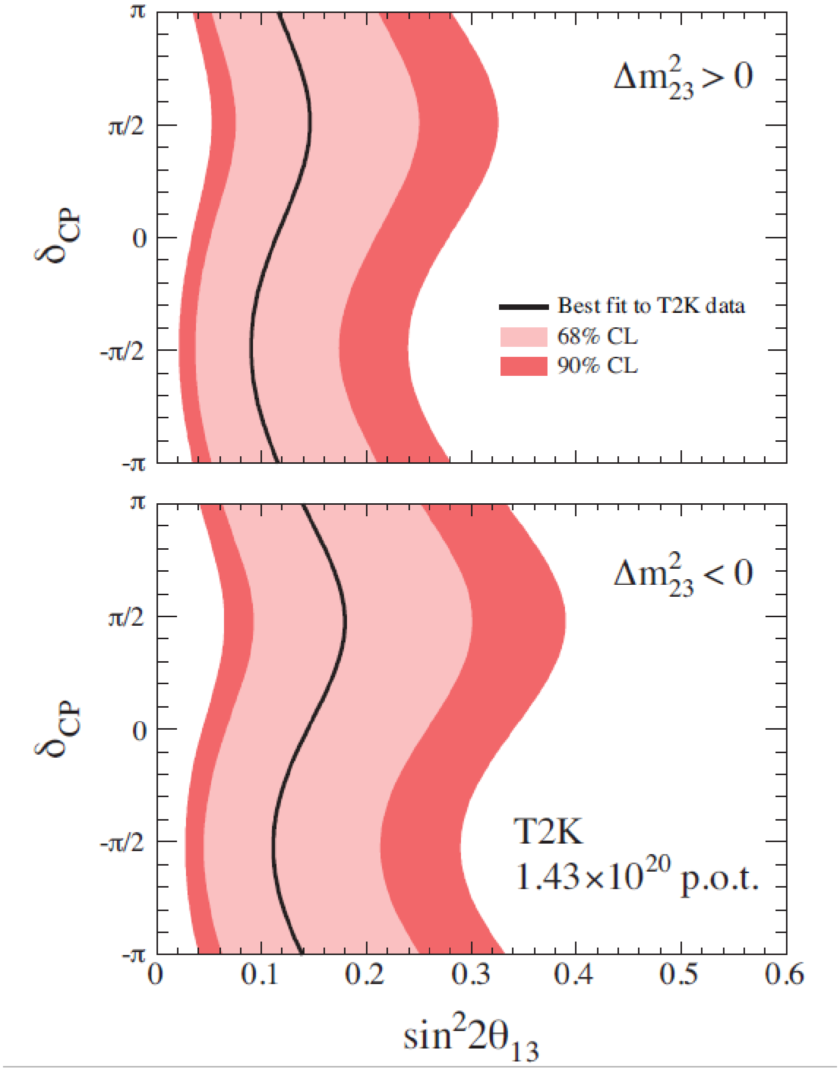

Figure A.2: 2011 T2K ν_e appearance results. This figure shows the Feldman-Cousins 90% and 68% CL for $\sin^2 2\theta_{13}$, as a function of δ_{CP} and choice of mass hierarchy: Normal (top) or Inverted (bottom). The best fit is shown as well. Figure from [145].

fit was conducted in the 3 LEM × 5 E bins used as a cross-check in this thesis. The final result is interpreted using a full Feldman-Cousins contour, shown in Figure A.3. Many small changes implemented in this analysis, such as rock muon corrections, were not used. The results were very similar to those presented in this thesis. A best fit of $\sin^2(2\theta_{13}) = 0.040$ (0.078) is found for the Normal (Inverted) Hierarchy, with an upper 90% CL of 0.12 (0.19). $\sin^2(2\theta_{13}) = 0$ is excluded at the 89% CL.

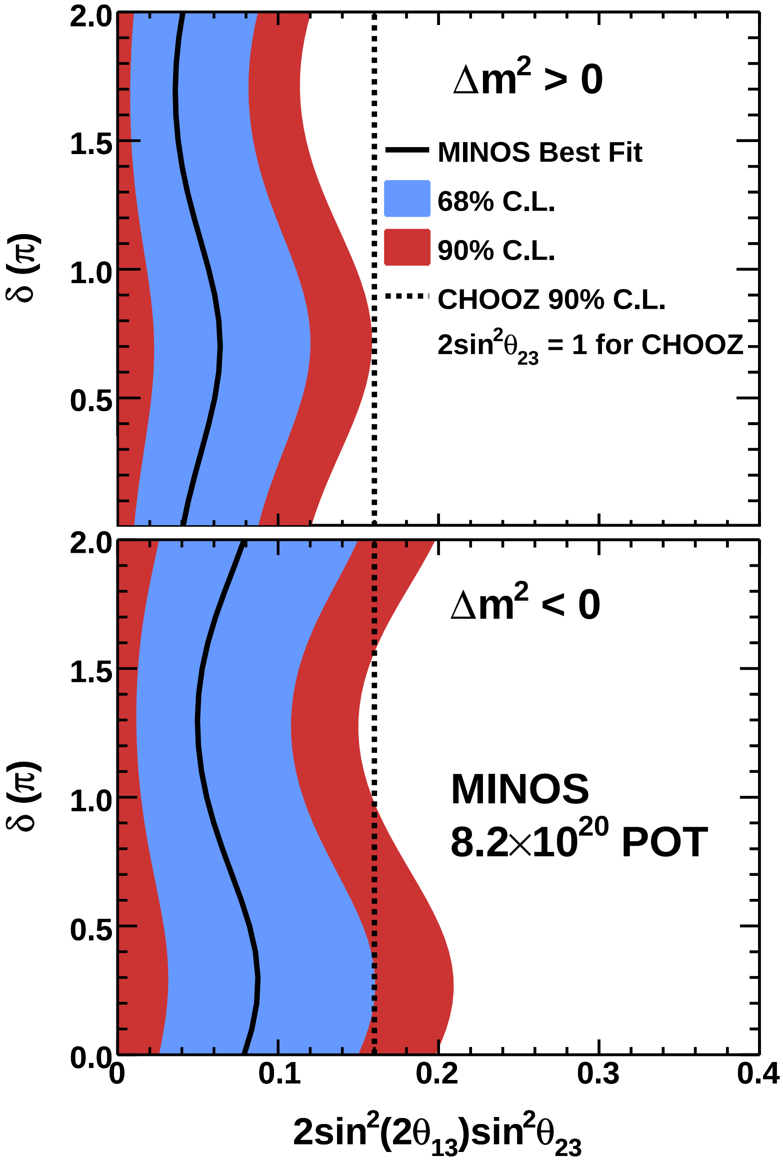

Figure A.3: 2011 MINOS ν_e appearance results. Feldman-Cousins 90% and 68% CL for $\sin^2 2\theta_{13}$, as a function of δ_C and choice of mass hierarchy: Normal (top) or Inverted (bottom). The best fit is shown as well. Figure from [51].

Bibliography

[1] J. Chadwick, Distribution in Intensity in the Magnetic Spectrum of the β-rays of Radium, Ver. Dtsch. Physik. Ges. **16**, 383-391 (1914).

[2] C. D. Ellis and W. A. Wooster, The Average Energy of Disintegration of Radium E, Proceedings of the Royal Society of London. Series A **117**(776), 109-123 (1927).

[3] W. Pauli. Letter to a physicists' gathering at Tubingen. December 4, 1930, reprinted in *Wolfgang Pauli, Collected scientific papers*, edited by R. Kronig and V. Weisskopf, vol. 2, p. 1313, New York, Interscience, 1964.

[4] J. Chadwick, Possible Existence of a Neutron, Nature **129**, 312 (1932).

[5] J. Cowan, C. L., F. Reines, et al., Detection of the Free Neutrino: a Conrmation, Science **124**(3212), 103-104 (1956).

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

[7] K. Kodama et al., Observation of tau-neutrino interactions, Phys. Lett. **B504**, 218-224 (2001).

[8] J. Davis, Raymond, D. S. Harmer, and K. C. Hoffman, Search for neutrinos from the sun, Phys. Rev. Lett. **20**, 1205-1209 (1968).

[9] J. N. Bahcall, Solar neutrinos. I: Theoretical, Phys. Rev. Lett. **12**, 300-302 (1964).

[10] J. N. Bahcall et al., Solar Models: current epoch and time dependences, neutrinos, and helioseismological properties, Astrophysical Journal, **555**, 990-1012 (2001).

[11] K. S. Hirata et al., Real time, directional measurement of B-8 solar neutrinos in the Kamiokande-II detector, Phys. Rev. **D44**, 2241-2260 (1991).

[12] W. Hampel et al., GALLEX solar neutrino observations: Results for GALLEX IV, Phys. Lett. **B447**, 127-133 (1999).

[13] D. N. Abdurashitov et al., The Russian-American gallium experiment (SAGE) Cr neutrino source measurement, Phys. Rev. Lett. **77**, 4708-4711 (1996).

[14] B. Pontecorvo, Inverse beta processes and nonconservation of lepton charge, JETP (34), 172-173 (1958).

[15] K. Nakamura et al. (Particle Data Group), JP G 37, 075021 (2010) and 2011 partial update for the 2012 edition (URL: http://pdg.lbl.gov)

[16] S. L. Glashow, Partial symmetries of weak interactions, Nucl. Phys. **22**, 579-588 (1961).

[17] A. Salam and J. C. Ward, Electromagnetic and weak interactions, Phys. Lett. **13**, 168-171 (1964).

[18] S. Weinberg, A model of leptons, Phys. Rev. Lett. **19**, 1264-1266 (1967).

[19] G. Arnison et al., Experimental observation of lepton pairs of invariant mass around 95-GeV/c^2 at the CERN SPS collider, Phys. Lett. **B126**, 398410 (1983).

[20] C. Conta, Production of W and Z in the UA2 experiment at CERN $\bar{p}p$ collider, Annals N. Y. Acad. Sci. **461**, 666-676 (1986).

[21] The ALEPH Collaboration. Precision electroweak measurements on the Z resonance. Phys. Rept., 427:257, 2006.

[22] C. S. Wu et al., Experimental Test of Parity Conservation in Beta Decay, Phys. Rev. **105**(4), 1413-1415 (1957).

[23] M. Thomson, "Methods and problems in experiments with neutrino beam (short and long baseline)," SAPP, Varenna, August 2011.

[24] Durham Database Group, HEPDATA: Reaction Data Database, URL http://durpdg.dur.ac.uk/HEPDATA/REAC.

[25] E. Majorana, Nuovo Cimento **14**, 171 (1937).

[26] M. Gell-Mann, P. Ramond and R. Slansky, in *Supergravity,* ed. by D. Freedman et al., North Holland (1979).

[27] T. Yanagida, Prog. Theor. Phys. **64**, 1103 (1980).

[28] E. Komatsu et al., Five-Year Wilkinson Microwave Anisotropy Probe (WMAP) Observations: Cosmological Interpretation, Astrophys. J. Suppl. **180**, 330-376 (2009).

[29] C. Kraus et al., Final Results from phase II of the Mainz Neutrino Mass Search in Tritium β Decay, Eur. Phys. J. **C40**, 447-468 (2005).

[30] H.V. Klapdor-Kleingrothaus et al., Phys. Lett. **B** 586, 198 (2004).

[31] S. Davidson et al.,Leptogenesis. Phys. Rept. **466**:105-177, 2008

[32] B. Pontecorvo, Mesonium and antimesonium, Sov. Phys. JETP **6**, 429 (1957).

[33] Z. Maki, M. Nakagawa, and S. Sakata, Remarks on the Unifed Model of Elementary Particles, Progress of Theoretical Physics **28**(5), 870-880 (1962).

[34] Q. R. Ahmad et al., Measurement of the Rate of $\nu_e + d \to p + p + e^-$ Interactions Produced by ^{8}B Solar Neutrinos at the Sudbury Neutrino Observatory, Phys. Rev. Lett. **87**(7), 071301 (2001).

[35] S. Abe et al. (the KamLAND Collaboration), Precision Measurement of Neutrino Oscillation Parameters with KamLAND, Physical Review Letters **100**(22), 221803 (2008).

[36] G. D. Barr, T. K. Gaisser, S. Robbins, and T. Stanev, Uncertainties in atmospheric neutrino fluxes, Phys. Rev. **D74**, 094009 (2006).

[37] K. S. Hirata et al., Experimental Study of the Atmospheric Neutrino Flux, Phys. Lett. **B205**, 416 (1988).

[38] T. J. Haines et al., Calculation of atmospheric neutrino induced backgrounds in a nucleon decay search, Phys. Rev. Lett. **57**, 1986-1989 (1986).

[39] M. Ambrosio et al., Matter effects in upward-going muons and sterile neutrino oscillations, Phys. Lett. **B517**, 59-66 (2001).

[40] M. C. Sanchez et al., Observation of atmospheric neutrino oscillations in Soudan 2, Phys. Rev. D **68**, 113004 (2003).

[41] Y. Ashie et al., A Measurement of Atmospheric Neutrino Oscillation Parameters by Super-Kamiokande I, Phys. Rev. D **71**, 112005 (2005).

[42] P. Adamson et al., Measurement of the neutrino mass splitting and flavor mixing by MINOS, Phys. Rev. Lett. **106**, 181801 (2011).

[43] M. Apollonio et al., Search for neutrino oscillations on a long base-line at the CHOOZ nuclear power station, Eur. Phys. J. **C27**, 331-374 (2003).

[44] Th. A. Mueller et al., Improved Predictions of Reactor Antineutrino Spectra, Phys. Rev. **C**83, 054615, (2011).

[45] G. Mention et al., The Reactor Antineutrino Anomaly, Phys. Rev. D **83**, 073006 (2011).

[46] L. Wolfenstein, Neutrino oscillations in matter, Phys. Rev. D **17**(9), 2369-2374, (1978)

[47] S. P. Mikheyev and A. Y. Smirnov, Resonance enhancement of oscillations in matter and solar neutrino spectroscopy, Sov. J. Nucl. Phys. **42**, 913-917 (1985).

[48] J. Boehm, A Measurement of Electron Neutrino Appearance with the MINOS Experiment, Ph.D. Thesis, Harvard (2009).

[49] E. K. Akhmedov et al., Series expansions for three-flavor neutrino oscillation probabilities in matter, Journal of High Energy Physics 2004(04), 078 (2004).

[50] J. P. Ochoa, A Search for Muon Neutrino to Electron Neutrino Oscillations in the MINOS Experiment, Ph.D. Thesis, Caltech (2009).

[51] P. Adamson et al, Improved Search for muon-neutrino to electron-neutrino oscillations in MINOS, arXiv:1108.0015v1 [hep-ex]

[52] D.G. Michael et al., The magnetized steel and scintillator calorimeters of the MINOS experiment, Nucl. Instrum. Meth. **A596**, 190-228 (2008).

[53] P. Adamson et al., A Study of Muon Neutrino Disappearance Using the Fermilab Main Injector Neutrino Beam, Phys. Rev. **D77**, 072002 (2008).

[54] A. G. Abramov et al., Nucl. Instrum. Meth. **A485**, 209 (2002).

[55] R. Zwaska et al., Beam-based alignment of the NuMI target station components at FNAL, Nucl. Instrum. Meth. **A568**, 548-560 (2006).

[56] P. Adamson, “Accelerator Update and Target Saga,” MINOS-doc-8243, 2010.

[57] B.J. Rebel, Neutrino-induced muons in the MINOS far detector, Ph.D.Thesis, Indiana University, Fermilab-Thesis-2004-33, 2004.

[58] C.L.F. Howcroft, Atmospheric neutrinos in the MINOS far detector, Ph.D. Thesis, University of Cambridge, Fermilab-Thesis-2004-71, 2004.

[59] A.S.T. Blake, A study of atmospheric neutrino oscillations in the MINOS far detector, Ph.D. Thesis, University of Cambridge, Fermilab-Thesis-2005-77, 2005.

[60] P.Adamson, et al., Nucl. Instr. and Meth. **A566**, 119 (2006).

[61] MINOS Collaboration, "Calibration Position Paper for Pre-Shutdown Data", MINOS-doc 3941, 2008.

[62] K. Lang et al., Characterization of 1600 Hamamatsu 16-anode photomultipliers for the MINOS Far detector, Nucl. Instrum. Meth. **A545**, 852-871 (2005).x

[63] P. Adamson, et al., Nucl. Instr. and Meth. **A492**, 325 (2002).

[64] P. Adamson, et al., Nucl. Instr.and Meth. **A521**, 361 (2004).

[65] N. Tagg, et al., Nucl. Instr. and Meth. **A539**, 668 (2005).

[66] K. Lang, et al., Nucl. Instr. and Meth. **A545**, 852 (2005).

[67] J.J.Hartnell, Measurement of the calorimetric energy scale in MINOS, Ph.D. Thesis, University of Oxford, Fermilab-Thesis-2005-51, 2005.

[68] P.L. Vahle, Electromagnetic interactions in the MINOS detectors, Ph.D. Thesis, University of Texas at Austin, Fermilab-Thesis-2004-35, 2004.

[69] M.A. Kordosky, Hadronic interactions in the MINOS detectors, Ph.D. Thesis, University of Texas at Austin, Fermilab-Thesis-2004-34, 2004.

[70] A. Ferrari, P. R. Sala, A. Fasso, and J. Ranft (2005), CERN-2005-010.

[71] R. Brun et al., GEANT Detector Description and Simulation Tool, CERN Program Library Long Writeup W5013 (1994).

[72] C. Zeitnitz and T. A. Gabriel, Nucl. Instr. Meth. A349, 106 (1994).

[73] Sacha Kopp, Zarko Pavlovic, and Patricia Vahle. "SKZP2: Update on Fitting the Beam MC to the ND Data." Minos-doc-1650,

[74] H. Gallagher. The NEUGEN neutrino event generator. Nuclear Physics B - Proceedings Supplements, **112**(1-3), 188-194, 2002.

[75] Dieter Rein and Lalit M. Sehgal. Neutrino-excitation of baryon resonances and single pion production. Annals of Physics, **133**(1, 79-153, 1981.

[76] H. Plothow-Besch. Program PDFLIB in CERN Program Library Pool W5051, 1991.

[77] Dieter Rein and Lalit M. Sehgal. Coherent π^0 production in neutrino reactions. Nuclear Physics B, **223**(1), 29-44, 1983.

[78] R. Merenyi et al., Determination of pion intranuclear rescattering rates in ν_μ-Ne versus ν_μ-D interactions for the atmospheric ν flux, Phys. Rev. D **45**(3), 743-751 (1992).

[79] PYTHIA/JETSET. http://home.thep.lu.se/ torbjorn/Pythia.html.y

[80] Serge Rudaz and Pierre Valin. Koba-Nielsen-Olesen scaling, its violation, and the structure of hadrons. Phys. Rev. D, **34**(7), 2025-2042 (1986).

[81] PhotonTransport and DetSim documentation: http://www-numi.fnal.gov/ computing/d120/releases/development/WebDocs/detsim/DetSim_Documentation.html

[82] T. Yang, G. Pawloski, MINOS-doc-5319-v2 2008.

[83] R. Patterson, "New results for $\nu_\mu \rightarrow \nu_e$ oscillations in MINOS." Wine and Cheese Talk, FNAL, April 9th 2010.

[84] S. Cavanaugh, A Search for Muon Neutrino to Electron Neutrino Oscillations in the MINOS Experiment, Ph.D. Thesis, Harvard (2010).

[85] P. Adamson et al., "Search for muon-neutrino to electron-neutrino transitions in MINOS," Phys.Rev.Lett. 103. 261802, 2009.

[86] P. Adamson et al., New constraints on muon-neutrino to elecron-neutrino transitions in MINOS, Phy.Rev.(Rapid Communications) **D82**, 051102 (2010).

[87] G.L. Fogli et al., Neutrino oscillations, global analysis and θ_{13}. Talk presented at NEUTEL 2009, XIII International Workshop on "Neutrino Telescopes" (Venice, Italy, March 10-13, 2009).

[88] X. Huang. "Nue Calibration." MINOS-doc-6702, 2011.

[89] B. Aharmim et al., An Independent Measurement of the Total Active B-8 Solar Neutrino Flux Using an Array of He-3 Proportional Counters at the Sudbury Neutrino Observatory. Phys. Rev. Lett. 101, 111301 (2008).

[90] J. Ling and T. Yang. “MINOS ANN based NuE Selection Methods.” MINOS-doc-6762, 2009.

[91] M. Thomsoon, “A new approach to the MINOS electron neutrino search.” MINOS-doc-1398, 2006.

[92] J. P. Ochoa, “The Monte Carlo Nearest Neighbors (MCNN) Selection.” MINOS-doc-5031, 2008.

[93] A. Holin, “MRCC correction method for the determination of nue BG.” MINOS-doc-4018-v1, 2008.

[94] N. Tagg, “Gain Calibration Systematics’.’ MINOS-doc-4865-v2, 2008.

[95] “Evaluation of systematic uncertainties on the measurement of Nue CC Appearance measurement.” J.Boehm, M. Sanchez, MINOS-doc-5387-v6, 2008.

[96] M. Orchanian, “Update on LEM Modification.” MINOS-doc-6372-v1, 2009.

[97] Numbers provided by Mhair Orchanian.

[98] R. Toner, “LEM Update: Preselection, Processing.” MINOS-doc-6377-v2, 2009.

[99] T. Vahle, “Nue PID on caldet electrons.” MINOS-doc-5546, 2009.

[100] R. Patterson, “The impact of background composition on nue sensitivity.” MINOS-doc-3760-v1, 2007.

[101] A. Hoecker, P. Speckmayer, J. Stelzer, J. Therhaag, E. von Toerne, and H. Voss, TMVA - Toolkit for Multivariate Data Analysis, PoS ACAT 040 (2007), arXiv:physics/0703039.

[102] D.K. MacKay. *Information Theory, Inference, and Learning Algorithms.* University of Cambridge Press, September 2003.

[103] G. Cybenko., Approximations by superpositions of sigmoidal functions, Mathematics of Control, Signals, and Systems, **2** (4), 303-314 (1989).

[104] R. Toner, “MLP LEM - SuperFOMs for alternate versions of the neural net.” MINOS-doc-7665-v1, 2010.

[105] T. Yang, C. Andreopoulos, H. Gallagher, K. Hoffmann, P. Kehayias, Eur.Phys.J. **C63**, 1-10 (2009).

[106] P. Allen et al., Nucl. Phys. B **181**, 385 (1981).

[107] A.A. Ivanilov et al., Yad. Fiz. **41**, 1520 (1985).

[108] H. Grassler et al., Nucl. Phys. B **223**, 269 (1983).

[109] D. Allasia et al., Z. Phys. C **24**, 119 (1984).

[110] W. Wittek et al., Z. Phys. C **40**, 231 (1988).

[111] D. Zieminska et al., Phys. Rev. D **27**, 47 (1983).

[112] C. Angelini et al., Phys. Lett. B **179**, 307 (1986).

[113] M. Derrick et al., Phys. Rev. D **17**, 1 (1978)

[114] A.M. Cooper-Sarkar, in *Neutrino '82*, Balatonfured, Hungary June 14-19, 1982.

[115] J. Boehm, H. Gallagher, T. Yang, Minos-doc-5392-v1, 2008.

[116] H. Gallagher, "Hadronization Model Uncertainties." Minos-doc-5443-v1, 2010.

[117] H. Gallagher, personal communication, April 13, 2010.

[118] J. Coelho, "Near Detector Beam Decomposition for the 3rd nue analysis." Minos-doc-7986, 2011.

[119] J. Coelho, "ND Decomposition Plots." Minos-doc-8091, 2011.

[120] L. Whitehead, "Extrapolation and Fitting in the 3rd Nue Analysis." MINOS-doc-7999, 2011.

[121] J. Boehm, "Measuring Nue CC Selection Efficiency and Systematic Error using the Muon Removal with MC Electron Addition Process." MINOS-doc-5385, 2009.

[122] A. Schreckenberger, "MRE 3rd Analysis Note." MINOS-doc-7980, 2011.

[123] M. Betancourt, "Cosmic Muon Background Studies." MINOS-doc-7985, 2011.

[124] M. Strait, "RAF 2010 CC deliverables: Position/Blessing/Summary." MINOS-doc-6870, 2010.

[125] M. Betancourt et al., "Systematic uncertainties for the third nue box opening." MINOS-doc-7976, 2011.

[126] T. Vahle, “Updated Normalization Systematic.” MINOS-doc-6854, 2010.

[127] The MINOS Calibration Group, “2009 Calibraiton Position paper on Runs I-II-III.” MINOS-doc-6717, 2010.

[128] G. Pawloski, T. Yang, “Crosstalk Tuning.” MINOS-doc-5319, 2008.

[129] S. Dytman, H. Gallagher, “New Intranuke model for Neugen - hA.” MINOS-doc-3484, 2007.

[130] M. Kordosky, “A procedure to re-weight events to account for uncertainties in final state interactions.” MINOS-doc-3449, 2007.

[131] S. Cavanaugh, “Tau Background in the Nue Appearance Analysis using Firebird-SUN.” MINOS-doc-5271, 2008.

[132] Hagiwara et al., Pseudoscalar form factors in tau-neutrino nucleon scattering, Phys. Lett. B **591**, 113-118 (2004).

[133] H. Gallagher, “Tau Background Uncertainty.” MINOS-doc-5442, 2008.

[134] A. Himmel, “The NuMI Beam Simulation with Flugg.” MINOS-doc-6316, 2009.

[135] C. Amsler et al., PL B667, 1 (2008) and 2009 partial update for the 2010 edition (http://pdg.lbl.gov)

[136] TMinuit package, see documentation here: http://root.cern.ch/root/html/TMinuit.html

[137] R. Toner, “Sensitivity Numbers for NueFit2D Integrated Likelihood code.” MINOS-doc-7893, 2011.

[138] L. Whitehead, “Multibin Fit for Nue.” MINOS-doc-7296, 2010.

[139] M. Orchanian, “First Look at MultiBin Sensitivities for LEM and Binning Optimization.” MINOS-doc-7699, 2010.

[140] M. Orchanian, “Next Round of Multibin Sensitivities for LEM.” MINOS-doc-7746, 2010.

[141] R. Toner, “A Likelihood fit for a 3LEMx5E bin fit.” MINOS-doc-8198, 2011.

[142] A. Holin, “Muon Removal Background Estimation Method.” Minos-doc-4141-v7, 2008.

[143] S. Cavanaugh, “Update on Far MRCC Data.” Minos-doc-6596-v1, 2009.

[144] G. J. Feldman and R. D. Cousins, Phys. Rev. D 57, 3873 (1998).

[145] K. Abe et al., Indication of Electron Neutrino Appearance from an Accelerator-Produced Off-Axis Muon Neutrino Beam, Phys. Rev. Lett. 107, 041801 (2011).

[146] G. L. Fogli et al., Evidence of theta(13)>0 from global neutrino data analysis, arXiv:1106.6028v2.

[147] F. Ardellier et al., Double Chooz: A search for the neutrino mixing angle θ_{13}, hepex/0606025 (2006).

[148] X. Guo et al., A precision measurement of the neutrino mixing angle θ_{13} using reactor antineutrinos at Daya Bay, arXiv:hep-ex/0701029v1 (2006).

[149] J.K. Ahn, et al, RENO: An Experiment for Neutrino Oscillation Parameter θ_{13} Using Reactor Neutrinos at Yonggwang, arXiv:1003.1391v1 [hep-ex] (2010).

[150] D. S. Ayres et al., NOvA proposal to build a 30-kiloton off-axis detector to study neutrino oscillations in the Fermilab NuMI beamline, arXiv:hep-ex/0503053 (2004).

List of Figures

List of Tables