Dissertations, Theses, and Masters Projects   Theses, Dissertations, & Master Projects

2011

A measurement of neutrino oscillations with muon neutrinos in the MINOS experiment

Stephen James Coleman
College of William & Mary - Arts & Sciences

Follow this and additional works at: https://scholarworks.wm.edu/etd

Part of the Physics Commons

A MEASUREMENT OF NEUTRINO OSCILLATIONS WITH MUON NEUTRINOS IN THE MINOS EXPERIMENT

Stephen James Coleman

Richmond, Virginia

Master of Science, College of William and Mary, 2007
Bachelor of Arts, University of Colorado, Boulder, 2005

A Dissertation presented to the Graduate Faculty
of the College of William and Mary in Candidacy for the Degree of
Doctor of Philosophy

Department of Physics

The College of William and Mary
May 2011

APPROVAL PAGE

This Dissertation is submitted in partial fulfillment of
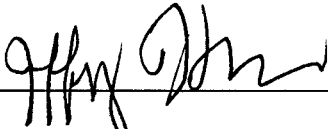the requirements for the degree of

Doctor of Philosophy

Stephen James Coleman

Approved by the Committee, January 2011

Committee Chair
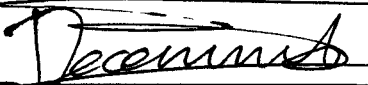Cornelia B. Talbot Term Distinguished Associate Professor Jeffrey Nelson, Physics

The College of William and Mary

Professor Todd Averett, Physics
The College of William & Mary

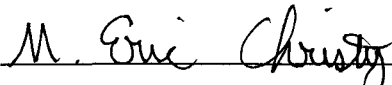Assistant Professor Wouter Deconinck,
Physics
The College of William & Mary

Assistant Professor Konstantinos Orginos,
Physics
The College of William & Mary

Assistant Professor M. Eric Christy, Physics
Hampton University

ABSTRACT PAGE

Experimental evidence has established that neutrino flavor states evolve over time. A neutrino of a particular flavor that travels some distance can be detected in a different neutrino flavor state. The Main Injector Neutrino Oscillation Search (MINOS) is a long-baseline experiment that is designed to study this phenomenon, called neutrino oscillations. MINOS is based at Fermilab near Chicago, IL, and consists of two detectors: the Near Detector located at Fermilab, and the Far Detector, which is located in an old iron mine in Soudan, MN. Both detectors are exposed to a beam of muon neutrinos from the NuMI beamline, and MINOS measures the fraction of muon neutrinos that disappear after traveling the 734 km between the two detectors. One can measure the atmospheric neutrino mass splitting and mixing angle by observing the energy-dependence of this muon neutrino disappearance. MINOS has made several prior measurements of these parameters.

Here I describe recently-developed techniques used to enhance our sensitivity to the oscillation parameters, and I present the results obtained when they are applied to a dataset that is twice as large as has been previously analyzed. We measure the mass splitting $\Delta m_{23}^2 = (2.32^{+0.12}_{-0.08}) \times 10^{-3}$ eV^2/c^4 and the mixing angle $\sin^2(2\theta_{32}) > 0.90$ at 90% C.L. These results comprise the world's best measurement of the atmospheric neutrino mass splitting. Alternative disappearance models are also tested. The neutrino decay hypothesis is disfavored at 7.2σ and the neutrino quantum decoherence hypothesis is disfavored at 9.0σ.

TABLE OF CONTENTS

Page

LIST OF FIGURES

LIST OF TABLES

Acknowledgements

I don't think anyone has an accurate expectation of graduate school when they first start it. I certainly could not have made it this far without the support from friends and family that I am lucky to have.

I would like to thank my parents for their support, starting way back with my first telescope (and then second, and then third). I know it was a gift from Uncle Scott, but you let me put it together by myself when I was five. Knowing that "science" was something that I could put my hands on certainly colored my worldview, then and now.

Thank you to my friends at W&M, for eating lunch with me at the same four places for five years, and for putting up with my jokes the entire time. And for getting me through homework sets. And for throwing baseballs outside when the weather's nice. Even prisoners get yard time.

Also, thanks to Aunt Cathy and Uncle Ted, for being in Williamsburg and for having spare vehicles.

Thank you to my fellow Young Minoans, especially at Oxford and Caltech, for being in convenient time zones for tech support. You know who you are. One of these days we'll be the ones in charge of experiments, and then we can finally use all the crazy acronyms we came up with over beers.

I am especially thankful to my advisor Jeff Nelson, for his guidance, patience, and friendship, and for being a great hockey defenseman. I sincerely hope to work with him again sometime. I wish him many years of success and future graduate students who are better at this than me.

Most importantly, I would like to thank my beautiful wife Carla, for being there.

There are known knowns; there are things we know we know. We also know there are known unknowns; that is to say we know there are some things we do not know. But there are also unknown unknowns – the ones we don't know we don't know.

– Donald Rumsfeld

CHAPTER 1

Introduction

1.1 Neutrinos

In the early part of the twentieth century, the theory describing nuclear β-decay was based on experimental observations of two particles in the final state. The β-decay model described the two-body decay of a neutron into an electron and a proton after about 15 minutes:

$$n^0 \rightarrow e^- + p^+ \tag{1.1}$$

The electron energy spectrum from a two-body decay should be mono-energetic, yet experiments showed that electrons emitted from β-decay had a continuous energy spectrum. Modifying existing theory to agree with experimental evidence presented an uncomfortable choice - either abandon the postulate of the conservation of momentum and energy or invent a third, undetected, particle produced in the decay to remove some of the energy.

The existence of neutrinos was famously predicted in 1930 by Wolfgang Pauli [1]. The new particle Pauli proposed (originally called a "neutron") was electrically neutral, spin-1/2, and had a mass similar to that of the electron. It was emitted from a nucleus along with a proton and electron, making β-decay a three-body process instead of a two-body process, and thereby explaining the continuous energy spectrum seen in nuclear β-decay.

The properties of the newly-proposed particle made it impossible to detect, a quality that Pauli considered unsettling. With this new and invisible particle (denoted by the symbol ν), energy and momentum conservation could be preserved.

$$n^0 \rightarrow e^- + p^+ + \nu \tag{1.2}$$

1.1.1 Enter Enrico Fermi

The particle proposed by Pauli was incorporated into existing theory by Enrico Fermi in 1934 [2, 3], and the particle was renamed the "neutrino," Italian for "little neutral one," to distinguish it from the neutral nucleon discovered by James Chadwick in 1932 [4]. The neutron is both neutral and spin-1/2, but is strong-interacting and too massive to be the particle that Pauli had proposed. Fermi calculated the matrix element for a single-point vertex between a neutron, proton, electron, and neutrino. The matrix element is

$$\mathcal{M} = \frac{G_F}{(\hbar c)^3} \bar{u}_n \gamma^\mu u_p \bar{u}_{\nu_e} \gamma_\mu u_e \tag{1.3}$$

where G_F is the effective coupling constant, u_x are spinors, and γ_μ are the Dirac matrices. Fermi also devised a way to determine the neutrino mass from the endpoint of the electron energy distribution measured from β-decay. Comparing his calculations to the β-decay experimental data available at the time, he concluded that the neutrino mass must either be zero or "in any case, very small in comparison to the mass of the electron."

1.2 The Early Years

Pauli initially expressed regret about his introduction of the particle that would come to be known as the neutrino. "I have done a terrible thing," he wrote in 1930, "I have postulated a particle which cannot be detected." It would be another 26 years, after the invention of the fission reactor, before experimental evidence for the existence of the neutrino would be published. Since then, extensive data have been collected about neutrino properties and their interactions.

As with all developments in physics, progress is only made when theory and experiment work in concert. Neutrino physics is no different, with experimental discoveries at times driving theory, and other times vice versa. Here I describe some of the major advancements in the field since Pauli and Fermi laid the groundwork.

1.2.1 Early experiments

The first experimental evidence for the existence of neutrinos came with the Savannah River experiment performed by Frederick Reines and Clyde Cowan, the results of which were published in 1956 [5]. After a fairly crazy idea, Project Poltergeist, was

abandoned (which involved dropping a large detector down a long shaft, in proximity to a detonating nuclear bomb), an inconclusive attempt to detect neutrinos was made with the Hanford Experiment [6]. After moving to Savannah River, Reines and Cowan achieved a detection with a convincing signal to background ratio of 4/1. This experiment detected anti-neutrinos emitted from fission in a nearby nuclear reactor. The detector was heavily shielded to reduce the number of background neutrons and photons within the detector. The neutrino interacted within the detector volume by inverse β-decay:

$$\bar{\nu} + p \rightarrow n + e^+ \tag{1.4}$$

The neutrino interaction signal was tagged by the coincidence detection of a prompt positron and a photon due to delayed neutron capture. Reines and Cowan made a series of attempts to detect these neutrinos, and collected data when the reactor was on and off to demonstrate that the neutrinos were indeed coming from the nuclear reactor. Reines won the Nobel Prize for Physics for this discovery in 1995.

Leon Lederman, Melvin Schwartz, and Jack Steinberger opened a new door in experimental neutrino physics by utilizing a particle accelerator to study neutrinos. Their experiment used a proton beam that was directed to strike a fixed target and produce pions ($\pi^\pm$), which then decayed into muons (μ) and neutrinos. The muons were stopped by a large amount of absorbing material and the surviving beam of neutrinos was aimed at spark chambers. The neutrinos interacted with matter, and flashes in the chambers indicated tracks of outgoing paticles, which were recorded with photographic plates [7]. Previously, neutrinos from β-decay had been observed with electrons leaving the inter-

action vertices. The neutrinos in this experiment were detected with only muons leaving the interaction vertices. The collaborators went to great lengths to prove that these interactions involved neutrinos produced in the pion beam and not neutrons or cosmic rays.

The logical conclusion was that there were actually two different types of neutrinos, ν_e and ν_μ, partnered with the known charged leptons, e and μ. The type of neutrino interacting, then, could be determined by the flavor of the lepton leaving the vertex. Lederman, Schwartz, and Steinberger shared the Nobel Prize for Physics for this discovery in 1988.

This experiment also observed six "showers" with no obvious outgoing lepton, and which they confirmed were not electron showers. They left the explanation of these showers to future experiments.

1.2.2   GSW Theory

The weak interaction model that Fermi proposed was of vector-vector form. The discovery of parity violation in the 1950's [8, 9] hinted that the vector-vector weak interaction was not correct. An equal axial component ($\gamma_\mu \gamma_5$) was needed in the weak matrix element to violate parity. This made the weak interaction of "vector-axial" form, or "V-A." The full theory for neutrino interactions came in the early 1960's from Glashow, Salam, and Weinberg (GSW) with the prediction of as-yet undiscovered new bosons mediating the weak force [10–12]. The full theory contains quarks and leptons, where the neutrinos interact only weakly. The new bosons in this model, the $W^\pm$ and the Z^0, coupled to neutrinos. Neutrino interactions tagged with an outgoing lepton are charged-current interactions, mediated by the $W^\pm$.

The GSW model predicted that the $W^\pm$ and Z^0 are massive, and therefore short-ranged, bosons [13]. It also predicted a new kind of interaction that had not yet been identified experimentally, the neutral-current interaction, mediated by the Z^0, which did not feature an outgoing lepton.

The GSW model achieved another milestone, unifiying the electromagnetic and weak forces into a single electroweak force. The GSW model with three generations of matter, combined with a model for the strong force (quantum chromodynamics, or QCD), constitute the Standard Model of particle physics.

1.3 Neutrinos in the Standard Model

In the Standard Model (SM), neutrinos are massless and interact only weakly. The SM Lagrangian describes two types of interactions for neutrinos. For each of the lepton species the neutrino couples to W bosons in Charged-Current interactions:

$$L_{CC} = -\frac{g}{\sqrt{2}} \sum_j \left(\bar{e}_{jL} \gamma^\mu W_\mu^- \nu_{jL} + e_{jL} \gamma^\mu W_\mu^+ \bar{\nu}_{jL} \right) \tag{1.5}$$

and to the Z^0 boson in Neutral-Current interactions:

$$L_{NC} = -\frac{g}{2 \cos \theta_W} \sum_\alpha \bar{\nu}_{jL} \gamma^\mu Z_\mu^0 \nu_{jL} \tag{1.6}$$

Weak interactions in the SM maximally violate parity (P) and charge-conjugation (C) but conserve CP. Maximal parity violation means that only left-handed neutrinos (or

(a) ν_μCC interaction   (b) ν_μNC interaction
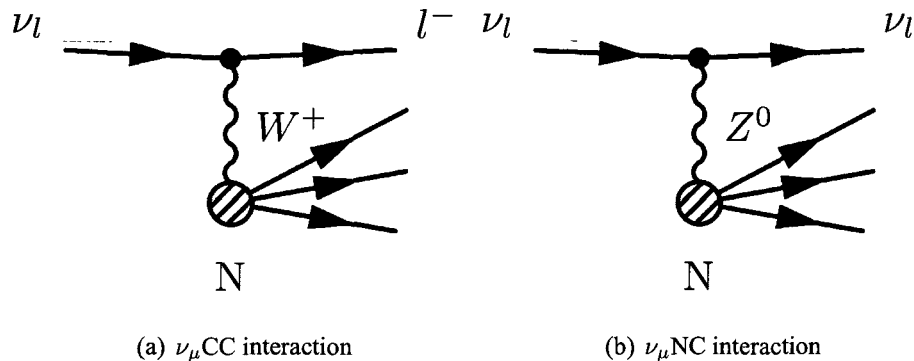
FIG. 1.1: Examples of Charged-Current and Neutral-Current interactions between neutrinos and nuclei. The $W^\pm$ vertex with the nucleus can either be quasi-elastic with the entire nucleus, resonance-producing with a nucleon, or deep-inelastic scattering off of a nucleon's down quark (up quark for an antineutrino).

left-handed anti-neutrinos) interact with charged fermions. Right-handed particles travel with their spin aligned in the direction of their momentum, and left-handed particles travel with their spin anti-aligned with the direction of their momentum.

1.3.1 Weak Charged-Currents

The Savannah River experiment, which first confirmed the existence of neutrinos, and all of the subsequent experiments leading up to the GSW model observed neutrinos via weak charged-current interactions. In these interactions, a neutrino exchanges a W^+ with a target. The W^+ has an electric charge $q = +1$ and the neutrino has $q = 0$, so the third participant at the νW^+ vertex must be a lepton of charge $q = -1$ in order to conserve charge and lepton number. This is shown in Figure 1.1(a).

The W^+ itself will transfer some momentum and charge to the target. If the transferred momentum is small, the interaction is quasi-elastic (QE) and a neutron in the target nucleus will convert into a proton with little recoil momentum. At higher momentum

transfer (q^2), the target nucleon (a proton or a neutron) will be converted into a Δ resonance (RES), which will decay into a nucleon and a π. The π may induce a hadronic shower, a cascade of hadrons induced by strong interactions with nuclei. At even higher q^2, the W^+ will interact directly with a d-quark in the target nucleus and will proceed to break up the nucleus with a large number of final-state particles in a deep-inelastic scattering interaction (DIS).

1.3.2 Weak Neutral-Currents

Neutrinos can also exchange a Z^0 boson with matter, looking similar to the charged-current case without the outgoing lepton. The Z^0 has $q = 0$, so the third participant at the νZ^0 vertex is another neutrino, as shown in Figure 1.1(b). These neutral-current (NC) interactions were first identified with the Gargamelle experiment at CERN in 1973 [14]. Gargamelle was a bubble chamber which held $12\,\mathrm{m}^3$ of freon, which was placed in a neutrino beam created from the CERN proton synchroton. Gargamelle first observed neutral-current interactions in the quasi-elastic regime (with little momentum transferred to the target nucleus), with $\nu_\mu + e^- \rightarrow \nu_\mu + e^-$. Gargamelle also ran with an anti-neutrino beam, observing $\bar{\nu}_\mu + e^- \rightarrow \bar{\nu}_\mu + e^-$, and measured the double-ratio of cross sections to be $(CC/NC)_{\bar{\nu}_\mu}/(CC/NC)_{\nu_\mu} \approx 2$.

The Lederman, Schwartz, and Steinberger experiment which first observed the ν_μ actually observed six NC events, but they did not identify them as such. They placed their apparatus in an electron beam to ensure the observed showers were not consistent with electron showers (which would invalidate their results by indicating that $\nu_\mu = \nu_e$). Once

they were sure the showers were not electromagnetic in origin, they essentially gave up on understanding these hadronic showers from NC interactions.

1.4 Modern experiments involving neutrinos

Fermi's theory of β-decay was the starting point for experimental measurements of neutrinos. He observed from the β-decay spectrum that the neutrino mass had to be much smaller than the electron mass, at $511\,\text{eV}$, contrary to the initial prediction of Pauli. To tell the story of neutrino measurements in the intervening time requires more information about the modern fundamental particle zoo. The Standard Model includes three generations of quarks, the u and d, the s and c, and t and b. Quarks are never observed singly, but are bound in groups, called hadrons. Quark-antiquark pairs are called mesons, while groups of three quarks are called baryons, like protons and neutrons (uud and udd, respectively). Likewise there are three generations of leptons, which include the charged leptons (with electric charge -1) and neutral leptons, or neutrinos.

The discovery of the τ lepton in 1974 by Perl *et al.* [15] indicated that there were in fact three generations of matter. Perl shared the Nobel Prize with Reines in 1995. By 1995, all three generations of quarks were in place, and only the ν_τ was left to complete the stable of fermions.

Experiments at SLAC and CERN looked at Z^0 decays to determine the number of generations of neutrinos which have masses less than half that of the Z^0 itself. The width of the Z^0 decay peak revealed that the number of light neutrinos was $N_\nu = 2.984 \pm 0.008$ [16]. This brought neutrinos in line with the three generations of quarks and charged

leptons. The third neutrino, the ν_τ, was directly observed by the DONUT collaboration at Fermilab in 2000 [17].

1.4.1 Neutrino scattering

As neutrinos interact only weakly, it follows that the neutrino cross-section should be very small. As with nuclear physics, an investigation of the neutrino cross-section is most easily carried out with fixed-target scattering experiments. Many experiments have carried out measurements of neutrino and anti-neutrino cross-sections, along a wide range of energies. Neutrinos are a unique probe for measuring nuclear structure, since they only interact weakly. Charged-current neutrino interaction are separated into three classifications. In quasi-elastic interactions (QE), the neutrino exchanges a W with a proton or a neutron, and the only two outgoing products are the neutrino's corresponding charged lepton and the recoil neutron or proton. If enough momentum is transferred to the struck nucleon, a resonance may be produced (RES), which will result in an extra pion in the final state, along with the charged lepton and neutron or proton. Finally, the neutrino may exchange a W with the struck nucleon's constituent quarks in a deep-inelastic scattering event. This produces a hadronic shower in the final state, along with the charged lepton.

The current knowledge of ν_μ cross-sections, in quasi-elastic, resonance production (with a single outgoing pion), and deep-inelastic scattering is shown compared to theoretical predictions in Figure 1.2. Identifying incoming neutrino energies in the few-GeV region are tricky, since neutrinos can interact through any one of these processes, with

varying amounts of particles below detection threshold, depending on the particular type

of detector being utilized. The relative cross-sections between ν_μ and $\overline{\nu}_\mu$ are shown in

Figure 1.3. Cross-section measurements obtained with the MINOS experiment are shown

in Figure 1.4.
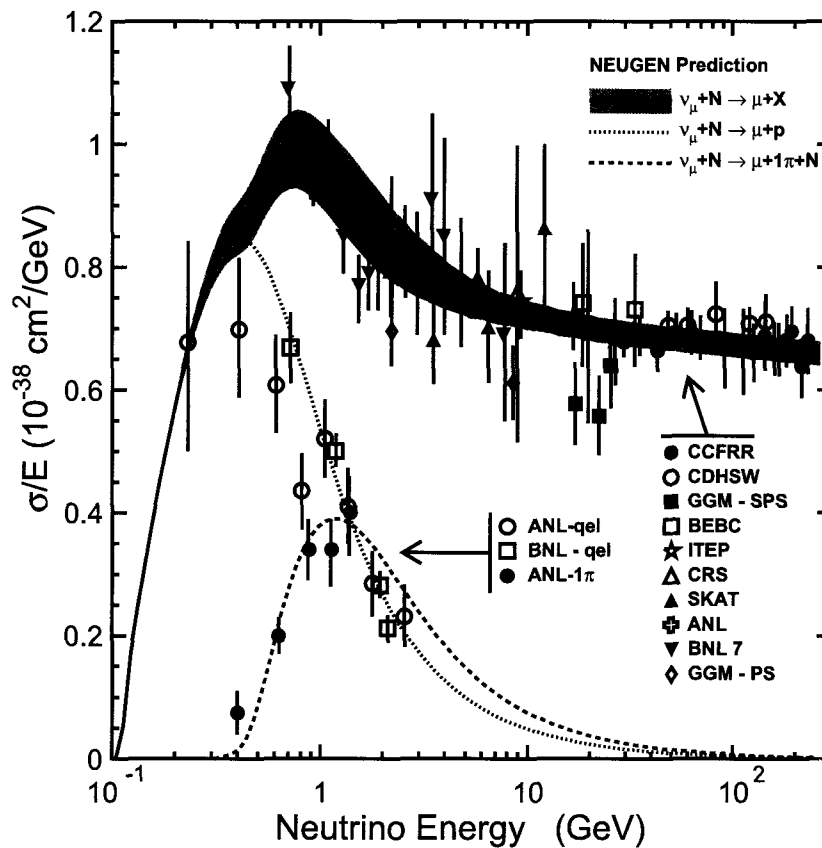
FIG. 1.2: The total cross-section for ν_μCC interactions per nucleon on an isoscalar target, as output from NEUGEN [18]. The solid line is the total cross-section and the green band is the total model uncertainty. The dotted line is the quasi-elastic component, and the dashed line is the resonance component of the cross-section [19].
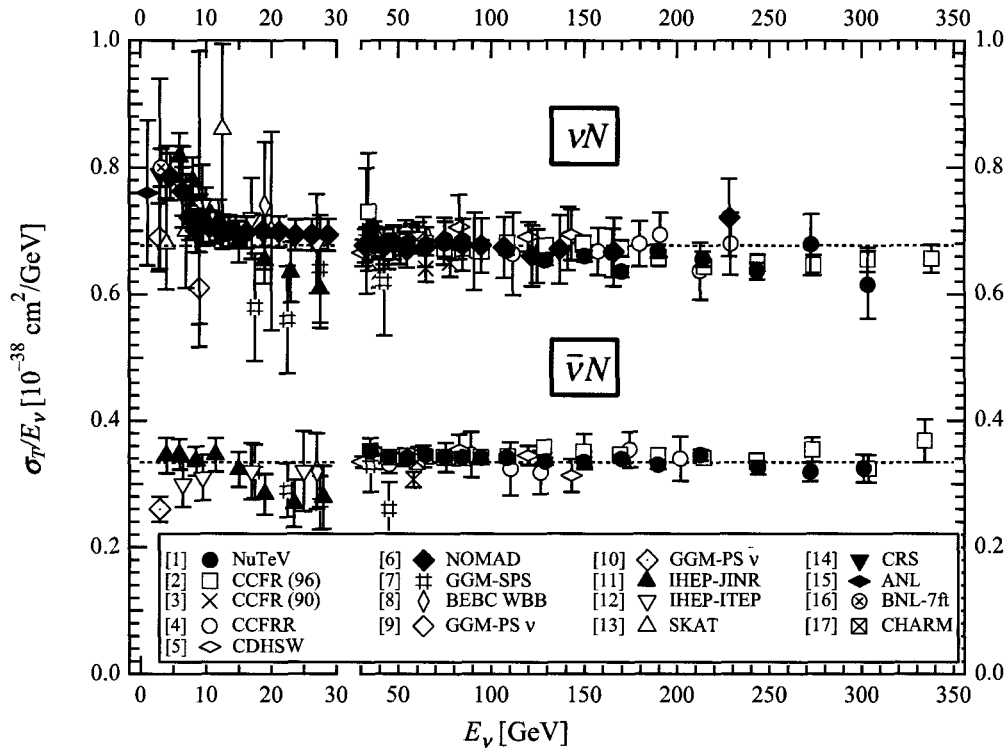
FIG. 1.3: σ_T/E_ν for the muon neutrino and anti-neutrino charged-current total cross section as a function of neutrino energy. The error bars include both statistical and systematic errors. The straight lines are the isoscalar-corrected total cross-section values averaged over 30-200 GeV. Note the change in the energy scale at 30 GeV. A full list of references is shown in [20].

1.4.2 The Solar Neutrino Problem

Ray Davis is renowned for his determination in measuring the flux of neutrinos coming from decays of ^{8}B in the Sun. Davis operated the Homestake experiment located in the Homestake gold mine in South Dakota. His tank of 0.6 kilo-ton of Chlorine-rich dry-cleaning fluid was located 2300 ft underground to minimize the incidence of cosmic ray-induced background events. Electron neutrinos (ν_e) from the Sun interacted within
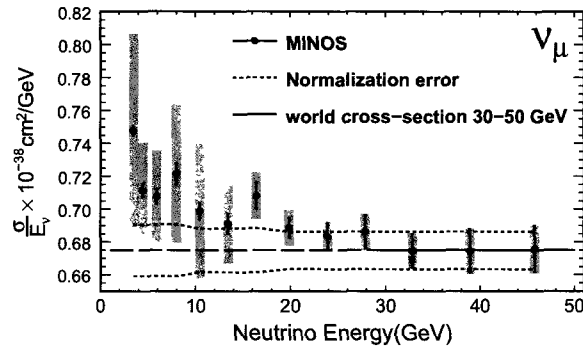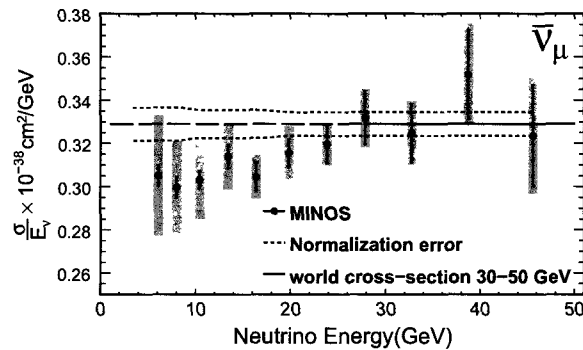
FIG. 1.4: Muon neutrino and muon antineutrino inclusive cross-sections (QE+RES+DIS) as measured by the MINOS experiment. For comparison, the world average is shown, along with the size of a 1.5% normalization systematic error on the MINOS result [21].

the tank and produced an argon isotope.

$$\nu + {}^{37}\text{Cl} \rightarrow {}^{37}\text{Ar} + e^{-} \tag{1.7}$$

Every few weeks, Davis would bubble helium through the tank to collect all of the argon isotopes produced in neutrino interactions. Counting the number of argon isotopes decaying gave the number of neutrinos that had interacted within the tank. The experiment required an incredible amount of patience, as the interaction rate was a paltry

0.4 interactions/day. This experiment published its first results in 1964 [22] and ceased operations in 1994 [23].

The number of neutrino interactions measured with this apparatus appeared to be roughly 1/3 the number that was expected from the solar models of the time and the known neutrino cross-sections, as calculated by astrophysicist John Bahcall, *et al.* This deficit of electron neutrinos was later confirmed by numerous experiments [24], and came to be known as the solar neutrino problem.

For discovery of solar neutrinos and the solar neutrino problem, Davis shared the Nobel Prize in Physics in 2002.

1.4.3 The atmospheric neutrino anomaly

Cosmic rays, highly-energetic particles of cosmic origin, interact in the upper atmosphere and produce a cascade of mesons which eventually decay into ν_μ, $\bar{\nu}_\mu$, and ν_e.

The cosmic ray neutrino flux was an important background to understand for proton decay experiments, which have very small signal-to-noise ratios. There were several proton decay experiments running in the 1980's which looked at the cosmic ray neutrino flux. In 1988, a proton decay experiment called Kamiokande, a water Cerenkov detector in Kamioka, Japan, published their measurement of ν_e and ν_μ fluxes. Kamiokande observed a $(56 \pm 7)\%$ deficit of ν_μ relative to ν_e [25].

Cosmic ray neutrinos are produced from mesons just like the neutrinos produced in

proton beams.

$$\pi^+ \rightarrow \mu^+ + \nu_\mu \tag{1.8}$$

$$\mu^+ \rightarrow e^+ + \bar{\nu}_\mu + \nu_e \tag{1.9}$$

and likewise for π^-. The flux ratio one would expect is $(\nu_\mu + \bar{\nu}_\mu)/(\nu_e + \bar{\nu}_e) \approx 2$, regardless of the actual pion flux. By observing this flux with many different experiments, the error on this ratio is only 2% below 10 GeV [25]. Cosmic rays provide a better test of neutrino disappearance, since the ratio of ν_μ to ν_e fluxes is self-calibrating, and not dependent on complicated solar models.

The early 1990's produced a flurry of conflicting results. The Kamiokande deficit was not corroborated by iron calorimeter experiments like Frejus [26] and NUSEX [27]. Another water Cerenkov experiment, IMB, observed a 2.6σ deficit of ν_μ [28]. It took ν_μ deficit measurements from MACRO [29] and Soudan 2 [30] before the atmospheric neutrino anomoly was widely believed to be anything but an undiscovered problem with water Cerenkov detectors. Masatoshi Koshiba, from the Kamiokande experiment, shared the Nobel Prize with Ray Davis in 2002.

1.5 Neutrino oscillations

The most successful hypothesis for neutrino disappearance is a mechanism called neutrino oscillations in a formalism proposed by Pontecorvo [31] and Maki, Nakagawa, and Sakata [32]. In this model, neutrinos are quantum mechanical wave packets with a

unique and well-defined mass. These are the neutrino mass eigenstates $|\nu_1\rangle$, $|\nu_2\rangle$, and $|\nu_3\rangle$.
The neutrino flavor states, the states which couple to the W and Z, are not mass eigenstates. The flavor states, labeled $|\nu_e\rangle$, $|\nu_\mu\rangle$, and $|\nu_\tau\rangle$, are related to the mass eigenstates by a unitary rotation matrix U

$$|\nu_\alpha\rangle = \sum_i U_{\alpha i}^* |\nu_i\rangle \tag{1.10}$$

As the neutrino propagates, its wave function evolves in space and time. Thus the time evolution operator acts on the state

$$|\nu_\alpha(x,t)\rangle = \sum_i U_{\alpha i}^* |\nu_i(x,t)\rangle \tag{1.11}$$

The probability of detecting a neutrino interacting as flavor β is

$$P(\nu_\alpha \rightarrow \nu_\beta) = |\langle \nu_\beta | \nu_\alpha(x,t)\rangle|^2 \tag{1.12}$$

$$= | \left(\sum_j \langle \nu_j | U_{\beta j} \right) \sum_i U_{\alpha i}^* |\nu_i(x,t)\rangle |^2 \tag{1.13}$$

$$= | \sum_i \sum_j U_{\beta j} U_{\alpha i}^* \langle \nu_j | \nu_i(x,t)\rangle |^2 \tag{1.14}$$

Before we square the right side of the equation, we must determine the effect of time-evolution on the state $|\nu_i\rangle$. Translating the state from position x_0 to x, we write the state as

$$|\nu_i(x,t)\rangle = e^{ip_i(x-x_0)}|\nu_i(t)\rangle \tag{1.15}$$

and then apply the time-evolution operator

$$|\nu_i(x,t)\rangle = e^{\imath p_i(x-x_0)}e^{-\imath \hat{H}(t-t_0)}|\nu_i\rangle \tag{1.16}$$

$$= e^{\imath p_i(x-x_0)}e^{-\imath E_i(t-t_0)}|\nu_i\rangle \tag{1.17}$$

where we have used the fact that $|\nu_i\rangle$ is an energy eigenstate and natural units ($\hbar = c = 1$).

Experiments have shown that the neutrino mass scale is very small (a fact that will be

discussed briefly in Section 1.8.3), so they are highly relativistic, and we can approximate

$(x - x_0) = c(t - t_0) = L$, the total distance the neutrino travels before being detected.

Since we know the neutrino mass is small, we can safely say that $m_i \ll p_i$, and using

the energy-momentum relationship in Special Relativity we can approximate

$$E_i^2 = p_i^2 c^2 + m_i^2 c^4 \tag{1.18}$$

$$E_i \approx p_i + \frac{m_i^2}{2p_i} \tag{1.19}$$

$$\approx p_i + \frac{m_i^2}{2E} \tag{1.20}$$

where E is the energy of the neutrino at production, which is common to all initial mass

eigenstates. We now have a wavefunction we can insert into Equation 1.14.

$$|\nu_i(x,t)\rangle = e^{\imath L(p_i - (p_i + \frac{m_i^2}{2E}))}|\nu_i\rangle \tag{1.21}$$

$$= e^{-\imath \frac{m_i^2 L}{2E}}|\nu_i\rangle \tag{1.22}$$

The full transition probability is then

$$P(\alpha \to \beta) = |\sum_i \sum_j U_{\beta j} e^{i\frac{m_j^2 L}{2E}} U_{\alpha i}^* e^{-i\frac{m_i^2 L}{2E}} \langle \nu_j | \nu_i \rangle |^2 \qquad (1.23)$$

$$P(\alpha \to \beta) = \delta_{\alpha\beta} - 4 \sum_{i<j}^{n} \mathrm{Re}[U_{\alpha i} U_{\beta i}^* U_{\alpha j}^* U_{\beta j}] \sin^2 \left(\frac{(m_i^2 - m_j^2)L}{4E} \right) \qquad (1.24)$$

$$+ 2 \sum_{i<j}^{n} \mathrm{Im}[U_{\alpha i} U_{\beta i}^* U_{\alpha j}^* U_{\beta j}] \sin \left(\frac{(m_i^2 - m_j^2)L}{2E} \right) \qquad (1.25)$$

The matrix U is the PMNS matrix, named after Pontecorvo, Maki, Nakagawa, and Sakata. In the case of three neutrino flavors,

$$\begin{pmatrix} \nu_e \\ \nu_\mu \\ \nu_\tau \end{pmatrix} = U_{PMNS} \begin{pmatrix} \nu_1 \\ \nu_2 \\ \nu_3 \end{pmatrix} \qquad (1.26)$$

Due to the unitarity of U, the values of the matrix elements U_{ij} are determined by four independent parameters, three mixing angles and one phase:

$$U_{PMNS} = \begin{pmatrix} 1 & 0 & 0 \\ 0 & c_{23} & s_{23} \\ 0 & -s_{23} & c_{23} \end{pmatrix} \begin{pmatrix} c_{13} & 0 & s_{13}e^{-i\delta_{CP}} \\ 0 & 1 & 0 \\ -s_{13}e^{-i\delta_{CP}} & 0 & c_{13} \end{pmatrix} \begin{pmatrix} c_{12} & s_{12} & 0 \\ -s_{12} & c_{12} & 0 \\ 0 & 0 & 1 \end{pmatrix}$$

$$(1.27)$$

where $c_{ij} = \cos(\theta_{ij})$ and $s_{ij} = \sin(\theta_{ij})$. (N.B. there are additionally Majorana phases

which are ignored here, but will be discussed later). The phase δ_{CP} is known as the Dirac phase. If CP-invariance holds ($\delta_{CP} = 0$), then U is real in addition to being unitary ($U^* = U$). Written out completely, the full PMNS matrix is

$$U_{PMNS} = \begin{pmatrix} c_{12}c_{13} & s_{12}c_{13} & s_{13}e^{-i\delta_{CP}} \\ -s_{12}c_{23} - c_{12}s_{13}s_{23}e^{i\delta_{CP}} & c_{12}c_{23} - s_{12}s_{13}s_{23}e^{i\delta_{CP}} & c_{13}s_{23} \\ s_{12}s_{23} - c_{12}s_{13}c_{23}e^{i\delta_{CP}} & -c_{12}s_{23} - s_{12}s_{13}c_{23}e^{i\delta_{CP}} & c_{13}c_{23} \end{pmatrix}$$
$$(1.28)$$

With four free parameters in U_{PMNS} and two mass splittings Δm_{32}^2 and Δm_{12}^2, there are six parameters in total describing full three-flavor neutrino oscillations in vacuum.

Many experiments measure the survival probability for a particular neutrino flavor, in which case the neutrino oscillation signal manifests itself as a deficit relative to an unoscillated flux prediction. The full three-flavor survival probability relevant for MINOS is $P(\nu_\mu \to \nu_\mu)$, which is

$$P(\nu_\mu \to \nu_\mu) = 1 - 4\sum_{i<j}^{3} \mathrm{Re}[|U_{\mu i}|^2 |U_{\mu j}|^2] \sin^2\left(\frac{(m_i^2 - m_j^2)L}{4E}\right)$$
$$+ 2\sum_{i<j}^{3} \mathrm{Im}[|U_{\mu i}|^2 |U_{\mu j}|^2] \sin\left(\frac{(m_i^2 - m_j^2)L}{2E}\right)$$

where the imaginary term disappears without imposing the requirement that U_{PMNS} is real. Defining $\Delta m_{ij}^2 = m_i^2 - m_j^2$, the survival probability is then

$$P(\nu_\mu \to \nu_\mu) = 1 - 4\sum_{i<j}^{3} |U_{\mu i}|^2 |U_{\mu j}|^2 \sin^2\left(\frac{\Delta m_{ij}^2 L}{4E}\right) \qquad (1.29)$$

(a) Normal Hierarchy   (b) Inverted Hierarchy
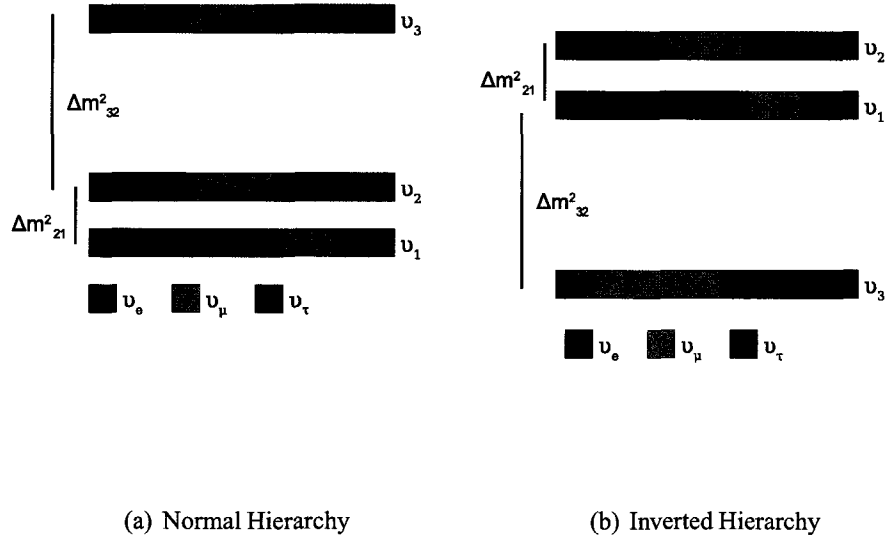
FIG. 1.5: The possible neutrino mass state spectra and the mass state flavor components, depending on whether or not the mass splittings are configured in the "Normal" or "Inverted" mass hierarchy. Mass increases from bottom to top.

1.5.1 The two-flavor approximation

One can rewrite the neutrino oscillation probability in terms of effective mass splittings, Δm_{eff}^2, instead of Δm_{32}^2 and Δm_{12}^2. This model is convenient since $\Delta m_{32}^2 \gg \Delta m_{12}^2$, as the full three-flavor oscillation model decouples into two two-flavor oscillation modes. The Δm_{eff}^2 for a disappearing flavor eigenstate is the weighted average of the two true mass splittings, weighted by that disappearing flavor eigenstate's fractional component in the other mass eigenstates.

$$\Delta m_{eff}^2|_\alpha \equiv \frac{|U_{\alpha1}|^2 \Delta m_{31}^2 + |U_{\alpha2}|^2 \Delta m_{32}^2}{|U_{\alpha1}|^2 + |U_{\alpha2}|^2} \tag{1.30}$$

so for a disappearing ν_μ, the mass splitting that is being measured is

$$\Delta m_{eff}^2|_\mu = \sin^2\theta_{12}\Delta m_{31}^2 + \cos^2\theta_{12}\Delta m_{32}^2$$

$$- \cos\delta_{CP}\sin\theta_{13}\sin(2\theta_{12})\cot\theta_{23}\Delta m_{21}^2$$

$$= \sin^2\theta_{12}\Delta m_{31}^2 + \cos^2\theta_{12}\Delta m_{32}^2 \tag{1.31}$$

where the last term, $\mathcal{O}(\sin\theta_{13})$, has been ignored, because θ_{13} is known to be very small [33].

The two-flavor atmospheric oscillation, in terms of muon neutrino survival probability, is

$$P(\nu_\mu \to \nu_\mu) \simeq 1 - \sin^2(2\theta_{atm})\sin^2\left(1.27\Delta m_{atm}^2\frac{L\ \text{(km)}}{E\ \text{(GeV)}}\right) \tag{1.32}$$

The other two-flavor approximation mode is solar neutrino oscillations. In terms of electron neutrinos, the ν_e survival probability is

$$P(\nu_e \to \nu_e) \simeq 1 - \sin^2(2\theta_{sol})\sin^2\left(1.27\Delta m_{sol}^2\frac{L\ \text{(km)}}{E\ \text{(GeV)}}\right) \tag{1.33}$$

The two-flavor approximation is convenient from an experimentalist's point of view, since it probes two fundamental constants with two controlable parameters. The ability of muon neutrinos to pass through large quantities of matter without interacting allows for long experimental baselines L, when measuring atmospheric neutrino oscillations, and neutrino beams with tunable energies E allow experiments to probe a large phase space of Δm_{atm}^2-$\sin(2\theta_{atm})$ values. For experiments with fixed L and measuring a wide range

Neutrino Oscillation Probability
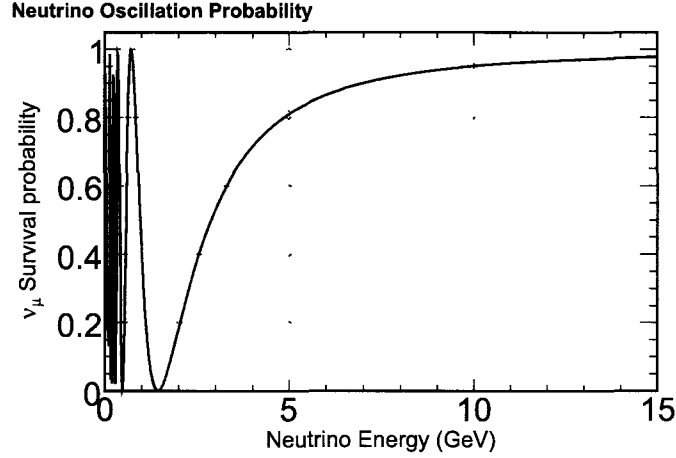
FIG. 1.6: An example for muon neutrino survival probability as a function of neutrino energy for two-flavor oscillations with baseline $L = 735\,\text{km}$, $\Delta m^2 = 2.43 \times 10^{-3}\,\text{eV}^2/\text{c}^4$, and $\sin^2(2\theta) = 1.0$. For these parameters, below 1 GeV the survival probability is in "fast oscillations", where the probability of detecting the ν_μ is 50%. The depth of the dip indicates the value of $\sin^2(2\theta)$ (1.0 in this case since, it is maximal), and the location of the dip in E scales linearly with Δm^2 if L is fixed.

of E, the ratio of an oscillated spectrum to an unoscillated spectrum gives an intuitive measurement of Δm^2_{atm} and $\sin(2\theta_{atm})$. The depth of the lowest point of the spectrum (called the oscillation dip) occurs where $1.27\Delta m^2_{atm}\frac{L}{E} = \pi/2$ and gives a measure of $\sin^2(2\theta_{atm})$. The location of the oscillation dip in E indicates the value for Δm^2_{atm} where $1.27\Delta m^2_{atm}\frac{L}{E} = \pi/2$. This is shown in Figure 1.6 for chosen values of $L, \Delta m^2_{atm}$, and $\sin^2(2\theta_{atm})$. Experiments measuring the solar oscillation parameters using neutrino fluxes from many nearby nuclear reactors, as will be discussed later, do not have a fixed baseline L, and instead treat L/E as their independent variable.

For the rest of this document, Δm^2 is meant to mean $\Delta m^2_{eff}|_\mu \approx \Delta m^2_{atm}$, though the actual difference between the two is below the precision of this experiment.

1.6 Matter Effects

The process for neutrino flavor oscillations described above holds when the neutrinos are propagating in a vacuum. In matter, neutrinos can coherently forward-scatter off of e, p, or n by exchanging any $W^{\pm}$ between the time of their creation and the time of their detection. This decouples the time-evolution from the wave function and essentially "resets the clock" on $|\nu_\alpha(t)\rangle$ to $|\nu_\alpha(0)\rangle$. The explanation of this effect is simplified by using the two-flavor assumption.

In normal matter, propagating ν_e and ν_μ will be affected differently by matter due to the difference in $\nu_e + e$ and $\nu_\mu + e$ scattering amplitudes. All neteutrino flavors states exchange Z^0 in NC interactions in the same way, and so the effect of coherent NC interactions in matter are the same for all neutrino flavors. Only the ν_e can interact with matter electrons coherently via a CC interaction. This effect was first described by Mikheyev, Smirnov, and Wolfenstein [34, 35] and is known as the MSW effect. A full derivation exists in many places (in particular, [24]), and only the basics are presented here.

The possibility of coherent interactions of neutrinos in matter introduces a new term to the Hamiltonian:

$$H = H_0 + H_{int} \tag{1.34}$$

The mass eigenstates (ν_1, ν_2) which were eigenstates of the vacuum Hamiltonian H_0 are not necessarily eigenstates once H_{int} is introduced. Translating this into the observable flavor basis, the ν_e term picks up an additional potential

$$\langle \nu_e | H_{int} | \nu_e \rangle = V - \sqrt{2} G_F N_e \tag{1.35}$$

where V is the matter potential associated with coherent interactions with p and n (so $\langle \nu_\mu | H_{int} | \nu_\mu \rangle = V$), G_F is the Fermi constant and N_e is the number density of electrons in the matter being traversed by neutrinos.

The net Hamiltonian H can be rotated back to the mass basis, where the mass eigenstates are no longer ν_1 and ν_2. The new mass eigenstates can be written in terms of the observable flavor states:

$$
\begin{pmatrix} \nu_1^m \\ \nu_2^m \end{pmatrix} = \begin{pmatrix} \cos\theta_m & -\sin\theta_m \\ \sin\theta_m & \cos\theta_m \end{pmatrix} \begin{pmatrix} \nu_e \\ \nu_\mu \end{pmatrix} \tag{1.36}
$$

Thus the two-flavor oscillation probability in Section 1.5.1 are modified, with

$$
\sin^2 2\theta_m = \frac{\sin^2 2\theta_{12}}{f_{MSW}} \tag{1.37}
$$

$$
\Delta m_m^2 = \Delta m_{12}^2 f_{MSW} \tag{1.38}
$$

and

$$
f_{MSW} = \sqrt{\left(\frac{2\sqrt{2} G_F N_e E}{\Delta m_{12}^2} - \cos(2\theta_{12}) \right) + \sin^2(2\theta_{12})} \tag{1.39}
$$

Thus the effect of matter on neutrinos is to alter the flavor composition of the mass eigenstates in an energy-dependent way.

Note that when

$$
N_e = \frac{\Delta m^2 \cos(2\theta_{12})}{2\sqrt{2} G_F E} \tag{1.40}
$$

then oscillations are maximal, and a resonance is produced between ν_e and ν_μ. The max-

imal oscillations produce "fast oscillations" for neutrinos with certain energies.

For antineutrinos, the potential attributed to $\overline{\nu_e} + e$ is different:

$$\langle \overline{\nu_e}|H_{int}|\overline{\nu_e}\rangle = V + \sqrt{2}G_F N_e \tag{1.41}$$

This implies that the neutrino and antineutrino mass eigenstates are different in matter. Measuring this difference is one way to solve the problem of the mass hierarchy and to measure the CP-violating phase δ [33].

1.7 Experimental Evidence for Oscillations

1.7.1 Solar neutrinos and the solution to the solar neutrino problem

The nuclear fusion process that powers the Sun produces a large flux of low-energy electron neutrinos. The two main fusion processes, the $p - p$ chain and the CNO cycle, fuse four protons into ^{4}He and create two neutrinos in the process:

$$4p \rightarrow^4 \text{He} + 2e^+ + 2\nu_e + \gamma \tag{1.42}$$

Other neutrinos come from the decay of semi-stable by-products of the fusion process within the sun. The isotope ^8B, for example, β-decays and produces a neutrino flux with a wide energy spectrum. Other decays produce mono-energetic fluxes of neutrinos from 2-body decays, as shown in Figure 1.7.

The ν_e produced in the fusion reaction at the core of the Sun must travel through the
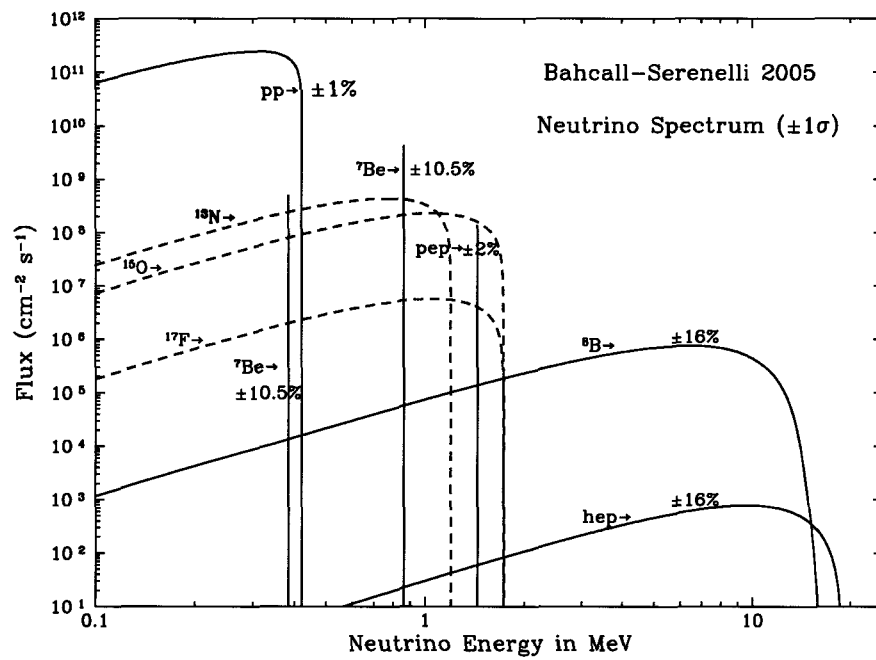
FIG. 1.7: The solar neutrino spectrum and its components, as predicted by Bachall et al.. Solar neutrinos are useful for study because there is a wideband flux from the $p - p$ chain, as well as monoenergetic peaks from specific β-decays within the sun [20].

radius of the Sun before being emitted. The Sun is dense with electrons, so the MSW effect described above can significantly alter the ν_e signal. The MSW resonance effect is energy dependent, and $p-p$ and ^{7}Be solar neutrinos are below the threshold for significant matter effects. the electron number density N_e of the sun such that ν_e from ^{8}B neutrinos, however, oscillate in the "fast oscillations" regime. The fraction of ν_e that are ν_1 is f_1 and likewise for ν_2 and f_2:

$$\langle P(\nu_e \to \nu_e)\rangle = f_1 \cos^2(\theta_{12}) + f_2 \sin^2(\theta_{12}) \tag{1.43}$$

For the ^{8}B neutrinos eminating from the sun, $f_2 = 0.9$ by the time they reach the vacuum of space. Solar neutrino experiments detecting ^{8}B neutrinos are thus observing a nearly pure ν_2 solar neutrino fluxes and can measure $|U_{e2}|^2$.

Several other experiments have measured fluxes of solar neutrinos from different fusion processes, based on the solar model and the observed neutrino energy. These experiments, such as BOREXINO [36], GALLEX [37], and SAGE [38], have observed smaller deficits of ν_e in different channels that are consistent with the expected values of f_1 and f_2 in Equation 1.43 due to oscillations.

In 2002 the Sudbury Neutrino Observatory (SNO), in conjunction with SuperK, made an important measurement of the total solar neutrino flux. SNO was a water-Cherenkov detector filled with pure heavy water (D_2O) which was sensitive to both the total solar neutrino flux (through NC interactions dissociating the deuteron, followed by observing the delayed neutron capture) and the ν_e flux (by observing normal CC inter-actions). SNO found that the rate of solar neutrinos in ν_e CC interactions was 1/3 the
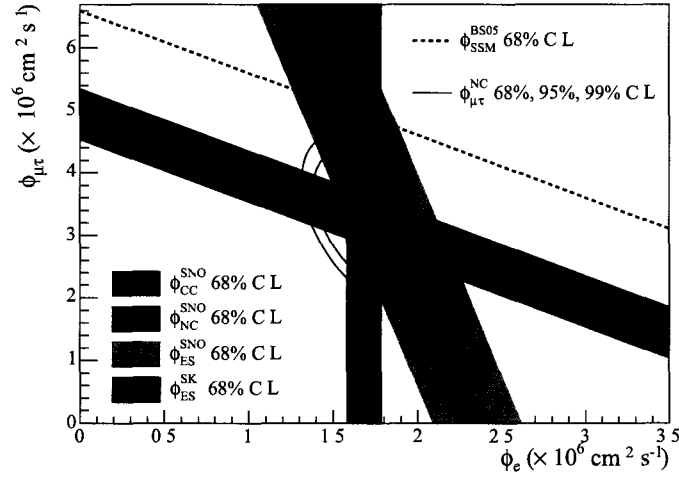
FIG. 1.8· Fluxes of ^{8}B solar neutrinos, $\phi(\nu_e)$, and $\phi(\nu_\mu \text{or } \nu_\tau)$, deduced from the SNO's CC, elastic scattering (ES), and NC results of the salt phase, where salts had been introduced to enhance the NC measurement [40]. The Super-Kamiokande ES flux is from [41]. The BS05(OP) standard solar model prediction Bahcall is also shown. The bands represent the 1σ error. The contours show the 68%, 95%, and 99% joint probability for $\phi(\nu_e)$ and $\phi(\nu_\mu \text{or } \nu_\tau)$. The figure was originally published in [40].

predicted rate, but that the rate of NC interactions matched the expectation [39]. Thus the total flux of neutrinos predicted by Bachall *et al.* was correct, but the MSW effect described above had converted the ν_e flux into a mostly ν_2 flux, where $\langle \nu_e | \nu_2 \rangle \approx 1/3$. The results from SNO are shown in Figure 1.8. This result corroborated a prediction made by the oscillation framework, and confirmed oscillations as the predominant explanation of solar neutrino disappearance.

1.7.2 Reactor Neutrinos

The "solar" neutrino oscillation parameters can also be measured with neutrinos emanating from nuclear reactors. Recall the very first neutrino detection experiment detected $\bar{\nu}_e$ emitted by the core of a nuclear reactor. Electron anti-neutrinos are created in β-decay

in the nuclear fission process.

$$n^0 \rightarrow p^+ + e^- + \bar{\nu}_e$$
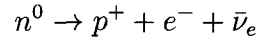
The neutrino can then be detected by the prompt positron produced in inverse β-decay, followed by the delayed capture of the neutron:

$$\bar{\nu}_e + p^+ \rightarrow e^+ + n^0$$
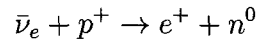
The KamLAND experiment in Japan used a detector containing 1 kilo-ton (kt) of liquid scintillator located near several nuclear reactors in Japan, Korea, and Russia. The "solar" neutrino flux of $\bar{\nu}_e$ was the sum of the fluxes of all nearby reactors, weighted by L/E, based on the reactor power output and baseline distance from the detector. The largest source of background, which was subtracted from the overall signal, was geoneutrinos, $\bar{\nu}_e$ from the decay of radioactive elements within the Earth. The KamLAND collaboration showed $\bar{\nu}_e$ disappearance over a range of L/E covering two oscillation maxima [42]. This confirmed solar neutrino oscillations with a man-made source. The oscillation signal is shown in Figure 1.9. The best-fit oscillation parameters for this data are $\Delta m_{21}^2 = (7.58^{+0.21}_{-0.20}) \times 10^{-5} \, \text{eV}^2/\text{c}^4$ and $\tan\theta_{12} = 0.56^{+0.14}_{-0.09}$.

The combination of measurements from KamLAND with the results from SNO also made a measurement of the mass hierarchy. KamLAND observed $\bar{\nu}_e$ passing through the Earth's crust, while SNO observed ν_e passing through the entire Earth at nighttime. By comparing the allowed parameter space for θ_{solar} measured by the two experiments, they
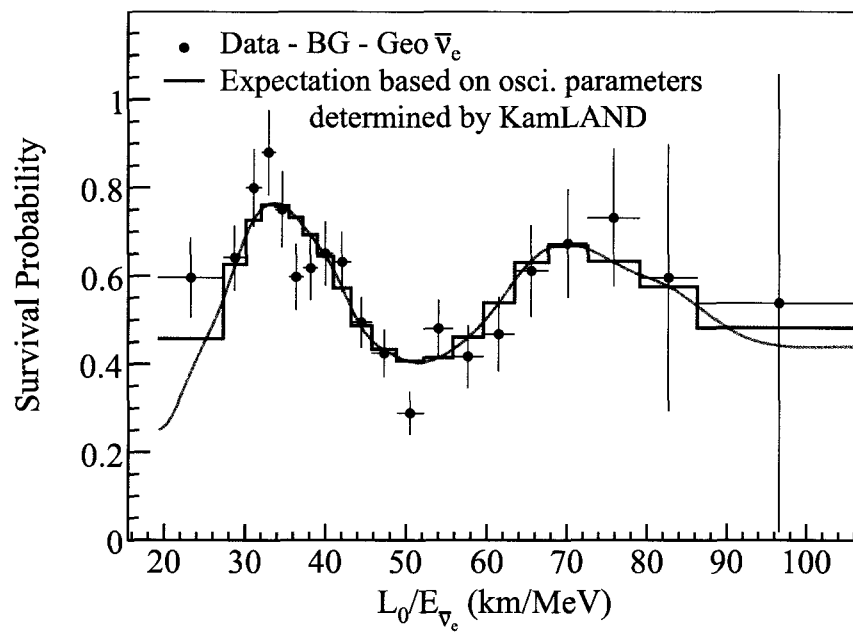
FIG. 1.9: The survival probability of $\overline{\nu_e}$, with geoneutrinos and other backgrounds subtracted, displayed as a function of L/E from different nearby reactors. The curves show the best-fit expectations for $\overline{\nu_e}$ oscillations [42].

could compare the relative MSW resonance effect between ν_e and $\overline{\nu_e}$. The combined results showed that $m_{\nu_1} > m_{\nu_2}$, which resolved one of two mass hierarchy ambiguities [43].

1.7.3 Atmospheric Neutrinos

The study of atmospheric neutrinos has progressed beyond the measurement of the ν_μ/ν_e flux ratio which illuminated the atmospheric neutrino anomaly. Cosmic rays strike the atmosphere isotropically and the neutrinos produced can travel through the Earth before interacting in a detector. For a detector in a fixed position on Earth, intersecting cosmic-ray neutrinos travel through a wide range of baselines. An experiment can determine the neutrino's baseline from their generated position in the atmosphere by reconstructing its zenith angle. An asymmetry between upward-going neutrinos and downward-going neutrinos indicates that the upward-going neutrinos are oscillating on a baseline of less than the diameter of the Earth.

The Super-Kamiokande experiment (SuperK), in Kamioka, Japan, has made the most accurate measurement of atmospheric neutrino rate as a function of zenith angle. The data is shown in Figure 1.10. This analysis found that $\Delta m^2_{atm} = \left(2.4^{+0.6}_{-0.5}\right) \times 10^{-3}\,\mathrm{eV^2/c^4}$ and $\sin^2(2\theta_{atm}) > 0.9$ at 90% confidence. As described in [33], the approximation that the atmospheric mass splitting $\Delta m^2_{atm} \approx \Delta m^2_{32}$ is valid within the precision of modern experiments.
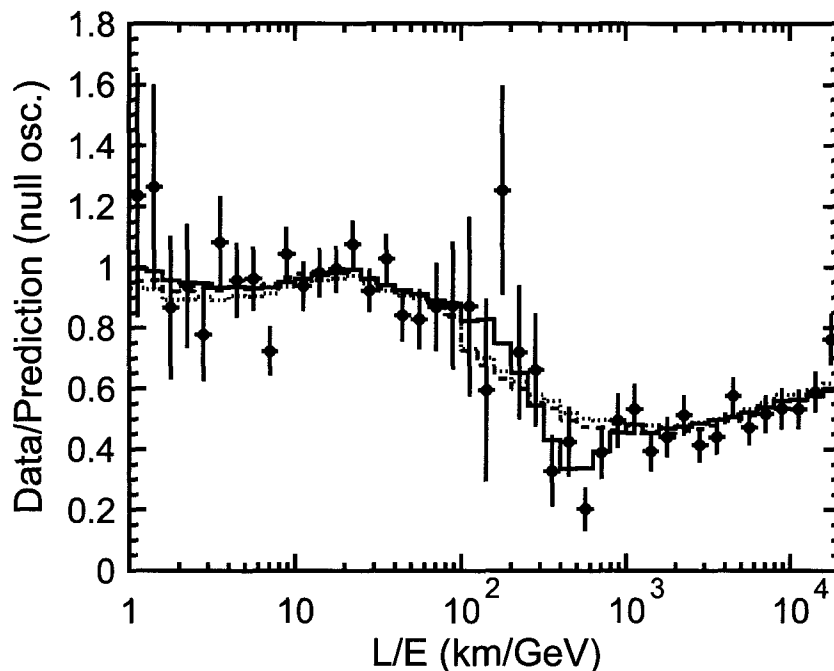
FIG. 1.10: Ratio of the atmospheric neutrino flux to that expected in the absence of oscillations as a function of the distance travelled divided by the incident neutrino energy, as measured by Super-Kamiokande. The black points are the data and the solid black line is the best oscillation fit. Also shown are the best fits to neutrino decay (dashed blue) and neutrino decoherence (dotted red) [44]. These two models were disfavored but not excluded by the data.

1.7.4 Accelerator Neutrinos

Modern accelerator neutrino beams use much the same technique that Lederman *et al.* used to discover the ν_μ [7]. A proton beam is used to expose a target creating charged mesons, mostly $\pi^\pm$, which then decay into $\mu^+(\mu^-)$ and $\nu_\mu(\bar{\nu}_\mu)$. This is a tertiary neutrino beam, since the neutrino is the daughter of the secondary.

Kamioka-to-Kamiokande (K2K) was the first experiment to detect neutrino oscillations with man-made neutrinos [45]. The K2K experiment used the 12 GeV proton synchotron in Tsukuba, Japan to produce low-energy neutrinos. These neutrinos were detected with the SuperK detector in Kamioka, Japan, 250 km away. A smaller 1 kT

water-Cherenkov and scintillator detector sat 300 m from the beam target and was used to characterize the neutrino flux. This was a counting analysis, where the total number of neutrino interactions in the SuperK detector was predicted, integrated over all energies. The actual number of neutrinos detected was 4.3σ lower than the prediction, consistent with $\nu_\mu \to \nu_x$ atmospheric neutrino oscillations.

The Main Injector Neutrino Oscillation Search (MINOS), the subject of this thesis, uses a tertiary neutrino beam originating with 120 GeV protons from the Main Injector at Fermilab. MINOS utilizes a pair of detectors, one 1 km away and one 735 km away from the beam target. MINOS is designed to study the atmospheric sector and to measure the atmospheric oscillation parameters with high precision. This experiment has already released precision measurements at various integrated beam exposures [46, 47]. The experimental apparatus will be described in more detail in Chapter 2.

OPERA is an experiment using blocks of an emulsion material to detect the τ exiting ν_τ charged-current interactions. It is located in a beam of ν_μ from 400 GeV protons from the CERN Super Proton Synchrotron, on the Cern-to-Gran-Sasso (CNGS) beamline. The detector itself is located in Gran Sasso, Italy. OPERA is located 730 km from the beam target, a similar baseline to MINOS. OPERA is currently the only experiment able to confirm that a disappearing ν_μ can in fact be detected as a ν_τ, strengthening the case for the oscillation hypothesis of neutrino disappearance. OPERA has recently observed the appearance of a single ν_τ in their detector, with 2.4σ significance over backgrounds, after accumulating 5.3×10^{19} protons-on-target [48].

1.7.5 Measuring θ_{13}

Experiments in the atmospheric and solar sectors measure θ_{12} and θ_{23}, but not θ_{13}. Given the known values and uncertainties on θ_{12} and θ_{23}, θ_{13} must be small. Recall that θ_{13} is the mixing angle between the ν_1 and ν_3 mass states, and from Figure 1.5 that the ν_1 mass state is composed of primarily ν_e, while ν_3 is split approximately evenly between ν_e, ν_μ, and ν_τ.

For reactor neutrino experiments, the value of θ_{13} is measured by observing a deficit of $\overline{\nu_e}$ relative to flux expectations from nearby reactors. The $\overline{\nu_e}$ survival probability is

$$P(\overline{\nu_e} \to \overline{\nu_e}) \simeq 1 - \sin^2(2\theta_{13}) \sin^2 \left(\Delta m_{atm} \frac{L \text{ (km)}}{E \text{ (GeV)}} \right) \tag{1.44}$$

The current best limit on θ_{13} was set by the CHOOZ reactor experiment in Chooz, France, which used a scintillator tank, with an external instrumented veto volume, to observe reactor neutrinos via inverse β-decay ($p + \overline{\nu_e} \to n + e^+$) [49]. The experiment ran with the reactor on and off, monitored the reactor power output (which correlates to neutrino flux), and did not observe a statistically significant deficit of $\overline{\nu_e}$. CHOOZ was able to place an upper limit on $\sin^2(2\theta_{13}) < 0.15$ at the 90% confidence level, with the Δm^2_{atm} from SuperK and MINOS described above.

In accelerator neutrino experiments, θ_{13} is determined through measuring sub-dominant oscillations of $\nu_\mu \to \nu_e$. The ν_e appearance probability is

$$P(\nu_\mu \to \nu_e) \simeq \sin^2(\theta_{23}) \sin^2(2\theta_{13}) \sin^2 \left(1.27 \Delta m_{atm} \frac{L \text{ (km)}}{E \text{ (GeV)}} \right) \tag{1.45}$$

Accelerator experiments have not achieved the sensitivity of CHOOZ, but there have been interesting results in accelerator experiments looking at $\nu_\mu \to \nu_e$ in a two-flavor model. The Liquid Scintillator Neutrino Detector (LSND) experiment at Los Alamos searched for $\bar{\nu}_\mu \to \bar{\nu}_e$ oscillations and saw a excess of ν_e events at low energy, consistent with a large $\Delta m^2 \approx 1\,\text{eV}^2$ [50]. They later corroborated their own result with a 3.8σ excess showing $\nu_\mu \to \nu_e$ [51] oscillations.

This large mass splitting, known as the LSND anomaly, is larger than the sum of the other two mass splittings, Δm_{21}^2 and Δm_{32}^2. This hinted at the existence of a possible fourth neutrino flavor to which the ν_μ could be oscillating. This fourth neutrino could not couple to the known leptons in a Charged-Current interaction, nor could it couple to the Z^0 due to the narrowness of the Z^0 decay width. The possible fourth neutrino would have to be *sterile*, non-interacting in matter. While the existence of a fourth, sterile, neutrino produces interesting implications, other experiments, such as KARMEN2 [52], NOMAD [53], and Bugey [54] did not corroborate this ν_e or $\bar{\nu}_e$ excess.

The MiniBooNE experiment was designed to precisely test this result. MiniBooNE used 8 GeV protons from the Booster ring at Fermilab to create a ν_μ beam aimed at a 0.8 kT mineral oil tank 541 m away [55]. MiniBooNE did not see any signifigant low energy excess, effectively ruling out the LSND result if $\nu_\mu \to \nu_e$ behaved in the same way as $\bar{\nu}_\mu \to \bar{\nu}_e$. To verify this, MiniBooNE later ran with $\bar{\nu}_\mu$ and did see a low-energy excess consistent with LSND [56]. The region of agreement is shown in Figure 1.11. The MiniBooNE experiment is still ongoing, and will run at least until 2013.

MINOS attempted to measure a ν_e appearance signature from its ν_μ beam, in excess
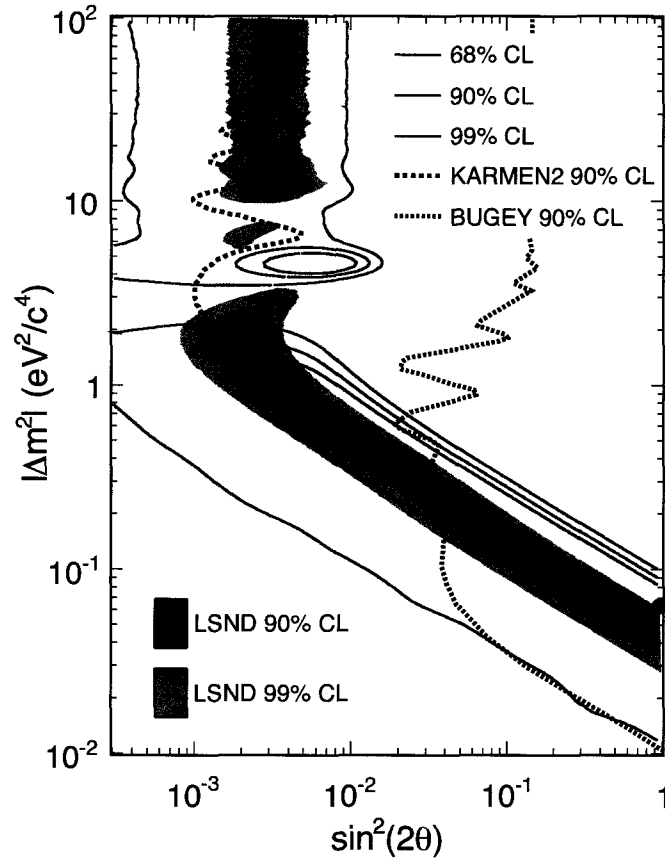
FIG. 1.11: MiniBooNE 68%, 90%, and 99% C.L. allowed regions for events with $E_\nu^{QE} >$ 475 MeV within a two neutrino $\nu_\mu \rightarrow \nu_e$ oscillation model. Also shown are limits from KARMEN [52] and Bugey [54]. The Bugey curve is a 1-sided limit for $\sin^2_{2\theta}$ corresponding to $\Delta\chi^2 = 1.64$, while the KARMEN curve is a "unified approach" 2D contour. The shaded areas show the 90% and 99% C.L. LSND allowed regions. The black dot shows the best fit point [56].

of the beam ν_e contamination [57]. MINOS observed a small 0.7σ excess consistent with

two-flavor oscillations. The exact value MINOS measures for θ_{13} depends on the value of

δ_{CP} due to matter effects, and the value of δ_{CP} is so far unknown. The allowed region for

MINOS and CHOOZ are shown in Figure 1.12.

Other experiments are in the works to measure θ_{13} more accurately. As their results

will not be known for some time, these experiments will be discussed in Chapter 9.

1.8   Status of oscillation parameters

The field of neutrino oscillation physics has made great advancements since the

Homestake experiment. The oscillation parameters Δm_{12}^2 and θ_{12}, as well as the sign

of Δm_{12}^2, have been measured to high precision from solar neutrino experiments and

from KamLAND. SuperK and MINOS have made the highest precision measurement of

Δm_{23}^2 and θ_{23}. CHOOZ and MiniBOONE have set limits on θ_{13}. The relative sizes of

Δm_{12}^2 and Δm_{23}^2 make the approximation in Equation 1.31 valid.

Δm_{12}^2	$(7.59 \pm 0.20) \times 10^{-5}\,\mathrm{eV}^2/\mathrm{c}^4$
Δm_{23}^2	$(2.43 \pm 0.13) \times 10^{-3}\,\mathrm{eV}^2/\mathrm{c}^4$
$\sin^2(2\theta_{12})$	0.87 ± 0.03
$\sin^2(2\theta_{23})$	> 0.92
$\sin^2(2\theta_{13})$	< 0.10

TABLE 1.1: The current measured values for parameters governing three-flavor oscillations.

With all of the data accumulated, some parameters of the neutrino oscillation model

still remain unknown. The unsolved problems themselves will be discussed here, and

their prospects for future experimental testing will be discussed in Chapter 9.
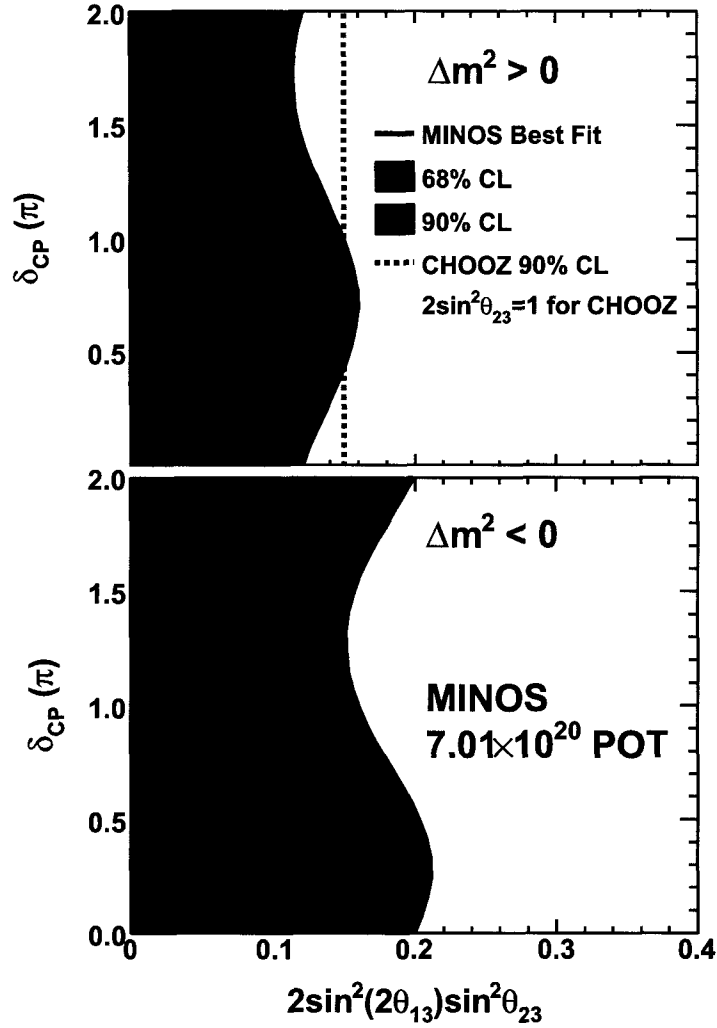
FIG. 1.12: Values of $2\sin^2(2\theta_{13})\sin^2\theta_{23}$ and δ_{CP} that produce a number of candidate events in the MINOS Far Detector consistent with the observation for the normal hierarchy (top) and inverted hierarchy (bottom). Black lines show those values that best represent the MINOS data. Red (blue) regions show the 90% (68%) C.L. intervals. The CHOOZ limit is drawn for $\Delta m^2_{23} = 2.43 \times 10^{-3}\,\mathrm{eV}^2$, $\sin^2(2\theta_{23}) = 1.0$ [57].

1.8.1 Mass Hierarchy

With two independent mass splittings, Δm_{12}^2 and Δm_{23}^2 ($\Delta m_{13}^2 = \Delta m_{12}^2 + \Delta m_{23}^2$) there are two possible mass spectra available to neutrino mass eigenstates. In the so-called Normal Hierarchy, the ν_3 mass state is more massive than the ν_2 and ν_1 states, so the sign of $\Delta m_{32}^2 \equiv m_3^2 - m_2^2$ is positive. The inverted hierarchy is a mass spectrum where ν_3 is less massive than either ν_1 or ν_2. In this case, Δm_{32}^2 is negative. These two hierarchies are shown in Figure 1.5. While the sign of Δm_{12}^2 was measured from the effect of MSW on solar neutrinos and antineutrinos, the sign of Δm_{23}^2 will have to be measured analogously for atmospheric neutrinos.

1.8.2 Dirac or Majorana?

Neutrinos are massless in the SM, only left-handed neutrinos (and right-handed anti-neutrinos) couple to charged leptons by weak interactions. Oscillations and related mass splittings imply non-zero neutrino masses. There are two ways for the neutrino to acquire mass in the SM with relatively small modifications, but oscillation experiments are not sensitive to tests of these modifications. They are noted here only for completeness.

Inserting right-handed neutrinos into the SM gives neutrinos masses through a Dirac mass term. The neutrinos then gain masses just like the quarks and charged leptons. The right-handed neutrino (and left-handed antineutrino) would have to be sterile, to explain why they have not yet been observed.

Since we know that all neutrinos are left-handed and all antineutrinos are right-handed, it is possible that neutrinos are their own anti-particle, with the differences we

observe between neutrinos and antineutrinos due simply to the difference in chirality. This would make neutrinos Majorana particles, and would imply that lepton number violation is possible in the form of neutrinoless double β-decay. Neutrinoless double β-decay is most simply a case where a nucleus emits two electrons and two neutrinos, but the neutrinos annihilate before leaving the nucleus since they are their own anti-particle.

It is possible that neutrinos have both Dirac and Majorana mass terms in the SM, a situation which leads to an intriguing description of the small size of the neutrino mass scale. In this case, there exists a very massive neutrino, out of the reach of LHC, which suppresses the masses of the other neutrinos. This is called the SeeSaw mechanism, and is a common feature of Grand Unified theories (a more thorough derivation can be found in [24]).

1.8.3 Mass Scale

Oscillation experiments only measure the differences in the masses of neutrino mass eigenstates, not the actual masses themselves, but the mass splittings measured from oscillations do place lower bounds on neutrino masses. In the limiting case that the $m_{\nu_1} = 0$ in the normal hierarchy, then $m_{\nu_3} = \Delta m_{12}^2 + m_{23}^2$.

Fermi knew from β-decay experiments in his time that the neutrino mass must be small. The mass could be measured with more precision today with the same method by observing the endpoint of the hard β spectrum in $^3\text{H} \rightarrow ^3 \text{He} + e^- + \bar{\nu}_e$. The current

experimental upper limit from this measurement is

$$m_\nu^2 = \sum_{i=1}^{3} |U_{\alpha i}|^2 m_{\nu_i} < 2.2\,\text{eV} \qquad (1.46)$$

at 95% C.L. [58], but future experiments such as KATRIN [59] will be sensitive to 0.5 eV.

There is also a cosmological bound based on anisotropy observed in the Cosmic Microwave Background, since very massive neutrinos ($\sum m_{\nu_i} > 1$ eV) would supress the observed anisotropy. The cosmological limit calculated from WMAP data is $\sum m_{\nu_i} < 0.3\,\text{eV}^2/\text{c}^4$ [60].

If neutrinos are Majorana, then the neutrinoless $\beta\beta$-decay rate is proportional to the effective Majorana mass of ν. The upper limit on neutrino mass from this reaction is $m_\nu < 0.34\,\text{eV}$ at 90% C.L. [61].

1.8.4 CP-Violating Phase

The CP-violating phase δ_{CP} manifests itself in an asymmetry in the MSW effect observed between ν_e and $\bar{\nu}_e$. It also plays in to the measurement of θ_{13}, as shown in Figure 1.12. Future long-baseline experiments such as T2K and NOνA could observe the CP-violating phase by comparing the effective θ_{13} and $\bar{\theta}_{13}$ after passing through hundreds of kilometers of matter.

1.9 Alternative disappearance models

Other models have been suggested to explain neutrino disappearance over long base-lines. Many disappearance models date back to early results from SuperK, which were too limited to discriminate between disappearance models. Additional data from other neutrino oscillation experiments have also tested these models, and found that they do not describe the energy-dependent neutrino disappearance data as well as the oscillation model. Two models, neutrino decay and neutrino quantum decoherence, were disfavored but not eliminated with additional data. We will test these two models with the MINOS analysis described in this thesis.

1.9.1 Neutrino Decay

At least some portion of neutrino disappearance could be attributed to a neutrino state decaying into a sterile state which no longer mixes with the others, assuming such a state exists. The decay could occur in addition to neutrino oscillations, and the survival probability in that case becomes [62, 63].

$$P_{\nu_\mu \to \nu_\mu} = \sin^4(\theta) + \cos^4(\theta)e^{\frac{-\alpha L}{2E}} + 2\sin^2(\theta)\cos^2(\theta)e^{\frac{-\alpha L}{2E}}\cos\left(\frac{\Delta m_{23}^2 L}{2E}\right) \qquad (1.47)$$

where two neutrino flavors are assumed and $\alpha = \frac{m_2}{\tau_2}$ where τ_2 is the decay constant. If the decay product is $\nu_2 \to \bar{\nu}_3 + J$, where J is a massless scalar, and if Δm_{23}^2 is large enough

that the third term averages to zero through rapid oscillations, we are left with

$$P_{\nu_\mu \to \nu_\mu} = \sin^4(\theta) + \cos^4(\theta) e^{\frac{-\alpha L}{2E}} \tag{1.48}$$

The Decay hypothesis was found to fit early Super-K atmospheric neutrino results. Kam-LAND disfavored this model by 2.8σ [42]. More recently, in 2008 efforts were made to use MINOS data to fit pure neutrino decay in Equation 1.48. This model was found to be disfavored by 3.7σ [47].

1.9.2 Quantum Decoherence

Another possible mechanism for neutrino disappearance is a decoherence introduced to the quantum mechanical wave function of neutrinos. Quantum decoherence is an effect that one would expect over very long baselines (i.e. neutrinos from supernovae), but to observe decoherence over terrestrial baselines requires an additional potential. Decoherence can be introduced to the neutrino wave packet due to interactions with Planck-scale quantum foam in some theories of quantum gravity [64].

Quantum decoherence introduced to the neutrino wave packet can affect the neutrino survival in addition to neutrino oscillations. In that case, the survival probability for a muon neutrino traveling a distance L and with an energy E in GeV is

$$P_{\nu_\mu \to \nu_\mu} = 1 - \frac{\sin^2(2\theta)}{2} \left[1 - e^{\frac{-\mu^2 L}{2E}} \cos\left(\frac{\Delta m^2 L}{2E}\right) \right] \tag{1.49}$$

In the limit that there are no oscillations ($\Delta m^2 = 0$), this expression reduces to

describe the effect of pure decoherence:

$$P_{\nu_\mu \to \nu_\mu} = 1 - \frac{\sin^2(2\theta)}{2} \left(1 - e^{\frac{-\mu^2 L}{2E}}\right) \tag{1.50}$$

The pure decoherence model was shown to fit data in the atmospheric sector for Super-K and K2K in 2003 [65]. KamLAND disfavored this model by 2.45σ [42]. More recently, the MINOS analysis mentioned in the previous section considered pure quantum decoherence [47], and the hypothesis was disfavored by 5.7σ.

1.10 Conclusion

The 20th century saw the development of particle physics from a simple model consisting of a single charged electron and a single charged proton to Quantum Chromodynamics and Electroweak unification. Experimental neutrino physics played a large part in the development of the latter, as the scientific revolution and fundamental physics have been a triumph of the 20th century.

In modern neutrino physics we have moved beyond the Solar Neutrino problem and the Atmospheric Neutrino Anomaly, and are working to quantify the parameters that govern neutrino disappearance. The neutrino oscillation hypothesis, with its mixing of flavor eigenstates and mass eigenstates, has proven to be the model which best describes the neutrino flavor transitions observed by experiments in the last fifty years. The experimental evidence began with the Homestake experiment, introducing the world to the Solar Neutrino Problem, which was subsequently solved with the observation of solar NC in-

teractions by SNO in 2002 and confirmed by KamLAND [40, 42]. Explaining neutrino disappearance in any form will require modification of the Standard Model, as oscillations demonstrate the need for a neutrino mass term in the SM Lagrangian.

A note on the notation used in the rest of this thesis - since the MINOS measurement makes use of the two-flavor approximation, we will use the shorthand $\Delta m^2 \equiv \Delta m_{32}^2$ and $\sin^2(2\theta) \equiv \sin^2(2\theta_{23})$ [33].

This thesis describes the analysis of the disappearance of accelerator-produced muon neutrinos over a long baseline, as measured by the MINOS experiment after collecting data over a period of four years. MINOS is an experiment that relies on a large and knowledgeable collaboration for its construction, maintenance, and analysis. Much of the work herein represents the work of the entire collaboration, but my efforts will be highlighted in this document.

Chapter 2 describes the apparatus that makes up the NuMI beam, the MINOS experiment and how it detects neutrino interactions. Chapter 3 describes the simulations of the beam and detectors which are used to characterize the detectors prior to oscillation analyses. Chapter 4 discusses the topology and properties of events in the two MINOS detectors. In Chapter 5, the improvements over the previously published MINOS analyses that are employed in the current analysis are described. Chapter 6 details the procedures used to extrapolate the NuMI beam over the MINOS baseline. Chapter 7 shows the blind analysis methods, the sensitivities, and statistical uncertainties of the data that is the subject of this thesis. Chapter 8 shows the ultimate results of the analysis, the measurement of the oscillation Δm^2 and $\sin^2(2\theta)$, and compares them to the values measured by other

experiments.
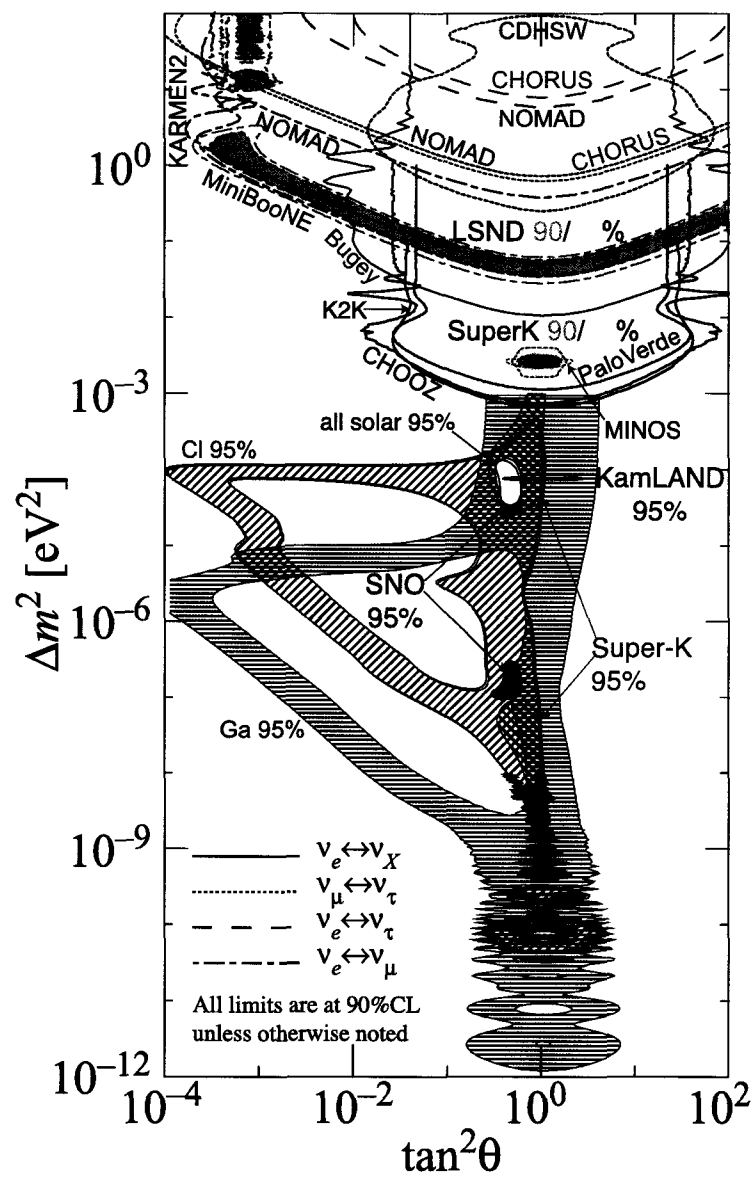
FIG. 1.13: The regions of squared-mass splitting and mixing angle favored or excluded by various proton decay, solar neutrino, and atmospheric neutrino experiments. This plot is from [20].

CHAPTER 2

The MINOS Experiment

The Main Injector Neutrino Oscillation Search (MINOS) is a long-baseline neutrino oscillation experiment with two detectors exposed to an intense neutrino beam produced by the Neutrinos from the Main Injector (NuMI) beam. The NuMI beam is located on the grounds of the Fermi National Accelerator Laboratory (Fermilab) in Batavia, IL, 50 miles west of Chicago, IL. The two detectors are both steel and scintillator sampling calorimeters, and though they have different sizes, they are designed to be functionally identical, with similar hardware components and software. The Near Detector is located on the grounds of Fermilab 1 km downstream from the NuMI beam's target, and the Far Detector is located 735 km away, in the Soudan Underground Laboratory in Soudan, MN. This chapter describes the design and operation of the NuMI beam and the MINOS detectors. For a detailed look at the detectors and their construction, see [66].

2.1 The NuMI Beam

The process of making a beam of muon neutrinos is essentially unchanged since the Nobel Prize winning experiment that first proved the existence of muon neutrinos in 1962 [7]. High-energy protons are smashed on a fixed target and the interactions produce pions and kaons. The pions and kaons decay into neutrinos and muons on a short timescale:

$$\pi^{\pm} \to \mu^{\pm} + \nu_{\mu}(\bar{\nu}_{\mu})$$
$$K^{\pm} \to \mu^{\pm} + \nu_{\mu}(\bar{\nu}_{\mu})$$

(2.1)

A large mass of rock stops the muons, and leaves only a beam of neutrinos.

A neutrino beam is focused by focusing the π and K off the target with a pair of electromagnetic focusing horns. These horns allow for specific ranges of neutrino energy to be selected by changing the relative positions of the target and horns. The focused mesons travel through an evacuated decay pipe to minimize secondary interactions before they decay.

The NuMI facility is designed to deliver a neutrino beam to the MINOS experiment. It is located on the grounds of Fermilab adjacent to the Main Injector accelerator ring, which was initially designed to deliver protons to the larger Tevatron accelerator ring for proton-antiproton collider experiments. A schematic of the NuMI facility is shown in Figure 2.1. A kicker magnet extracts protons with momentum 120 GeV/c from the Main Injector every 2.2 seconds. The beam of protons is bent 58 mrad below the horizontal to account for the curvature of the earth when aiming for the Soudan Underground Laboratory. A spill of NuMI protons strike a fixed graphite target located 350 m downstream

from the extraction point [19]. The spill of protons on the target lasts $10\,\mu s$ and comes in five or six batches from the Main Injector, depending on whether or not antiprotons are being produced for the Tevatron collider ring at Fermilab.

2.1.1 The Target

The target is comprised of 47 fins of graphite which are 15 mm tall, 5.4 mm wide, and 20 mm deep, in the direction of the beam. The fins are aligned edgewise with respect to the beam. The target is enclosed within an aluminum vacuum vessel filled with gaseous helium and with beryllium windows [19]. The thin edge of the target is presented to the beam to minimize secondary interactions of π and K in the target. The total thickness of the target represents 2.4 interaction lengths for the incident protons. The graphite fins are continuously water cooled through pipes running along the top and bottom edges of the target assembly.

Upstream of the target assembly there are several pieces of equipment which protect the target and monitor the beam. A baffle sits upstream of the target assembly and protects it and the other downstream equipment from the proton beam if it is mis-steered. The baffle is a hollow cylinder of graphite and is 1.5 m long with an 11 mm inner diameter. Just upstream of the first target fin is a Budal monitor, which is a fin in a horizontal orientation and used to align the proton beam vertically. A toroid measures the current of protons moving through its center by magnetic induction to measure the total exposure of the experiment to neutrinos. The exposure is expressed in units of protons-on-target (POT). The upstream toroid has been determined to be accurate to $\pm 1\%$ [19].

Plan View

Extraction magnets

Main Injector Accelerator

Target Hall

Evacuated Decay Volume

Access Tunnel

Hadron Absorber

Muon Alcoves

target

focusing horns

Near Detector Hall

-ν beam

Elevation View

Ground Level

Surface Building

Surface Building

V108 Bend

Carrier Tunnel

V118 Bend

Target Hall

Evacuated Decay Volume

Hadron Absorber

Muon Alcoves

Service Shaft

Near Detector Hall

-ν beam

250 m

FIG. 2.1: Plan and elevation views of the NuMI beam facility. Taken from [19].

51

The intense beam of protons damages the graphite fins over time. There have been 2 targets used during the duration of the MINOS experiment described in this thesis. It is only possible to swap out targets during breaks in runs. Run periods are defined by the roughly annual shutdown of the Fermilab facilities for repairs and upgrades. One target was used in Run I, after which a different target was installed and used in the Run II and Run III periods. Targets are swapped when the degradation is too great to continue running. In Runs II and III, the target degradation of the second target is modeled and included in the beam systematic error, as will be described in Chapter 3.

2.1.2 Electromagnetic Focusing Horns

The entire target and baffle assembly is placed on rails, allowing it to move relative to two conical electromagnetic focusing horns, in line with the target and beam. The electromagnetic focusing horns are designed to focus π and K of one charge sign coming off the target into a beam, while defocusing π and K of the opposite charge sign. The focusing horns can be pulsed with a current of magnitude between $0\,\mathrm{kA} < I < 200\,\mathrm{kA}$, timed with each beam spill from the Main Injector. The focusing horns are water cooled with spray nozzles located around the horn assembly. A drawing of the horn assembly is shown in Figure 2.2. The polarity of the horns may be reversed to select the charge of the focused π and K. The "forward" horn current selects π^+ and K^+(and hence ν_μ), while the "reverse" horn current selects π^- and K^- (and hence $\bar{\nu}_\mu$). When π^+ are focused, some π^- travel through the center of the horn necks and are not deflected. These "neck-to-neck" pions produce a background of $\bar{\nu}_\mu$'s in a ν_μ beam. The inverse occurs when

negatively charged mesons are focused.
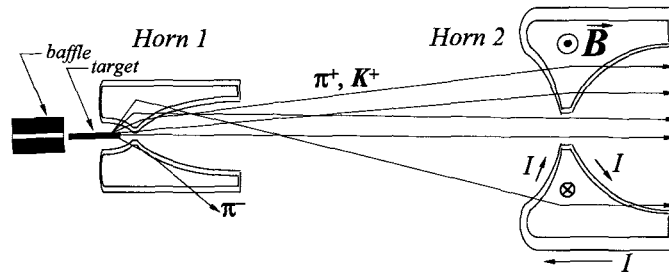
FIG. 2.2: A sketch of the NuMI focusing horns in "forward" mode. The vertical scale is exaggerated for display purposes. Horn 1 is 3 m long has a radius of 16.2 cm for the outside of the outer conductor. Horn 2 is also 3 m long and has a radius of 35.87 cm for the outside of the outer conductor [67]. Drawing from [19].

2.1.3 Beam Configurations

Varying either the location of the target relative to the horn or varying the current in the horn selects a kinematic range of pions from the target, and therefore a different neutrino flux in the beam. This was initially a design feature of the MINOS experiment, but in practice the configuration of target location and horn current are rarely changed, due to radiation damage of the wheels on the target cart. The nominal beam configuration used for the oscillation analysis is for the target's end to be inside the horn, 10 cm from the narrowest part of the horn 1, and for the horn current to be 185 kA. Short runs of many configurations are used in a beam tuning procedure, which is discussed in Section 3.4. Other available configurations and their flux profiles are shown in Table 2.1 and Figure 2.3. While the higher-energy beam configurations produce a larger overall neutrino flux, experimental evidence released after the start of construction of the NuMI beam [68] showed that the neutrino energy region of interest for oscillations was below

Configuration	Target Position (cm)	Horn Current (kA)	Total Exposures ($\times 10^{18}$ POT)
	10	0	10.36
	10	170	1.42
LE	10	185	9.93
	10	200	1.34
pME	100	200	1.12
	150	200	1.72
pHE	250	250	3.08

TABLE 2.1: Example target/horn configurations for the NuMI beam, and their beam exposure used in the beam fits, described in Section 3.4. The target position is measured relative to the neck of Horn 1. The Low Energy (LE) beam is the primary configuration used in the oscillation analysis, with a small contribution from the pseudo-High Energy beam configuration. For primary configurations, the exposure used in the beam fit is a subset of the total exposure used in the oscillation analysis [69].

4 GeV, where the LE beam configuration produces the largest flux.

2.1.4 Decay Pipe

After passing through the horns, the focused pions and kaons enter an evacuated decay pipe in which they can decay into muons and neutrinos with only a small probability of secondary interactions. The decay pipe is 675 m long, 2 m in diameter, and is made of steel surrounded by between 2.5 m and 3.5 m of concrete. The upstream end of the decay pipe has a 1 cm-thick Aluminum window 0.5 m in diameter. The outer diameter of this window, which mates to the decay pipe walls, is 2.3 cm-thick steel. The decay pipe was evacuated to 0.5 Torr between commissioning in 2005 and Fall of 2007. After that point, the decay pipe was filled with 1 atm of Helium, due to concerns about the corrosion of the aluminum window on the upstream end of the pipe. The addition of helium in the decay pipe changes the neutrino flux by introducing additional nuclear targets for pion scattering before their decay. The overall flux at higher energies (above 10 GeV) is lower
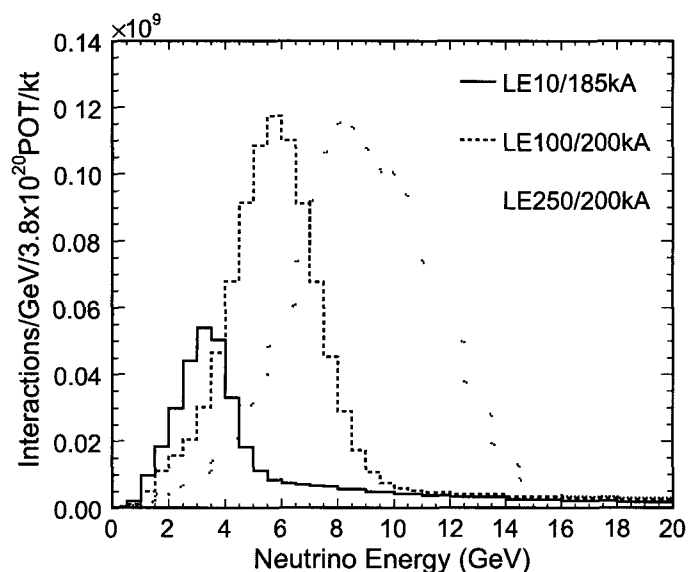
FIG 2 3   The Near Detector specra for three different beam configurations   The majority of the data taken for the analysis described in this thesis was in the LE10 configuration LE100 is another name for pME, and LE250 is another name for pHE

by about 5%, and the focusing peak shifted about 0.5 GeV lower [70]. The helium was

factored into later Monte Carlo simulation models of the flux, as will be described in later

chapters.

2.1.5   Muon Monitors

At the end of the decay pipe, surviving muons must travel through a beam dump

composed of water-cooled steel and aluminum, followed by steel and concrete blocks.

After the beam dump, there is 240 m of rock before the beam intercepts the near detector.

The majority of these muons are stopped in the rock. There are two muon monitors 12 m

and 30 m downstream of the back end of the decay pipe. A third muon monitor sits behind

the beam dump.

2.2 The MINOS Detectors

There are two detectors in the MINOS experiment, both of which are magnetized tracking and sampling calorimeters made of planes of steel and scintillator sandwiched together. The detectors are designed to measure the energy of neutrinos participating in Charged-Current neutrino interactions by reconstructing outgoing muons and hadronic showers originating from the struck nuclei. The steel planes act as inactive absorbing material and the scintillator acts as an active sampling calorimeter for any hadronic showers resulting from the initial interaction and a tracking spectrometer for the muons. The detectors are magnetized to contain the muons within the detector, to identify the charges of the contained muons, and also to measure the momentum of exiting muons. The two detectors are designed to behave similarly, although the Far Detector is roughly 5 times more massive than the Near Detector. A drawing showing the relative sizes of both detectors is shown in Figure 2.4. This section will describe their design and performance.

2.2.1 Steel

Both the Near and Far detectors use planes of steel for both nuclear targets for neutrino interactions and passive absorbers for the products of those interactions. The steel is also the medium that carries the magnetic field, and provides a structure for the detector. The planes used in both detectors are made from 1004 alloy low-carbon, hot-rolled steel that were manufactured at Bethlehem Steel in Indiana. The iron itself is from Minnesota, and the steel planes for the Near Detector were cut in Iowa (ND) and Minnesota (FD). The planes in the Near Detector are 2.56 cm thick solid planes, while each plane in the

Far Detector is constructed of 8 smaller plates welded together to make a 2.56 cm thick octagonal plane. This piecewise construction was necessary for the Far Detector to accommodate the size and shape of the mineshaft used to convey the pieces of the detector from the surface to the Soudan Underground Laboratory.

The steel was made in batches called heats, and the heats of steel comprising the Near Detector are a subset of heats in the Far Detector. The same heats were used to ensure similar density and magnetic properties between the two detectors. The planes in the Near Detector were measured with an ultrasound probe and were found to be 2.563 ± 0.002 cm thick, averaged over all planes. The planes in the Far Detector were measured with the same ultrasound probe and found to be 2.558 ± 0.005 cm thick. The density of the steel was measured to be 7.85 ± 0.03 g/cm^3, with no systematic difference in the density of the steel making up the two detectors [66].

(a) Near Detector   (b) Far Detector
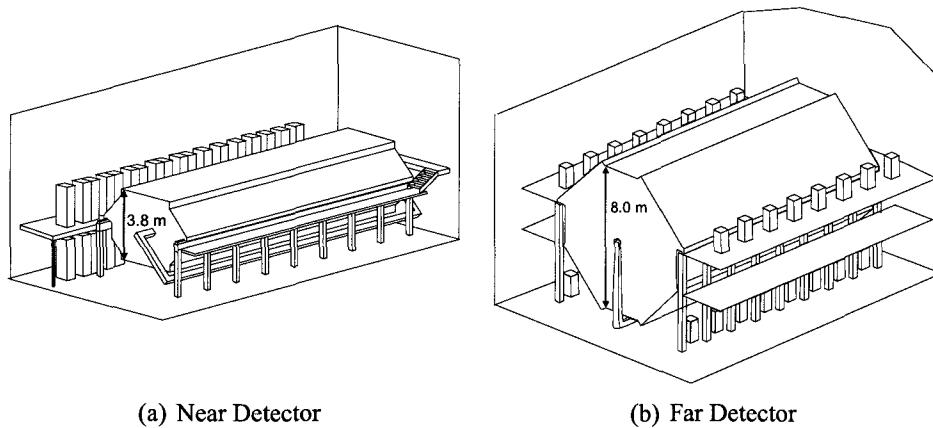
FIG. 2.4: The Near Detector and one supermodule of the Far Detector. The full Far Detector consists of two supermodules, placed end to end. Drawings taken from [19].

2.2.2 Scintillator

The scintillator planes record the passage of ionizing particles, and are used for both calorimetry and collecting tracking information. Planes of scintillator are composed of scintillator modules, which are themselves composed of scintillator strips glued together and held in place with an 0.5 mm-thick Aluminum outer skin. The strips are doped with organic fluors PPO and POPOP so that UV photons are emitted from the scintillating plastic when ionized by a particle passing through a strip [66].

The scintillator strips used in both the Near and Far detectors are made of extruded polystyrene, which were manufactured at Itasca Plastics in (Batavia, IL). The strips are co-extruded with a 0.25 mm thick layer of Titanium Dioxide (TiO_2) doped polystyrene on the outside for internal reflectivity and light-tightness. There is a 2.3 mm deep by 2.0 mm wide channel in the long edge of the strip which contains a wavelength-shifting (WLS) fiber. Once the WLS fiber has been laid in the channel, a piece of reflective aluminized Mylar tape is applied along the length of the strip to both hold the fiber in place and to maintain a high level of internal reflectivity and light-tightness within the strip.

Photons emitted in the scintillator will internally reflect within the strip until they are absorbed by a WLS fiber. The WLS fiber shifts the wavelength of the light absorbed from 420 nm to 470 nm. The re-emitted 470 nm photon is emitted in a different direction, allowing some of the light to be captured within the fiber by total internal reflection. Through total internal reflection within the fiber, the photon is guided to the end of the strip. A diagram of the scintillator is shown in Figure 2.5.

MINOS SCINTILLATOR STRIP

UP TO 8m

REFLECTIVE SEAL

TiO2 LOADED POLYSTYRENE CAP

CLEAR POLYSTYRENE
SCINTILLATOR

WLS FIBER

10mm

41mm
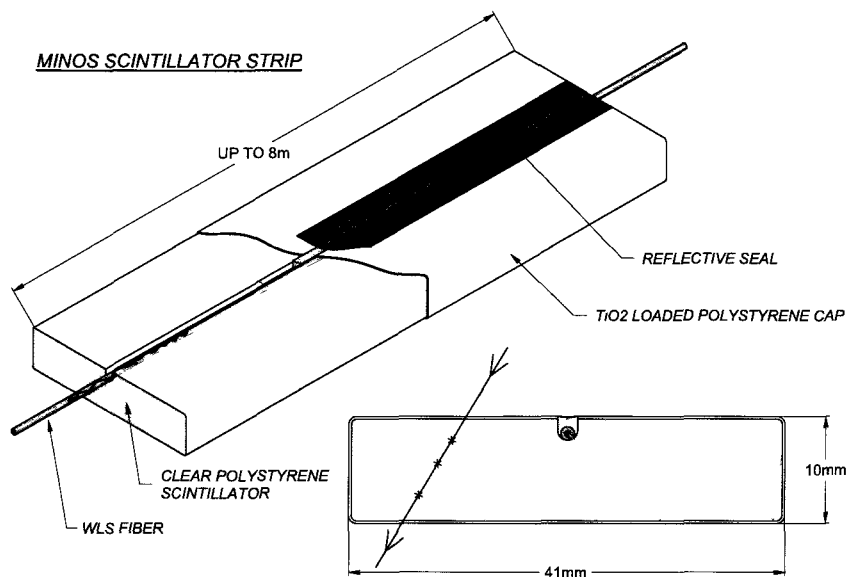
FIG. 2.5: A diagram of the scintillator used in both MINOS detectors. Light emmitted by the scintillator when an ionizing particle passes through is reflected many times within the TiO_2 coating, and some eventually reaches the WLS fiber. Some of the wavelength-shifted light ($<$ 5%) is then directed down the fiber by total internal reflection [66].

2.2.3 Scintillator modules

Scintillator strips, as described above, are arranged in modules and planes to be sandwiched between steel absorber planes in the detector. The u and v directions are defined to be orthogonal $\pm\pi/4$ radians from the horizonal and vertical, called x and y, respectively. Planes of scintillator are encased in and glued to an aluminum skin 0.5 mm thick for light-tightness and structural stability. A set of scintillator strips encased in aluminum is refered to as a module.

The strips in the Near Detector can be anywhere from to 2.5 to 6 m long. In the Near Detector, light yield is high enough that only one end needs to be read out. A mirror is glued to the other end with optical epoxy to collect more photons and improve the strip efficiency. The strips in the Far Detector can be anywhere from 3.4 m to 8 m, and are

read out at both ends. The sum of the signal at both ends of a strip is the total light yield recorded for a hit in that strip. The sum of the signal at both strip ends varies by only 25% along the length of the strip. [66].

2.3 The Near Detector

The Near Detector is located in the Near Detector Hall on the grounds of Fermilab. The cavern housing the Near Detector is located 100 m underground so that it intersects the neutrino beam, which is already underground and is aimed 58 mrad below horizontal. Its overburden is equivalent to 225 meters under water (meters-water-equivalent, mwe). It consists of 282 steel planes sandwiched with scintillator planes. The Near Detector planes are placed with a 5.95 ± 0.37 cm plane-to-plane pitch. There are four different kinds of planes in the Near Detector, full planes and partial planes, each of which may be either u or v view. The detector has a square hole offset 55.8 cm from the center of each plane for the current-carrying coil, which is used to magnetize the detector.

2.3.1 Calorimeter and Spectrometer

The first 120 planes of the detector comprise the MINOS Near Detector calorimeter. The calorimeter is meant to measure the energy of the products of the neutrino interactions within the detector with good spatial and calorimetric resolution. The calorimeter region contains mostly partially instrumented planes of scintillator. Partially instrumented planes are about 3 m × 3 m and are intended for measuring hadronic showers. The first and sixth planes of every ten planes are fully instrumented. The fully instrumented planes are
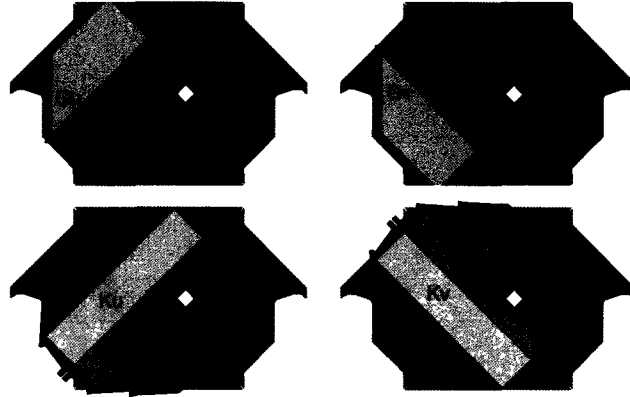
FIG. 2.6: A diagram of full and partial scintillator planes in the Near Detector, constructed of different type of modules, in both u and v orientation [66]

intended to reconstruct muon tracks. A diagram of full and partial scintillator planes in the Near Detector can be seen in Figure 2.6. Strips are oriented 45° from the horizontal and vertical axes (x and y, respectively). They are oriented along the orthogonal unit vectors u and v. The unit vector defined through the depth of the detector is z. The pattern of Full v-view (FV), Full u-view (FU), Partial v-view (PV) and Partial u-view (PU) is: FU-PV-PU-PV-PU-FV-PU-PV-PU-PV.

Hits in planes 1-20 identify particles created from interactions upstream of the detector, so this section is used as an upstream veto. Planes 21-60 are taken to be the target region, where interactions are likely from beam neutrinos interacting within in the target region of the calorimeter. These interactions are identified by an event vertex, a point within the target region where particles that are products of an interaction originate. Planes 61-120 are used to contain and measure the energy deposition of showers from interactions in the calorimeter. Neutrino interactions that occur this deep in the detector are rejected, since the hadronic showers from these interactions may not be fully contained

within the calorimeter.

Planes 121-281 comprise the MINOS spectrometer, as hadronic activity from interactions in the target region is minimal this far downstream. This section is used solely to track muons from CC interactions. There are no partially instrumented planes in the spectrometer, only alternating FV and FU scintillator planes, one for every five planes of steel.

2.3.2 Near Detector Readout

Hits in the Near Detector are read out by Hamamatsu 64-anode (M64) photomultiplier tubes (PMT's) housed singly in light-tight steel enclosures. The enclosures contain clear fiber bundles channeling photons from the WLS fibers to the PMT pixels. The number of photoelectrons (PE's) emitted at the photocathode in the PMTs is recorded. A drawing of this enclosure can be found in Figure 2.7. Due to the high event rate in the Near Detector, the electronics need to be fast and minimize dead-time. Each strip in the calorimeter is read out with its own individual pixel. In the spectrometer, four channels are electronically summed (or multiplexed) and read out as a single channel. The strips associated with each summed channel are 1 m apart, allowing the four-fold ambiguity of a single pixel to be solved in software by considering the location of upstream and downstream hits along a muon track.

The Near Detector PMT's are sampled by an analog-to-digital converter (ADC) at a rate of 53 Mhz with a threshold of 0.3 photoelectrons. The Near Detector uses a combination of a multi-ranging integrated circuit and an 8-bit ADC to achieve the dynamic

range of a 16-bit ADC with a constant calibrated (linearized) output error of 0.5%. One

photoelectron corresponds to about 106 ADC counts in the Near Detector [66].
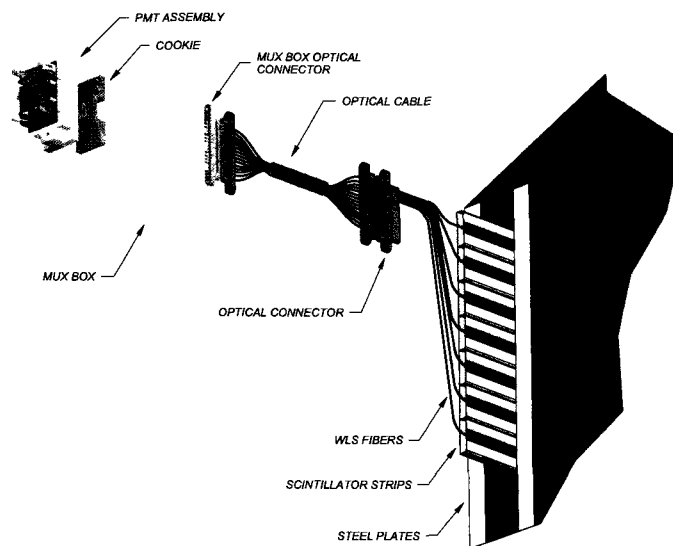
FIG. 2.7: A drawing of the interface between scintillator planes and WLS fibers to clear fibers, which route light to the pixels of the PMT. Drawing is taken from [66].

2.4 The Far Detector

The Far Detector is located in the Soudan Underground Laboratory in Soudan, MN. The cavern housing the Far Detector is 705 m underground (2070 mwe). The detector is octagonal, is 8 m across, and has a total mass of 5.4 kT. The Far Detector planes consist of eight steel planks welded together, as shown in Figure 2.8. There are only full scintillator planes in the Far Detector, with 192 strips per plane. The Far Detector consists of 484 active planes broken into two supermodules. Supermodule 1 (SM1) is upstream of Super-module 2 (SM2) and contains 249 active planes. SM2 contains 237 active planes. There is a 1.4 m air gap between the two supermodules. The Far Detector planes have the same
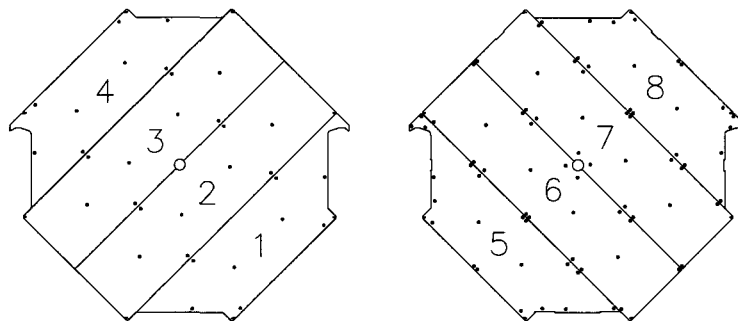
FIG. 2.8: The eight 0.5 in.-thick pieces that are welded together to create a single 1 in.-thick plane of the Far Detector. Black dots indicate weld points [66].

5.95 cm plane-to-plane pitch as the Near Detector, with a standard deviation of 0.35 cm. In the center of each plane is a circular hole for the current-carrying coil, which is used to magnetize the detector. Each supermodule has its own magnetizing coil, with the coil return vertical with respect to the floor of the lab and looping beneath the detector. Both of the coils are continuously water cooled throughout the length of the detector.

2.4.1 Veto Shield

Scintillator modules are suspended on the top and to the sides of the Far Detector, parallel to the z-axis, to act as a veto shield for cosmic rays. A minimum-ionizing cosmic ray muon passing through the veto shield and the volume of the Far Detector is tagged as a cosmic ray muon. The cosmic ray muon rate at the Far Detector is about 0.5 Hz [66]. Two layers of scintillator planes are suspended horizontally above the entire length of the Far Detector. The signals from the veto shield are read out at both ends by the same front-end electronics and data acquisition system as the rest of the Far Detector. The PMTs used to read veto shield channels are set to a higher dynode threshold (1-2 PE) to reduce the rate

of false-positives. A drawing of the positions of the veto shield modules can be found in Figure 2.9.
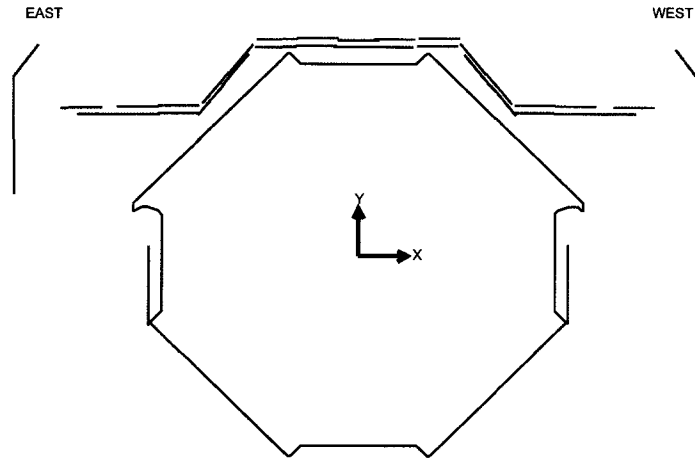
FIG. 2.9: A drawing of the locations of the scintillator modules in the Far Detector veto shield. Drawing is taken from [66].

2.4.2 Far Detector Readout

Signals in the Far Detector are read out by Hamamatsu 16-anode (M16) PMT's housed in light-tight steel boxes, with three PMT's per box. Clear fiber bundles channel photons from an interface with WLS fibers to the PMT's pixels. Each pixel records the optical sum of eight channels with a quantum efficiency > 13%. The channels read by an individual pixel are from geometrically distinct locations, allowing software to solve the eight-fold ambiguity of a single pixel by considering neighboring hits in the same event. The pixel-to-strip pattern is different on both ends of strip readout.

The event rate is low enough in the Far Detector that many channels can be allocated to a single high-speed ADC in order to reduce overall costs. Sixteen channels are summed on a front-end board that includes a charge-sensitive preamplifier for each channel, as

well as an output switch. Each channel can be switched to a read to a remote ADC, one at a time. A triggering system withholds digitization unless there are 2 signals from different PMT's in a 400 ns window before being sent to the ADC. This reduces overall dead time due to dark noise, light in in the PMT's resulting from background radioactivity and thermal emission. The ADC is 14-bit and reads out at a rate of 10 MHz and with a threshold of 0.3 photoelectrons [66].

2.4.3 Mapping strip-to-strip efficiency

Prior to detector assembly, modules were mapped with a γ source to record strip-to-strip differences in light output, and also to record dead channels from damaged or poorly-glued fibers. In the Far Detector about 0.16% of all 191k channels can be considered damaged (defined as outputting less light than 50% of the average strip light output) [66]. This is shown in Figure 2.10, to the left of the Gaussian fit to healthy fibers. Once installed, strip attenuation as a function of hit location along the strip can be measured with minimum-ionizing particles (MIPs) from cosmic rays, as shown in Figure 2.11. By summing the signal read out at both ends, the attenuation along the entire 8 m strip is only a 20% effect, instead of an 80% effect.

2.4.4 Data Aquisition

The Data Acquisition (DAQ) system consists of computers in both the Near and Far Detector laboratories which record the response from the detectors' ADC's. The two DAQ's are functionally identical small farms of PC's. The DAQ's process algorithms on-
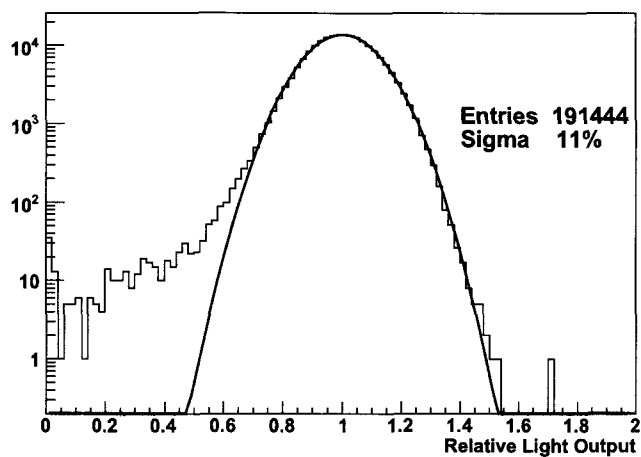
FIG. 2.10: The distribution of light output from all far detector strips for a 662 keV γ-source at the strip center [66]. Strips on the low side of the distribution are either damaged or broken.
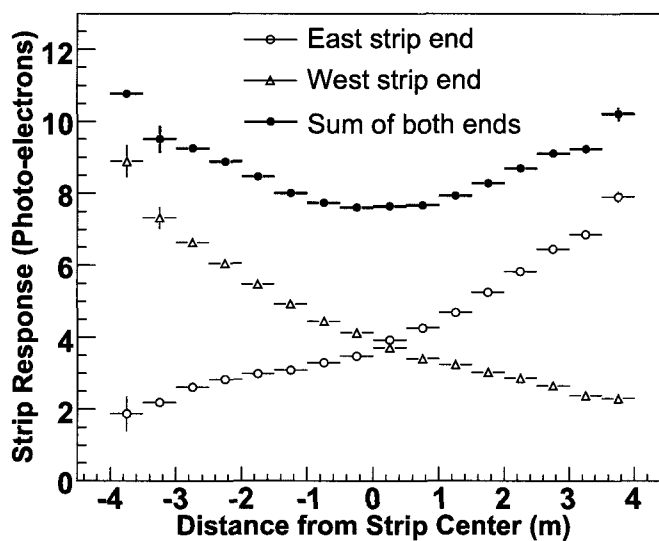
FIG. 2.11: Average light output from in-situ Far Detector strips as a function of distance from their center for normally-incident MIPs. The data shown are from stopping cosmic ray muons, for which containment criteria cause lower statistical precision at the ends of the strips [66].

line to select events of interest to be stored, and also to process calibration runs with the Light Injection (LI) system to record detector response for later offline calibration (the LI system will be discussed in Chapter 4).

There are a number of software triggers determining which signals the DAQs records. The first trigger is a timing trigger. In the Near Detector, the gate is opened when a beam spill trigger is signaled from the NuMI beam. The Near Detector records the timestamp of the spill via signals from the global-positioning-system (GPS). The GPS timestamp is then sent to the Far Detector via the internet to record the remote spill trigger. The Far Detector receives GPS signals from an antenna on the surface that passes through cables down to the laboratory. There is a 64 ns uncertainty on the timing at the Far Detector due to uncertainties in hardware delays [71]. Fake spill times are also generated at random intervals to sample detector activity and to record cosmic ray events [66]. Another trigger requires that four in five contiguous planes record at least one hit, and that there must be activity in at least twenty planes. In the Far Detector, an additional trigger requires at least 1500 ADC counts summed across five different planes, deposited in at least six hits [66]. One photoelectron corresponds to about 75 ADC counts in the Far Detector.

The DAQ transfers all data output to the Fermilab mass storage facility. The output data rate from the Near and Far Detectors is 20 kB/s and 10 kB/s, for trigger rates of 4 Hz and 30 Hz, respectively.

2.5 CALDET

Prior to the data-taking phase of MINOS, a smaller detector was placed in a test beam at CERN to calibrate the calorimetric sampling of the steel and scintillator configuration of the MINOS detectors. The Calibration Detector (CALDET) was constructed of 1 m × 1 m steel planes sandwiched with scintillator planes comprised of 1 cm thick scintillator strips. The steel absorber in CALDET was not magnetized. The read-out ends of the CALDET detector were designed to couple to two different sets of front-end electronics, which were identical to the electronics used in either the Near Detector or Far Detector. This detector was placed in a test beam at CERN exposed to e, π, and p beams of varying momenta. The CALDET calibration will be discussed in Chapter 4.

2.6 Magnetic Field

Both the Near and Far Detectors are magnetized to contain and measure muon tracks. The μ^- from ν_μ Charged-Current interactions are contained by the toroidal magnetic field in each detector so that a momentum measurement can be made of the muon's total ionization energy loss. If the muon can not be contained by the magnetic field or if it exits out of the back of the detector, then the muon track's curvature can give a measurement of muon's momentum. Each detector is magnetized by a current traveling through the center of the detector, and both are magnetized to similar field strength, close to the point of magnetic saturation in steel. The details of the magnetization for each detector is described in this section.

2.6.1 Field Strength Modeling

The steel in the MINOS detectors is an alloy of iron, carbon, and other trace metals. This alloy is ferromagnetic. The current-carrying coil running through the center of each detector induces a magnetic field within the steel, in accordance with the steel's permeability and geometry.

The magnetic field within the steel is related to the applied field and the properties of the steel. From a de-magnetized state, driving the current increases H (in units of Amperes/meter), which increases the net magnetic moment per unit volume M(also in units of Amperes/meter), which in turn increases the magnetic field B (in units of Tesla) within the steel. Ferromagnets are non-linear media, so both M and the permeability μ are functions of H.

$$B = \mu(H)(H + M(H)) \tag{2.2}$$

For a small increase in H, there is a large increase in M, such that H is negligible and $B \approx M$, up to a point. Saturation occurs when $\mu(H)$ becomes constant above some value of H. At saturation, increasing H does not increase the overall magnetization, M. On a microscopic level, at the saturation point all of the magnetic domains are already aligned with the magnetic field, and so no increase in M is possible. The relation between the applied H and the induced B is called a *hysteresis loop*, or a B-H curve.

The magnetic field maps used in MINOS simulations and reconstruction are important to the overall systematic error. Since the magnetic field is within the steel, it is difficult to measure directly (say, with a Hall probe). Instead, steel is tested for its B-H

curve and modeled with software.

The magnetic field strength of the detectors is calculated with a Finite Element Analysis (FEA) performed with ANSYS software [72]. The FEA models the shape of the steel plane, with the apropriate current running through the magnetizing coil, and calculates the magnetic field strength and direction at a particular localized element in the presence of all of the other elements surrounding it. In the Near Detector, the ANSYS geometric model is simple because Near Detector planes are solid, but in the Far Detector the model must include the small airgaps between the eight 0.5 in.-thick planks that are arranged to create the 1 in.-thick plane (see Figure 2.8).

Prior to 2008, the magnetic field maps used by MINOS had been generated using a sample of steel produced at the foundry prior to the main production run generating the actual steel in the MINOS detectors. Due to concerns about these field maps, a new study was comissioned to generate new field maps using steel that had been cut from heats of actual MINOS steel.

2.6.2   Field map generation and validation

Six samples of the steel used in both detectors were chosen to have their B-H curves measured. Five of these were from different steel heats spanning the 39-heat production run of the MINOS steel. One sample was a duplicate sample from that set of five samples, in order to gauge the reliability of the measured curves.

These samples were cut down to small rings with inner diameter 42.4 mm, outer diameter 50.8 mm, and width 6.35 mm. The rings were cut with a wire-EDM process to
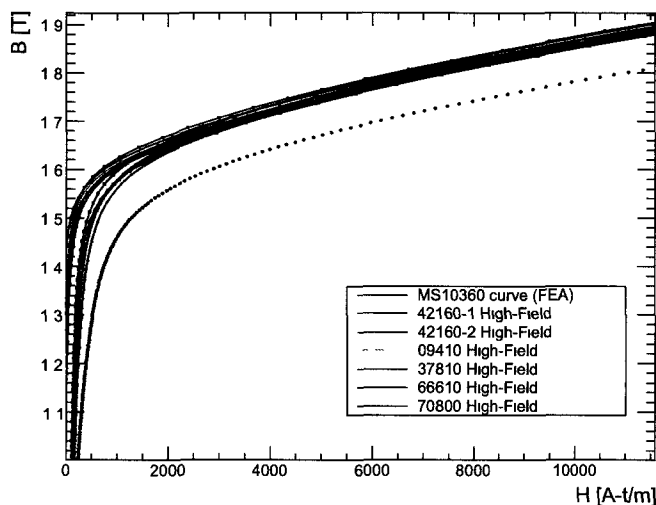
FIG. 2.12: B-H curves for five steel samples of steel used in the construction of the Near Detector and the Far Detector. The MS10360 curve was from a sample of pre-production steel not used in the actual detectors.

reduce the possibility of work-hardening the material and altering its magnetic properties. These rings were wound with primary (H) and secondary (B) turns of copper wire and connected to a KJS Associates SMT-600-5 Computer Automated Soft Magnetic Hysteresigraph System, which measured the B-H curves of these steel samples by magnetic induction in accordance with ASTM A773. These B-H curves are shown in Figure 2.12 and are directly compared to older steel in Figure 2.13.

The median B-H curve of the 5 samples that were measured was fed to the FEA as the B-H curve for a generic element within the steel. The newly generated field maps are shown in Figures 2.15 and 2.16. Compared to the old field maps generated with pre-production steel, the Near Detector magnetic field strength is 4% larger on average, but 12% larger in the 1 m-radius cylinder that is the fiducial region. The Far Detector magnetic field strength is 11% larger on average, 12% larger in the 3.74 m-radius cylinder
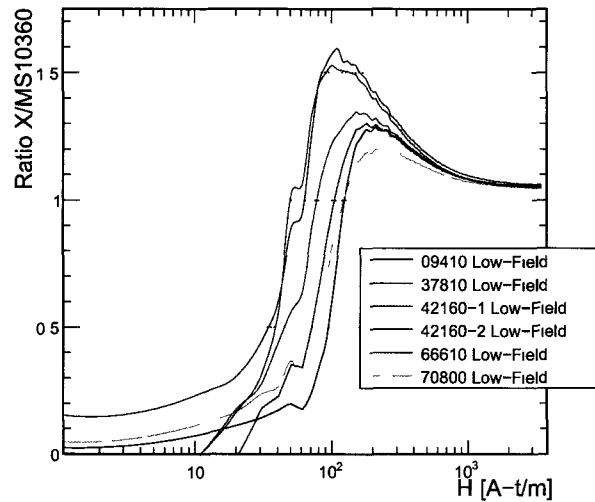
FIG. 2.13: The dispersion of the measured BH curves relative to that of the pre-production steel sample. All new samples had been degaussed to a higher degree than the old sample, evident at H, and saturated at a 4% larger field at high H.

that is the fiducial region.

Measuring the magnetic field of the detectors *in situ* is one method of validating the generated magnetic field maps. The BDOT system was included in the initial designs of the detectors in order to make this measurement. All of the planes in the Near and Far Detectors have 50 turns of wire looped between the coil hole and an edge of the detector. Ramping the magnetizing current from 0 A to full power (40 kA-turns for the Near Detector, 15.2 kA-turns for the Far Detector SM's) induces a current in this loop of wire by magnetic induction. Loops are oriented at eight different angles in certain places, with respect to the face of the detector, to measure the azimuthal symmetry of the magnetic field. As a part of the BDOT system, the two ends of each loop were connected to an ADC channel and recorded with a LabView-based DAQ. The intent of this system was
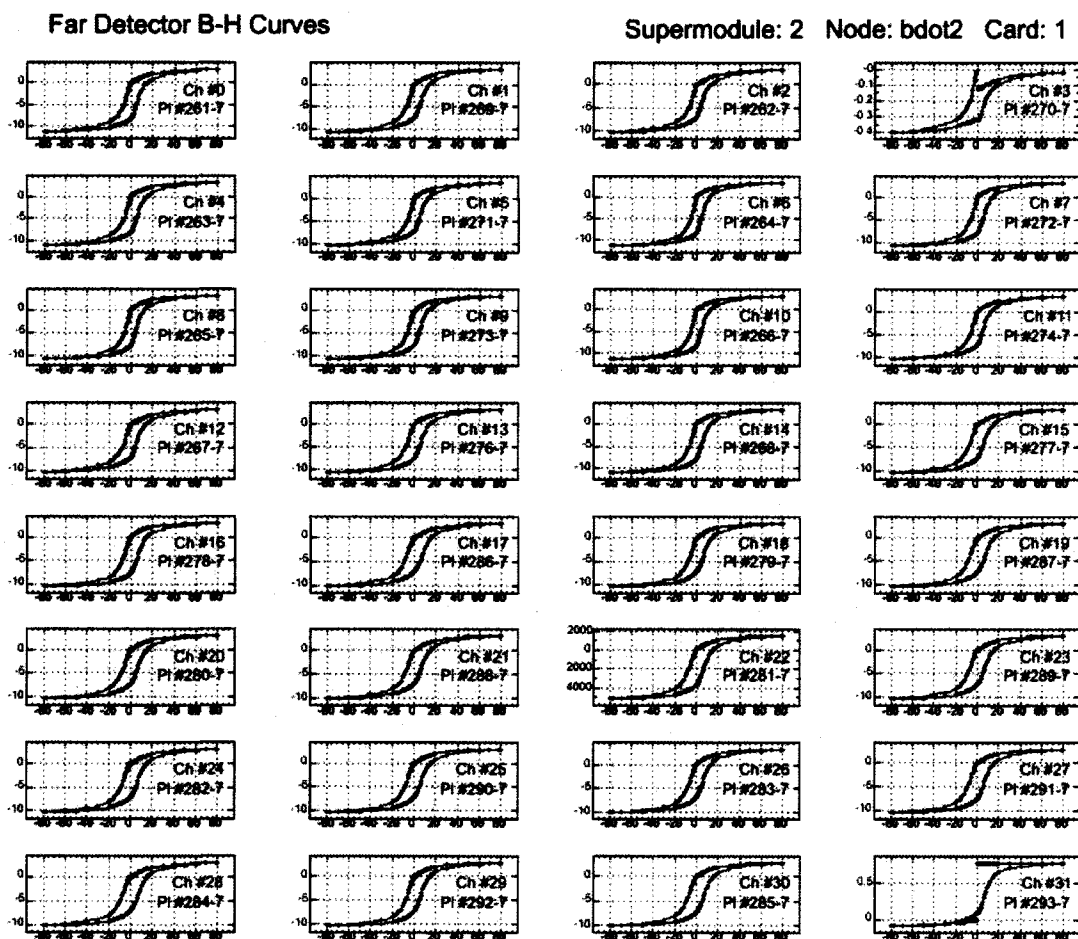
FIG. 2.14: B-H curves measured and recorded with the BDOT system for various planes and coil positions in SM2 of the Far Detector.

to measure the B-H curves of several planes and sampling several regions per plane. An example of the B-H curves measured from 32 different locations is shown in Figure 2.14.

An analog method for measuring B-H curves was used to validate the BDOT system. The ends of the BDOT loop were connected to a precision "magnetic integrator," which charges a capacitor with the current induced by the magnetization of the detector. The voltage across the capacitor was measured at several stages in the magnetization of the

detector to give a coarse measurement of the B-H curve of the plane being measured.

The charge collected on the magnetic integrator's capacitor is proportional to the integral of the magnetic flux moving through the BDOT loop. This can be compared to the field map produced by the ANSYS model by integrating the value of the magnetic flux density perpendicular to a slice of the detector in the position of the BDOT loop.

The analog magnetic integrator was used to test 15 planes in the Near Detector. The analog measurements agreed with the newly generated Near Detector field maps to within $\pm 1.7\%$. The magnetic integrator was also used to test 24 configurations in the Far Detector, on 15 different planes. Every plane that was tested was measured in the same BDOT loop orientation, for comparison. The analog measurements agreed with the generated Far Detector field maps also to within $\pm 1.7\%$ [73].

The field maps were further verified by comparing muon momentum calculated from range (ionization energy loss) and curvature for muons stopping in the Near Detector, in both data and simulations. With old field maps generated from pre-production steel properties, the momentum of stopping muons (with $E < 6\,\mathrm{GeV}$) is calculated from both range and curvature. The double ratio of $(P_{\mathrm{range}}/P_{\mathrm{curvature}})^{Data}/(P_{\mathrm{range}}/P_{\mathrm{curvature}})^{MC}$ for muons with was found to be 0.95 in simulations generated and reconstructed with old field maps. This 5% difference disappeared when the magnetic field strength was scaled up uniformly by 13% in detector simulations [66, 74]. This scale factor test is only a toy, due to the non-linear relationship between H and B, especially in saturated regions.

The same study was re-run with the simulations generated and reconstructed with the new fieldmaps. The double ratio $(P_{\mathrm{range}}/P_{\mathrm{curvature}})^{Data}/(P_{\mathrm{range}}/P_{\mathrm{curvature}})^{MC}$ improved

to 1.01. This means that the corrected uncertainty on muon momentum from curvature is 1%, relative to the uncertainty on muon range, which is 2%. The total uncertainty on muon momentum energy scale is taken to be the fully-correlated sum of these two uncertainties. The resulting 3% error is a great improvement over the previous 7% error [74].

The triumph of this validation scheme is that two independent methods were used to determine the 13% offset of the magnetic field strength necessary to bring the magnetic field map in agreement with the detector steel. These efforts reduced one of the largest systematic errors in the MINOS analysis by more than 50%.

2.6.3 Field maps

In the Near Detector, the current-carrying coil is carried through a square hole offset 55.8 cm from the center of each plane. The ND coil consists of eight turns, with the coil return at 45° on the shorter side of the plane. The coil is continuously water cooled throughout the length of the detector. The coil carries a 40 kA-turn current to magnetize the detector. The average magnetic field of the fiducial region is 1.286 T. The magnetic field map for a generic ND plane is shown in Figure 2.15.

The two SM of the Far Detector are magnetized separately. The magnetic field map for a generic FD plane located in the middle of a supermodule is shown in Figure 2.16.
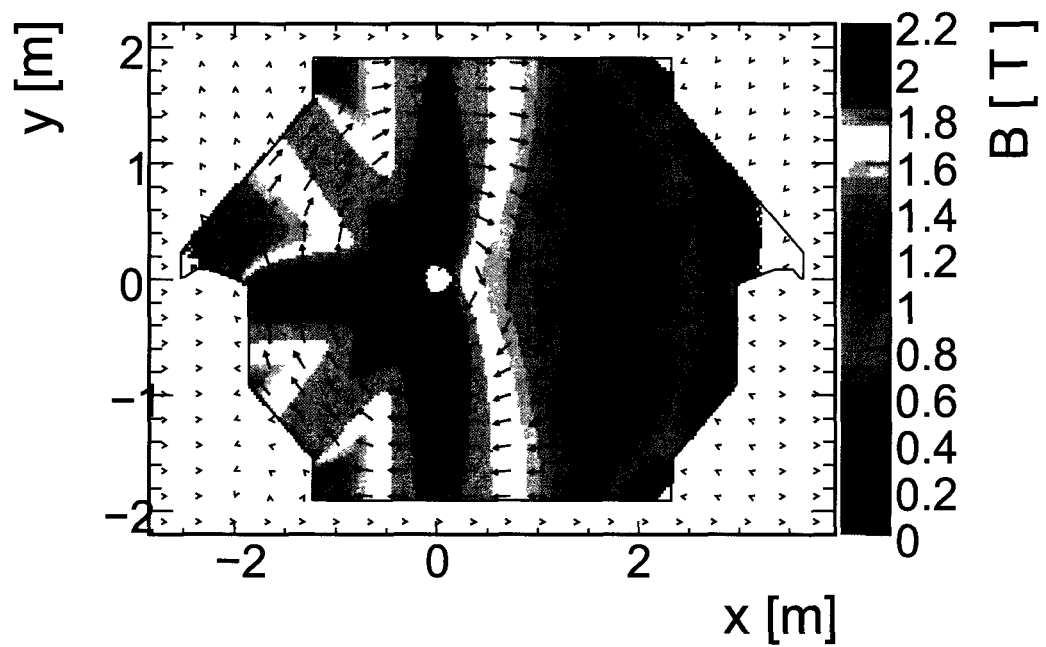
FIG. 2.15: FEA model of the Near Detector magnetic field for a generic interior plane after measurement of the magnetic properties of detector steel. Values below 0.1 T are omitted from this map. Arrows indicate the direction of the magnetic field B, and the color scale indicates magnetic field strength ($|B|$).
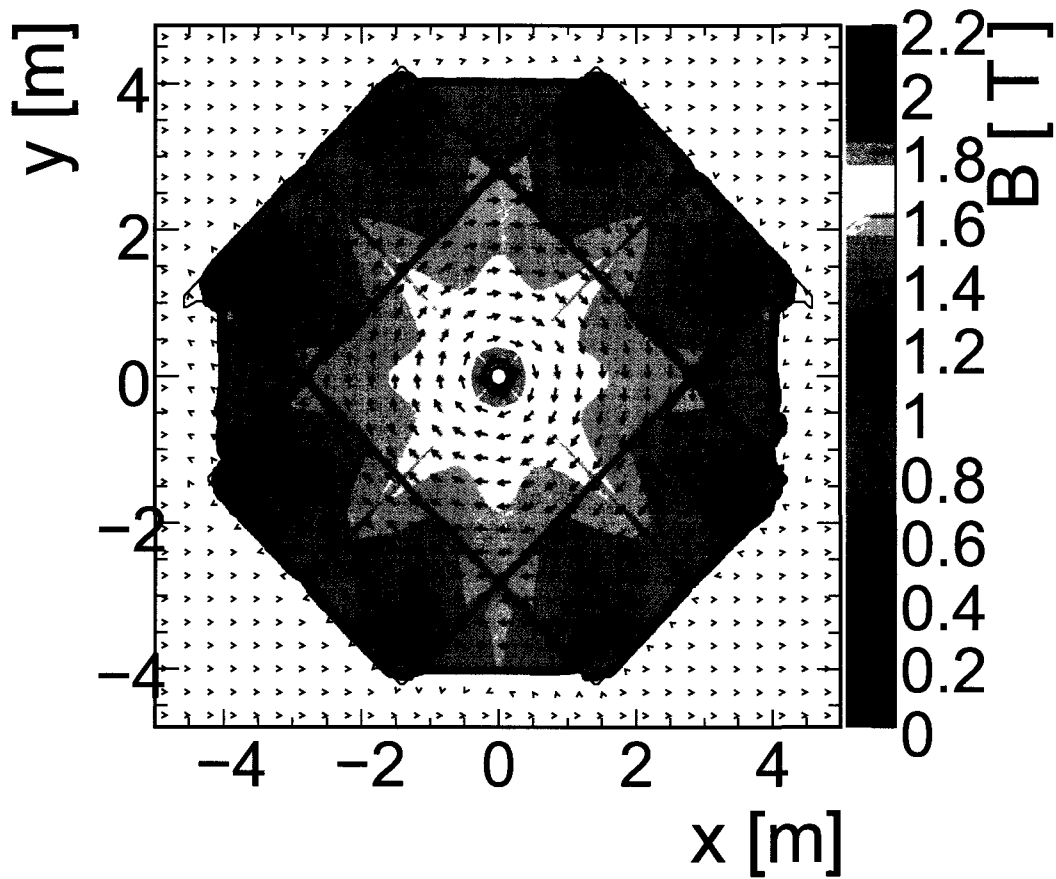
FIG. 2.16: FEA model of the Far Detector magnetic field strength for a generic interior plane after measurement of the magnetic properties of detector steel. Values below 0.1 T are omitted from this map. Arrows indicate the direction of the magnetic field B, and the color scale indicates magnetic field strength ($|B|$).

2.6.4   End Effects

The plane model used to determine the generic magnetic field consisted of a single plane with an infinitely long current-carrying wire passing through the coil hole, to model a plane deep inside the detector or supermodule. For the planes of steel nearest to the ends of the detector, there is an additional contribution to the magnetic field from the coil return arms. These planes will have different magnetic field maps which must be added to simulations and reconstruction software.

An ANSYS FEA model was created with 15 planes exposed to the current from the magnetizing coil making a right angle to follow the return arm along the face of the outermost plane. Separate models were generated for each detector. This model creates 30 separate magnetic field maps for the Near Detector and 30 field maps for a generic Far Detector SM. An example is shown in Figure 2.17 for an earlier model with only 12 planes in the Far Detector.

The file size of 15 magnetic field maps is prohibitively large for use in simulations and reconstruction. An piecewise-linear interpolation scheme including the end plane, third plane, and the nominal interior plane was found to accurately model the end effects of all 15 planes with small error. The residual RMS field errors were less than 5 gauss for intermediate planes between the simulated end planes and the interpolation scheme [75]. The interpolated set of magnetic field maps are included in the simulation and reconstruction software packages at the upstream and downstream ends of the Near Detector and both supermodules of the Far Detector.
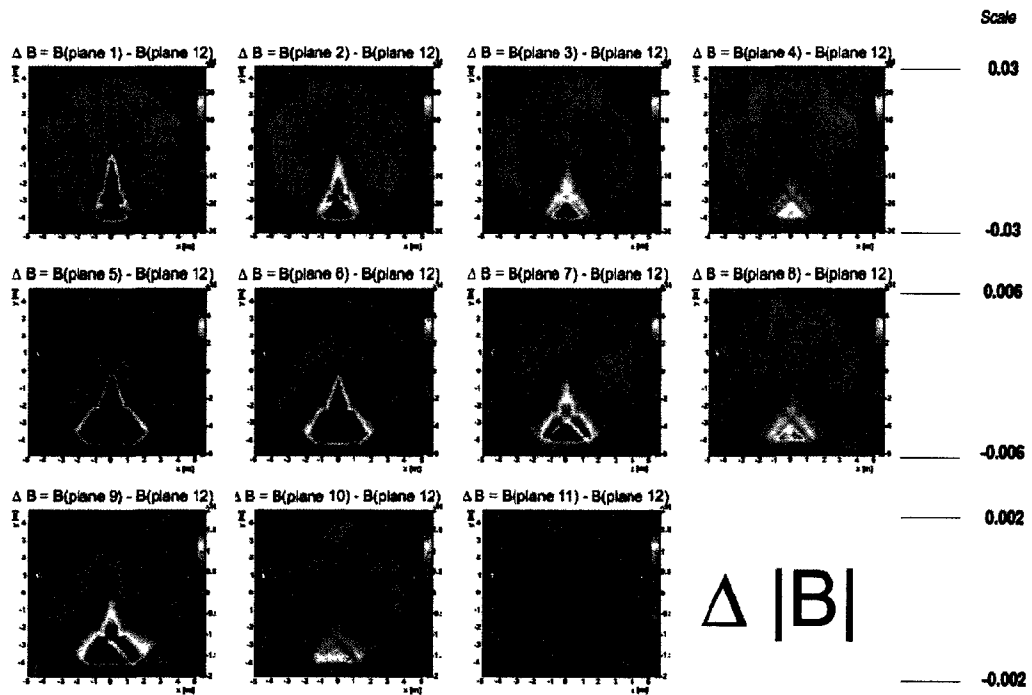
FIG. 2.17: The effect of the return arm of the current-carrying coil on the magnetic field map of the end of a Far Detector supermodule, expressed as the $\Delta|B| = |B_{\text{plane}}| - |B_{\text{nominal}}|$. The top left field map is for the outermost plane, and the bottom right field map is for the innermost plane.

CHAPTER 3

Monte Carlo Simulations

Many of the processes involved in particle physics are probabilistic. Given the low event rate and overall yield expected in both detectors, simulations are needed to validate our understanding of the physical processes and the properties of the beam and detectors. Simulations of this nature are performed with the Monte Carlo method.

The Monte Carlo method, also called the Metropolis method, dates back to calculations made within the context of the Manhattan Project [76], and refers to a class of statistical computational analyses. Within particle physics, Monte Carlo is used to describe simulations of ensembles of interactions, each with its own probability density function (PDF), which are repeated a large number of times in order to sample the phase space available to the process in question. The method relies on random numbers, which are "thrown," typically with a normal distribution between 0 and 1. This is best illustrated with an example, which is relevant to the simulations described in the rest of this chapter.

Imagine a 120 GeV proton incident on a 90 cm long piece of graphite. The proton

has a small probability of interacting within the first 1 cm of the graphite piece, and that interaction has a small probability of producing some number of hadrons. The Monte Carlo simulation steps the proton through the piece of graphite, "throwing" a random number and comparing it to the PDF of interaction for each step. If the PDF indicates that a particular interaction process has a 10% chance of occurring, and the random number thrown is below 0.1 (which, for a normal distribution between 0 and 1 would happen 10% of the time), the interaction process is said to have occurred and the outgoing hadrons are tabulated. Repeat this simulation many millions of times and you will have tabulated outgoing hadrons due to interactions throughout the length of the target, with roughly $1/e$ of them occurring within the first interaction length. This is the procedure to produce a simulated primary flux of hadrons emanating from the NuMI target.

Many physics models exist for the simulation of high energy physics experiments. These packages range from libraries with a Unix philosophy ("Do one thing, do it well") to fully featured suites able to take into account multiple models, materials, and detector geometries. Four main packages are used in the simulations in this thesis: ROOT [77], NEUGEN [18], FLUKA [78], and GEANT [79].

3.1 William & Mary Farms

The High-Energy Physics group at the College of William and Mary has two computing clusters that were utilized for producing Monte Carlo simulations. These clusters, at one time or another, were used to simulate every step in the modeling of the MINOS experiment. Their specifications are described here.

3.1.1   Nova cluster

The Nova cluster consists of eight Dell PowerEdge 1750 dual core computing nodes, one disk server, and one PowerEdge 2650 head node with a 3 GHz CPU. All of the computers run Scientific Linux 3.3. This cluster was used to generate much of the target simulation and hadron transport simulations eventually used in [47].

3.1.2   Zaphod cluster

The Zaphod cluster consists of 108 Dell PowerEdge 1950 computing nodes, two large 13 TB disk servers, and two head nodes (hence, Zaphod), which are PowerEdge 2950 2.5 GHz, with 16 GB RAM. Each computing node has a dual quad-core 2.5 GHz CPU and 8 GB RAM. The cluster is shared between the W&M Experimental High Energy Physics group and the Lattice QCD group. This cluster was used to simulate events in the Near and Far Detectors. This cluster was used extensively for this purpose, logging 500k+ CPU-hours in service to the MINOS collaboration in 2008.

3.2   NuMI Flux

Before we are able to simulate neutrino events within the MINOS detectors, we must first predict the flux of neutrinos coming from the beam. The flux of hadrons coming off the target is simulated first, then the pions and kaons are allowed to decay in the decay pipe, producing neutrinos.

3.2.1 Target

The first piece of the NuMI beam to be simulated is the interaction of 120 GeV protons from the Main Injector interacting within the NuMI target. The target hadronization simulation is performed with the FLUKA simulation package. The simulation outputs the hadron multiplicity and their 4-vectors immediately off the graphite target, and also simulates re-interaction of π and K within the length of the target. It also includes the apparatus surrounding the target, including the beryllium windows, the cooling lines, and the helium gas filling the target volume. The FLUKA software package is updated frequently with bug fixes and improved experimental constraints. The FLUKA05 version of the software package was used for past MINOS analyses [19], but the FLUKA08 version produced the target hadron flux for the analysis described in this thesis.

3.2.2 Decay Pipe

The output of the FLUKA simulation is fed to FLUGG [80], which transports the π and K through the focusing horns and decay pipe, and also simulates their decay. FLUGG is a modified version of FLUKA which combines the physics of the FLUKA libraries with the geometry of GEANT4. GEANT allows for easier configuration of the complex geometries involved in the NuMI beamline downstream of the target, while FLUKA is more trusted with accurate hadronic interaction modeling.

The hadronic flux from FLUGG was extensively validated against past simulations, produced with GEANT3, and real data from the Near Detector, as well as the muon and hadron monitors downstream of the decay pipe [80]. The differences between FLUGG

and the older simulations were substantial, so the validation worked backwards to try and reproduce the FLUKA results by removing newer components. This validation study was able to reproduce the older GEANT3 flux by removing the updated geometry and hadronic modeling.

3.3 Detector

Upstream of the detector, there are only a few instruments available to monitor the real beam for comparison and validation of the simulations that are used. The two MINOS detectors, however, contain hundreds of thousands of channels and provide more opportunity to study the neutrino flux produced by the beam. The downside is, with so many channels, the detectors are complicated beasts which require complicated simulations. The neutrino interaction itself is a probabilistic process, as is the intranuclear rescattering within the struck target nucleus, multiple scattering of the exiting muon, and the interactions of any electromagnetic or hadronic showers produced in the neutrino interaction.

3.3.1 Neutrino interactions

The neutrino interactions within the MINOS detectors are simulated with a custom set of libraries called GMINOS. GMINOS generates the neutrino interactions with NEU-GEN 3.5.5 neutrino interaction model [18]. Simulating an event yield requires an estimate of not only the flux, but also the neutrino cross-section. MINOS uses a model within NEUGEN that is a modified combination of Bodek-Yang [81] and Rein-Seghal [82] models. NEUGEN also accounts for intranuclear resecattering of secondary hadrons with

INTRANUKE [83].

Hadronic and secondary showers are generated with the GEANT3 library GCALOR. The muon (if the interaction is CC) and the shower hadrons are propegated through the volume of the detector with GEANT3.

3.3.2 Detector simulation

Truth hits in the detector from GMINOS are read by PhotonTransport, a simulation of the scintillator. PhotonTransport generates photons in the scintillator and models absorption, re-emission, and transport through the WLS fibers to the PMT readout end. The photon signal in each strip is multiplied by the inverse of the strip's calibration constant from a random time within the data taking period. This is so that the calibration procedure, which is applied to both data and simulated detector readout, returns correctly-calculated simulated signals, distributed over the entire run. The calibration procedure will be described in Chapter 4.

From there, DetSim, a simulation of the readout electronics, takes the transported photons and simulates the PMT photoelectron amplification, ADC digitization, and triggering. The final simulated output has the same format as true raw data and can be calibrated and reconstructed with the same software, reducing the possibility of bias.

The entire simulation process aims to achieve good data/Monte Carlo model agreement, so as to be useful for understanding the changes in the data expected for various changes in the underlying physics of the detector, beam, and neutrino interactions. Not only does the detector simulation need to reflect the true behavior of the physical detector,

but it needs to be fed a simulated neutrino flux that reflects the physical neutrino flux as closely as possible to produce the correct kinematic variable distributions.

3.3.3 Monte Carlo version

Each of the steps of the full simulation has undergone many revisions and upgrades. The versions of simulation software are assigned codenames after vegetables, in alphabetical order. The GMINOS version Daikon represents a larger rewrite of GMINOS over the previous version (Carrot), which is the simulation version used in [46]. The major features of Daikon07 include production with updated FLUGG flux files, updates to NEUGEN, and the updated magnetic field maps used to propagate muons (introduced in Daikon03).

3.4 Beam Tuning

Upstream of the simulation of neutrinos within the detector volumes, neutrinos are simulated as daughter particles in the two-body decay of $\pi^\pm$ and $K^\pm$ in the decay pipe. Simulating the right population of secondary pions and kaons in the decay pipe requires knowledge of the pions and kaons as they come off the target and are focused or defocused. Model uncertainties in the software stream generate a significant uncertainty in the overall neutrino flux expectation. Empirically, comparison between Near Detector neutrino data and raw Near Detector simulations display substantial differences in the higher-energy edge of the focusing peak in the Low Energy beam configuration. In other beam configurations, the data and simulations agree in the same energy range, indicating that the discrepancy is not due to mis-modeling of detector acceptance or neutrino cross-

sections. A beam fitting process is constructed to constrain the Monte Carlo simulations and produce a more accurate neutrino flux prediction by utilizing the Near Detector data. The beam tuning minimizes the effect of the beam model uncertainties on the final oscillation analysis. The correlations between remaining systematic uncertainties are also tabulated, to produce a single 1σ error band for all beam systematics. For a full discussion of the beam tuning procedure, see [84] and [69].

The Near Detector is used to measure the neutrino flux from many different beam configurations. The configurable beam was discussed in Chapter 2, and the beams and integrated exposures accumulated were shown in Table 2.1. These beam configurations were also simulated as described in this chapter, and a multi-variable fit of the simulated flux was constructed to achieve good agreement with the observed Near Detector data. Penalty terms are constructed to constrain the π^+/π^- ratio to both FLUKA05 simulated results and experimental results from the NA49 experiment [85]. This fit is simultaneous across all beam configurations and separate runs. The best-fit values of the model parameters are used to assign an importance weight to each simulated pion, which is propagated to the Near and Far Detectors.

The hadron production off of the target is the most important factor in the beam tuning procedure. There is little experimental data to constrain the hadron production models for Monte Carlo simulations at NuMI proton energies and with thick graphite targets. Experiments such as NA49 have data for only thin targets [85], where incident protons pass through less than one interaction length of material. An experiment at Fermilab, Main Injector Particle Production (MIPP) [86] did take data with a NuMI target, but the data is

not yet available for validation.

Sixteen parameters for tuning the target hadron production come from the BMPT parameterization [87]. The BMPT parameterization is:

$$\frac{d^2N}{dx_F dp_T} = [A + Bp_T]\,e^{-Cp_T^{3/2}}$$

(3.1)

The functions $A(x_F)$, $B(x_F)$, and $C(x_F)$ are themselves warped linearly with a total of six parameters for π^+ and six for K^+ to produce the importance weight W for ν_μ. The warping for pions is

$$A'(x_F) = (par[0] + par[1]x_F)A(x_F)$$

(3.2)

$$B'(x_F) = (par[2] + par[3]x_F)B(x_F)$$

(3.3)

$$C'(x_F) = (par[4] + par[5]x_F)C(x_F)$$

(3.4)

and likewise for kaons and parameters $par[6]$ through $par[11]$. The importance weight for positive π/K is defined as

$$W(\pi^+/K^+, p_T, p_Z) \equiv \frac{[A' + B'p_T]exp(-C'p_T^{3/2})}{[A + Bp_T]exp(-Cp_T^{3/2})}$$

(3.5)

There are two additional parameters to define a linear correlation between ν_μ weights and $\bar{\nu}_\mu$ weights from π^-, as a function of x_F, and two more for K^-. This brings the total to sixteen parameters for tuning the target hadron production.

There are eight other parameters included in the fit to account for uncertainties in
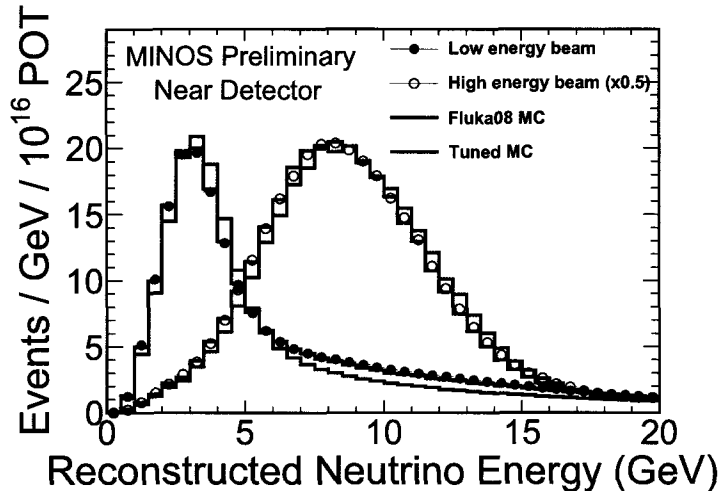
FIG. 3.1: The effect of the beam tuning on the Near Detector energy spectrum for the two beam configurations used in the neutrino oscillation analysis described in this thesis. FLUKA08 produces a raw ν_μ flux simulation, which is reconstructed in the Near Detector. Data from many beam configurations are used to produce the tuned beams (red). The tuned beam shows better agreement with the measured Near Detector data (circles).

the flux measured in the Near Detector, including focusing and target degradation. The 24-parameter fit is performed using MINUIT minimization software [88]. The agreement achieved between data and tuned simulations is shown in Figure 3.1.

3.4.1 Beam tuning error

The 16 tuning parameters for hadron production encapsulate the total beam modeling uncertainty. The remaining eight nuisance parameters reflect uncertainties in focusing, target degradation, and detector background effects. Below follows a brief description of the uncertainties and their origins:

- **Focusing** - The error on focusing is due to horn mis-alignment and the current distribution in the horns due to the skin depth of the current in the horns. There are two

parameters that parameterize focusing effects.

- **Neutrino energy scale parameter** - A single parameter affecting the fully-correlated track and shower energy scale. A 1σ shift corresponds to a 5% shift in neutrino energy scale.

- **Target decay** - The second NuMI target was in place for Runs II and III, and likely eroded due to the extreme radiation it was exposed to. Target degradation is modeled by comparing nominal NuMI simulations with simulations where the 7^{th} and 8^{th} fins have been removed. The locations of fins 7 and 8 correspond to the location of maximum shower energy. Two parameters, one for each of Runs II and III, account for a linear interpolation between simulations with and without these two succeptible fins.

- **NC contamination** - The flux × cross-section measurement for CC neutrino interactions suffers from contamination from CC-like NC events. There is a 30% 1σ error on the size of the NC background expectation for both ν_μ and $\bar{\nu}_\mu$ events in the Near Detector. This is considered separately from the NC contamination in the oscillation analysis, described in Chapter 7.

- **$\nu_\mu/\bar{\nu}_\mu$ Cross-section** - A 30% error on $\nu_\mu/\bar{\nu}_\mu$ cross-section ratio is allowed below 25 GeV at 1σ to account for poor worldwide data constraints.

Earlier efforts to tune the NuMI flux used nuisance parameters reflecting the uncertainties of the exact location of the target with respect to the focusing horns. The tuning described here uses new survey data to fix the target position parameters [69].

(a) Near Beam Error Band

(b) Far Beam Error Band

(c) Far/Near beam error
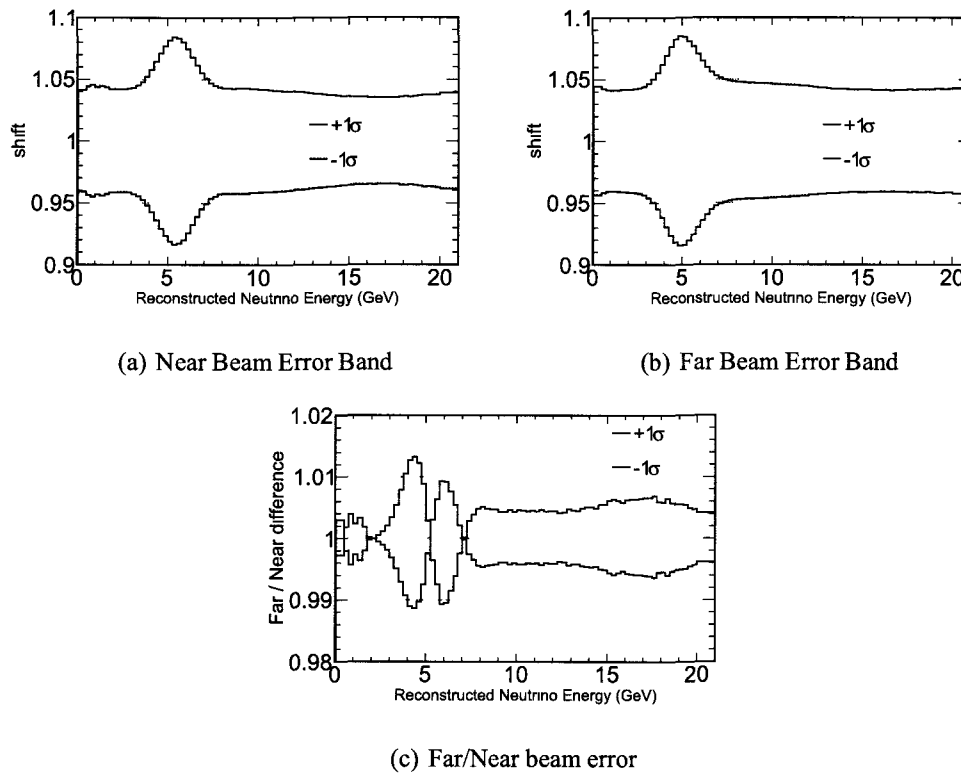
FIG. 3.2: The $\pm1\sigma$ error bands due to the beam fitting procedure, as a function of reconstructed neutrino energy in both the Near and Far Detectors. The peak is largest for the high falling edge of the focusing peak. Below, the ratio of these errors, showing the effect of the extrapolation, which is described in Chapter 6.

Fitting all twenty-four parameters yields best-fit values as well as 1σ errors for each parameter. It is unnecessary to report values for all of these systematic errors separately, so instead we use the correlations between all of these errors produce a single $\pm1\sigma$ error band for the neutrino flux. The total $\pm1\sigma$ error band is shown in Figure 3.2.

3.5 Conclusion

The MINOS experiment relies heavily on beam and detector simulations to under-
stand the raw data recorded in both the Near and Far Detectors. The hadronic production,
hadronic decay, and detector response are all generated with separate Monte Carlo sim-
ulation programs, using models to predict the outcome of many probabilistic processes.
The simulated neutrino flux is tuned with a 24-parameter fit which incorporates seven
different beam configurations and π^+/π^- data from the NA49 experiment. This tuning
method is robust, and the total residual systematic error associated with beam modeling
is small.

CHAPTER 4

Detector modeling, calibration, and

data reconstruction

The NuMI beam creates muon neutrinos which interact in the Near and Far detectors in either CC or NC interactions. The detectors need to be able to provide enough information about the interactions that the energy of the interacting neutrino can be reconstructed. This chapter describes the process of converting raw signals in both detectors into neutrino energy spectra so that the experiment can measure neutrino oscillation parameters.

4.1 Signal

The raw data recorded by the MINOS detectors consists of pulse heights recorded by PMT's with nanosecond timing. Moving from these hits to an oscillation measurement requires several steps, beginning with the calibration of individual hits, then associating

94

the hits into individual events and reconstructing all relevant physics quantities. Monte Carlo simulations are constructed to behave like real data, but contain the truth information about their initial generation. This allows for adjustments to be made in the simulated data in order to achieve better overall agreement between data and simulations. With data and simulations in the same format, they can both be processed in the manner described in this chapter.

4.2 Muon Tracks

Relativistic heavy charged particles passing through matter lose energy by ionizing the surrounding material and by exciting atomic nuclei. The rate of energy loss is given by Groom, *et al*, which updates and tabulates energy loss for muons [89]. These tables include updates to the classic Bethe-Bloch theory of ionization energy loss [20].

The Bethe-Bloch equation appears in several forms using approximations appropriate for certain conditions. The full expression for the rate of energy loss is

$$-\frac{dE}{dx} = Kz^2\frac{Z}{A}\frac{1}{\beta^2}\left[\frac{1}{2}\ln\frac{2m_ec^2\beta^2\gamma^2T_{max}}{I^2} - \beta^2 - \frac{\delta(\beta\gamma)}{2}\right] \quad (4.1)$$

where

$$T_{max} = \frac{2m_ec^2\beta^2\gamma^2}{1 + 2\gamma m_e/M + (m_e/M)^2} \quad (4.2)$$

is the maximum kinetic energy that can be imparted to a free electron in a single collision, z is the charge of the incident particle (in units of electron charge), I is the mean excitation energy of the medium, Z and A are the atomic mass and atomic number of the medium,

m_e is the mass of the electron, and M is the mass of the incident particle. K is a collection of several constants: $K/A = 4\pi N_A r_e^2 m_e c^2 \approx 0.307 \, \text{MeVg}^{-1}\text{cm}^2$ for $A = 1$, where r_e is the charge radius of the electron and N_A is Avogadro's number.

The $\delta(\beta\gamma)$ term is the density effect correction to ionization energy loss, which can be neglected below $\beta\gamma \approx 100$. Other radiative processes can be ignored in the energy range relevant to MINOS muons, pions, and protons, such as bremsstrahlung and e^+e^- pair production, which only contribute significantly to muon energy loss when the muon energy is $\approx 400 \, \text{GeV}$ [20].

The energy loss as a function of $\beta\gamma$ reaches a minimum between $2 < \beta\gamma < 4$ and plateaus at higher energies with $\frac{dE}{dx}$ only slightly higher than $\frac{dE}{dx}\big|_{min}$. Particles with $\beta\gamma$ near or slightly above the minimum energy loss are called Minimum Ionizing Particles, or MIPs. The muons produced in CC interactions with NuMI neutrinos are relativistic and typically minimum-ionizing. Once particles lose enough energy such that they are below minimum-ionizing, they then lose more energy per Δx traversed.

Muons in the MINOS detectors, then, are detected when they ionize within scintillator strips. Scintillator planes are sandwiched between planes of steel absorber, so to first order the energy of the muon can be determined by simply counting the number of planes the muon passes through before stopping, using the expected stopping distance from Groom. This approximation would be exact only in the case of normally incident muons whose path length in steel is L ($L = nd_0$, where n is the number of planes traversed and d_0 is the average thickness of the steel planes). In reality, muons can carry some transverse momentum from the vertex of a CC interaction, and the curvature of the

muon's path in the detectors' magnetic field yields path-lengths per plane of steel $d > d_0$. The momentum of a muon which stops within the detector has good energy resolution from dE/dx in the steel ($\sigma_{E_{trk}} = 2\%$), with small source of systematic error arising from the uncertainty of the total path length of the muon within the steel planes and the uncertainty of the thicknesses of each plane.

Some muons observed in the detectors exit the instrumented region before stopping, making a calculation of its energy loss from Bethe-Bloch impossible from its total path length. For these muons, the track curvature is used to estimate the muon momentum. The momentum from track curvature is determined with a Kalman filtering technique [90, 91]. The Kalman filter constructs predictions for five parameters from strip-to-strip within the detector and checks those predictions against the next hit. The five variables are u,v, $\frac{du}{dz}$, $\frac{dv}{dz}$, and q/p, where u and v and z are the spatial coordinates relative to the orthogonal scintillator planes, q is the charge of the lepton being tracked (either +1 or -1), and p is its momentum. The first four variables are known from point to point, and last variable is the product of the filter. The muon track is tracked with the Kalman filter through its entire length, taking both the ionization energy loss and magnetic field into account as it makes predictions. As a muon passes through steel and loses energy, the radius of curvature of the track in a constant magnetic field changes, and is predicted by the Kalman filter.

The Kalman filter also produces an error matrix based on the accuracy of its predictions, with an estimate of its uncertainty, $\sigma_{q/p}$. A significant contribution to the Kalman error is the effect of multiple scattering of the muons off of nuclei in the steel. These small random perturbations make the momentum measurement from curvature less precise than

the range-based measurement.

4.3 Electromagnetic and Hadronic Showers

Neutrino-induced hadronic showers are not resolved on a strip-by-strip, particle-by-particle basis in the MINOS detectors. Instead, calorimetric techniques are used to measure shower energy. The calibration of calorimetric response to reconstructed shower energy depends on many things, including the type of shower. The full calibration algorithm is summarized later in this chapter.

4.3.1 Hadronic Showers

The momentum transfer to a nucleus from either a CC or NC interaction is often enough to produce a hadronic shower large enough to be visible in the MINOS detectors. The charged secondaries from the interaction (i.e. n, p, π, or more exotic particles) lose energy through Bethe-Bloch processes, but also interact strongly with nuclei, which produces more low-energy hadrons. Depending on the π^0 production in the shower, these hadronic showers may have smaller electromagnetic showers embedded in them, from $\pi^0 \rightarrow \gamma\gamma$. Many of the hadronic interactions with the matter, such as pion absorption by nuclei within the steel, are invisible to the detector. Some fraction of the momentum is carried away by neutrons, which are poorly contained in the detector. A typical hadronic shower from an NC event is shown in the middle of Figure 4.1.
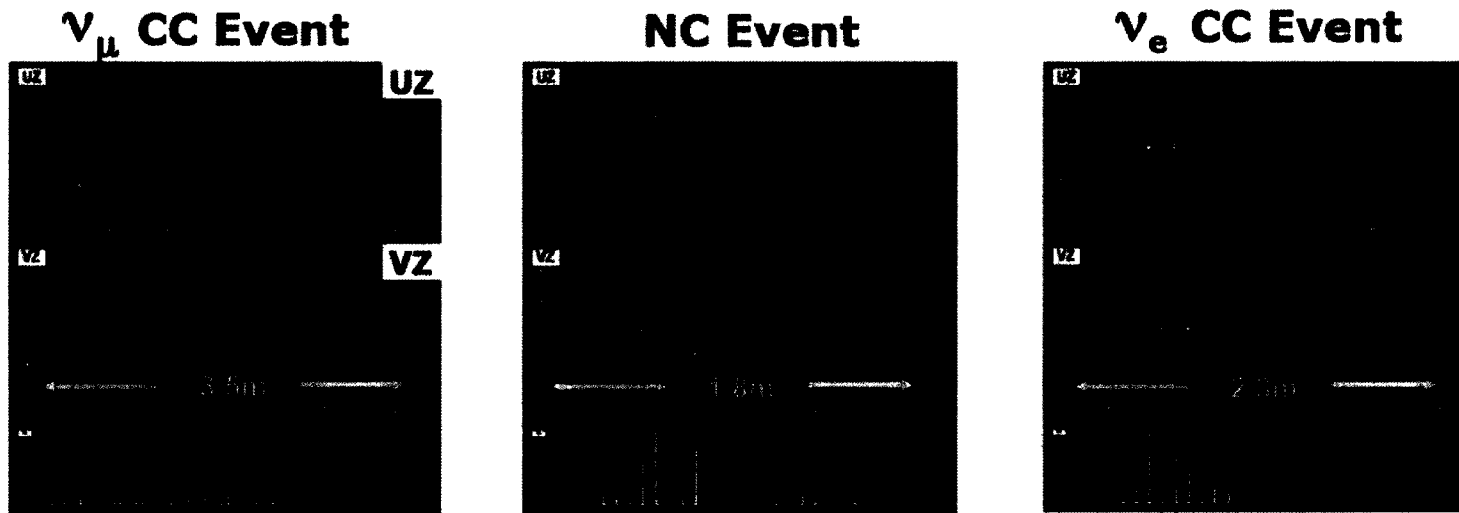
FIG. 4.1: Examples of different types of events as recorded by the MINOS detectors, shown in two spatial views, uz and vz, and with recorded pulse time along the bottom. On the left, a ν_μ CC interaction with a long muon track with a small hadronic shower at its vertex. In the middle, an NC interaction, which produces a hadronic shower. On the right, a ν_e CC interaction, inducing an electromagnetic shower induced by the outgoing electron. The bottom panels display pulse-height as a function of time for the displayed events.

4.3.2 Electromagnetic Showers

Muons are 500 times more massive than electrons, so electromagnetic showers induced by Bremsstrahlung radiation are negligible in MINOS. Electrons appear in MINOS from CC interactions of ν_e, from the ν_e contamination in the NuMI beam and possibly from $\nu_\mu \rightarrow \nu_e$ oscillations. Electrons exhibit a different $\beta\gamma$-E dependence than muons, and so they do Bremsstrahlung photons which pair-produce e^+e^- at energies represented in the NuMI beam. The resulting e^+e^- pair can themselves Bremsstrahlung or annihilate, producing more photons, etc. A high energy electromagnetic particle interacting with matter will produce an electromagnetic cascade, or electromagnetic shower. A typical electromagnetic shower from a ν_e CC event is shown on the right in Figure 4.1.

4.3.3 CALDET calibration

The small CALDET detector was placed in a test beam at CERN and exposed to e, π, and p beams of varying momenta. The calorimetric electromagnetic shower response and hadronic shower response in the MINOS detectors are different because hadrons are interacting strongly, while electrons interact electromagnetically. The measured response to each species is shown in Figure 4.2.

CALDET data was used to extensively validate shower models for possible use in MINOS simulations [92]. The model chosen was GCALOR [93], as it agreed well with the response seen from π^+. Neither of the two models tested, GCALOR or GEISHA [94] agreed particularly well with π^- in CALDET, and the 5.7% spread between the data and predictions from both GCALOR simulations and GEISHA simulations was taken to be the

flat 5.7% hadronic energy scale systematic uncertainty [47].
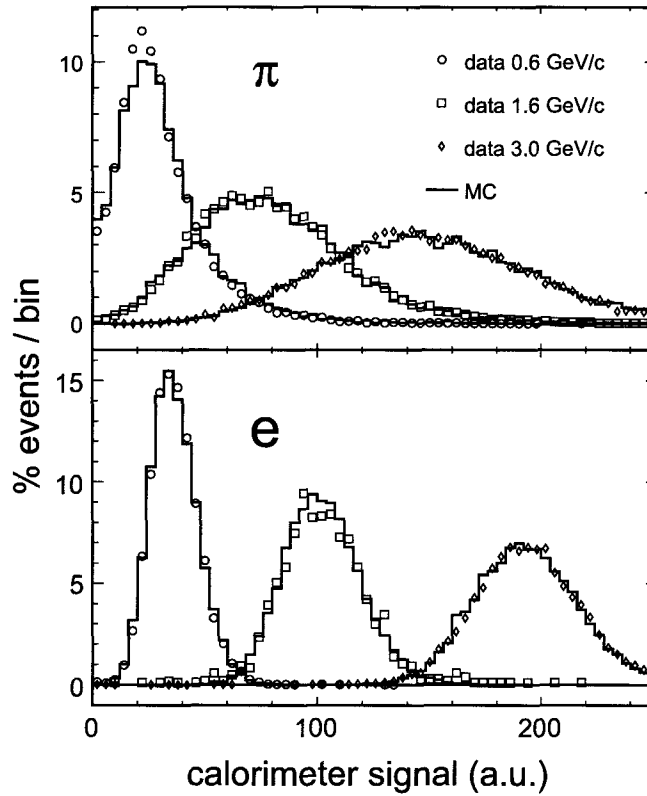
FIG. 4.2: Calorimetric response in CALDET from π and e at three momenta. The calorimeter signal scale is in arbitrary units. GCALOR simulations for π and e showers at these three momenta are also shown, as well as CALDET data (open points) [66].

4.4 Reconstruction

Reconstruction software takes the energy depositions and their locations recorded in the detector and reconstructs showers and tracks. In the Near Detector, there are several neutrino interactions producing secondary particles in a single spill. Hits that are in close proximity in space and in time are assumed to be associated with a single neutrino interaction, and collections of hits separated in space and time are sliced into separate events.

The Far Detector event rate is so low that slicing is not performed. Once events are sliced, the reconstruction procedure for events in both the Near and Far Detectors is the same.

Events associated with a single neutrino interaction are then constructed into tracks and showers. The events are given to the Kalman filter described above, which walks plane-by-plane to group hits together which produce the longest track. If the track exits, the Kalman filter assigns the track a momentum derived from the track curvature in the magnetic field.

Hits that are not associated with the track are then grouped into clusters, and groups of clusters are considered to be related in a single shower. It is possible for a single event to have more than one reconstructed shower, but the most energetic shower is taken to be the primary shower. The energy of a neutrino is then taken to be the sum of the momentum of the longest track and the energy of the primary shower.

$$E_\nu = E_{\text{track}} + E_{\text{shower}} \tag{4.3}$$

It is also useful to define a variable, called kinematic y, which is the fraction of neutrino energy which goes into the shower:

$$y \equiv \frac{E_{shw}}{E_{trk} + E_{shw}} \tag{4.4}$$

Studies have shown [74] that the reconstruction efficiency is improved by removing any hit depositing less than two photoelectrons from consideration in reconstruction. These low-energy hits are considered to be effects of cross-talk between PMT pixels which are

not well modeled in the Monte Carlo simulations.

4.5 Calibration

A raw photomultiplier signal is converted into a calibrated signal via a series of multiplicative factors [66]. The total conversion factor for the raw pulse height (Q_{raw}) depends on several factors:

- **Location**. The total PE yield for a MIP will be different along the length of the strip relative to the read-out end due to attenuation in the WLS fiber. The attenuation correction is A, and depends on strip i and location x.

- **Channel**. The total PE yield depends on the efficiency of the WLS fiber and the PMT. The strip-to-strip correction is S, and is a function of strip i and time t.

- **Time**. The PMT response is time-dependent on a short time scale due to temperature and high voltage fluctuations, and the total photon yield for a MIP changes as scintillator ages on a long time scale. D is the drift correction.

- **Linearity**. Strip-by-strip functional corrections L are used to linearize PMT response with pulse-height.

- **Scale factor**. M is an overall scale factor that converts the corrected pulse-height into the same standardized energy unit (muon energy unit, MEU), comparable between the Near and Far Detectors, as well as CALDET.

The calibrated signal is found by calculating

$$Q_{cal} = Q_{raw} \times D(t) \times L(i, Q_{raw}) \times S(i, t) \times A(i, x) \times M \qquad (4.5)$$

The calibration procedure is not perfect, and some energy scale uncertainty remains. The individual detectors display differences in calorimetric response relative to the absolute energy scale. These are primarily due to spatial variations in detector response not accounted for with $A(i, x)$. The differences between data and simulations are assigned as detector-dependent systematic uncertainties [95].

4.5.1 Light Injection System

The PMT's that read the signal from the WLS fibers are run in proportional mode, where the PMT signal is proportional to the number of photoelectrons, rather than trigger mode, which records the same signal above any threshold. The pulse height measured from the WS fibers carries information about the energy loss of particles passing through the scintillator. One tool for the calibration of individual strips and PMT pixels is the Light Injection (LI) system. The LI system consists of UV Light Emitting Diodes (LED's) housed in rack-mounted "pulser boxes" which are connected to the read-out ends of scintillator strips. The LI fibers illuminate the WLS fibers and a PIN photodiode with a well-known energy spectrum and intensity. The Data Aquisition system records the correlation between the intensity of the LI pulse and the number of photoelectrons recorded by the PMT. The LED intensity is tuned to produce roughly 50 PE per pulse. The PIN

photodiode is also illuminated to correct for the aging of the LED's and optical fibers themselves over time. There is an LI system for both the Near and Far detectors, and the systems for the two detectors are nearly identical.

The results from LI testing map the linearity of the instrumentation and the PMT and gain stability. The LI LED's in both detectors are flashed several hundred times per hour in both detectors. There are functionally identical LI systems in the Near and Far Detectors, as well as in the CALDET detector.

4.5.2 Cosmic Rays

Both Near and Far Detectors are exposed to muons produced from cosmic ray interactions in the upper atmosphere. Many of these muons have energies that make them MIPs in the MINOS detectors. Stopping cosmic rays can be used to calibrate between Monte Carlo simulations and data by comparing the dE/dx of stopping muons to the theoretical ionization energy loss. Samples of stopping cosmic ray muons are collected between beam spills within the nominal fiducial volume, to ensure event containment. Aligning the minima of the dE/dx distributions provides an absolute energy scale calibration. The agreement in the minimum ionizing dip is shown in between cosmic ray muon data and Bethe-Bloch energy loss model in Figure 4.3. The 2% uncertainty on the exact location of the minimum translates to a 0.2% error on the energy deposition [66].
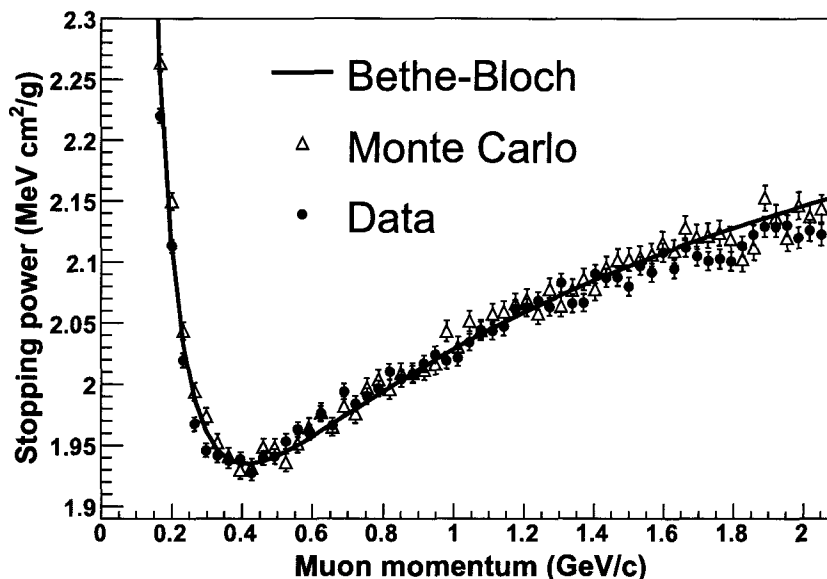
FIG. 4.3: Stopping power for muons in data taken from the Far Detector. Theoretical calculation from Bethe-Bloch and Monte Carlo simulations are also shown. Bethe-Bloch is shown for reference only, as muon energy loss in MINOS is modeled by [89]. Minimum ionization for muons is found to be 0.4 GeV/c. Plot from [66].

4.6 Timing

The NuMI beam spills are 2.2 s apart, last $12\,\mu$s, and protons are extracted from the Main Injector in either five or six batches, depending on the Tevatron. To minimize cosmic ray backgrounds, the Far Detector oscillation analysis only accepts beam spills for $14\,\mu$s around the expected spill time. The neutrino time of flight between FNAL and Soudan is 2.449 ms [71]. The timing of the observed events relative to expected spill time are shown for the Near Detector in Figure 4.4, and for the Far Detector in the Appendix.
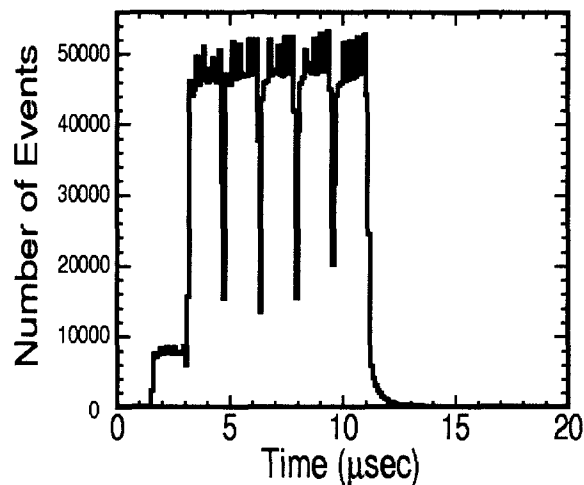
FIG. 4.4: The spill timing seen in the Near Detector, relative to the expected spill time. Spills in all six time "buckets" are visible. Spill times outside of a window of expected arrival time $-2 < t < +14\,\mu s$, where t is the expected arrival time, are rejected [19].

4.7 Event selection

A pre-selection battery of cuts is applied to data in both the Near and Far Detectors as a first pass on selection purity to remove most cosmic ray muons and some Neutral-Current events. Pre-selection cuts also mitigate geometric effects within the detectors by imposing fiducial volumes within the two detectors. Fiducial volume imposition ensures containment of neutrino interaction location and hadronic showers for precise energy resolution.

Since cosmic ray events and beam data quality are not concerns for Monte Carlo simulated events, the preselection cuts are divided in to two classes: those that are applied to all events, MC or data, and those that are applied only to the data. The preselection cuts that are applied to all events are listed here, along with their purpose.

- **Require a track** (`ntrk> 0`). All CC interactions involve an outgoing muon, so this cut removes a substantial NC background.

- **Track has been fit** (`trkfitpass==true`). The track fit has successfully converged.

- **Fiducial Volume**. The vertex of the interaction should be within this volume, which is different between the two detectors due to the magnetic field differences and acceptance requirements. Diagrams of the fiducial volume definitions can be found in Figure 4.5.

    - **Near Detector** A cylinder with $r < 1\,\text{m}$, where $r = 0$ is the center of the beam spot on the Near Detector, offset $1.48\,\text{m}$ from the center of the coil hole in x and $0.24\,\text{m}$ in y. The cylinder is not exactly a right cylinder. The ND fiducial volume follows the beam path, which means the top and bottom edges of the cylinder slant down $3.3°$ from horizontal. The front edge of the fiducial cylinder is $1\,\text{m}$ back downstream the front face, and the cylinder extends for $4\,\text{m}$.

    - **Far Detector** A torus in each SM with $0.5\,\text{m} < r < \sqrt{14}\,\text{m}$, where $r = 0$ is the center of the coil hole. Also, the first four planes of each SM are excluded, as well as the last eight planes of SM1 and the last twenty planes of SM2. The xz properties of the FD fiducial volume are shown in Figure 5.7.

The second class of preselection cuts are applied to data in the Near and Far detectors to remove backgrounds not present in Monte Carlo and to ensure data quality with respect to the NuMI Beamline. The extrapolation method described in Chapter 6 requires neutrinos from pions that are focused by the electromagnetic focusing horns. NuMI spills

which deposit too many protons in the upstream baffles induce large uncertainties in the overall neutrino flux. Spills where data quality is not assured are therefore rejected.

- `goodBeamToUse==true`. This variable ensures that the NuMI spill was within all acceptable parameters. These include the current in the magnetic focusing horns, the status of POT counting, and size of the beam spot.

- `coilIsOK==true`. This variable requires that the detectors be fully magnetized so that muon momentum can be measured from its track curvature.

- `isLI==false`. Ensuring Light Injection calibration signals do not pulse right before the spill, which would mimick highly energetic strip hits.

- `dirCosNu>0.6`. This checks the angle of the beginning of the muon track with respect to the vertex position and the z-axis. If the cosine of the angle is less than 0.6, the track is likely from a cosmic ray neutrino instead of a beam neutrino.

- `-2<GoodTimeToNearestSpill<+12` μs. This short window controls the live time of the detector, minimizing the cosmic ray background. This leaves the spill window open for 2 μs on either side of the 10 μs NuMI spill duration.

4.7.1 The primary kNN selection algorithm

The spill timing cuts eliminate cosmic ray background events, so the remaining detector interactions consist of CC and NC beam events. Some CC events are obvious, due to the long muon tracks leaving the interaction vertex. Some NC events are obviously NC,

due to the lack of a muon track. The challenge is differentiating between CC events with shorter muon tracks and NC events with an energetic pion punching through the shower.

Separation of CC and NC events is achieved through selection criteria. The separation can be evaluated by calculating the selection efficiency and purity. Purity is defined as

$$P = \frac{(\# \text{ True CC Events })}{(\# \text{ Selected Events })} \tag{4.6}$$

and efficiency is defined as

$$E = \frac{(\# \text{ Selected CC events })}{(\# \text{ Total true CC events })} \tag{4.7}$$

Applying selection criteria harshly can increase the purity of the selected sample, but can remove true CC events, hurting the efficiency of the selected sample and increasing the size of the statistical error. The oscillation sensitivity is optimized when the product of purity and efficiency is maximized.

In the case of this analysis, two further selection algorithms are used, one applied to all events and is tuned to maximize oscillation sensitivity, and one applied to events with $E < 5\,\text{GeV}$ and is tuned to maximize sensitivity to alternative disappearance hypotheses. Both of these selectors are based on the same statistical algorithm, the kNN method, that outputs a single-valued particle identification parameter, called the PID, which separates different event types in to different parameter value ranges. Since MINOS suffers from one significant background, the PID is valued between 0 and 1, with NC-like events assigned smaller PID values and CC-like events assigned larger PID values.

The kNN algorithm, short for "k Nearest-Neighbors," takes the values of n variables in a Monte Carlo training sample, containing reconstructed and truth information, and populates an n-dimensional space of reconstructed event properties. A test event is placed in the space and compared to its k nearest neighbors in the n-dimensional space. If the sought-after variable, say interaction type, is in truth CC for $\frac{k}{4}$ of those neighbors, then the test event is assigned a kNN parameter value of $\frac{k}{4} \div k = 0.25$. The kNN parameter value ranges from 0 to 1, representing the fraction of k nearest neighbors that agree. A cut on the kNN parameter is used to select samples enhanced on signal events and depleted in background events.

This kNN method was pioneered for use in MINOS by R. Ospanov for the oscillation analysis of Runs I and II [74]. The input variables and the optimum values for k and n are tuned to maximize the product of selection purity and efficiency. He found the optimum values to be $k = 80$, where the optimum variables contained in the training sample are:

- **Number of planes in event**. CC interactions produce muons with long tracks in the detector. NC events produce hadronic showers which do not extend as far in the z-direction.

- **Mean energy deposited per strip**. Total pulse-height in an event, measured in MIPs, divided by the total number of strips participating. Muons are minimum-ionizing particles, so longer tracks with many strip hits and little energy are separated in this variable from hadronic showers. Pions will knock out protons, which move slowly and deposit lots of energy in each strip they pass through.

- **Signal fluctuation parameter**. The number of low pulse-height strips per high pulse-

height strips in and around the track, where high and low pulse-height are defined relative to the mean pulse hight of the strip. Pions are likely to be absorbed towards the end of their tracks.

- **Transverse profile parameter.** The pulse height of identified track strips divided by the number of near-track strips over threshold. This indicates the separation of the μ-like track from the rest of the event. Pion tracks will show fluctuations along the length of the track due to pion-nucleon scattering.

These variables were only calculated for the last 80% of a track. The intent was to remove all noise from hadronic showers. The cut value of the output CC/NC separation variable, called `roID`, was optimized to maximize the oscillation sensitivity of the sample. Distributions of these variables are shown in Figure A.2. Events are rejected if their primary kNN variable is less than 0.25.

4.7.2 The Secondary kNN selection algorithm

For the analysis of Runs I, II, and III, a secondary selection criterion is applied to those events that fail the primary selection criteria. This secondary selection is based on a second kNN filled with three different variables aimed primarily at characterizing short muon tracks with low energy. Low energy events, below 5 GeV populate a region where neutrino disappearance model discrimination is most sensitive. Instead of maximizing `purity` $\times$ `efficiency` for Monte Carlo simulated events, this cut was tuned to maximize $\Delta\chi^2$ of pure decay and pure decoherence predicted spectra for high-statistics fake data [96]. The new variables for this kNN are calculated for 100% of the track length, unlike

the kNN calculation described above. The new variables are shown in Figure 4.7, and are

- **Number of planes in event.** In this case, the first 20% of the track has not been

  removed, as in the primary kNN selection variable.

- **Pulse height in the last 5 planes of the track.** Pulse height significantly larger than

  that of a minimum-ionizing particle is indicative of the track undergoing nuclear inter-

  actions.

- **Scattering variables (2).** Two variables quantifying the smoothness of the track.

  muon-like pion tracks undergo nuclear interactions and scatter more than true muons.

  A Pearson coefficient is constructed in the $u - z$ and $v - z$ views to calculate the

  scattering variable P:

$$P = \frac{0.01}{1.01 - \rho} \tag{4.8}$$

where

$$\rho = \frac{\sum_i j_i z_i}{N \sigma_j \sigma_z} \tag{4.9}$$

where N is the number of hits and (j,z) is the position of the hit (either u or v is

substituted for j), and σ_j and σ_z are the standard deviations of the position variables.

4.8 Data/Monte Carlo Agreement

Simulations are necessarily abstractions of the real physical world. A perfect Monte

Carlo simulation would yield identical physical distributions, within statistical errors, for

variables that we are concerned with in the experiment. We expect low statistical significance for these variables in the Far Detector, due to the small expected event rate, but the Near Detector records several orders of magnitude more events. Data and Monte Carlo quantities can be compared with little statistical noise in the Near Detector. Low level Far Detector variables, with no oscillation sensitivity, can be directly compared to simulations prior to oscillation analysis. Data/Monte Carlo agreement for variables used in the selection are shown in Appendix A. Observed differences are the basis for many of the systematic uncertainties, but overall the level of agreement is sufficient such that that the oscillation analysis is not affected.

4.9 Conclusion

The MINOS detectors detect neutrinos through the products of neutrino interactions. Computer software is used to process patterns of hits in the plastic scintillator into reconstructed muon tracks and hadronic showers. The energy of detected neutrinos is reconstructed from the sum of energies deposited by daughter particles from neutrino interactions in the detector. For showers, the energy deposited in the detector is proportional to the amount of light deposited in the scintillator, where the scintillator response is calibrated with an *in situ* light injection system and the detector response is calibrated by measurements made with CALDET in calibration beams of particles. The hadronic model used in the simulations is accurate to at least 5.7% based on agreement between CALDET data and and simulations. The uncertainty is energy dependent, and is larger at lower energies.

Track energy is calculated by the muon's integrated ionization energy loss in the case of fully contained events or by the muon's curvature in the detector's magnetic field if the muon exits the detector. Track energy from ionization energy loss is calibrated with stopping cosmic ray muons. The track energy scale uncertainty is 2% when the energy is determined by range, and the energy scale uncertainty is 3% when the energy is determined from curvature, or 1% relative to the range error, fully correlated.

A number of cuts are implemented to remove signals which do not originate with CC interactions from the neutrinos produced by the NuMI beam. Two parameters, based on the kNN statistical algorithm, are used to produce a particle classification variable in order to select ν_μ CC events. The cut on the primary kNN PID is tuned to maximize the product of the purity and efficiency of the selected events. The cut on the secondary kNN PID is designed to reclaim low energy events rejected by the primary kNN selection, and is tuned to maximize the sensitivity to alternative neutrino disappearance models. With calibrations, data quality, and PID cuts applied to real data, we observe that there is good agreement between data and simulations in the Near Detector.
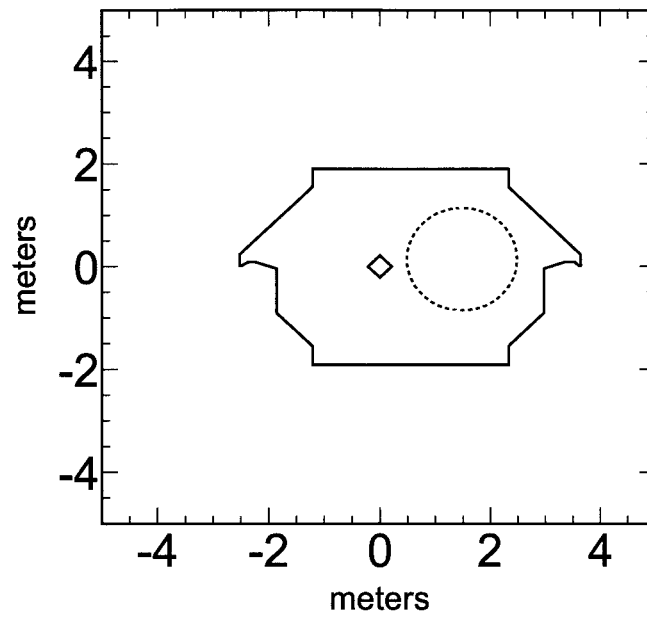
(a) Near Detector
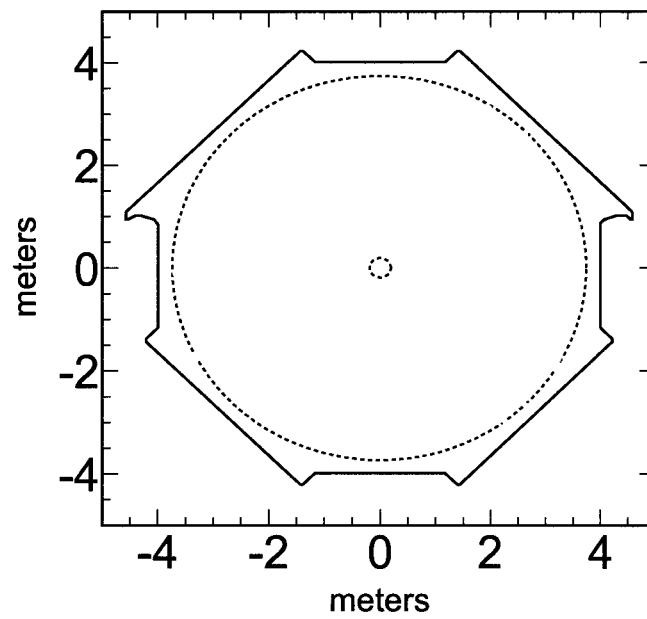
(b) Far Detector

FIG. 4.5: The xy view of the fiducial volumes of both detectors are indicated by the blue dashed circles. Only events with vertices contained within these volumes are considered for the analysis.

(a) Number of active planes

(b) Mean energy deposited per strip

(c) Signal fluctuation parameter

(d) Transverse profile parameter
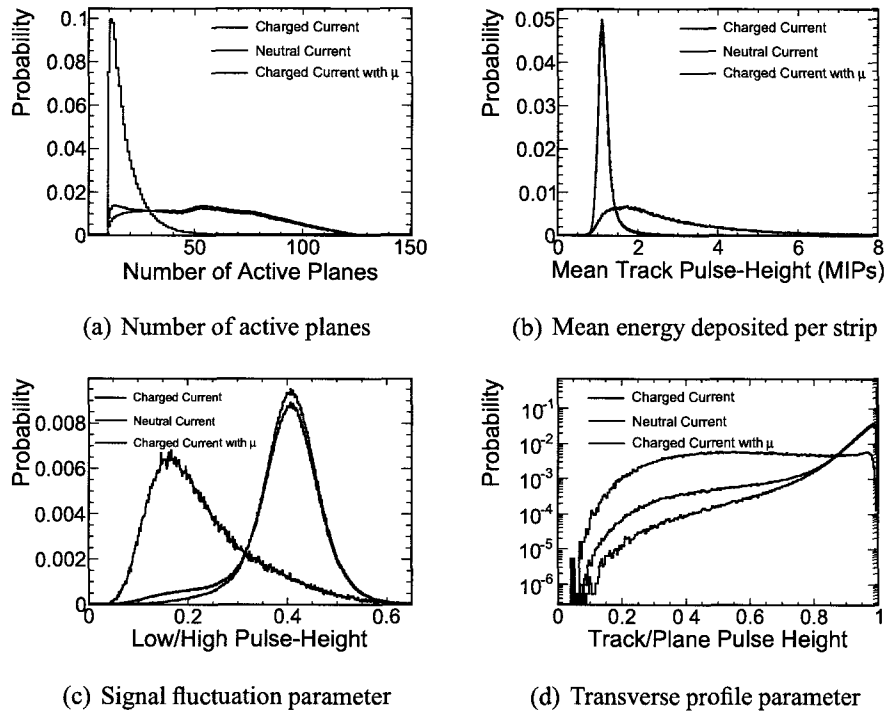
FIG. 4.6: Variables used in the primary kNN selection variable algorithm. Histograms are area-normalized filled with simulated Near Detector events. True ν_μ CC events and NC events are shown separately, to show the separation power of each variable [96].

(a) Track end pulse height

(b) Scattering variable
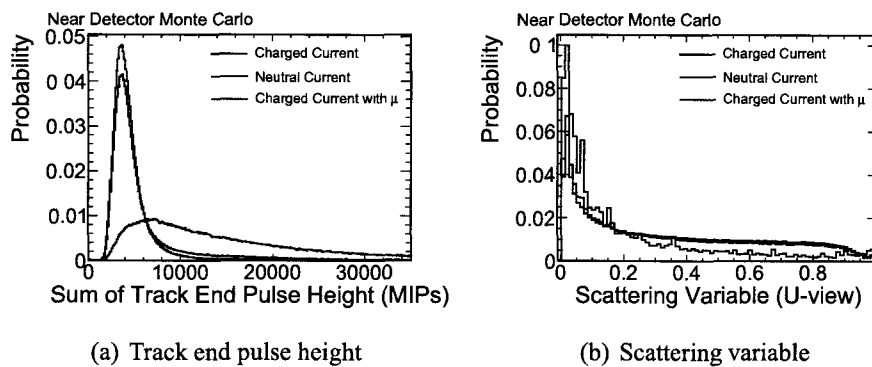
FIG. 4.7: Variables used in the secondary kNN selection variable algorithm, which is tuned to select low-energy events that are rejected by the primary kNN selection cuts. True ν_μ CC events and NC events are shown separately, to show the separation power of each variable [96].

CHAPTER 5

Neutrino Oscillation Analysis

The data analyzed in this thesis was taken between May 2005 and June 2009. The data accumulated between these dates contains an integrated beam exposure of 7.2×10^{20} POT separated into three separate runs, as shown in Table 5.1.

The MINOS detectors have been taking data for five years, and the experiment is near the end of its operational lifetime. Much of the time and effort that MINOS requires is spent on improving analysis methods and understanding systematic errors. This chapter describes the analysis methods used to measure the oscillation parameters for a dataset composed of Runs I, II, and III with ν_μ Charged-Current events. Since the last publication

Configuration	Start Date	End Date	Total POT ($\times 10^{20}$)
I LE	5/20/2005	2/26/2006	1.269
I pHE	6/11/2006	8/13/2006	0.153
II LE	9/12/2006	7/17/2007	1.943
III LE	11/18/2007	6/13/2009	3.881

TABLE 5.1: The total Far Detector beam exposure and run time for MINOS beam analysis runs.

of MINOS results, the experiment has accumulated more than twice the amount of data and made many improvements to the analysis algorithms.

5.0.1 Blind Analysis

The MINOS collaboration has chosen to perform a blind analysis, to eliminate the possibility of bias due to prior experimental measurements of neutrino oscillation parameters. Blinding entails obscuring the disappearance signal in the Far Detector by hiding part of the Far Detector data. A fraction of the Far Detector data set is open for the purpose of assuring data quality. A blinding algorithm has been applied to an energy-dependent fraction of the Far Detector data that is open so that any attempted oscillation analyses of this subset of data would be fruitless. All Monte Carlo simulations are open, as well as the entire Near Detector data set.

Every time an analysis group within MINOS wishes to analyze a set of Far Detector data, known as "opening the box," the proposed analysis must be frozen and subject to collaboration approval. The frozen analysis components include the event reconstruction, simulations, event selection, and the fit method. The systematic uncertainties and their effect on the overall systematic error must also be evaluated. Once the analysis has been documented and concerns from fellow collaborators have been addressed in accordance with the collaboration bylaws, the group may open the box.

5.1   Summary of prior MINOS oscillation results

The analysis described in this thesis is the primary mission of the MINOS experiment. There have been two analyses of cumulative datasets published prior to the opening described in this document. These publications are released after the conclusions of experimental runs. Results from Run I were published in December 2006 with an accumulated exposures of 1.27×10^{20} POT at the Far Detector [97]. This preliminary measurement found $\Delta m^2 = \left(2.74^{+0\ 44}_{-0\ 26}\right) \times 10^3 \mathrm{eV}^2/\mathrm{c}^4$ and $\sin^2(2\theta) > 0.87$ at 90% C.L.

Results from Runs I + II were published in 2008 with an accumulated exposure of 3.2×10^{20} POT [47]. This was the first analysis to utilize the kNN statistical methods to define a PID. A single PID was used, which was identical to the primary kNN described in Chapter 4. The beam was extrapolated with the beam matrix method and cross-checked with the Far/Near extrapolation. These extrapolation methods will be described in Chapter 6. Three of the largest systematics were included as nuisance parameters in the deterimination of the oscillation parameters. These systematics were implemented as scale factors on the size of the NC background, the absolute hadronic energy scale, and the total event rate normalization. The oscillation parameters were found to be $\Delta m^2 = \left(2.43^{+0\ 13}_{-0\ 13}\right) \times 10^{-3} \mathrm{eV}^2/\mathrm{c}^4$ and $\sin^2(2\theta) > 0.9$ at 90% C.L. with a $\chi^2/\mathrm{DOF}=90.2/97$. The best fit point in $\sin^2(2\theta)$ was unphysical, with $\sin^2(2\theta) = 1.066$. The alternative neutrino disappearance models of neutrino decay and neutrino decoherence were disfavored at 3.7σ and 5.7σ, respectively.

The allowed region for oscillation parameters Δm^2 and $\sin^2(2\theta)$ found by these two analyses are shown on the right of Figure 5.1. These were the most sensitive measure-
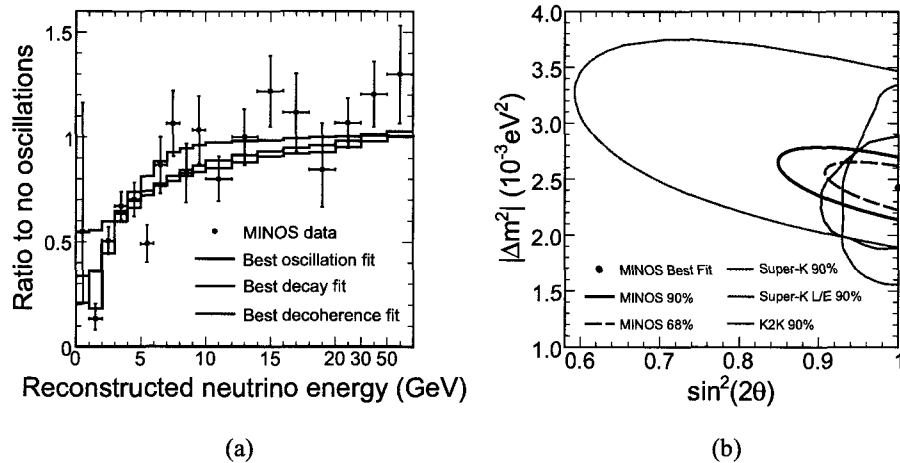
FIG. 5.1: The ratio of data to unoscillated expectation for Far Detector data consisting of Runs I and II. On the right, the allowed region for oscillation parameters, compared to other experimental measurements [47].

ments of Δm^2 at the time. This is compared to oscillation measurements made with Run I alone in Figure 5.2.

5.2 Analysis Improvements

The results of the analysis from Runs I + II presented a quandary for the MINOS experiment. The signal in the region of the oscillation minimum was less than the expected background. The resulting best-fit point was so far unphysical, where $\sin^2(2\theta) > 0$, that the $\sin^2(2\theta)$ limit in the physical region was suppressed - in a sense the result was lucky. The limits were not guaranteed to improve by taking more data. To make the next run worthwhile, improvements had to be made to the analysis to extract as much oscillation information as possible.
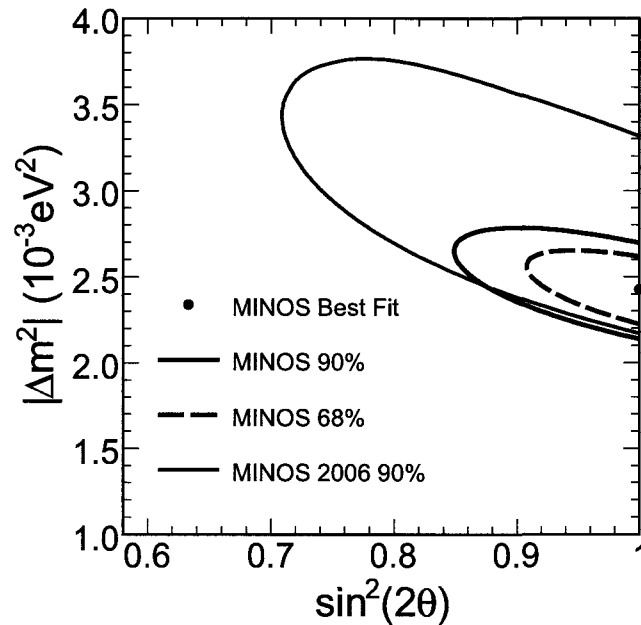
FIG. 5.2: MINOS oscillation results published in 2006 (red, 90% C.L. only) and 2008 (black, 68% and 90% C.L.), with 1.3×10^{20} POT and 3.2×10^{20} POT, respectively.

A number of new approaches have been explored between the box openings for Runs II and III. The effectiveness of new techniques have been evaluated by comparing oscillation sensitivities with and without these methods in place. The sensitivity calculation and gain from these new methods will be shown in Chapter 7, but in this chapter the methods themselves will be described.

A new algorithm is now used to determine shower energy. The event selection method has been augmented to maximize alternative model discrimination, as described in Chapter 4. Event energy resolution information is now used to improve oscillation sensitivity. Antineutrino-like events are now considered in the overall fit to further improve oscillation sensitivity. Finally, events recorded outside the fiducial volume of the Far De-

tector are now also considered in the overall fit. The details of these improvements are described below.

5.3   Augmented event selection

The secondary event selection, described in Chapter 4, is designed to recapture low-energy ν_μ events which are rejected by the primary selection algorithm. Since the maximal discrepancy between the alternative disappearance models occurs at low energies, the augmented selection contributes to the discrimination between these models. These low-energy events have a lower purity, as shown in the comparison of the selection algorithms in Figure 5.3. This was found to leave the oscillation sensitivity unchanged.
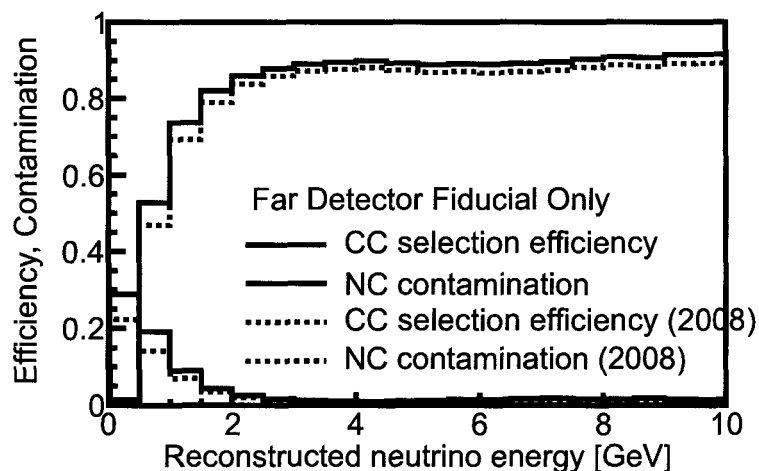
FIG. 5.3: The effect of the secondary selection algorithm on the Far Detector selection efficiency and purity, compared to the efficiency and purity with the primary selector alone, as was used in the 2008 MINOS analysis [47].

5.4 Shower Energy from a kNN algorithm

Calorimetric shower energy estimation, which was the method used for shower energy estimation in past publications, yields poor energy resolution at low energies. For this analysis, the shower energy was obtained from a kNN algorithm. The kNN algorithm compares n variables describing hadronic showers in data to simulated showers in an n-dimensional space. The value for shwEn for data showers is assigned to be the mean of the true shower energies of the k nearest-neighbors in the n-dimensional space.

The variables chosen to populate the kNN parameter space were three ($n = 3$) reconstructed quantities that correlate with shower energy [98]. They are:

- nplaneshw. The number of planes struck by the primary shower. This is strongly correlated with the shower energy.

- trkShwEnNearDW. The sum of de-weighted shower energies within 1 m of track vertex. The deweighted shower energy calculation alters the response of the detector by changing the relative importance of the number of strips vs. the total number of photoelectrons recorded. The deweighting function is shower energy-dependent [90].

- shwEnCor+((nshw>1)*shwEnCor2). Sum of all reconstructed showers, including secondary reconstructed showers, if any.

The optimized value for k with these variables was found to be $k = 400$.

This algorithm produces shower energies closer to the true E_{shw} in Monte Carlo. It does introduce an energy bias at low energies, introduced by the presence of physical boundaries within the n-dimensional space. For very low-E_{shw} showers, the values of the

kNN variables are close to zero. The 400 nearest neighbors in the n-dimensional space will not surround the low-E_{shw} event isotropically, and so will bias the mean E_{shw} to that of events with kNN variable values further from the physical boundary.

This bias is corrected with a shower weighting procedure corresponding to the polynomial fit shown in Figure 5.4. The improvement in resolution is shown in Figure 5.5 integrated over all values for E_{shw}, and broken up into 500 MeV true energy bits in Figure 5.6. At low energies, below 500 MeV, the energy resolution attained with the kNN algorithm is 50% better than the calorimetric shower energy estimation [98].
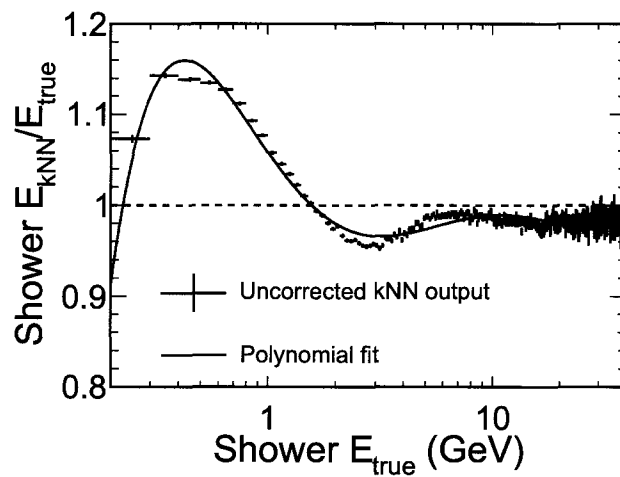
FIG. 5.4: Shower energy bias from kNN shower energy estimator and the polynomial fit used to correct this bias. Taken from [98].
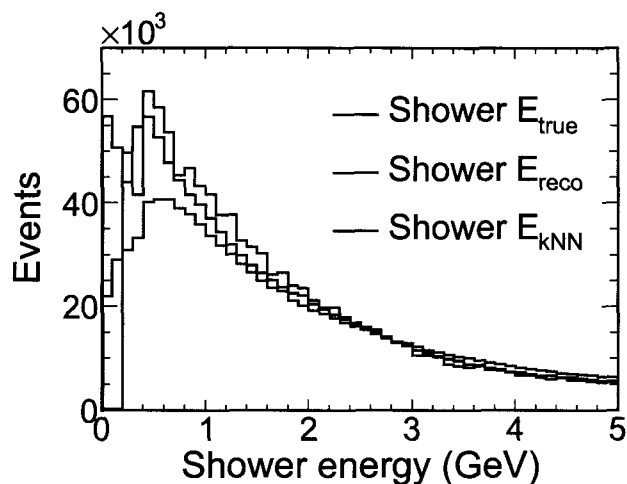
FIG 5 5 Companson of vanous methods of hadronic shower energy estimation, applied to simulated events The black line is the simulated shower energy The blue line is the calorimetric shower energy and the red line is the shower energy as estimated by a kNN algorithm The kNN shower energy estimation matches the true simulated shower energy more closely than the calorimetric shower energy estimation method [98]
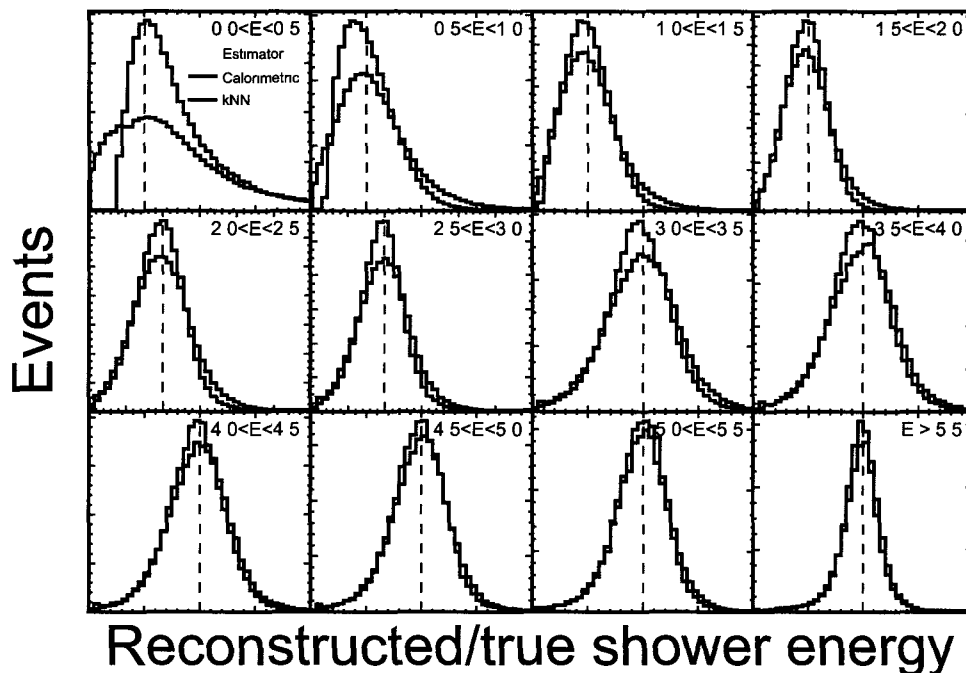
FIG 5 6 Companson of calorimetric shower energy resolution and kNN shower energy resolution in 500 MeV energy bits This is shown prior to the energy bias correction, and the dashed vertical line shows the location of the peak of calorimetric shower energy distribution, illustrating the misalignment of the peaks with the two methods prior to correction [98]

5.5 Rock and Anti-Fiducial events

The neutrino beam is several kilometers wide by the time it reaches the Far Detector. Some neutrinos from the NuMI beamline interact in the rock upstream and around the Far Detector, and the muons from CC interactions in the rock can punch into the detector. In addition, the mass of the detector outside of the fiducial volume yields a non-negligible number of neutrino CC events. These Rock and Anti-Fiducial (RAF) events are included as a separate sample. For a complete discussion of the RAF analysis, see Reference [99].

The true energy of an incoming neutrino that interacts via a Charged-Current interaction in the rock around the detector (*a.k.a.* a Rock event) is $E_\nu = E_{track} + E_{shower}$. For rock events, little or none of the hadronic shower is seen by the Far Detector. Likewise for neutrino events occuring in the anti-fiducial region, the shower energy is often poorly contained, yielding poor shower energy resolution. For these reasons, shower energy is ignored for both samples in the RAF analysis.

The anti-fiducial region of the Far Detector is large and different sub-regions are sensitive to different event pathologies. The Far Detector is broken up into six geometric regions, and each is predicted separately to account for these differences. These six regions are shown in Figure 5.7 and are described below:

- **Front face.** The fiducial region excludes events with vertices in the first 4 planes in order to exclude the background of muons from neutrino interactions upstream of the detector.

- **Rock-like edge.** The Far Detector is instrumented outside the fiducial region, defined

to be $r < \sqrt{14}$ m, but shower and track containment yield low-resolution event information. The RAF analysis ignores shower energy. Some of the events in this region posess true vertices in the rock around the detector, and some fall within the detector but outside of the fiducial volume.

- **Detector-like edge**. The first strip hit in an edge event defines whether the event was Rock-like or Detector-like. If the first strip is near the outside edge of the detector it is indicative of an event with a vertex in the rock, but if the first strip is several strips inside, the event vertex is likely in the anti-fiducial region itself.

- **SM gap**. The fiducial region excludes events with vertices in the last 8 planes in SM1 and first 4 planes in SM2 because it is unlikely that either hadronic showers or muon tracks originating here would be contained.

- **Edge of SM gap**. The hollow cylinder region outside $r = \sqrt{14}$ m and also within the set of planes defined to be the SM gap above is defined to be a separate region.

- **Back planes**. Like the SM gap, The fiducial region excludes events with vertices in the last 20 planes because it is unlikely that either hadronic showers or muon tracks originating here would be contained.

Each of these regions comprises a different fiducial mass, so each records a different event rate. There are two binning schemes applied to energy distributions for events in these regions. The high-rate regions are the detector-like and rock-like edge, and the front face. These regions have the following 28-bin binning scheme: one bin for 0–0.75 GeV, 0.25 GeV bins from 0.75 GeV up to 4 GeV, 0.5 GeV bins up to 6 GeV, 1 GeV bins up

to 10 GeV, 5 GeV bins up to 30 GeV, one bin for 30–45 GeV, and one bin for all events

above 45 GeV. The lower-rate regions are the SM gap, SM gap edge, and back planes.

These regions have the following 9-bin binning scheme: 0–1 GeV, 1–1.5 GeV, 1.5–2 GeV,

2–3 GeV, 3–4 GeV, 4–6 GeV, 6–9 GeV, 9–15 GeV, and 1 bin for all events above 15 GeV

[100]. These six regions add 111 degrees of freedom for each experimental run, and is

meant to be fit for oscillations simultaneously with data and predictions from fiducial

events.
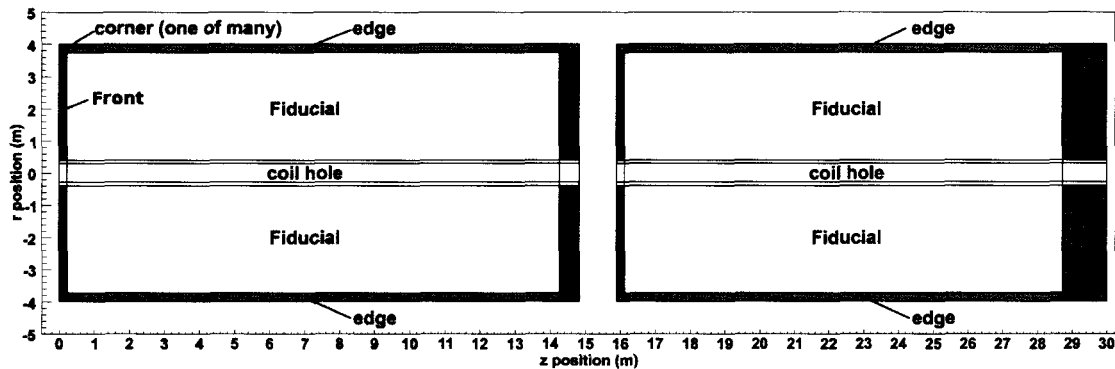
FIG. 5.7: A diagram of all of the regions used in the RAF analysis. The detector edge region is further broken into two samples, one rock-like and one detector-like, based on the likely location of the event vertex. Figure taken from [100].

5.6 Resolution information

The hadronic shower energy estimation and Neutral-Current backgrounds remain

problematic and affect the precision of the $\sin^2(2\theta)$ measurement. NC background events

are reconstructed with missing energy, due to the exiting neutrino, and feed down to lower

energies, filling in the oscillation dip. Poor shower energy resolution tends to smear out

reconstructed neutrino energies. A smeared spectrum in the Far Detector also tends to fill

in events in the oscillation dip, degrading the sensitivity to $\sin^2(2\theta)$.

The effect of these problems on the overall oscillation sensitivity can be mitigated by taking into account the estimated energy resolution of each event. If the muon stops in the detector then the muon momentum from dE/dx and the energy resolution is very good. If the muon exits the detector, then the energy resolution is somewhat poorer. The portion of neutrino energy determined by hadronic shower reconstruction has significantly poorer energy resolution.

A parameterized resolution function was calculated separately for the reconstructed energies of hadronic showers, contained tracks, and exiting tracks. The parameterization was derived from studies comparing the true energies of Monte Carlo simulated events to the energies assigned to the same events by the reconstruction software. The resolution is defined as the Gaussian width of the distribution of the difference between reconstructed and true event energies, as a function of reconstructed energy. Using this parameterization, an estimate of energy resolution can be calculated event-by-event.

Just as there are two components contributing to the reconstructed energy of every event, $E_{total} = E_{track} + E_{shower}$, there are two components contributing to the resolution of every event. The components come from the energy resolution of the muon track σ_{track} and the energy resolution of the hadronic shower σ_{shower}, in units of GeV,

$$\sigma_{total} = \sqrt{\sigma_{track}^2 + \sigma_{shower}^2} \tag{5.1}$$

The energy resolution of a muon track depends on the muon's containment, since this determines whether the muon momentum is calculated from dE/dx or from the muon
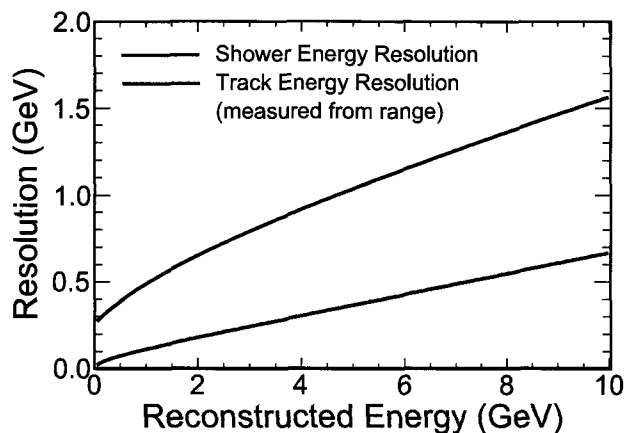
FIG. 5.8: The reconstructed energy-dependent resolution parameterization for contained tracks and showers. The resolution parameterization of uncontained tracks depends on both reconstructed momentum from track curvature and from the track fitting uncertainty.

track curvature.

$$\sigma^2_{\text{track}} = \begin{cases} \sigma^2_{\text{range}} = (0.0512\,\text{GeV}^{1/2})^2 E_{\text{range}} + ((0.0692)E_{\text{range}})^2 \\ \\ \sigma^2_{\text{curv}} = (1.341\,\text{GeV}^{1/2})^2 p^2 \sigma_{q/p} + (p^2 \sigma_{q/p})^2 \end{cases} \tag{5.2}$$

The parameterization of the shower energy resolution depends only on the shower energy itself:

$$\sigma^2_{\text{shower}} = (0.181\,\text{GeV})^2 + (0.425\,\text{GeV}^{1/2})^2 E_{\text{shower}} + (0.075 E_{\text{shower}})^2 \tag{5.3}$$

The resolution parameterization for contained tracks and hadronic showers is shown in Figure 5.8. Distributions of track and shower resolutions as a function of neutrino energy are shown in Figures 5.9 and 5.10.
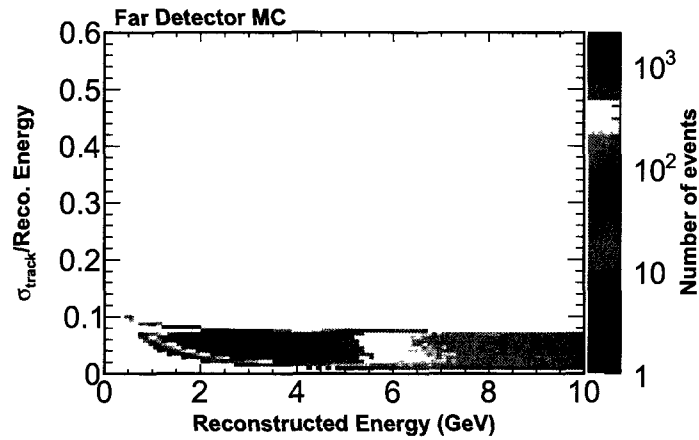
(a) Contained events
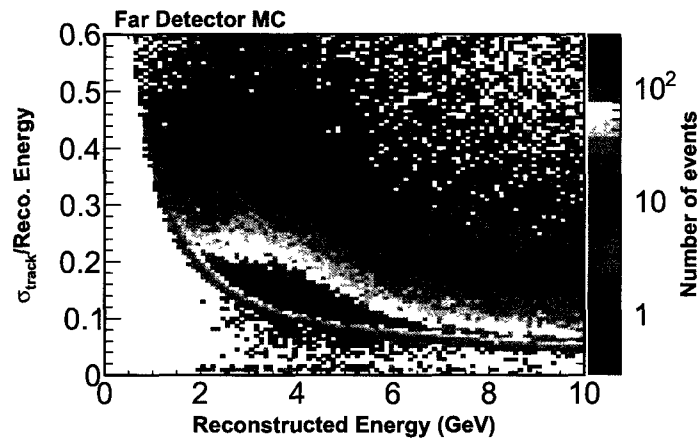
(b) Exiting events

FIG. 5.9: The fractional energy resolution contribution from reconstructed tracks to the total parameterized energy resolution. This is separated in to samples of events with μ tracks contained within the detector 5.9(a), with μ momentum measured from range, and events with μ tracks exiting the detector 5.9(b), with μ momentum measured from curvature. These are shown on the same σ_{track}/Reco. Energy scale to show the relative energy resolution of both types of μ tracks.
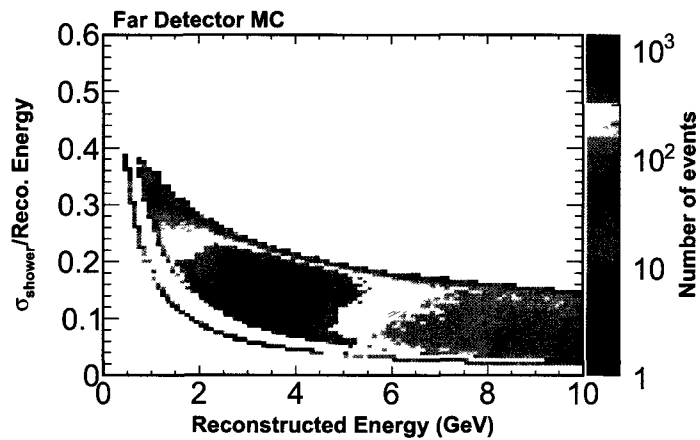
FIG. 5.10: The fractional energy resolution contribution from reconstructed showers to the total parameterized energy resolution.

The fractional energy resolution, σ_{Total}/E_{reco} is the metric used to separate events based on their energy resolution. For every value of reconstructed neutrino energy there is a distribution of σ_{Total}/E_{reco} in Monte Carlo. Boundaries are placed on σ_{Total}/E_{reco} for this energy to break this distribution into a number of samples, each containing the same number of simulated events. An example of this calculation is shown in Figure 5.11. This is performed for every bin in the Far Detector reconstructed neutrino energy spectrum. The samples are then analyzed separately, and are called bins of resolution. For the oscil-
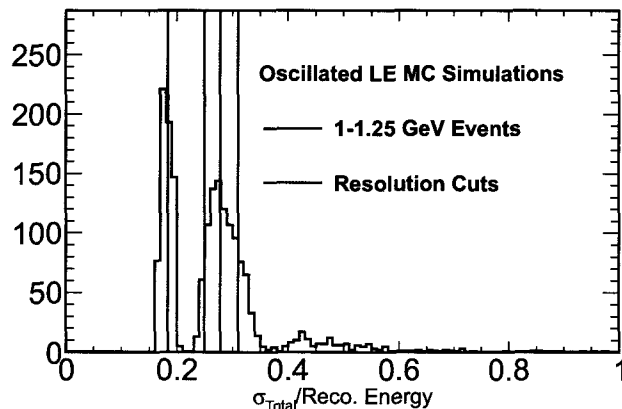
FIG. 5.11: Fractional resolution for 1-1.25 GeV neutrino energy bin. The blue lines indicate placement of cuts to split this energy bin into five quantiles based on energy resolution, each of which has an equal number of events. This calculation is performed for each neutrino energy bin.

lation analysis, Far Detector events are split into five bins of resolution. The distribution of σ_{Total}/E_{reco} and the boundaries between resolution bins are shown in Figure 5.12.

Different Far Detector energy spectra predictions are made for the different bins of resolution, which are all identical before any neutrino disappearance function is applied. When the oscillation measurement is made, the Far Detector data events are split into bins of resolution using the same set of boundaries, and the five Far Detector predictions are
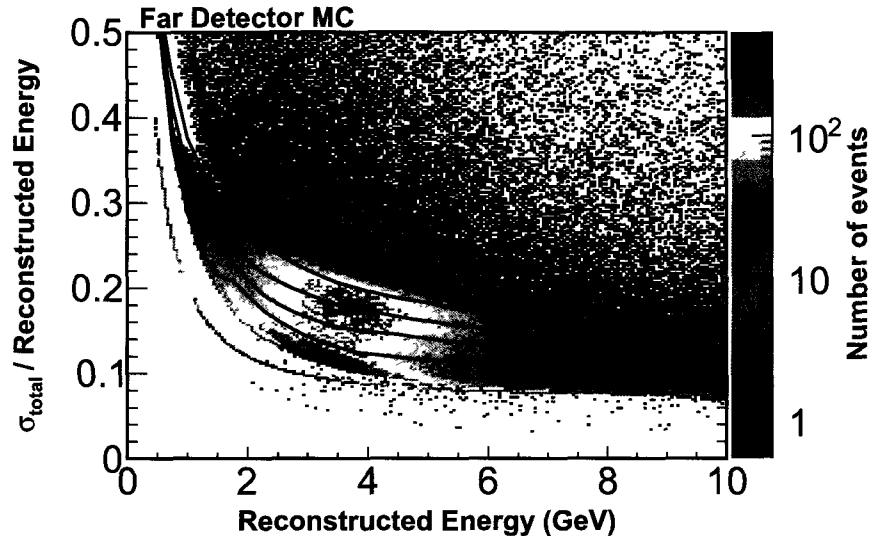
FIG. 5.12: Distribution of σ_{Total}/E_{Reco} for Run III MC simulations. The four black lines indicate where the boundaries are between the five bins of resolution. The resolution variable allows for separation of well-measured quasi-elastic events from less precise DIS, and resonance events, as well as background NC events that pass CC/NC selection cuts.

fit simultaneously to the five data spectra. A precise description of how this technique improves oscillation sensitivity will have to wait until Chapter 6, but essentially the most precise resolution bins are able to pull the best fit point toward the true value while the backgrounds are quarantined in the resolution bin with the poorest energy resolution.

Many binning schemes were investigated using different numbers of resolution bins and in different configurations. Studies were performed with high-statistics Monte Carlo fake data. In these studies, the configuration and the resolution bin boundaries, as a function of energy, were determined from high-statistics Monte Carlo simulations. Fake data was then split up according to these boundaries, and the chosen number of resolution bins were fit simultaneously to produce a statistical sensitivity. The sensitivities produced from this procedure were compared to the nominal case of one resolution bin (that is to say, with no consideration of resolution, as in prior analyses).

Fitting with five resolution bins produced a significantly improved sensitivity, while fitting with ten resolution bins showed only a marginal improvement beyond that. This is consistent with earlier studies of this technique [90]. Splitting the data into too many groups risks statistical complications. Sub-dividing in to too many resolution bins, the number of events per energy bin falls to the level where bins no longer obeys Gaussian statistics. When resolution bins have too few events, statistical fluctuations can lead to unphysical oscillations parameters ($\sin^2(2\theta) > 1$), which leads to predicted bin weights of less than 0 events.

Based on this study, the number of expected data events produce significantly better sensitivity with the fewest statistical side-effects with five resolution bins. The two extreme bins predominately represent different interaction types. Bin 0, containing the 20% of events with the most precise estimated resolution, contains mostly CC quasi-elastic ν_μ interactions with stopping μ tracks. Bin 4, containing the 20% of events with the poorest estimated energy resolution, contains almost all of the NC background, as well as poorly-resolved high-y events. No weighting scheme is employed when fitting these five resolution bins for oscillations.

In practice, the events in bin 0 also contain some badly reconstructed events with muon tracks that appear to stop within the detector. These events are typically ones that enter the coil hole, where there is no scintillator, and do not punch through the coil hole on the far side. Thus the muon track is reconstructed to be shorter than it truly is, so they tend to be events with low-biased reconstructed energies. To mitigate this, tracks with endpoints located within 40 cm of the Far Detector coil hole have their track resolution

assigned from curvature, even if they are considered stopping, to inflate their fractional resolution and remove them from bin 0.

5.7 Analyzing $\bar{\nu}_\mu$'s

To further improve the oscillation sensitivity of this analysis, events with defocused muons are extrapolated and fit along with the five resolution bins. Defocused muons are identified by muon tracks with charge> 0. These can be broken down further into three categories:

- Beam $\bar{\nu}_\mu$'s from π^- decays that traveled neck-to-neck in the focusing horns, and so were not defocused away from the beam. These neutrinos tend to have higher energies.

- Beam ν_μ's from π^+ decays. Some μ^-'s from ν_μ CC interactions have low momentum and are susceptible to multiple scattering, obfuscating their charge sign in the detector's magnetic field.

- CC-like NC events with a pion mis-identified as a muon. Pions of both negative and positive charges are produced in hadronic showers, so the NC background is relevant for both ν_μ and $\bar{\nu}_\mu$ analyses.

Including the sample of events with positive reconstructed charge represents a 12% increase in the total number of events. Of this, 24% is ν_μ and 69% is $\bar{\nu}_\mu$. The relative composition of the $\bar{\nu}_\mu$-like sample is shown in Figure 5.13. The expected rate of $\bar{\nu}_\mu$-like events per POT varies by run, based on the change in the position of the target relative to the horns in runs I and II, and the addition of helium in the decay pipe in Run III. The
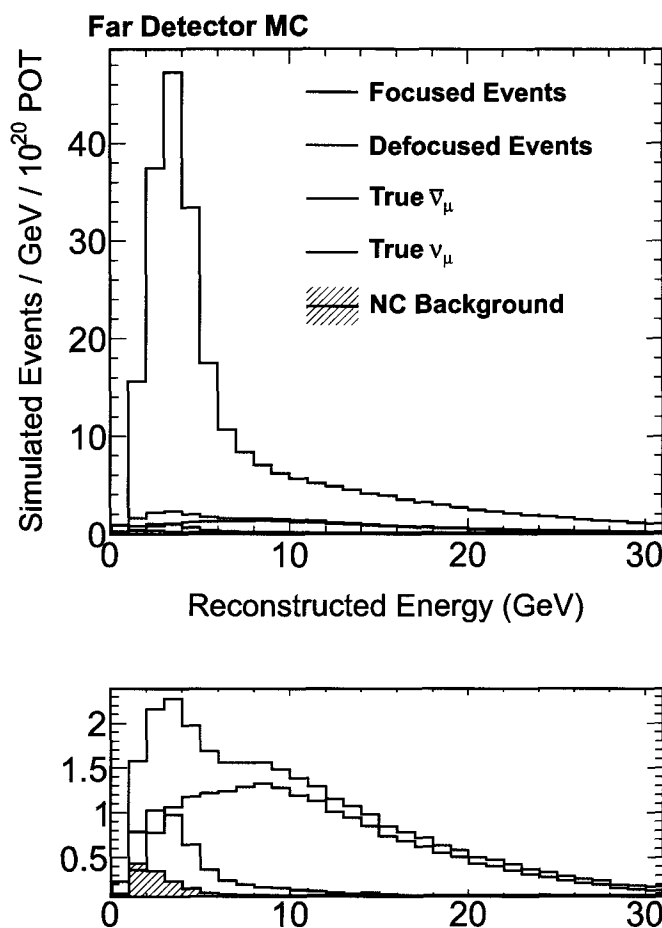
FIG. 5.13: Comparison of negative curvature and positive curvature spectra simulated in the Far Detector with a neutrino beam. Below, the positive curvature spectrum and its components are enlarged.

degradation of the target in Runs II and III also introduces a small effect in the $\bar{\nu}_\mu$-like spectrum. From Monte Carlo, the expectation per run is shown in Table 5.2.

The exact amount of ν_μ expected to be misidentified with positive reconstructed charge depends on the accuracy of simulations. To try and place a systematic uncertainty on this value, the sample of charge> 0 events in the Near Detector was subjected to an additional cut on track angle to make a purer estimate of the size of the ν_μ signal. The track angle is measured by projecting the μ direction measured at the vertex and

Run	Total charge< 0	$\bar{\nu}_\mu$	ν_μ	NC	Total charge> 0
Run I	321.99	27.70	9.04	3.01	39.75
Run I pHE	118.90	2.70	0.28	1.27	4.25
Run II	481.19	41.87	13.41	4.53	59.80
Run III	944.81	85.31	26.57	8.45	120.33

TABLE 5.2: Expected number of events in charge> 0 sample by species from Monte Carlo simulations by run in absence of oscillations. All runs are scaled to their respective exposure. For comparison, the expected number of events in the charge< 0 sample is shown.

comparing it to the last track hit. This separates μ^- and μ^+ with higher significance than the Kalman filter, since it measures focusing or defocusing as a function of radial deflection. Data and simulations were compared with this variable and the maximum discrepancy between the two was found to be 40%. The maximum discrepancy occured at energies above the true ν_μ focusing peak, so this 40% is a conservative estimate of the uncertainty on this sample [101].

The inclusion of positive curvature events introduces an additional complication when implementing resolution binning. Initially, the positive curvature events were lumped in with the negative curvature events and the net sample was broken up into five resolution bins. The resolution bin boundaries were trained over a sample of both positive and negative curvature events, with the relative proportion determined by the beam-weighted Monte Carlo. Almost all of the positive curvature events wound up in Bin 5, and the extra background yielded a statistical sensitivity that was worse than eliminating positive curvature events altogether.

We investigated separating samples based on charge sign and then splitting each of those into five resolution bins (for ten total resolution bins per run). In this case, the resolution function is trained on focused events only, and the resolution boundaries were

used for both ν_μ and $\bar{\nu}_\mu$-like events separately. This is not the optimum scenario, as the

$\bar{\nu}_\mu$ energy spectrum has a different shape due to the de-focusing of π^+ in the focusing

horns and the different y-distributions associated with antineutrinos (Figure 5.14). This
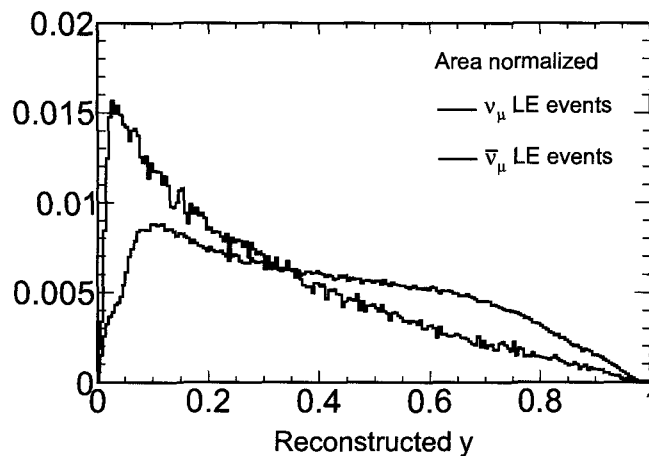
FIG. 5.14: An area-normalized comparison of the MC simulated y-distributions for ν_μ and $\bar{\nu}_\mu$ CC interactions.

model with ten resolution bins produced a poorer oscillation sensitivity. This is due to

the fact that the $\bar{\nu}_\mu$-like data spectrum at this exposure does not contain enough events

to avoid the problem with statistical complications. The total $\bar{\nu}_\mu$-like expectation is too

small to avoid statistical fluctuations leading to negative Far Detector event predictions

when oscillations are applied. The optimum oscillation sensitivity was found with five

resolution bins for focused events and one sample containing all $\bar{\nu}_\mu$-like events.

CHAPTER 6

Extrapolation and fitting

The MINOS detectors are designed to sample a flux of muon neutrinos before and after those neutrinos have traveled a long distance. Determining the parameters describing neutrino disappearance through oscillations, or any other muon neutrino disappearance model, requires a precise prediction in the absence of any flux modification. As has been discussed in previous chapters, there are large systematic uncertainties associated with the neutrino flux from the NuMI beam. The Near Detector provides a way to characterize the beam and inform the flux simulations, reducing flux-related systematic uncertainties.

The Near Detector measures the neutrino flux times the neutrino interaction cross-section, and produces an energy spectrum of the NuMI beam. The reconstructed neutrino energy spectrum measured in the Near Detector is not a direct prediction of the energy spectrum one expects to measure in the Far Detector, given the difference in fluxes both detectors are exposed to. This chapter describes the methods used to predict the Far Detector spectrum, given the spectrum measurement made in the Near Detector.

141

6.1 The Need for Extrapolation

In MINOS, the Near Detector is located 1 km downstream of the NuMI target, and less than 400 m downstream of the end of the decay pipe, which marks the end of neutrino production in the NuMI beam. The Near Detector can sample the ν_μ flux and beam composition (and contributes to the calculation of beam weights, as described in Chapter 4) before any oscillations have taken place. The Near Detector and Far Detector are functionally similar but their relative angular sizes and location relative to the beam origin means that the two detectors have differences in their relative neutrino acceptances. This means that the energy spectrum measured in the Near Detector is not a direct prediction of the energy spectrum expected in the Far Detector.

Neutrinos from the NuMI beam are created along the entire 675 m length of the decay pipe from decaying pions (see Chapter 2). The Near Detector lies on the beam axis and detects neutrinos originating from pion and muon decays along the length of the decay pipe and from a range of parent decay angles. From the point of view of the Near Detector, the NuMI beam looks like a distributed line source of neutrinos. From 735 km away, in the point of view of the Far Detector, the NuMI beam looks like a point-source of neutrinos. This matters in constructing the Far Detector flux prediction because the Near Detector will be exposed to neutrinos with different kinematic ranges than the Far Detector. This effect must be taken out if we wish to use the Near Detector to predict the Far Detector flux.

For example, consider pions which travel co-linearly with the incident NuMI protons (the $\hat{z}$ direction). The 1 m-radius fiducial region of the Near Detector accepts neutrinos

with decay angles anywhere from 0.006 degrees at the beginning of the decay pipe to 0.19 degrees at the end of the decay pipe (in the lab frame). The Far Detector accepts neutrinos with an opening angle of at most 3×10^{-4} degrees. This example is complicated by the fact that parent pions carry some transverse momentum and travel with some opening angle relative to the initial proton beam.
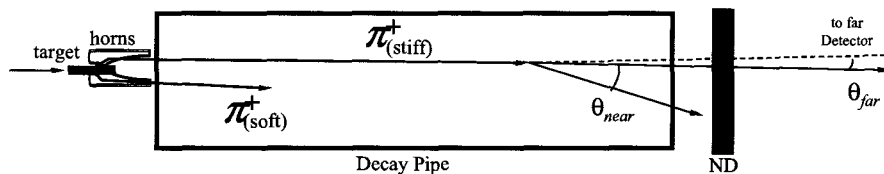
FIG. 6.1: A cartoon showing the relative angular size of the two detectors with respect to the NuMI target and the need for extrapolation. Taken from [19].

The neutrino flux seen in either the Near or the Far detector, then, has a dependence on the decay kinematics of the parent pion. A pion with a significant transverse momentum (p_T), will more likely produce a neutrino with a large opening angle, with respect to $\hat{z}$, that will appear in only the Near Detector. A pion with a large p_T is unlikely to produce a neutrino that intersects both detectors, as shown in the diagram in Figure 6.1.

Various extrapolation methods have been developed to predict the Far Detector neutrino energy spectrum given the Near Detector energy spectrum [46]. More than one extrapolation method is used for each analysis to cross-check the predicted spectrum's validity prior to looking at the data. The beam matrix extrapolation method is the primary extrapolation method used for both the ν_μ and $\overline{\nu}_\mu$ analyses [102]. The Far/Near ratio method is used as a cross-check. This chapter describes the implementation of these two extrapolation methods in MINOS, which were used in this analysis.

6.2 The Far/Near method

The simplest extrapolation method MINOS uses to extrapolate between the Near and Far Detectors is the Far/Near method. This method simply reweights the number of events in the Near Detector data by the ratio of the number of simulated Far Detector events to the simulated Near Detector events. The predicted Far Detector number of events in reconstructed energy bin i is

$$F_i^{\text{Prediction}} = N_i^{\text{Data}} \times \frac{F_i^{MC}}{N_i^{MC}} \tag{6.1}$$

where F represents Far Detector energy spectra and N represents Near Detector energy spectra. The simulated spectra, F_i^{MC} and N_i^{MC}, are filled with the reconstructed energies of selected simulated events. The ratio, F^{MC}/N^{MC}, encapsulates the beamline geometry, as coded into FLUKA, the detector efficiency, as modeled in GMINOS (GEANT), and the overall normalization differences between the two detectors due to the detectors' relative solid angles with respect to the beam.

One can think of the Far/Near method in two ways. The true Near Detector data is being reweighted by the expected spectral shape difference modeled in simulations. One could also note that, in the limit that $N^{MC} = N_i^{Data}$, $F_i^{\text{Prediction}}$ is identical to F_i^{MC}. In this sense, we are reweighting the Far Detector simulation by the data/MC differences observed in reconstructed neutrino energy in the Near Detector.

The Far/Near method assumes that the relationship between true neutrino energy and reconstructed neutrino energy is the same in both detectors. The prediction process
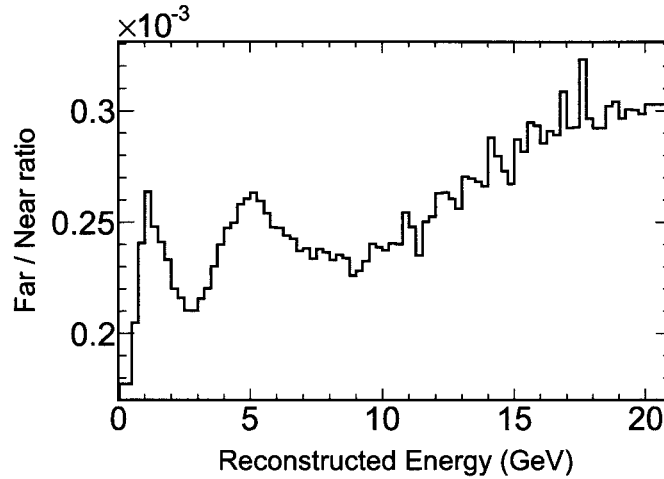
FIG. 6.2: The Far/Near ratio for simulated events in Run III. The Near Detector reconstructed energy spectrum is reweighted bin-by-bin by this ratio to produce a Far Detector predicted spectrum.

must be carried out on a set of one dimensional spectra (histograms) with no selection efficiency or purity corrections, *i.e.* backgrounds are lumped in to each sample.

$$F_i^{MC} = F_i^{MC,\, \text{True}\, \nu_\mu} + F_i^{MC,\, \nu_\mu - \text{like NC}} \tag{6.2}$$

and likewise for N_i^{MC}, N_i^{Data}, and $F^{\text{Prediction}}$. Selection efficiency does not matter and no correction needs to be applied, as the prediction includes all of the backgrounds.

This simplicity has the advantage of making a computationally simple prediction in the absence of any disappearance phenomenon. To make a prediction with disappearance phenomenon, the histograms must be filled event-by-event from MC. Filling event-by-event allows the disappearance weight to be calculated on the true neutrino properties.

Every extrapolation method has a different sensitivity to systematic errors. In the Far/Near method, the sensitivity to neutrino interaction cross-section errors are small, as

the neutrino interaction rates in the FD and ND Monte Carlo simulations cancel exactly in the ratio, since both detectors are made of the same material. The Far/Near extrapolation does suffer from overall larger systematic errors than the beam matrix method, though, because it assumes that the detectors have the same energy resolution and same selection efficiency, and that these things are well modeled.

6.2.1   Predicting the Far Detector spectrum with oscillations

Measuring a deficit of muon neutrinos requires knowledge of the expected flux in the absence of the disappearance mechanism. To determine the properties of the disappearance mechanism, you must insert your model and tune its parameters such that your prediction best matches the observation. The predicted flux described above can be modulated with the two-flavor oscillation model, and the oscillation parameters can be tuned until the data and oscillated prediction are in agreement.

In the two-flavor approximation, muon neutrinos oscillate into tau neutrinos while traversing the MINOS baseline. The tau neutrinos do interact in the Far Detector but tau leptons produced in CC interactions are not identifiable in the Far Detector. Tau leptons are very short lived. The problematic τ decay mode for MINOS is $\tau^- \rightarrow \mu^- \bar{\nu}_\mu \nu_\tau$, which occurs about 17% of the time [20]. Other decays look like electromagnetic or hadronic showers, which can sometimes mimic μ tracks. If the τ is reconstructed as a μ, the reconstructed energy will be biased to lower values because of the energy carried away by the ν_τ.

A true modeling of $\nu_\mu \rightarrow \nu_\tau$ oscillations includes the ν_τ which could be mis-

identified as ν_μ in the Far Detector. The ν_τ flux depends on the magnitude and shape of the true oscillations so, unlike the NC background, the size of the background will vary with oscillations.

The probability of observing either a ν_μ or ν_τ at the Far Detector in the two-flavor approximation is

$$P_{\nu_\mu \to \nu_\mu} = 1 - \sin^2(2\theta) \sin^2 \left(1.27 \Delta m^2 \frac{L}{E} \right) \tag{6.3}$$

$$P_{\nu_\mu \to \nu_\tau} = 1 - P_{\nu_\mu \to \nu_\mu} = \sin^2(2\theta) \sin^2 \left(1.27 \Delta m^2 \frac{L}{E} \right) \tag{6.4}$$

To predict the background in MINOS from ν_τ appearance, we first need a simulated ν_τ flux. Rather than simulate a ν_τ beam, the simulated ν_μ flux is re-used. The variable in simulations containing the true particle identification is flipped from ν_μ to ν_τ, representing a 100% transition of ν_μ to ν_τ. The ν_τ flux is then simulated and reconstructed in the Far Detector with ν_τ cross-sections in GMINOS. The ν_τ events which pass the event selection algorithm comprise the maximum possible ν_τ background. This ν_τ sample then has oscillations applied, the inverse of whatever oscillation function is being applied to the ν_μ sample, to determine the predicted ν_τ appearance spectrum due to $\nu_\mu \to \nu_\tau$ oscillations.

The total oscillated prediction is:

$$F_i^{\text{Prediction}} = \frac{N_i^{Data}}{N_i^{MC}} \sum_{j}^{\text{MC events}} \left(O_{j,i} F_i^{MC} + (1 - O_{j,i}) F_i^{MC, \nu_\mu - \text{like} \, \tau} \right) \tag{6.5}$$

where $O_{j,i}$ is the oscillation probability for true energy j of an neutrino associated with that neutrino's Far Detector reconstructed energy i. Applying the oscillations in simu-

lated true energy requires looping over the entire Monte Carlo sample for each pair of oscillation parameters to be tested.

6.3 The Beam Matrix Extrapolation Method

The Beam Matrix method makes fewer assumptions about the relationships between the Near and Far Detector acceptances, and uses knowledge of pion decay kinematics to predict the neutrino flux at the Far Detector from the measured Near Detector reconstructed energy spectrum. The general extrapolation strategy is to convert the Near Detector data energy spectrum into a neutrino flux that can be extrapolated, extrapolate that flux, then convert the extrapolated flux into an energy spectrum prediction at the Far Detector.

A series of corrections must first be applied to the Near Detector data to estimate the true ν_μ flux from the beam in the ND. A matrix is then filled wich relates the ND flux to the flux which is expected 734 km away. The elements of this matrix are derived from the two-body decay kinematics of pions simulated in the target hall and decay pipe. Finally, a series of corrections must be applied to the Far Detector flux to convert it into a predicted energy spectrum.

6.3.1 Corrections

The conversion of Near Detector data, in units of reconstructed energy, into a flux requires several corrections, all of which are derived from simulations. Before examining the extrapolation procedure in full, these corrections are defined here.

- **Purity Correction** - The purity of the sample of events is the energy-dependent fraction of events that are NC background events surviving event selection.

$$P_i = \left(\frac{\text{True CC Events}}{\text{Selected Events}} \right)_i \tag{6.6}$$

where i is a bin of reconstructed energy. Purity corrections are calculated separately for the Near Detector P_i^N and the Far Detector P_i^F.

- **Efficiency Correction** - The selection efficiency represents the energy-dependent fraction of CC events which remain after event selection.

$$E_j = \left(\frac{\text{Selected CC events}}{\text{Total true CC events}} \right)_j \tag{6.7}$$

where j is a bin of true energy. Efficiencies are calculated separately for the Near Detector E_j^N the Far Detector E_j^F.

- **Reco to True** - Chapter 4 discussed how track and shower reconstruction can reconstruct neutrino energies that are not accurate. Monte Carlo simulations, which retain information about their true generated energies, can quantify the relationship between reconstructed and true neutrino energies so the effect may be taken out. A matrix is constructed relating true generated neutrino energies to the reconstructed energy values found by the reconstruction software. Multiplying a pure ν_μ reconstructed energy spectrum by this matrix returns a pure ν_μ true energy spectrum. The elements are normalized such that multiplying the matrix by the reconstructed energy of a single event

returns a distribution of integral 1.0. An example is shown in Figure 6.3.

The matrix is denoted by M_{ij}^N, where i is the energy bin in reconstructed energy and j is the energy bin in true energy. The reco-to-true matrix is only calculated for the Near Detector.

- **True to Reco** - The inverse of the reco-to-true conversion, the true-to-reco matrix, smears a true energy spectrum by the detector resolution to produce a reconstructed energy spectrum, M_{ji}^F. The true-to-reco matrix is only calculated for the Far Detector.

- ν_μ **CC cross-section** - The ν_μ CC cross-section has been measured by many other experiments, as described in Chapter 1. MINOS uses the MODBYRS-4 model within NEUGEN [18], with its associated error band. Here the total CC cross-section, X_j is used (The sum of quasi-elastic, resonance, and deep-inelastic scattering cross-sections) as a function of true energy. The cross-section and related error is shown in Figures 1.2 and 1.3.

- **Detector mass** - The mass of material confined by the boundaries of the fiducial volume determines the number of events expected given a flux and cross-section. The detector mass is determined by the mean plane thickness, density, and total detector-specific fiducial volume. The fiducial volume masses are m_{ND} and m_{FD} [66].

6.3.2 Pion decay

Most of the neutrinos detected in the MINOS detectors come from two-body decays of pions and kaons in the evacuated NuMI decay pipe. Before explaining how the Beam

Matrix elements themselves are populated, it is worthwhile to describe the kinematics of pion decays in the NuMI decay pipe.

The parent with four-momentum q_π and mass m_π decays into a muon with four-momentum q_μ and mass m_μ and a neutrino with four-momentum q_ν and negligible mass.

$$q_\mu = q_\pi - q_\nu \tag{6.8}$$

Squaring both sides,

$$
\begin{aligned}
q_\mu^2 &= (q_\pi - q_\nu) \cdot (q_\pi - q_\nu) \\
m_\mu^2 &= q_\pi^2 + 2q_\pi \cdot q_\nu + q_\nu^2 \\
&= m_\pi^2 + 2(E_\pi^* E_\nu^* - \vec{q_\pi^*} \cdot \vec{q_\nu^*}) \\
&= m_\pi^2 - 2(m_\pi E_\nu^* - 0)
\end{aligned}
\tag{6.9}
$$

where center-of-mass variables are denoted with a *. Solving this for E_ν^* in the center-of-mass frame gives E_ν^* in invariant terms:

$$E_\nu^* = \frac{m_\mu - m_p}{2m_\pi} \tag{6.10}$$

so the neutrino is mono-energetic in the center-of-mass frame.

Transitioning to the lab frame and using the conservation of momentum we know

that

$$q_\pi^* \cdot q_\nu^* = \qquad q_\pi \cdot q_\nu \qquad (6.11)$$

$$E_\pi^* E_\nu^* - \vec{q_\pi^*} \cdot \vec{q_\nu^*} = \qquad E_\pi E_\nu - \vec{q_\pi} \cdot \vec{q_\nu}$$

$$m_\pi E_\nu^* - 0 = \quad E_\pi E_\nu - |\vec{q_\pi}||\vec{q_\nu}| \cos \theta_\nu$$

where θ_ν is the decay angle relative to $\hat{z}$, the direction of q_π, called $\hat{z}$. The energy of a daughter neutrino is then

$$E_\nu = \frac{m_\pi E_\nu^*}{E_\pi - |\vec{q_\pi}| \cos \theta_\nu} \qquad (6.12)$$

We can put this in natural units by recalling that $E = \gamma m$ and $\frac{p}{m} = \gamma \beta$:

$$E_\nu = \frac{E_\nu^*}{\gamma_\pi (1 - \beta_\pi \cos \theta_\nu)} \qquad (6.13)$$

Thus E_ν depends on the boost, and decay angle relative to the parent pion. Pion decays are isotropic in the center-of-mass frame.

$$\frac{dN}{d \cos \theta^*} = \frac{1}{2} \qquad (6.14)$$

Lorentz-transforming this into the lab frame,

$$\frac{dN}{d \cos \theta} = \frac{dN}{d \cos \theta^*} \frac{d \cos \theta^*}{d \cos \theta}$$

where

$$\frac{d\cos\theta^*}{d\cos\theta} = \frac{1}{\gamma_\pi^2(1 - \beta_\pi\cos\theta)^2}$$

The angular distribution of neutrinos from the beam is then [1]

$$\frac{dN}{d\cos\theta} = \frac{1}{2\gamma_\pi^2(1 - \beta_\pi\cos\theta)^2} \tag{6.15}$$

We can quantify the difference between neutrino fluxes at both detectors by using decay kinematics and beam detector geometry available to us from beam Monte Carlo simulations. Simulations reveal geometric effects that are difficult to model algebraically. Higher-energy pion decays tend to pile up toward the end of the decay pipe, because they pass through the volume of the evacuated decay pipe in a shorter time. These higher-energy pions decaying at the end of the decay pipe can produce neutrinos with large decay angles that are visible to the Near Detector but not the Far Detector. Similar aperture affects occur in the target hall and at the front of the decay pipe.

6.3.3 Beam Matrix

The beam matrix is a collection of weights which relates the flux of the neutrino beam from 1 km downstream of the target to 735 km downstream of the target. Construc-

---

[1] The two-body decay described above holds for the isotropic decays of π and K in the beam, but the daughter muons themselves can decay:

$$\mu^\pm \to e^\pm + \nu_e(\overline{\nu_e}) + \overline{\nu}_\mu(\nu_\mu)$$

Muon decays are not isotropic, due to the conservation of angular momentum. The parent pions of these muons are all spinless, and due to the left-handed nature of the neutrino, all μ^+, which are emitted from the pion decay back-to-back with the neutrino, must also have left-handed helicity. The overall contribution to the neutrino flux is small, and while the correct angular distribution is modeled in the beam matrix, it is neglected here.

tion of the beam matrix begins with simulations of neutrino parents $\pi^{\pm}$, and $K^{\pm}$ in the decay pipe.

As shown above, the flux of neutrinos is

$$\frac{dN}{d\cos\theta} = \frac{1}{2\gamma^2(1 - \beta\cos\theta)^2} \tag{6.16}$$

The calculation of each parent meson decay is repeated ten times toward different random locations within the Near Detector, and the true neutrino energy produced in the Near Detector is recorded. Sets of parents producing mono-energetic neutrinos in the Near Detector (within a small energy range) are collected. The Far Detector, 735 km away, represents a negligible solid angle, so these parents are only decayed a single time toward one point, which represents the effective size of the Far Detector from the beam's point of view. The energy of the neutrino at the Far Detector is weighted by the probability of the particular decay angle θ necessary to intersect the Far Detector.

$$P_{FD} \propto \frac{1}{r^2} \frac{dN}{d\cos\theta} \bigg|_{\cos\theta = \cos\theta_D} \tag{6.17}$$

where r is the distance from the decay point to the Far Detector and θ_D is the single decay angle that intercepts the Far Detector. The beam matrix is normalized such that a single event in the Near Detector yields a distribution of Far Detector energies with integral 1. A very thorough derivation of the beam matrix calculation can be found in Reference [103]. Beam matrices for neutrinos and antineutrinos can be seen in in Figure 6.4.

6.3.4 Beam Matrix Extrapolation Procedure

The Far Detector prediction begins with the Near Detector data energy spectrum. The purity correction is first applied to the Near Detector data energy spectrum to remove contributions from NC and wrong-sign contamination events. The result is a pure ν_μ Charged-Current event spectrum in the Near Detector, as selected by the PID. This is multiplied by the normalized reco-to-true smearing matrix, which converts the Near Detector reconstructed energy spectrum to a true neutrino energy spectrum. The Near Detector efficiency correction is then applied to correct for true CC events which are removed in the event selection process. The spectrum is divided by the ν_μ CC total cross-section and the fiducial region mass gives the total number of neutrinos passing through the fiducial volume. Dividing this number by the total beam exposure in units of protons-on-target (POT) yields the true ν_μ flux/POT at the Near Detector.

It is at this point that the beam matrix weights are applied to extrapolate the flux from the Near Detector to the Far Detector.

The same series of corrections which were applied to the Near Detector data and returned a flux are now applied in reverse order to the Far Detector flux to return a prediction. The Far Detector flux is multiplied by the detector mass, total exposure, and CC cross-section, to give the total number of neutrinos expected to interact within the Far Detector fiducial volume. The true ν_μ CC spectrum is scaled down by FD selection efficiency (see Figure 6.6), removing CC events which may be identified as NC events.

With the Far Detector flux in true energy at this point, oscillations or a different neutrino disappearance model may be incorporated, modulating the prediction in order

to better match the real data. The Far Detector flux is then converted into reconstructed energy with the Far Detector true-to-reco smearing matrix. The Far Detector NC expectation are added back in with the Far Detector purity correction. The result is a Far Detector reconstructed energy spectrum prediction which we can directly compare to the Far Detector data. For a visual interpretation of this procedure, see the flow chart in Figure 6.5.

The number of Far Detector events predicted in a particular energy bin may be calculated in a manner analagous to the Far/Near method. For FD reconstructed energy bins i, ND reconstructed energy bins j, FD true energy l, ND true energy k, the predicted weight in each FD reconstructed energy bin is:

$$F_i = \frac{m_{FD}}{m_{ND}} \sum_{j,k,l} N_j^{DATA} P_j^N M_{jk}^N (E_k^N)^{-1} (X_k)^{-1} B_{kl}^{NF} X_l O_l E_l^F M_{li}^F (P_i^F)^{-1} \tag{6.18}$$

where the terms were defined in Section 6.3.1.

6.3.5 Direct vs. Indirect extrapolation methods

The Far/Near method shares many properties with the beam matrix method. They both use beam and detector simulations to translate an observed Near Detector flux to the Far Detector. One could alternatively predict the Far Detector flux by tuning the neutrino cross-section and beam models within uncertainties so that Near Detector simulations match the Near Detector data, and then apply the same tuning to the Far Detector simulations. Various indirect methods (i.e. not "data-driven") of this variety have been explored within MINOS [46].

For example, the "NDFit" method attempted to fit the Near Detector simulations to the Near Detector data with a log-likelihood fit, and included Gaussian nuisance parameters that were then applied for several model uncertainties. The best-fit values of these parameters were then applied to the Far Detector simulations to predict the Far Detector spectrum. The model uncertainties considered were two cross-section parameters, track and shower absolute energy scales, and overall event rate normalization. The "2DFit" method was similar, but it attempted to fit the Far Detector data in two dimensions, in both reconstructed energy and the kinematic y.

For a thorough description of these models, see Reference [19]. These two methods were found in the past to be more sensitive to systematic errors than the two direct extrapolation methods [46, 104]. Due to this vulnerability, these two indirect methods have not been considered in the analysis the data set described in this thesis.

6.3.6 Accounting for disappearance

The background from ν_τ interactions is estimated in a similar manner as was described above with Far/Near. An all-ν_τ flux is predicted from the beam matrix, which carries a small selection efficiency ($< 20\%$). As oscillations are applied to the ν_μ spectrum in true energy, inverse oscillations are applied to the ν_τ as in Equation 6.4. The true ν_τ are then passed through a separate true-to-reco matrix which reflects the missing energy associated with ν_τ interactions. The ν_τ true-to-treco matrix and 100% oscillated spectrum is shown in Figure 6.8 for Run III as an example.

6.3.7 Extrapolating rock and anti-fiducial events

The standard beam matrix extrapolation philosophy does not hold for neutrinos interacting in the rock around the detector because the standard procedure depends on the mass of the target volume, and for RAF events the target volume is an infinite mass of rock rather than the fiducial volume of the Far Detector. The extrapolation procedure has been modified to allow prediction of the RAF visible energy spectral prediction. To predict RAFs, the procedure shown in Figure 6.5 is the same up to the point where the Far Detector flux is determined. The nominal Far Detector Monte Carlo is then corrected in bins of true neutrino energy by the ratio of the nominal flux to the predicted flux [100]. This is performed for each of the defined RAF regions.
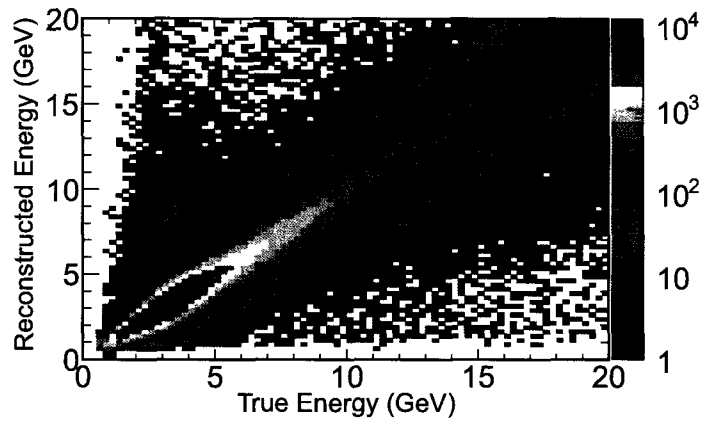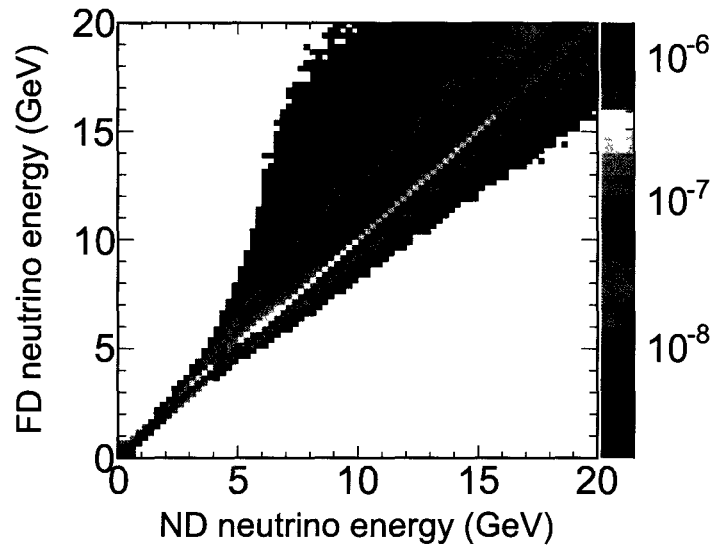
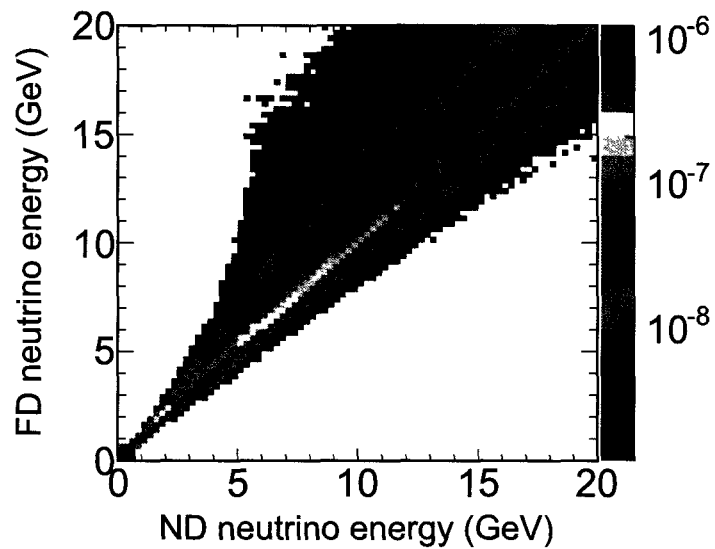(a) Near Detector Reco vs. True
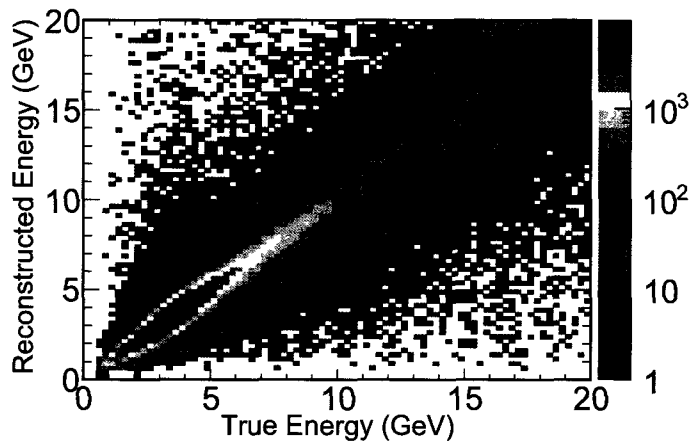
(b) Far Detector Reco vs. True

FIG. 6.3: Truth smearing matrices used for conversion of reconstructed CC ν_μ energy to true CC ν_μ energy for Monte Carlo simulations of Run III. The matrix 6.3(a) is used before extrapolation with a beam matrix and the matrix 6.3(b) is used after extrapolation.

(a)

(b)

FIG. 6.4: Beam matrices used to extrapolate the NuMI beam for Run III. On top, the beam matrix relating the flux measured by the Near Detector to the flux expected in the Far Detector. The color scale is logarithmic, and represents the relative normalization between the two detectors, due to $1/r^2$ effects. On bottom, the beam matrix used to extrapolate beam anti-neutrinos.
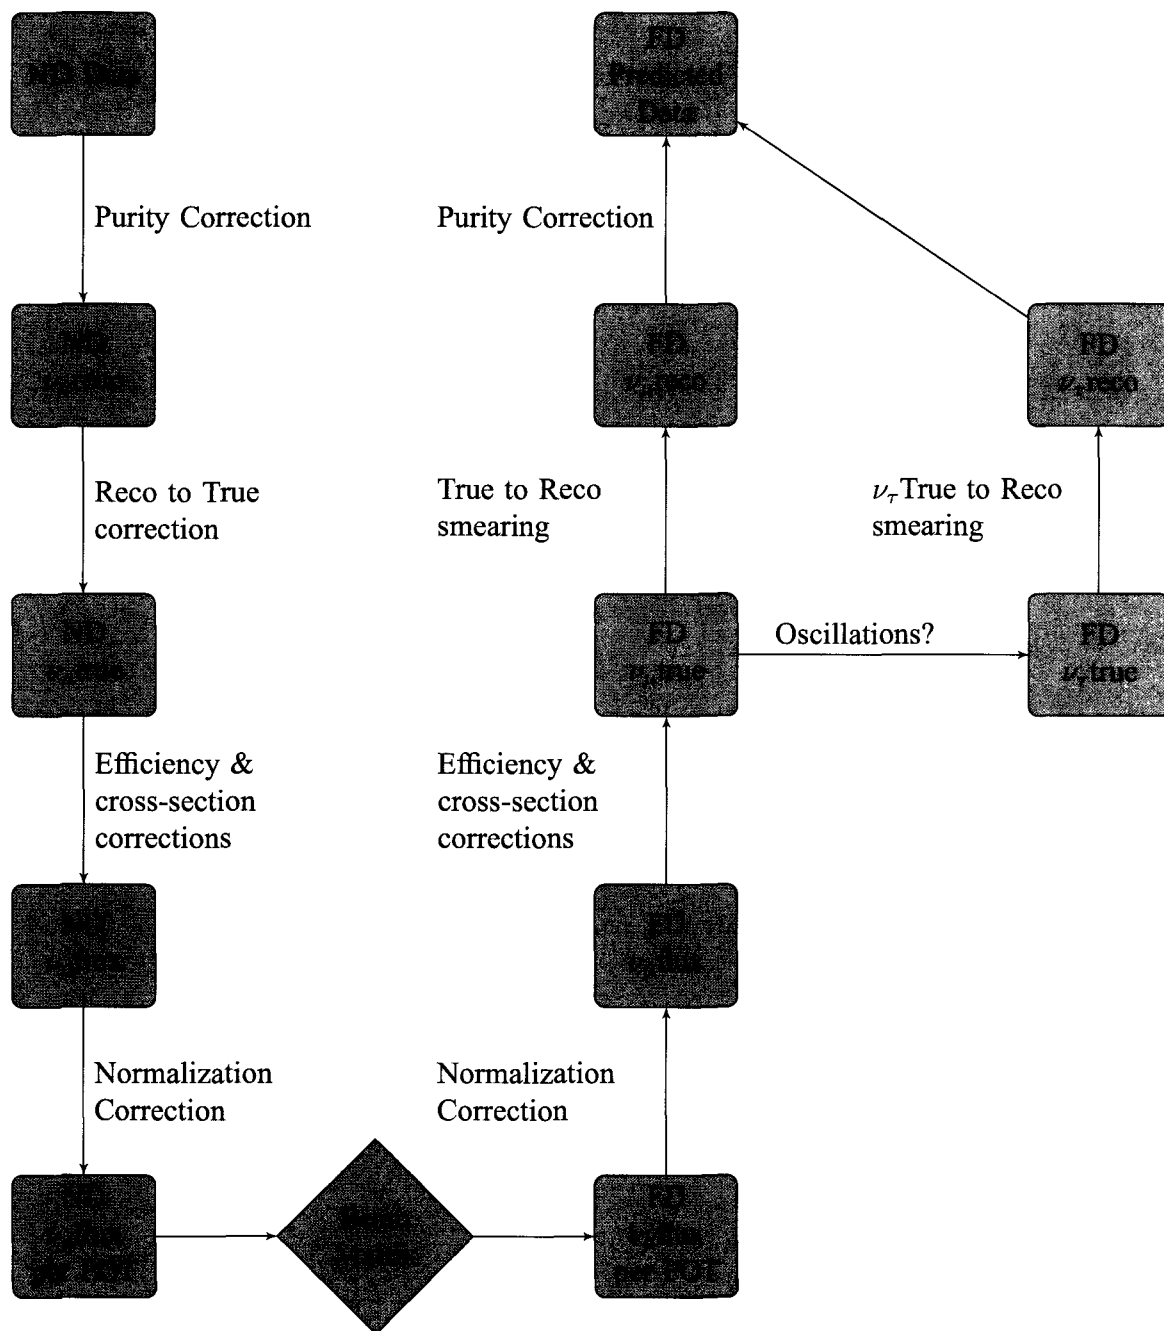
FIG. 6.5: The Matrix Method flow chart. Boxes are individual spectra created in the process, and arrows represent corrections derrived from Monte Carlo simulations applied to each spectrum.
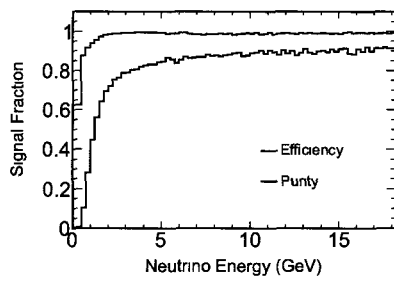
(a) Run 1
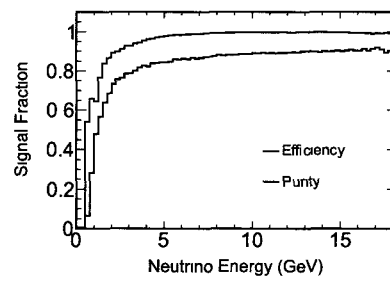
(b) Run 1 pHE
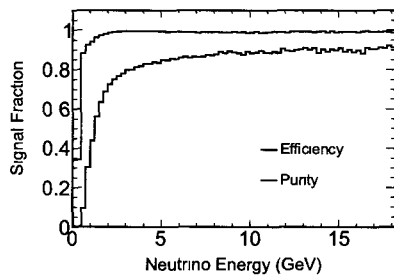
(c) Run 2
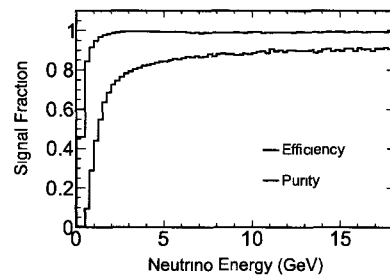
(d) Run 3

FIG. 6.6: The selection efficiency and purity in Monte Carlo samples by run for events in the Far Detector. Purity is a function of reconstructed neutrino energy and efficiency is a function of true neutrino energy.
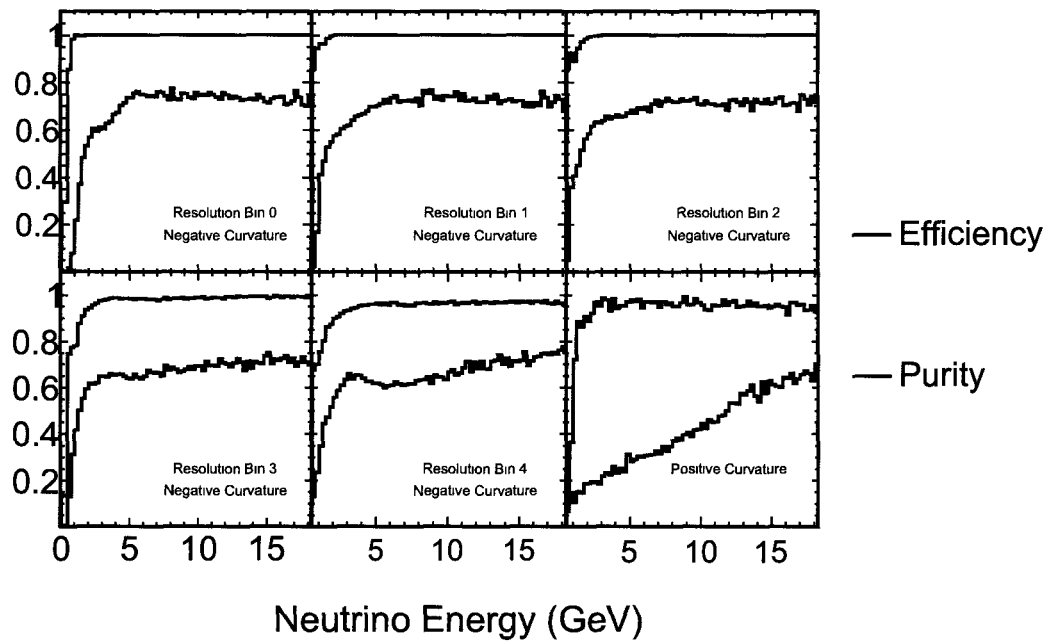
FIG. 6.7: The purity and efficiency of the event selection algorithm, shown separated in to five resolution bins. The purity is highest in bin 0, which containins events with the best resolution.

(a)

(b)
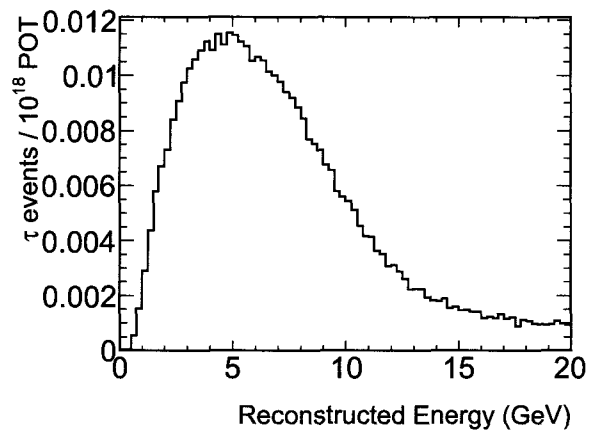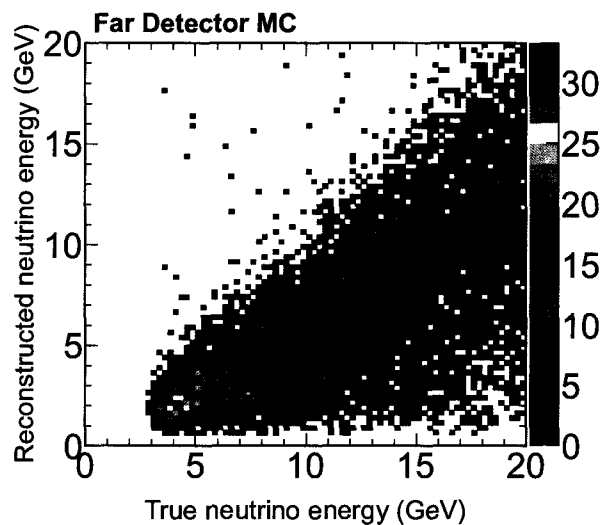
FIG. 6.8: The pure ν_τ spectrum before inverse oscillations are applied to model the background from ν_τ. The focusing peak shape is different for ν_τ events because of the reconstructed energy resoution for these events, as the ν_τ produced in the prompt τ-decay carries away some of the energy. The reconstructed energy resolution for ν_τ CC events is shown in the smearing matrix 6.8(b). The spectrum shown is from simulations of Run III.

6.4 Fitting

The Far Detector data is compared to a Far Detector predicted spectrum that may be altered to incorporate one of the neutrino disappearance models. Goodness of fit between these two spectra is calculated by a log-likelihood calculation. With N_E energy bins, the log-likelihood function is:

$$\chi^2 = -2\ln L = 2\sum_j^{N_E} \left(N_j^{MC} - N_j^{Data} + N_j^{Data} \ln\left(\frac{N_j^{Data}}{N_j^{MC}}\right) \right) + \sum_k^{N_{Syst}} \left(\frac{a_k^2}{2\sigma_k^2}\right) \quad (6.19)$$

where a_k are the systematics considered as nuisance parameters in the fit, with 1σ errors σ_k.

In the extrapolation methods described above, oscillation parameters can be chosen and oscillations may be applied to the Far Detector prediction, reducing the number of expected muon neutrinos. Once oscillations have been applied to the prediction, the value for χ^2 is re-calculated between prediction and data. The set of oscillation parameters which yield the smallest value for χ^2 can be found with a searching algorithm, such as MINUIT, or with a grid search over a wide space of parameters.

The oscillation parameters governing neutrino disappearance are not the only parameters that can be considered in this minimization. The effects of systematic uncertainties can play a role here as well. For example, in both the Matrix Method and the Far/Near method the size of the NC background is predicted from simulations. The uncertainty of the NC fraction allows for a range in the size of this background, and the size of the NC

background can change the value of χ^2, even with the oscillation parameters fixed.

We can reduce the magnitude of certain systematic uncertainties on the overall fit by fitting these systematics as extra nuisance parameters. Allowing the fit to minimize χ^2 for parameters of a particular model is called profiling (though this procedure is commonly and mistakenly called "marginalizing"). When a grid search is performed, the values of the nuisance parameters are allowed to float, within physical boundaries, so that the value of χ^2 at any point in parameter-space is minimized.

The results of the χ^2 minimization are reported in the 2-dimensional space of oscillation parameters Δm^2 and $\sin^2(2\theta)$. The procedure yields a pair of oscillation parameters which minimize the value of χ^2, but must be presented with a confidence interval representative of the significant statistical errors expected with this experiment. The confidence interval is reported in terms of the 68% and 90% confidence level (CL) contours, representing the boundaries in parameter-space containing the results that would be found 68% or 90% of the time on repeated experiments. In practice the statistical errors in MINOS dominate the identified systematic errors, and so the smaller systematic errors are reported separate from the CL contours.

Generally, the coverage of the 90% CL contour increases when systematics are included in the fit, as a coordinate just outside of the contour on a statistics-only fit can inflate the nuisance term to minimize the value of χ^2 to fit within 90% once systematic shifts are allowed.

6.4.1 Fitting with resolution binning

Because the Far Detector data is being split up five ways, this is equivalent to fitting five separate experiments, each with fewer events. Allowing the oscillation fit to wander far into the unphysical region with one-fifth the statistics leads to negative bin weight predictions, which are ignored by the likelihood function. This leads to plateaus in the unphysical region for individual resolution bin χ^2 surfaces. Resolution bins with a large proportion of background in the oscillation region do not see this phenomenon.

The FD data binning scheme for this fit is: 1 bin of 0.5 GeV between 0 GeV and 0.5 GeV, 78 bins of 0.25 GeV between 0.5 GeV and 20.0 GeV, 10 bins of 1 GeV between 20 GeV and 30 GeV, 10 bins of 2 GeV between 30 GeV and 50 GeV, and one bin of 150 GeV between 50 GeV and 200 GeV. This is 100 bins altogether. This binning scheme is used for each resolution bin as well as the positive curvature bin, so for each run there are 600 degrees of freedom (DOF). The positive curvature sample in the pHE run is insignificant in this scheme, so that run only has 500 DOF. The total fiducial sample over three LE runs (600 DOF each) and 1 pHE run (500 DOF) has 2300 bins, but fitting in a two parameter space reduces the number of DOF to 2298.

The full log-likelihood function, taking into account the number of resolution bins and $\bar{\nu}_\mu$-like events, is:

$$\chi^2 = -2\ln L = 2 \sum_i^{N_{Res}+1} \sum_j^{N_E} \left(N_{ij}^{MC} - N_{ij}^{Data} + N_{ij}^{Data} \ln\left(\frac{N_{ij}^{Data}}{N_{ij}^{MC}}\right)\right) + \sum_k^{N_{Syst}} \left(\frac{a_k^2}{2\sigma_k^2}\right)$$

(6.20)

where a_k are the systematics considered as nuisance parameters in the fit, with 1σ errors

σ_k. There is now a sum over $N_{Res} + 1$, which is the number of resolution bins plus the one sample of $\bar{\nu}_\mu$-like events.

6.4.2 Why resolution binning works

Even though all five resolution bins yield the same predicted reconstructed energy spectrum in the absence of oscillations, the underlying true spectra are different for each resolution bin. For comparison, plots of MC reconstructed energy vs. true energy for the best resolution bin, Bin 0, and the poorest resolution bin, Bin 4, are shown in Figure 6.10(a) and Figure 6.10(b), respectively. Removing the NC background and poorly-resolved events from Bins 0, 1, and 2 allows the smearing matrices for these bins to be more diagonal. In these bins the oscillation dip is able to be resolved more precisely. The poorest-resolution Bin 4 has some sensitivity to oscillations, but at worst adds a flat χ^2 contribution to each energy bin when all resolution bins are fit simultaneously. The difference in oscillation dip resolution between bins is shown in Figure 6.9.

Neutrino Disappearance
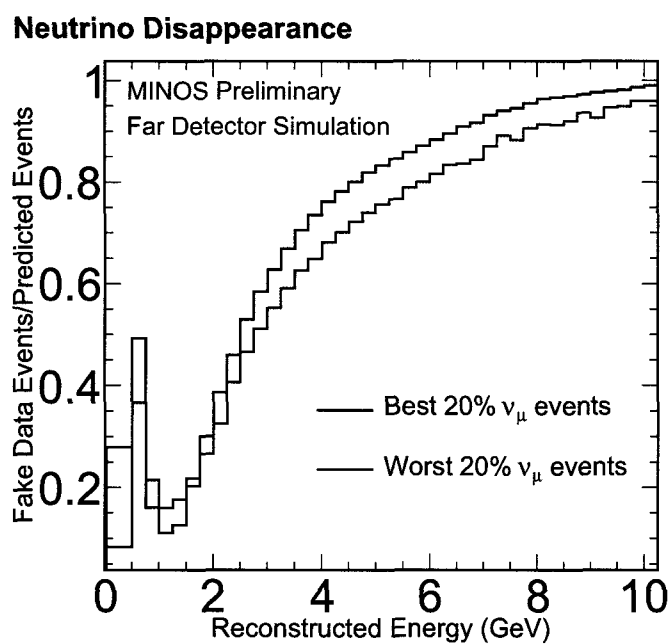
FIG. 6.9: The oscillation dip measurements with two different resolution bins. Bin 0 contains events reconstructed with the most precise energy resolution. Bin 4 contains events reconstructed with the poorest energy resolution. With oscillated fake data, the sample in bin 0 is able to resolve the oscillation dip deeper than bin 4.
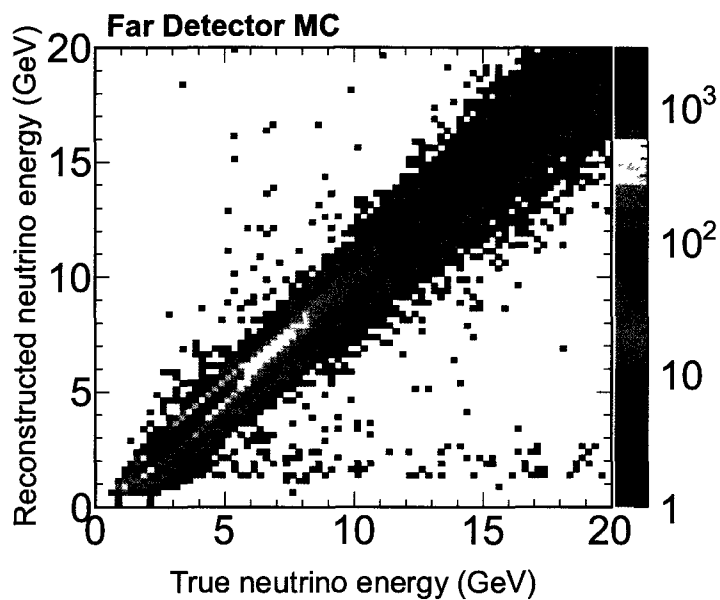
(a) Bin 0
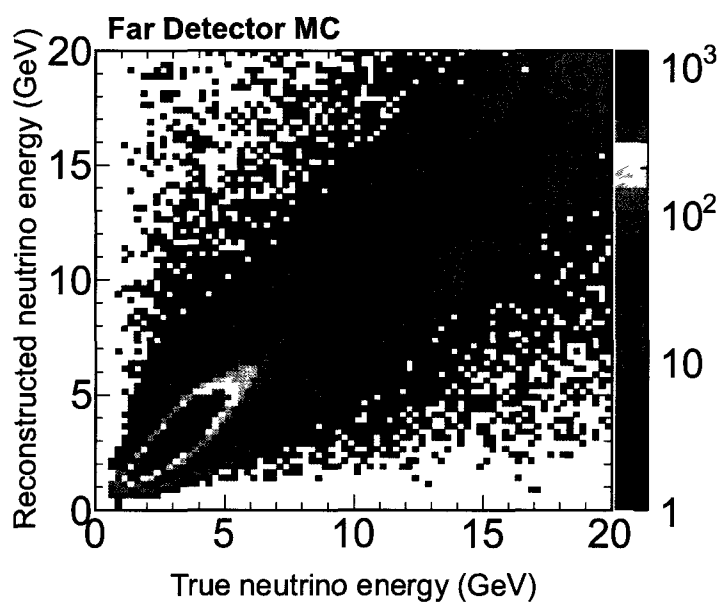
(b) Bin 4

FIG. 6.10: The smearing matrix, similar to fig.6.3, for resolution bin 0, the bin with the 20% of events with the lowest and 20% poorest σ_{Total}/reconstructed neutrino energy. The color scale represents the number of events in each bin, and is logarithmic.

6.4.3 Alternative shape-only fit

An alternate log-likelihood function can be constructed that separates the contributions from shape by resolution bin and overall normalization across all resolution bins. This log-likelihood function is:

$$-\ln L = N^{MC} - N^{Data} \ln N^{MC} + \sum_{i}^{n_{Res}} \sum_{j}^{n_E} \left(\mu_j^i - n_j^i \ln \mu_j^i \right) + \sum_{k}^{n_{syst}} \frac{a_k^2}{2\sigma^2} \qquad (6.21)$$

where μ_j^i is the normalized prediction in energy bin j and resolution bin i and n_j^i is the normalized data in the same bin. The $N^{MC} - N^{Data} \ln N^{MC}$ term is integrated over all resolution bins and all energy bins. This log-likelihood function was considered to mitigate the effect of systematics that affect different resolution bins differently. For example, since most NC events reside in bin 4, a shift in the NC background expectation changes the spectrum in the Near Detector and in bin 4. No improvement was seen by using this alternate log-likelihood function when systematics were considered.

6.4.4 Fitting Frameworks

Two frameworks are used to fit the data against various neutrino disappearance models. The NuSystFitter algorithm, part of the MINOS software framework NtupleUtils, is able to fit spectra to different models very quickly but is unable to include nuisance parameter terms for certain types of systematics. NuSystFitter uses MINUIT to search for best-fit oscillation parameters and to profile over systematic uncertainty nuisance parameters. Systematic uncertainties involving energy resolution are problematic in with

`NuSystFitter`. These systematics uncertainties, like hadronic shower energy scale, migrate events back and forth between energy bins. `MINUIT` returns jagged, discontinuous χ^2 surfaces with these systematics, which are not representative of the true confidence level coverage. This led to the development of a second algorithm to handle these nuisance parameters.

The second algorithm, called `GhostFitter`, is a stand alone package that fits data and includes systematic shifts in a different way. The `GhostFitter` characterizes various systematics with systematically shifted templates. A template is a two-dimensional Far Detector predicted neutrino energy spectrum from fake data, where the two dimensions are reconstructed energy and true neutrino energy. Templates are produced for nominal simulations, as well as simulations that have been systematically shifted $\pm 1\sigma$ and $\pm 2\sigma$ in each particular systematic uncertainty. The `GhostFitter` algorithm interpolates between these five templates for non-integer σ values of systematic shifts.

The `GhostFitter` algorithm is a recent development, but is the primary fitting algorithm for the eventual analysis of Far Detector data. Since it is new, both the `NuSystFitter` and `GhostFitter` will be used to fit the Far Detector data with statistical errors only, as well as the nuisance parameters that are well-behaved within both algorithms. The `GhostFitter` results are considered satisfactory if they lie within the statistical error. The `NuSystFitter` will be used to evaluate the systematic errors due to all systematics, as described in Chapter 7.

CHAPTER 7

Sensitivities and systematic errors

Using the techniques described in Chapter 5 and the Beam Matrix method described in Chapter 6, we may now make a prediction of the Far Detector neutrino energy spectrum, in the absence of oscillations, and calculate our statistical sensitivity to neutrino oscillations.

7.1   Far Detector Prediction

The Near Detector neutrino energy spectrum that is used to predict the Far Detector expectation was shown as Figure 3.1. It is shown here broken down by run in Figure 7.1. The total exposure of the Near Detector data set is shown in Table 7.1. Corresponding predictions for Far Detector spectra are shown in Figure 7.2. Both the Far/Near extrapolation and the beam matrix method are used to predict the Far Detector spectrum for fiducial events. This is shown in Figure 7.3. The two extrapolations differ by $\pm 5\%$ below 5 GeV,

Run	MC Exposure ($\times 10^{18}$ POT)	Data Exposure ($\times 10^{18}$ POT)
Run I LE	49.82	128.56
I pHE	3.311	15.62
II LE	66.98	181.60
III LE	102.1	359.55
Total	222.2	685.33

TABLE 7.1: Total beam exposure in the Near Detector, by run, in both data and Monte Carlo simulations.
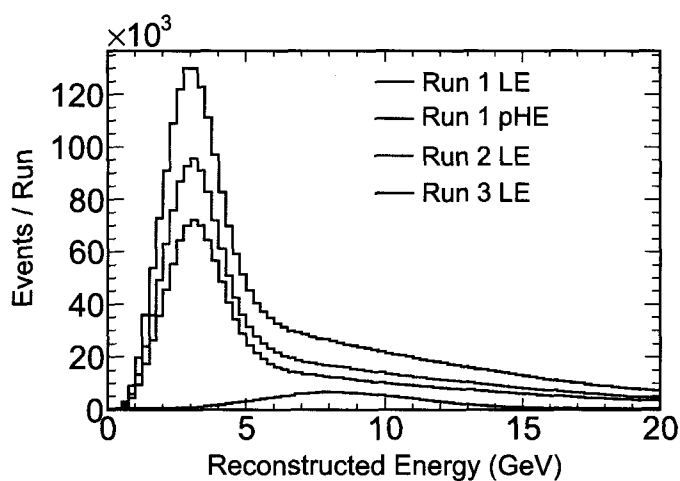
FIG. 7.1: The Near Detector data accumulated by run. These are the spectra that are extrapolated, using either the Far/Near method or the beam matrix method, to predict the Far Detector energy spectrum.

which is the behavior we expect between the two methods [105]. Here the Far/Near method is predicted without resolution binning, while the beam matrix method prediction is the sum of all resolution bin predictions. The sum of resolution bin predictions before oscillations is identical to the prediction without resolution bins.
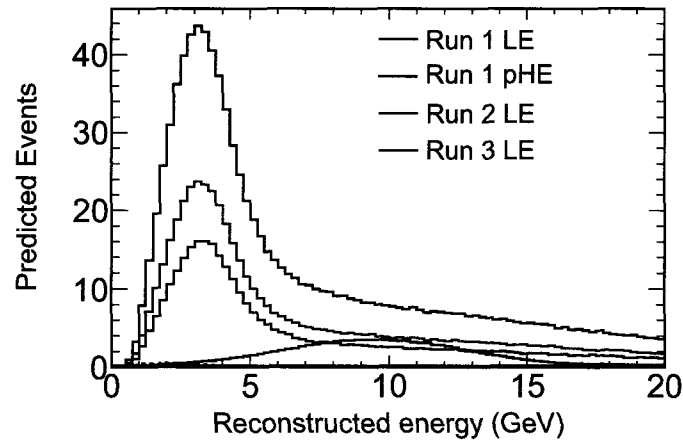
FIG. 7.2: The Far Detector predicted spectra, predicted with the beam matrix method, for each run period. NC background expectation per run period is plotted with dashed lines, with line color corresponding to the run.
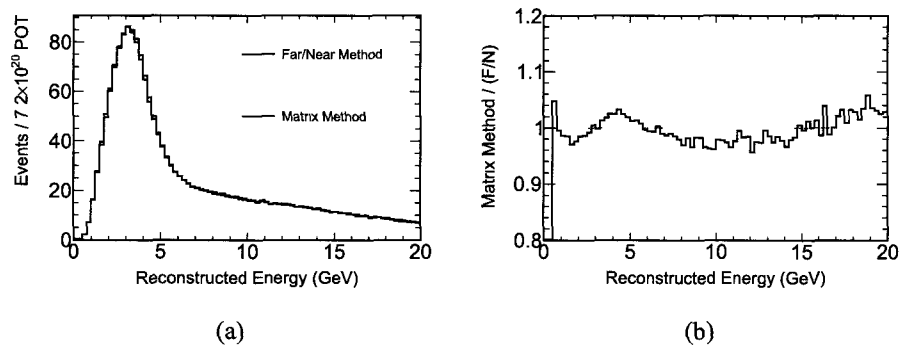
FIG. 7.3: Far Detector predictions from the Far/Near extrapolation and Matrix Method extrapolation, for Runs I, II, III, and Run I pHE. For simplicity, resolution binning is not used in either prediction shown here. The $\approx 5\%$ differences near the focusing peak are consistent with previous comparisons [105].

7.2 Backgrounds

The MINOS experiment produces a clean signal with few backgrounds. The single largest background is Neutral Current interactions which mimic high-y Charged-Current interactions. The neutrino Neutral Current cross-section is the same for all three active neutrino flavors, and so the NC expectation is only a function of neutrino flux, regardless of oscillations. The Neutral Current background is minimized through event selection criteria, and the remaining background is predicted through the extrapolation process and removed from the final result, as described in Chapter 6. NC events are modeled by removing muon tracks from CC events, leaving behind only hadronic showers. Comparisons with these events between data and Monte Carlo show an excess of $(6 \pm 15)\%$ in data [106].

In the scheme of resolution binning, NC events tend to be sequestered in Bin 4, the resolution bin containing events with the poorest energy resolution. This is because NC events look like high-y CC events, and the hadronic shower component of events has the poorest energy resolution. The NC events which pass selection cuts contain short π^+ and π^- tracks, and contribute roughly the same number of background events to the ν_μ-like and $\bar{\nu}_\mu$-like samples. Since the overall flux is smaller for the $\bar{\nu}_\mu$-like sample considered in this analysis, NC events are a more significant background. This was apparent in Figure 5.13. The predicted number of events for each resolution bin is shown in Table 7.2.

Tau appearance events are the next largest background in the MINOS detectors, and are the only other background that this oscillation analysis models in the fit. Tau neutrino appearance was also described in Chapter 6. The tau neutrino background expectation

is predicted as a function of neutrino flux and the oscillation parameters during the fit. For the Far Detector prediction in the absence of oscillations, the predicted number of ν_τ events is approximately 0 in the ν_μ beam. With the neutrino oscillation parameters measured in [47], and ν_μ oscillating into ν_τ, the expectation is 3.2 events for the total beam exposure in runs I, II, and III.

The Far Detector is placed far underground to minimize the flux of cosmic ray muons within the detector volume. Cosmic ray muons have a rate of 0.2 Hz at the depth of the Far Detector. The cosmic ray muon background is further minimized with a 14 μs timing cut around NuMI spill times, which minimizes the total livetime of the detector. With 2.892×10^7 spills [107], the Far Detector livetime susceptible to cosmic ray events is to 405 seconds, so the total number of expected cosmic ray events is roughly 80. A further cut is made on the angle of a muon track with respect to the beam (described in Chapter 4), further reducing the significance of the background. The expected background was evaluated in [47] and found to be < 0.5 events at 68% C.L. With roughly double the amount of spills, we can assume a negligible expected cosmic ray neutrino background of < 1 event.

Rock muons can also be a background in the Far Detector if the muons pass through the anti-fiducial region without recording a hit in the scintillator. Lacking the appearance of a hadronic shower in the detector, they would appear to be a muon created from a quasi-elastic CC interaction. To quantify the rock muon background expectation, rock muon MC is passed through the fiducial region's selection criteria. Since quasi-elastic interactions are typically high-resolution events, the background tends to populate Bin 0,

Resolution Bin	NC Bknd.	Rock Bknd.	ν_τ Bknd.
Bin 0	0.09	4.96	0.39
Bin 1	0.27	0.79	0.48
Bin 2	0.65	0.36	0.54
Bin 3	6.33	0.38	0.66
Bin 4	12.6	0.50	0.81
Positive Curvature	21.10	1.11	0.33
Total	41.0	8.1	3.2

TABLE 7.2: Expected backgrounds in each resolution bin for Runs I, II, III, and pHE, a total of 7.2×10^{20} POT. The ν_τ background is calculated using oscillation parameters from [47].

the resolution bin containing high-resolution events. The Rock background is calculated by running special rock simulated events through the normal set of fiducial volume and muon angle cuts, and then scaling the result by the ratio of the simulated beam exposure to the data beam exposure. The expectation for 7.2×10^{20} POT is shown in Table 7.2.

7.3 Statistical Sensitivity

Given the relatively small number of expected events in the MINOS experiment, it is important to estimate the expected statistical error to compare to the total systematic error. It provides a metric for measuring the gain for each of the analysis improvements described above. Also, if MINOS ever reaches the point where statistical error is smaller than the systematic error, then there is little to gain from taking more data. As much effort has been spent to reduce the systematic uncertainties on the oscillation result, the analysis presented in this thesis is statistically limited.

To calculate the expected statistical uncertainty, Monte Carlo simulated fake data is oscillated with values for Δm^2 and $\sin^2(2\theta)$ that are near the expected true values. These

oscillations are applied as a function of the simulated events' true energy. The simulated fake data exposure is 1000 times larger than the true exposure, so the simulated fake data energy spectrum is scaled down to the same exposure as the data sample, to produce a Far Detector data-like energy spectrum with minuscule statistical errors. Fake data generated in this manner is used to calculate sensitivities and evaluate systematic errors, as discussed later in this chapter.

7.3.1 One-dimensional statistical sensitivity

The goal of this experiment is to measure two parameters, Δm^2 and $\sin^2(2\theta)$, but it is important to examine our sensitivity to each separately. For this we evaluate the one-dimensional statistical sensitivity for each of our variables. The one-dimensional sensitivity defines the statistical error on each oscillation parameter individually.

The value for χ^2, as defined in Equation 6.20, is calculated between the oscillated prediction and the oscillated fake data for many steps along Δm^2. At each step in Δm^2, MINUIT is allowed to profile over values of $\sin^2(2\theta)$ as a nuisance parameter, finding the minimum possible χ^2. The minimum for each step is recorded, and the range of Δm^2 where $\Delta\chi^2 < 1.0$, relative to the minimum χ^2 calculated, constitutes the 1σ sensitivity of the experiment to Δm^2. A similar procedure is carried out to calculate the sensitivity of the experiment to $\sin^2(2\theta)$.

The $\Delta\chi^2$ sensitivities for Δm^2 and $\sin^2(2\theta)$ are shown in Figure 7.4, comparing the statistical sensitivity expected from the total 7.2×10^{20} POT dataset to the statistical sensitivity expected with the implementation of resolution binning, and also the inclusion

of $\overline{\nu}_\mu$-like events. These represent a 1.8% improvement in the Δm^2 statistical sensitivity

and a 5.7% improvement in $\sin^2(2\theta)$ statistical sensitivity.
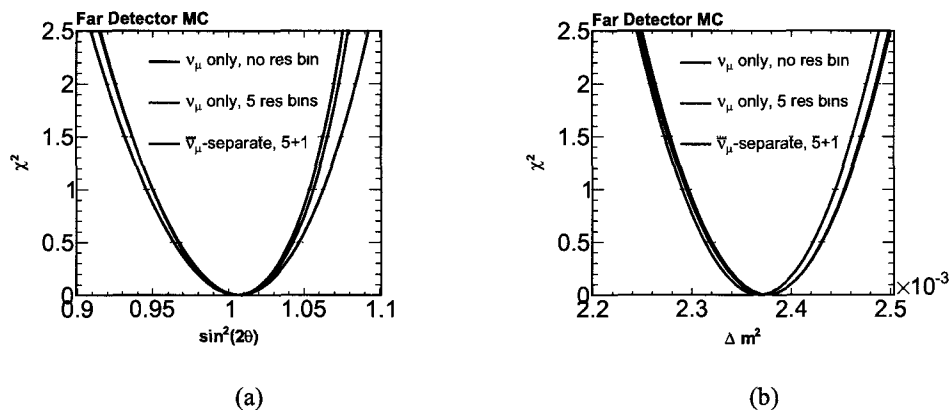

(a)   (b)

FIG. 7.4: The one-dimensional statistical sensitivities for the measurements of Δm^2 and $\sin^2(2\theta)$ with a 2008-style analysis, with an implementation of resolution binning for ν_μ-like events, and then with the inclusion of $\overline{\nu}_\mu$-like events as a separate sample. Fake data oscillations defined at $\Delta m^2 = 2.43 \times 10^{-3} \text{eV}^2/\text{c}^4$ and $\sin^2(2\theta) = 1.0$.

The total statistical sensitivity is shown in Figure 7.5. Assuming the same oscillation

parameters as measured in [47], the expected 1σ statistical error is

$$\delta(\Delta m^2) = ^{+0.12}_{-0.11} \times 10^{-3} \text{eV}^2/\text{c}^4 \text{ and } \delta(\sin^2(2\theta)) = \pm 0.05.$$

7.3.2 Two-dimensional statistical sensitivity

To produce a two-dimensional statistical sensitivity, the fake data set described above

is subject to a grid search over Δm^2 and $\sin^2(2\theta)$ to produce a χ^2 surface. Relative to

the minimum of the surface, isolines of $\Delta\chi^2$ where the values are 2.3 or 4.61 yield 68%

or 90% confidence level contours, respectively. These contours describe the sensitivity of

the experiment to the two oscillation parameters, given that the true physics parameters

are close in $(\Delta m^2, \sin^2(2\theta))$-space. This two-dimensional statistical sensitivity procedure
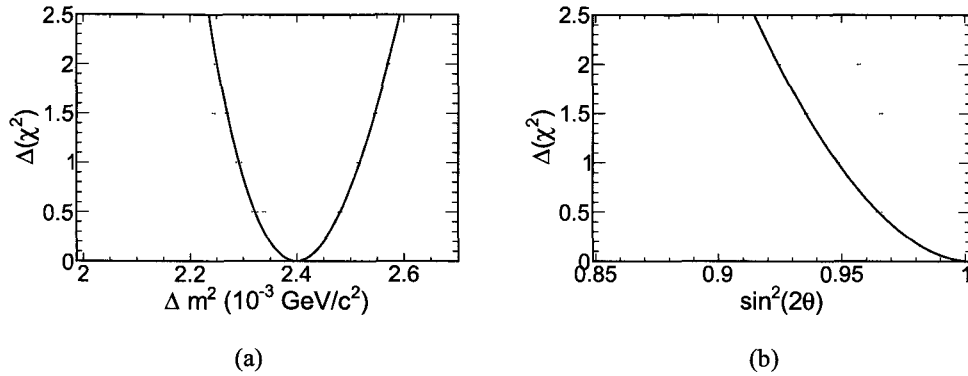
FIG. 7.5: The one-dimensional statistical sensitivities for the measurements of Δm^2 and $\sin^2(2\theta)$ achieved with all of the techniques developed in Chapter 5. The expected statistical error are found to be $\delta(\Delta m^2) =^{+0.12}_{-0.11} \times 10^{-3} \mathrm{eV}^2/\mathrm{c}^4$ and $\delta(\sin^2(2\theta)) = \pm 0.05$. Fake data oscillated at $\Delta m^2 = 2.43 \times 10^{-3} \mathrm{eV}^2/\mathrm{c}^4$ and $\sin^2(2\theta) = 1.0$.

is carried out many times, with the techniques described in Chapter 5 turned on in series. In this way we can see the effectiveness of each of these analysis improvements.

The gain in oscillation sensitivity achieved when resolution binning is used and when $\overline{\nu}_\mu$-like events are fit is shown in Figure 7.6. The oscillation sensitivity of RAF events as a separate sample are shown in Figure 7.7. The result, Figure 7.8, shows the net gain in sensitivity to the two oscillation parameters presented in this document. The effect of the secondary event selection and the inclusion of $\overline{\nu}_\mu$-like events is too small to be shown on this plot.
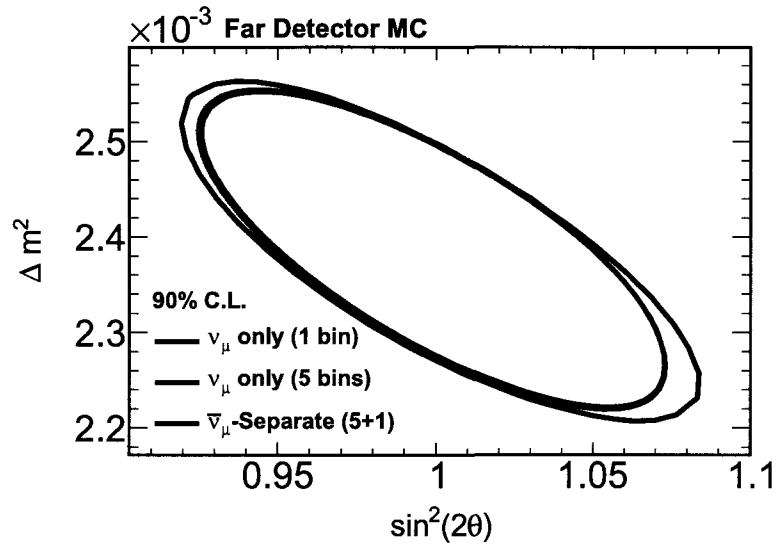
FIG. 7.6: The oscillation sensitivity without the improvements described in Chapter 5 (black contour), shown with the oscillation sensitivity when resolution binning is utilized (blue contour), and then with resolution binning and the inclusion of $\bar{\nu}_\mu$-like events (red contour).
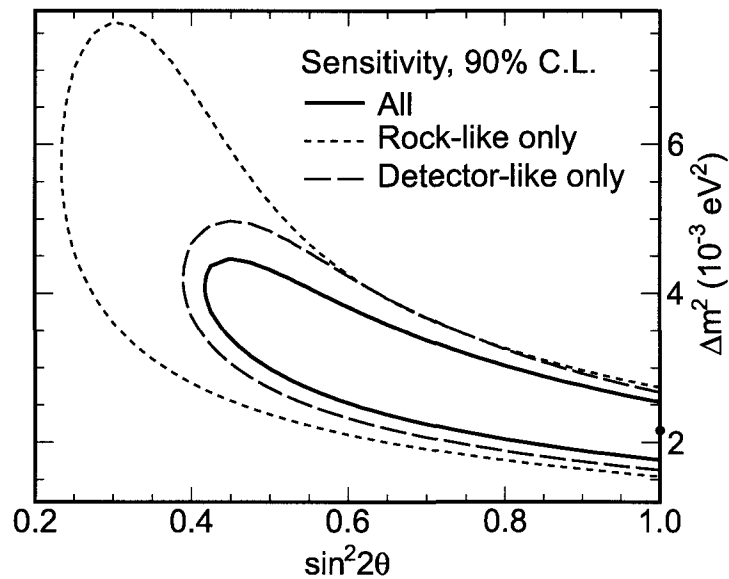
FIG. 7.7: Oscillation sensitivity of RAF sample, shown separately for events with vertices in the rock (red contour) and in the detector anti-fiducial regions (black dashed). The solid black line shows the combined oscillation sensitivity of these two samples [99]

FIG. 7.8: The improvement of statistical sensitivities as each of the analysis techniques described in this chapter are implemented in succession.

7.4 Systematics

Because of the truth-unfolding process involved in the beam matrix extrapolation method, it is difficult to propagate systematic errors algebraically. The solution for determining the systematic error associated w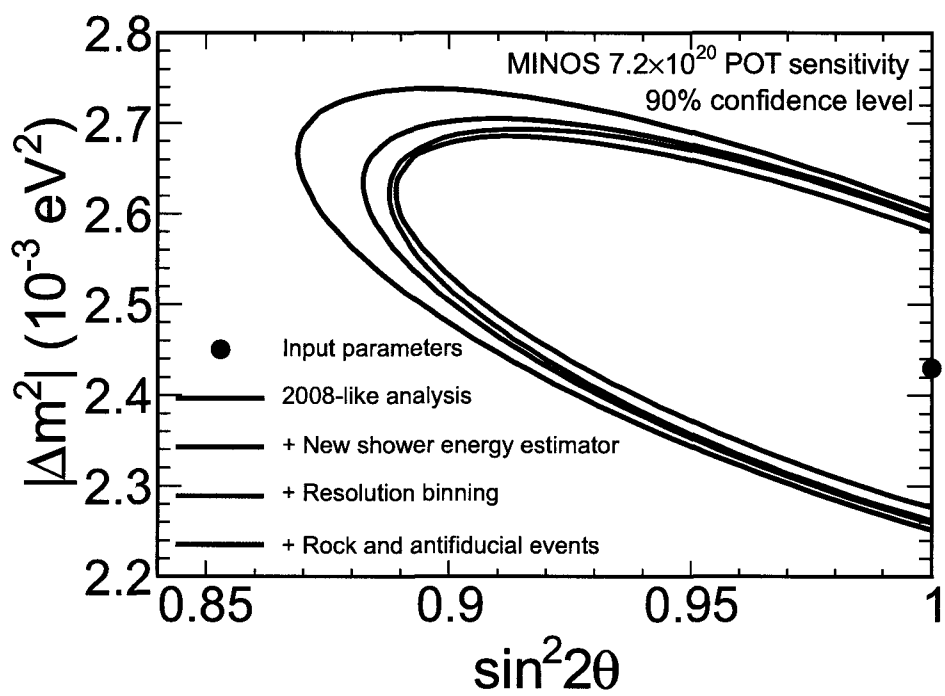ith the final result is to fit systematically shifted fake data and record the deflection of the best-fit point from the true oscillation parameters used to generate the fake data. The fake data is shifted $\pm 1\sigma$ for each systematic uncertainty. The systematic uncertainties that have the largest effect on the fit are then included as nuisance parameters for the final fit.

The systematic uncertainties and their values are:

- **Normalization 1.6%.** A identical scale factor is applied to every energy bin in the Far Detector fake data spectrum. This comes primarily from a 1.3% selection bias between the Near and Far Detectors, as determined from visual scans of events [108]. The remainder comes from fiducial mass biases, related to the spatial definitions in Monte Carlo and in data, and the steel thickness measurements described in Section 2.2.1.

- **NC Background Normalization 20%.** Scaling the predicted NC background in the Near and Far Detectors, fully correlated. The size of the uncertainty are determined from studies of data and Monte Carlo described in Section 7.2. The real NC background after selection cuts have been applied is energy-dependent and dominant at energies $< 5\,\text{GeV}$, so this is a conservative estimate.

- **Shower Energy Normalization.** This is an energy-dependent error that includes a 5.7% error band from hadronic energy calibration, as well as an energy-dependent

error band from NEUGEN. The NEUGEN errors are related to nuclear effects that are not well-modeled, and are larger at low E_{shw}. While this error was evaluated with calorimetric shower energy, the kNN shower energy estimation relies on calorimetric shower energy variables for training. The relative difference between correlated shifts in the Near and Far Detectors is shown in Figure 7.9.
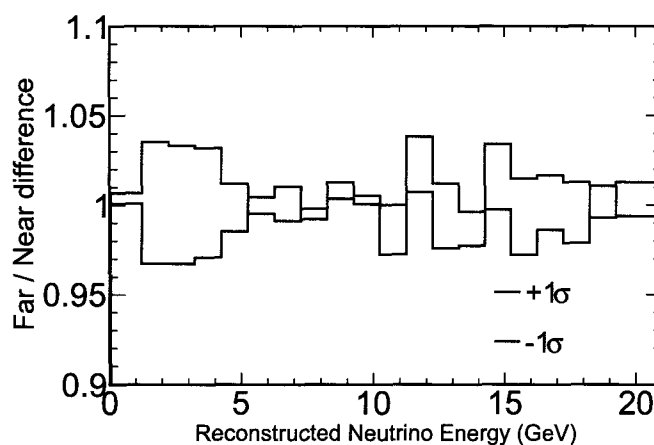
FIG. 7.9: The $\pm 1\sigma$ error band for shower energy scale systematic, expressed as a relative difference between Near Detector and predicted Far Detector spectra. This sytematic error has the largest net effect on the final oscillation fit.

- **Near Detector Shower Energy Estimation 1.9%.** A scale factor applied to the shower energy of events in the Near Detector. This comes from the difference between simulated and observed calorimetric response, relative to the absolute shower energy scale, as described in Section 4.5 and in [95].

- **Far Detector Shower Energy Estimation 1.1%.** A scale factor applied to the shower energy of events in the Far Detector, as described above, in Section 4.5, and in [95].

- μ **Track Energy estimation.** The 1σ error is estimated to be a 2% shift in muon

momentum from range and 1% shift in muon momentum from curvature, relative to the range error. These are taken to be fully correlated. These values come from studies done with range and curvature agreement in data and Monte Carlo [74], and were discussed in Chapter 4.

- **Beam Parameterization.** The 1σ error band on the beam tuning, as discussed in Chapter 3 and shown in Figure 3.2.

- $\bar{\nu}_\mu$ **Charge ID 40%.** The relative size of the sample of true ν_μ reconstructed with `charge> 0`. Conservatively, the entire 40% discrepancy between data and simulations,as described in Section 5.7, is assigned across all energies. Since $\bar{\nu}_\mu$ are assumed to oscillate with similar oscillation parameters to ν_μ, this large uncertainty does not manifest itself as a large effect on the final result.

- **Neutrino cross-sections.** Uncertainties in neutrino cross-sections manifest themselves as uncertainties in overall event yield, given a neutrino flux. A change in the ν_μ-nucleon cross-section would ideally cancel out in a two-detector experiment, but the truth-unfolding and smearing process in the beam matrix extrapolation method leaves a residual effect. Uncertainties exist with the total neutrino cross-section, as well as quasi-elastic, resonance, and DIS exclusive channels. The QE and RES channels are accessed by altering the value of M_A in the dipole approximation of the axial form factor [84].

$$F_A(q^2) \propto \frac{1}{\left(1 - \left(\frac{q}{M_A}\right)^2\right)^2} \tag{7.1}$$

One can change relative QE and RES cross-sections by modeling different values for

M_A^{QE} and MA_A^{RES}. The DIS region is accessed by warping the transition region between resonance and DIS interactions. Systematic uncertainties within the NEUGEN model used in Monte Carlo generation (Section 3.3.1) have been evaluated by the collaboration by fitting to the NEUGEN model to available data [109].

- ν_μ **CC Cross-Section 3.5%.** A scale factor on the total normalization of the ν_μ CC cross-section, applied to fluxes in both the Near and Far Detectors.

- $\bar{\nu}_\mu$ **CC Cross-Section 4%.** Same as above, but for $\bar{\nu}_\mu$ interactions only.

- M_a^{QE} **Cross-Section 15%.** Scaling the value of M_a^{QE} by $\pm 15\%$.

- M_a^{RES} **Cross-Section 15%.** Scaling the value of M_a^{RES} by $\pm 15\%$.

- M_a^{QE} $\bar{\nu}_\mu$ **Cross-Section 8%.** Increasing the value of M_a^{QE} by $\pm 15\%$ for $\bar{\nu}_\mu$'s only.

- M_a^{RES} $\bar{\nu}_\mu$ **Cross-Section 8%.** Increasing the value of M_a^{RES} by $\pm 15\%$ for $\bar{\nu}_\mu$'s only.

- **NEUGEN parameters** There are three parameters within the NEUGEN model that warp the resonance/DIS transition region. These parameters affect the multiplicty of the recoil system and are called "KNO multiplicity" parameters within NEUGEN. We carry that terminology in the following tables.

The shifts in the best-fit oscillation parameters induced by systematically shifted fake data indicate the affect and importance of a particular shift. The shifts of best-fit points for $\pm 1\sigma$ systematic shifts applied to the fiducial sample are collected in Fig.7.10. When the shifts are applied to the RAF sample and fit simultaneously with the fiducial sample,

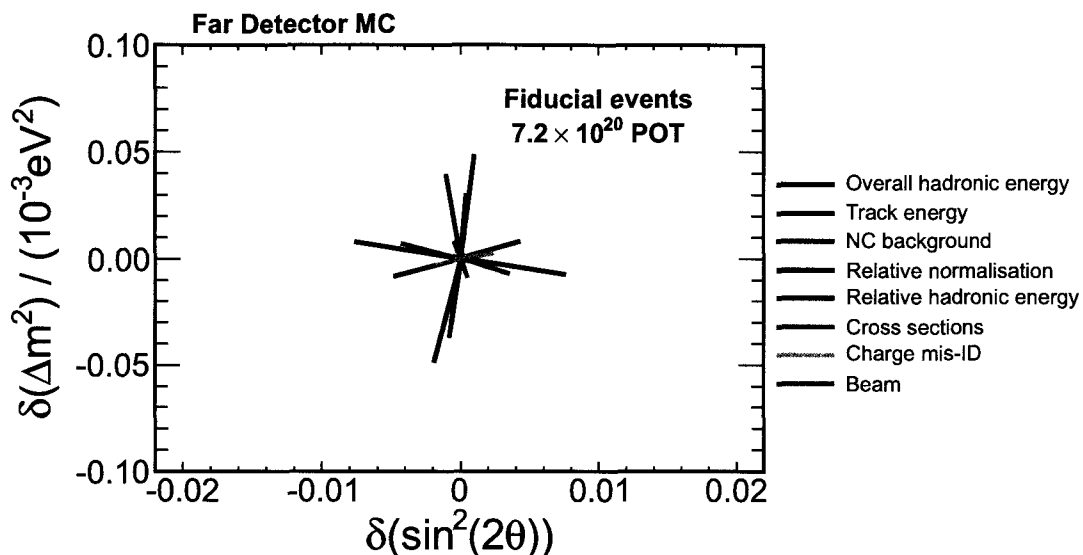the outcome is shown in Tab. 7.3 and graphically in Fig. 7.12.

FIG. 7.10: Graphical representation of systematic errors for oscillation measurements with only fiducial events. The lines indicate the shift in the oscillation best-fit point for high-statistics fake data when specific systematic shifts are applied.

The systematic uncertainties that apply to the RAF samples are assumed to be identical between each sub-region and fully correlated. The systematic uncertainties described above apply to the RAF samples. There are three additional systematic uncertainties that apply to the RAF sample only [99].

- **Rock Normalization 0.9%.** A normalization scale factor on all rock events, but not Anti-Fiducial events, taken from the uncertainty in the Z/A ratio of the rock surrounding the Far Detector.

- **Rock Cross-section 1%.** A scale factor on non-DIS and non-^{56}Fe cross-sections, taken from the uncertainty in the density of the rock surrounding the Far Detector.

- **Detector Edge.** 1σ shift on the strip alignment on the edges of the detector, altering

Systematic	Shift	Δm^2 $(10^{-3} eV^2)$	$\sin^2(2\theta)$	$\delta(\Delta m^2)$ $(10^{-3} eV^2)$	$\delta(\sin^2(2\theta))$
Truth	-	2.295	0.985	-	-
Best Fit	-	2.28485	0.98932	-0.01015	+0.00432
μ Track Energy	$+1\sigma$	2.32467	0.98822	+0.03982	-0.0011
	-1σ	2.24776	0.98846	-0.03709	-0.00085
Normalization	$+1.6\%$	2.25470	0.98863	-0.03015	-0.00068
	-1.6%	2.31553	0.98972	+0.03068	+0.00041
NC Background	$+20\%$	2.29254	0.98168	+0.00769	-0.00764
	-20%	2.27723	0.99690	-0.00762	+0.00758
Absolute Shower Energy	$+1\sigma$	2.33349	0.99028	+0.04864	+0.00096
	-1σ	2.23618	0.98737	-0.04867	-0.00195
ND Shower Energy	$+1.9\%$	2.28247	0.98556	-0.00238	-0.00376
	-1.9%	2.28713	0.99295	+0.00228	+0.00363
FD Shower Energy	$+1.1\%$	2.29250	0.99163	+0.00765	+0.00232
	-1.1%	2.27683	0.98691	-0.00802	-0.00241
Total CC Cross-section	$+3.5\%$	2.28355	0.99059	-0.00130	+0.00128
	-3.5%	2.28625	0.98794	+0.00140	-0.00138
M_A^{QE}	$+15\%$	2.28574	0.99147	+0.00089	+0.00216
	-15%	2.28277	0.98734	-0.00208	-0.00198
M_A^{Res}	$+15\%$	2.27956	0.99228	-0.00529	+0.00296
	-15%	2.29021	0.98616	+0.00536	-0.00316
kno Multiplicity 2	$+0.1$	2.28084	0.99080	-0.00401	+0.00148
	-0.1	2.28909	0.98773	+0.00424	-0.00159
kno Multiplicity 3	$+0.2$	2.28485	0.98932	+0.0000	+0.0000
	-0.2	2.28544	0.98916	+0.00059	-0.00016
$\bar{\nu}_\mu$Cross-section	$+4\%$	2.28414	0.98943	-0.00071	+0.00011
	-4%	2.28556	0.98912	+0.00071	-0.00020
$\bar{\nu}_\mu$QE Cross-section	$+8\%$	2.28452	0.98937	-0.00033	+0.00005
	-8%	2.28517	0.98924	+0.00032	-0.00008
$\bar{\nu}_\mu$Res Cross-section	$+8\%$	2.28432	0.98939	-0.00053	+0.00008
	-8%	2.28537	0.98918	+0.00052	-0.00014
$\bar{\nu}_\mu$kno Multiplicity 2	$+0.2$	2.28374	0.98951	-0.00111	+0.00019
	-0.2	2.28579	0.98910	+0.00094	-0.00022
Beam tuning	$+1\sigma$	2.29309	0.98884	+0.00824	-0.00048
	-1σ	2.28578	0.98980	-0.00907	+0.00048
$\bar{\nu}_\mu$Wrong-Sign	$+40\%$	2.28178	0.98750	-0.00307	-0.00182
	-40%	2.28739	0.98168	+0.00254	+0.00236

TABLE 7.3: Systematic errors for the combined analysis, Runs I, II, and III, LE and pHE, extrapolated NQ and PQ events.
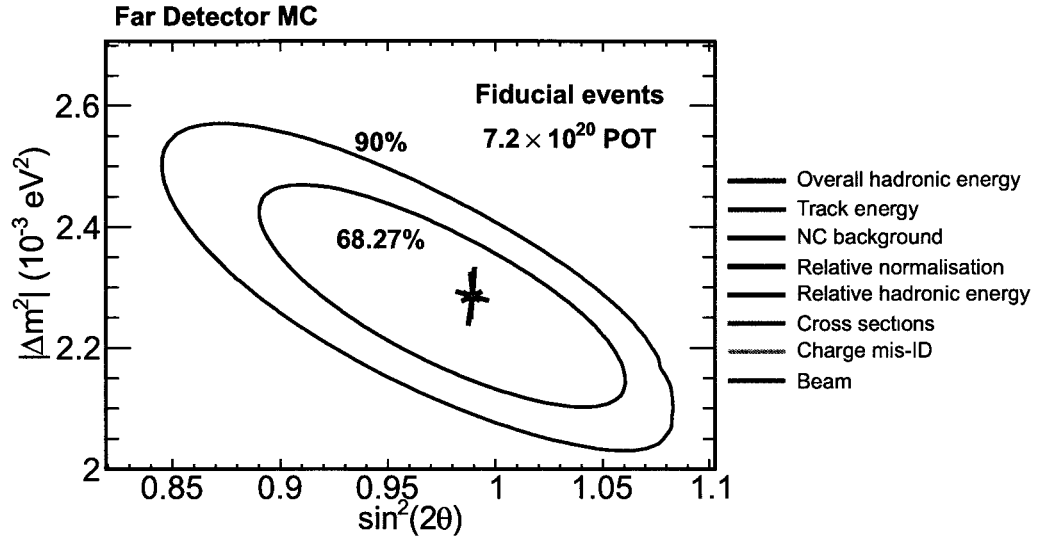
FIG. 7.11: Systematic errors shown with statistics-only sensitivity for 68% (red contour) and 90% (blue contour) C.L. for fiducial events only.

the relative populations of events defined to be Rock or Anti-Fiducial. The strips are shifted 2.5-10 mm, with a Gaussian distribution.

The four systematics that produce the largest shifts in the oscillation parameters are included as nuisance parameters in the fit. These are the same systematics included as nuisance parameters in the prior analysis [47]. The shifts are the overall Normalization, the size of the NC background, the overall shower energy uncertainty, and the track energy uncertainty. In particular, the shower energy and track energy uncertainties are difficult to include as nuisance parameters, since they redistribute events in different energy bins and shift the location of the focusing peak. The interpolation scheme of the GhostFitter algorithm is meant to compensate for this. These nuisance parameters are not utilized in generating the tables and sensitivities shown in this chapter.

Systematic	Shift	Δm^2 $(10^{-3}eV^2)$	$\sin^2(2\theta)$	$\delta(\Delta m^2)$ $(10^{-3}eV^2)$	$\delta(\sin^2(2\theta))$
Truth	-	2.295	0.985	-	-
Best Fit	-	2.28522	0.98933	-0.00978	+0.00433
μ track energy	$+1\sigma$	2.33329	0.99011	-0.04807	-0.00078
	-1σ	2.24016	0.98790	+0.04506	+0.00143
Normalization	$+1.6\%$	2.24410	0.98937	+0.04112	-0.00005
	-1.6%	2.32763	0.98878	-0.04241	+0.00055
NC Background	$+20\%$	2.29008	0.981228	-0.00486	+0.00810
	-20%	2.28105	0.998324	+0.00417	-0.00900
Absolute Shower Energy	$+1\sigma$	2.33682	0.98961	-0.05160	-0.00028
	-1σ	2.23403	0.98932	+0.05119	0.0000
ND Shower Energy	$+1.9\%$	2.28417	0.98620	+0.00105	+0.00312
	-1.9%	2.28655	0.99350	-0.00133	-0.00417
FD Shower Energy	$+1.1\%$	2.29143	0.99222	-0.00621	-0.00290
	-1.1%	2.27906	0.98756	+0.00616	+0.00176
Total CC Cross-section	$+3.5\%$	2.28456	0.99135	+0.00066	-0.00203
	-3.5%	2.28615	0.98838	-0.00093	+0.00095
M_A^{QE}	$+15\%$	2.26795	0.99314	+0.01727	-0.00382
	-15%	2.30015	0.98692	-0.01493	+0.00241
M_A^{Res}	$+15\%$	2.28620	0.99329	+0.00262	-0.00396
	-15%	2.28761	0.98623	-0.00239	+0.00310
kno Multiplicity 2	$+0.1$	2.28344	0.99173	+0.00178	-0.00240
	-0.1	2.28732	0.98806	-0.00210	+0.00126
kno Multiplicity 3	$+0.2$	2.28534	0.98993	-0.00012	-0.00061
	-0.2	2.28547	0.98979	-0.00025	-0.00047
$\bar{\nu}_\mu$Cross-section	$+4\%$	2.28423	0.99013	+0.00099	-0.00081
	-4%	2.28639	0.98963	-0.00117	-0.00041
$\bar{\nu}_\mu$QE Cross-section	$+8\%$	2.28496	0.99004	+0.00026	-0.00071
	-8%	2.28561	0.98988	-0.00039	-0.00056
$\bar{\nu}_\mu$Res Cross-section	$+8\%$	2.28379	0.99016	+0.00143	-0.00083
	-8%	2.28491	0.98989	+0.00031	-0.00057
$\bar{\nu}_\mu$kno Multiplicity 2	$+0.2$	2.28155	0.99043	+0.00367	-0.00111
	-0.2	2.28831	0.98963	-0.00309	-0.0030
Beam tuning	$+1\sigma$	2.29587	0.98930	-0.01065	-0.00030
	-1σ	2.27389	0.99053	+0.01133	-0.00121
$\bar{\nu}_\mu$Wrong-Sign	$+40\%$	2.28307	0.98820	+0.00215	+0.00112
	-40%	2.28695	0.99229	-0.00173	-0.00297
Rock Cross-section	$+1\%$	2.29495	0.99199	-0.00973	-0.00266
	-1%	2.29777	0.99177	-0.01249	-0.00245
Rock Z/A Ratio	$+0.9\%$	2.29083	0.99234	-0.00561	-0.00301
	-0.9%	2.30178	0.99144	-0.001656	-0.00212
Detector Edge	$+1\sigma$	2.28561	0.98860	-0.00039	+0.00072
	-1σ	2.28557	0.98891	-0.00035	+0.00041

TABLE 7.4: Systematic errors for the combined analysis, Runs I, II, and III, LE and pHE, PQ events and RAF data.

Shift	2008 PRL		2010 "PRL-style"		2010 Fiducial		2010 Fiducial+RAF	
	$\delta(\Delta m^2)$	$\delta(\sin^2(2\theta))$	$\delta(\Delta m^2)$	$\delta(\sin^2(2\theta))$	$\delta(\Delta m^2)$	$\delta(\sin^2(2\theta))$	$\delta(\Delta m^2)$	$\delta(\sin^2(2\theta))$
Shower Energy	0.052	0.004	0.047	0.007	0.049	0.002	0.051	< 0.001
Rel. Shower Energy	0.027	0.006	0.008	0.005	0.008	0.004	0.006	0.004
Normalization $(4\% \rightarrow 1.6\%)$	0.081	0.001	0.031	0.0	0.031	0.001	0.042	< 0.001
NC Bknd. $(50\% \rightarrow 20\%)$	0.021	0.016	0.012	0.009	0.008	0.008	0.005	0.006
μ Momentum	0.032	0.003	0.036	0.002	0.040	0.001	0.047	0.001
σ_ν (sum in quadrature)	0.006	0.004	0.009	0.005	0.007	0.004	0.020	0.007
SKZP	0.010	0.0	0.009	0.001	0.009	0.0	0.011	0.001
$\bar{\nu}_\mu$wrong-sign	-	-	-	-	0.003	0.002	0.002	0.002
RAF-only errors	-	-	-	-	-	-	0.004	0.003
Total	0.104	0.017	0.070	0.014	0.072	0.010	0.085	0.013

TABLE 7.5: Comparisons of Fiducial-only systematics for Runs I, II, and III, compared to the table published in [47]. The values from 2008 have been fit with NC background, track energy, absolute shower energy, and normalization as nuisance parameters, decreasing their size. All of the other errors quoted in the table are statistics-only. The 2010 "PRL-style" column includes Run III, but does not use resolution binning, PQ events, or RAF events. For all 2010 analyses, Relative Shower energy error is quadrature sum of shifts in both Near and Far detectors. Details about the reduction of errors to fill this table can be found in [110].
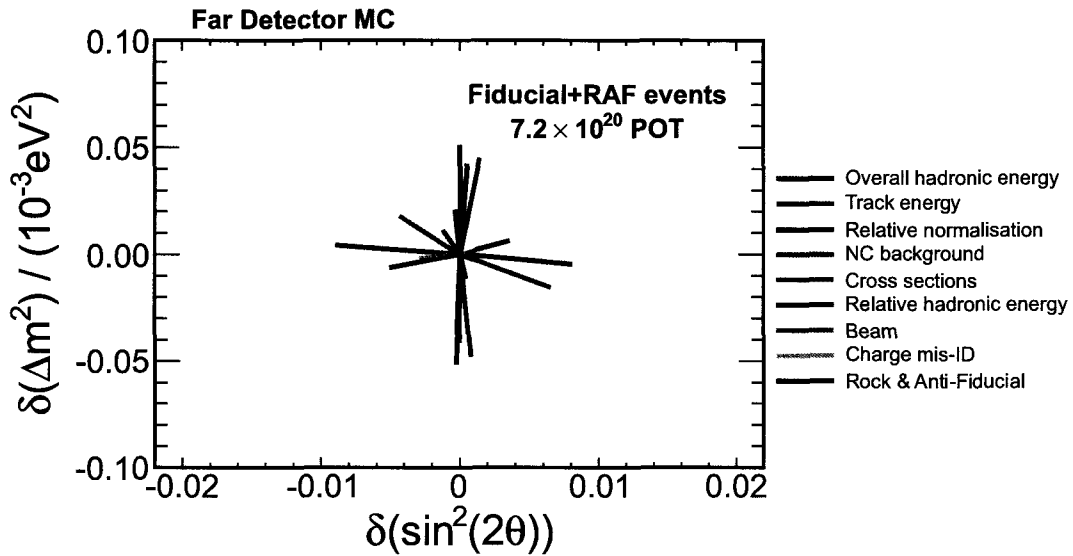
FIG. 7.12: Graphical representation of systematic errors for oscillation measurements with fiducial and RAF events. The lines indicate the shift in the oscillation best-fit point for high-statistics fake data when specific systematic shifts are applied.

7.4.1 Alternative disappearance models

Prior to opening the box, we can also evaluate how these new samples included in the analysis affect our discrimination to the alternative disappearance models discussed in Chapter 1. High-statistics Monte Carlo fake data is generated with oscillations, as above, and fit with the two models shown in Equation 1.48 and Equation 1.50. When fitting with the oscillation model, the value for χ^2 is very close to zero, so the value of χ^2 for these alternative models indicates the discriminating power available. Each alternative model contains two parameters: α and $\sin^2(2\theta)$ for Decay, and μ^2 and $\sin^2(2\theta)$ for Decoherence. Though both parameters are allowed to float when finding the best fit, in practice the value of χ^2 incurs a large penalty when straying from $\sin^2(2\theta) = 1.0$.

By comparing the values of χ^2 for these fits with the oscillation fits, we can determine how much discrimination power each sample gains. This is shown in Table 7.6.
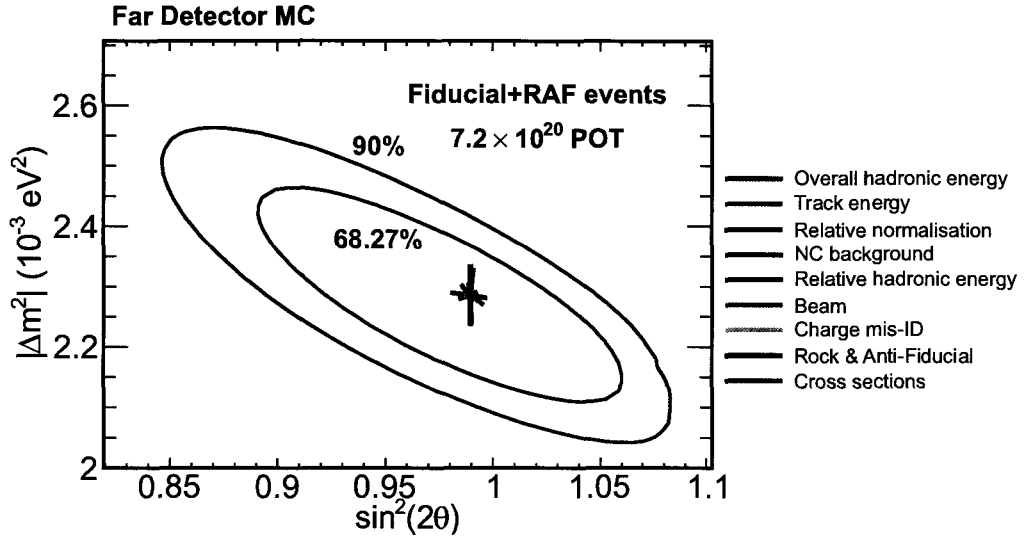
Far Detector MC

FIG. 7.13: Systematic errors shown with statistics-only sensitivity for 68% (red contour) and 90% (blue contour) C.L. for both fiducial and RAF events.

Here we express the discrimination in terms of standard deviations, with

$$n\sigma = \sqrt{\chi^2_{\text{model}} - \chi^2_{\text{oscillations}}} \qquad (7.2)$$

These systematics are intrinsic to the beam modeling, detector acceptance, and the state of our knowledge of the underlying physics at this point in time, not to the particular neutrino disappearance model that is being fit. Though this analysis is centered on measuring parameters describing neutrino oscillations, these systematic uncertainties are still present for the two alternate disappearance models we are fitting, pure decay and pure decoherence. The systematic error on fits to these two models can be seen in Figure 7.14 and Figure 7.15 for an analysis considering only events with vertices within the fiducial volume. The systematic errors on fits including fiducial and RAF events are shown in

	ν_μ-only	+ Res. Binning	+$\overline{\nu}_\mu$-like	+Sec. selection
		Decay		
α	1.125×10^{-3}	1.127×10^{-3}	1.118×10^{-3}	-
χ^2	29.783	31.690	33.671	≈ 34.2
$N\sigma$	5.457	5.629	5.803	≈ 5.9
		Decoherence		
μ^2	9.593×10^{-4}	9.614×10^{-4}	9.529×10^{-4}	-
χ^2	34.297	36.212	38.371	≈ 39.4
$N\sigma$	5.856	6.018	6.194	≈ 6.3

TABLE 7.6: Alternate disappearance model discrimination improvement with resolution binning and the inclusion of $\overline{\nu}_\mu$-like events. Calculated with high-statistics fake data oscillated at $\Delta m^2 = 2.42 \times 10^{-3} eV^2/c^4$ and $\sin^2(2\theta) = 1.0$ with a simulated 7.2×10^{20} POT. The discrimination gain reported for the secondary selection is inferred from [96].

Figures 7.16 and 7.17. The 1σ systematic error, when fitting with fiducial events and RAF events, is $\pm 0.71\sigma$ for neutrino decay and $\pm 0.69\sigma$ for neutrino decoherence. These systematic errors will be subtracted from the final model exclusion calculations.

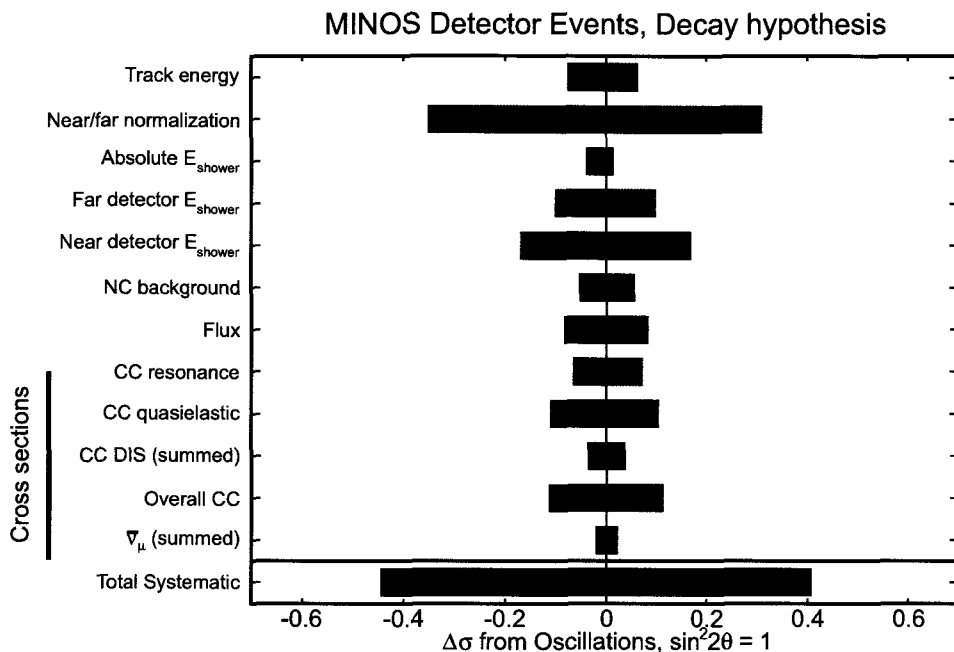## MINOS Detector Events, Decay hypothesis

FIG. 7.14: 1σ systematic uncertainties and their affect on the pure decay neutrino disappearance model when high-statistics fake data fiducial events are considered.

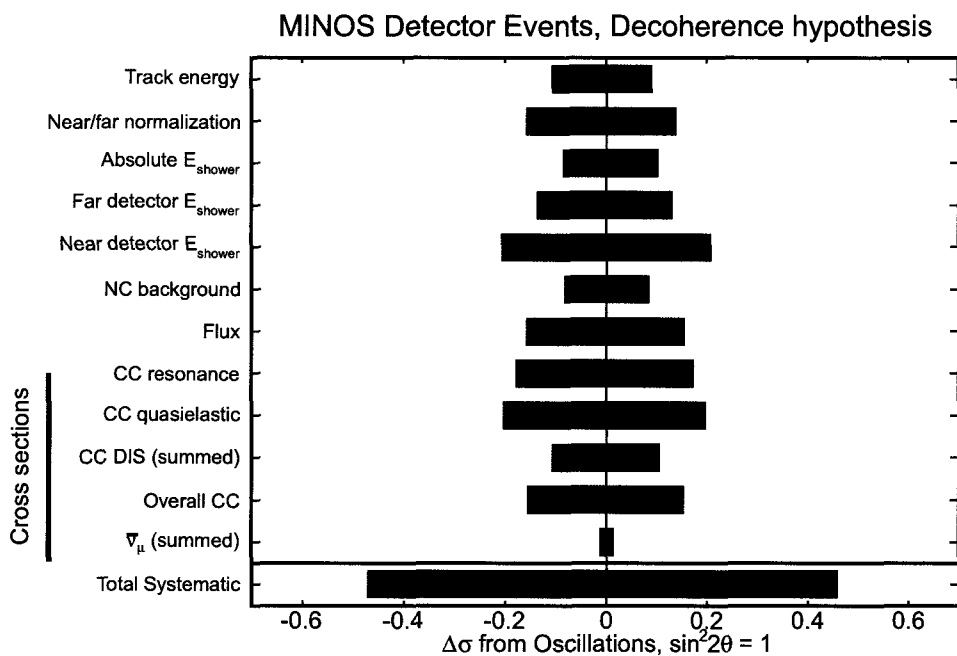## MINOS Detector Events, Decoherence hypothesis

FIG. 7.15: Systematic uncertainties and their affect on the pure decoherence neutrino disappearance model when high-statistics fake-data fiducial events are considered.
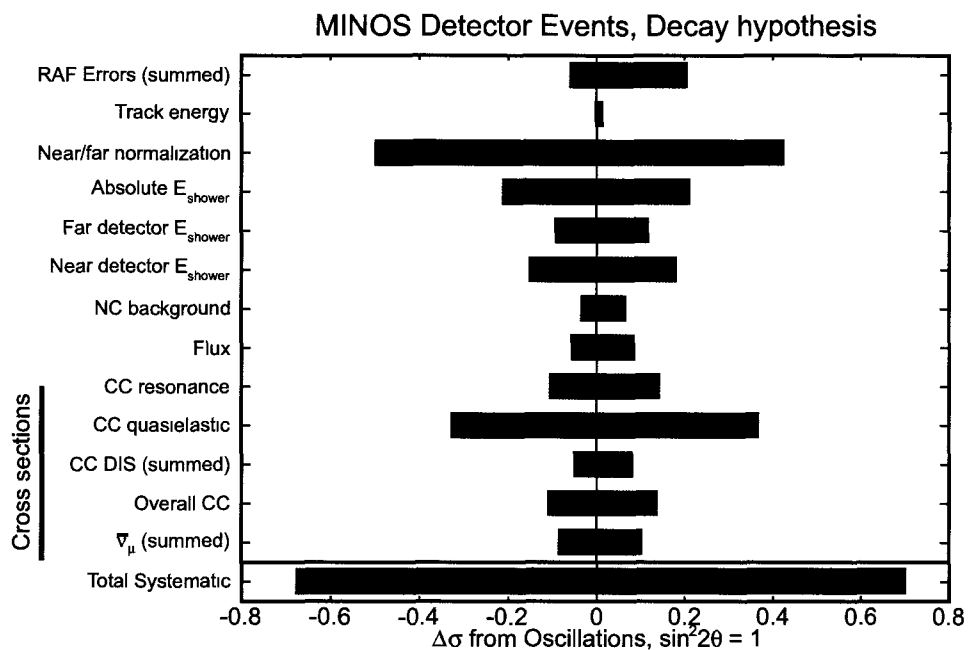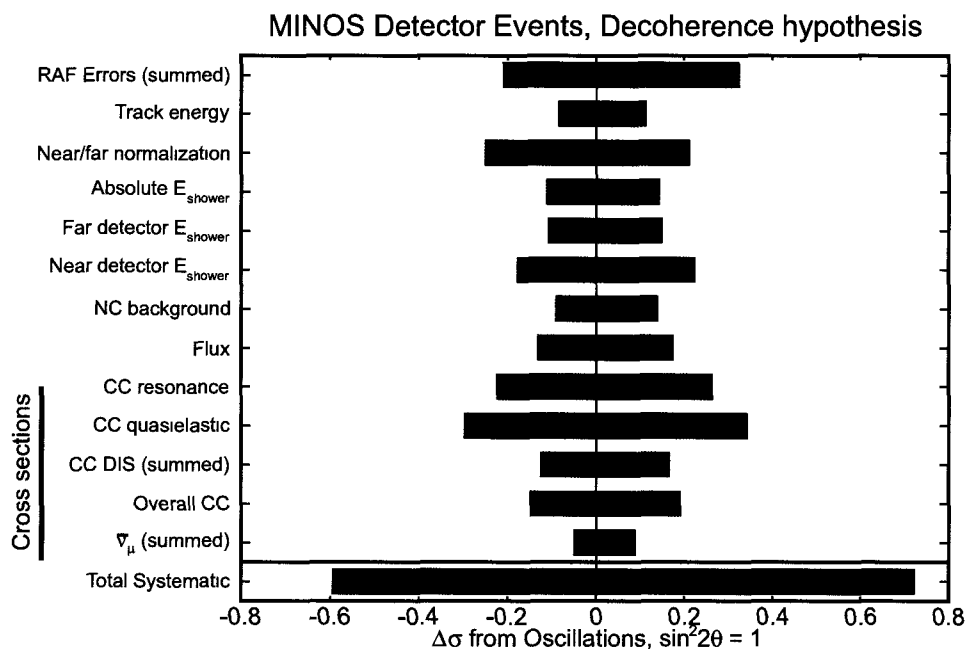
MINOS Detector Events, Decay hypothesis

FIG. 7.16: 1σ systematic uncertainties and their affect on the pure decay neutrino disappearance model when high-statistics fake data fiducial events and RAF events are considered.

MINOS Detector Events, Decoherence hypothesis

FIG. 7.17: 1σ systematic uncertainties and their affect on the pure decoherence neutrino disappearance model when high-statistics fake data fiducial events and RAF events are considered.
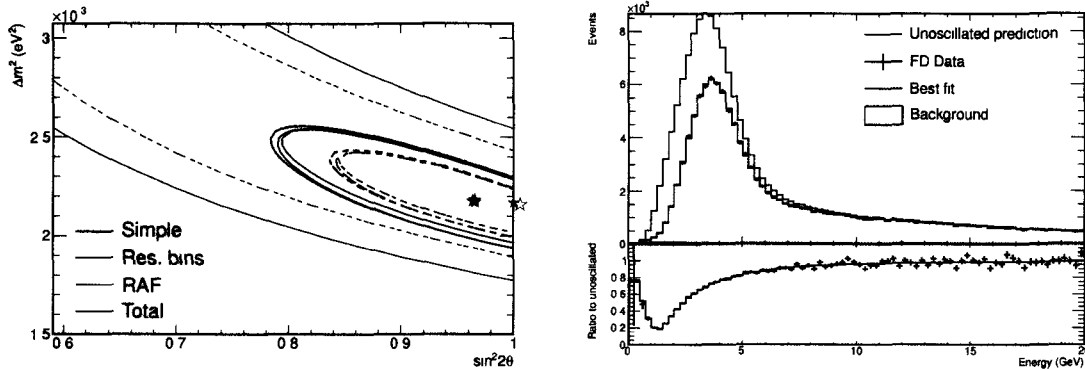
FIG. 7.18: The best-fit spectrum and 68% and 90% C.L. contours for the MDC performed prior to box opening. On the left, contours have been generated ignoring resolution bining (called "simple"), with resolution binning, with RAF events, and finally combined into a single fit [111].

7.5 Mock Data Challenge

A final test of the extrapolation and fit procedure is to examine Mock Data in a Mock Data Challenge (MDC). A large sample of fake data is independently generated with characteristics like oscillation parameters unknown to all but one person in the collaboration. Prior to fitting the mock data, it was agreed that the new extrapolation procedure is validated by fitting the MDC correctly to within 1σ of the statistical uncertainty.

The total exposure of the mock data set is equal to 100 times the exposure of the dataset being analyzed, and is separated equally into 100 datasets so that each may be fit individually. Statistical fluctuations within each energy bin will yield 100 different best fit points. This provides an opportunity to test the procedure used to calculate 68% and 90% C.L. sensitivities, by comparing the best-fit points of the experiments to the sensitivity. The best fit was found to be $\Delta m^2 = 2.18 \times 10^{-3} \text{eV}^2/\text{c}^4$ and $\sin^2(2\theta) = 0.966$. The true parameters were revealed to be $\Delta m^2 = 2.1704 \times 10^{-3} \text{eV}^2/\text{c}^4$ and $\sin^2(2\theta) = 0.9756$. This was within the 68% contour, as proscribed, and the analysis was ready to open the
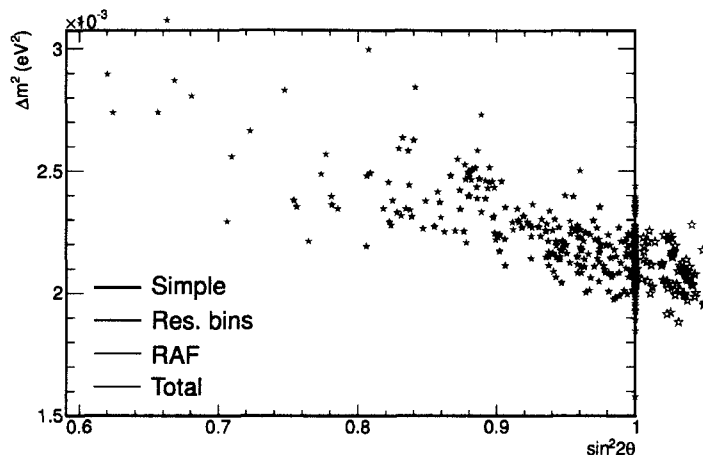
FIG. 7.19: The distribution of best fit points for 100 fake experiments in the Mock Data Challenge. The individual rock runs which would not fit with an unconstrained $\sin^2(2\theta)$ were re-fit with a $\sin^2(2\theta) \leq 1$ constraint, and pile up on the boundary. See Figure 7.18 for an explanation of the naming convention for these samples [111].

box and look at Far Detector data.

7.6 Conclusion

Many new features have been added to the Charged-Current neutrino oscillation analysis to augment the effect of the doubling of the exposure. The new shower energy estimator, resolution binning, and $\overline{\nu}_\mu$-like events all improve sensitivity in Δm^2 and $\sin^2(2\theta)$. Fitting these fiducial results simultaneously with RAF events improves our sensitivity in Δm^2 by 12%. The new selection criteria and inclusion of positive-curvature events reclaim low energy events which, along with our other analysis improvements, improve our model discrimination between oscillation hypothesis and alternative disappearance models. The level of the systematic uncertainties is such that the limiting factor on the measurements made is the size of the statistical errors. The new techniques have

been tested on mock data, and the results passed within the defined specifications of the test. Numerous comparisons of data and Monte Carlo simulated distributions were compared, to ensure data quality and that the simulations used model reality. These have been relegated to Appendix A.

CHAPTER 8

Results

At this point, having performed checks on data quality and performed numerous checks on the analysis structure, development was frozen pending the collaboration's approval. A blessing package was presented to the collaboration for review [112] detailing the analysis, the checks, and the procedure to be followed once the box had been opened. The collaboration agreed to allow this analysis to look at unblinded Far Detector data.

8.0.1 Checks against previous results

Some of the Far Detector data analyzed here had already been examined and published in 2008 [47]. The first step upon opening the box is to re-examine Runs I and II with the new methods described above. As with the mock data challenge, the fit is defined to be acceptable if it is contained within the 68% C.L. contour obtained from the same data in [47].

Some change is expected between the two fits. The new shower energy estimator and

201

the new selection criteria should change the Far Detector reconstructed energy spectra for these two runs. The new selection criteria, in particular, allows more NC background events into the data sample, filling in the oscillation dip and moving the best-fit point away from maximal mixing.

The data from Runs I and II were refit and found to be $\Delta m^2 = 2.481 \times 10^{-3} \text{eV}^2/\text{c}^4$ and $\sin^2(2\theta) = 0.914$. As an additional check, the shower energy and selection were reverted to be identical to those used in the previous analysis, and the old results were recovered. The fit lies within the 68% C.L. contours of the previous result, so the test was considered satisfactory and the analysis moved ahead with the full analysis of all three runs.

The RAF events were also checked against a previous result, which had measured the oscillation parameters using only the data from Run I [113]. Fitting Run I with the modern version of the RAF analysis yielded a best-fit result that was only 0.49 units of χ^2 away from the previous measurement [99].

8.0.2 Event selection performance

We apply each of our selection criteria in succession to the data so we can understand where all of the data is cut away. Table 8.1 shows the size of the data sets by run as they are reduced.

Note that the secondary selection algorithm does indeed increase the number of events we select at low energies. The events gained with the secondary selector are shown in Figure 8.2. The new PID selection gains a total of 62 events, most of the events
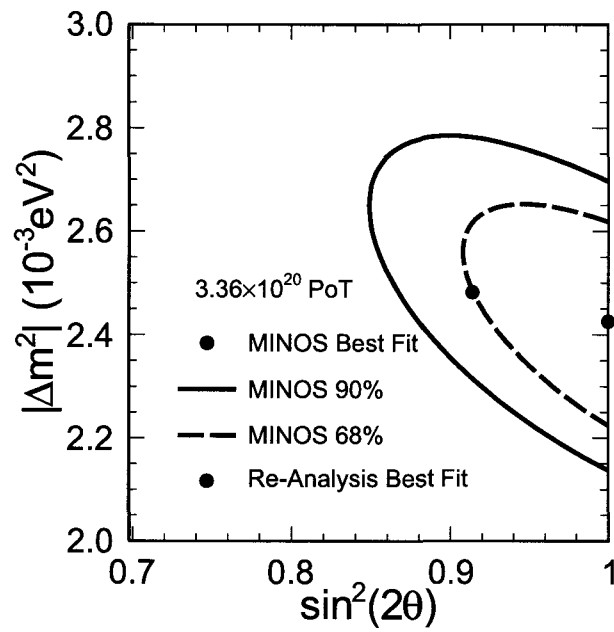
FIG. 8.1: Re-analyzing Runs I and II with the new shower energy estimator and new event selector (and no resolution binning), the best fit point moves away from maximal mixing, as expected.

gained have energies below 10 GeV, which is a region sensitive to alternative disappearance model discrimination. The distribution of the selected events amongst resolution bins is shown in Table 8.2.

	Run I LE	Run I pHE	Run II LE	Run III LE	Total
Raw Data	8846	1616	8911	13622	32995
ntrk> 0	891	331	1401	2604	2623
trkfitpass	886	330	1397	2595	5208
In Fid Vol	433	178	694	1382	2687
Data Quality	428	172	682	1354	2636
Track Angle	415	171	665	1306	2557
PID	318	129	511	1037	1986
Negative Curvature	293	120	459	902	1774
Positive Curvature	25	9*	52	135	212

TABLE 8.1: Number of events surviving preselection cuts, first tabulated in [114]. *The Run I pHE positive curvature sample is not used in the overall analysis nor included in the total, as it was deemed insignifigant before the box opening.

Resolution Bin	Run I LE	Run I pHE	Run II LE	Run III LE	Total
2008 Analysis	282	118	448	-	848
2010 Analysis	318	120	511	1037	1986
Bin 0	63	25	106	168	362
Bin 1	52	28	84	193	357
Bin 2	59	25	87	150	321
Bin 3	60	16	89	186	351
Bin 4	59	26	93	205	383
Positive Curvature	25	-	52	135	212

TABLE 8.2: Numbers of events in the Far Detector recorded in each of the three runs, and in each of the resolution bins used for the present analysis.
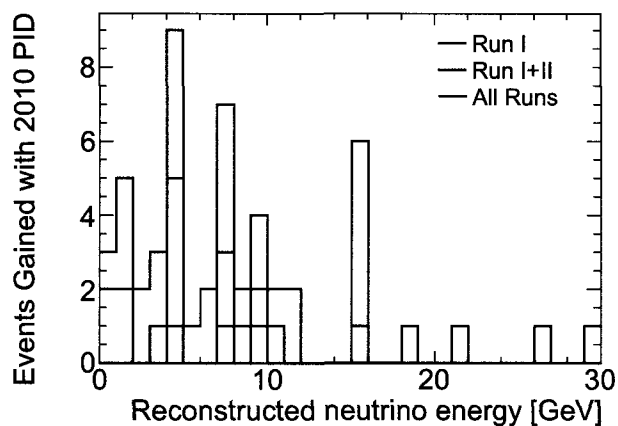
FIG. 8.2: The reconstructed energies of events gained by using the secondary selection algorithm. These events were rejected as NC-like by the primary selection algorithm [96].

8.1 Oscillation Fit

As described in Chapter 7, two separate code bases are used to fit the Far Detector data, SystFitter and GhostFitter. Only the GhostFitter code is able to fit with all four systematic errors as nuisance parameters. The ultimate result from this measurement was defined, prior to opening the box, to be a simultaneous fit between data and beam matrix extrapolated fiducial events and RAF events fit performed with the GhostFitter, with four nuisance parameters.

Both are used to fit data with statistical errors and with simple nuisance parameters to provide an extra cross-check.

8.1.1 SystFitter vs. GhostFitter

Normalization and NC background systematics are available to the SystFitter algorithm as nuisance parameters for fiducial events, so these can be fit one-at-a-time for comparison between the two fitters. The values for all of the nuisance parameters are shown in Table 8.3. The differences between the two fits are consistent within statistical errors. Best-fit values for nuisance parameters are in opposite directions because of the methods used to evaluate the shifts. The SystFitter applies systematic shifts to the Far Detector data, while the GhostFitter applies systematic shifts to the oscillated prediction.

8.1.2 Fitting fiducial events

The Far Detector data is shown in Figure 8.3 compared to the un-oscillated prediction. The ratio of the data and best-fit spectra to the un-oscillated prediction is shown in Figure 8.4. The Far Detector data is broken up by each resolution quantile, as well as the positive curvature sample, in Figure 8.5. The best-fit oscillation parameters for the fiducial sample is Δm^2=2.072 $\times$ 10^{-3} eV^2/c^4 and sin$^2(2\theta)$ consistent with maximal mixing.

8.1.3 Fitting RAF Events

Care must be taken when fitting the RAF sample on its own, due to the nature of the RAF sensitivity. The RAF sensitivity to sin$^2(2\theta)$ is very broad, and when fit alone tends to run to unphysical values where sin$^2(2\theta) \gg 1$. Applying a constraint that sin$^2(2\theta) \leq 1$,

Run	SystFitter				GhostFitter			
	$\Delta m^2 (\times 10^{-3} \text{ eV}^2/\text{c}^4)$	$\sin^2(2\theta)$	NC	Norm.	$\Delta m^2 (\times 10^{-3} \text{ eV}^2/\text{c}^4)$	$\sin^2(2\theta)$	NC	Norm.
Fiducial Only	2.316	1.002	-	-	2.330	0.995	-	-
+Norm.	2.329	1.003	-	+0.8%	2.342	0.997	-	+0.8%
+NC	2.318	1.000	-10%	-	2.330	0.995	+6%	-
RAF Only	2.072	1.000	-	-	2.090	1.000	-	-
Fiducial + RAF	2.285	1.000	-	-	2.298	0.995	-	-

TABLE 8.3: Comparison of results from SystFitter and GhostFitter. Nuisance parameter shifts carry opposite sign between the two because one shifts fake data and the other shifts predictions. The differences between the two fits are within statistical errors.
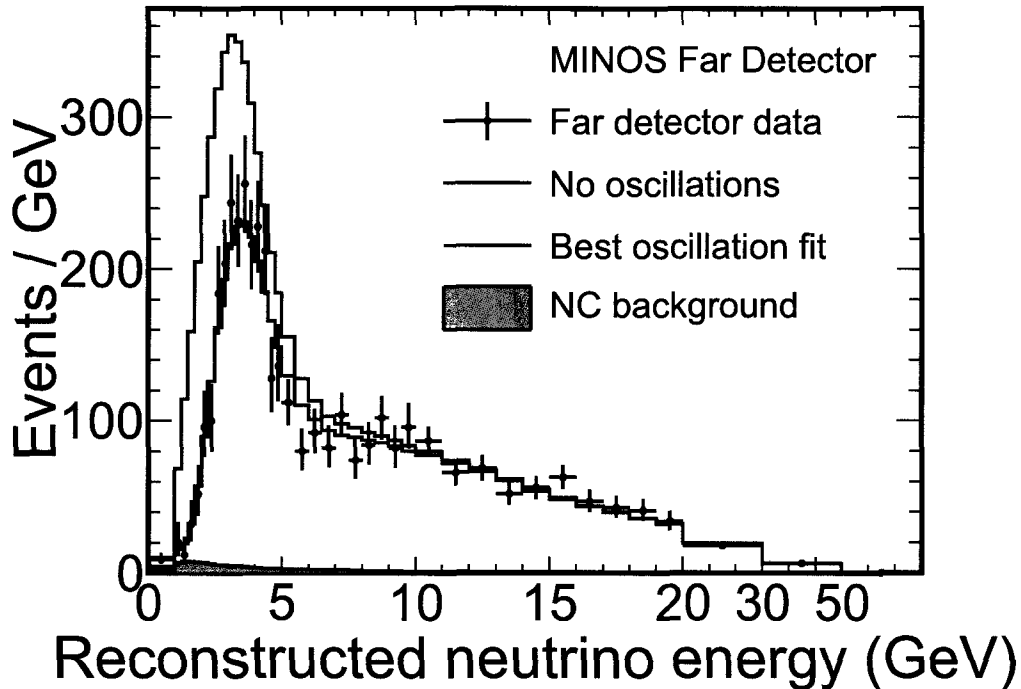
FIG. 8.3: The complete Far Detector fiducial sample energy spectrum from Runs I, II, III, and Run I pHE. The red histogram is the Far Detector prediction with no oscillations and the black crosses are the observed Far Detector data. The grey histogram shows the expected NC background. The best fit line (in blue) is shown for fiducial events only.

the best fit is at $\Delta m^2 = 2.072 \times 10^{-3}$ eV^2/c^4 and $\sin^2(2\theta) = 1.0$ with χ^2/DOF= 515.5/444.

The 90% C.L. contours for RAF events and fiducial events are compared in Figure 8.7.

As expected by observing the RAF oscillation sensitivity, the RAF sample contains little

information on the value of $\sin^2(2\theta)$, but does help to constrain the value of Δm^2.

8.1.4 Fitting Fiducial + RAF Events

When fitting the fiducial events and RAF events simultaneously, the constraint that

$\sin^2(2\theta) \leq 1$ may be lifted, as the fiducial sample closes the contour on the high $\sin^2(2\theta)$

side. A summary of the best fit values from all of the different samples, without con-
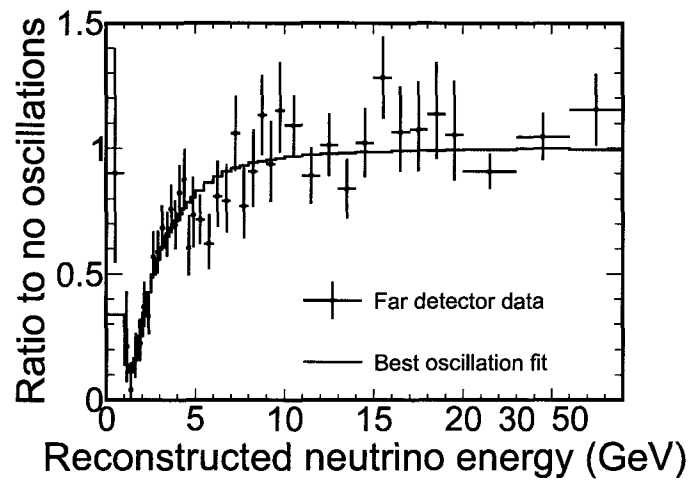
FIG. 8.4: Ratio of Far Detector data (crosses) for Runs I, II, and III, to the un-oscillated predicted spectrum. The blue line shows the best fit to the oscillation disappearance model. A result with no oscillations would be flat at 1.
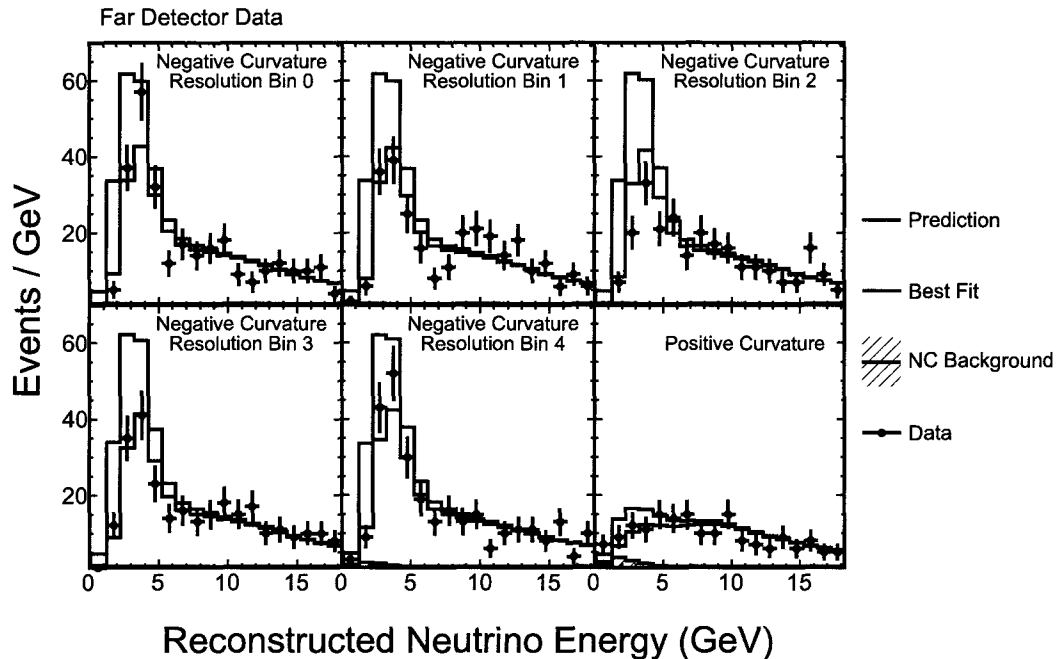
FIG. 8.5: Far Detector data spectra broken out by resolution bins, summed over all runs. The red line shows the prediction in the absence of oscillations and the black crosses show the data. The blue line is the best-fit spectrum. The small NC background is cross-hatched.
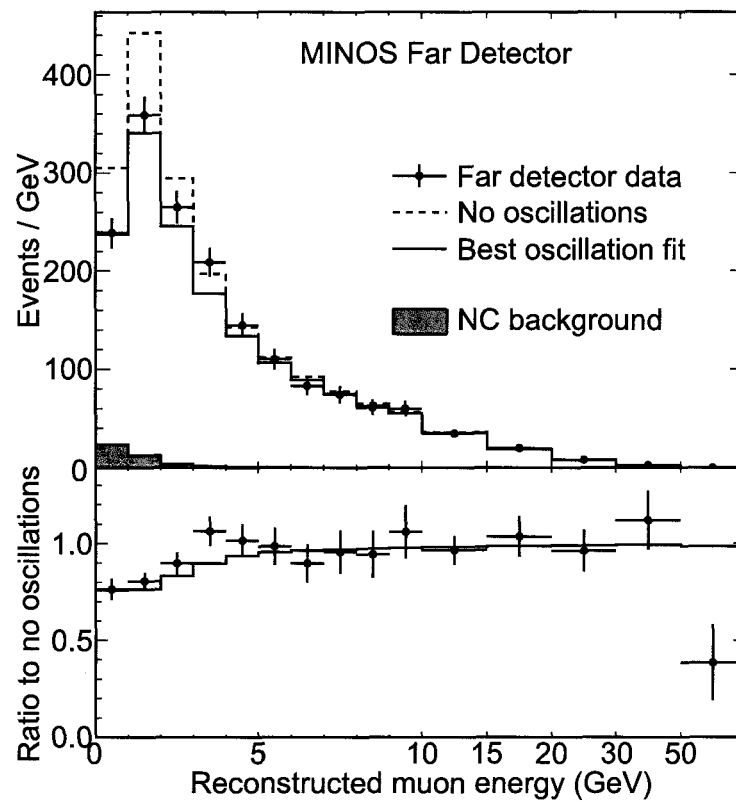
FIG. 8.6: The Far detector prediction for RAF events, summed over all geometric regions. Predictions with no oscillations are shown in the dashed line, while data are in points and the best-fit line is solid black. Below, the background-subtracted ratio of the data to the unoscillated predicted spectrum [115].
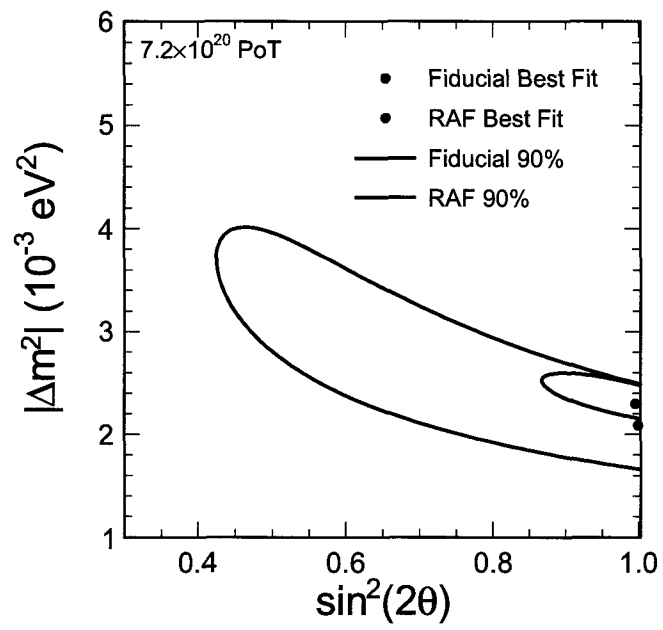
FIG. 8.7: The 90% C.L. contours for RAF events and fiducial events, separately. RAF events and fiducial events are fit simultaneously to produce the final result for this analysis.

Run	Δm^2 $(10^{-3}\,\mathrm{eV}^2)$	$\sin^2(2\theta)$	χ^2/DOF
Runs I+II LE, 2008-style	2.43	1.0	90/97
Runs I+II LE	2.452	0.9232	219/198
+ pHE	2.448	0.9232	311.0/298
+ Run III	2.280	1.022	409.5/398
+ PQ	2.297	1.006	701.1/698
+ ResBins	2.317	1.002	2119.5/2298
+ RAF	2.285	1.000	2636.5/2742

TABLE 8.4: Statistics-only best fit values using the SystFitter algorithm. The 2008-style analysis refers to [47], where calorimetric shower energy and only a single selection algorithm are used.

sideration of any systematic uncertainties, is shown in Table 8.4. Table 8.5 shows the best-fit oscillation parameters fit with single systematic nuisance parameters, as well as the full fit with all four nuisance parameters. None of the nuisance parameters is pulled significantly away from their nominal in the full fit. With the full fit including RAF events and all four nuisance parameters, the best-fit oscillation parameters are Δm^2=2.314 × 10^{-3} eV^2/c^4 and $\sin^2(2\theta) = 1.001$ with χ^2/DOF=2633.3/2742. At 90% C.L., we can say that $\sin^2(2\theta) > 0.90$ at 90% C.L. We see from Figure 8.8 that the statistical error on the Δm^2 measurement is $\delta(\Delta m^2) =^{+0.12}_{-0.08}$, and the systematic error from Table 7.5 is $\delta(\Delta m^2) =^{+0.09}_{-0.09}$.
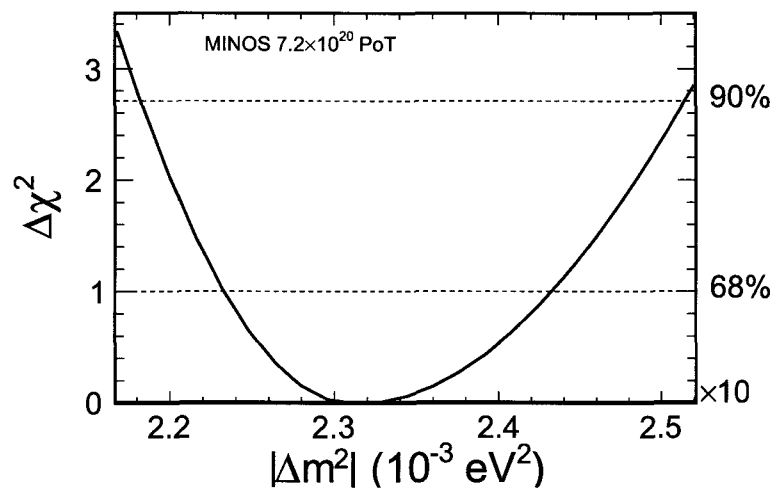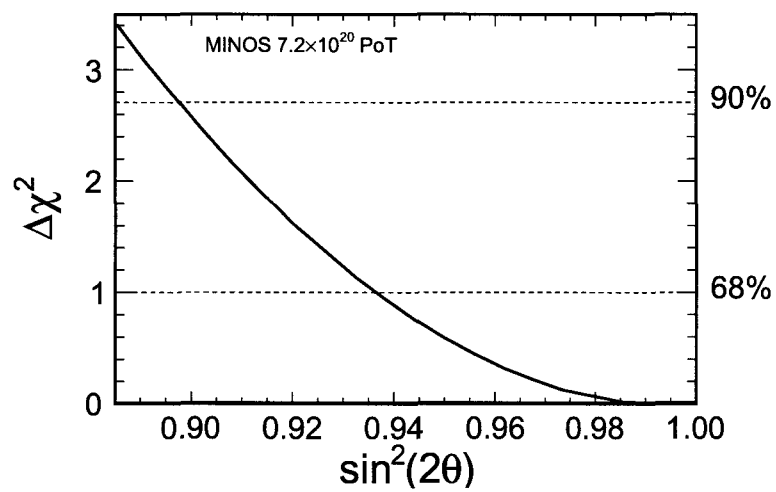
FIG. 8.8: One-dimensional projections of the oscillation parameter fit results to Runs I, II, and III Far Detector data.

	$\Delta m^2 (\times 10^{-3} \text{ eV}^2/\text{c}^4)$	$\sin^2(2\theta)$	NC	Norm.	E_{shw}	E_{trk}
Fiducial Only	2.330	0.995	-	-	-	-
	2.330	0.997	+6%	-	-	-
	2.342	0.997	-	+0.8%	-	-
	2.344	1.00	-	-	-0.28σ	-
	2.331	0.995	-	-	-	-0.02σ
	2.346	1.00	+5%	+0.8%	-0.29σ	$+0.13\sigma$
Fiducial + RAF	2.298	0.995	-	-	-	-
	2.298	0.997	+6%	-	-	-
	2.310	0.995	-	+0.7%	-	-
	2.310	0.995	-	-	-0.29σ	-
	2.302	0.993	-	-	-	-0.07σ
	2.314	1.001	+5%	+0.6%	-0.3σ	$+0.07\sigma$

TABLE 8.5: Values for nuisance parameters with the GhostFitter.

8.1.5   Statistical likelihood - or "How likely is our data?"

The complexity of this analysis has grown substantially, and it is worth while to look at the statistical likelihood of our goodness-of-fit to be sure there is nothing pathological affecting our result. Our χ^2 calculation incorporates a large number of bins, many of which may have < 1 events predicted or zero measured in them. An accurate assessment of the statistical likelihood of the fit is the location of our χ^2 value relative to the distribution of a large number of statistically fluctuating simulated data sets.

A high-statistics Far Detector fake data set is produced with oscillation parameters equal to those that have been measured with the MINOS data. The exposure of the fake data set is scaled down, so the data spectrum has the same integral as the data spectrum, but has little to no statistical fluctuations. The bin weights in the Far Detector spectra are then randomly fluctuated about a Poisson distribution and are fit with `SystFitter` to produce a pair of measured oscillation parameters that has been smeared by statistical fluctuations, but with true oscillation parameters that we know. This process is repeated ten thousand times, and each time the χ^2 value of the fit is stored. The distribution of these values of χ^2 for fits with fiducial events only are shown in Figure 8.9, along with a marker marking the location of the χ^2 value for the fit to data. The $\chi^2/\text{DOF} = 2119.5/2298$ measured in data is a better fit than 66% of the random fits. This indicates that the measurement with fiducial events has the benefit of favorable statistical fluctuations. When RAF's are included, the pendulum swings the other way - the $\chi^2/\text{DOF} = 2633.3/2742$ is better than only 41% of random fits. This means that the statistical fluctuations were relatively unfavorable within the RAF sample, and the measurement could have been better given

the oscillation sensitivity.

8.2 Fitting Alternative Models

In addition to oscillations, the data was fit against two alternative models, neutrino decay and quantum decoherence, which are described in Chapter 1. Each model is used to warp the no-oscillation Far Detector prediction, which is then compared to the real Far Detector data with the log-likelihood equation (Equation 6.19). The same four systematics are included for the alternative models as the oscillation fit; Near/Far normalization, the size of the NC background, and the calibration uncertainties for track and shower energy estimates. The effect of each of the systematic uncertainties on the alternative models is shown in Chapter 7, particularly Figure 7.16 and Figure 7.17.

For the case of pure neutrino decay (with no neutrino oscillations), there are two parameters, $\sin^2(2\theta)$ and α, where α is the decay constant in Equation 1.48. For the case of pure neutrino decoherence, again with no oscillations, there are two parameters, $\sin^2(2\theta)$ and μ^2, where μ^2 acts as an effective mass introduced by the extra interaction potential required to induce decoherence on the distance scale of the MINOS baseline (Equation 1.50).

The fit to the decay hypothesis has best-fit decay parameters $\alpha = 2.22 \times 10^{-3}$ with a $\chi^2/\text{DOF} = 2165.8/2298$ with fiducial events only, and also $\alpha = 2.22 \times 10^{-3}$ with $\chi^2/\text{DOF} = 2696.1/2742$ with fiducial and RAF events. The fiducial fit is 6.8σ from the oscillation hypothesis, while the fiducial and RAF fit to neutrino decay is 7.9σ from oscillations. We must subtract the 0.71σ systematic error calculated in Chapter 7 from this,
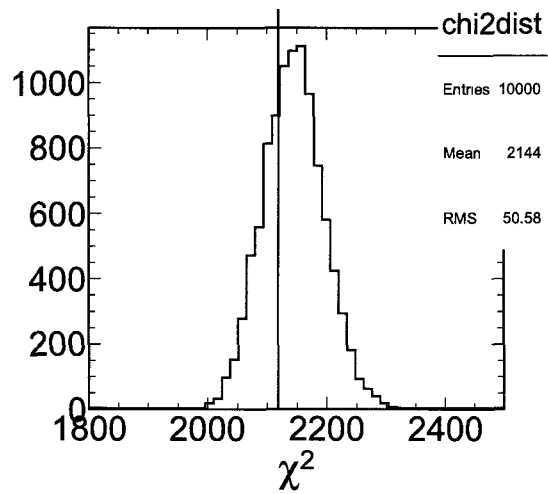
(a) Fiducial fits only
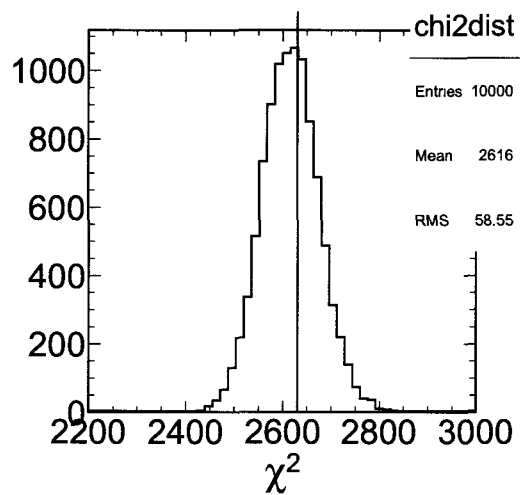
(b) Fiducial+RAF fits

FIG. 8.9: The distribution of χ^2 values for 10,000 statistically fluctuated fits in the Far Detector. The blue line indicates the location of the χ^2 value for the relevant fit to data events.
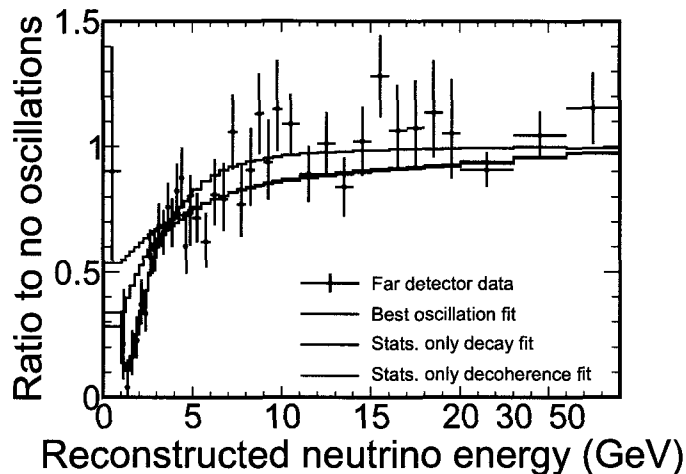
FIG. 8.10: Ratio of Far Detector data to the un-oscillated predicted spectrum, compared to the best fit spectra for three disappearance models.

so the final result is that we disfavor neutrino decay by 6.1σ for the fiducial events and by 7.2σ when fitting fiducial and RAF events simultaneously.

The fit to the decoherence hypothesis has best fit parameters $\mu^2 = -2.07 \times 10^{-3}$ with $\chi^2/\mathrm{DOF}=$ 2197.6/2298 with fiducial events only, and also $\mu^2 = -2.07 \times 10^{-3}$ with $\chi^2/\mathrm{DOF}=$ 2727.1/2742 with fiducial and RAF events. the fiducial fit is 8.8σ from oscillations, while the fiducial and RAF fit to decoherence is 9.7σ from oscillations. In both of these cases, $\sin^2(2\theta)$ acts to constrain the Near/Far normalization, which is fixed by the high energy tail. This can plainly be seen in Figure 8.10. We must subtract the 0.69σ systematic error calculated in Chapter 7 from this, so the final result is that we disfavor the decoherence hypothesis by 8.1σ with fiducial events and 9.0σ when fitting fiducial and RAF events simultaneously.

CHAPTER 9

Conclusion

The data collected in Runs I, II, and III have been analyzed and found to support the neutrino oscillation hypothesis. We measure $\Delta m^2 = (2.32^{+0.12}_{-0.08}(\text{stat.})^{+0.09}_{-0.09}(\text{syst.})) \times 10^{-3} \text{eV}^2/\text{c}^4$ and $\sin^2(2\theta) > 0.90$ at 90% C.L. The addition of Run III to Runs I and II, pure neutrino decay model disfavoring increases from 3.7σ to 6.1σ when considering only events with vertices in the fiducial region. Similarly, the addition of Run III increases the exclusion of pure quantum decoherence from 5.7σ to 8.2σ when fitting only fiducial events. By including rock and anti-fiducial events, pure neutrino decay model disfavoring increases to 7.2σ, and the pure quantum decoherence model disfavoring increases to 9.0σ.

This is a milestone achievement for the MINOS experiment. The two remaining viable models that could explain neutrino disappearance have been disfavored at greater than 6σ. MINOS has also made the world's most precise measurement of Δm^2_{atm}, beating a measurement that MINOS made in 2008 in what is likely to be the last neutrino mass splitting measurement the experiment will make. The comparison of the MINOS allowed
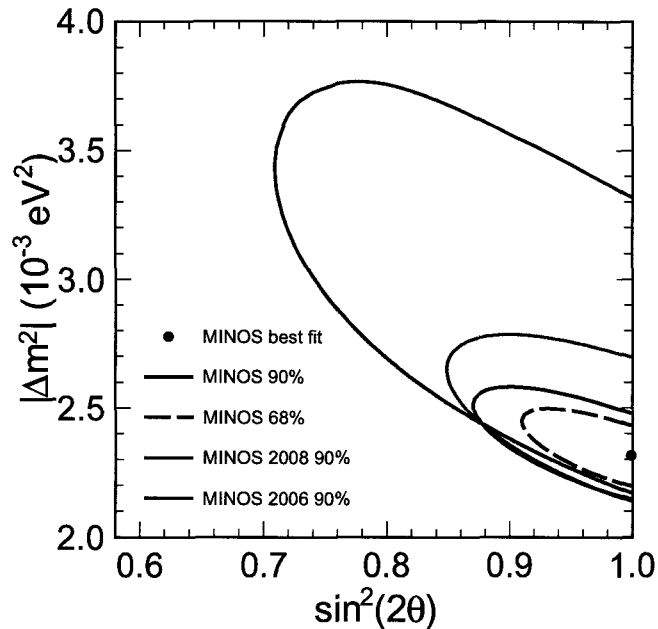
FIG. 9.1: The 90% C.L. contours for each of the results published by MINOS. MINOS 2006 and MINOS 2008 in blue and black refer to oscillation results measured in [46] and [47], respectively.

region to measurements made by other experiments is shown in Figure 9.2. Looking at the progression of MINOS results in Figure 9.1, there is also a cautionary tale here, that observing high Δm^2 values in preliminary results is an effect of statistics, not physics. Recent measurements of the antineutrino oscillation parameters $\Delta \bar{m}^2$ and $\sin^2(2\bar{\theta})$ in MINOS have also shown that exciting preliminary results do not last [116].

9.1 Future Experiments

With these results from MINOS, two of the three mixing angles in the PMNS matrix have been measured to high precision. The next generation of long-baseline experiments
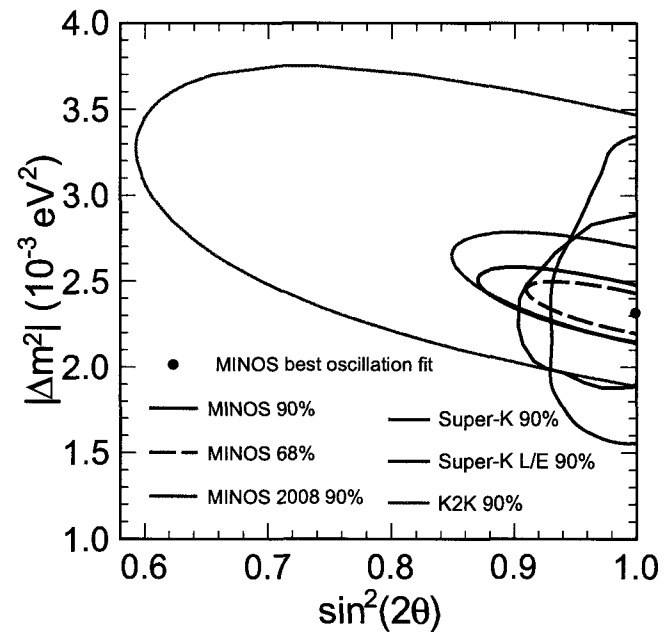
FIG. 9.2: The allowed region of oscillation parameters as measured by several different experiments, compared to the most recent measurement made by MINOS. SuperK contours are from the the most recently published results [44].

are designed to measure θ_{13}. NOνA is a long-baseline experiment located in northern Minnesota in a position off-axis from the NuMI beam at Fermilab. T2K is another off-axis long-baseline experiment, utilizing the SuperK detector at Kamiokande and a beam at J-PARC in Tokai, Japan. Both of these experiments use narrow-band ν_μ beams to look for ν_e appearance, but will also measure $\nu_\mu \rightarrow \nu_\tau$ oscillations as well. These experiments will contribute precise measurements of Δm_{32}^2, as well as θ_{13}. If they choose to run for a long period of time in $\bar{\nu}_\mu$ mode and measure $\bar{\nu}_e$ appearance, they may begin to set limits on the value of δ_{CP} as well. These experiments, along with many other double β-decay experiments and reactor experiments, will make accurate measurements with neutrinos to try and improve the Standard Model, as well as further human knowledge about the physical world. They may even turn up some surprises.

APPENDIX A

Data/MC Validation

The MINOS experiment relies heavily on Monte Carlo simulations to understand what is happening within the two detectors. This appendix displays the key Data/MC validation plots for both the Near and Far Detectors. The comparison occured prior to examining the oscillation signal in the Far Detector data, and the differences were deemed to be unlikely to affect the final result.
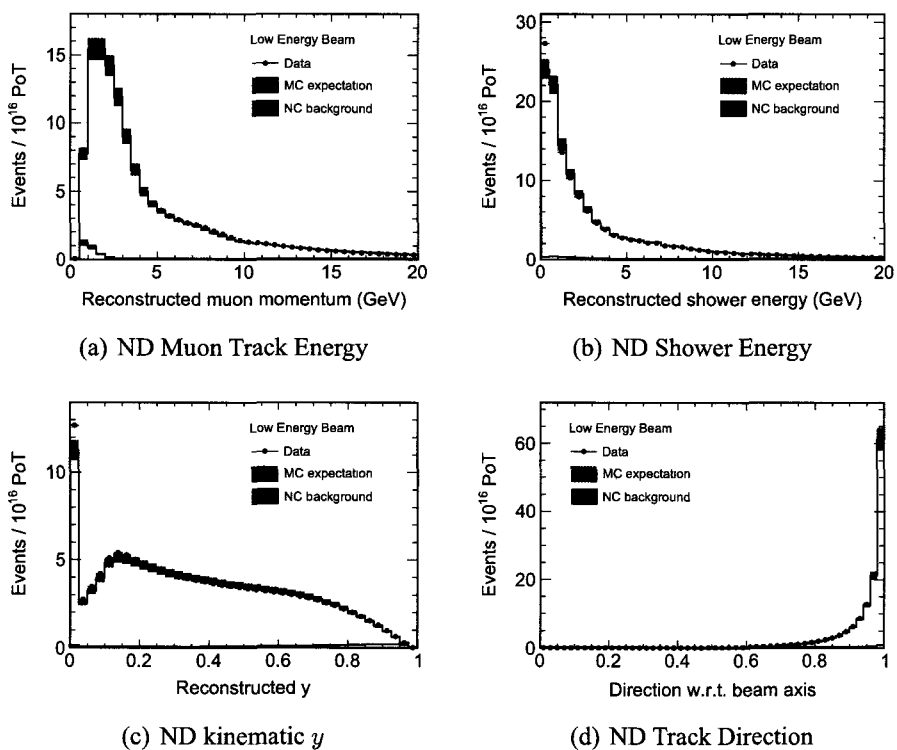
A.1 Near Detector

(a) ND Muon Track Energy

(b) ND Shower Energy

(c) ND kinematic y

(d) ND Track Direction

FIG. A.1: Reconstructed kinematic distributions for events in the Near Detector with Data (points) and MC (red line). Expected NC background is also shown (blue line). Shaded bands indicate MC statistical error bars [117].

(a) Number of active planes

(b) Mean energy deposited per strip

(c) Signal fluctuation parameter

(d) Transverse profile parameter
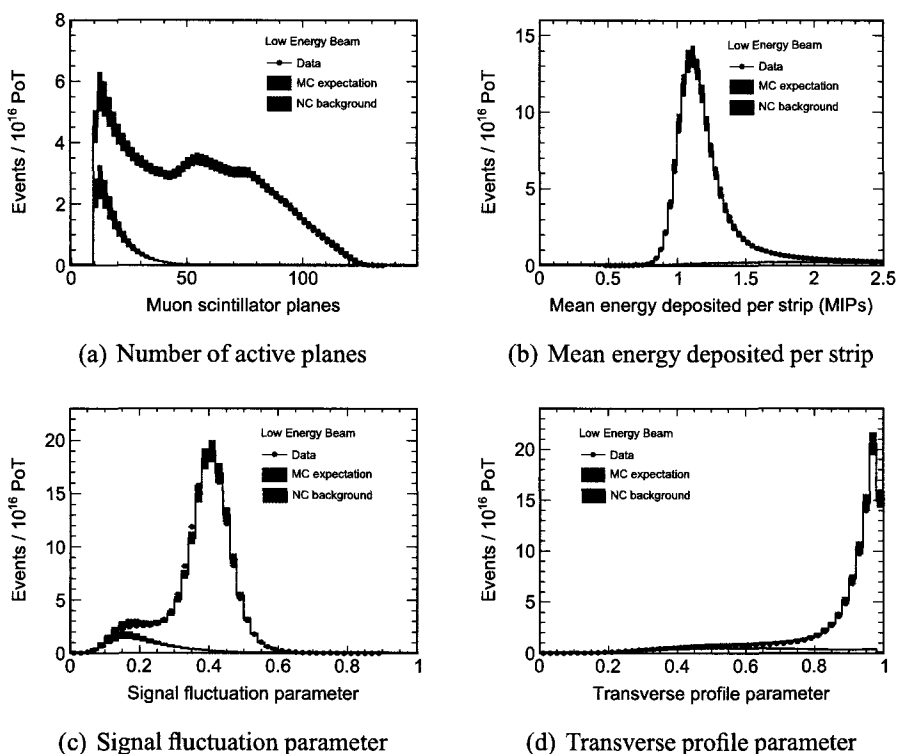
FIG. A.2: Data/MC agreement between variables used in the primary kNN selection variable algorithm. Near Detector data are shown with MC expectations, along with the MC statistical error band [96].

(a) Track end pulse height

(b) Scattering variable
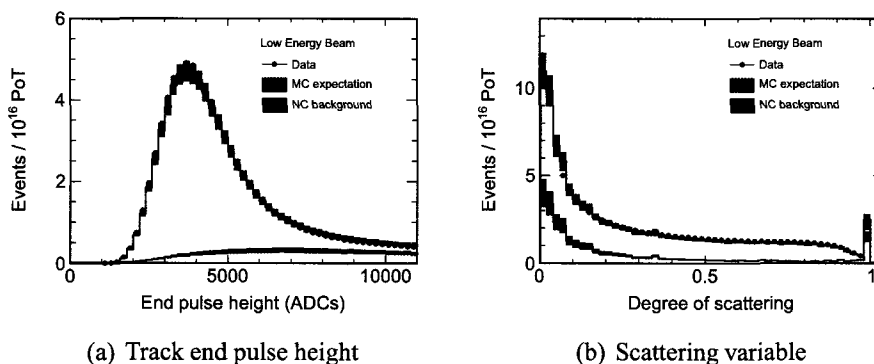
FIG. A.3: Data/MC agreement with variables used in the secondary kNN selection variable algorithm. Near Detector data are shown in points with MC expectation in red and NC background expectation in blue, along with the MC statistical error band [96].
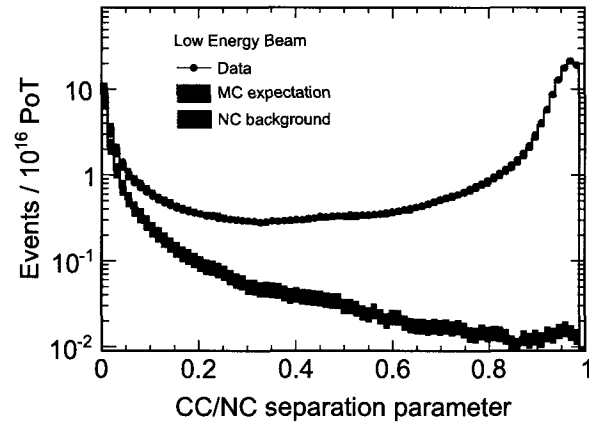
(a) Primary kNN
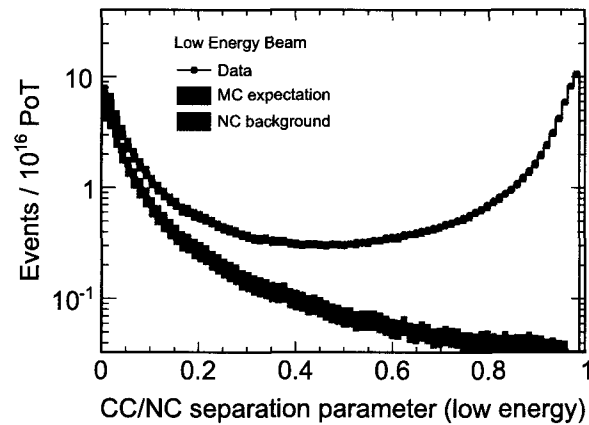
(b) Secondary kNN

FIG. A.4: Near Detector data/MC agreement for the two selection algorithms, the primary kNN and the secondary kNN. Near Detector data is shown with black points, while MC expectation is shown in red and NC expectation is shown in blue, with MC statistical error bars shaded in. A CC/NC separation parameter value of 0 indicates an event is maximally NC-like, and a value of 1 indicates an event is maximally CC-like. An event is accepted in the oscillation analysis if it has a primary selector value > 0.26 or a secondary selector value > 0.51. Note both plots are on log scales [117].
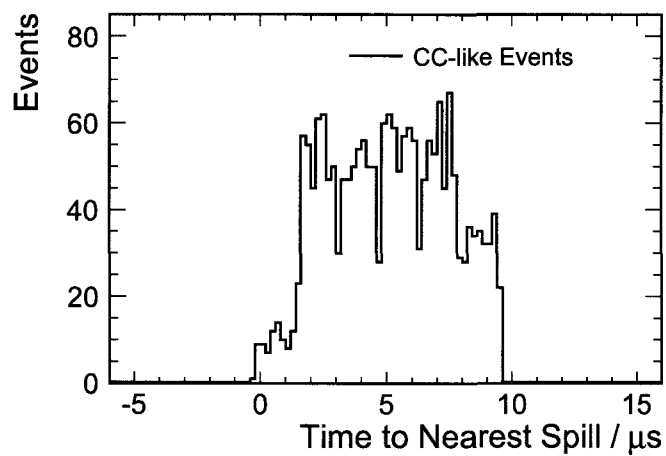
FIG. A.5: Far Detector event timings for Runs I, II, and III. The events fall in six "buckets," just as they are delivered to the NuMI beam from the Main Injector. All six buckets fit within the 12 μs spill timing window [118].

A.2 Far Detector

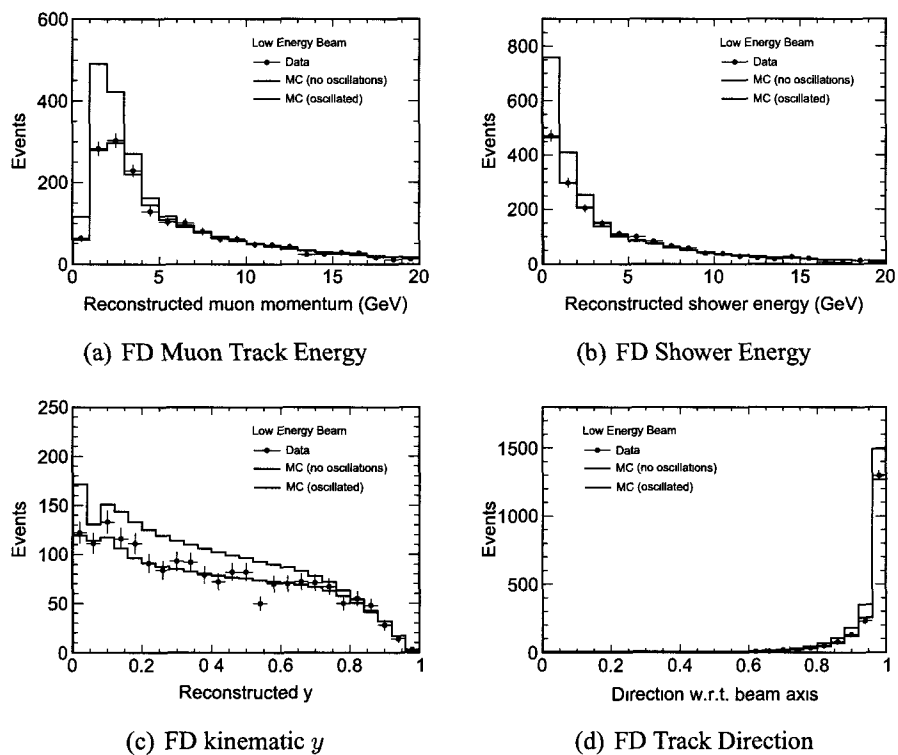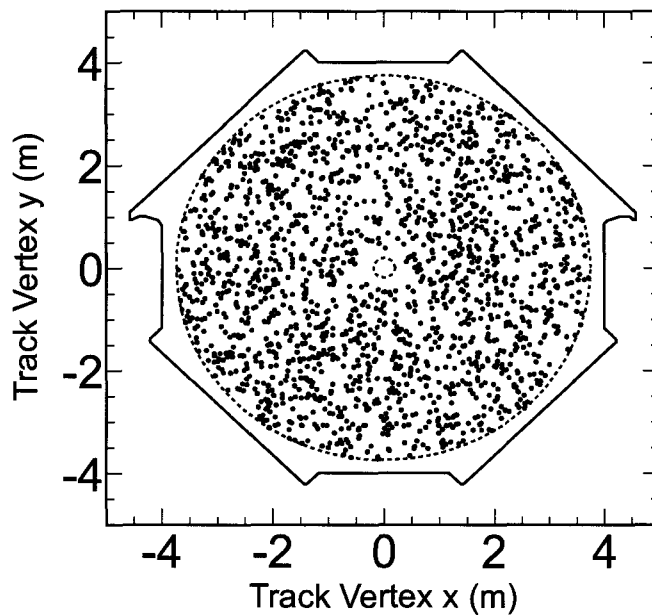(a) FD Muon Track Energy

(b) FD Shower Energy
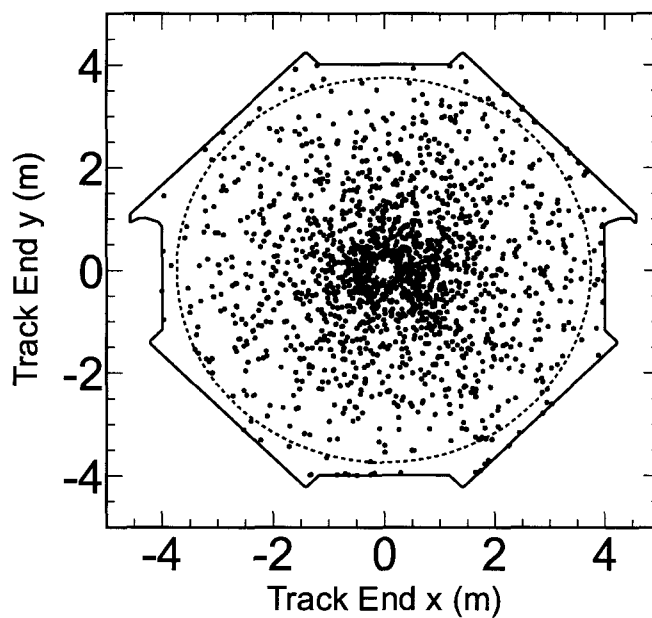
(c) FD kinematic y

(d) FD Track Direction

FIG. A.6: Reconstructed kinematic distributions for events in the Far Detector with Data (points) and MC (grey line). MC expectation with oscillations are shown in red. Shaded bands indicate MC statistical error bars [118].
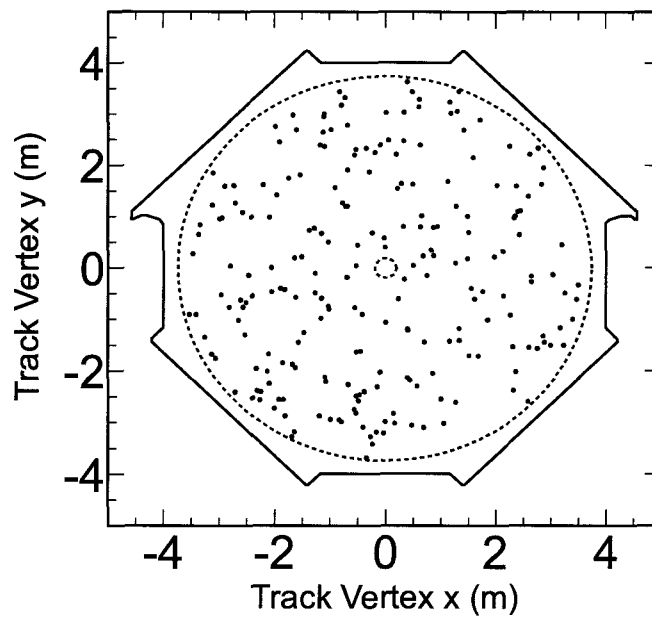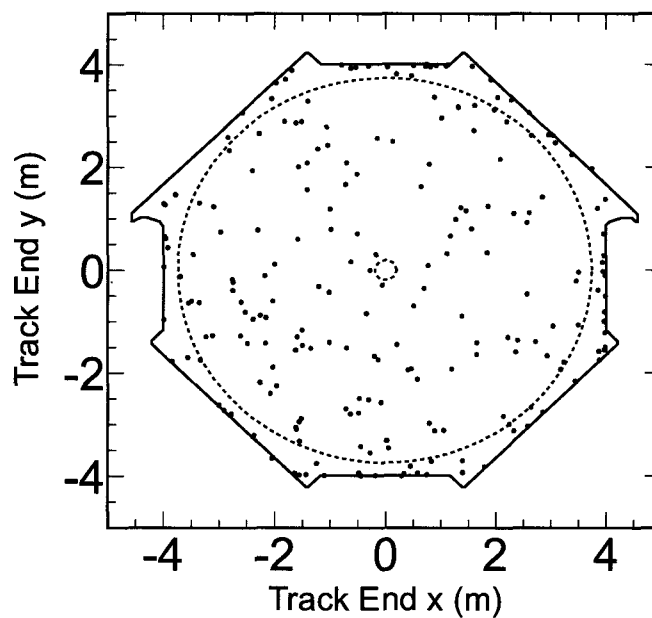
(a) Track Vertices

(b) Track Endpoints

FIG. A.7: The distribution of selected event vertices and endpoints for events in the Far Detector x and y for all runs. All events shown here had tracks reconstructed with negative curvature. The uniform distribution of track vertices indicates that there are no readout problems present, and the clustering of event endpoints around the magnetizing coil hole represents the focusing effect of the magnetic field.

(a) Track Vertices

(b) Track Endpoints

FIG. A.8: The distribution of selected event vertices and endpoints for tracks with positive curvature in the Far Detector x and y for all runs. The uniform distribution of track vertices indicates that there are no readout problems present, and the defocusing effects on μ^+ are clearly visible along the outside edge of the detector.

(a) FD muon track vertex r

(b) FD muon track vertex x

(c) FD muon track vertex y
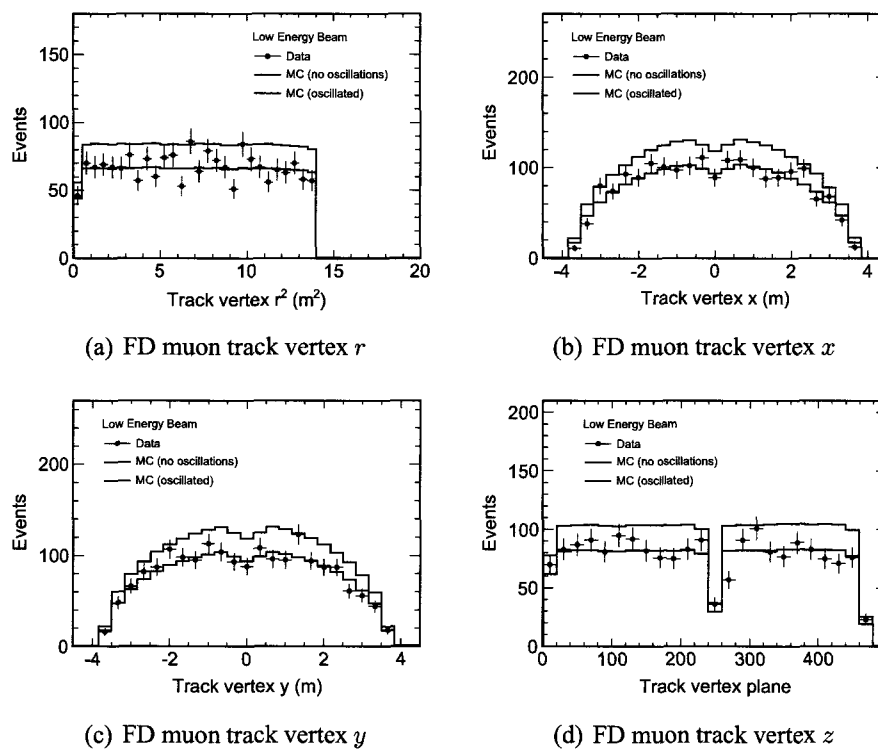
(d) FD muon track vertex z

FIG. A.9: Distributions of endpoints of reconstructed tracks for events in the Far Detector with Data (points) and MC (grey line). MC expectation with oscillations are shown in red. The variable r is defined as $r \equiv \sqrt{x^2 + y^2}$ [118].

(a) FD muon track end r

(b) FD muon track end x

(c) FD muon track end y

(d) FD muon track end z

FIG. A.10: Distributions of endpoints of reconstructed tracks for events in the Far Detector with Data (points) and MC (grey line). MC expectation with oscillations are shown in red. The variable r is defined as $r \equiv \sqrt{x^2 + y^2}$.
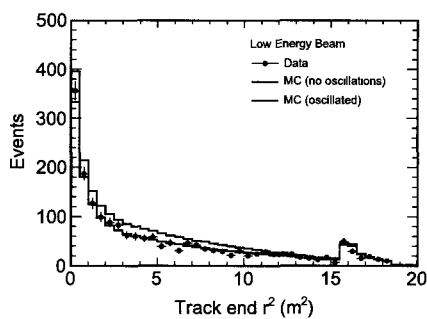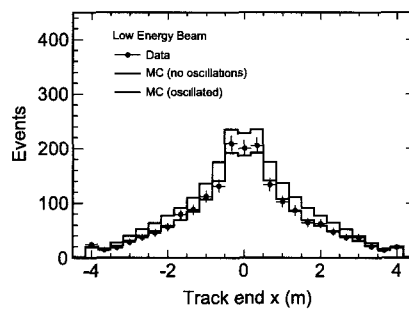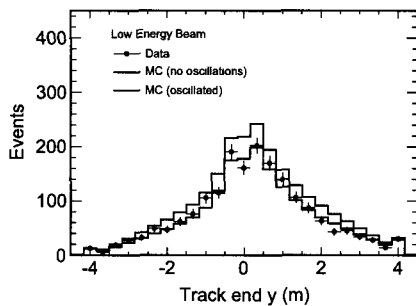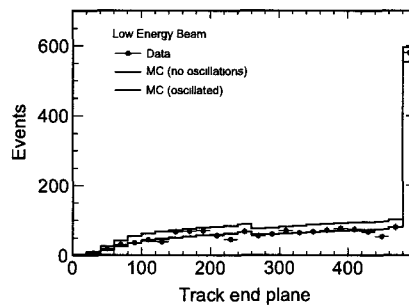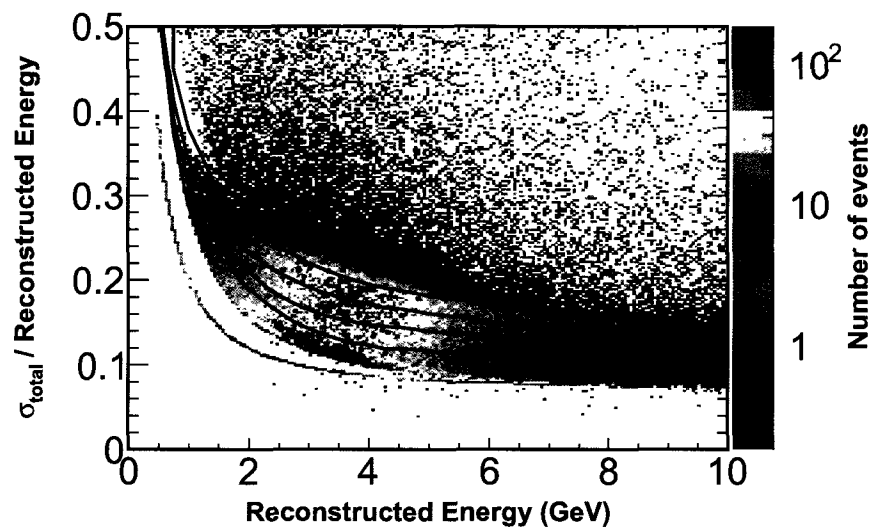
(a) Run 1
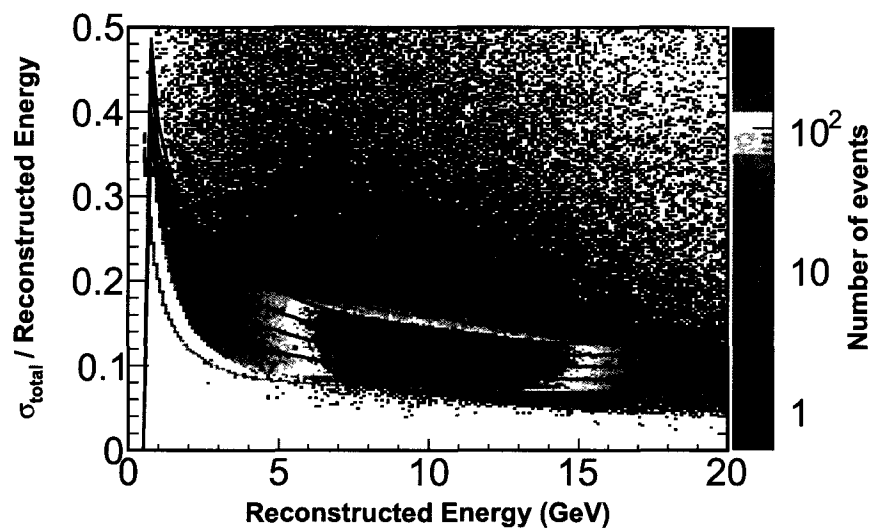
(b) Run 1 pHE

FIG. A.11: Estimated resolution of data events in Run I, both LE and pHE beam configurations, compared to the resolution cuts determined from MC. Monte Carlo population is in the background on a log color scale. Black dots are data points.
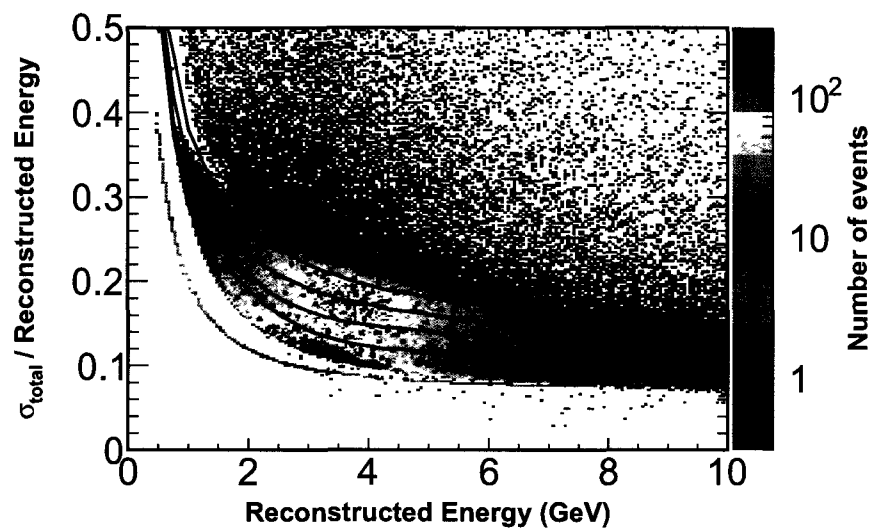
(a) Run 2
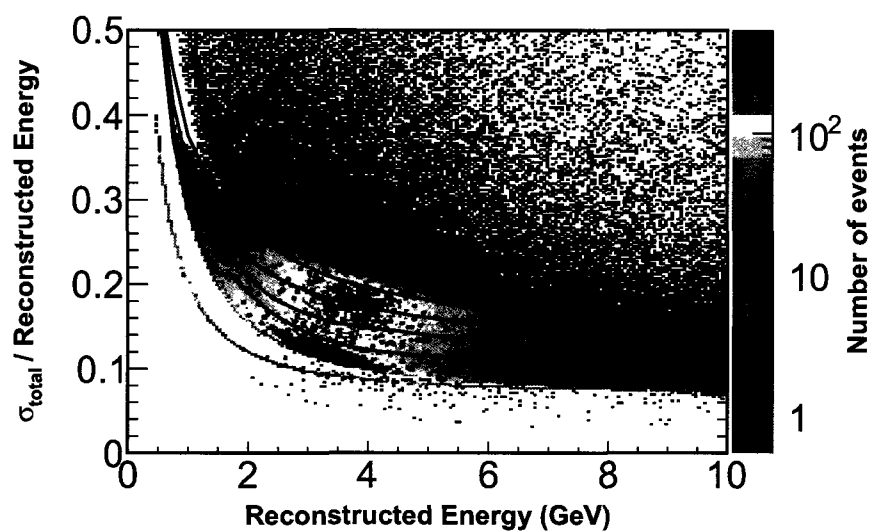
(b) Run 3

FIG. A.12: Resolution of data events in runs II and III, compared to the cuts determined from MC. Monte Carlo population is in the background on a log color scale. Black dots are data points.

BIBLIOGRAPHY

[1] Wolfgang Pauli. Letter to Tubingen meeting. December 1930.

[2] E. Fermi. Versuch einer theorie der β-strahlen (towards the theory of β-rays). *Z. Phys.*, 88:161–177, 1934.

[3] Fred L. Wilson. Fermi's theory of beta decay. *American Journal of Physics*, 36(12):1150–1160, 1968.

[4] J. Chadwick. Possible existence of a neutron. *Nature*, 129:312, 1932.

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

[6] F. Reines and C. L. Cowan. Detection of the free neutrino. *Phys. Rev.*, 92(3):830–831, Nov 1953.

[7] G. Danby, J-M. Gaillard, K. Goulianos, L. M. Lederman, N. Mistry, M. Schwartz, and J. Steinberger. Observation of high-energy neutrino reactions and the existence of two kinds of neutrinos. *Phys. Rev. Lett.*, 9(1):36–44, Jul 1962.

[8] T. D. Lee and C. N. Yang. Question of parity conservation in weak interactions. *Phys. Rev.*, 104(1):254–258, Oct 1956.

[9] C. S. Wu, E. Ambler, R. W. Hayward, D. D. Hoppes, and R. P. Hudson. Experimental test of parity conservation in beta decay. *Phys. Rev.*, 105(4):1413–1415, Feb 1957.

[10] Sheldon L. Glashow. Partial-symmetries of weak interactions. *Nuclear Physics*, 22(4):579 – 588, 1961.

[11] A. Salam and J.C. Ward. Electromagnetic and weak interactions. *Physics Letters*, 13(2):168 – 171, 1964.

[12] Steven Weinberg. A model of leptons. *Phys. Rev. Lett.*, 19(21):1264–1266, Nov 1967.

[13] R. P. Feynman and M. Gell-Mann. Theory of the Fermi interaction. *Phys. Rev.*, 109(1):193–198, Jan 1958.

[14] F.J. Hasert et al. Search for elastic muon-neutrino electron scattering. *Physics Letters B*, 46(1):121 – 124, 1973.

[15] M. L. Perl et al. Evidence for anomalous lepton production in $e^+ - e^-$ annihilation. *Phys. Rev. Lett.*, 35(22):1489–1492, Dec 1975.

[16] Precision electroweak measurements on the Z resonance. *Physics Reports*, 427(5-6):257 – 454, 2006.

[17] K. Kodama et al. Observation of tau neutrino interactions. *Physics Letters B*, 504(3):218 – 224, 2001.

[18] H. Gallagher. The NEUGEN neutrino event generator. *Nuclear Physics B - Proceedings Supplements*, 112(1-3):188 – 194, 2002.

[19] P. Adamson et al. A Study of Muon Neutrino Disappearance Using the Fermilab Main Injector Neutrino Beam. *Phys. Rev.*, D77:072002, 2008.

[20] K. Nakamura et al. (Particle Data Group). *J. Phys G*, 37(07521), 2010.

[21] P. Adamson et al. Neutrino and antineutrino inclusive Charged-Current cross section measurements with the MINOS Near Detector. *Phys. Rev. D*, 81(7):072002, Apr 2010.

[22] R. Davis. Solar neutrinos. II: Experimental. *Phys. Rev. Lett.*, 12:303–305, 1964.

[23] B. T. Cleveland et al. Measurement of the solar electron neutrino flux with the Homestake chlorine detector. *Astrophys. J.*, 496:505–526, 1998.

[24] M. C. Gonzalez-Garcia and Michele Maltoni. Phenomenology with massive neutrinos. *Phys. Rept.*, 460:1–129, 2008.

[25] K. S. Hirata et al. Experimental study of the atmospheric neutrino flux. *Physics Letters B*, 205(2-3):416 – 420, 1988.

[26] Ch. Berger et al. Study of atmospheric neutrino interactions with the Fréjus detector. *Physics Letters B*, 227(3-4):489 – 494, 1989.

[27] M. Aglietta et al. Experimental study of atmospheric neutrino flux in the NUSEX experiment. *Europhys. Lett*, 8:611, 1989.

[28] R. Becker-Szendy et al. Electron- and muon-neutrino content of the atmospheric flux. *Phys. Rev. D*, 46(9):3720–3724, Nov 1992.

[29] S. Ahlen et al. Atmospheric neutrino flux measurement using upgoing muons. *Physics Letters B*, 357(3):481 – 486, 1995.

[30] W. W. M. Allison et al. Measurement of the atmospheric neutrino flavour composition in Soudan 2. *Physics Letters B*, 391(3-4):491 – 500, 1997.

[31] B. Pontecorvo. Inverse beta processes and nonconservation of lepton charge. *Sov. Phys. JETP*, 7:172–173, 1958.

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

[33] Hiroshi Nunokawa, Stephen Parke, and Renata Zukanovich Funchal. Another possible way to determine the neutrino mass hierarchy. *Phys. Rev. D*, 72(1):013009, Jul 2005.

[34] L. Wolfenstein. Neutrino oscillations in matter. *Phys. Rev. D*, 17(9):2369–2374, May 1978.

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

[36] C. Arpesella et al. Direct measurement of the ^{7}Be solar neutrino flux with 192 days of Borexino data. *Phys. Rev. Lett.*, 101(9):091302, Aug 2008.

[37] W. Hampel et al. GALLEX solar neutrino observations: results for GALLEX IV. *Physics Letters B*, 447(1-2):127 – 133, 1999.

[38] V.N. Gavrin et al. Solar neutrino results from SAGE. *Nuclear Physics B - Proceedings Supplements*, 77(1-3):20 – 25, 1999.

[39] Q. R. Ahmad et al. Measurement of day and night neutrino energy spectra at SNO and constraints on neutrino mixing parameters. *Phys. Rev. Lett.*, 89(1):011302, Jun 2002.

[40] B. Aharmim et al. Electron energy spectra, fluxes, and day-night asymmetries of ^{8}b solar neutrinos from measurements with nacl dissolved in the heavy-water detector at the sudbury neutrino observatory. *Phys. Rev. C*, 72(5):055502, Nov 2005.

[41] S. Fukuda et al. Determination of solar neutrino oscillation parameters using 1496 days of Super-Kamiokande-I data. *Physics Letters B*, 539(3-4):179 – 187, 2002.

[42] S. Abe et al. Precision measurement of neutrino oscillation parameters with kamland. *Phys. Rev. Lett.*, 100(22):221803, Jun 2008.

[43] B. Aharmim et al. Electron energy spectra, fluxes, and day-night asymmetries of ^{8}B solar neutrinos from measurements with NaCl dissolved in the heavy-water detector at the Sudbury Neutrino Observatory. *Phys. Rev. C*, 72(5):055502, Nov 2005.

[44] Y. Ashie et al. Evidence for an oscillatory signature in atmospheric neutrino oscillations. *Phys. Rev. Lett.*, 93(10):101801, Sep 2004.

[45] M. H. Ahn et al. Measurement of neutrino oscillation by the K2K experiment. *Phys. Rev. D*, 74(7):072003, Oct 2006.

[46] D. G. Michael et al. Observation of muon neutrino disappearance with the MINOS detectors and the NuMI neutrino beam. *Phys. Rev. Lett.*, 97:191801, 2006.

[47] P. Adamson et al. Measurement of Neutrino Oscillations with the MINOS Detectors in the NuMI Beam. *Phys. Rev. Lett.*, 101:131802, 2008.

[48] N. Agafonova et al. Observation of a first ν_τ candidate in the OPERA experiment in the cngs beam, 2010.

[49] M. Apollonio et al. Search for neutrino oscillations on a long base-line at the CHOOZ nuclear power station. *The European Physical Journal C - Particles and Fields*, 27:331–374, 2003. 10.1140/epjc/s2002-01127-9.

[50] C. Athanassopoulos et al. Evidence for neutrino oscillations from muon decay at rest. *Phys. Rev. C*, 54(5):2685–2708, Nov 1996.

[51] C. Athanassopoulos et al. Results on $\nu_\mu \to \nu_e$ neutrino oscillations from the LSND experiment. *Phys. Rev. Lett.*, 81(9):1774–1777, Aug 1998.

[52] B. Armbruster et al. Upper limits for neutrino oscillations $\bar{\nu}_\mu \to \bar{\nu}_e$ from muon decay at rest. *Phys. Rev. D*, 65(11):112001, Jun 2002.

[53] P. Astier et al. Search for $\nu_\mu \to \nu_e$ oscillations in the NOMAD experiment. *Physics Letters B*, 570(1-2):19 – 31, 2003.

[54] B. Achkar et al. Search for neutrino oscillations at 15, 40 and 95 meters from a nuclear power reactor at Bugey. *Nuclear Physics B*, 434(3):503 – 532, 1995.

[55] A. A. Aguilar-Arevalo et al. Search for electron neutrino appearance at the $\Delta m^2 \sim 1$ eV^2 scale. *Phys. Rev. Lett.*, 98(23):231801, Jun 2007.

[56] A. A. Aguilar-Arevalo et al. Event excess in the MiniBooNE search for $\bar{\nu}_\mu \to \bar{\nu}_e$ oscillations. *Phys. Rev. Lett.*, 105(18):181801, Oct 2010.

[57] P. Adamson et al. New constraints on muon-neutrino to electron-neutrino transitions in MINOS. *Phys. Rev. D*, 82(5):051102, Sep 2010.

[58] C. Kraus et al. Final results from phase II of the Mainz neutrino mass search in tritium β decay. *European Physical Journal C*, 40(4):447–468, 2005.

[59] Ch. Weinheimer. KATRIN, a next generation tritium β-decay experiment in search for the absolute neutrino mass scale. *Progress in Particle and Nuclear Physics*, 48(1):141 – 150, 2002.

[60] D. N. Spergel et al. Wilkinson Microwave Anisotropy Probe (WMAP) three year results: Implications for cosmology. *Astrophys. J. Suppl.*, 170:377, 2007.

[61] H.V. Klapdor-Kleingrothaus et al. Latest results from the HEIDELBERG-MOSCOW double beta decay experiment. *European Physical Journal A*, 12(2):147–154, 2001.

[62] V. Barger, J. G. Learned, S. Pakvasa, and T. J. Weiler. Neutrino decay as an explanation of atmospheric neutrino observations. *Phys. Rev. Lett.*, 82(13):2640–2643, Mar 1999.

[63] V. Barger, J. G. Learned, P. Lipari, M. Lusignoli, S. Pakvasa, and T. J. Weiler. Neutrino decay and atmospheric neutrinos. *Physics Letters B*, 462(1-2):109 – 114, 1999.

[64] E. Lisi, A. Marrone, and D. Montanino. Probing possible decoherence effects in atmospheric neutrino oscillations. *Phys. Rev. Lett.*, 85(6):1166–1169, Aug 2000.

[65] G. L. Fogli, E. Lisi, A. Marrone, and D. Montanino. Status of atmospheric neutrino $\nu_\mu \to \nu_\tau$ oscillations and decoherence after the first K2K spectral data. *Phys. Rev. D*, 67(9):093006, May 2003.

[66] D. G. Michael et al. The Magnetized steel and scintillator calorimeters of the MINOS experiment. *Nucl. Instrum. Meth.*, A596:190–228, 2008.

[67] K Anderson et al. NuMI facility technical design report. *Tech. Rep. FERMILAB-DESIGN-1998-01*, 1998.

[68] Takaaki Kajita. Atmospheric neutrino results from Super-Kamiokande and Kamiokande: Evidence for nu_mu oscillations. *Nucl. Phys. Proc. Suppl.*, 77:123–132, 1999.

[69] M. Dorman. Beam fit position paper. *MINOS-DOCDB-7146*, 2010.

[70] A. Himmel. The FLUGG beam simulation. *MINOS-DOCDB-5803*, 2009.

[71] P. Adamson et al. Measurement of neutrino velocity with the MINOS detectors and NuMI neutrino beam. *Phys. Rev.*, D76:072005, 2007.

[72] A. Inc. Ansys commercial software. Southpointe 275 Technology Drive,Canonsburg, PA 15317.

[73] S. Coleman. Magnetics update. *MINOS-DOCDB-3680*, 2007.

[74] R. Ospanov. *A measurement of muon neutrino disappearance with the MINOS detectors and the NuMI beam*. PhD thesis, University of Texas, Austin, 2008.

[75] D. Damiani. Momentum analysis of cosmogenic muons in the MINOS detectors, 2006. William & Mary B.S. Honors Thesis.

[76] Nicholas Metropolis and S. Ulam. The Monte Carlo method. *Journal of the American Statistical Association*, 44(247):pp. 335–341, 1949.

[77] Rene Brun and Fons Rademakers. ROOT – An object oriented data analysis framework. *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment*, 389(1-2):81–86, 1997. New Computing Techniques in Physics Research V.

[78] A. Ferrari, P.R. Sala, A. Fasso, and J. Ranft. FLUKA: A multi-particle tranport code (program version 2005). *CERN-2005-010*, 2005.

[79] S. Agostinelli et al. G4–a simulation toolkit. *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment*, 506(3):250 – 303, 2003.

[80] A. Himmel. The NuMI beam simulation with FLUGG. *MINOS-DOCDB-6316*, 2009.

[81] A. Bodek and U. K. Yang. Unified approach for modelling neutrino and electron nucleon scattering cross sections from very high Q^2 to $Q^2 = 0$. *AIP Conf. Proc.*, 721:358–362, 2004.

[82] Dieter Rein and Lalit M. Sehgal. Neutrino-excitation of baryon resonances and single pion production. *Annals of Physics*, 133(1):79 – 153, 1981.

[83] Steven Dytman. Final state interactions in neutrino-nucleus experiments. *Acta Phys. Polon.*, B40:2445–2460, 2009.

[84] M. Dorman. *Cross Section Measurements for Quasi-Elastic Neutrino-Nucleus Scattering with the MINOS Near Detector*. PhD thesis, University College London, 2010.

[85] C. Alt et al. Inclusive production of charged pions in p+C collisions at 158 GeV/c beam momentum. *Eur. Phys. J. C*, 49(4):897–917, 2007.

[86] Rajendran Raja. The Main Injector particle production experiment (MIPP) at Fermilab. *Journal of Physics: Conference Series*, 9(1):303, 2005.

[87] Bonesini, M., Marchionni, A., Pietropaolo, F., and T. Tabarelli de Fatis. On particle production for high energy neutrino beams. *The European Physical Journal C - Particles and Fields*, 20:13–27, 2001. 10.1007/s100520100656.

[88] F. James and M. Roos. Minuit - a system for function minimization and analysis of the parameter errors and correlations. *Computer Physics Communications*, 10(6):343 – 367, 1975.

[89] D. Groom, N. Mokhov, and S. Striganov. Muon stopping power and range tables 10 MeV-100 TeV. *Atomic Data and Nuclear Data Tables*, 78(2):183 – 356, 2001.

[90] J. Marshall. *A study of muon neutrino disappearance with the MINOS detectors and the NuMI neutrino beam*. PhD thesis, University College London, 2008.

[91] R. Frühwirth. Application of Kalman filtering to track and vertex fitting. *Nuc. Inst. and Meth. A*, 262(2-3):444 – 450, 1987.

[92] Michael Alan Kordosky. *Hadronic interactions in the MINOS detectors*. PhD thesis, University of Texas, Austin, 2004. FERMILAB-THESIS-2004-34.

[93] C. Zeitnitz and T.A. Gabriel. The GEANT-CALOR interface and benchmark calculations of ZEUS test calorimeters. *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment*, 349(1):106 – 111, 1994.

[94] Geisha. *RWTH Aachen report, PITHA*, 1985.

[95] J. DeJong et al. 2009 position paper on calibration of runs I-II-III. *MINOS-DOCDB-6717*, 2010.

[96] J. Ratchford. A k-Nearest Neighbor based particle identification for the 2010 Charged Current analysis. *MINOS-DOCDB-7116*, 2010.

[97] Brian J. Rebel. First MINOS results with the NuMI beam. *Nucl. Phys. Proc. Suppl.*, 168:195–198, 2007.

[98] C. Backhouse. Estimating shower energies using a kNN-based method. *MINOS-DOCDB-6868*, 2010.

[99] M. Strait. *A Measurement of Oscillation Parameters using Antifiducial Charged Current Events in MINOS*. PhD thesis, University of Minnesota, Twin Cities, 2010.

[100] M. Strait. Position paper on the inclusion of rock and anti-fiducial events in the Charged Current analysis. *MINOS-DOCDB-6870*, 2010.

[101] J. Evans. Wrong-sign background systematic for the CC analysis. *MINOS-DOCDB-7203*, 2010.

[102] J. Evans, P. Rodrigues, J. Hartnell, and A. Himmel. The UK Beam Matrix Method. *MINOS-DOCDB-4298*, 2008.

[103] J. Evans. *Measuring Antineutrino oscillations with the MINOS experiment*. PhD thesis, University of Oxford, 2008.

[104] Reuben Phillip Litchfield. *Neutrino induced events in the MINOS detectors*. PhD thesis, University of Oxford, 2008. FERMILAB-THESIS-2008-71.

[105] A. Blake. Comparison of predicted Far Detector spectra. *MINOS-DOCDB-4257*, 2008.

[106] J. Ratchford. A data-driven estimate to the NC background. *MINOS-DOCDB-7115*, 2010.

[107] A. Blake. Spill counting at the Far Detector. *MINOS-DOCDB-7785*, 2011.

[108] A. Himmel. CC handscan results. *MINOS-DOCDB-5613*, 2009.

[109] H. Gallagher et al. Updated cross-section model uncertainties for the Charged-Current analysis. *MINOS-DOCDB-2989*, 2007.

[110] S. Coleman. Distilling systematic errors for the 2010 CC analysis. *MINOS-DOCDB-7514*, 2010.

[111] C. Backhouse. MDC Results. *MINOS-DOCDB-7180*, 2010.

[112] J. Evans et al. Charged Current analysis blessing package. *MINOS-DOCDB-7217*, 2010.

[113] A. McGowan. *Observation of deficit in NuMI neutrino-induced rock and non-fiducial muons in MINOS Far Detector and measurement of neutrino oscillation parameters*. PhD thesis, University of Minnesota, Twin Cities, 2007.

[114] J. Ratchford. CC Event Selection By Cut Table. *MINOS-DOCDB-7350*, 2010.

[115] P. Adamson et al. Measurement of the neutrino mass splitting and lepton flavor mixing by MINOS. *Submitted to Physical Review Letters*, 2011.

[116] P. Adamson et al. Measurement of antineutrino oscillations with a $\bar{\nu}_\mu$ beam in the MINOS experiment. *To be submitted to Physical Review Letters*, 2011.

[117] A. Blake. Near Detector Data/MC Plots for CC Analysis. *MINOS-DOCDB-7660*, 2010.

[118] A. Blake. Near Detector Data/MC Plots for CC Analysis. *MINOS-DOCDB-7659*, 2010.