

**A Measurement of Neutrino Oscillation Parameters
Using Anti-fiducial Charged Current Events in MINOS**

A DISSERTATION
SUBMITTED TO THE FACULTY OF THE GRADUATE SCHOOL
OF THE UNIVERSITY OF MINNESOTA
BY

Matthew Levy Strait

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

Professor Daniel P. Cronin-Hennessy, Faculty Advisor

September 2010

© 2010 Matthew Levy Strait

This thesis is distributed under the terms of the Creative Commons Attribution License, which permits unrestricted use, distribution, and reproduction in any medium, provided the original author and source are credited.

Abstract

The Main Injector Neutrino Oscillation Search (MINOS) observes the disappearance of muon neutrinos as they propagate in the long baseline Neutrinos at the Main Injector (NuMI) beam. MINOS consists of two detectors. The near detector samples the initial composition of the beam. The far detector, 735 km away, looks for an energy-dependent deficit in the neutrino spectrum. This energy-dependent deficit is interpreted as quantum mechanical oscillations between neutrino flavors. A measurement is made of the effective two-neutrino mixing parameters $\Delta m^2 \approx \Delta m_{23}^2$ and $\sin^2 2\theta \approx \sin^2 2\theta_{23}$. The primary MINOS analysis uses charged current events in the fiducial volume of the far detector. This analysis uses the roughly equal-sized sample of events that fails the fiducial cut, consisting of interactions outside the fiducial region of the detector and in the surrounding rock. These events provide an independent and complementary measurement, albeit weaker due to incomplete reconstruction of the events. This analysis reports on an exposure of 7.25×10^{20} protons-on-target. Due to poor energy resolution, the measurement of $\sin^2 2\theta$ is much weaker than established results, but the measurement of $\sin^2 2\theta > 0.56$ at 90% confidence is consistent with the accepted value. The measurement of Δm^2 is much stronger. Assuming $\sin^2 2\theta = 1$, $\Delta m^2 = (2.20 \pm 0.18[\text{stat}] \pm 0.14[\text{syst}]) \times 10^{-3} \text{ eV}^2$.

Contents

List of Tables	vi
List of Figures	vii
1 Introduction	1
2 Neutrino Oscillations	4
2.1 History	4
2.2 Theoretical Interest	6
2.3 General Neutrino Oscillation Theory	7
2.4 Three Neutrino Case	9
2.5 Two Neutrino Case	10
2.6 MINOS Measurements	10
2.6.1 ν_μ Disappearance	11
2.6.2 ν_τ Appearance	13
2.6.3 ν_e Appearance	13
3 The NuMI Beam	15
3.1 Overview	15
3.2 Near and Far Flux	17
3.3 Beam Angle	19
3.4 History of Run Conditions	20
3.4.1 Target Decay	22
4 The MINOS Detectors	23
4.1 General Detector Design	23
4.2 Steel Planes and Field Coils	26

4.3	Scintillator Planes	26
4.4	Light Collection and Readout	29
4.5	Calibration and Aging	31
4.6	Far Detector Triggering	31
4.7	Veto Shield	32
4.8	Event Characteristics and Terminology	33
5	Anti-fiducial Muons	36
5.1	Fiducial Definition	36
5.2	Anti-fiducial Regions	37
5.3	Event Characteristics	40
6	Monte Carlo	45
6.1	Beam Simulation	45
6.2	Neutrino Interactions	46
6.3	Particle Propagation	47
6.4	Detector Response	48
6.4.1	Correction to Strip Lengths	49
7	Rock Model	51
7.1	Cavern Dimensions	51
7.2	Rock Composition and Density	52
7.3	Independence of Flux and Density	57
7.3.1	Caveat: Radial Dependence	58
7.3.2	Caveat: Nuclear Composition	58
7.3.3	Caveat: Density and Decay	59
7.3.4	Caveat: Event Shape	59
7.4	Monte Carlo Density Study	59
7.5	Model of Cavern Geometry	60
8	Selection Cuts	65
8.1	Basics	65
8.2	Cosmic Rejection	66
8.3	Particle Identification	70

8.4	Performance	73
8.5	Data	76
9	Near-to-Far Extrapolation	78
9.1	Initial Beam Fit	79
9.2	Near Detector Data \rightarrow Flux	80
9.3	Flux \rightarrow Far Detector Prediction	81
9.4	Oscillation	83
10	Oscillation Analysis	84
10.1	Region Separation	84
10.1.1	Motivation	84
10.1.2	List of Regions	87
10.2	Rock/Detector Separation on the Edge	88
10.2.1	Horizontal and Vertical Edges	89
10.2.2	Diagonal Edges	90
10.2.3	Veto Shield	90
10.3	Angle	91
10.4	Binning and Fitting	92
10.4.1	Alternative Binning	95
10.5	Mock Data Challenge	96
11	Results	99
11.1	Checks	99
11.2	Alternative Binning	109
11.3	Contribution to MINOS Result	109
12	Systematic Errors	111
12.1	Shared Systematics	111
12.1.1	Relative Near/Far Normalization	111
12.1.2	Muon Track Energy	112
12.1.3	Subdominant Systematics	114
12.2	Systematics Unique to this Analysis	116
12.2.1	Cross Sections	116

12.2.2	Z/A	122
12.2.3	Alignment	122
12.2.4	Correlations Between Systematics	125
12.3	Evaluation of Systematics	128
12.3.1	Incorporation of Systematics into Fit	129
12.3.2	Nuisance Parameter Fits	129
12.3.3	Fit for This Analysis	132
12.4	Conclusions	134
13	Summary, Conclusions, and Future	135
13.1	Potential to Increase Statistical Power	136
13.2	Improvements to Systematics	137
13.3	Antineutrino Beam Analysis	138
13.4	Application to Other Experiments	139
	Bibliography	141

List of Tables

3.1	Summary of beam conditions	22
6.1	Values of NEUGEN deep inelastic scattering parameters	47
7.1	Measured composition of Ely greenstone	54
7.2	Density data from far detector core sample	64
8.1	Prediction of selected events by true interaction and selected track	70
8.2	Predicted events by region and interaction type	73
8.3	Effect of cuts on data	76
8.4	Selected events by region	77
10.1	Monte Carlo predicted events by region	87
11.1	Summary of fit results by run	100
12.1	Systematics for NEUGEN deep inelastic scattering parameters	115
12.2	Elemental compositions of the detector and rock	117
12.3	Mineral composition of Ely greenstone	121
12.4	Evaluation of all systematic errors	127

List of Figures

2.1	Masses of particles in the standard model	6
2.2	Loci of possible values of $\sin^2 \theta_{23}$ and $\sin^2 \theta_{13}$	12
3.1	The NuMI beamline	16
3.2	NuMI beam flux at the near detector	17
3.3	Pion decay kinematics causing difference in the near and far spectra	18
3.4	Differences in near and far neutrino spectra	18
3.5	Effect of beam angle on rock events	20
3.6	Effect of varying conditions among Runs I, II and III	21
4.1	NuMI baseline from Fermilab to the Soudan Mine	24
4.2	North face of the far detector	24
4.3	South face of the near detector	25
4.4	Far detector steel plane design	26
4.5	Scintillator strip design	27
4.6	Far detector scintillator modules	28
4.7	Near detector scintillator modules	29
4.8	Change in the median response to cosmic rays over time	31
4.9	Far detector veto shield	32
4.10	A typical in-detector event	35
5.1	Regions of the anti-fiducial volume	36
5.2	Detector cross section showing all scintillator strips	38
5.3	Detail of the coil hole region	39

5.4	Detail of scintillator strips on diagonal edges	39
5.5	Detail of scintillator strips on vertical and horizontal edges	39
5.6	Frequency of events by region	40
5.7	Some Monte Carlo rock events	41
5.8	Some Monte Carlo detector events	42
5.9	Typical selected data events	43
5.10	Less typical selected data events	44
6.1	Simulation of light loss at strip ends	48
7.1	MINOS far detector cavern cross section, as designed	53
7.2	MINOS far detector cavern top view, as measured	53
7.3	Plan view of the MINOS and Soudan 2 halls	55
7.4	Photo of a core sample box	55
7.5	Density data from far detector core sample	56
7.6	Number of Monte Carlo events as a function of rock density	60
7.7	Monte Carlo rock events as a function of x , y , and z position	61
8.1	Timing of selected events	67
8.2	Angle and transverse momentum cosmic rejection cuts	68
8.3	2010kNN PID for non-back and back events	72
8.4	PID selection efficiency for detector-like and rock-like events	74
8.5	Components of the selected sample by interaction type	75
9.1	Effect of beam tuning on near detector data/Monte Carlo agreement	79
9.2	The near-to-far beam matrix	80
9.3	Far detector Monte Carlo correction functions	82
10.1	Spectra of neutrino and muon energies for rock and detector events	85
10.2	Spectra of neutrino energy for front face and edge rock events	86
10.3	Rock/detector separation selection on the edges	89
10.4	Energy/angle correlations for rock and detector events	93

10.5	Sensitivities as fit is refined	93
10.6	2-dimensional energy/angle binning scheme	94
10.7	Sensitivity contours for rock and detector separately	96
10.8	Best fit points of mock data sets	97
10.9	Mock data result near the truth point	97
10.10	All mock data sets summed	98
11.1	Contours by run	100
11.2	Contours for all runs combined	101
11.3	$\Delta\chi^2$ by run at maximal mixing	101
11.4	Events in all regions by energy	102
11.5	Events in all regions by angle	103
11.6	Events by region and energy	104
11.7	Events by region and angle	105
11.8	Data events in histograms as fit	106
11.9	χ^2 as a function of $\sin^2 2\theta$	107
11.10	Contours for rock and detector separately	108
11.11	Comparison of results to 2007 results	108
11.12	Comparison of full and simplified analyses	109
11.13	Contribution to overall MINOS result	110
12.1	Systematic error on track energy	114
12.2	Error on the flux estimate	117
12.3	Spectrum of rock events showing deep inelastic scattering fraction	118
12.4	Comparison of Fermi gas and spectral function models	119
12.5	Transverse alignment of outermost strips	123
12.6	Summary of systematic errors	129
12.7	Summary of systematic errors for the simplified analysis	130
12.8	Nuisance parameter fit	132
13.1	Comparison of Δm^2 measurements	135

Chapter 1

Introduction

The MINOS experiment consists of a beam of muon neutrinos that passes through two detectors separated by 734 km. The *near detector* observes the initial beam spectrum. The *far detector* observes the spectrum after some neutrinos have oscillated into other flavors. Since these other flavors are not readily detectable in MINOS, the far detector observes a deficit compared to what would be expected without oscillations. The oscillation parameters Δm^2 and $\sin^2 2\theta$ can be extracted from the magnitude and energy dependence of this deficit. By using two similar detectors, a variety of systematic errors relating to the knowledge of the beam composition, neutrino interaction cross sections, and detector efficiencies cancel out in the near-to-far ratio. Chapter 2 covers the theory of neutrino oscillations, while chapters 3 and 4 describe the beam and detectors.

Muon neutrino charged current events generally result in a muon and a hadronic shower. In the fiducial volume — that is, away from the edges of the detector — each of these is well reconstructed and the neutrino energy found by summing their energies. This analysis measures the oscillation parameters by using the set of beam ν_μ and $\bar{\nu}_\mu$ charged current events in the MINOS far detector in which the reconstructed muon track begins outside the fiducial volume, that is, near the edge of the detector. These include both muons from neutrino interactions in the surrounding rock and from interactions around the exterior of the detector; they are collectively referred to as *anti-fiducial events*. The primary characteristic of events in this analysis is that, in general, only part of the neutrino energy is reconstructed.

For neutrino interactions on the edge of the detector, the muon is well measured, but hadronic showers are lost or only partially contained. For interactions in the rock, usually only the muon is seen and only its *remaining* energy is reconstructed; the distance it has traveled before arriving at the detector is unknown. Chapter 5 details the characteristics of these events. In order to predict the characteristics of events originating outside the detector, Monte Carlo models are crucial. Chapter 6 gives an overview of the MINOS Monte Carlo and chapter 7 specifically goes into detail about the rock model.

Only muon tracks are used in this analysis. Rock events' showers are typically entirely lost in the rock. Events in the exterior of the detector often have showers, but the containment is not modeled well enough in the Monte Carlo to warrant their use. The analysis gets much of its power from categorization of events based on location in the detector. For instance, since the beam is directed north, tracks starting on the front (south) face are nearly always rock interactions, whereas tracks starting in the back (north) of the detector are nearly all detector interactions. The fit is improved because the events in which the reconstructed muon energy more closely correlates to neutrino energy are isolated. The angle of tracks is also used in the fit, as it correlates with neutrino energy at a fixed track energy. A maximum likelihood fit is performed to determine the oscillation parameters and confidence contours. Chapter 8 explains the event selection algorithms, chapter 9 explains the procedure for extrapolating the spectrum of events at the near detector to a prediction at the far detector, and chapter 10 gives the fitting and binning methods.

The strength of this analysis is in measuring Δm^2 . Despite only partial reconstruction of events, it is possible to have significant sensitivity to Δm^2 since, at a given value of the mixing angle, it is primarily the count of events that determines Δm^2 . Contrariwise, to measure $\sin^2 2\theta$ at high precision it is necessary to have good energy resolution. Since nearly every event in this analysis has an unknown amount of lost energy, the sensitivity to this parameter is quite weak. Nevertheless, it can be used as a consistency check with other measurements. The anti-fiducial sample is statistically independent of all other MINOS samples and can be combined with the fiducial analysis for an overall improved result. A similar number of events are selected in the fiducial and anti-fiducial samples. This analysis measures Δm^2 with somewhat less than half the sensitivity as the fiducial analysis for the

same exposure: $\pm 0.18 \times 10^{-3} \text{ eV}^2$ compared to $\pm 0.08 \times 10^{-3} \text{ eV}^2$. Summed, the two analyses have a 8% better sensitivity in this parameter than the fiducial analysis alone, with only statistical errors considered. Due to correlated systematics, the overall improvement is a lesser 4%. The improvement on $\sin^2 2\theta$ is negligible. Chapter 11 shows these results.

This analysis relies heavily on Monte Carlo and cannot use the near detector as much as the fiducial analysis does for cancellation of systematic effects. There are two major concerns: (1) the composition of the rock and its nuclear cross sections and (2) the geometry of the outer edges of the detector. These systematics and others shared with the fiducial analysis are explored in chapter 12. Chapter 13 concludes.

Chapter 2

Neutrino Oscillations

2.1 History

Neutrinos were first postulated by Pauli in 1930 to explain how the electrons created by beta decay could have a wide spread of energies and still conserve energy even though no other particle was observed to be emitted [1]. We now call the type of neutrino that he considered an electron antineutrino. These were first directly observed in 1956 by Reines and Cowan [2].

In 1962, Lederman, Schwartz and Steinburger discovered the muon neutrino by showing that the neutrinos in the process $\pi \rightarrow \mu \nu$ behave distinctly from those created in nuclear beta decay [3]. With the discovery of the tau in 1975, it was natural to expect that there was an accompanying tau neutrino. Its existence was confirmed by direct observation in 2000 by the Direct Observation of Nu-Tau (DONuT) experiment at Fermilab [4].

Neutrino oscillations were first proposed in 1957 by Pontecorvo [5]. Since at that time only electron neutrinos were known, he considered the oscillation $\nu \leftrightarrow \bar{\nu}$. In 1962, Maki, Nakagawa and Sakata (MNS), with knowledge of the muon neutrino, proposed mixing between ν_e and ν_μ [6]. After the discovery of the tau, MNS's scheme was extended to the case of three neutrino mixing [7]. For these contributions, the neutrino mixing matrix is usually known as the Pontecorvo-Maki-Nakagawa-Sakata (PMNS) matrix.

The first experimental hint of the reality of neutrino oscillations occurred in an experiment by Davis in 1967 which searched for the reaction $^{37}\text{Cl} \nu_e \rightarrow ^{37}\text{Ar} e^-$ initiated by solar neutrinos. The reaction was not initially observed, with an upper limit on the rate set at $0.3 \times 10^{-35}/\text{s}/\text{atom}$, in conflict with the prediction of $(2.0 \pm 1.2) \times 10^{-35}/\text{s}/\text{atom}$ [8]. Subsequent solar neutrino measurements observed finite rates, but generally with significant deficits compared to the solar model predictions [9]. For many years it was not clear whether the source of the discrepancy lay in the solar model, in the neutrinos, or in the experiments.

A similar deficit of neutrinos was observed in atmospheric showers initiated by cosmic rays. Pions in these showers decay via $\pi \rightarrow \mu \nu_\mu \rightarrow (e \nu_e \nu_\mu) \nu_\mu$; this is predicted to be the predominant source of atmospheric neutrinos. The prediction is therefore that there are 2 muon neutrinos per electron neutrino in these showers. Experiments beginning in the 1960s sought to measure atmospheric neutrinos. Once sufficient statistics were gathered, it began to be apparent that there was a tendency for both the $\nu_\mu : \nu_e$ ratio and the absolute ν_μ flux to fall short of predictions. Perhaps the first strong hint pointing towards the oscillation hypothesis came in the early 1980s at the Irvine-Michigan-Brookhaven (IMB) experiment, which reported a 2.5σ deficit in the number of decaying (low energy) muons in their detector [10]. A much more significant deficit in both the ratio and the absolute ν_μ flux was reported in 1988 by the Kamiokande collaboration [11]. In 1998, Super-Kamiokande announced definitive evidence for the effect and for the first time fit their results under the hypothesis of $\nu_\mu \leftrightarrow \nu_\tau$ neutrino oscillations [12, 13, 14, 15]. The Soudan 2 experiment, formerly located in the cavern adjacent to MINOS, confirmed this measurement [16]. The KEK* to Kamioka (K2K) experiment again observed the same effect, this time using a well-controlled source of neutrinos, an accelerator-produced beam of ν_μ [17].

The solar neutrino situation was resolved in 2001 by the Sudbury Neutrino Observatory (SNO) experiment which, in agreement with the oscillation hypothesis, showed that there is a component of non- ν_e neutrinos in the solar flux and that the total flux agrees with the standard solar model prediction [18, 19].

**Kō Enerugī Kasokuki Kenkyū Kikō* (High Energy Accelerator Research Organization)




Figure 2.1: Masses of particles in the standard model. Approximate allowed region for neutrino masses is shown in wavy lines. It is bounded from above by cosmological bounds on the sum of neutrino masses and from below by the magnitude of Δm_{23}^2 . Without a complete picture of the PMNS matrix and without a determination of the mass hierarchy (see section 2.4), it is difficult to assign specific bounds on individual flavors [22].

2.2 Theoretical Interest

To date, neutrino oscillations provide the only direct method capable of resolving a non-zero neutrino mass. As three of the twelve fundamental fermions in the standard model, we must know these masses to have a full understanding of the Universe. We know that the neutrino masses are much smaller than those of all other massive particles (see figure 2.1). In fact, the ratio between the heaviest allowed neutrino mass and that of the next lightest particle, the electron, is *larger* than that between the electron’s mass and the mass of the heaviest known particle, the top quark. If their masses were all zero, we could easily accept this as being simply the way things are, but to be non-zero yet discrepantly small calls for some explanation. One attempt, the *see-saw mechanism*, uses the small masses of the observed neutrinos to predict a spectrum of very heavy neutrinos [20]; these neutrinos could be a component of the as-yet-unobserved dark matter known to exist in the Universe (see, e.g. [21]).

Neutrino oscillations may also have a hand in explaining the matter-antimatter asymmetry in the Universe. The usual assumption is that the Universe started out matter-antimatter symmetric and then evolved into its current asymmetric state. The most straightforward method for producing this evolution is to satisfy the three *Sakharov conditions* [23]:

1. There must be interactions that violate baryon and lepton number
2. These interactions must violate C and CP symmetry
3. They must have occurred out of thermal equilibrium in the early Universe

No observations have been made of baryon or lepton number violation, but it is axiomatic under the assumption of initial symmetry.[†] The conditions of C and CP violation allow processes that create baryons and leptons to proceed faster than their inverse processes. The third condition ensures that this speed difference has an impact; without it, the inverse processes would still bring the Universe into equilibrium with net baryon and lepton numbers of zero. The second condition is of interest here. CP violation has been observed in quarks, but it is too small to explain the observed asymmetry. There exists the possibility for large CP violation in leptons. The measurement presented in this thesis does not directly probe CP violation, but helps to lay the groundwork for future experiments that can; see section 2.6.3 below on ν_e appearance.

2.3 General Neutrino Oscillation Theory

Neutrino oscillations arise because the flavor eigenstates of neutrinos are not the same as the mass eigenstates. The flavor eigenstates determine how neutrinos interact, but the mass eigenstates determine how they propagate. The two sets of eigenstates are related by:

$$|\nu_\alpha\rangle = \sum_i U_{\alpha i}^* |\nu_i\rangle, \quad (2.1)$$

where Greek letters are used to denote flavor states ($\alpha = e, \mu, \tau$), Latin letters are used to denote mass states ($i = 1, 2, 3$), and U is the unitary PMNS mixing matrix. This formula is general; it does not depend on the number of neutrino flavors nor the weak interaction properties of the participating neutrinos.

When a neutrino is created, it is a flavor eigenstate and thus it is, in general, a superposition of N mass eigenstates, where N is the number of neutrinos that participate in mixing.

[†]Unless one can find a mechanism by which the antimatter of the Universe has been segregated to such an extent that it is unobservable today.

The propagation of the mass eigenstates follows the relation

$$|\nu_i(t)\rangle = e^{-i(Et-|\vec{p}|L)}|\nu_i(0)\rangle \quad (\text{with } c = \hbar = 1), \quad (2.2)$$

where L is the distance traveled. The propagation of the initial state $|\nu_\alpha\rangle$ is described by

$$|\nu_\alpha(t)\rangle = \sum_i U_{\alpha i}^* e^{-i(Et-|\vec{p}|L)} |\nu_i(0)\rangle = \sum_i \sum_\beta U_{\alpha i}^* U_{\beta i} e^{-i(Et-|\vec{p}|L)} |\nu_\beta\rangle. \quad (2.3)$$

Conceptually, the heavier mass states lag behind the lighter ones. Each flavor component of each mass state interferes with the corresponding flavor component of the other mass states. This interference changes the probability of measuring each flavor state. The probability of measuring the state $|\nu_\beta\rangle$ after a time t is

$$|\langle\nu_\beta|\nu_\alpha(t)\rangle|^2 = \left| \sum_i U_{\alpha i}^* U_{\beta i} e^{-i(Et-|\vec{p}|L)} \right|^2. \quad (2.4)$$

Since in any realistic experiment, the neutrinos are ultrarelativistic (in the NuMI beam, $\gamma \approx 10^9-10^{12}$), we can use the approximations

$$E = \sqrt{|\vec{p}|^2 + m_i^2} \approx |\vec{p}| + \frac{m_i^2}{2|\vec{p}|} \approx |\vec{p}| + \frac{m_i^2}{2E}, \quad t \approx L \quad (2.5)$$

to let $\exp(-i[Et-|\vec{p}|L]) \approx \exp(-imL/2E)$. Then by using the unitarity condition and quite a bit of algebra we get [22]

$$\begin{aligned} P_{\alpha \rightarrow \beta} = |\langle\nu_\beta|\nu_\alpha(L)\rangle|^2 &= \delta_{\alpha\beta} - 4 \sum_{i>j} \Re(U_{\alpha i}^* U_{\beta i} U_{\alpha j} U_{\beta j}^*) \sin^2 \frac{\Delta m_{ij}^2 L}{4E} \\ &\quad + 2 \sum_{i>j} \Im(U_{\alpha i}^* U_{\beta i} U_{\alpha j} U_{\beta j}^*) \sin \frac{\Delta m_{ij}^2 L}{2E}, \end{aligned} \quad (2.6)$$

where $\Delta m_{ij}^2 \equiv m_i^2 - m_j^2$. Restoring c and \hbar to the phases gives $\Delta m^2 L c^3 / 4E\hbar$. It is conventional to rearrange this to get

$$\frac{\Delta m^2 L c^3}{4E\hbar} = \left(\frac{10^6 \text{ eV m}}{4\hbar c} \right) \frac{\Delta m^2}{\text{eV}^2} \frac{L}{\text{km}} \frac{\text{GeV}}{E} \approx 1.267 \frac{\Delta m^2}{\text{eV}^2} \frac{L}{\text{km}} \frac{\text{GeV}}{E}. \quad (2.7)$$

In this form, it is convenient to plug in the mass squared difference, baseline length and neutrino energy for modern neutrino experiments.

2.4 Three Neutrino Case

There has been much theoretical and experimental interest in the possibility that the three known flavors of neutrinos mix with one or more *sterile* neutrinos that do not participate in weak interactions. Both the LSND [24] and MiniBooNE [25] experiments have presented evidence of $\bar{\nu}_\mu \rightarrow \bar{\nu}_e$ oscillations at small L/E indicating a value of Δm^2 incompatible with the assumption that there are only three neutrino flavors. Nevertheless, the existence of additional flavors is not yet well-established, so I will assume from here on that only three neutrinos are relevant for oscillations. It is likely that even if there are, in fact, additional flavors, the following formalism is a useful approximation. The mixing matrix U can be parametrized by three angles and one CP-violating phase, δ :

$$\begin{pmatrix} c_{12}c_{13} & s_{12}c_{13} & s_{13}e^{-i\delta} \\ -s_{12}c_{23} - c_{12}s_{23}s_{13}e^{i\delta} & c_{12}c_{23} - s_{12}s_{23}s_{13}e^{i\delta} & s_{23}c_{13} \\ s_{12}s_{23} - c_{12}c_{23}s_{13}e^{i\delta} & -c_{12}s_{23} - s_{12}c_{23}s_{13}e^{i\delta} & c_{23}c_{13} \end{pmatrix}, \quad (2.8)$$

where $s_{ij} \equiv \sin \theta_{ij}$ and $c_{ij} \equiv \cos \theta_{ij}$. Two Majorana phases have been omitted since they do not contribute to oscillations. It is known experimentally that θ_{12} and θ_{23} are large, with $\sin^2 2\theta_{23} > 0.965$ at 90% confidence and $\sin^2 2\theta_{12} = 0.861^{+0.027}_{-0.021}$ [19, 26]. On the other hand, θ_{13} is small, with $\sin^2 2\theta_{13} < 0.19$ at 90% confidence [27]. Nothing is known about the value of δ . Current theory does not give any compelling suggestions as to whether or not θ_{23} is exactly $\pi/4$, θ_{13} exactly zero, or what the value of δ is likely to be.

The mass squared difference $|\Delta m_{23}^2| \approx |\Delta m_{13}^2|$ is known to be much larger than $|\Delta m_{12}^2|$ with the former being $2.43 \pm 0.13 \times 10^{-3} \text{ eV}^2$ and the latter $7.59^{+0.20}_{-0.21} \times 10^{-5} \text{ eV}^2$ [19, 28]. It is not known whether the nearly degenerate pair (ν_1, ν_2) is more or less massive than the lone ν_3 .[‡] If ν_3 is the heaviest, the mass hierarchy is termed *normal*, since in that case the electron neutrino has the smallest mass and the tau neutrino the largest. In the other

[‡]By definition ν_1 is lighter than ν_2 .

case, the mass hierarchy is termed *inverted*. There is no consensus on which spectrum to expect. There is a potential for this ambiguity to be resolved in the next generation of long-baseline experiments by observation of the difference in oscillations in neutrinos and anti-neutrinos [29, 30] or via the combination of precision measurements from both long-baseline and reactor experiments [31].

2.5 Two Neutrino Case

If only two neutrinos participate in mixing, the situation simplifies considerably. This is often a good approximation of experimental situations due to the smallness of θ_{13} and the fact that $|\Delta m_{23}^2| \gg |\Delta m_{12}^2|$. The two-neutrino mixing matrix has only one free parameter:

$$U = \begin{pmatrix} \cos \theta & \sin \theta \\ -\sin \theta & \cos \theta \end{pmatrix}. \quad (2.9)$$

And the probabilities become

$$P_{\alpha \rightarrow \beta, \alpha \neq \beta} = \sin^2 2\theta \sin^2 \frac{\Delta m^2 L}{4E}, \quad (2.10)$$

$$P_{\alpha \rightarrow \alpha} = 1 - \sin^2 2\theta \sin^2 \frac{\Delta m^2 L}{4E}. \quad (2.11)$$

2.6 MINOS Measurements

The two mass splittings $\Delta m_{13}^2 \approx \Delta m_{23}^2$ and Δm_{12}^2 set the frequencies of superimposed oscillations. The latter controls oscillations much slower than the former, and so as a good approximation the slow oscillations can be ignored for a short enough baseline; MINOS falls in this category. In this case, the oscillation probabilities adopt an effective form identical to equation 2.10, but in which ν_β may be a linear combination of ν_e and ν_τ . MINOS primarily observes ν_μ disappearance; the details of how this measurement is interpreted are discussed below. Searches for ν_e and ν_τ appearance are also possible; a brief discussion of these follows.

2.6.1 ν_μ Disappearance

For ν_μ disappearance, the above approximation allows MINOS data to be fit to a model with only two free mixing parameters: $\sin^2 2\theta_{\text{effective}} \equiv \sin^2 2\theta$, defined as the depth of the oscillation dip, and $\Delta m_{\text{effective}}^2 \equiv \Delta m^2$, defined by the L/E at the bottom of the oscillation dip, as per equation 2.10. In this form, the sign of Δm^2 has no effect, and so without loss of generality, I will drop the absolute value signs.

To report physical results, we must translate these effective two-neutrino parameters into the physical three-neutrino parameters. The MINOS error on Δm^2 is comparable to the value of Δm_{12}^2 , so it is interesting to ask whether Δm^2 is closer to $|\Delta m_{13}^2|$ or $|\Delta m_{23}^2|$. The exact form for muon neutrino disappearance for any number of neutrinos is

$$P_{\mu \rightarrow \mu} = 1 - 4 \sum_{i>j} |U_{\mu i}|^2 |U_{\mu j}|^2 \sin^2 \frac{\Delta m_{ij}^2 L}{4E}. \quad (2.12)$$

For the 3-neutrino case, as shown in [31], if terms of order $(\Delta m_{12}^2/\Delta m_{13}^2)^2$ are dropped, the Δm^2 measured by MINOS is

$$\Delta m^2 = r |\Delta m_{13}^2| + (1-r) |\Delta m_{23}^2|, \quad (2.13)$$

where

$$r \equiv \frac{|U_{\mu 1}|^2}{|U_{\mu 1}|^2 + |U_{\mu 2}|^2}. \quad (2.14)$$

In other words, Δm^2 is a weighted average of the two physical splittings where the weights are the probability of finding a ν_μ to be in the corresponding mass state. If $\theta_{13} = 0$, $r = \sin^2 \theta_{12}$. In this case, given $\sin^2 2\theta_{12} = 0.861$, $r = 0.31$. For $\sin^2 2\theta_{23} = 1$ and $\sin^2 2\theta_{13} = 0.1$, r ranges between 0.18–0.47 depending on the value of δ .

Likewise, the effective $\sin^2 2\theta$ measured by MINOS is as a good second order approximation, in the limit of $|\Delta m_{23}^2| = |\Delta m_{13}^2|$ [26]:

$$\sin^2 2\theta = 4 \sin^2 \theta_{23} \cos^2 \theta_{13} (1 - \sin^2 \theta_{23} \cos^2 \theta_{13}). \quad (2.15)$$




Figure 2.2: Loci of possible values of the physical parameters $\sin^2 \theta_{23}$ and $\sin^2 \theta_{13}$ for several values of the measured effective parameter $\sin^2 2\theta$. The CHOOZ limit for $\sin^2 \theta_{13}$ is shown. The region to the right is excluded at 90% confidence.

Therefore, $\sin^2 2\theta$ should be interpreted as a measurement of $\sin^2 \theta_{23} \cos^2 \theta_{13}$, where

$$\sin^2 \theta_{23} \cos^2 \theta_{13} = \frac{1 \pm \sqrt{1 - \sin^2 2\theta}}{2}. \quad (2.16)$$

(It is tempting to call the right side of this equality “ $\sin^2 \theta$ ”, but this is not a physical quantity and could only be defined as $\sin^2 \left(\frac{1}{2} \arcsin \sqrt{\sin^2 2\theta} \right)$, which obscures the two-fold ambiguity manifestly displayed above.)

To put it another way, if one had a measurement of $\sin^2 \theta_{13}$, say from a reactor experiment, as well as MINOS’s measurement, $\sin^2 \theta_{23}$ is determined by (see figure 2.2)

$$\sin^2 \theta_{23} = \frac{1 \pm \sqrt{1 - \sin^2 2\theta}}{2(1 - \sin^2 \theta_{13})}. \quad (2.17)$$

In the absence of knowledge of θ_{13} , even a perfect measurement of $\sin^2 2\theta$ gives quite limited information about θ_{23} . For instance, for $\sin^2 2\theta = 1$, $\sin^2 \theta_{23}$ can lie anywhere from 0.5 to 1. Fortunately, we know from CHOOZ that $\sin^2 \theta_{13} < 0.04$ (90% confidence) [27]. At such small values equation 2.17 becomes approximately linear in $\sin^2 \theta_{13}$:

$$\sin^2 \theta_{23} = \frac{1 \pm \sqrt{1 - \sin^2 2\theta}}{2} (1 + \sin^2 \theta_{13}). \quad (2.18)$$

From this form it is clear that for a fixed $\sin^2 2\theta$ near 1, a 1% change in $\sin^2 \theta_{13}$ gives a $\sim 0.5\%$ change in $\sin^2 \theta_{23}$.

2.6.2 ν_τ Appearance

In the limit of $\theta_{13} = 0$ and $|\Delta m_{23}^2| \gg |\Delta m_{12}^2|$, ν_μ transition purely into ν_τ . These ν_τ have a charged current interaction threshold of ~ 3.5 GeV and so in MINOS such interactions are rare, but not totally negligible. In some fraction of these interactions, a muon is produced either by tau decay or as part of a hadronic shower. This muon is generally much lower energy than the incident neutrino and so these events form a small low-energy background to ν_μ disappearance.

If ν_τ appearance could be confirmed by MINOS, this would be a significant accomplishment. It would both confirm the interpretation of ν_μ disappearance and could set stronger limits on sterile neutrino scenarios. Identification of ν_τ charged current events in MINOS, however, is extremely difficult. The detector is orders of magnitude too coarse to see the tau itself. Tau neutrino events in which the tau decays via $\tau \rightarrow \mu \nu_\mu \nu_\tau$ or $\tau \rightarrow e \nu_e \nu_\tau$ are essentially indistinguishable from ν_μ and ν_e charged current events. One must therefore look for decays into hadrons, which strongly resemble neutral current events. The estimated total number of ν_τ charged current interactions in the far detector fiducial volume over all running is roughly 15–20, whereas the number of neutral current events is about 30 times larger. This makes a search extremely difficult, and no results have been announced to date. The OPERA[§] experiment, which has been built specifically to observe ν_τ appearance, has recently announced its first candidate ν_τ event [32]. Over the experimental run ~ 10 such events are expected, with a background of less than 1. Given this, it seems unlikely that MINOS will be able to significantly add to world knowledge of this process.

2.6.3 ν_e Appearance

If θ_{13} is not zero, a small fraction of the ν_τ appearance is replaced by ν_e . Unlike ν_τ events, these have effectively no interaction threshold, but they are also less likely to produce a muon and be selected in this analysis. Overall, the effect of non-zero θ_{13} is to reduce the

[§]Oscillation Project with Emulsion-tRacking Apparatus

background for ν_μ disappearance, but it is a small effect on a small background and is neglected. A tiny fraction of ν_e appearance is also expected from oscillations driven by θ_{12} . This effect is much smaller than the upper limit of ν_e appearance due to θ_{13} and so is also neglected.

Observation of ν_e appearance in MINOS is possible for sufficiently large values of θ_{13} . This measurement is of great importance since it can allow access to the CP violating parameter δ . From the form of the PMNS matrix (equation 2.8) it is clear that if θ_{13} vanishes, the value of δ becomes inaccessible. While the search for ν_e appearance in MINOS is, like that for ν_τ appearance, complicated by background from neutral current events, the higher potential event rate and characteristic features of electromagnetic showers initiated by the primary electron in ν_e charged current interactions make it feasible. MINOS has announced new limits on $\sin^2 2\theta_{13}$ that are slightly more restrictive than the CHOOZ result [33].

Chapter 3

The NuMI Beam

The Neutrinos at the Main Injector (NuMI) beamline [34] was built for the MINOS experiment. It produces a high intensity, well understood flux of muon neutrinos. Since 2009 it has also served the Main Injector Experiment for ν -A (MINER ν A) [35], which measures neutrino cross sections, and has been used briefly by several other groups. In the future it will be upgraded for use by the NuMI Off-axis ν_e Appearance (NO ν A) experiment [29].

3.1 Overview

The NuMI beam is produced at Fermilab using 120 GeV protons from the Main Injector. Every 2.2 seconds a $10\ \mu\text{s}$ pulse of $\sim 3 \times 10^{13}$ protons is directed at a segmented graphite target of total length 95.38 cm. This pulse is called a *beam spill* or simply a *spill*. Referring to these spills, NuMI beam exposure at both detectors is always quoted in *protons on target* (POT). The number of POT/spill varies depending on conditions at Fermilab. In particular, NuMI receives $\sim 20\%$ more protons per spill if the Tevatron is not running. Due to these variations, the spill may or may not have a constant intensity over its $10\ \mu\text{s}$ duration, or may be slightly shorter or longer.

Interactions of the protons with carbon creates secondary particles, mostly pions and kaons. Pions of the desired charge sign and momentum are collimated by two magnetic “horns” and decay into muons and neutrinos as they traverse a 675 meter long, 2 meter




Figure 3.1: The NuMI beamline, not to scale.

diameter pipe. The pipe is pointed downwards toward Soudan, Minnesota. Muons resulting from pion and kaon decay, undecayed secondaries, and the significant fraction of protons that do not interact in the target are monitored at the end of the decay tunnel by ionization chambers known as the *hadron monitor*. Three *muon monitors* behind successively longer lengths of absorbing material monitor the muon flux. See figure 3.1.

Because the oscillation parameters were poorly known during design of the beamline, NuMI was designed with the capability of providing several beam energies by varying the position of the target and/or the downstream horn. Since Δm^2 turned out to be at the very bottom of the expected range, MINOS nearly always uses the *low energy* (LE) configuration in which the second horn is 10 m downstream of the first, the target is 10 cm upstream of its reference point on the first horn, and the horn current is 185 kA. A small amount of production data was also taken with the downstream horn in the same position, but with the target at 250 cm, and with 200 kA horn current. This is known as the *pseudo-high energy* (pHE) configuration. (The “real” high energy configuration involves moving the downstream horn, but this is much more difficult. Due to the low value of Δm^2 , it has never been considered worthwhile.) The flux from these two configurations is shown in figure 3.2. Various other configurations have been used for short testing periods; these are not used directly in analysis.

In any configuration, there is some population of secondaries that leave the target traveling along the axis of the horns with small transverse momentum. These particles are unaffected by the magnetic field and so enter the decay volume regardless of charge sign or energy. These generally produce neutrinos quite a bit higher in energy than the LE




Figure 3.2: NuMI beam flux at the near detector in terms of number of events expected in the near detector. The LE and pHE configurations are shown. For each, the nominal Monte Carlo is shown in grey while the Monte Carlo tuned by observation is shown in black. Note how above ~ 17 GeV the two spectra become the same. Here the neutrino parents are unfocused, low p_T , secondaries. Adapted from reference [28].

beam focusing peak, so the high-energy tail due to this effect is quite significant. In the LE configuration, 87% of ν_μ originate from π^+ decay, 13% from K^+ two-body decay, and $<0.1\%$ from all other contributors, including K^- , K^0 , μ , and three-body K^+ decay [36]. Interactions at the far detector, in the absence of oscillations, are expected to be 91.9% ν_μ , 7.0% $\bar{\nu}_\mu$, 0.9% ν_e , and 0.1% $\bar{\nu}_e$. Since it originates from unfocused secondaries, the peak of the $\bar{\nu}_\mu$ spectrum is much higher energy than the ν_μ , at about 9 GeV.

3.2 Near and Far Flux

The far detector sees the beam as originating from a point source and sees only neutrinos with extremely small angles to the beam axis, but the near detector sees a line source and samples a larger range of angles. This leads to somewhat different observed spectra in the two detectors.

The effect comes in two parts. First, pions — particularly those decaying near the end of the decay volume — see the near detector subtending a large angle. Larger angle decays of




Figure 3.3: Pion decay kinematics causing difference in the near and far spectra. The shown pion can decay sending a neutrino at a variety of angles. Only very small angles can reach the far detector, but a significant spread of angles can reach the near detector.

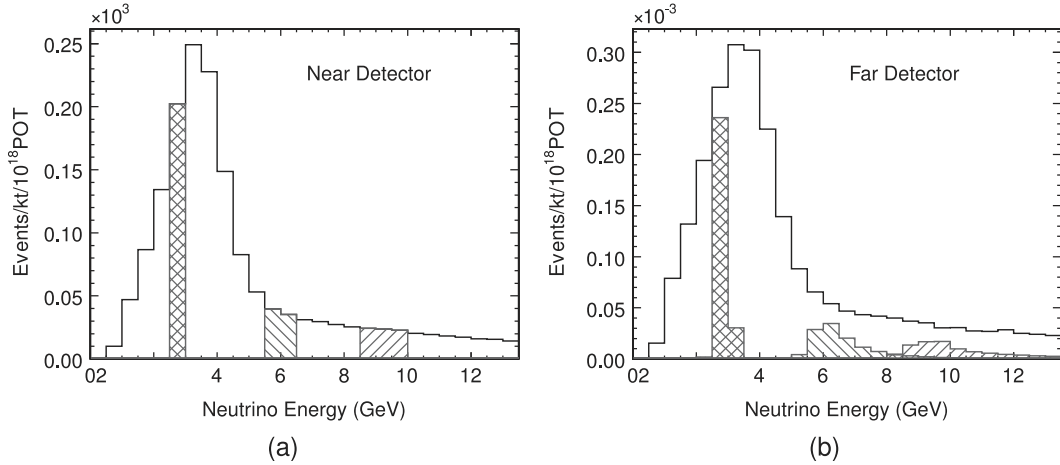


Figure 3.4: Differences in neutrino spectrum between (a) the near detector and (b) the far detector as a result of pion kinematics and angular acceptance. The shaded bands show, for three energy ranges at the near detector, what energy neutrinos result from the same decays at the far detector. Adapted from reference [36].

these pions producing slightly lower energy neutrinos can therefore reach the near detector, but not the far detector (see figure 3.3). The MINOS detectors cannot resolve the angle of the incident neutrino at this level and neither can it be judged by the transverse position, since the pion itself can have a variety of initial transverse positions in the decay volume. Because the location of the pion decay is a function of pion energy, the difference in angular acceptance between the two detectors is also a function of energy.

Second, higher energy pions with sufficiently small transverse momentum result in decays in which the neutrino nearly always intersects the near detector. As the pion energy increases above this point, the fraction of decays resulting in a neutrino that intersects the far detector increases without any increase at the near detector.

The overall effect is that the spectrum at the near detector has a somewhat stronger focusing peak and the spectrum at the far detector has somewhat more high energy tail (see figure 3.4). The method of handling these differences is given in chapter 9.

3.3 Beam Angle

Due to the curvature of the Earth, the NuMI beam is angled downwards at Fermilab and upwards at Soudan. The Earth is not spherical and it is rotating, so the angle between the beam and local *down* — i.e. the direction that objects fall — is not the same at the two sites. This is the relevant angle since the detectors are installed on floors that are level with respect to local down. Taking the beam to be the line that connects the NuMI target with the center of the far detector, the beam is pointed downwards 3.34022° at Fermilab and upwards 3.27646° at Soudan. The locations of the target and the far detector were determined primarily using by the Global Positioning System. In the case of the far detector, an inertial survey was used to find the position of the detector with respect to the top of the mine shaft. The inertial survey is the least accurate component of the measurement, with an error of 0.7 m on each coordinate, or 5×10^{-5} degrees. The error, being much smaller than the size of the detector, is negligible.

Precision knowledge of the angle at the far detector is irrelevant for reconstruction, since the angular resolution of the detector is on order of 1 degree. Nor does it affect any fiducial analysis, since it does not significantly change the character of events originating in the detector. However, it does have a significant effect on Monte Carlo generation of rock events, since the probability that a muon produced in the rock above the far detector can be scattered downwards and detected decreases as the beam angle increases.

In previous analyses, the Monte Carlo used a far detector beam angle that was 0.07° larger than the correct figure, while in other parts of MINOS analysis code, a variety of other angles in the neighborhood of 3 – 3.3° were in use. I noticed this and worked to track down an authoritative source giving the above precise values. These angles are now used for Monte Carlo generation. This correction is not large, but even such a small change in angle has a measurable effect. Rock events entering the top of the detector are about one third as




Figure 3.5: Effect of angle on the ratio between rock events entering the top of the far detector and all events entering either the top or the bottom. The Monte Carlo was run with four angles. In order, these were zero, the correct angle, a realistic intermediate angle, and the old angle used in the Monte Carlo. On the left all four are shown; the fit is done on this plot. On the right, the plot is zoomed in on the realistic angles and the same fit is shown. The top (bottom) of the detector is defined as the top (bottom) horizontal edge and two upper (lower) diagonal edges.

common as those entering the bottom. The corrected figure increases the top:bottom ratio by 1.7% (see figure 3.5). Since the angle changes the effective surface area of the detector, it also modifies the overall flux of reconstructable rock muons, but this is a smaller effect.

3.4 History of Run Conditions

The first beam data used in this analysis was taken on 20 May 2005. The following period, called Run I, was taken in the LE configuration and continued until 26 February 2006. During this period the decay pipe was evacuated to 0.5 Torr to minimize scattering of secondaries.

After a shutdown period, the beam was run in the pHE configuration from 11 June 2006 to 13 August 2006. This period is called Run IpHE or simply “the pHE running”.

In order to return the beam to LE mode, it was necessary to replace the target since the motion system of the first target had failed. With the new target in place, Run II data was taken in LE mode from 12 September 2006 to 17 July 2007. The second target was later found to have been misaligned by ~ 1 cm longitudinally such that it was at 9 cm relative to the first horn rather than the nominal 10 cm [37]. The effect of this was to reduce the focusing peak by $\sim 5\%$ (see figure 3.6).




Figure 3.6: Effect of varying conditions between Runs I, II and III. Shown is the Monte Carlo prediction of all reconstructed events at the far detector in the fiducial volume, normalized by POT. Run II has a reduced focusing peak due to target misalignment. Run III has a reduced focusing peak due to helium in the decay pipe. Compared to Run I, the high energy tail (above 8 GeV) in Run II is 1.1% larger. In Run III it is 3.1% larger.

During the following shutdown the decay pipe was filled with helium at ~ 1 atm due to concerns that the beryllium window at the upstream end might fail, causing a catastrophic implosion. The following running, Run III, makes up the majority of the data used in this analysis. It began on 18 November 2007 and continued until 13 June 2009. The addition of helium to the decay pipe modifies the spectrum in several ways. Protons that pass through the target can interact with the helium, producing unfocused secondaries; this generally increases neutrino production and has the most pronounced effect on the high energy tail. Meanwhile, properly focused pions can be scattered out of the pipe before decaying, so the beam peak is suppressed. Run III is the last data used in this analysis; see table 3.1 for a summary of the exposures and run conditions for Runs I–III.

Running resumed in September 2009 in antineutrino mode (Run IV) with a third target and then switched back to neutrino running for March–July 2010 (Runs V and VI). After a shutdown and another target replacement, MINOS running resumed (Run VII) in neutrino mode in Aug 2010. MINOS is expected to run until the summer of 2012, but no firm plans for its end have been made as of this writing.

	POT/10 ²⁰	Comments
Run I	1.2693	First target, at 10 cm
Run IpHE	0.1531	Same target, at 250 cm
Run II	1.9427	Second target, at 9 cm
Run III	3.8809	Helium in the beam pipe
Total LE	7.0929	

Table 3.1: Summary of data used in this analysis. While the absolute intensity of the proton beam striking the NuMI target is known only to 1.0%, this error is almost entirely due to calibration and is strongly correlated between runs [36, 38].

3.4.1 Target Decay

A gradual degradation of the NuMI target due to radiation damage was observed over the course of Runs II and III by monitoring the number of near detector events as a function of protons on target. Because the first target was used for only 1.3×10^{20} POT as compared to the second target's 6.1×10^{20} POT, the same effect was not evident in Run I. However, it is expected that this target also degraded over time. A mechanical inspection of the target to verify the damage has not yet been performed due to high radiation levels.

The effect of the target damage is roughly the same as if the target were getting shorter over time via loss of material at the upstream end: more protons pass through without interacting and so the neutrino beam intensity drops. The focusing peak suffers more than the high energy tail since the horn configuration is optimized for interactions in an intact target; over the course of Runs II and III its intensity dropped $\sim 10\%$.

Chapter 4

The MINOS Detectors

The MINOS detectors are a pair of magnetized steel and solid scintillator tracking calorimeters. A small unmagnetized calibration detector was also built to characterize the response of the main detectors to a variety of particles. The two primary detectors and the calibration detector are described in detail in reference [39]. Here I will briefly describe the near and far detectors.

The MINOS near detector is installed at Fermilab in a cavern 100 m below the surface (225 m water equivalent). 240 m of rock separate it from the end of the beam's 675 m decay volume. The MINOS far detector is installed in the Soudan Underground Laboratory in northern Minnesota (part of the Soudan Underground Mine State Park and former site of the Soudan Iron Mine), 705 m below the surface (2070 m water equivalent), 735.3 km from the NuMI target. Figure 4.1 shows the baseline.

4.1 General Detector Design

The near and far detectors are designed to be very similar in order to minimize systematic errors associated with near-to-far extrapolation. They are not, however, identical. Most of the differences between the two detectors were necessary due to the much higher rate of interactions at the near detector, which sees a flux roughly 10^6 times larger than the far detector. These differences are detailed in the detector element sections below.



Figure 4.1: NuMI baseline from Fermilab to the Soudan Mine.




Figure 4.2: North face of the far detector, showing (A) a steel plane (B) the veto shield (C) the field coil and (D) an electronics rack. Figure taken from reference [39].




Figure 4.3: South face of the near detector, showing (A) a steel plane (B) the field coil and (C) an electronics rack. Figure taken from reference [39].

Each detector is constructed of a series of planes of 2.54 cm steel and 1.0 cm solid scintillator, separated from each other by air gaps of 2.4 cm. The steel provides mass for neutrino interactions and also supports the magnetic field. Each plane has a hole in it to admit the field coil. The curvature of muons in the field allows charge sign determination and momentum measurement. The field polarity is chosen to focus the desired sign of muon (μ^- during normal neutrino running) towards the center of the detector, making it more likely that they are fully contained. The scintillator is made of up strips oriented 45° to horizontal, with planes alternating between the two possible orientations. In a right-handed coordinate system in which the beam direction is $+z$, the two orientations are known as $u \equiv (y + x)/\sqrt{2}$ and $v \equiv (y - x)/\sqrt{2}$. Scintillator light is captured by wavelength-shifting fibers and read out using multipixel photomultiplier tubes.

The far detector has a total of 486 planes, 484 of which are instrumented with scintillator. Its total mass is 5.4 kt. It is divided into two *supermodules* of 249 and 237 planes respectively, separated by a 1.15 m air gap. Far detector planes are nearly regular octagons 8 m across. See figure 4.2. The near detector has a total of 282 planes, 153 of which are instrumented. Its total mass is 980 t. Near detector planes are “squashed octagons”, 6.2 m wide and 3.8 m tall [36]. See figure 4.3.

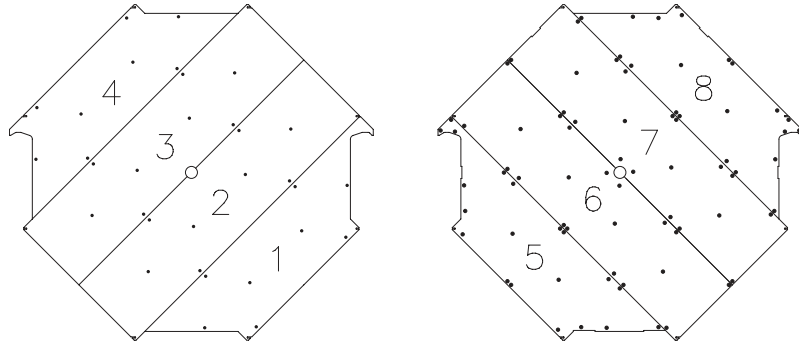


Figure 4.4: Far detector steel planes are built from two 1.27 cm layers plug-welded together. Each layer is made from four plates of steel. The dots show holes for plug-welds and handling. Figure taken from reference [39].

4.2 Steel Planes and Field Coils

In the far detector, steel planes are made from two layers of 1.27 cm thick plates plug-welded together (see figure 4.4). Each layer is composed of 4 plates. This design was the result of the size and weight limits imposed by the mine shaft. The steel thickness is uniform at a level of 0.62%. At the near detector each steel plane is a solid plate of 2.54 cm steel with thickness variation of $\sim 0.3\%$.

The steel planes are magnetized by a field coil. The near detector has one such coil while the far detector has an independent coil for each of the two supermodules. The near detector coil hole is centered vertically, but is off center horizontally so that the beam spot (and therefore the fiducial region) is comfortably away from it. At the far detector, where the beam spot is kilometers wide, the coil hole is centered horizontally and vertically. In all cases, the return leg of the coil is outside the detector. The near detector has an field of 1.17 T at the center of the fiducial volume, and the far detector has an average field of 1.27 T [36].

4.3 Scintillator Planes

The scintillator planes are composed of a number of modules containing 16*, 20, or 28 scintillator strips each. These strips are 4.1 cm \times 1.0 cm extruded polystyrene loaded with

*Here reference [40] disagrees with [39], which says 14; I take the former as more authoritative.




Figure 4.5: Scintillator strip design. Charged particles passing through the strips create scintillation light. This light reflects diffusely from the internal surfaces of the strip. Fiber epoxied in a groove in the strip absorbs this light and re-emits it isotropically. The fraction of re-emitted light which satisfies the condition for total internal reflection is transported down the fiber to the module end. Adapted from reference [39].




Figure 4.6: Far detector scintillator modules. Modules types A and B have 28 strips of varying length, while C–F have 20 that are all the full 8 m long. Figure taken from reference [39].

scintillator fluors PPO[†] and POPOP[‡] and coextruded with a 0.25 mm reflective coating of polystyrene containing 15% TiO₂. Each strip has a 2.3 mm×2.0 mm groove in the center of the 4.1 cm side into which is epoxied 1.2 mm diameter wavelength-shifting fiber. Scintillation light produced by charged particles passing through a scintillator strip is multiply reflected by the coating. The fraction of this light absorbed by the wavelength-shifting fiber is re-emitted isotropically at a longer wavelength and that which is captured by total internal reflection is channeled to the end of the strip. See figure 4.5.

Every scintillator strip end is cut at a right angle to its length. This means that in modules of varying scintillator strip length, the scintillator boundary forms a sawtooth pattern. Scintillator strips that intersect the coil hole are similarly cut at right angles to make room for the coil. The fiber is not cut here, but rather is routed around the coil so that these strips have the same light collection efficiency as unbroken strips.

There are no design differences between the near and far detectors for the fiber and scintillator themselves. However, the different sizes and shapes of the two detector necessitated different packaging of the scintillator into modules as shown in figures 4.6 and 4.7. At the far detector, every steel plane is covered ~99% with scintillator except for the first in each

[†]2,5-di(phenyl)-1,3-oxazole; C₁₅H₁₁NO.

[‡]5-Phenyl-2-[4-(5-phenyl-1,3-oxazol-2-yl)phenyl]-1,3-oxazole; C₂₄H₁₆N₂O₂.




Figure 4.7: Near detector scintillator modules. There are 16 types in all: 6 for partially instrumented planes and 10 for fully instrumented planes. The beam center strikes the near detector halfway between the coil hole and the left edge of the scintillator. Figure taken from reference [39].

supermodule, which is uninstrumented. At the near detector, every fifth steel plane has nearly complete scintillator coverage, these are known as *fully instrumented*. The southern (upstream) 120 planes, other than those divisible by five, have scintillator coverage near the beam center; these are known as *partially instrumented*. The northern (downstream) planes not divisible by five are completely uninstrumented. This more-instrumented upstream region is known as the *calorimeter*; the remainder of the detector is the *muon spectrometer*.

4.4 Light Collection and Readout

The wavelength-shifting fibers of each module are collected together in a manifold at each of the module ends and interfaced with clear fiber which transports the light to the phototubes. At the near detector, only one end of the modules is instrumented; at the other end, the fibers are mirrored. At the far detector, both ends are instrumented.

At the near detector, each clear fiber corresponds to one pixel on a 64 pixel phototube. In the 120 planes at the south end, each phototube pixel is read separately. In the sparsely instrumented muon spectrometer region, groups of four phototube pixels are summed before being read out. The grouping is chosen so that the summed strips are about 1 m apart; this allows for several muons to be tracked without ambiguity in one beam spill.

At the far detector, groups of 8 fibers are multiplexed onto each pixel of 16 pixel phototubes. Within each multiplexed group, the fibers correspond to strips that are about 1 m apart. At the other end of the module, the pattern of groupings is different in a fashion that allows unambiguous demultiplexing in software at reconstruction time, provided that events do not have a u or v extent of more than 1 m on any plane. Very rarely ($\lesssim 0.1\%$ of the time), optical cross-talk will cause the wrong solution to be found. Some of these events can be identified by looking for tracks that have sudden jumps of ~ 1 m transversely where the track angle on each side of the jump is identical.

The single photoelectron dark noise rate is 4 Hz/pixel at the near detector and is 25 Hz/pixel at the far detector. At the far detector, a pixel is read out if a signal of at least $1/3$ of a photoelectron, giving a noise rate of ~ 400 Hz/pixel. The majority of noise does not satisfy the triggers (see section 4.6); the rate of noise events that are written out is 30 Hz at the far detector.

The far detector was designed with the requirement that a minimum ionizing particle normally incident at the extreme end of a strip should produce an average of at least 1.0 photoelectron in the phototube at the other end. The mean measured response to such a particle at the center of an 8 m far detector strip is 3.8 photoelectrons.

The far detector timing resolution is 2.3 ns. This was designed with the goal of being sufficient to distinguish upward-going muons from atmospheric events from downward-going cosmic rays. While sufficient for this purpose, it is not overdesigned. Tracks must be at least ~ 20 planes long for a measurement to be made of their direction with reasonable confidence. Since the most interesting part of the beam muon spectrum is below this level, timing is not used for beam analyses. This lack of timing resolution somewhat complicates cosmic ray rejection, as will be discussed in chapter 8. The near detector timing resolution is 5 ns. Here the design requirement was that multiple events in a beam spill be well separated by the reconstruction.

The efficiency of far detector planes for detecting muons is very high. By looking for gaps in cosmic ray tracks, I measured it to be 99.55% on average. That is, a muon track typically has to be 150 planes long before any plane fails to register a hit. This automatically includes all sources of inefficiency, including the muons slipping through gaps between strips. This




Figure 4.8: Change in the median response to through-going cosmic rays over the period of data used in this analysis (20 May 2005 through 13 June 2009). The zero point is arbitrary. Short term changes are mainly due to temperature variations in the detector hall. The long-term downward trend is believed to be caused primarily by scintillator aging. Figure adapted from [42].

efficiency is the average over Runs I–III. It is slowly dropping over time — at the end of Run III it was 99.4% [41].

4.5 Calibration and Aging

Each module-end manifold is outfitted with a light-injection module. These periodically illuminate the wavelength-shifting fibers in the area between scintillator and where they interface with the clear fibers. Response to this light characterizes the performance of the fibers, phototubes, and electronics over time. The response of all detector components, including the scintillator, is calibrated using cosmic muons. The detector response is generally stable with small short-term fluctuations and a long-term downwards trend in light output of $\sim 1.4\%/year$ (see figure 4.8).

4.6 Far Detector Triggering

During beam operation, a signal is sent over the Internet from Fermilab to the far detector for each beam spill. During these *spill triggers*, all detector activity is read out beginning $40\ \mu s$ before the arrival of the first neutrinos and ending $60\ \mu s$ after this time. In addition to normal spill triggers, *fake spills* are issued periodically in which the detector reads out




Figure 4.9: Far detector veto shield configuration. Figure taken from reference [39].

identically, but in the absence of beam. These can be used to characterize backgrounds (see section 8.2).

In the absence of a spill trigger, the detector is read out if significant activity is registered. The most common of such out-of-spill trigger conditions is satisfied if out of any group of 5 adjacent planes, at least 4 have one or more hit. These hits must occur close together in time, with no more than 150 ns elapsing between any pair of time-adjacent hits. Using the same timing requirement, the far detector also triggers if any 4 adjacent planes record a sum of at least 1500 raw ADC counts (about 20 photoelectrons) with at least 6 total hits in these planes, or if 20 planes anywhere in the detector have activity. These conditions allow nearly all beam events to be preserved even in the case of network problems that prevent the spill trigger from arriving in time. The beam timestamps are, in this case, stored at Fermilab and later reconciled with the data taken during the network outage. These triggers also record cosmic rays and atmospheric neutrino events.

4.7 Veto Shield


To improve the rejection of cosmic rays in atmospheric neutrino analyses, the far detector (but not the near detector) is surrounded above and partially on the east and west sides by

a veto shield made of the same scintillator as is used in the main detector (see figure 4.9). Strips in veto modules are aligned with the z -axis of the detector. The scintillator is read out with 8 adjacent strips per phototube pixel. While they are read from both ends, the strip-to-pixel pattern is the same at each end, so, unlike in the body of the detector, no demultiplexing can be performed. Cosmic rays entering the detector typically pass through 1–3 veto layers, with the exception of those that enter the front or back faces of the detector.

4.8 Event Characteristics and Terminology

At the near detector, each typical beam spill results in 20–40 neutrino interactions that cause detector activity — some in the detector, some in the surrounding rock. An event display without timing information shows a confusing mess of many overlapping interactions; from this, the collection of activity during the spill has been termed a *snarl*. By extension, the data read out at the far detector during a beam spill is also called a “snarl”, even though most contain no neutrino interactions and the probability of more than one is negligible. The data recorded during a cosmic trigger, which can contain multiple tracks — although usually not snarled — is also called a snarl.

The primary signal in MINOS is a muon neutrino charged current interaction on an iron nucleus:



The muon is generally observed as a long clean track, while the hadronic debris, X , is observed as a cluster of hits near the beginning of the track. In general, no structure can be clearly discerned in these clusters. In muon disappearance analyses such as this one, any such cluster is termed a *shower* and no attempt is made to reconstruct its contents. Tracks and showers are the only kinds of reconstructed objects. A snarl may have any number of tracks and any number of showers.

While there is rarely a clear joining of the shower to the track at a specific point (i.e. the neutrino interaction location), the beginning of a reconstructed track is termed the *vertex*. This notation is used universally, including cases in which no visible vertex, in the geometric sense of the word, is expected, such as low-energy quasielastic interactions, rock muons, and

cosmic rays. The end of a track is called the *end*. Showers are also assigned a vertex near the center of their upstream extent. This is done in isolation without consulting the end points of overlapping or nearby tracks.

After the reconstruction code has identified all tracks and showers, it attempts to group related sets of these into *events*. An event may have as little as one track or one shower in it. Events are assigned their own *event vertex*, which is the best estimate of the neutrino interaction point — generally this is very similar to the primary track vertex. At the near detector only, the first step is to use timing information to separate groups of tracks and showers into *time slices* (or simply *slices*). At the far detector no slicing is performed, but multiple events may be reconstructed if there are several spatially related groups of tracks and/or showers. Generally this is only useful for cosmic rays, as the probability of two beam events occurring within the same trigger is negligible. Occasionally, a neutrino interacting in the rock will produce secondaries that enter the detector a significant distance from each other and are assigned to separate events. However, in this analysis, only a single muon is used per snarl, so this case does not cause any trouble.

The energy of a track is calculated in two ways: by range and by curvature in the magnetic field. For tracks that stop in the detector, range is the more accurate measurement. If a track exits the detector, curvature must be used. All tracks are assumed to be muons for both calculations. The consequence of this is that tracks formed by other particles, such as pions or protons, which lose energy more quickly, are assigned incorrectly low energies by range and incorrectly high energies by curvature (although usually only muons form tracks sufficiently long for the curvature measurement to be meaningful).

The energy of showers in previous MINOS ν_μ analyses was estimated calorimetrically by the simple sum of the energies of the individual hits [36]. The 2010 MINOS ν_μ analysis uses a more sophisticated algorithm which takes into account other details of the shape of the shower [43]. While the anti-fiducial analysis uses only tracks at the far detector, both the fiducial and anti-fiducial analyses depend on the full reconstruction of neutrino events in the near detector to measure the neutrino flux and thereby to predict the spectrum of events at the far detector (see chapter 9).

An example event is shown in figure 4.10. The left pane shows the *xy* view, including the




Figure 4.10: A typical in-detector event (run 36649, snarl 73152). See text for details.

scintillator plane boundary, reconstructed track hits (closely set red discs) and reconstructed shower hits (purple circles). Track hits are accurately reconstructed in 3D. Shower hits are not — they are drawn using their raw u (v) position and the shower vertex's v (u), such that the hits of a shower form an X-shape. The center and right panes show the uz and vz views, with showing both raw and reconstructed hits. Raw hits above 2 photo-electrons are shown as dark blue discs. Those below are shown as light green discs; most are crosstalk and appear in a distinctive pattern paralleling the true track. Hits reconstructed as part of tracks or showers have the smaller red discs (track) or purple circles (showers) superimposed on them. This event has one inward-curving (negatively charged) track with energy by curvature of (9.7 ± 1.1) GeV and energy by range of (5.7 ± 0.3) GeV. Tracks this long are almost exclusively muons. By eye, the track appears to be exiting; the discrepancy between the two energies confirms this. While this event is anti-fiducial, the (3.3 ± 1.0) GeV reconstructed shower appears to be well contained. This is likely a (13.0 ± 1.5) GeV beam ν_μ charged current interaction.

Chapter 5

Anti-fiducial Muons

The purpose of this analysis is to recover events that fail the fiducial cut. In order to optimize sensitivity to the oscillation parameters, the antifiducial volume is logically divided into a number of regions in which the events have different characteristics. Most importantly, these divisions tend to separate events due to interactions in the rock from those due to interactions in the anti-fiducial volume of the detector.

5.1 Fiducial Definition

The fiducial volume in the far detector, illustrated in figures 5.1 and 5.2, was chosen to ensure hadronic shower containment. Hadronic showers extend predominantly forward of




Figure 5.1: Regions of the anti-fiducial volume. The lines within each supermodule show the fiducial cuts.

the neutrino interaction point and typically have a transverse extent no larger than 20 cm. To capture showers longitudinally, the vertex of the longest track in the event is required to be outside the first four planes at the south end of the detector, the eight planes at the north end of supermodule 1, the four planes at the south end of supermodule 2 and the 20 planes at the north end of the detector. For transverse containment, the vertex must be within $0.16 \text{ m}^2 < r^2 < 14 \text{ m}^2$. The fiducial volume thus extends to within ~ 8 cm of the edge of full scintillator coverage around the coil hole at its nearest point and ~ 25 cm from the exterior edge of the detector (see figures 5.3, 5.4 and 5.5). In both cases, the actual distances to the edge of scintillator varies with position around the cut since the edges are not circular. Circular cuts were chosen for convenience; most likely a square cut around the coil and an octagonal cut near the outside would perform somewhat better. In particular, near the outside corners of the detector there are substantial regions much farther than 25 cm from the edge which are nevertheless excluded from the fiducial volume.

5.2 Anti-fiducial Regions

The anti-fiducial volume consists of all regions of the detector excluded by the fiducial cut. While the fiducial volume is, by design, treated in analysis as a homogeneous entity, analysis of events in the anti-fiducial volume is greatly enhanced by dividing it up into a number of logical regions. These regions are defined by the various fiducial cuts:

- The front (south) face (hereafter sometimes called just the *front*)
- The detector regions near supermodule gap, north and south sides (when it is unambiguous, I will call this the *gap*).
- The back
- The outside radial edge (the *edge*)
- The region near the coil hole (called just the *coil hole* when unambiguous)
- Regions that are anti-fiducial due to failing both a longitudinal and a transverse fiducial cut, such as the edge of the front (*corners*)




Figure 5.2: Far detector cross section. Scintillator strips in U and V are shown along with the inner and outer fiducial cuts (white dashed circles). Detail of the anti-fiducial geometry is in following figures.




Figure 5.3: Detail of the coil hole region. The scintillator strips in U (black) and V (grey) are shown. The four concentric circles are, from largest to smallest: The fiducial cut (40 cm radius), the boundary of the steel planes (30 cm), the module boundary (19.7 cm) and the coil itself, including its cooling jacket (12.5 cm).




Figure 5.4: Detail of a diagonal edge (rotated 45°) showing scintillator strips in U (black) and V (grey) and the fiducial cut (red arc). All four diagonal edges have the same geometry (with U and V interchanged as appropriate) up to unintended differences in alignment. Note that the strips perpendicular to the edge overhang the last parallel strip by about a strip width.




Figure 5.5: Detail of the bottom edge. As with diagonal edges, all four horizontal/vertical edges nominally have the same geometry. These edges are made of A and B modules. Within each module, each strip is a strip width shorter/longer than the last, making a sawtooth at a constant x or y position. At the A/B module boundary, the step is 2.9 cm instead. This causes one view to extend 2.4 cm past the other on each side, which has a visible effect on which is more likely to be hit first by an incoming muon. This and preceding strip diagrams are accurate up to the level of alignment uncertainty discussed in section 12.2.3.

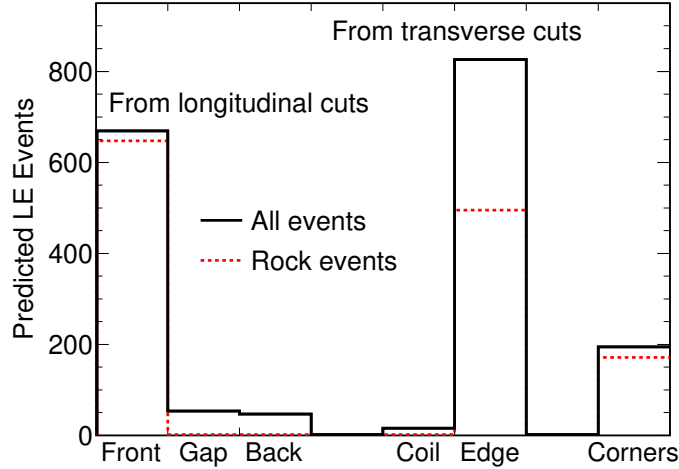


Figure 5.6: Frequency of events in each region predicted for all LE running at the measured oscillation parameters.

The front face has a large cross sectional area to accept rock events, but having only 4 steel planes it has little mass for neutrino interactions. The supermodule gap, coil hole, and the back are shielded from rock events and so most events there are in-detector interactions. The edge has both a large area exposed to rock events and a substantial mass for neutrino interactions, ~ 0.8 kt, so it has a large number of both types of events. See figure 5.6. The large bulk of all anti-fiducial events are in either the edge, the front, or the front+edge corner. Event regions are discussed in more detail in chapter 10.

5.3 Event Characteristics

Most non-muon products of neutrino interactions quickly decay or come to rest in the rock, and so rock events typically consist of a single muon. About half of the time the reconstruction finds a shower as well, but usually it is very low energy and either due to noise or the interaction of the muon with the detector somewhere along its track. Detector events, on the other hand, usually have a significant reconstructed shower. See figures 5.7, 5.8, 5.9, and 5.10 for examples.

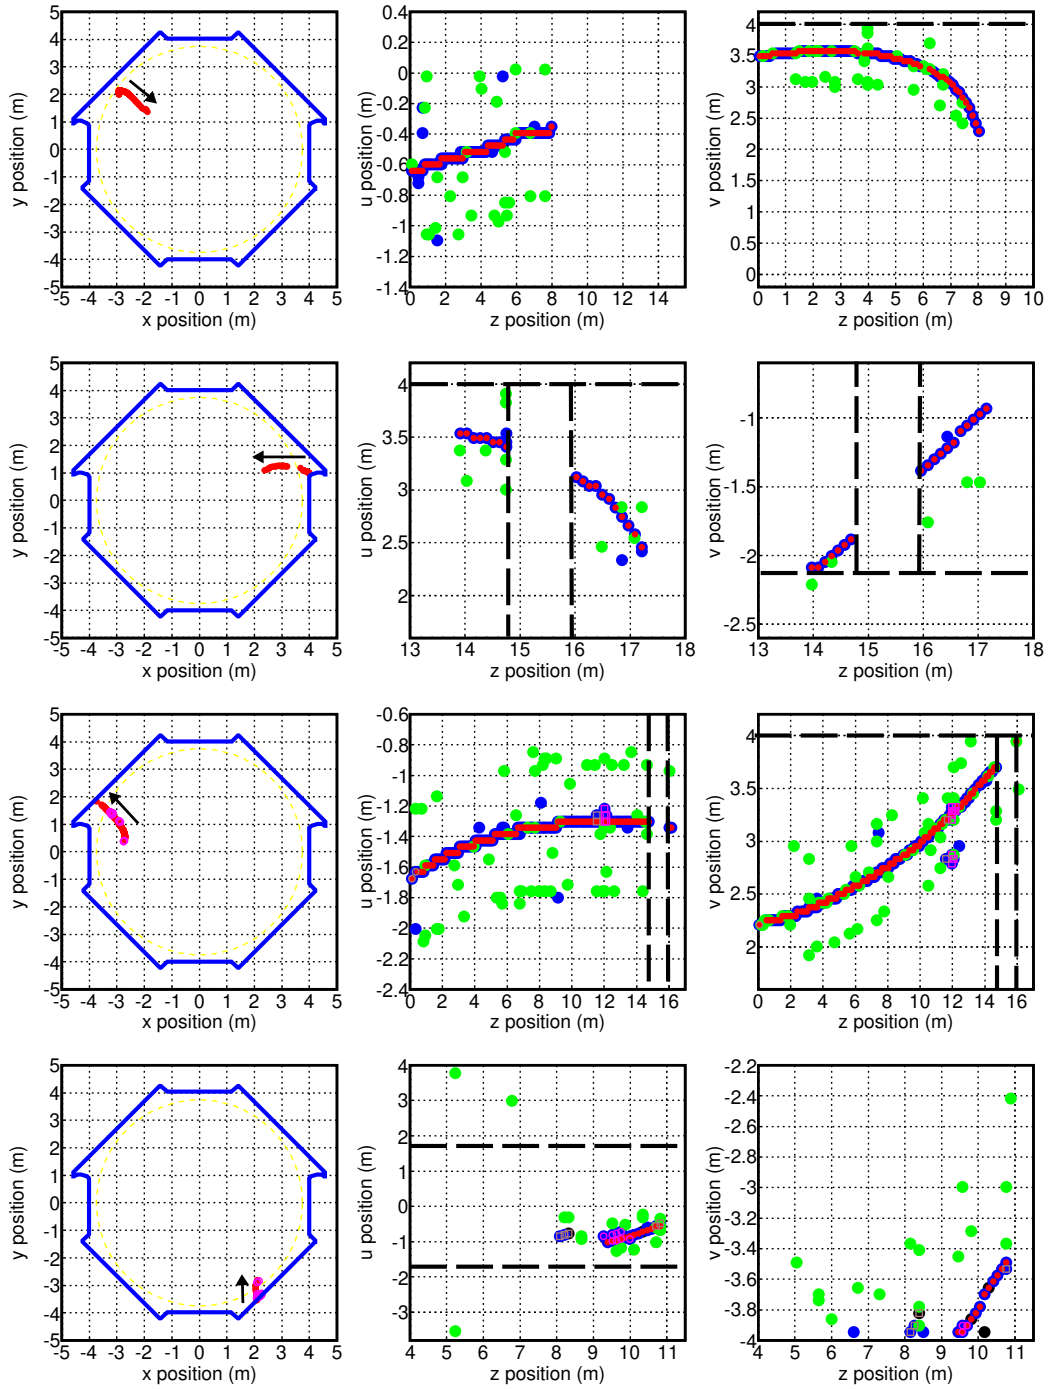


Figure 5.7: Monte Carlo rock events. For a full description of the event display, see the caption of figure 4.10. Arrows indicate muon direction. Detector edges are shown as dashed lines. These are somewhat approximate when near an edge with variable strip lengths. Top: a front face event. Second: an edge event near the supermodule gap. Third: a front face event with three showers, none of which include hadronic activity. Bottom: an edge event with a hadronic shower remnant: a pion that enters at $z = 8$ m.

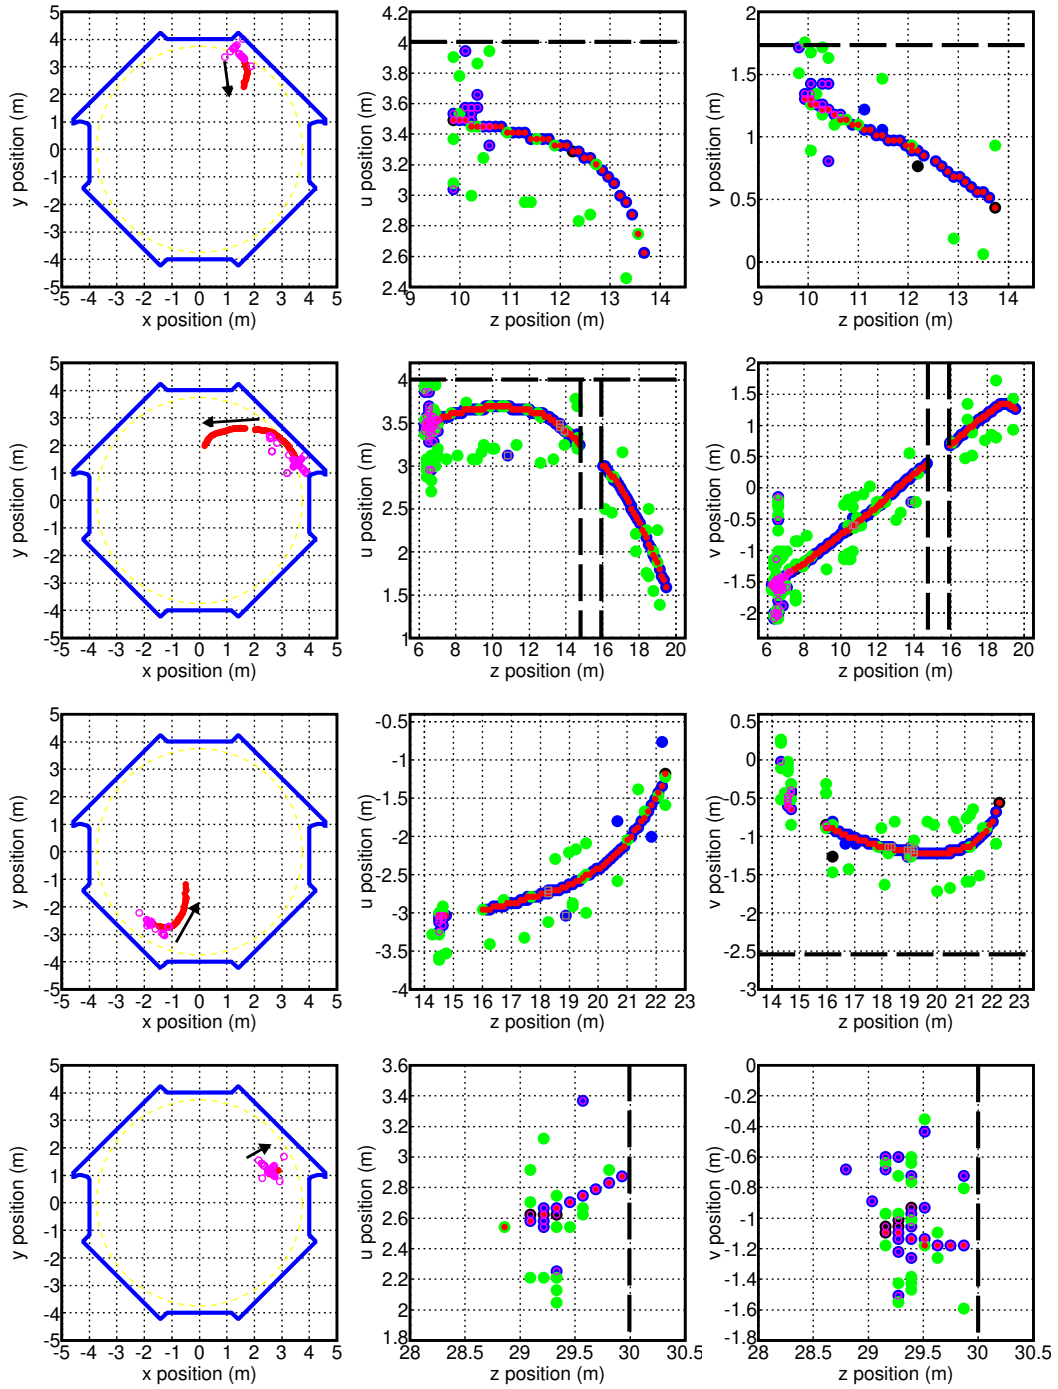


Figure 5.8: Monte Carlo anti-fiducial detector events. Top and second: edge events. Note that there is no ambiguity that they originate inside the detector. It is evident that part of the hadronic shower has probably been lost in each case. Third: event in the south side of the supermodule gap. Bottom: back event.

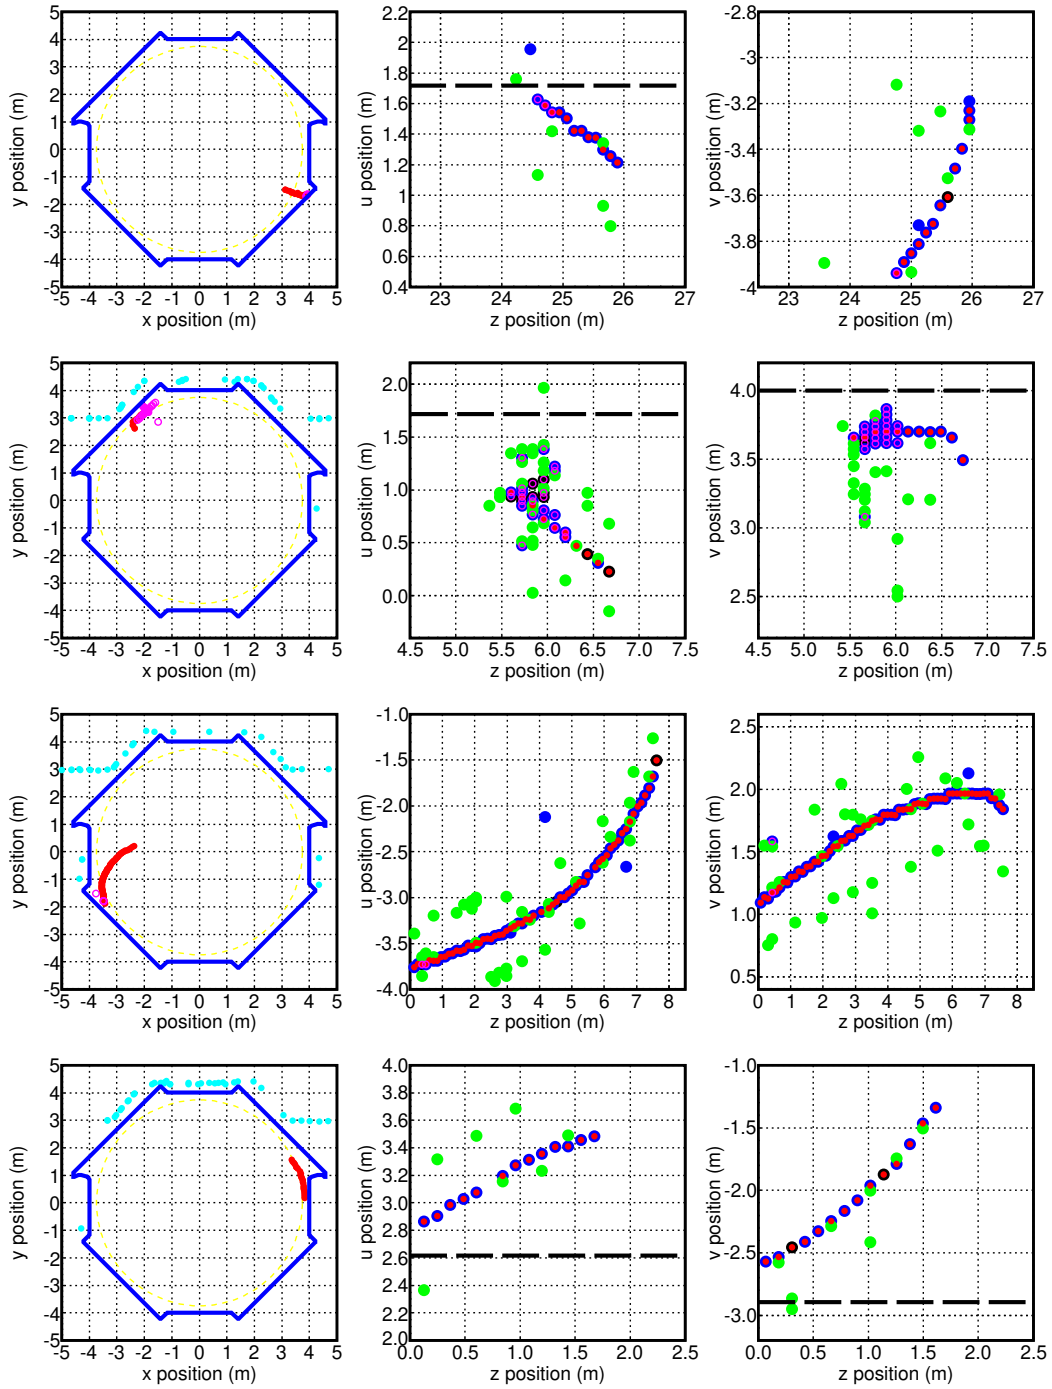


Figure 5.9: Typical selected data events. Top: rock-like edge event (run 39971, snarl 59065). The hits “outside” the detector in the uz view are either noise or activity caused by another particle at a different v position. Second: detector-like edge event (40076/145684). The shower here is evidently fully contained. Third and bottom: two front events which both happen to be near the edge (39993/228876, 39996/59590).

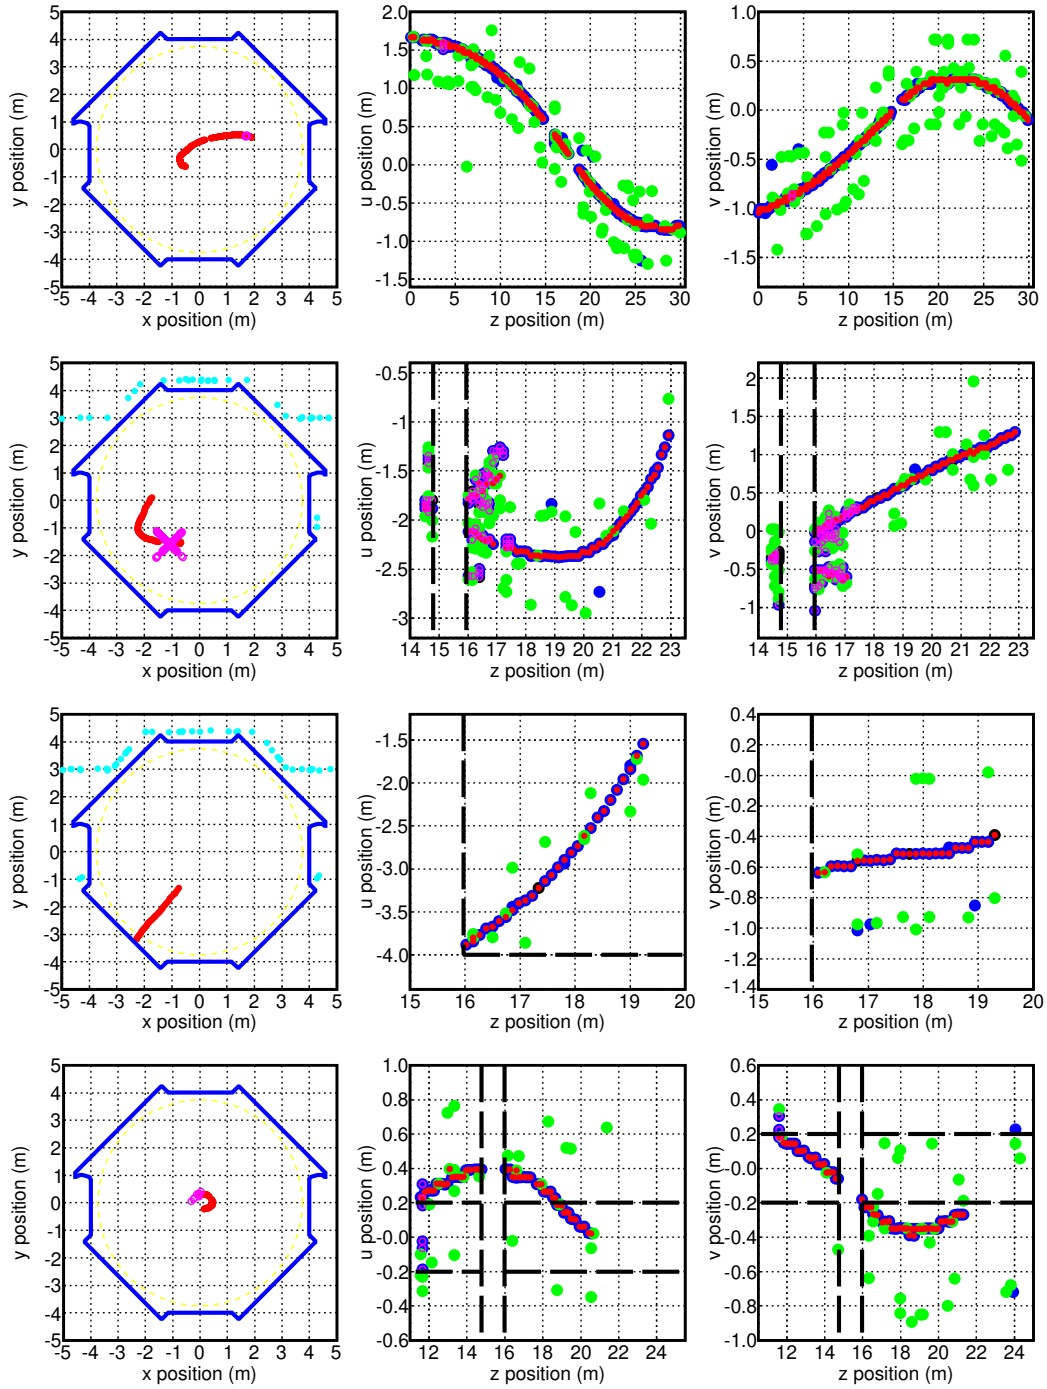


Figure 5.10: Less typical selected data events. Top: front face event traversing the entire detector length (run 39999, snarl 147692). Tracks like this are $\gtrsim 20$ GeV and so do not affect the fit. Second: supermodule gap south event (39971/194019). Note large shower and second track. Third: supermodule gap north edge event (39990/27787). This is almost certainly a rock event. Bottom: coil hole event (40488/68104). Evidently the muon reenters the coil hole at $z \approx 21$ m and perhaps exits again at $z \approx 24$ m immediately before stopping.

Chapter 6

Monte Carlo

This analysis uses Monte Carlo extensively in order to predict the spectrum of events at the far detector, to optimize event selection and analysis strategy, and to evaluate systematics. Due to the nature of the data to be analyzed, a limited number of data-driven cross checks of these results are possible, so it is important for the Monte Carlo to be as complete and correct as possible. The MINOS Monte Carlo consists of simulations of (1) the beam, (2) neutrino interactions, including propagation of initial interaction products through the nucleus as applicable, (3) propagation of interaction products through the bulk material of the rock and detector, and (4) the detector response to energy deposition in scintillator.

6.1 Beam Simulation

To produce the flux of neutrinos for each detector, a common beam simulation is used. The geometry of the target, horns, target hall and decay volume is modeled using GEANT4 [44, 45]. Interactions of primary 120 GeV protons with the NuMI target are handled by FLUKA [46, 47]. The FLUGG package [48, 49] integrates these two simulations so that FLUKA handles only interactions and GEANT4 handles all of the geometry. Secondaries coming off of the target are propagated through the target hall and decay pipe, where they may re-interact with material in this volume, or decay. For each decay producing a neutrino, the location of the decay and the identities of the child particles are saved. The

specifics of the decay kinematics are not generated at this time.

The production of hadrons from interactions of 120 GeV protons on carbon is not well measured. Therefore, it is not expected that the flux prediction is sufficiently accurate for direct use in the oscillation analysis. Rather, it is reweighted using near detector data before being compared to far detector data (see section 9.1). Although MINOS has gone through several generations of beam Monte Carlo, the information about the flux from the near detector has not been fed back into the hadron production simulation. Instead, we always start with the dead-reckoned simulation and apply the full correction after the fact. Effects of target degradation (see section 3.4.1) are not simulated directly, but folded into the reweighting based on the near detector data. Each of the 4 run periods (Runs I, IpHE, II and III) are separately simulated and then reweighted, with all time-varying effects averaged over the course of each run.

6.2 Neutrino Interactions

For each neutrino parent decay in the beam pipe, a probability is calculated using the decay kinematics that the neutrino reaches the near (far) detector. Proportional to this probability, neutrinos are chosen to be propagated through each detector with their energies and trajectories determined by the kinematics of the decay that produced them. The neutrino interactions in and around the detectors are simulated using NEUGEN [50], which uses the package INTRANUKE for pion rescattering and charge exchange within the parent nucleus. NEUGEN simulates 5 categories of events:

- Quasi-elastic: the axial mass M_A is set to 0.99 GeV by default.
- Resonance production: the axial mass M_A is set to 1.12 GeV by default.
- Deep inelastic scattering: Here there are 16 free parameters — the Koba-Nielsen-Olesen multiplicities — that set the fraction of the cross section which should be simulated as deep inelastic scattering. These are defined as r_{ijk} , where $i = 1$ (charged current), 2 (neutral current); $j = 1$ (νp), 2 (νn), 3 ($\bar{\nu} p$), 4 ($\bar{\nu} n$); and $k = 2, 3$

Parameter	Value(
112 — charged current, νp , multiplicity 2	0.1
122 — charged current, νn , multiplicity 2	0.3
132 — charged current, $\bar{\nu} p$, multiplicity 2	0.3
142 — charged current, $\bar{\nu} n$, multiplicity 2	0.1
212 — neutral current, νp , multiplicity 2	0.1
222 — neutral current, νn , multiplicity 2	0.3
232 — neutral current, $\bar{\nu} p$, multiplicity 2	0.3
242 — neutral current, $\bar{\nu} n$, multiplicity 2	0.1
113, 123, 133, 143 — charged current, multiplicity 3	1.0
213, 223, 233, 243 — neutral current, multiplicity 3	1.0

Table 6.1: NEUGEN parameters r_{ijk} . See text for details.

for the multiplicity of the final state [51]*. The values used for these parameters are shown in table 6.1.

- Coherent π^0 production. This is a rare process, 0.6% of events, which does not significantly affect this analysis.
- Inverse muon decay: $\nu_\mu e \rightarrow \mu \nu_e$. This accounts for 0.02% of events and does not significantly affect this analysis.

6.3 Particle Propagation

Particles leaving the nucleus are tracked through the cavern and detector geometry by GEANT4 (see chapter 7 for a detailed discussion of the rock model). Further interactions and decays are possible. In particular, charged particles typically produce large numbers of low energy bremsstrahlung photons and delta rays. Particles above a configurable cut-off energy are saved in a fixed size array of length 4000. If this array is exceeded, further particles are not saved, although they continue to be simulated. This can cause trouble when, at analysis time, one is trying to understand complex rock events. However, in most

*NEUGEN documentation does not explain these parameters. However, its successor, GENIE uses the same concept [52].




Figure 6.1: Fraction of light removed by the Monte Carlo for a hit near a strip end as compared to a hit in the bulk of the detector. The kink at 0.2m is a deliberate (although not very physically motivated) feature meant to quickly bring the value up to unity.

cases, the relevant aspects of the event are preserved. Energy deposition in scintillator is saved and later passed into MINOS-specific code that generates the detector response.

6.4 Detector Response

The simulation of the detector response to energy deposition in scintillator is handled by custom code within the MINOS framework. Energy deposition in scintillator is converted into blue photons, which are converted into green photons in the wavelength-shifting fibers, then into photoelectrons. Finally, the phototube response is converted into the simulated data stream. For each event, a set of calibration constants is chosen at random from spills during that run period (I, IpHE, II or III). This averages the gradually diminishing detector response over each run period.

Most interesting to this analysis is the handling of blue photons in the scintillator. The MINOS framework has a switch that allows the user to select whether to (1) explicitly produce and propagate individual photons or (2) use predetermined distributions based on the length of the path through the scintillator and the particle's dE/dx . In standard production running, which represents all samples used in this analysis, the second approach is always used since it is much faster. Here, the amount of light available for green photon production in fibers must be reduced near the ends of strips to represent the fact that some

blue photons are lost out of the (unmirrored) strip ends (see figure 6.1). This calculation takes into account only the average distance of the hit from the end of the strip. For particles passing through both very near a strip end and at a large angle with respect to z , this approximation breaks down. However, for beam muons, this is a very good approximation. To give a sense of the accuracy, a particle passing through a strip at a 45° angle such that it enters 1 cm from the end and exits 0 cm from the end has its photon count reduced as though it passed straight through at 0.5 cm. The reduction is underestimated in this case by 1.0%, which is still quite small even for this moderately extreme case.

6.4.1 Correction to Strip Lengths

While using cosmic rays to investigate strip alignment on the complex horizontal and vertical edges (see figure 5.5), I noticed a discrepancy between the apparent lengths of strips in data and Monte Carlo. In data, this length must be inferred by the extent of locations in a plane in which track vertices are reconstructed. Fortunately, the reconstruction code has no knowledge of strip lengths and so is not biased in this matter. It simply uses the transverse positions of hits in each view — i.e. those that are determined by which strip is hit — to find a three dimensional track. (The downside to this strategy is that if a track is low quality, perhaps overlapping other detector activity or even reconstructed out of several unrelated particles, a vertex can be reconstructed well outside the detector.)

With $\sim 70\text{M}$ cosmic events recorded at the time of this study, each centimeter of each strip has been crossed by about 60 tracks. This allows the ends to be mapped well enough to notice data/Monte Carlo discrepancies on the order of a centimeter. The discrepancy was approximately a strip and so was easily visible. Ultimately, the mine crew was asked to autopsy a spare module. They used a tape measure to check the strip lengths and confirmed the effect. The main cause of the mismodeling is that the longest strip of the B modules had been assumed to be a full 8 m, with strip lengths beginning to decrease starting with the next strip. In fact, the longest B strip is already shorter than full length. The autopsy also revealed the idiosyncratic step at the A/B boundary, which is shorter than the steps within each module.

As the tape measurements were much more accurate than the cosmic ray study, these

were used to correct the Monte Carlo. For technical reasons, the model of the detector through which particles are transported and their energy deposited was not changed. Rather, at the step in which energy depositions are converted into scintillation light, a special check is made to prevent the simulated scintillator that should not have been there from generating any light. (This trick works because the scintillator strips were always too long and never too short.)

A very slight mismodeling persists in that the non-existent scintillator material still removes energy from particles that pass through it. However, in comparison with other approximations in use in the Monte Carlo, this is completely negligible. For instance, none of the fiber manifolds or electronics present in the same general area are simulated at all. All concerns about material near the edge of the detector are ameliorated by noting that their presence or absence is approximately equivalent to moving the cavern wall inward or outward slightly. As will be shown in section 7.3, this has a very small effect on the data.

This correction removed 0.47 m^2 of scintillator from each plane in the Monte Carlo model, or 0.90%. For rock events entering the front face, this corrects the prediction of the event rate in direct proportion by decreasing the effective cross section that the detector presents to the beam. The surface area of the edge of the detector is also reduced by 0.83%, which again directly impacts the number of rock events in this region. The biggest impact is in the number of in-detector edge events. In the edge, there is 8.32 m^2 of scintillator per plane and so this reduction is fractionally quite substantial: 5.3%. Note that the fact that the steel planes have not changed size is irrelevant at first order. As will be explained in section 7.3, any detector steel encountered by neutrinos and/or muons prior to their first scintillator crossing acts in the same fashion as the rock and has no net effect on the observed muon flux.

From the direct measurement of the A/B strip lengths and clear documentation about the specifications of the lengths of the 8 m strips (required to be 7.996–8.000 m long), we know all of the strip lengths with an error of roughly 2 mm. This gives an error on the above statements about event rates in various regions. However, this error is quite small. Even for the most strongly impacted event category, in-detector edge events, it is only $\pm 0.4\%$. As will be discussed later, this is negligible compared to other normalization systematics.

Chapter 7

Rock Model

An accurate model of the rock surrounding the far detector is critical for this analysis. Relevant features of this Monte Carlo include the shape and size of the cavern, the density and composition of the rock, the extent of the rock volume simulated, and the details of any special approximations used in modeling rock interactions. Since the number of data-driven cross checks that can be performed to test the rock model are limited, a variety of Monte Carlo studies have been performed to bolster confidence in the model.

7.1 Cavern Dimensions

The far detector cavern was designed and excavated for MINOS. It is a rough rectangular box with the detector near the south end. The long axis of the cavern is aligned with that of the detector; both point at Fermilab (but are level, not pointed down at the beam angle). The ceiling of the cavern is curved for mechanical strength. Neither walls nor ceiling were required to be flat on a small scale, but rather only to form a cavern at least as big as the collaboration contracted for, i.e. no rock could remain inside the neatline. They are formed from the rock remaining after blasting, stabilized with a layer of shotcrete (sprayed-on concrete). This gives each an irregular egg carton-like shape with variations at the level of 30 cm.

Since the Monte Carlo model does not include the shotcrete, but simply uses rock

material all the way up to the wall, it is interesting to know how thick the shotcrete is in order to determine the impact of this approximation. Estimates of its thickness range from 3 to 15 cm. Similarly, the floor is poured concrete. It was required to have a minimum depth everywhere of 10 cm. A depth close to this is typical, but in some regions it is closer to 1 m deep [53].

The design height of the cavern is 11.3 m at the east and west walls and 13.6 m at the center. The distance from the southernmost steel plane of the detector to the south wall of the cavern was designed to be 13.6 m (see figure 7.1). After construction, the mine crew measured it to be 13.6 ± 0.3 m, with accuracy limited by the egg carton shape. The width of the cavern gradually increases from south to north, with the detector centered east-west on average (see figure 7.2). At the south end of the detector, I measured width to be 14.3 ± 0.4 m.

7.2 Rock Composition and Density

The rock surrounding MINOS is known as Ely greenstone. The elemental composition is known from measurements elsewhere in the region, see table 7.1.

To precisely determine the density, I traveled to Hibbing, MN in August 2008 where the core drilled prior to excavation of the MINOS hall is stored. The core was drilled by CNA Consulting Engineers in September and October of 1995. It is 7.3 cm in diameter and 142 m long. It begins to the north of the MINOS hall in the preexisting tunnel that connects the old mining tunnel to the Soudan 2 hall (see figure 7.3). At the north end, it is angled 8° upwards; it curves downwards to approximately level by the end. The core hole is visible in the cavern today about $2/3$ of the way up the center of the south wall. The extent of the core beyond this point has some uncertainty, but the best estimate is 32 m. The core was extracted in sections that range from very short pieces to ones ~ 30 cm long. See figure 7.4 for a representative box.

I measured the density by weighing pieces of core and then finding their volume through water displacement. I weighed them on a triple beam balance which can be read to the nearest 0.01 g, but which has a temperature dependence which limits overall accuracy to




Figure 7.1: MINOS far detector cavern cross section, as designed. The exact shape of the walls and ceiling are schematic only.




Figure 7.2: MINOS far detector cavern top view, as measured. South is to the right. The thin rectangles indicate the southernmost planes of each supermodule. While precise corners are shown to indicate the boundaries of each width region, in reality only a gradual widening is visible. Measurements are accurate only to the nearest 30 cm or so due to the wall shape.

Element	Greenstone (%)	Monte Carlo Concrete (%)
O	45.86±0.53	53.0
Si	23.7±0.16	33.0
Fe	8.56±0.35	2.0
Al	7.97±0.21	4.2
Ca	6.41±0.30	6.3
Mg	3.95±0.23	—
Na	1.90±0.12	1.5
Ti	0.64±0.05	—
K	0.34±0.09	—
H	0.26±0.04	—
C	0.20±0.11	—
Mn	0.15±0.01	—
P	0.06±0.01	—

Table 7.1: Measured composition of Ely greenstone. The central values are used in the Monte Carlo. Errors are derived from the spread of values over 15 samples recorded in references [54, 55]. For comparison, the Monte Carlo composition of concrete is given. The Z/A of concrete differs from greenstone's by less than the error from greenstone's composition.

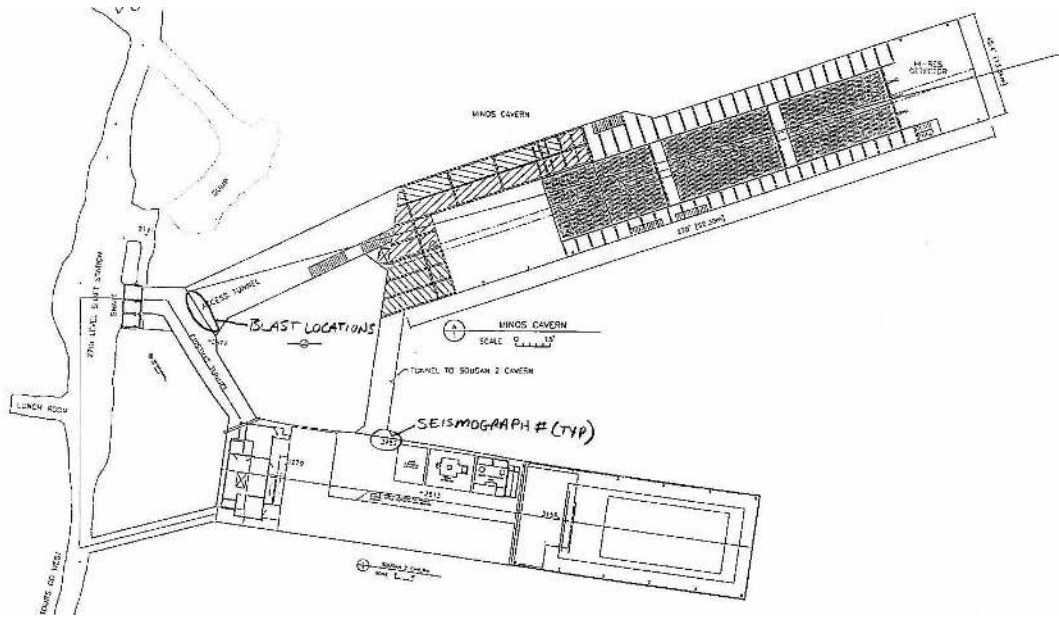


Figure 7.3: Plan view of the MINOS and Soudan 2 halls. (Three MINOS super-modules are schematically shown since this diagram predates construction.) The core is the line from the wall marked “blast locations” through and out the south end of the MINOS hall.



Figure 7.4: Photo of MINOS core sample box 51, ~18 m south of the south wall. CNA logged this rock as metabasalt.




Figure 7.5: Raw density data. The south end of the detector is at +14 m. Fitting to a constant (solid line) gives a good fit. The best linear fit (dotted line) is consistent with flat.

about 1 part in 500.* Since this error is small compared to that from the volume measurement, I did not log the temperature. To find the volume of a sample, I filled a 1000 mL graduated cylinder partially with water, immersed the sample and recorded the difference in water level. (All samples are odd shapes, so it is impossible to simply calculate the volume.) The cylinder is marked in 10 mL increments, but these are 3.5 mm apart, so it is possible to read it to the nearest mL or so.

Table 7.2 gives the raw data. Locations are estimated from sample placement inside the boxes and are correct to the nearest 30 cm. Measurements stop a few meters south of the detector due to lack of time to examine further samples. There is no reason to believe from the CNA log that the remaining relevant rock would be significantly different. Most of the rock between the end of my measurements and the north end of the detector is the same as that at the south end (metabasalt and meta-andesite).

The density data is consistent with the density being constant over the length of the measured core; see figure 7.5. The best fit is $2.769 \pm 0.007 \text{ g/cm}^3$. The rock *in situ* has small air gaps, estimated to be $0.3 \pm 0.3\%$ of the volume [56]. Taking this into account yields the final result of $2.757 \pm 0.012 \text{ g/cm}^3$. The central value is used in the Monte Carlo.

*Based on my extensive study of the same balance for another purpose. The manufacturer simply quotes “0.1 g readability”, referring to the tick marks.

7.3 Independence of Flux and Density

As a good approximation, the number of muons reaching the detector from beam events in the rock is invariant with respect to the density of the rock, the shape of the detector hall, and the presence of objects in the hall. More generally, none of these should have an effect on the flux or spectrum of any particle reaching any point on the detector. The fundamental reason is that both the number of neutrino interactions and the stopping power of a material for any particle are linear with density [57].

This argument has several caveats. It assumes that the flux of neutrinos is uniform transverse to the beam direction, and that there is negligible multiple Coulomb scattering and particle decay. It also assumes that all materials have the same neutrino cross sections per unit mass. Finally, it only applies to the total flux of each type of particle; the character of individual multi-particle events can vary. These are discussed in detail below.

The reasoning is as follows. Consider a very wide monodirectional neutrino beam traversing material with arbitrary density variations. The interactions of the neutrinos produce secondary particles with a variety of energies and momenta. Consider a point in this material and a particular particle type produced somewhere upstream with a given initial momentum such that it passes through this point before stopping. To find the flux of these particles at the chosen point, trace back along the direction of propagation until enough material has been encountered to stop such a particle. This gives the amount of material available for production. The number of particles produced along this length is proportional to the integrated density while the length itself is inversely proportional to the integrated density, so there is no dependence on density profile.

Now it is established that the flux of particles of each particular species and *initial* momentum reaching some point is not dependent on the density profile. These particles will reach the point with a spread of *final* momenta which is distributed in a fixed way between zero and the initial momentum (the distribution is roughly uniform for muons in the context of MINOS, but this is not necessary). Thus the flux of particles and particle momenta at the point is also not dependent on the density profile.

As an example, consider a point on the face of the MINOS far detector. If we wish to

know the flux of muons arriving here with $\vec{p} = (0.0, 1.0, 5.9)$ GeV, i.e. with an energy of 6 GeV and pointing 9° upwards, it does not matter, within the range of applicability of this argument, whether we have modeled the cavern precisely enough to know whether these muons come out of the south wall or out of the floor, nor does it matter what the density of the surrounding rock is, nor whether the density of the concrete floor is the same as that of the rock beneath, nor if we have correctly modeled objects in the hall between the rock and the detector.

7.3.1 Caveat: Radial Dependence

This argument does not apply to the near detector, since there is a strong radial dependence on neutrino flux. This severely limits the usefulness of near detector rock events. At the far detector, the flux is constant to about 1 part in 200,000 in the region that can produce events causing detector activity [58].

7.3.2 Caveat: Nuclear Composition

If the nuclear composition is changed, this alters the fluxes. As a reasonable approximation, the number of events is proportional to A , while the stopping power is proportional to Z . Therefore, the rock's Z/A has a roughly linear effect on the flux of secondaries. Ely greenstone's Z/A is known to be 0.498 ± 0.004 from geological samples (see section 12.2.2), and so the ratio of production to stopping power is known, at best, to 0.8%.

A secondary effect of differing compositions is to modify the cross sections even at constant density and Z/A . The mechanism of most MINOS events, particularly the higher energy events from the rock, is deep inelastic scattering, in which the overall composition of the nucleus is not important. Cross sections for quasielastic and resonance events are expected to depend significantly on isotopic composition rather than simply scaling with the number of protons and neutrons. This effect is unmeasured for the major elements in the rock (see table 7.1). The MINOS Monte Carlo uses the relativistic Fermi gas model — considered obsolete by theorists — which predicts only tiny variations from simple scaling by neutron and proton number. Modern models predict larger variations. About 21% of rock events are affected by this. Since it affects non-deep inelastic events only, it is energy

dependent and could partially mask the oscillation signal. See section 12.2.1 for further discussion.

7.3.3 Caveat: Density and Decay

If the secondary particles are allowed to decay, this effectively adds the same constant to the stopping power of any medium irrespective of density, which removes the proportionality between density and stopping power. However, at MINOS very few muons decay in flight so this effect is very small. Even a 0.5 GeV muon, which is at the lower limit of detection, has $\gamma c\tau = 3$ km, whereas 90% of the muons arriving at MINOS below 10 GeV have travelled less than 50 m. According to the Monte Carlo, the average probability for a rock muon that would otherwise reach the detector to decay is 7×10^{-4} for neutrinos interactions below 10 GeV and 4×10^{-4} overall. However, pions and kaons produced in the rock will be more strongly affected as their decay-in-flight probability is typically 1–10%. The effect is to reduce the flux for lower densities since the secondaries must travel farther on average before reaching the detector. Since this analysis selects only muons, only the backgrounds are affected.

7.3.4 Caveat: Event Shape

The argument says nothing about the shape of individual events. A cavern model with the walls closer to the detector will produce multi-particle rock events that are more compact as they enter the detector than one where the walls are further away. The MINOS Monte Carlo's cavern is narrower than the actual cavern along the entire length of the detector. However, it also does not model the various objects in the hall between the walls and the detector. These effects should tend to cancel. In any case, this analysis does not analyze event shape, so it is only relevant insofar as it complicates muon tracking.

7.4 Monte Carlo Density Study

To verify the smallness of the effects listed above, I used standard unoscillated MINOS Monte Carlo to generate rock events at the far detector for each of three realistic densities:

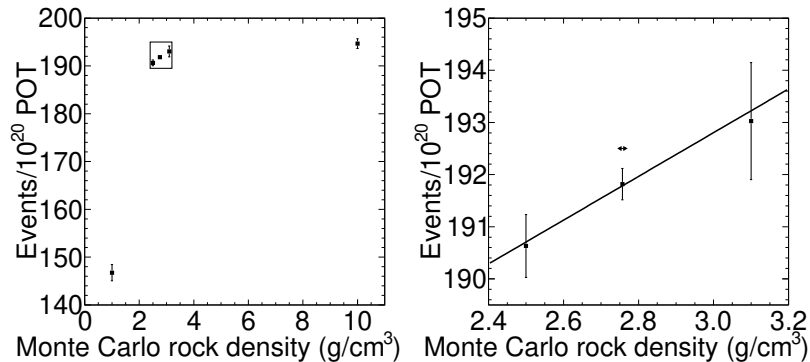


Figure 7.6: Number of Monte Carlo charged current events in rock with a reconstructed track in the far detector as a function of rock density. Left: three realistic densities and two extreme densities (*note suppressed zero*). Right: reasonable densities only with a linear fit. The double headed arrow above the central point shows the uncertainty on the rock density, 0.012 g/cm^3 . The lengths of the error bars reflect the sizes of the samples generated.

2.5 g/cm^3 , 2.757 g/cm^3 (the measured value), and 3.1 g/cm^3 , as well as two extreme densities: 1 g/cm^3 and 10 g/cm^3 . By looking at the distribution of true neutrino interaction vertices, I verified for each density that the amount of rock volume modeled was sufficient to include all events that could survive to reach the detector.

Within the realistic densities, I observed a weak dependence between the rock density and the number of rock charged current events resulting in a reconstructed track. For each g/cm^3 increase in density, the number of events increased by 4.2 ± 2.0 per 10^{20} protons-on-target (see figure 7.6). While this is only a 2.1σ separation from zero, it is in the expected direction. The two extreme cases show the trend continuing (non-linearly, since clearly at zero density the number of events would be zero). The uncertainty from this effect is 0.4 ± 0.2 events. In comparison, the uncertainty due to Z/A is about 40 times larger.

7.5 Model of Cavern Geometry

In the MINOS Monte Carlo, the far detector hall is modeled as a perfect rectangular box. Transverse to the beam direction (i.e. roughly east-west) the width of the hall is 14 m; the detector is perfectly centered. While the real hall widens towards the north, the Monte Carlo hall does not. The hall is 12.2 m tall, with the detector center 4.73456 m above the floor, putting the bottom of the detector's steel planes 0.73456 m above the floor. The hall

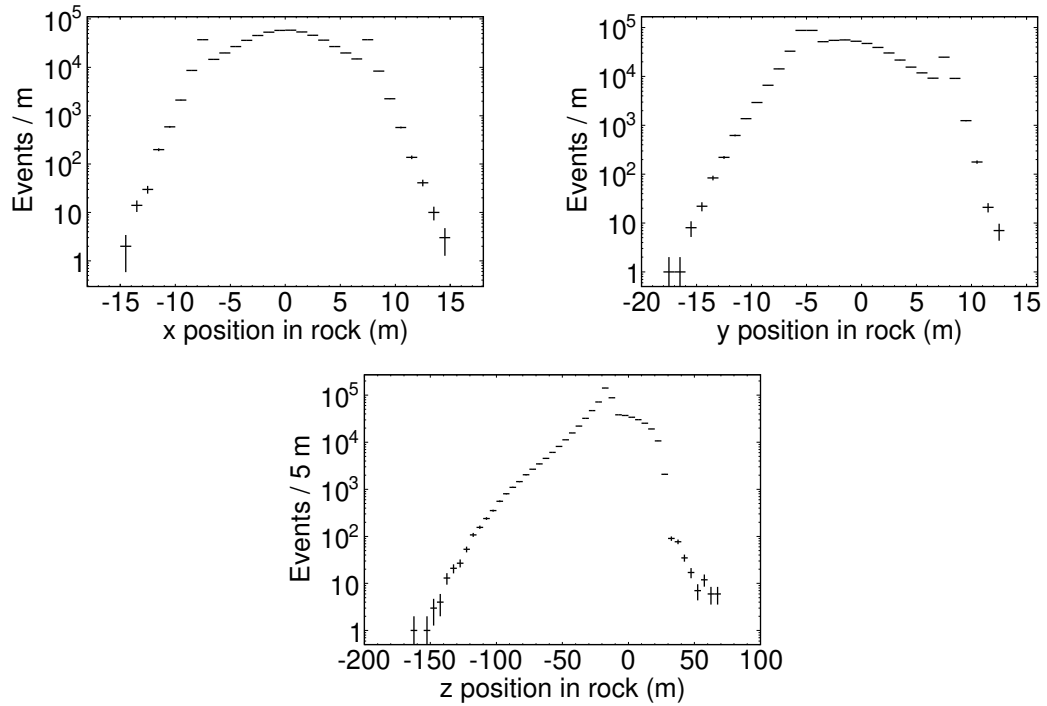


Figure 7.7: Monte Carlo rock events resulting in detector activity as a function of the x , y , and z position of the neutrino interaction. The “ears” at ± 7 m in x are where the rock begins, and similarly in y . The z distribution peaks at the south wall (-14 m) and falls off sharply past the north end of the detector. The simulation ends at $+70$ m in z . These plots include all rock Monte Carlo used in the analysis: $\sim 7 \times 10^5$ events. It is evident that no more than ~ 10 additional events would have been generated with an infinite volume.

extends 14.022 m south of the front of the detector's first steel plane and 38.3 m past the north end of the detector.

Enough rock is simulated around the detector hall to ensure that no significant number of interactions that could result in detector activity are lost (see figure 7.7). This volume had initially been tuned to allow for fast generation of rock events interesting only as a small background to fiducial analyses. Studies done for the Run I version of this analysis [59] motivated the use of a larger volume for use in analyzing rock events. Given the higher statistics now available, I again extended the volume to ensure that there could be no concern about the effects of the Monte Carlo boundary.

Quite a lot of time is spent simulating rock events that never make it out of the rock. However, this is mitigated by an optimization in the Monte Carlo which attempts to detect when a particle cannot possibly reach the detector without a full simulation. It does this by considering whether the particle would get there if it were aimed directly for the detector and lost energy at the lowest possible rate (minimum ionization) the whole time. To account for fluctuations in energy loss, an additional safety factor is added on top of this, so that its energy must fall short of what is calculated above by some tunable fraction. If it fails this test, it is killed immediately, saving time. I studied this optimization and found that it incorrectly prevents only about 1 particle per 10^{23} POT from reaching the detector, or ~ 0.006 particles for the Run I–III exposure. About 60% of the run time of the Monte Carlo is saved via this technique. It is also notable that no matter what the Monte Carlo strategy at the far detector, the time necessary to generate near detector Monte Carlo always dominates due to the much higher event rate there.

The concrete liner and floor of the cavern are simulated as rock. The concrete is sufficiently thin and its composition sufficiently similar to the rock to justify this approximation (see table 7.1). Likewise, objects in the hall between the detector and the rock are not simulated on the assumption that they represent a small column of material to incoming particles and that their Z/A is similar to the rock's (see section 7.3.2).

In one case, this approximation appeared to be strained. While measuring the cavern, I discovered that there are two $\sim 1 \text{ m}^3$ stacks of lead bricks upstream of the far detector on the floor of the cavern. Lead's Z/A is 0.40, only 79% of the rock's. Therefore lead will

produce a larger flux of muons reaching the detector. However, as they only intersect about 4% of the beam and are only 1 m long in the beam direction, the enhancement of the muon flux is $\sim 0.04\%$ (somewhat less than one event), much smaller than the systematic error on the flux normalization. Therefore, no effort was made to include the effect in this analysis. The only other common elements with $Z/A < 0.45$ are tungsten and tin; neither of these is present in large quantities in the detector hall.

Position (m)	Mass (g)	Volume (mL)	Density (g/mL)	Type
-32.0	148.3	55	2.71±0.05	metabasalt
-32.0	151.6	56	2.72±0.05	metabasalt
-29.0	173.7	62	2.79±0.04	metabasalt
-29.6	150.4	55	2.75±0.05	metabasalt
-26.5	132.8	49	2.71±0.06	metabasalt
-28.0	169.5	63	2.70±0.04	metabasalt
-21.9	604.1	216	2.800±0.013	metabasalt
-22.9	362.0	130	2.79±0.02	metabasalt
-21.6	208.1	76	2.75±0.04	metabasalt
-17.4	259.6	95	2.73±0.03	meta-andesite
-16.5	597.9	215	2.78±0.02	meta-andesite
-14.0	539.3	195	2.77±0.03	meta-andesite
-11.9	202.8	71	2.84±0.04	meta-andesite
-9.4	381.6	140	2.73±0.02	meta-andesite
-6.7	460.3	169	2.73±0.02	meta-andesite
-3.7	187.1	67	2.79±0.04	meta-andesite
-1.2	170.2	61	2.79±0.05	silicic tuff
+1.8	204.3	71	2.87±0.04	silicic tuff
+2.4	151.4	54	2.80±0.05	silicic tuff
+6.1	202.3	74	2.75±0.04	mylonite
+7.6	273.7	101	2.71±0.03	mylonite
+9.8	164.1	59	2.81±0.05	metabasalt

Table 7.2: Density data. Location is relative to the south wall, with north being positive. Relative locations are accurate to a few tens of centimeters, but the absolute position is less certain. Rock types are as logged by CNA. Each rock type is a subset of Ely greenstone.

Chapter 8

Selection Cuts

The goal of this analysis is to analyze beam muon (anti)neutrino charged current interactions. To do this, we must reject cosmic ray events and beam backgrounds, primarily from neutral current interactions. The distinguishing characteristic of charged current ν_μ and $\bar{\nu}_\mu$ events is the production of a muon, and so the selection criteria focus on identification of muon tracks.

8.1 Basics

Beam spills are only considered if both the beam and the detector are operating normally. This analysis uses the same criteria to determine beam and detector status as previous MINOS analyses [36]. Selected events must have a reconstructed track. The longest such track must have its vertex outside the fiducial volume (see section 5.1). This automatically ensures that events selected by this analysis are fully independent of the fiducial charged current analysis.

No charge cut is applied to tracks; this analysis accepts both positive and negative muons and fits them together under the assumption of CPT conservation and lack of non-standard interactions. This is motivated primarily by the desire for increased statistics: The Monte Carlo predicts that 9.3% of selected events are produced by charged current interactions of the $\bar{\nu}_\mu$ component of the beam. In addition, since the anti-fiducial muon spectrum peaks

at the very bottom of MINOS's resolution, many selected tracks are not long enough for good charge sign determination. Requiring, for example, a 2σ sign determination would cut 9.6% of otherwise accepted events. Finally, by neither cutting nor binning events based on charge, there is no systematic error associated with charge determination as is present for the 2010 fiducial analysis [43].

All cuts were optimized for the LE beam; the same set of cuts was used for both the LE and pHE data. In principle, the pHE data would benefit from separately optimized cuts, but the potential gain was judged to be far too small to be worth the complication. The figure of merit for a set of cuts was the error on Δm^2 if the true values of the mixing parameters were $\Delta m^2 = 2.46 \times 10^{-3} \text{ eV}^2$ and $\sin^2 2\theta = 0.98$. This is a point near the center of both the MINOS and Super-K contours. A point away from maximal mixing was deliberately chosen to avoid any complications arising near the physical boundary. The motivation to optimize only Δm^2 is that this analysis has a much better sensitivity for it than for the mixing angle. The error on Δm^2 for a set of cuts was found by running the full fitting procedure on Monte Carlo. Systematic errors were not included at this stage. However, cuts were adjusted away from their mathematically optimal position towards safer values if they appeared likely to cause undue exposure to systematics; see, for instance, section 8.3.

8.2 Cosmic Rejection

The cosmic ray rate at the MINOS far detector is 0.5 Hz. Since this analysis explicitly includes events appearing to enter the top and side faces of the detector, it has a larger exposure to cosmic rays than a fiducial analysis in which the exterior of the detector acts as a shield. Nor is the veto shield useful for cosmic rejection, since events from outside the detector are part of the signal. However, as in a fiducial analysis, the beam timing is a powerful discriminant against cosmics. The same timing cut is used as in the MINOS fiducial analysis: events are accepted in which the first track hit occurs between -2 and $+12 \mu\text{s}$ after the beginning of the neutrino pulse reaches the far detector. This includes the $10 \mu\text{s}$ length of the proton spill plus a $2 \mu\text{s}$ buffer on either side. See figure 8.1. Since all relevant muons are highly relativistic, no modification of this cut is needed to include




Figure 8.1: Selected events in LE running as a function of time since the beginning of the beam spill. The timing cut is shown as dashed lines.

events even very deep in the rock — these muons reach the detector at the same time as the neutrinos do. Since the beam is pulsed with a period of 2.2 s, this timing cut removes all but 6.4×10^{-6} of cosmics.

Within the spill itself, spatial information is used for further cosmic rejection. The MINOS fiducial analysis requires that tracks be within a particular angle of the beam direction. This analysis uses a different strategy, driven by both the desirability to accept large-angle muons from the rock, since they are the ones more likely to come from low energy neutrino interactions, and the additional susceptibility to cosmic background due to accepting events around exterior of the detector. The rules are as follows:

1. The track is accepted if $p_T/\text{GeV} < 20(\cos\theta - 0.5)$, where p_T is the transverse momentum and θ is the angle between the track and the beam (see figure 8.2).
2. If the transverse momentum of the track is less than 1 GeV, it is accepted.
3. If the track vertex is in the bottom three edges, away from the front, supermodule gap, or back *and* the track end is in the fiducial volume, it is accepted regardless of angle and transverse momentum.
4. Any track not satisfying at least one of these conditions is rejected.




Figure 8.2: Two of the three spatial cosmic rejection cuts. Events above and to the left of the blue dashed line are rejected. Red crosses are cosmic data; black dots are LE spill Monte Carlo. Each is shown at an exposure of 10 times Runs I–III.

The first rule accepts the bulk of the signal. It is similar to that used by the Run I anti-fiducial analysis [59], but modified to accept high energy, low angle events as well.* The rule makes no distinction between upward and downwards-going tracks because all tracks in the spill window are assumed by the reconstruction code to be travelling northwards. Since the timing resolution of the detector is not very good and most of these tracks *are* going northwards, this was found to yield better results than determining track direction from timing. This means that a cosmic going south and down is reconstructed as going north and up. Therefore to reject cosmics, we cannot do the seemingly obvious thing of only cutting out downwards-going tracks.

The second rule accepts the tail of low energy, high angle beam events that carry a relatively large amount of oscillation signal per event. The third rule accepts events very unlikely to be cosmics that are not accepted by the first two rules.

To evaluate the cosmic background accepted by these cuts, I used a large sample of far detector cosmic data. This data is recorded without the requirement that events be within

*While accepting more high energy events is not very helpful for determining oscillation parameters, it is beneficial for discriminating between models of neutrino disappearance. In particular, the neutrino decay model predicts a more gradual disappearance fall-off at high energy than oscillation does. This is a topic of the combined fiducial+anti-fiducial analysis [28, 43].

a beam spill window. Scaling the accepted events in this sample down to the live time of beam spills ($14\ \mu\text{s}$ for 28 917 009 spills), 1.6 events were accepted by the first rule alone, 1.1 by the second rule alone, 0.5 by both the first and second rules, and 0.04 by the third rule alone, for a total of 3.3.

Use of out-of-spill cosmic data gave a high statistics result, but could differ from beam data because (1) out-of-spill events are processed through a somewhat different reconstruction algorithm, not assuming that all tracks are northwards and (2) events are only written out if triggered by detector activity, which could introduce a bias. To compensate for the first effect, I swapped all variables referring to track vertices and ends for southwards tracks to make them look like they were spill-reconstructed. This is not perfect, since the track momentum from curvature depends on direction and I did not recalculate it. The effect from triggering bias is almost certainly negligible given the high efficiency of triggering on muon tracks.

To simultaneously check both effects, I looked at Run I-III spill data with an anti-selection on the spill itself. The spill data consists of $100\ \mu\text{s}$ *spill windows* in which all detector activity is written out. Of this, the spill itself is only $\sim 10\ \mu\text{s}$. Within each spill window, the beginning of the spill is time zero and occurs $40\ \mu\text{s}$ after the beginning of the window (see figure 8.1). I selected events with time either between -39 and $-4\ \mu\text{s}$ or between $+14$ and $+59\ \mu\text{s}$. This avoids the first and last $1\ \mu\text{s}$ in case there are edge effects and has a $4\ \mu\text{s}$ buffer around the spill, including a $2\ \mu\text{s}$ buffer around the usual timing cut. There were 19 accepted events. Taking the ratio between the size of this time window and that of the timing cut used in the analysis, this gives 3.3 ± 0.8 , consistent with the high statistics cosmic data.

These cuts reject 0.75% of beam Monte Carlo events, predominantly either below 4 GeV or exiting tracks too straight for the curvature to be measured. The median energy of the remaining cosmic background is 1.5 GeV. Despite peaking near the most important part of the spectrum, it is so small as to have a negligible effect. This number of events is 0.22% of the expected signal, much smaller than the 1.54% normalization systematic error (see section 12.1).

		True μ^\pm			True other
			μ^-	μ^+	
True ν_μ or $\bar{\nu}_\mu$ CC	1727	ν_μ	1560	7.1	68
		$\bar{\nu}_\mu$	0.5	160	
True NC or other CC		6.2			41

Table 8.1: Monte Carlo prediction of selected rock and anti-fiducial events in Runs I-III, LE beam only, oscillated using the best fit point. Note that identification of truth track particle is somewhat complicated by the Monte Carlo's limitation of storing only 4000 particles per event. In complex rock interactions, the truth particle can be misidentified as the last-stored ancestor of that particle.

8.3 Particle Identification

An event is accepted if it passes one or both of two PIDs (particle identifications) developed to distinguish muon tracks from non-muon tracks. Each of these uses the *k-nearest-neighbors* (kNN) algorithm [60]. In this algorithm, the candidate event is compared to a corpus of Monte Carlo events that includes both signal and background. In an N -dimensional space of reconstructed quantities, the nearest k Monte Carlo events to the candidate event are found, where k is some number tuned to produce good results. The output of the PID is the fraction of these Monte Carlo events that are signal.

The first PID (*2008kNN*) was developed by Rustem Ospanov for use in the 2008 MINOS analysis [61]. It identifies ν_μ charged current events by determining whether the longest track in the event is a muon. No other reconstructed objects are considered. (Previous MINOS analyses have used shower characteristics as part of their method for distinguishing charged current and neutral current events [36], but studies have shown that focusing on the track yield superior results.) It uses a 4 variables: the track length, the mean pulse height of hits in the track, the spread of pulse heights in the track, and the transverse profile of the track (muons typically hit one scintillator strip per plane, but hadrons often cause a wider pattern of energy deposition). It uses the 80 nearest neighbors and so outputs a value in the range 0–1 (inclusive), quantized in units of 1/80, where 1 is most muon-like. If an event has no track, the PID will not consider the event at all and outputs a negative number. Similarly, if the track is too short — defined as having hits in fewer than 5 planes in either the U or V view — it is rejected entirely.

The second PID (*2010kNN*) was developed by Jasmine Ratchford for use in the 2010 MINOS analysis [62]. It similarly operates by attempting to identify the longest track as muon-like or non-muon-like. It is optimized for accepting the very short tracks that 2008kNN does not attempt to classify, but also returns an answer for all longer tracks. It will consider a track as long as it hits at least 5 planes in total and at least 2 planes in each view. (In order to be reconstructed at all, a track must have at least 3 planes in total and at least 1 plane in each view.) It uses 4 variables: the length of the track, the sum of the pulse heights in the last 5 planes, the amount of scatter of the track in the U view and the same in the V view. 2010kNN also uses 80 nearest neighbors and therefore has the same output format as 2008kNN.

In the absence of systematic error and assuming that the Monte Carlo is a perfect model, the sensitivity to Δm^2 is maximized if all events with tracks are accepted, i.e. if no PID cut is made. However, clearly neither of these assumptions is correct, so instead I have looked for the optimum non-null cuts. Different optimum values of the PIDs are found depending on whether $\sin^2 2\theta$ or Δm^2 is optimized. Since Δm^2 depends mainly on the count of events, it benefits from a higher efficiency and does not suffer very much as background is admitted. On the other hand, $\sin^2 2\theta$ is dependent on energy resolution since it is a measure of the depth of the oscillation deficit. It benefits from accepting fewer events with higher purity so as to avoid both non- $\nu_\mu/\bar{\nu}_\mu$ -CC events and $\nu_\mu/\bar{\nu}_\mu$ -CC events where the primary track is not a true muon. Since the MINOS reconstruction assumes all tracks are muons, it dramatically underestimates the energy of pion and proton tracks; these form a low energy background obscuring the oscillation maximum.

For each of three categories of events — rock-like, detector-like, and back events — a different set of PID cuts is used. This optimizes the selection given the different background levels in each category.

- **Rock-like events:** This category consists of events on the front face and those around the radial edge that are selected as rock-like. The rock is an efficient filter against non-muon particles, so a much higher fraction of tracks arising from true rock interactions are muons as compared to tracks arising from true detector interactions. Therefore it is beneficial to use loose PID cuts here. The sensitivity to Δm^2 is optimized by




Figure 8.3: 2010kNN for non-back (left) and back (right) events. The solid black histograms are true ν_μ charged current events; the dashed red histograms are all other interaction types, predominantly neutral current events. Higher numbers mean more muon-like. The bin below zero indicates events not considered because the track was too short. 2008kNN follows a similar pattern.

selecting events with $2010\text{kNN} > 0.02$ or $2008\text{kNN} > 0.09$. (Each of the numerical values falls between the quantization of the PIDs, 0.025 in each case, so there are never any ambiguities from events with PID values *equal* to the cutoff values.)

- **Detector-like events:** This category consists of events in the supermodule gap, the edge of the supermodule gap, the coil hole, and those in the radial edge that are selected as detector-like. In these regions a tighter PID cut is applied. Here an event is selected if $2010\text{kNN} > 0.5$ or $2008\text{kNN} > 0.25$. This cut was chosen to match the cut used in the 2010 fiducial analysis.
- **Back events:** Tracks reconstructed in the back of the detector are, by definition, short. This means that 2008kNN, which requires 10 planes, frequently rejects them outright. 2010kNN recovers many of these. However, both 2008kNN and 2010kNN tend to give any such tracks values near zero both because the tracks are short and because they have lost information about the track end. In general, both are distributed in a very different way in this region than in others; see figure 8.3. The sensitivity to Δm^2 optimizes if events with $2010\text{kNN} > 0.01$ or $2008\text{kNN} > 0.14$ are selected.

	Rock-like	Detector-like	Back
All Selected	1409	360	73
Signal	1379	351	66
Neutral Current	27 (4.0 μ)	7.4 (0.5 μ)	8.1 (0.1 μ)
ν_τ appearance	1.3 (0.8 μ)	0.5 (0.4 μ)	0.3 (0.1 μ)
Beam ν_e	1.9 (0.3 μ)	0.3 (0.03 μ)	0.2 (0.01 μ)
Total non- μ background	25	7.3	8.4
Total μ background	5.1	0.9	0.2

Table 8.2: Monte Carlo beam prediction for the LE running, evaluated at the best fit point, of selected events broken down by the three regions in which different PID cuts are made. For the neutral current, ν_τ and ν_e backgrounds, the number of events containing a muon (being therefore irreducible) is shown.

8.4 Performance

The PID selection cuts have been tuned for high *efficiency* — the fraction of true signal events that are selected — at the expense of *purity* — the fraction of selected events that are truly signal.

Of rock-like events in which the longest reconstructed track is a muon which entered the detector with an energy of at least 1.5 GeV, 99.9% are selected (see figure 8.4). However, not all of these are primary muons of a charged current interaction. When the primary muon enters the detector with at least 1.5 GeV and is reconstructed, it is selected with 99.3% efficiency. The 0.7% difference is accounted for by the case where a track other than the primary muon is the longest and is rejected by the PID cuts. No attempt is made to examine other tracks in the event. Detector-like events are often complicated by large hadronic showers and so both efficiencies are lower and do not plateau until around 4 GeV (again see figure 8.4). Above this, they are 98.4% and 93.6%, respectively.

The muon selection efficiency is a measure of how many events could be recovered if the PID cuts were to be relaxed, or if an improved track classification method were to be developed. The difference between the two measures of efficiency gives the signal that could perhaps be partially recovered with a more sophisticated strategy that considers reconstructed objects other than the longest track.

The major selected background is neutral current interactions in which the longest track




Figure 8.4: PID selection efficiency, broken down into detector-like events (left), rock-like events (right), as a function of muon energy at the first scintillator hit (top) and as a function of neutrino energy (bottom). The efficiency of selecting muons (thick grey lines) is shown along with the efficiency of identifying $\nu_\mu/\bar{\nu}_\mu$ charged current interactions (thin black lines). The latter has a lower efficiency due to the case in which the longest reconstructed track is not a muon and therefore (in general) fails the PID. In all cases, the denominator includes only events in which a track was reconstructed.




Figure 8.5: Components of the selected sample for LE beam. The dominant background for both rock and detector samples is neutral current events in which the selected track is not a muon. 14% of the background consists of muon tracks produced by something other than a $\nu_\mu/\bar{\nu}_\mu$ charged current interaction, most commonly high energy neutral current events in the rock. There is also a very small background component from ν_τ appearance and intrinsic beam ν_e .

	MC	Oscillated MC	Data	Cosmics
LE Preselected	2415	2087	2354	1039
After cosmic rejection	2336	2023	2136	31
After PID	2199	1815	1907	19
Below 3 GeV	1063	766	813	11
pHE Preselected	157	149	152	29
After cosmic rejection	152	145	140	0
After PID	144	130	129	0
Below 3 GeV	65	55	59	0

Table 8.3: Effect of cuts on data. Preselected data is required to have been taken during good detector and beam conditions, have a track, have the longest track's vertex in the anti-fiducial volume, and to satisfy the timing cut. The cosmics are those from the sidebands of the spill window described in section 8.2. The MINOS 2008 fiducial result was used to oscillate the Monte Carlo. The events below 3 GeV are shown since this is the region with the bulk of the oscillation signal.

is not a muon. Also significant is an irreducible background of true muons arising from non- $\nu_\mu/\bar{\nu}_\mu$ charged current interactions. The majority of these are again from neutral current events. There are also very small backgrounds from ν_τ appearance and intrinsic beam ν_e . (If there is also ν_e appearance, the number of selected events is very nearly unchanged due to reduced ν_τ appearance.) See figure 8.5 and table 8.2.

After all cuts, the overall efficiency for selecting $\nu_\mu/\bar{\nu}_\mu$ charged current interactions is 91%, counting only events in which there is a reconstructed track. Of selected events, 2.7% are background, of which 88% are neutral current interactions.

8.5 Data

After all cuts, 361 events were selected in Run I, 559 in Run II, 987 in Run III, and 129 in the pHE running for a total of 2036 compared to a Monte Carlo expectation of 2344 in the absence of oscillations or 1944 assuming the previous MINOS result's best fit point. See tables 8.3 and 8.4. Since more than 1944 events were observed, it is expected that the fit value of Δm^2 will be lower than the previous result.

Region	LE Data	pHE Data
Coil	15	1
Front	850	53
Gap	56	2
Edge of gap	41	2
Back	85	5
Edge, rock-like	611	45
Edge, detector-like	249	21
Total	1907	129

Table 8.4: Selected events by region, separated into LE and pHE running.

Chapter 9

Near-to-Far Extrapolation

In an ideal two-detector experiment, the flux at the two detectors would be identical, or would be different only by a factor of the distance squared. In a realistic experiment, however, non-trivial corrections must be made to account for differences in the flux between the near and far detectors. As illustrated in figures 3.3 and 3.4, geometric differences — primarily the proximity of the near detector to the decay volume — give rise to these differences. In order to predict the spectrum of events at the far detector, the following steps are used:

1. Using the near detector data, preliminarily reweight the near detector Monte Carlo using a fit to hadron production model parameters.
2. Using the reweighted near detector Monte Carlo as a guide, convert the near detector data into a neutrino flux.
3. Apply geometrical corrections to find the far detector flux corresponding to the near detector flux. Reweight the far detector Monte Carlo accordingly to produce an unoscillated prediction.
4. Apply oscillations.




Figure 9.1: Comparison of Run III near detector data to nominal Monte Carlo and to the Monte Carlo with beam weights applied. The agreement is much improved after tuning, but still has systematic differences. The other runs follow a very similar pattern. Figure adapted from [63].

9.1 Initial Beam Fit

The untuned beam Monte Carlo produces a flux of neutrinos that differs from the observed flux at the level of $\sim 20\%$ near the focusing peak and as much as 60% in the high energy tail. This is due to large uncertainties in hadron production for 120 GeV protons on a thick carbon target. Of primary concern is the $\pi : K$ ratio and the distribution of longitudinal and transverse hadron momenta (p_z, p_T). Once in the decay pipe, modeling of the pions and kaons is orders of magnitude more accurate, as their lifetimes and decay kinematics are well understood. A non-negligible amount of additional uncertainty results from re-interactions of these hadrons with the walls of the decay pipe and with the helium in the pipe during Run III.

To correct this, the MINOS beam group provides a set of weights as a function of the neutrino parent hadron's species ($\pi^\pm, K^\pm, 0$), p_z and p_T with respect to the beam axis (see figure 9.1). These weights are derived from fits done to near detector data in a variety of special-running beam configurations, chosen to cover $p_z p_T$ space in order to constrain




Figure 9.2: The transfer matrix relating each true ν_μ energy at the near detector to a range of energies at the far detector. A separate matrix is defined for each neutrino species due to their different production distributions.

the production parameters. The fit also simultaneously allows for variations in the horn calibration, the amount of target decay, overall neutrino energy scale, relative $\bar{\nu}_\mu$ charged current cross section and neutral current background [63].

While this fit does not bring the Monte Carlo perfectly in line with the data, it provides a better starting point for the next step. It also produces a more accurate estimate of the characteristics of the neutrino parent hadrons; this is useful for systematic studies. For instance, in the MINOS antineutrino analysis, one systematic concerns the fraction of neutrinos produced by hadrons themselves produced in the beam pipe rather than in the target [64]. No similar systematics are considered for this analysis, but there is a systematic error due to the total uncertainty of this reweighting (see chapter 12).

9.2 Near Detector Data \rightarrow Flux

The near detector data is now converted into an estimate of the neutrino flux at the near detector. The usual MINOS approach for doing this, which this analysis uses, is called the *beam matrix method*. The beam-fit-weighted Monte Carlo is used to construct three

correction histograms that are applied to the data [65, 66, 67]:

1. A 1-dimensional histogram of *purity*: the fraction of selected events in each reconstructed energy bin that are true $\nu_\mu/\bar{\nu}_\mu$ charged current interactions.
2. A 2-dimensional histogram that gives in each bin the probability of the given true energy for the given reconstructed energy.
3. A 1-dimensional histogram of *efficiency*: the fraction of interactions in each true energy bin that are selected.

When these three are applied in sequence to a histogram of near detector data binned in reconstructed energy, the result is a histogram containing an estimate of the number of true $\nu_\mu/\bar{\nu}_\mu$ charged current interactions per bin of true energy. Dividing by the cross sections and fiducial mass gives an estimate of the flux at the near detector. The near detector flux is then related to the far detector flux via a 2-dimensional histogram, the *beam matrix* (see figure 9.2).

9.3 Flux \rightarrow Far Detector Prediction

At this point, the fiducial analysis reverses all of the steps to arrive at a prediction for far detector selected events in each reconstructed energy bin. This approach takes advantage of the similarities between the two detectors. Mismodeling at any step — the selection efficiency, for instance — will largely cancel. (Thus what we call “the flux” above is not claimed to be the real flux, but rather the real flux multiplied by $\text{purity}_{\text{true}}/\text{purity}_{\text{estimated}}$ and so forth.)

For rock and anti-fiducial events, the direct reversal of the matrix method at the far detector — unfolding the flux into a spectrum of reconstructed events — does not work since near detector fiducial volume events do not have a one-to-one correspondence with far detector rock and anti-fiducial events. If the two detectors were truly identical, we could use the same procedure, but with the near detector rock and anti-fiducial samples. However, the relevant details of the near and far detector are quite divergent: (1) the rock at the near detector has a different composition (2) the cavern shape is different at the two detectors




Figure 9.3: Correction functions used to reweight Runs I, II, and III Monte Carlo as a function of true neutrino energy. In the ν_μ spectrum, energies above and below the beam peak are generally weighted up compared to the peak itself. The pHE running also has a reweighting function not displayed here since it has a substantially larger variation.

(3) the anti-fiducial volume of the near detector has a very different geometry than the far detector's (4) the beam at the near detector is narrow with a strong radial dependence, but isotropic at the far detector.

For these reasons, rock and anti-fiducial events are predicted by using the ratio of flux derived above and the nominal flux at the far detector to directly reweight raw far detector Monte Carlo events as a function of energy and neutrino type.* The ν_μ and $\bar{\nu}_\mu$ correction functions are displayed in figure 9.3.

9.4 Oscillation

Oscillated predictions are produced by weighting each ν_μ and $\bar{\nu}_\mu$ charged current event by $1 - \sin^2 2\theta \sin^2(\Delta m^2 L/4E)$. Generally it is assumed that the same oscillation parameters apply to both neutrinos and antineutrinos. However, the two can be weighted using different values of $\sin^2 2\theta$ and Δm^2 if desired.

No oscillation weight is assigned to $\nu_e/\bar{\nu}_e$ charged current events, since Δm_{12}^2 -driven oscillations are too slow to have any impact to the result. The mean $\nu_e \rightarrow \nu_{\mu,\tau}$ oscillation probability is $\sim 10^{-4}$. Given that only 1.3% of the beam is ν_e , a rough estimate of the combined number of additional ν_μ and ν_τ events at the far detector due to these oscillations is 0.01 over all runs. Neutral current events are unaffected by oscillation and so also are not reweighted at this step.

To account for $\nu_\tau/\bar{\nu}_\tau$ appearance from $\nu_\mu/\bar{\nu}_\mu$, a sample of tau-only Monte Carlo weighted by the inverse probability $\sin^2 2\theta \sin^2(\Delta m^2 L/4E)$ is added to the normal beam Monte Carlo. This is a very small effect contributing only 2.3 selected events at the best fit point (1.2 from the rock, 1.1 from the detector) over all runs. The possibility of ν_e appearance is ignored in this analysis, but it could be handled similarly by replacing a small fraction of the ν_τ appearance. This approximation changes the background estimate by $\lesssim 0.2$ events.

*While the mechanics are different, in principle this approach could be applied to the prediction of the reconstructed fiducial spectrum as well and should produce the same answer. The difference lies only in whether the efficiencies, purities, resolutions, and cross section are explicitly constructed from the Monte Carlo as histograms/matrices or whether the Monte Carlo is used directly.

Chapter 10

Oscillation Analysis

In a fiducial analysis, the fit is performed by binning events by reconstructed neutrino energy and comparing this spectrum to the prediction for various sets of oscillation parameters. The full neutrino energy is not available in this analysis, however, the same procedure can be done binning in muon energy instead. Individual muon energies do not map directly onto neutrino energies, but naturally the prediction is still modified by oscillations. The fit can be improved if events are separated further by variables that, at a given muon energy, correlate with the neutrino energy. The track angle and the region of the detector containing the track vertex are therefore also used in the fit.

10.1 Region Separation

10.1.1 Motivation

We measure the initial energy of muons originating in the detector, but only the remaining energy of muons originating in the rock. While neither gives the total neutrino energy, the detector muons' energy is more strongly correlated to neutrino energy, so total sensitivity is increased by separating these categories. In addition, the spectra of neutrinos in the two categories is different. The large volume of upstream rock provide a large reaction mass for the beam, but only high energy events deep in the rock produce particles that reach the

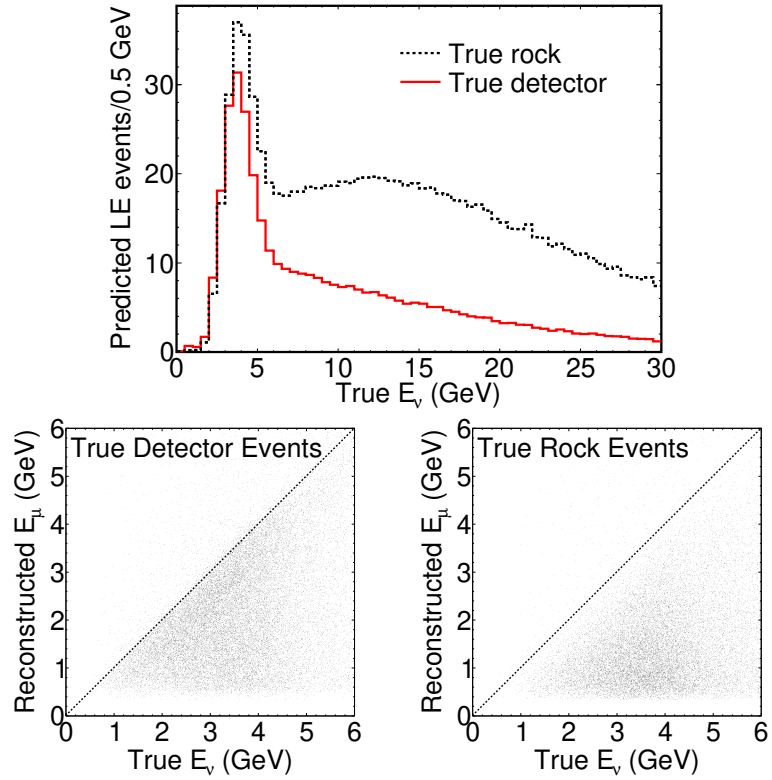


Figure 10.1: Top: True neutrino spectrum (LE beam) for selected anti-fiducial events, separated into true rock events and true detector events. Bottom: Relationship between neutrino energy and reconstructed muon energy for detector and rock events. The diagonal line in each case shows $E_{\nu, \text{true}} = E_{\mu, \text{reconstructed}}$.




Figure 10.2: Comparison of true neutrino energy (LE beam) for selected rock-like events on the front face and the edge. The front face events tend to be higher energy since it is easy for them to reach this face from deep in the upstream rock.

detector (see figure 10.1). So separating rock events from detector events is beneficial for this reason also.

Within rock events, further separation is beneficial. The detector’s front face has a large acceptance for high energy muons produced deep in the upstream rock. In contrast, the detector’s radial edge has a small cross section for muons travelling along the beam axis and so receives a larger fraction of muons from nearby, lower energy, interactions (see figure 10.2). In addition, rock-like events in the two regions have different fractions of detector event contamination. Since there is so little detector mass in the front face, events there are 98% rock-pure. In contrast, the more complex geometry of the edges means that rock-like events there have a much larger detector contamination of $\sim 10\%$.

As with rock events, not all detector events are created equal. The bulk of anti-fiducial detector events occur around the radial edge, away from the front, supermodule gap and back. These are high quality events since they can be well separated from rock events and have good muon energy resolution. However, the statistical power of the various minor regions — near the supermodule gap, around the coil hole, and the back of the detector — is still significant. These regions have different resolutions and backgrounds so are separated into different histograms.

Region	Unoscillated	At Best Fit	NC background
Coil	20	17	1.3%
Front	878	820	1.4%
Gap	70	57	1.6%
Edge of gap	41	37	0.8%
Back	94	74	9.4%
Edge, rock-like	665	590	3.3%
Edge, detector-like	328	265	1.8%
Total	2097	1860	2.5%

Table 10.1: Monte Carlo predicted selected events by region for Runs I–III, LE only, unoscillated and oscillated at the best fit point. The fraction of selected events that are true neutral current in the oscillated sample is given. It is high for the events in the back of the detector due to a large number of exiting tracks. In these events, PID variables that rely on pulse height at the track end lose power.

10.1.2 List of Regions

As shown in figure 5.1, the anti-fiducial volume can be divided into as many as $5 \times 3 - 1 = 14$ regions: {front, longitudinally fiducial, south gap, north gap, back} \otimes {coil, radially fiducial, edge}, excluding the fiducial volume itself. However, some of these regions are very small and some are very similar to each other. Therefore, these groupings are used:

1. The longitudinally fiducial coil is a low-statistics, high-resolution, detector-pure region, and so is kept separate. The coil regions that overlap other regions are distributed below.
2. All front regions (coil, radially fiducial, edge) are grouped together. The edge and coil events here are more similar to the other front events than to the other edge/coil events because the vast majority of them are rock events starting in the first plane.
3. Both sides of the radially fiducial supermodule gap are characterized by being quite detector-pure. The gap+coil events are also included here. While the north side has noticeably more rock event contamination, the statistics here are fairly low and not very much is to be gained from treating them separately, so they are grouped together.

4. Both sides of the edge of the supermodule gap are likewise similar. Again the north side has a higher concentration of rock events, but even more so than above the statistics are quite low, so the two are grouped together.
5. The back is a relatively low-statistics region characterized by poor track energy resolution. As with the front face, all three radial regions are grouped together.
6. Longitudinally fiducial edge events. Instead of being combined with other regions, these are *split* into two categories:
 - (a) Rock-like
 - (b) Detector-like

When it is not ambiguous, these 7 categories will be called, respectively: the coil, the front, the gap, the edge of the gap, the back, the rock-like edge and the detector-like edge. The number of events expected in each is given in table 10.1.

10.2 Rock/Detector Separation on the Edge

In all regions other than the edge, it is either the case that the vast majority of events are true rock (such as the front face), or that the vast majority are true detector (such as the supermodule gap), or that the statistical power of the region is too small to warrant an attempt to separate rock events from detector events (such as the back). On the edge, however, the distribution of events is roughly half and half, with a large number of total events.

The nearer an event vertex is to the edge, the more likely, in general, it is to be a rock event. In this region, rock events effectively form a background to the detector events and vice versa. However, since the detector sample is of higher quality event-by-event, it is more important to achieve a pure detector sample than a pure rock sample. Fortunately, the event characteristics tend to lead to this result in any case. Rock muons cannot realistically appear deep into the detector in a way that resembles a detector event, but detector events can easily occur at the very edge of the detector and resemble a rock event.




Figure 10.3: Left: Rock/detector separation on horizontal and vertical edges. When $H \leq 57$, an event is called rock-like and otherwise detector-like. The Monte Carlo is oscillated at the best fit parameters. Right: The same on diagonal edges. When the parallel strip number is 0 or 1, an event is called rock-like and otherwise detector-like.

10.2.1 Horizontal and Vertical Edges

A strip-based selection is used to separate rock events from detector events. Since the detector has two types of edges, two different selections are required. As shown in figure 5.4, the horizontal and vertical edges are formed by the ends of strips in the A and B modules of each plane. Rock events are separated from detector events by looking at the first U strip hit and the first V strip hit in the track. Adding or subtracting the strip numbers appropriately* gives a single number H which is a measure of how close the track vertex is to the edge. For instance, on the bottom edge, moving a vertex horizontally increases the U (V) strip number and decreases the V (U) strip number while the sum H stays constant. Moving towards the edge by $\sqrt{2}$ strip widths reduces H by 2 since both U and V numbers are reduced by 1. (For Monte Carlo events, a small correction is applied while calculating H to account for the difference between the ideal Monte Carlo geometry and the real detector; see section 12.2.3.) If $H \leq 57$, the event is classified as rock-like and otherwise detector-like. This choice was made by scanning across possible cut values and choosing the one with the best sensitivity to Δm^2 .

*The strip numbers run 0–55 or 136–191 on these edges.

$H = \{1 + u + v$ (bottom), $383 - u - v$ (top), $192 - u + v$ (west), $192 + u - v$ (east) $\}$

10.2.2 Diagonal Edges

The diagonal edges have an entirely different geometry (see figure 5.4). They are formed in one view by strips running parallel to the edge, with the shortest strip in an A module forming the edge itself. In the other view, the edge is formed by the ends of full length (8 m) strips. These overhang the A module in the first view by about one strip width.

Rock events are separated from detector events by looking at the first hit in the view whose strips are parallel to the edge. The *parallel strip number* is defined as the strip with the first hit, the outermost strip being numbered 0. A track is selected as rock-like if its parallel strip number is 0 or 1. A large majority of tracks with parallel strip number = 0 are rock events; see figure 10.3. Tracks with parallel strip number = 1 are somewhat more likely to be detector events than rock events, but a better overall sensitivity is obtained by lumping these in with the rock events since the rock background is high. While in principle better discrimination could be achieved by also using the overhanging perpendicular strips (rock muons usually hit these first), by selecting in this manner the exposure to alignment uncertainties is minimized (see section 12.2.3).

10.2.3 Veto Shield

Some rock muons entering the edge pass through the veto shield. It seems like using this information to separate rock events from detector events would be both beneficial and more straightforward than the above procedures. Unfortunately, the detector's Monte Carlo model has no veto shield. This adds quite a bit of complication to any attempt to use it in an analysis.

To see if this effort would be worthwhile, I looked at the gain in sensitivity that would result if the shield could be used to achieve *perfect* rock/detector separation on the edges it covers. For the bottom 3 edges I marked events as rock-like or detector-like according to the above discrimination techniques, while for the top 5 edges I instead used the Monte Carlo truth information. The gain in sensitivity to Δm^2 at maximal mixing was only 0.6%. This can be explained by two observations. First, the above methods already achieve quite good separation. Second, due to the beam angle, fewer events enter the detector's top than its

bottom. Were the bottom of the detector also shielded, the gain would be proportionally somewhat greater, although still not large. Realistically, the shield would give an even smaller benefit than found in this idealized study. Therefore, it is not used in this analysis.

10.3 Angle

Since this analysis does not use shower energy, any available variable that correlates with neutrino energy is useful. Essentially the only variable left is the track direction. The relevant quantity is the angle of the track with respect to the beam. This correlates with neutrino energy at a fixed reconstructed track energy.

Two competing effects are present. First, for lower energy events, the initial muon angle correlates positively with neutrino energy at a fixed value of muon energy, since in general a larger muon angle means a low-elasticity interaction in which a relatively large amount of energy went into the recoiling hadronic system. Second, for higher energy events, the muon angle *anti-correlates* with neutrino energy because the higher the neutrino energy, the more forward-boosted the interaction.

The first effect is important primarily for detector events. Here the high energy tail is relatively small (see figure 10.1) and so the main challenge is to sort out the energy of neutrinos inside the focusing peak. For example, as shown in figure 10.4, 1.0–1.5 GeV muons are most likely from ~ 2 GeV neutrinos if their angle is below 10° and most likely from ~ 5 GeV neutrinos otherwise.

The second effect is important primarily for rock events. Here the main challenge is to separate the high energy tail (magnified by the large rock volume) from the focusing peak. The most important consideration is the lack of information about the initial muon momentum. Muons arrive from the rock with an energy that is, roughly speaking, a uniform random number between their initial momentum and zero. The reconstructed rock muon spectrum peaks at 0.75 GeV (falling off below this primarily because of the detector's efficiency limitation), so many different initial energy muons all arrive with the same low momentum. As seen in figure 10.4, out of low reconstructed energy muons, almost all of those from the focusing peak have an angle greater than 10° . Low energy, small angle

muons, in contrast, originate more than half of the time from interactions above 10 GeV.

To take advantage of this information, events are, for most regions, binned in both energy and angle. The precise binning scheme is described below. Figure 10.5 shows the improvements to the sensitivity that each set of refinements brings.

10.4 Binning and Fitting

Here follows the full detail of the binning used in this analysis. There are 4 run periods and 7 detector regions. In each run period, the binning is identical. The binning below 5 GeV was optimized for LE running. It is applied to pHE without reoptimization since the pHE running has such a small amount of statistical power that the overhead of using a different scheme would not be worthwhile. The binning above 5 GeV was chosen to match that used in the fiducial analysis. This was for ease of explanation more than anything else, given that the bulk of the oscillation information is below 5 GeV. The following four binnings are used for the 7 regions:

1. For the three high-statistics regions — the front, rock-like edge, detector-like edge — events below 5 GeV are binned in track energy and angle. Since there are regions of energy/angle space with very low statistics, a non-uniform binning scheme is used (displayed in figure 10.6) that prevents having too little Monte Carlo in any given bin. Above 5 GeV, events are binned in energy alone. Up to 20 GeV, 0.25 GeV bins are used, then 1 GeV bins up to 30 GeV, 2 GeV bins up to 50 GeV, a single bin for 50–200 GeV and finally a bin for events above 200 GeV. No beam events should actually be above 200 GeV, so this final bin primarily isolates poorly reconstructed tracks. (134 bins each)
2. For the gap, the same 2-dimensional binning is used as above for 0–5 GeV tracks. Since this region has lower Monte Carlo (and data) statistics, a more sparse binning is used above 5 GeV: 0.5 GeV bins up to 10 GeV, then 5 GeV bins up to 30 GeV, a single bin for 30–200 GeV and a bin for all events above 200 GeV. (68 bins)




Figure 10.4: Energy/angle correlations in unoscillated LE Monte Carlo. Left: True rock muons' reconstructed angle correlates negatively with neutrino energy at a fixed reconstructed track energy. Right: Detector muons' angle correlates positively with neutrino energy, except that under 10° there is also a significant contribution from highly boosted events.




Figure 10.5: Effect that each improvement of the binning scheme has on the fit. Each contour is a sensitivity assuming maximal mixing and $\Delta m^2 = 2.43 \times 10^{-3} \text{eV}^2$, the MINOS 2008 best fit point. (Solid contours appear in the same order as in the caption.) Note that $\sin^2 2\theta$ benefits disproportionately as the analysis sophistication increases. This is because the direct effect of most of the improvements is to move high energy neutrino events out of low muon energy bins, where they act as a background to measuring the depth of the deficit.




Figure 10.6: 2-dimensional binning for front, gap, and rock-like edge and detector-like edge histograms. Above 5 GeV, events are binned in energy alone as described in the text. In each bin is shown the total predicted events at the best fit oscillation parameters for all LE running summed over the four regions that use this binning. The color scale shows events/GeV/degree while the numbers give the absolute number of events per bin.

3. The edge of the gap suffers from quite low statistics and from having both many rock events (24%) and many detector events (76%). The coil also has too-low statistics for 2-dimensional binning. Therefore they are binned in energy alone, using 0.5 GeV bins up to 10 GeV, 5 GeV bins up to 30 GeV, a single 30–200 GeV bin and a bin for all events above 200 GeV. This is the same binning as for the gap, but without use of angle. (26 bins each)
4. Track energy in the back is nearly useless. Since nearly all tracks exit after a short distance, neither range nor curvature is a reliable measure in most cases. A study showed that substantially more power was available in the angle. Since the back is a fairly low statistics region, it is binned in angle alone, using 1° bins in 0–40°, 5° bins in 40–60°, a single 60–90° bin, and a bin for angles above 90° (possible for tracks pointing almost straight down since the beam is angled slightly upwards, but not a practically important case — only five out of 1.2M Monte Carlo events managed it). (46 bins)

Put together this is 568 bins per run for a grand total of 2272 bins. The fit is performed using a grid search over Δm^2 and $\sin^2 2\theta$. The best fit point is that at which the likelihood

function [22]

$$-2 \log_e \lambda(\Delta m^2, \sin^2 2\theta) \equiv 2 \sum_{i=1}^N \left[m_i - d_i + d_i \log_e \frac{d_i}{m_i} \right] \quad (10.1)$$

is minimized, where $m_i \equiv m_i(\Delta m^2, \sin^2 2\theta)$ is the number of events predicted by the Monte Carlo to be in the bin i and d_i is the number of data events in that bin. Note that if $d_i = 0$, $d_i \log(d_i/m_i) = 0$ and that m_i must be non-zero for all bins with a non-zero number of data events.

In the limit of large d_i , $-2 \log \lambda$ follows a χ^2 distribution and can be used to determine the goodness-of-fit. However, in this case, the mean number of events per bin is less than one, and so the absolute value of $-2 \log \lambda$ has no such clear meaning. However, we still expect it to be in the neighborhood of the number of bins if the fit is good and so it is often loosely referred to as χ^2 ; I will follow this convention. Relative values of χ^2 evaluated away from the best fit point give the confidence interval. The contour at which $\Delta\chi^2 = 1$ (2.71) gives 68.27% (90%) confidence for a 1 parameter fit and $\Delta\chi^2 = 2.3$ (4.61) gives the same for a 2 parameter fit.

The sensitivity of this analysis is shown in figure 10.7. Assuming maximal mixing and Δm^2 in the neighborhood of the MINOS 2008 best fit point (or the eventual best fit of this analysis), 1 parameter sensitivity to Δm^2 is $\pm 0.18 \text{ eV}^2$. The statistical power in Δm^2 is split between rock-like and detector-like events, with detector-like events contributing somewhat more. In $\sin^2 2\theta$, the detector-like events strongly dominate the statistical power due to their much higher energy resolution.

10.4.1 Alternative Binning

In order to integrate more easily with the fiducial analysis a second, dramatically less complicated, binning scheme was also used to fit the data. This scheme uses 6 region histograms instead of 7, with coil hole events dropped entirely and the other regions defined as above. Each region is binned in energy alone using the following scheme:

1. Front face, rock-like edge and detector-like edge histograms (high statistics): One bin for 0–0.75 GeV, 0.25 GeV bins up to 4 GeV, 0.5 GeV bins up to 6 GeV, 1 GeV bins up




Figure 10.7: Sensitivity for rock-like regions (front face, rock-like edge, and supermodule gap edge), detector-like regions (detector-like edge, supermodule gap, coil and back) and their sum, the complete analysis. The truth point is shown as a filled disc. It has been chosen after the fact to be the same as the best fit point for ease of comparison with the result.

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. (28 bins each)

2. Gap, gap edge and back histograms (low statistics): 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 a bin for all events above 15 GeV. (9 bins each)

This is 111 bins per run for a total of 444 bins. This simplification sacrifices 2.7% sensitivity to Δm^2 at maximal mixing and 3.7% sensitivity to $\sin^2 2\theta$.

10.5 Mock Data Challenge

Before examining the full far detector data set, two types of tests were run to test for problems in the code and for the possibility of biases being introduced by the analysis method. First, during the optimization procedure, the fit was run by comparing the full Monte Carlo set to the same set Monte carlo set uniformly scaled down the size of data and oscillated. If all is well, the input oscillation parameters should be exactly recovered by the fit. The disadvantage to this approach is that fractional events are used in the Monte Carlo




Figure 10.8: Distribution of best fit points of 100 mock data sets. About half of the fits are at $\sin^2 2\theta = 1$. These are shown with a small amount of random scatter so that each point can be seen individually.




Figure 10.9: Example of a mock data result near the truth point.

histograms that are functioning as the data. This could conceivably mask problems or give misleading results.

The second, more demanding, test was known as the *mock data challenge*. Here, 100 sets of Monte Carlo were specially generated with the same exposure as the data set. They were oscillated using parameters known only to the Monte Carlo group. Fitting these sets tests whether the analysis functions properly when the data comes in discrete (integer-valued) events and has realistic statistical scatter. The protocol of keeping the input parameters secret tests whether the fit can return the correct answer even when the experimenters do not know the correct answer up front. Also important is there is no overlap of events between the “data” and the Monte Carlo it is compared against. Once all of the mock data was fit and the results presented to the collaboration, the input parameters were revealed.

The fits to the mock data were examined in three ways. First, the distribution of the 100




Figure 10.10: Result of fitting the sum of all 100 mock data sets.

best fit points should be clustered around the input parameters. The clustering should follow the shape of the sensitivity contours generated using scaled down Monte Carlo. Second, the contours of each individual mock data fit should resemble the sensitivity contours if its best fit point is near the input parameters. Third, if all 100 mock data sets are added together and treated as a single experiment with 100 times the exposure, the best fit should be quite close to the input parameters.

The anti-fiducial analysis passed each of these tests. As shown in figure 10.8, the best fit points are clustered around the input parameters of $\Delta m^2 = 2.1704 \times 10^{-3} \text{ eV}^2$, $\sin^2 2\theta = 0.9756$. The tail of best fits at substantially lower $\sin^2 2\theta$ follows the shape of the sensitivity contours as shown in figure 10.5. Likewise, fits of individual mock data sets produced similar contours when their best fits were near $\sin^2 2\theta = 1$; see example in figure 10.9. When all 100 sets were summed, representing an exposure of 7.25×10^{22} POT, the resulting fit agreed well with the input (see figure 10.10).

Chapter 11

Results

For all runs combined, the oscillation parameters are found to be $\sin^2 2\theta = 1_{-0.29}$ (greater than 0.56 at 90% confidence) and $\Delta m^2 = 2.16_{-0.18}^{+0.62} \times 10^{-3} \text{ eV}^2$. The best fit point and errors in this chapter are statistical only. Both will be modified by systematics; see chapter 12. Contours for individual runs and for the whole data set are shown in figures 11.1 and 11.2. There is good agreement amongst the four runs (see table 11.1 and figure 11.3) and also good agreement with published results.

This analysis has very poor sensitivity to $\sin^2 2\theta$ compared to previously published results, but it has significant power in Δm^2 . Therefore it is appropriate to take the world best value of $\sin^2 2\theta$ as a given and perform a single parameter analysis of Δm^2 . It is clear that the true value of $\sin^2 2\theta$ is near maximal, with the most accurate measurement coming from Super-Kamiokande: >0.965 at 90% confidence [26]. At $\sin^2 2\theta = 1$, this analysis finds $\Delta m^2 = (2.16 \pm 0.18) \times 10^{-3} \text{ eV}^2$. For non-maximal $\sin^2 2\theta$ above ~ 0.9 , $\Delta m^2 = [2.16 + 1.6(1 - \sin^2 2\theta) \pm 0.18] \times 10^{-3} \text{ eV}^2$. For $\sin^2 2\theta$ far from maximal, the dependence becomes non-linear, as is evident in the contour plots.

11.1 Checks

The distribution of events in energy and angle in each detector region and run have been checked and no pathologies found. Figure 11.4 shows the distribution of all events by energy;

Run	I	IpHE	II	III	Combined
$\sin^2 2\theta$	0.62	—	1.0	0.80	1.0
Δm^2 (10^{-3} eV^2)	3.16	—	1.75	2.90	2.16
$\Delta m^2, \sin^2 2\theta \equiv 1$	$2.18^{+0.44}_{-0.45}$	$2.05^{+1.62}_{-2.05}$	$1.75^{+0.35}_{-0.32}$	$2.43^{+0.26}_{-0.26}$	2.16 ± 0.18
χ^2 of best fit	507.34	317.82	555.01	570.25	1954.16
$\chi^2_{\text{maximal}} - \chi^2_{\text{best}}$	0.49	0.23	0.00	0.27	0.00

Table 11.1: Best fit points with $\sin^2 2\theta$ free, and the best Δm^2 values with $\sin^2 2\theta$ constrained to be 1. One parameter errors are given on Δm^2 for the second case. The pHE contour is not closed and so the 2 parameter best fit is not given. See figures 11.1 and 11.2 for two parameter contours. χ^2 generally increases if $\sin^2 2\theta$ is not allowed to float. The amount of this increase is given, showing that no run's best fit is significantly different from maximal mixing.




Figure 11.1: Statistical-only contours for Runs I, IpHE, II and III separately and the combined contour. (The allowed region for the Run IpHE includes nearly all of the plotted area.) The best fit point for the combined fit is shown.




Figure 11.2: Statistical-only contours and best fit point for all runs combined.




Figure 11.3: Top: With $\sin^2 2\theta$ fixed at 1, runs' $\Delta\chi^2$ values with respect to their individual best fit points. Bottom: The same information in a stacked histogram. The minimum of the sum of the four runs gives the overall best fit. Note that the pHE running contributes almost nothing to the fit near the best fit point, but becomes significant in the strong exclusion of large values of Δm^2 .




Figure 11.4: Events in all regions by energy. Left: LE Runs I, II, and III combined. Right: pHE. In this and following figures, the unoscillated Monte Carlo prediction is shown in blue (upper line) and the best fit in red (lower line).

Figure 11.5 does the same for angle. Figures 11.6 and 11.7 break these up by region for LE data. Figure 11.8 shows the 2-dimensional energy/angle distributions for LE data.

Particularly encouraging is the fact that the normalization in the high energy bins, which for Monte Carlo is dead-reckoned from near detector data and far detector Monte Carlo, is completely consistent with the data in all regions. This is a good check because the high energy tail is largely unaffected by oscillations. To quantify the level of agreement, consider figure 11.4. Over 10 GeV for LE beam, the Monte Carlo prediction without oscillations is 406.8 events and at the best oscillation fit, 401.5. In this region, 396 data events were selected, which is in excellent agreement. However, this is not a very strong constraint. The normalization of the Monte Carlo could change 3.6% upwards or 6.3% downwards before it disagreed with the data at the 1σ level. Section 12.3.1 details how stronger constraints are obtained; the agreement is maintained.

All MINOS beam results to date* have found a best fit for $\sin^2 2\theta$ at its maximal value of 1. This has been cause for some concern, since if the true value were 1, half of fits should be off-maximal. We have not had so many results as to rule out the possibility that this is simply chance, but nevertheless, it is a good idea to look into the matter. A possible cause could be that we are estimating the resolution of the detector to be worse than it is.

*Run I fiducial [36], Runs I+II fiducial [28], Runs I+II+III fiducial [43], Run I anti-fiducial [59], and this analysis. I am not considering individual Runs within a combined result for this purpose.




Figure 11.5: Events by angle. Left: LE Runs I, II, and III combined. Right: pHE.

At maximal mixing we expect the oscillation minimum to be filled in to a certain extent from adjacent energy bins. If we observe fewer events at the minimum than this, it appears that the mixing is “greater than maximal”. This naturally will happen half of the time via statistical fluctuations, but will appear more than half of the time if the estimate of the amount of fill-in is overly pessimistic. Other mechanisms of overestimating the background near the oscillation minimum will have the same effect. The neutral current background peaks at low energy, for instance.

One simple check that could reveal a serious problem is to look at the slope of χ^2 at $\sin^2 2\theta = 1$. On average, this should be zero, since it is positive for off-maximal fits and negative for maximal fits. For non-pathological maximal fits, the slope should be small, and for this analysis it is: see figure 11.9. MINOS’s resolution has been extensively studied for the 2010 fiducial analysis [43] and no problems have been found, so we will continue to ascribe these results to chance.

The rock-like and detector-like regions have been fit separately to check for pathologies in the rock sample (or, less likely, the detector sample); see figure 11.10. Both samples have reasonable data contours that relate in a straightforward way to the sensitivities shown in figure 10.7. The best fit for detector-like events is off-maximal, while the best fit for rock-like events is maximal. Like the sample as a whole, its χ^2 slope at $\sin^2 2\theta = 1$ is shallow, but not completely vanishing, which pulls the contours in compared to the sensitivity.

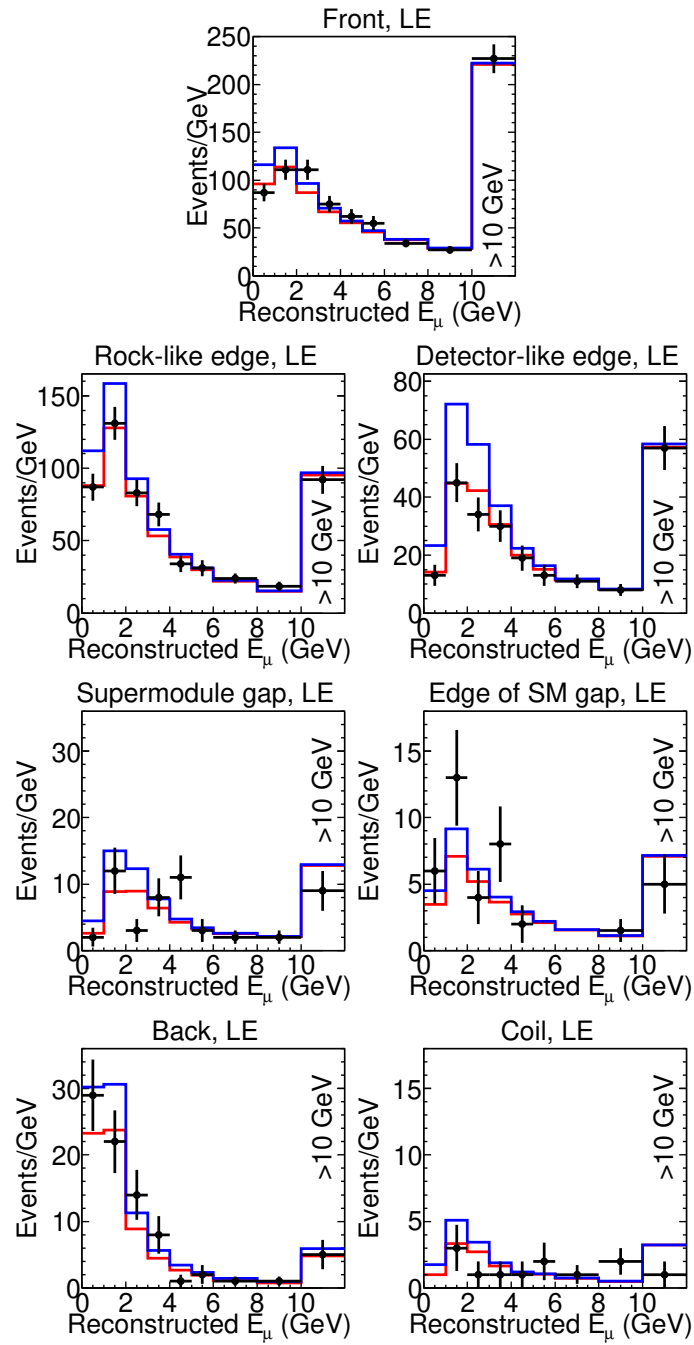


Figure 11.6: Events in each region by energy, LE running only, Runs I, II, and III combined.

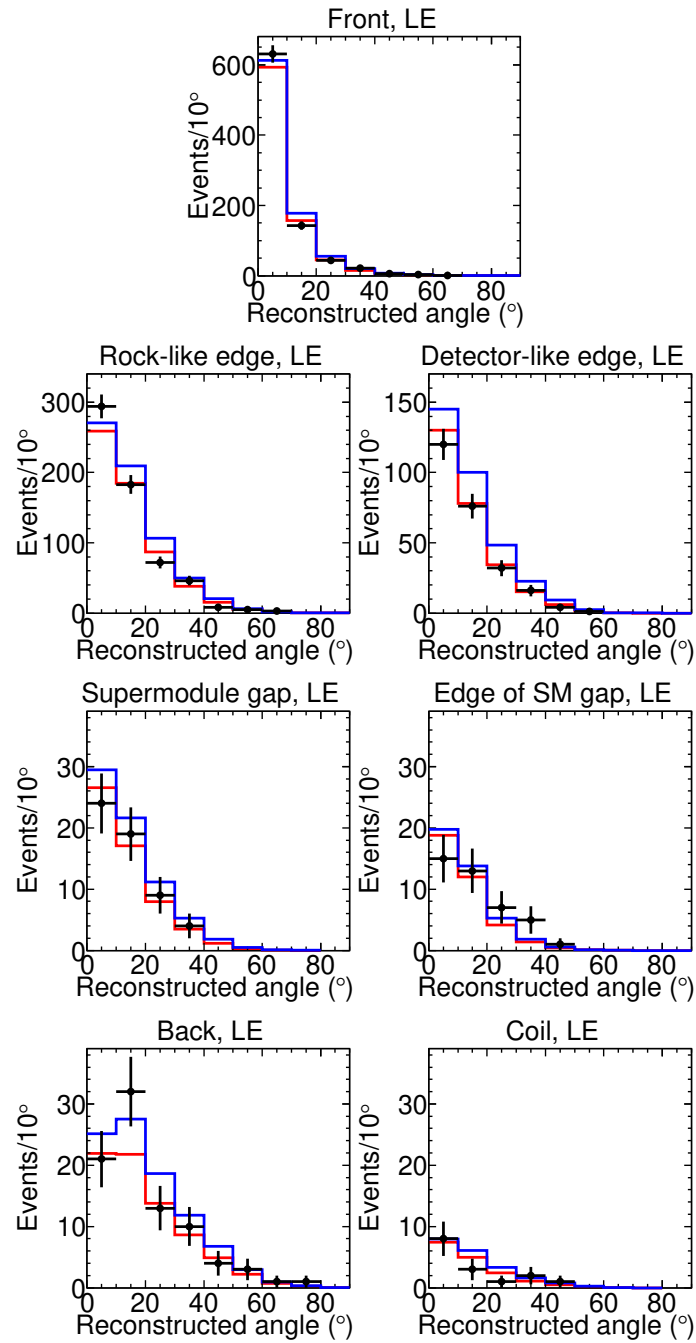


Figure 11.7: Events in each region by angle, LE running only, Runs I, II, and III combined.

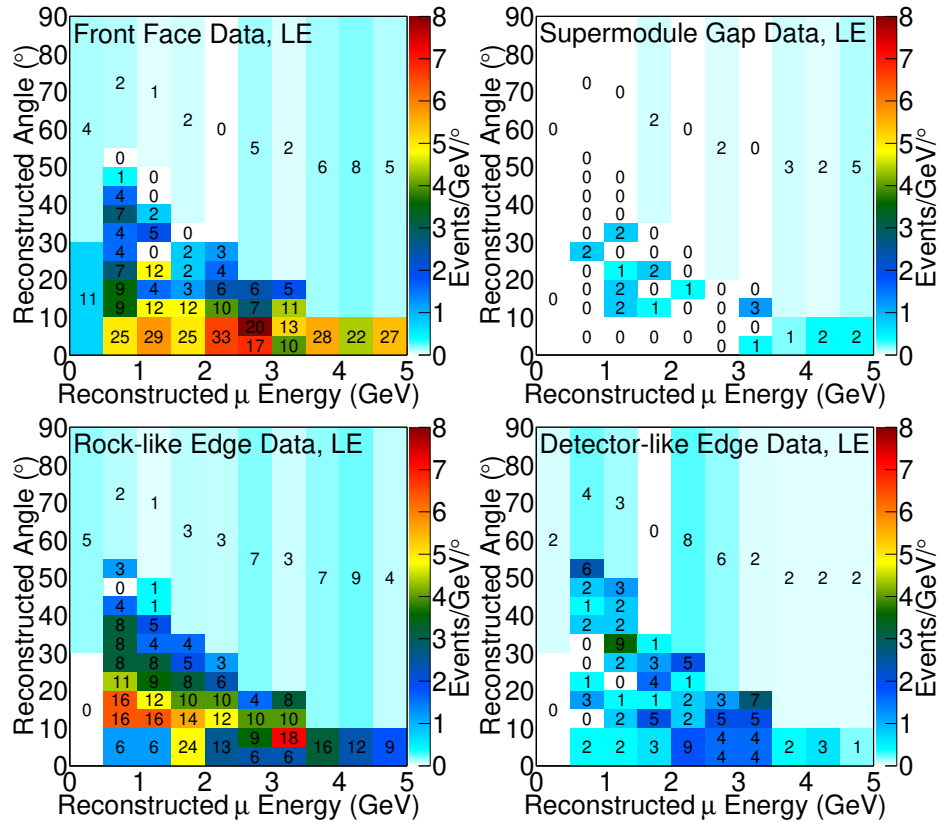


Figure 11.8: Events below 5 GeV in their energy/angle fit histograms, LE (combined Runs I, II and III).




Figure 11.9: χ^2 as a function of $\sin^2 2\theta$. At each point, the best fit Δm^2 for that $\sin^2 2\theta$ is used. The slope at maximal mixing is small; a parabola fit to the end of the curve has its minimum at 1.08.

The Run I anti-fiducial data has been previously analyzed [59]. This result should agree with the previous result, although since the analysis method has changed, perfect correspondence is not expected. Figure 11.11 shows the results by both analyses for Run I as well as this analysis for all data. The agreement is good. In this analysis, the Run I contours have bulged out somewhat as a result of the best fit being found off-maximal. This is not particularly meaningful since, as given in table 11.1, there are only 0.49 units of χ^2 between the new best fit and maximal.

At maximal mixing, the two analyses agree well, with the new analysis being shifted down slightly in Δm^2 . The difference at maximal mixing is $0.14 \times 10^{-3} \text{ eV}^2$, about 30% of the single-parameter Δm^2 error for either analysis. Since nearly every component of the analysis has been modified in some fashion — including the definition of the fiducial volume — it is difficult to assign a particular reason for this shift. However, the reduction in Δm^2 would generally indicate either a decrease in the predicted normalization or an increase in the number of selected events. Two instances of the former can be identified. First, the strip length correction discussed in section 6.4.1 reduces the normalization for the main category of detector events by 5.3% and rock events by $\sim 0.9\%$. Second, the Monte Carlo rock event cross section has been adjusted downwards (see section 12.2.1), which accounts for an additional reduction in rock event normalization of 0.3%. Without these corrections, this analysis would fit about $0.22 \times 10^{-3} \text{ eV}^2$ higher, so this is more than enough to account for the shift in Δm^2 .




Figure 11.10: Statistical-only contours for rock-like regions (front face, rock-like edge, and supermodule gap edge), detector-like regions (detector-like edge, supermodule gap, coil and back) and their sum, the complete result. The best fit point for rock-like regions alone is shown as an open circle, the detector-like regions alone as a square, and the combined fit as a filled disc. Compare to figure 10.7




Figure 11.11: Comparison of this analysis to the 2007 Run I anti-fiducial analysis [59]. The best fit point for that analysis is shown as an open square. The same data reanalyzed here is shown with the best fit point as a closed square, and the entire data set with the best fit as a disc.

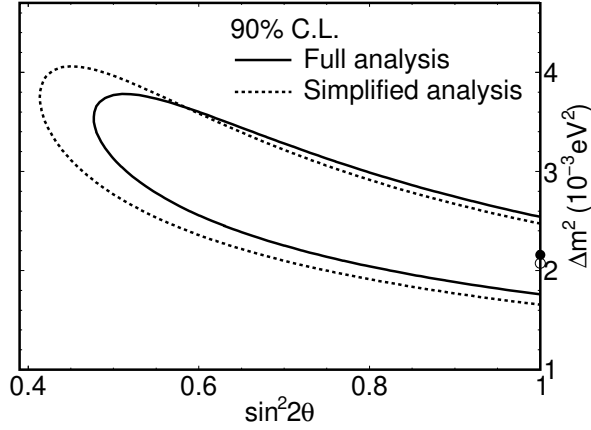


Figure 11.12: Comparison of the full analysis described in section 10.4 and the simplified one described in section 10.4.1. The simplified analysis is somewhat weaker because it drops the coil hole events (1% of the statistics), uses coarser bins, and does not use track angles in the fit. The best fit points are given as a filled disc (full analysis) and an open disk (simplified analysis).

11.2 Alternative Binning

The fit was also performed using the alternate energy-only binning scheme described in section 10.4.1. The result was $\sin^2 2\theta = 1_{-0.35}$ (greater than 0.50 at 90% confidence) and $\Delta m^2 = 2.07^{+0.77}_{-0.19} \times 10^{-3} \text{ eV}^2$, or at maximal mixing $\Delta m^2 = (2.07 \pm 0.19) \times 10^{-3} \text{ eV}^2$, statistical errors only. This is a somewhat weaker result that agrees well with the main result. The contours resulting from each analysis are compared in figure 11.12.

11.3 Contribution to MINOS Result

This alternate binning was used to combine the fiducial and anti-fiducial analyses. The result is shown in figure 11.13. Compared to the fiducial result alone, the combined result has a slightly lower best fit Δm^2 . Both fiducial and anti-fiducial samples' best fit is at maximal mixing. The combined result tightens the error on Δm^2 by 4%. The decrease in the overall MINOS error on $\sin^2 2\theta$ is tiny: $\sim 0.2\%$.




Figure 11.13: Contribution of this analysis to the overall MINOS result. The fiducial-only best fit is shown as a filled disc and the combined best fit as an open circle. These contours include systematics as nuisance parameters (see chapter 12).

Chapter 12

Systematic Errors

Systematic effects are a significant, but not overwhelming, component of the error in this analysis. Two major categories of systematics were considered in this analysis. First are systematics that affect both the near and the far detector and are shared with the fiducial analysis. Second are systematics that affect the far detector only and are unique to this analysis. By design, the first category tends to cancel at first order. The second category deals with our knowledge of the far detector edges and rock. These are first order effects with no cancellation from the near detector.

12.1 Shared Systematics

The systematics considered by the fiducial analysis are as follows. They are assumed to be independent of each other and fully correlated between the two detectors unless stated otherwise. Two of these systematics are dominant (see table 12.4 and figures 12.6 and 12.7): the relative near/far normalization and the muon track energy.

12.1.1 Relative Near/Far Normalization

The error on the overall normalization of events in the far detector relative to the near detector is 1.54% for detector events and 1.51% for rock events. This is primarily due to a 1.3% relative near/far track finding efficiency systematic, limited by the statistics of a

hand-scan of 12000 events [68]. Also contributing are errors of 0.7% on events migrating in/out of the near detector fiducial volume, 0.32% on the far detector live time, 0.2% on the relative near/far scintillator mass, and 0.2% on the relative near/far steel mass. The mass of detector components does not affect rock events, and so for these events the total error is slightly smaller.

12.1.2 Muon Track Energy

An error of 2% is assigned for track energy measured by range. This is the result of studies that found the detector mass model to be good to 1%, the Monte Carlo truth/reconstruction agreement to 1%, and the agreement between the energy loss model used in software agreed with published range tables to 1% [69]. Summed in quadrature this gives 1.7%. The group responsible for this systematic has used several methods to constrain the track energy and found tension between them at the 2% level [70]. Therefore they take 2% as their final number.

The curvature systematic is derived from a study comparing the energy by range and that from curvature for stopping tracks. The output of this study a measure of the uncertainty on the magnetic field strength, which was found to be 2.8% [71]. This is related to the error on track energy by curvature *relative to* energy by range. The error depends on the containment fraction and the number of planes used in the measurement as follows [72]:

Suppose a muon has true momentum p_0 at the first plane in a reconstructed track. This track consists of N planes and a constant momentum Δp is lost at each plane. At each plane, the measured curvature is an estimate of the momentum $p_0 - n\Delta p$, where n is the number of planes crossed so far. Given Δp , we can add $n\Delta p$ to this to estimate the momentum at the track vertex. Ideally, the mean of these measurements is, on average, the true energy:

$$\langle p_{\text{curvature}} \rangle = \frac{1}{N} \sum_{n=1}^N [(p_0 - n\Delta p)_{\text{curvature}} + (n\Delta p)_{\text{range}}] = \frac{1}{N} \sum_{n=1}^N p_0 = p_0 \quad (12.1)$$

But if the magnetic field used in the estimate is wrong by a factor $1 + \sigma_B$, instead we obtain

$$\langle p_{\text{curvature}} \rangle = \frac{1}{N} \sum_{n=1}^N [(1 + \sigma_B)(p_0 - n\Delta p) + n\Delta p] \quad (12.2)$$

$$= \frac{1}{N} \sum_{n=1}^N [(1 + \sigma_B)p_0 - n\sigma_B\Delta p] \quad (12.3)$$

$$= (1 + \sigma_B)p_0 - \frac{\sigma_B}{2}(N + 1)\Delta p \quad (12.4)$$

And so the fractional error is

$$\sigma_{p_{\text{curvature}}} = \langle p_{\text{curvature}} \rangle / p_0 - 1 = \sigma_B \left[1 - \frac{1}{2}(N + 1) \frac{\Delta p}{p_0} \right] \quad (12.5)$$

$$= \sigma_B \left(1 - \frac{N + 1}{2N} f \right), \quad (12.6)$$

where f is the fraction of the track that is contained, ideally p_{range}/p_0 . Conceptually what is happening is that as a track loses energy, each subsequent momentum measurement from curvature depends less on the field (first term of equation 12.1) and more on the range (second term). As shown in figure 12.1, a typical error is 2%. Since the field was calibrated using track momentum by range, this error must be added in quadrature with the range error. A muon that stops right at the detector edge, and is therefore labeled as exiting, has the smallest possible error of about 2.5%. Truly exiting tracks have errors evenly distributed between this and the maximum value of about 3.6%.

The track energy systematic has a direct effect on the value of Δm^2 . Because the typical event has somewhat more than half of its energy in the muon tracks, it effectively sets a $\sim 1.5\%$ scale error on all energy measurements. There is no near/far cancellation, and since Δm^2 is measured by observing the energy of maximum disappearance, it contributes at first order to the error on Δm^2 .

For detector interactions, this is the whole story. However, for rock interactions, there is an additional first order effect. If we under-estimate the energy of tracks in the near detector, we infer a lower energy neutrino spectrum. Therefore, we predict that less rock is within muon range of the far detector, since the amount of relevant rock is roughly proportional to neutrino energy. Therefore not only do we predict a shifted spectrum of muons, but also a




Figure 12.1: Systematic error on track energy measured by curvature due to a 3% systematic error on the magnetic field strength. Left: Error relative to the systematic error on track energy measured by range. Right: Total systematic error on curvature measurements. The dashed line shows the possible situations for a typical 2 GeV track. (The track angle modifies the energy lost per plane; this is the average for beam events.)

reduced normalization of that spectrum. Under-estimating the track energy decreases our estimate of Δm^2 both because we reconstruct the dip lower than it should be and because we simulate too few rock muons, thereby making the deficit in data look smaller.

To estimate the size of the effect for this analysis, take the systematic to be 2.5% on all tracks. Detector events are affected only by the direct energy scale and so have a 2.5% error. Rock events have a 2.5% direct energy scale error and an additional 2.5% normalization error. From the evaluation of the normalization systematic (see below), we know that a 1% normalization error results in a 2.7% error on Δm^2 . So this gives an additional 6.8% error for a total of $\sim 9\%$. Rock events have about 75% the statistical power of anti-fiducial detector events (see figure 10.7), so taking the weighted average, we would expect about a 5% total error from both categories. The full treatment gives 4.3%.

12.1.3 Subdominant Systematics

Each of these systematics has less than half the effect on Δm^2 as either of the above two. In a simple quadrature sum of systematic effects, all of these combined would add only 12% over the two dominant systematics.

- Neutral current background: 19% scale on neutral current events. This represents the combination of the error on the neutral current cross section and hadronic shower modeling. The former changes the total number of events, whereas the latter causes

Parameter(s)	Default	Systematic 1	Systematic 2	Systematic 3
112	0.1	± 0.1		
122	0.3	± 0.1		
132	0.3	± 0.1		± 0.2
142	0.1	± 0.1		± 0.2
212	0.1	± 0.1		
222	0.3	± 0.1		
232	0.3	± 0.1		
242	0.1	± 0.1		
113, 123, 133, 143	1.0		-0.2	
213, 223, 233, 243	1.0		-0.2	

Table 12.1: Defaults and systematics for NEUGEN parameters r_{ijk} . Three independent systematics are used: (1) multiplicity 2 events (2) multiplicity 3 events, and (3) charged current $\bar{\nu}$ multiplicity 2 events. See text for details.

events to migrate across the PID cut; both effects are modeled with a simple scaling of selected events.

- Absolute shower energy scale: Fully near/far correlated scaling of the reconstructed shower energy by $(6.6 + 3.5 \exp(-E/1.44 \text{ GeV}))\%$, where E is the shower energy. This error is based on calibration detector measurements and rests primarily on two considerations: (1) the level of agreement between data and Monte Carlo for pion energy measurement after tuning the Monte carlo parameters to match the data as well as possible and (2) the uncertainty in the energy of the test beam [42].
- Relative detector shower energy scale: Independent 1.9% (1.1%) scalings on shower energy in the near (far) detector. For this analysis, only the near detector error matters since far detector showers are not used.
- Total charged current cross section: 3.5% scale on the total charged current cross section. An additional independent 4% scale is applied to antineutrino events.
- Three systematics based on varying NEUGEN parameters:
 - Quasi-elastic cross section: 15% overall scale on the axial mass M_A (central value 0.99 GeV) and an independent 8% scale on antineutrino events. This results

in both a change to the normalization and the shape of the quasi-elastic cross section.

- Resonance cross section: As above, a 15% overall scale on the axial mass M_A (central value 1.12 GeV) and an independent 8% scale on antineutrino events.
- Koba-Nielsen-Olesen multiplicities (see section 6.2). Three independent sets of shifts are made; within each set all parameters are shifted in concert (also shown in table 12.1):
 - * ± 0.1 shifts on all multiplicity 2 events (r_{ij2})
 - * -0.2 shift on all multiplicity 3 events (r_{ij3}). The shift is one-sided because the parameters are by default set to their maximum value of 1.0.
 - * ± 0.2 shifts on multiplicity 2 charged current $\bar{\nu}$ events (r_{132}, r_{142}).

Multiplicities above 3 are not considered since they make no significant contribution in any case.

- Flux: As a function of energy, the error on the initial reweighting of the Monte Carlo (see section 9.1). The error is separately defined for each neutrino species and displayed in figure 12.2.

12.2 Systematics Unique to this Analysis

Three additional source of error unique to this analysis were considered: The error on the cross section due to the differing elemental composition of the rock as compared to the detectors, the normalization error due to the rock’s average Z/A , and the error due to alignment uncertainties in selecting events as rock-like or detector-like on the radial edge.

12.2.1 Cross Sections

The rock around the far detector has a markedly different composition than the MINOS detectors. While the detectors are mainly iron, the rock is mainly oxygen and silicon (see table 12.2). Because of this, cross section uncertainties do not cancel to the same degree




Figure 12.2: Error on the flux estimate. This error is fully correlated between the two detectors; like other flux and/or cross section errors, it cancels at first order. In order of increasing error, the thick blue line is ν_μ , the dotted red $\bar{\nu}_\mu$, the dot-dashed black ν_e , and the dashed purple $\bar{\nu}_e$. These errors are valid for all three LE runs. Note the peak in the ν_μ error around 5 GeV; this is the falling edge of the focusing peak.

	Steel	Scintillator	Far detector	Far rock
O	—	—	—	$(45.86 \pm 0.53)\%$
Si	0.29%	—	0.27%	$(23.7 \pm 0.16)\%$
Fe	98.53%	—	94.05%	$(8.56 \pm 0.35)\%$
Al	—	—	—	$(7.97 \pm 0.21)\%$
Ca	—	—	—	$(6.41 \pm 0.30)\%$
Mg	—	—	—	$(3.95 \pm 0.23)\%$
Na	—	—	—	$(1.90 \pm 0.12)\%$
Ti	—	—	—	$(0.64 \pm 0.05)\%$
K	—	—	—	$(0.34 \pm 0.09)\%$
H	—	14.29%	0.65%	$(0.26 \pm 0.04)\%$
C	0.18%	85.71%	4.07%	$(0.20 \pm 0.11)\%$
Mn	1.00%	—	0.95%	$(0.15 \pm 0.01)\%$
P	0.01%	—	0.01%	$(0.06 \pm 0.01)\%$

Table 12.2: Elemental compositions of the detector and the surrounding rock. The far detector is 95.5% steel and 4.5% scintillator by mass, with negligible contributions from other materials. No errors on the makeup of the detector are given, both because their effects cancel at first order and because they are small compared to the errors on the rock composition.




Figure 12.3: Predicted rock events at the far detector without oscillations. The upper histogram includes all event types; the lower, shaded histogram shows only non-deep inelastic scattering events. In the region around the oscillation maximum, ~ 1.5 GeV, deep inelastic scattering events are the minority.

for rock events as for detector events. Instead, there is a systematic on the ratio of cross sections $\langle\sigma_{\text{detector}}\rangle/\langle\sigma_{\text{rock}}\rangle \approx \sigma_{Fe}/\langle\sigma_{\text{rock}}\rangle$.

At high energy, nearly all events are due to deep inelastic scattering. The cross section for this type of interaction is nearly independent of nuclear effects and so we can be confident that $\sigma_{Fe}/\langle\sigma_{\text{rock}}\rangle$ is approximated well enough by the Monte Carlo. At lower energy, quasielastic and resonance production events dominate (see figure 12.3). The cross sections of these events are less well understood. While measurements in this energy range exist, they are primarily from bubble chambers in which the target material is very different from the far detector rock: hydrogen, deuterium [73, 74, 75, 76], neon [77], or hydrocarbon/halogenated hydrocarbon mixtures [78, 79, 80].

There are a few more recent measurements that come closer. IHEP-JINR* used a target of aluminum and liquid scintillator for neutrinos in the range 3–30 GeV [81]. NOMAD† used primarily carbon (64%) and oxygen (22%) for neutrinos in the range 2.5–40 GeV. This comes very close to the case of interest. However, their lowest energy bin spanned a large range, 2.5–6.0 GeV and in this bin, their total error is $\pm 4.7\%$ [82]. So the needed precision is not there. The currently running MINERνA experiment will greatly improve this situation [35], but for now, we must depend to a significant degree on theoretical models to provide cross sections in this range.

*Institute for High Energy Physics-Joint Institute for Nuclear Research

†Neutrino Oscillation Magnetic Detector

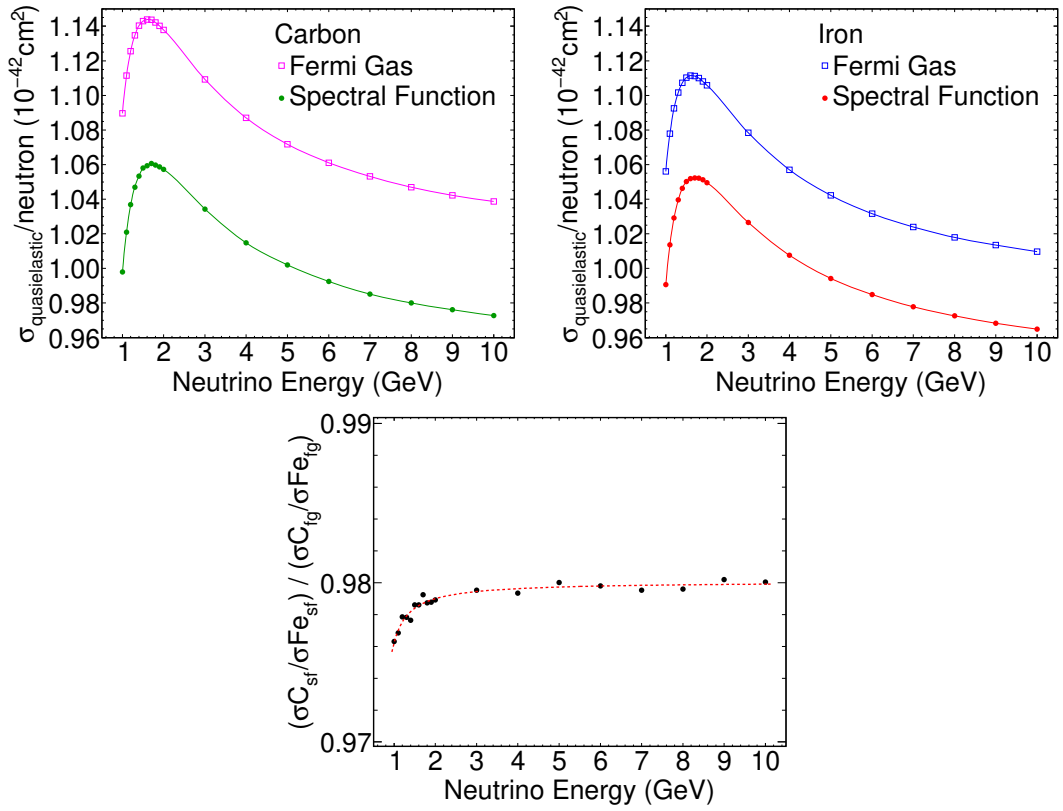


Figure 12.4: Top: Spectral function model of quasi-elastic cross sections for carbon and iron as compared to the Fermi gas model. The Fermi gas carbon (iron) model used a Fermi momentum of 220 (260) MeV and a binding energy of 34 (36) MeV [83]. Bottom: Double ratio constructed from the above cross sections with the raw model data as circles and the fit (equation 12.8) as a dashed line.

The MINOS Monte Carlo uses the relativistic Fermi gas model to generate cross sections [50]. However, this model does not account for nuclear effects except in a rudimentary way. The newer *spectral function* model improves upon this [84], but is not available within NEUGEN and so we must apply corrections after the fact. For interactions in the detector, data/Monte Carlo differences in cross section largely cancel since the near detector and far detector are made of the same materials. However, if the cross sections in the rock are mismodeled, the data/Monte Carlo comparison will be wrong by the double ratio

$$R_{\text{rock, detector}} \equiv \frac{\langle \sigma_{\text{rock data}} \rangle / \langle \sigma_{\text{detector data}} \rangle}{\langle \sigma_{\text{rock MC}} \rangle / \langle \sigma_{\text{detector MC}} \rangle}. \quad (12.7)$$

Spectral function models predict $\sim 10\%$ lower quasielastic and resonance production cross sections as compared to the Fermi gas model. Since it is lower for both detector elements and rock elements, most of the difference cancels. However, the spectral function model differs more strongly for smaller nuclei, and so the double ratio above is not 1. Using model predictions [83, 85] for ^{12}C , I found the double ratio as a function of neutrino energy to be well fit by:[‡]

$$R_{\text{C, Fe}} = 0.9801 - 0.001489/(E_\nu - 0.6096). \quad (12.8)$$

This is very close to 0.98 for all relevant energies (see figure 12.4). Approximating the detector as iron and using the assumption that R scales linearly towards unity as a function of Z , the Monte Carlo is reweighted accordingly. Since R is not sensible below ~ 0.61 GeV, for $E_\nu < 0.75$ GeV, the weight at 0.75 GeV is used. Besides preventing negative event weights, this has no practical effect since only 0.002% of rock events satisfy this condition. In the end, the total effect of this reweighting is to reduce the normalization of the affected rock events by an average of 1.37% for LE beam and somewhat less for pHE.

Since the spectral function model is not experimentally tested for the relevant interactions, and because of the *ad hoc* nature of the above correction, a 1% normalization systematic is applied to all events that use the correction.

[‡]It is not meant to be implied that these numbers are good to four digits. These are simply the central values used in the code.

Sample	1	2	3	4	5	6	7	8	9	10	11	12	13	14
SiO ₂	48.15	49.34	49.55	49.65	49.72	50.14	50.47	50.90	51.06	51.20	51.30	51.73	52.94	51.95
TiO ₂	0.84	0.79	0.98	0.78	0.89	1.59	—	1.33	1.44	1.47	0.76	0.78	—	1.03
Al ₂ O ₃	13.80	14.68	16.05	15.44	16.76	13.54	18.48	16.05	13.85	14.91	14.29	15.28	14.70	12.58
Fe ₂ O ₃	2.34	4.42	2.45	3.22	1.92	2.77	2.13	2.33	1.79	2.18	4.46	3.41	2.52	0.90
FeO	10.20	8.04	8.32	7.96	7.33	11.81	7.74	8.66	10.86	9.16	6.76	7.30	7.80	8.77
MnO	0.22	0.22	0.22	0.18	0.16	0.30	—	0.18	0.22	0.25	0.19	0.15	—	0.15
MgO	8.59	6.78	4.66	7.42	7.62	5.07	6.90	6.15	5.31	5.45	7.53	6.72	4.49	8.90
CaO	8.97	10.42	11.94	9.56	9.35	8.25	6.61	9.30	10.66	7.60	9.54	9.40	6.56	7.00
Na ₂ O	2.70	1.90	2.04	1.64	3.14	2.95	2.58	2.06	2.49	2.78	1.69	3.83	3.09	2.79
K ₂ O	0.10	0.12	0.62	0.32	0.71	0.25	0.30	0.76	0.25	0.00	0.16	0.76	0.04	1.38
H ₂ O	3.81	2.80	2.34	3.02	1.57	2.95	0.00	1.74	1.83	3.93	2.83	2.86	0.00	2.81
CO ₂	0.00	0.17	1.39	0.05	0.10	0.06	—	0.00	0.33	0.74	0.00	—	4.86	1.02
P ₂ O ₅	0.16	0.08	0.09	0.13	0.09	0.24	—	0.20	0.13	0.14	0.12	—	—	0.00

Table 12.3: Mineral compositions (in percent by mass) of 14 samples of Ely greenstone [54, 55]. From these numbers, the elemental composition in table 12.2 is derived. No errors are published on a sample-by-sample basis, but were extracted from the spread of the samples. Blank entries in this table were blank in the references.

12.2.2 Z/A

Our value for the far detector rock's average Z/A is derived from mineral composition measurements of samples of Ely greenstone taken elsewhere in the region (see table 12.3). As discussed in section 7.3.2, the rock's Z/A has a first order effect on the muon flux since increasing the average A increases the interaction cross section proportionally (as a good approximation) while increasing the average Z reduces the range of charged particles also proportionally. The measurements tell us that Ely greenstone's Z/A is 0.498 ± 0.004 , or 0.81%, where the error is set by the number of samples. Since none of these samples are from in or around the far detector cavern itself, I have arbitrarily added an additional 10% to the error to account for variations across the geographical region, bringing $\sigma_{Z/A}$ to 0.89%. This is implemented as a simple normalization shift of rock events.

12.2.3 Alignment

A significant portion of the statistical power of this analysis comes from the separation of rock events from detector events on the radial edge of the far detector, as described in section 10.2.1. This strip-based separation is vulnerable to uncertainties in the detector alignment which could cause events to migrate in or out of the region considered rock-like. These alignment uncertainties include non-ideal transverse and longitudinal strip positions, non-ideal strip lengths, and strip rotations, both within and out of the xy plane.

The Monte Carlo model of the detector uses perfect alignment for all planes. This is sufficient for the diagonal edges, where events are selected as rock-like if the first hit in the view in which the strips are parallel to the edge is in one of the first two strips. In this case, the strip length and longitudinal alignment are irrelevant, since the action is occurring away from the end of the strip. On the other hand, transverse alignment could have an effect. Suppose the strips of a given plane were placed farther out than the planes upstream of it. This plane is more likely to be first hit by a rock muon, and it is more likely that the hit will be on the second or third strip of this plane, since they are less shielded than if all alignments were perfect.

While these transverse positions vary by up to 40 mm in the detector as a whole with an




Figure 12.5: Transverse alignment of outermost strips. Top: Raw alignments of every plane's strip 0. A higher value means that the strips are farther out from the center of the detector in reality than in the Monte Carlo. Left: Histogram of the alignments above. Right: Histogram of *relative* alignments between adjacent planes. Note the decrease of width due to correlation in the position of adjacent planes evident in the top figure.

RMS of 7.9 mm, nearby ones have correlated positions. The RMS of the difference between adjacent ones is only 6.7 mm; see figure 12.5. Since the strip width is 41.1 mm and the spacing between planes with parallel strips is 118.8 mm, this would appear likely to make a very small difference. Strip rotations in the xy plane could have a similar effect (in opposite directions at the opposite ends of a rotated strip), but these rotations are small compared to the transverse alignment differences. I used a small sample of Monte Carlo that incorporates surveyed strip transverse alignments to check whether there is a significant change in the distribution of first parallel strip hits. None is apparent and so no correction is made nor systematic assigned to this selection.

On the horizontal and vertical edges, the situation is more complex. Here, since the selection uses strip ends as well as using two (usually adjacent) planes to construct the discrimination variable H (see section 10.2.1), the strip transverse alignment, longitudinal alignment and length all can directly change this number. The relevant quantity for transverse alignment is the variation (compared to the Monte Carlo model) in locations of nearby strips in adjacent planes; this is 7.1 mm (RMS). The longitudinal positions of the strips are unmeasured; there are not enough cosmics to map them (although there are enough to tell that they are not *grossly* misaligned). It is estimated that the variation is 5 mm. The strips were specified to be between 7996 mm and 8000 mm long. While it is not recorded what the lengths actually are, we have no reason to believe the spec was not met, so the strip half-length uncertainty is ~ 1 mm. Added in quadrature, the total RMS variation is 10.1 mm.

Therefore, an after-the-fact correction is applied to the Monte Carlo to account for reality being less ideal in this fashion. The Monte Carlo output does not contain enough information to determine on an event-by-event basis when a hit should be added, removed or shifted to another strip. However, the correct distribution is obtained by randomly modifying H as is appropriate for a 10.1 mm RMS total shift. First a uniform random number in the range 0–41 mm is chosen to represent where transversely in the strip the hit occurred. A Gaussian random number with $\sigma = 10.1$ mm is chosen to determine the movement of the hit. If the hit is moved outside the strip, H is incremented or decremented appropriately.

Because the 5 mm variation in longitudinal alignment is a rough estimate, a systematic is applied in which this is varied on the range 2.5–10 mm such that a lesser or greater number

of Monte Carlo events are moved between the rock-like and detector-like samples. This turns out to be a very small systematic, $\pm 0.0015 \text{ eV}^2$ assuming maximal mixing, which has no significant contribution to the total error.

12.2.4 Correlations Between Systematics

In order to return correct results using the method described in this chapter, the systematics used in the fit must be uncorrelated. Let us look at where each comes from. First the two dominant errors:

- Normalization (section 12.1.1): Several sources, but by far the most important is the data/Monte Carlo hand-scan study. Includes a tiny contribution from the detector mass model.
- Track energy systematic (section 12.1.2): Equal contributions from: (1) The detector mass model, (2) Monte Carlo/data agreement studies, and (3) agreement with range tables. Has some contribution from calibration of the magnetic field strength as determined by range/curvature agreement.

While both are somewhat dependent on the Monte Carlo model, they depend on different aspects of it. The normalization is dependent on the modeling of track finding efficiency, which is a function of the detector efficiency and the reconstruction algorithm. The track energy is dependent on the model of muon range and is largely independent of both the detector response model and the reconstruction code. The detector mass model comes into both effects, but for the track energy systematic, the relevant quantity is the *absolute* mass of the detector components, while for the normalization systematic, it is the *relative* near/far mass. These should be completely independent to any relevant level of precisions. It therefore seems reasonable to treat these systematics as uncorrelated.

Two of the subdominant systematics have related sources (see section 12.1.3):

- Neutral current background: Cross sections and hadronic shower modeling.
- Absolute shower energy: Predominantly hadronic shower modeling and test beam energy calibration.

Both are dependent on hadronic shower modeling, but it is justifiable to nevertheless treat them as uncorrelated. First, neither systematic rests entirely on our shower modeling — both include significant components not shared by any other systematic used in the fit. Second, they depend on different aspects of the shower modeling. The neutral current background systematic is dependent on how often a string of hits in a shower will resemble a muon track, and so is primarily affected by modeling of shower shape. The shower energy systematic is, on the other hand, primarily dependent on the model of energy deposition by pions. These are no doubt linked to some extent, but they have no one-to-one connection. Third, each of these systematics taken alone has less than half the effect on Δm^2 than either of the two dominant systematics, as judged by their 1σ shift effects (see table 12.4). Naively summing errors in quadrature therefore suggests that no matter how we treat them, they will be a roughly a 10–15% effect on the total systematic error, so extreme precision is not needed.

Of the other systematics, many are concerned with cross sections. These are certainly correlated to some degree. For instance, if we are underestimating the quasi-elastic cross section, it is probably either the case that we are also underestimating the resonance cross section — because the same small number of experiments measured both of them — or that we are overestimating the resonance cross section given that their sum must fit within the limits of the total cross section. A quadrature sum of all cross section systematics (see table 12.4) gives $\pm 0.037 \times 10^{-3} \text{ eV}^2$, which when added in quadrature with the statistical error of $0.18 \times 10^{-3} \text{ eV}^2$ gives $0.184 \times 10^{-3} \text{ eV}^2$ — a very small effect. A fully rigorous treatment of the correlations would find a value either somewhat larger or smaller than this. In the absolute worst case of full correlation requiring a linear sum of all cross section systematics, we would get $\pm 0.080 \times 10^{-3} \text{ eV}^2$ which added to the statistical error would inflate $0.18 \times 10^{-3} \text{ eV}^2$ to $0.20 \times 10^{-3} \text{ eV}^2$. Clearly this is not at all realistic, but the relatively modest impact even at this extreme shows that further work to pin down the exact correlations is not needed.

Systematic	Δm^2 shift (10^{-3} eV^2)	
	Full	Simplified
Muon track energy	+0.093 -0.093	+0.102 -0.103
Normalization	+0.090 -0.089	+0.093 -0.093
Absolute E_{shower}	+0.043 -0.042	+0.048 -0.047
Near detector relative E_{shower}	+0.0097 -0.0098	+0.011 -0.011
Neutral current background	+0.037 -0.037	+0.036 -0.036
Total CC cross section	+0.0061 -0.0066	+0.0062 -0.0067
Quasi-elastic cross section	+0.023 -0.026	+0.021 -0.023
Resonance cross section	+0.021 -0.030	+0.020 -0.027
DIS multiplicity 2 cross section	+0.016 -0.018	+0.016 -0.017
DIS multiplicity 3 cross section	-0.0029	-0.0028
$\bar{\nu}$ CC cross section	+0.0022 -0.0022	+0.0025 -0.0025
$\bar{\nu}$ quasielastic cross section	+0.00003 -0.00003	+0.00002 -0.00001
$\bar{\nu}$ resonance cross section	+0.00001 -0.00001	+0.00002 -0.00002
$\bar{\nu}$ DIS multiplicity 2 cross section	+0.0014 -0.0014	+0.0015 -0.0015
Flux	+0.021 -0.023	+0.022 -0.023
Rock Z/A	+0.030 -0.030	+0.032 -0.031
Rock non-DIS cross section	+0.0076 -0.0075	+0.0084 -0.0084
Edge alignment	+0.0015 -0.0015	+0.0017 -0.0017
Quadrature Sum	+0.151 -0.152	+0.160 -0.162

Table 12.4: Effects of all systematic errors considered in this analysis. Errors have been evaluated for both the full analysis and the simplified energy-only version; each is evaluated at the respective best fit. See also figures 12.6 and 12.7. Note dominance of the top two errors. The sum of all errors in quadrature is given for illustration, although this is overly conservative (see section 12.3.1).

12.3 Evaluation of Systematics

To judge the effect of individual systematics, each is varied by $\pm 1\sigma$ and the effect on the best fit point observed. The results of this procedure are shown in table 12.4 and figures 12.6 and 12.7. Those systematics that affect both the near and far detectors are evaluated by performing the entire extrapolation and fitting procedure with the shifts applied as appropriate. Other systematics can be evaluated by considering the far detector alone.

For the first category, consider for instance the evaluation of the quasielastic charged current cross section systematic error. The near detector data is converted to a neutrino flux using the higher (lower) value of this cross section. A smaller (larger) flux is produced in each energy bin in order to account for the observed number of events. This flux is extrapolated to the far detector and used to generate the predicted spectrum of events. (For practical reasons, the existing Monte Carlo is reweighted rather than generating a new set for each systematic.) This prediction again uses the shifted cross section and so it uses a larger (smaller) cross section for fewer (more) neutrinos. This produces nearly the same spectrum of events except for two effects. First, for a higher (lower) quasielastic cross section, the energy balance in simulated far detector events will be shifted towards tracks (showers).[§] This will modify efficiencies and, for rock events in particular, may significantly change the event rate itself by allowing more (fewer) events to reach the detector. Second, there is not a one-to-one correspondence between energies at the near and far detector. Each energy bin in the near detector is extrapolated to a generally asymmetric distribution in the far detector (see figure 9.2). This smearing introduces an additional, second order, shift in the spectrum.

The systematics that do not require the extrapolation to be repeated are the near/far normalization (which could alternatively be applied at the near detector only) and the three systematics in section 12.2. In the fiducial analysis, the shower energy reconstruction at the far detector also falls into this category. These systematics have no effect on the neutrino flux prediction and can be applied directly to the far detector Monte Carlo.

[§]There are alternative methods of running the extrapolation in which the measured elasticity is itself extrapolated, but the standard method used here discards such details of near detector events.




Figure 12.6: Summary of systematic errors, evaluated at $\Delta m^2 = 2.16 \times 10^{-3} \text{eV}^2$, $\sin^2 2\theta = 1$ with $\sin^2 2\theta$ held constant. See also table 12.4.

12.3.1 Incorporation of Systematics into Fit

The obvious method for determining the total systematic error is to add the results of the 1σ shifts above in quadrature. This results in a very conservative number because in general the data itself can carry some information about the values of the systematics. In order to take advantage of this, one or more systematics can be included in the fit using the method of nuisance parameters with penalty terms.[¶]

12.3.2 Nuisance Parameter Fits

For a nuisance parameter fit, the quantity

$$\chi^2 = \sum_{j=1}^M \frac{(s_j - s_{j_0})^2}{\sigma_j^2} + 2 \sum_{i=1}^N \left[m_i - d_i + d_i \log_e \frac{d_i}{m_i} \right] \quad (12.9)$$

[¶]This approach has been used by previous MINOS analyses such as in the 2006 and 2008 ν_μ -disappearance results [28, 86]. The method of including systematics in the fit is covered in [22, 87] and the use of penalty terms in [88, 89]. A very nice discussion is also to be found in [90].




Figure 12.7: Summary of systematic errors for the simplified analysis, evaluated at $\Delta m^2 = 2.07 \times 10^{-3} \text{ eV}^2$, $\sin^2 2\theta = 1$ with $\sin^2 2\theta$ held constant. See also table 12.4 Note that the cross section errors are a bit smaller than for the main analysis; without angle binning, the fraction of events due to each interaction type is less important. However, the two dominant errors are larger since the fit rests entirely on normalization and track energy.

is minimized, where, in the first term, M is the number of systematics, s_j is the value of the j th systematic parameter, s_{j_0} is that parameter's nominal value, and σ_j is its previously established error. The parameters s_j are called *nuisance parameters* or sometimes *constraint terms*. The second term of 12.9 is the same as in the statistics-only fit, equation 10.1, except that now the Monte Carlo prediction in each bin is a function not only of the physical parameters of interest, but also each of the systematic parameters s_j .

For this analysis, the fit will be performed with $\sin^2 2\theta$ fixed at 1 and only a single nuisance parameter, the normalization, and so the parameter space has 2 dimensions. The 2010 MINOS result, which will include this work, will fit over Δm^2 , $\sin^2 2\theta$, and four systematics. Therefore I will discuss the procedure in full generality, where χ^2 is a function of P physical parameters (p_1, \dots, p_P) and M systematic parameters (s_1, \dots, s_M) .

This fit procedure is equivalent to adding the χ^2 surfaces of $1 + M$ independent experiments, where M of these established the nominal values s_{j_0} and errors σ_j of each systematic parameter and the last is the main measurement. The systematics' χ^2 surfaces constrain only one parameter, s_j , and are flat in all others. The surface from the main measurement simultaneously measures all $P + M$ parameters, although in general provides only comparatively weak constraints on the values of the systematics. The power of the nuisance parameter method is that it incorporates both sources of information and automatically handles correlations in the effects of the systematics. In order for this to work, the systematics themselves have to be uncorrelated: see section 12.2.4.

This method produces a best fit in the space of all parameters $(p_1, \dots, p_P, s_1, \dots, s_M)$, in which the values of the physical parameters p are, in general, distinct from the statistics-only result (where all s_j are fixed at their nominal values). This result treats all $P + M$ parameters on an equal footing, so a $P + M$ -dimensional contour can be generated showing the correlations between all parameters. In practice, this is only presentable if $P + M$ is 2 or perhaps 3. To present contours in only the parameters of interest, say $(\Delta m^2, \sin^2 2\theta)$, we can project the boundary of the full contour onto the $(\Delta m^2, \sin^2 2\theta)$ plane using the appropriate 2-parameter $\Delta\chi^2$ value for the desired confidence level (e.g. $\Delta\chi^2 = 4.61$ gives the 90% contour). This is analogous to the method of extracting a single parameter confidence band from a 2-parameter χ^2 surface [22]. In practice, this is accomplished by scanning through $(\Delta m^2, \sin^2 2\theta)$ -space and at each point minimizing χ^2 as a function of the systematic parameters. The resulting 2-parameter χ^2 surface is the final result. The same can be done, of course, for only one parameter, as will be shown below.

While the value of each systematic at the best fit point is not fundamentally interesting, we expect that each is near its nominal value. In particular, if the far detector data provided *no* additional information about a given parameter, the fit should return exactly its nominal value. It would be surprising to find that the best fit of any systematic parameter is as much as σ_j away from nominal as this would mean that the far detector data is calibrating it better than the intentional calibration. Since this is unlikely, any such finding should be closely examined for mistakes. For the same reason, we would not expect the best fit point in the physical parameters to be substantially modified from the statistics-only fit.




Figure 12.8: Nuisance parameter fit using only normalization and Δm^2 . The square shows the statistics-only best fit. The shorter error bars are statistics-only; the longer ones have the 1σ normalization shifts added in quadrature. The solid contour shows $\Delta\chi^2 = 1$ for the nuisance parameter fit; the star is the best fit point. The horizontal tangents give the 1σ statistical+systematic errors for Δm^2 . Also shown is the result of dropping the penalty term. The inner contour is $\Delta\chi^2 = 1$ for comparison with the penalized contour. The outer contour is the 2-dimensional 68% confidence level. Encouragingly, this encloses both best fit points. See text for more details.

By using the data as an additional constraint on the values of the systematics, the resulting errors in the physical parameters are smaller than would be found if the systematics were simply added in quadrature. In addition, this procedure handles correlations between the effects of the systematics automatically. For instance, the track and shower energy systematics both affect the flux estimation in roughly the same way. This means they can float in opposite directions with only a small effect on the predicted spectrum of events at the far detector, but if they move in the same direction, the total energy will be changed and data will tend to disfavor the resulting prediction. One can check that the best fit point follows this pattern as expected.

12.3.3 Fit for This Analysis

The use of several systematics in the fit is technically complex and is the subject of ongoing work in the MINOS collaboration. For most systematics, the shift must be simultaneously applied at both detectors using the strategy discussed above for 1σ shifts. This method

requires that all the near detector data be read and reinterpreted to account for the systematic shift each time a set of systematic parameters is chosen. It therefore takes on order of an hour to produce one shifted prediction. Clearly this is impractically slow for finding the minimum χ^2 in a many-parameter space. A technique that achieves a good approximation with a reasonable amount of computer time is under development [91, 92]. For this analysis, I will use a single systematic which is trivial to apply and need be applied only to the far detector, the overall near/far normalization uncertainty. Subdominant systematics are then conservatively added in quadrature to the result of the nuisance parameter fit.

Since this fit is performed in only two dimensions, Δm^2 and the normalization, it is easy to present: see figure 12.8. To illustrate the power of the nuisance parameter method, the results of the naive method — fixing the normalization at its nominal value and adding the effects of 1σ shifts in quadrature — are also shown. This gives (with all other systematics ignored for the moment):

$$\Delta m^2 = (2.16 \pm 0.18[\text{stat}] \pm 0.09[\text{syst}]) \times 10^{-3} \text{ eV}^2 = (2.16 \pm 0.20[\text{total}]) \times 10^{-3} \text{ eV}^2.$$

The nuisance parameter fit shifts Δm^2 up slightly and results in a somewhat smaller statistical+systematic error:

$$\Delta m^2 = (2.20 \pm 0.19) \times 10^{-3} \text{ eV}^2.$$

We can loosely define the effect of the normalization systematic as $\sqrt{\text{total}^2 - \text{stat}^2}$, giving $0.065 \times 10^{-3} \text{ eV}^2$ in this case. This is an improvement over the $0.09 \times 10^{-3} \text{ eV}^2$ found via the 1σ shift method. If we wish, we can also read off the best fit and errors on the normalization: It is $(0.80 \pm 1.38)\%$ above nominal. Its error is slightly improved compared to the input error of 1.53%. To see where this comes from, the contours are shown that we would obtain if we had no other knowledge of the normalization and let it float freely in the fit: Δm^2 is pulled up considerably and has larger total errors. From this plot alone, the normalization systematic is restricted to being $4.4\% \pm 3.1\%$ over nominal (the horizontal limits of the $\Delta\chi^2 = 1$ contour). The best fit for the normalization is the weighted average of this measurement and the *a priori* $0 \pm 1.53\%$. Note that even if we arrived here with no previous knowledge of the normalization — effectively setting the normalization systematic

to infinity — the total error on Δm^2 would only be inflated by $\sim 25\%$.

12.4 Conclusions

The quadrature sum of all non-normalization systematics is $\pm 0.12 \times 10^{-3} \text{ eV}^2$. Adding the $0.065 \times 10^{-3} \text{ eV}^2$ from above gives a total systematic of $0.14 \times 10^{-3} \text{ eV}^2$. So the final result, assuming $\sin^2 2\theta = 1$, is:

$$\begin{aligned}\Delta m^2 &= (2.20 \pm 0.19[\text{stat} + \text{normalization syst}] \pm 0.12[\text{other syst}]) \times 10^{-3} \text{ eV}^2 \\ &= (2.20 \pm 0.18[\text{stat}] \pm 0.14[\text{syst}]) \times 10^{-3} \text{ eV}^2 \\ &= (2.20 \pm 0.23[\text{total}]) \times 10^{-3} \text{ eV}^2\end{aligned}$$

Chapter 13

Summary, Conclusions, and Future

The sample of beam-produced muon neutrino interactions in the rock surrounding the MINOS far detector and in the detector's anti-fiducial volume have been analyzed. This analysis has produced a statistically independent measurement of the oscillation parameters associated with muon neutrino disappearance. The value of $\sin^2 2\theta$, found to be greater than 0.56 at 90% confidence, is consistent with previous results. Δm^2 was found to be $(2.20 \pm 0.18[\text{stat}] \pm 0.14[\text{syst}]) \times 10^{-3} \text{ eV}^2$. This is also in good agreement with established results, and a strong enough measurement that when combined with the MINOS fiducial




Figure 13.1: Comparison of this result with the final K2K result [17], the latest Super-Kamiokande results [93], the 2010 MINOS fiducial result [43]. The fiducial+anti-fiducial combined result is also shown. For MINOS, $\sin^2 2\theta = 1$ is assumed; others are shown as reported.

result, the overall errors on Δm^2 are improved by 4% (see figure 13.1). This analysis benefited from a detailed and precise understanding of the rock and edges of the far detector. Statistical power was gained through separation of events into rock-like and detector-like, further separation by detector region, and binning by both muon energy and angle.

13.1 Potential to Increase Statistical Power

This analysis already breaks the sample of anti-fiducial muons into many subsamples, exploiting the varying backgrounds, resolutions and energies to gain sensitivity. More such divisions are possible that each bring along small improvements in sensitivity, but are currently not taken advantage of either because of limited Monte Carlo statistics or simply to avoid the additional complexity.

One simple improvement that could be implemented, at least in the limit of infinite Monte Carlo, is to separate regions of the anti-fiducial volume that are lumped together currently. For instance, this analysis puts the back and the edge of the back together, but the edge has a higher concentration of rock events and therefore a different neutrino spectrum. The same idea holds for the two sides of the supermodule gap; the north side has a higher rock muon concentration. Similarly, in the rock-like edge, a separation could be made between the top and bottom of the detector, since due to the beam angle, the bottom has a larger concentration of rock events. However, all of these together will only improve the fit by a few percent, so it is not clear the additional complexity is worthwhile.

In the same vein as the fiducial analysis's resolution binning, this analysis could gain a bit of sensitivity by separating tracks measured by range and those measured by curvature. However, most events contributing significantly to the fit use range and so this is estimated to improve the fit less than a percent. Another $\sim 1\%$ improvement can be gained by putting positively and negatively charged tracks in separate histograms. This follows mainly from the very different peak energy of the neutrino and antineutrino spectra in the neutrino beam.

A more obviously useful refinement would be the inclusion of shower energy for detector-like events. Many of these events, particularly the more valuable low energy ones, have showers that appear by eye to be well-contained. Even those that are not completely

contained could be useful if we had confidence in both the shower modeling and the Monte Carlo's modeling of the detector edge. A rough study shows that including shower energy for detector-like events could improve sensitivity by $\sim 4\%$ overall. In order to use shower energy, it would be important to use the currently in-development next-generation MINOS Monte Carlo that incorporates surveyed plane and strip alignments. This would allow correct generation of shower shape, particularly on the complicated horizontal and diagonal edges. Only the strip transverse alignments are surveyed currently, but it is possible that with increased cosmic statistics, the location of strip ends could be better constrained than they are now.

13.2 Improvements to Systematics

Each of the above statistical refinements comes with the risk of increased exposure to systematic error. This analysis in its current form may well be systematically limited before the end of MINOS running. Therefore, time might be better spent reducing systematics. Preliminary work on including several additional systematics in the nuisance parameter fit suggests that without any other modification, the systematic error could be reduced from the current $0.14 \times 10^{-3} \text{ eV}^2$ to $0.12 \times 10^{-3} \text{ eV}^2$. To reduce this further, work could be done to reduce the largest two systematics: the track energy and the relative near/far normalization. Each of these has possible, but difficult, avenues for improvement.

The majority of the effect of the track energy systematic comes from the near detector, since it affects the prediction of the far detector energy spectrum. Now that MINER ν A is running in the near detector hall simultaneously with MINOS, it is possible that it could be used to constrain our track energy measurements. There are two methods for doing this. First, tracks low enough in energy that they stop in MINER ν A ($\lesssim 1 \text{ GeV}$) can be compared to similar tracks in MINOS. Second, tracks that traverse both detectors can be analyzed for how well the two segments agree. While MINER ν A, which has no magnetic field, cannot independently measure the energy of exiting muons, the correspondence can be tracked by dividing the sample of these muons into those starting at different depths in MINER ν A. If the two detectors are in agreement about track energy, the spectrum should be the same

for different starting points. But if MINOS is overestimating track energy compared to MINER ν A, the spectrum will shift up for tracks less contained in MINER ν A. Assuming MINER ν A's track energy systematic is substantially smaller than 2%, our systematic might be significantly improved via this sort of study. Since MINER ν A calibration and analysis is currently still in the very early stages, it is not clear how feasible any of this is.

To improve the relative near/far normalization, the hand-scan study used to produce the current error could be repeated at higher statistics. Perhaps as much as four times more far detector data will eventually be available than was used originally, so without changing the methodology it would be possible to reduce the most significant part of this error by a factor of 2. This would reduce the total normalization error from 1.54% to 1.10%. It's not at all clear, however, that the enormous human effort required would be worthwhile, especially because for the fiducial analysis the benefit is somewhat smaller (although not negligible).

Even under a very optimistic scenario, the normalization systematic is certain to remain dominant, and the track energy systematic nearly so. However, at this point the subdominant systematics would begin to have noticeable impact, so bear consideration. Some of the cross section systematics should be reducible with hardly any effort in the coming years as MINER ν A returns results. The neutral current systematic is also straightforward to reduce: simply tighten the PID cuts. Presumably in this scenario, the small loss of statistical power will be worthwhile. All told, it is likely that the total systematic could be reduced to about two thirds of its current value, which would keep this analysis statistically limited for the remainder of MINOS running.

13.3 Antineutrino Beam Analysis

Much interest in the possibility of either CPT-violation or non-standard interactions has been generated lately by MINOS's Run IV antineutrino analysis, which found a best fit point $\sim 2\sigma$ separated from MINOS's neutrino best fit point. This analysis has used only fiducial events so far. Although more antineutrino running is planned, the statistics of this sample are fairly low at the moment and, due to the lower antineutrino cross section,

will stay lower than the neutrino sample. The addition of rock and anti-fiducial muons to the antineutrino measurement should benefit it in the same proportion as this analysis benefits the neutrino measurement. Since the antineutrino beam has considerable neutrino contamination at high energy, a reasonably tight charge cut will have to be applied, but this should not reduce the statistics more than 10%.

13.4 Application to Other Experiments

The upcoming $\text{NO}\nu\text{A}$ experiment, designed to search for ν_e appearance, will also perform a ν_μ disappearance measurement [29]. The $\text{NO}\nu\text{A}$ far detector is conceptually very similar to MINOS and so many of the techniques of this analysis would most likely be easy to translate to a $\text{NO}\nu\text{A}$ rock and anti-fiducial muon analysis.

$\text{NO}\nu\text{A}$ will have a front face with 4 times the surface area of MINOS's. Since it is on the surface of the Earth and will see ~ 1 cosmic muon per beam spill, it is not clear that beam events entering the top or sides could be used. However, the bottom edge will be usable and has a surface area also 4 times that of MINOS's bottom edge. Because it is off-axis, the event rate will be reduced in roughly the same proportion, so the number of rock events will be, to first order, the same. However, the reason for putting it off-axis — to produce a narrow neutrino spectrum with a highly suppressed high-energy tail — will be greatly beneficial for the sensitivity of a rock muon analysis as there will be dramatically less feed-down from high energy events far upstream.

The geometry of the $\text{NO}\nu\text{A}$ detector is considerably simpler than that of MINOS, being a square rather than an octagon with all cells the same length. So the bottom edge will be like the MINOS diagonal edges and similarly easy to handle. As a back-of-the-envelope calculation, it seems likely that a rock and anti-fiducial analysis would benefit $\text{NO}\nu\text{A}$ in about the same proportion as this analysis does for MINOS.

Several other long-baseline ν_μ -beam experiments are running now or in the planning

stages. These include T2K^{*}, OPERA[†] and LBNE[‡]. It is possible that the techniques presented here could have some relevance to these as well, however without personal experience on any of these experiments, I will not attempt a sensitivity estimate.

^{*}Tokai To Kamioka [30]

[†]Oscillation Project with Emulsion-tRacking Apparatus [32]

[‡]Long Baseline Neutrino Experiment. This is a proposal for a beam from Fermilab to Homestake [94]. The collaboration attempted to choose a less-generic name through popular vote, but failed to make one stick.

Bibliography

- [1] W. Pauli, Offener Brief an die Gruppe der Radioaktiven bei der Gauvereins-Tagung zu Tübingen, [Open letter to the group of radioactive people at the Gauverein meeting in Tübingen], 1930.
- [2] C.L. Cowan et al., *Science* 124 (1956) 103.
- [3] G. Danby et al., *Phys. Rev. Lett.* 9 (1962) 36.
- [4] DONuT, K. Kodama et al., *Phys. Lett. B*504 (2001) 218, hep-ex/0012035.
- [5] B. Pontecorvo, *Sov. Phys. JETP* 7 (1958) 172.
- [6] Z. Maki, M. Nakagawa and S. Sakata, *Prog. Theor. Phys.* 28 (1962) 870.
- [7] S.M. Bilenky and B. Pontecorvo, *Phys. Rept.* 41 (1978) 225.
- [8] R. Davis, Jr., D.S. Harmer and K.C. Hoffman, *Phys. Rev. Lett.* 20 (1968) 1205.
- [9] Super-Kamiokande, J.P. Cravens et al., *Phys. Rev. D* 78 (2008) 032002.
- [10] T.J. Haines et al., *Phys. Rev. Lett.* 57 (1986) 1986.
- [11] Kamiokande-II, K.S. Hirata et al., *Phys. Lett. B*205 (1988) 416.
- [12] Super-Kamiokande, Y. Fukuda et al., *Phys. Rev. Lett.* 81 (1998) 1562, hep-ex/9807003.
- [13] Super-Kamiokande, Y. Fukuda et al., *Phys. Rev. Lett.* 82 (1999) 2644, hep-ex/9812014.
- [14] Super-Kamiokande, Y. Fukuda et al., *Phys. Lett. B*433 (1998) 9, hep-ex/9803006.
- [15] Super-Kamiokande, Y. Fukuda et al., *Phys. Lett. B*436 (1998) 33, hep-ex/9805006.

- [16] W.W.M. Allison et al., Phys. Rev. D 72 (2005) 052005.
- [17] K2K, M.H. Ahn et al., Phys. Rev. D74 (2006) 072003, hep-ex/0606032.
- [18] SNO, Q.R. Ahmad et al., Phys. Rev. Lett. 87 (2001) 071301.
- [19] SNO, A.L. Hallin, Prog. Part. Nucl. Phys. 64 (2010) 199.
- [20] T. Yanagida, Prog. Theor. Phys. 64 (1980) 1103.
- [21] K.S. Babu, D. Eichler and R.N. Mohapatra, Physics Letters B 226 (1989) 347 .
- [22] Particle Data Group, C. Amsler et al., Phys. Lett. B667 (2008) 1.
- [23] A.D. Sakharov, Pisma Zh. Eksp. Teor. Fiz. 5 (1967) 32.
- [24] A. Aguilar et al., Phys. Rev. D 64 (2001) 112007.
- [25] MiniBooNE, A.A. Aguilar-Arevalo et al., Observed Event Excess in the MiniBooNE Search for $\bar{\nu}_\mu \rightarrow \bar{\nu}_e$ Oscillations, hep-ex/1007.1150, 2010.
- [26] R. Wendell et al., Phys. Rev. D 81 (2010) 092004.
- [27] CHOOZ, M. Apollonio et al., Phys. Lett. B466 (1999) 415, hep-ex/9907037.
- [28] MINOS, P. Adamson et al., Phys. Rev. Lett. 101 (2008) 131802, 0806.2237.
- [29] NO ν A, Proposal to Build a 30 Kiloton Off-Axis Detector to Study $\nu_\mu \rightarrow \nu_e$ Oscillations in the NuMI Beamline, 2005.
- [30] T2K, Y. Itow et al., The JHF-Kamioka neutrino project, hep-ex/0106019, 2001.
- [31] H. Minakata et al., Phys. Rev. D74 (2006) 053008, hep-ph/0607284.
- [32] OPERA, N. Agafonova et al., Phys. Lett. B691 (2010) 138, 1006.1623.
- [33] MINOS, P. Adamson et al., New constraints on muon-neutrino to electron-neutrino transitions in MINOS, hep-ex/1006.0996, 2010.
- [34] S.E. Kopp, The NuMI Neutrino Beam at Fermilab, physics/0508001, 2005.
- [35] MINERvA, K.S. McFarland, Nucl. Phys. Proc. Suppl. 159 (2006) 107, physics/0605088.

- [36] MINOS, P. Adamson et al., Phys. Rev. D77 (2008) 072002, 0711.0769.
- [37] R. Armstrong et al., Constraining Hadron Production Spectra in NuMI Beam by a Fit to the MINOS Near Detector Data, MINOS internal note 4166, version 5, 2008.
- [38] P. Adamson et al., Measuring the Number of Protons-on-Target (POT) in the NuMI Beamline, MINOS internal note 1491, version 2, 2006, See page 20: the uncertainty due to toroid noise is negligible being $< 0.03\%$.
- [39] MINOS, D.G. Michael et al., Nucl. Instrum. Meth. A596 (2008) 190, 0805.3170.
- [40] Robert Hatcher, personal communication, 2010.
- [41] M. Strait, Anti-fiducial analysis: Statistical optimization of edge, MINOS internal note 6811, version 3, 2010.
- [42] MINOS Calibration Group, 2009 Position Paper on Calibration of Runs I-II-III, MINOS internal note 6717, version 2, 2010.
- [43] MINOS, P. Adamson et al., Phys. Rev. Lett. forthcoming (2010[?]).
- [44] J. Allison et al., Nuclear Science, IEEE Transactions on 53 (2006) 270 .
- [45] S. Agostinelli et al., Nucl. Instrum. Meth. A506 (2003) 250 .
- [46] A. Ferrari et al., FLUKA: A multi-particle transport code (Program version 2005), CERN-2005-010.
- [47] G. Battistoni et al., AIP Conf. Proc. 896 (2007) 31.
- [48] M. Campanella et al., CERN Report ATLAS Internal Note ATL-SOFT 98-039 (1998).
- [49] M. Campanella et al., CERN Report ATL-SOFT-99-004 (1999).
- [50] H. Gallagher, Nucl. Phys. Proc. Suppl. 112 (2002) 188.
- [51] Hugh Gallagher, personal communication, 2010.
- [52] C. Andreopoulos et al., Nucl. Instrum. Meth. A614 (2010) 87, 0905.2517.
- [53] William Miller, personal communication, 2010.

- [54] A.P. Ruotsala and S.P. Tufford, Chemical Analyses of Igneous Rocks (Minnesota Geological Survey, 1965).
- [55] P.K. Sims, Geology of Minnesota: a centennial volume, in honor of George M. Schwartz (Minnesota Geological Survey, 1972).
- [56] Chris Loughton, personal communication, 2008.
- [57] M.C. Goodman, An Argument About Rock Muons, MINOS internal note 5589, version 1, 2009.
- [58] B. Armstrong and M. Ishitsuka, Simulation study of geomagnetic field in the decay pipe, MINOS internal note 4706, version 1, 2008.
- [59] A.M. 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, 2007.
- [60] T. Hastie, The elements of statistical learning: data mining, inference, and prediction (New York: Springer, 2009), Section 13.3 is on kNN.
- [61] R. Ospanov, A measurement of muon neutrino disappearance with the MINOS detectors and NuMI beam, PhD thesis, University of Texas, 2008.
- [62] J. Ratchford, *Forthcoming*, PhD thesis, University of Texas at Austin, 2011(?).
- [63] M. Dorman, Beam Fit Position Paper, MINOS internal note 7146, version 1, 2010.
- [64] MINOS, P. Adamson et al., Search for muon antineutrino flavor change in the MINOS experiment, MINOS internal note 6427, version 13, 2010.
- [65] N. Saoulidou, Far Detector Oscillation Analysis Using the Beam Matrix Method, MINOS internal note 1531, version 4, 2006.
- [66] J. Evans et al., The UK Beam Matrix Method, MINOS internal note 4298, version 3, 2008.

- [67] S.J. Coleman, Extrapolating with Antineutrinos and Resolution Bins, MINOS internal note 6874, version 16, 2010.
- [68] J.M. Paley, Hand-scan Estimation of the Reconstruction-Related Systematic Uncertainty in the Near/Far Normalization of the MINOS CC Analysis, MINOS internal note 5613, version 2, 2009.
- [69] R. Hatcher et al., Range/Energy Task Force Position Paper, MINOS internal note 3134, version 2, 2007.
- [70] James A Musser, personal communication, 2010.
- [71] R. Ospanov, Final report on the Near Detector Magnetic Field Calibration, MINOS internal note 3584, version 1, 2007.
- [72] Justin Evans, personal communication, 2010.
- [73] S.J. Barish et al., Phys. Rev. D16 (1977) 3103.
- [74] S.J. Barish et al., Phys. Rev. D19 (1979) 2521.
- [75] S.J. Barish et al., Phys. Lett. B66 (1977) 291.
- [76] N.J. Baker et al., Phys. Rev. D25 (1982) 617.
- [77] C. Baltay et al., Phys. Rev. Lett. 44 (1980) 916.
- [78] Gargamelle Neutrino Propane, S. Ciampolillo et al., Phys. Lett. B84 (1979) 281.
- [79] O. Erriquez et al., Phys. Lett. B80 (1979) 309.
- [80] D.S. Baranov et al., Phys. Lett. B81 (1979) 255.
- [81] V.B. Anikeev et al., Z. Phys. C70 (1996) 39.
- [82] NOMAD, Q. Wu et al., Phys. Lett. B660 (2008) 19, 0711.1183.
- [83] Omar Benhar and Jaroslaw Nowak, personal communication, 2010.
- [84] A.M. Ankowski and J.T. Sobczyk, Phys. Rev. C74 (2006) 054316, nucl-th/0512004.

- [85] A.M. Ankowski, O. Benhar and N. Farina, Phys. Rev. D 82 (2010) 013002.
- [86] MINOS, D.G. Michael et al., Phys. Rev. Lett. 97 (2006) 191801, hep-ex/0607088.
- [87] W.A. Rolke, A.M. Lopez and J. Conrad, Nucl. Instrum. Meth. A551 (2005) 493, physics/0403059.
- [88] G.D. Orebi Gann, An Improved Measurement of the ^8B Solar Neutrino Energy Spectrum at the Sudbury Neutrino Observatory, PhD thesis, University of Oxford, Merton College, 2008, Discussion of systematics is in section 7.5.
- [89] J.R. Wilson, A Measurement of the ^8B Solar Neutrino Energy Spectrum at the Sudbury Neutrino Observatory, PhD thesis, University of Oxford, Jesus College, 2004, Discussion of systematics is in section 5.1.
- [90] S. Oser, Physics 509 lecture 12: Propagating Systematic Uncertainties, 2008, from a course at the University of British Columbia.
- [91] J.S. Mitchell, Position Paper on Oscillation Fitting for the 2010 CC Analysis, MINOS internal note 6862, version 10, 2010.
- [92] J.S. Mitchell, *Forthcoming*, PhD thesis, University of Cambridge, 2011(?).
- [93] T. Yasuo, Results from Super-Kamiokande, Talk given at the XXIV International Conference on Neutrino Physics and Astrophysics (Neutrino 2010): <http://indico.cern.ch/contributionDisplay.py?contribId=224&confId=73981>, 2010.
- [94] LBNE Collaboration, LNBE Website at Fermilab, <http://lbne.fnal.gov>, 2010.