

**An Atmospheric Muon Neutrino Disappearance
Measurement with the MINOS Far Detector**

A THESIS

**SUBMITTED TO THE FACULTY OF THE GRADUATE SCHOOL
OF THE UNIVERSITY OF MINNESOTA**

BY

Jeremy Peter Gogos

**IN PARTIAL FULFILLMENT OF THE REQUIREMENTS
FOR THE DEGREE OF
Doctor Of Philosophy**

December, 2007

UMI Number: 3291500

Copyright 2007 by
Gogos, Jeremy Peter

All rights reserved.

UMI[®]

UMI Microform 3291500

Copyright 2008 by ProQuest Information and Learning Company.
All rights reserved. This microform edition is protected against
unauthorized copying under Title 17, United States Code.

ProQuest Information and Learning Company
300 North Zeeb Road
P.O. Box 1346
Ann Arbor, MI 48106-1346

© Jeremy Peter Gogos 2007
ALL RIGHTS RESERVED

Acknowledgements

“Ph.D.” Three letters. Two periods. Eight years. I obviously have a lot of people to thank.

I would like to immediately recognize two crucial contributors to my thesis. I was fortunate enough to end up with two advisers, Jon Urheim and Peter Litchfield, both of whom have my heartfelt gratitude for their insight, guidance, and forgiving temperament in dealing with such an incorrigible pupil. I am sorry that you had to put up with me for so long!

I have also been a pupil in more classes than I can remember, at the University of Minnesota, Worcester Polytechnic Institute and at Mesabi Community College. I have been impressed by my professors’ depth of knowledge, enthusiasm, and sincere desire to see their students succeed (the occasional unavoidable curmudgeon withstanding). As such, I want to thank Benjamin Bayman, William Zimmerman, Mikhail Voloshin, Serge Rudaz, Earl Peterson, Daniel Cronin-Hennessy, Yuichi Kubota, Ken Heller, Marvin Marshak, and Arkady Vainshtein of the U of M. From WPI, I will always fondly remember Suzanne Weekes, Harold Hilsinger, Van Bluemel, Adriaan Walther, Germano Iannacchione, Stephen Pierson, Stephen Jasperson, Thomas Keil, and Louis Colonna-Romano. Finally, at Mesabi, Georgia Suoja, Kerry Duncan, Ken Pontinen, and Bob Phillips all have my heartfelt thanks.

Never have I seen someone thank his or her elementary, junior high, or high school

teachers. I find it odd that the efforts of these teachers go unrecognized, as anyone who pursues a Ph.D. probably had academic interests and talents that were nurtured early on. For fostering my abilities and encouraging my education, I extend my gratitude to Margaret Schley, Richard Moore, David Johnson, Raymond Sandon, John Husnik, David Kriska, Roger Ezell, Gary Arndt, Cathy Priest, Tony Nigro, Robert Kochevar, Richard Paciotti, Becky Peterson, Blanche Barber, Margaret Peterson, Cammy Newman, and Heather Larsen.

Beyond superlative teachers, I was also financially fortunate; I won the Alworth Scholarship, and the monetary support of the Marshall H. and Nellie Alworth Memorial Fund was a huge boon and greatly appreciated.

Next, I want to thank the numerous people I have come to know at the U of M. Ben, everyone at Minnesota on MINOS or *NOVA* owes you a huge debt for your tireless work to maintain our computing resources, and your unprompted tirades about whatever esoterica happened to be angering you at the moment always made time fly. Sujeewa, I sincerely enjoyed all of our conversations, whether they were about physics or teaching me to curse in Sinhala; I already miss your banshee's shriek of a laugh. Erik, thanks for spitting on the back of my head while I was driving you from Ely to Tower. Dipu, you introduced me to the game of snooker and showed me how unfashionable I really am. Bernie, thanks for not pressing charges against me when I tried to set you on fire. Emily, if you weren't a raging libertarian, I would have asked for your hand in marriage. Steve, it was nice to have an identical hybrid around. Tackleberry, it was good to know someone with administrative clout. Tao, your visits to our office always reminded me that I should be working as hard as you. Xifeng, thanks to you and your wife for the authentic Chinese food. To Sue, David, Leon, Kregg, and Pete, you all helped me many times and I really appreciate it. The physics department staff, in particular Dave Holets, Judy Soine, Bobbi Eich, Ron Huhn, and Graham Allan, go above and beyond

the call of duty to help socially and administratively inept physicists like me. Thanks to Maribel for supporting local music.

I should also thank the MINOS collaboration, in particular the whole Cambridge University group, Brian Rebel, and the mine crew at Soudan.

Few people can endure school without outside distraction, and I have to give credit to the people who helped me stay sane. Mom, Sam, Georgia, Freddy, and Geoffrey, thank you for driving me around in Worcester, and for feeding me from time to time. Justin, Jeremey, and Cory, we bring the rock. Serena and Angie, all the bakery has been delicious. Heather, Meagan, Nicole, Meagan, Zoe, Abby, and Meagan, I am truly sorry; you all deserved better.

Dedication

With love, I dedicate this thesis to my grandparents, Frank and Faye Sherman. They raised me as their son, toiled so that I would never have to, and forsook their dreams so I could pursue mine. What they gave me I can never repay.

An Atmospheric Muon Neutrino Disappearance Measurement with the MINOS Far Detector

by Jeremy Peter Gogos

Under the supervision of Jon Urheim and Peter Litchfield

ABSTRACT

It is now widely accepted that the Standard Model assumption of massless neutrinos is wrong, due primarily to the observation of solar and atmospheric neutrino flavor oscillations by a small number of convincing experiments. The MINOS Far Detector, capable of observing both the outgoing lepton and associated showering products of a neutrino interaction, provides an excellent opportunity to independently search for an oscillation signature in atmospheric neutrinos. To this end, a MINOS data set from an 883 live day, 13.1 kt-yr exposure collected between July, 2003 and April, 2007 has been analyzed. 105 candidate charged current muon neutrino interactions were observed, with 120.5 ± 1.3 (statistical error only) expected in the absence of oscillation. A maximum likelihood analysis of the observed $\log(L/E)$ spectrum shows that the null oscillation hypothesis is excluded at over 96% confidence and that the best fit oscillation parameters are $\sin^2 2\theta_{23} = 0.95_{-0.32}$ and $\Delta m_{23}^2 = 0.93_{-0.44}^{+3.94} \times 10^{-3} \text{ eV}^2$. This measurement of oscillation parameters is consistent with the best fit values from the Super-Kamiokande experiment at 68% confidence.

Contents

Acknowledgements	i
Dedication	iv
Abstract	v
List of Tables	xi
List of Figures	xii
1 Introduction	1
2 Tour of Neutrino Physics and Neutrino Oscillations	3
2.1 Neutrino History Timeline	3
2.2 The Standard Model	5
2.2.1 Neutrino Production and Interactions	7
2.3 Oscillation Physics	12
2.3.1 Two Flavor Oscillation Phenomenology	14
2.4 Atmospheric Neutrino Summary	18
2.5 The Atmospheric Neutrino Problem	19
2.5.1 Kamiokande	22

2.5.2	Fréjus	23
2.5.3	IMB	23
2.5.4	Soudan 2	24
2.5.5	MACRO	26
2.5.6	Super-Kamiokande	26
2.5.7	Beam Experiments	29
2.6	Outstanding Issues	30
2.6.1	Number of Neutrino Species	31
2.6.2	θ_{13}	31
2.6.3	Mass Hierarchy	33
3	The MINOS Detector and Infrastructure	35
3.1	Detector Panorama	36
3.2	Plane Design	38
3.2.1	Scintillator	39
3.3	Shield Design	41
3.4	Data Acquisition	42
3.4.1	Optical Summing	44
3.4.2	VA Chip	47
3.4.3	VA Frontend Board	48
3.4.4	VMM	48
3.4.5	VARC	49
3.4.6	ROPs, BRPs, TPs, Oh My	51
3.5	Calibration	52
3.5.1	Charge Calibration	52
3.5.2	Timing	53
3.5.3	CalDet	54

3.5.4	MINOS Module Mapper	54
3.5.5	Strip-to-Strip Calibration	55
3.5.6	MEU Calibration	55
3.6	Event Reconstruction	57
3.6.1	Noise Filtering	59
3.6.2	Demultiplexing	59
3.6.3	Slicing	60
3.6.4	Clustering	60
3.6.5	Shower Finding	60
3.6.6	Tracking	61
3.6.7	Track Fitting	61
3.6.8	Event Assembly	62
3.7	Neutrino Variables and Resolutions	62
3.7.1	The Neutrino Variables L and E	62
3.7.2	Resolutions	66
4	Monte Carlo	69
4.1	Monte Carlo Components	69
4.1.1	Atmospheric Neutrino and Cosmic Ray Muon Fluxes	70
4.1.2	Neutrino Event Generation	71
4.1.3	Particle Propagation	71
4.1.4	Detector Response	72
4.2	Monte Carlo Samples	72
4.3	Comparison to Data	73
4.3.1	Cosmic Ray Muons	73
4.3.2	Atmospheric Neutrinos	76

5	Data Samples and Event Selection	79
5.1	Initial Data Sample	79
5.2	Event Selection	80
5.2.1	Quality Control	81
5.2.2	Kinematic Cuts	82
5.2.3	Fiducial Cuts	82
5.2.4	Neural Network	84
5.2.5	Beam Timing	94
5.2.6	Shield	96
5.3	Final Data Set	99
6	Oscillation Fit	107
6.1	Bayesian Interlude	109
6.2	Background	111
6.3	Fit Results	113
6.3.1	Sensitivity	123
6.4	Goodness of Fit	123
6.5	The Null Oscillation Scenario	128
7	Systematic Error Analysis	130
7.1	Uncertainties in Neutrino Production	131
7.1.1	Primary Flux	131
7.1.2	Hadron Production	132
7.2	Neutrino Cross-Section	133
7.3	Background Estimation	134
7.4	Neural Net	134
7.5	Muon Momentum and Hadronic Energy Reconstruction	140

7.6	Cumulative Systematic Effect	140
7.7	Physical Boundary	144
8	Conclusions	147
8.1	Comparisons	147
8.2	The Future	152
8.2.1	Sample Splitting	152
8.3	Summary	154
	Appendix A. Low $\log(L/E)$ Data Events	155

List of Tables

2.1	Fields and their associated bosons and coupling constants	6
3.1	The upper and lower 32 bit data words for a digit	50
4.1	MC Samples	72
5.1	The breakdown of shield efficiency for three similarly sized ranges of net output	99
5.2	The number of data and MC events left after each cut.	100
7.1	Systematic error summary table	144

List of Figures

2.1	Neutral current ν interaction with an up quark	8
2.2	Charged current ν_μ interaction	9
2.3	Total charged current neutrino cross section	10
2.4	A charged current neutrino interaction showing the hadronic shower	12
2.5	Flux of primary cosmic radiation	20
2.6	Total atmospheric neutrino flux	21
2.7	The Soudan 2 confidence limits	25
2.8	Super-Kamiokande's zenith angle distribution	27
2.9	Super-Kamiokande's $\nu_\mu \rightarrow \nu_\tau$ oscillation allowed region	28
2.10	Super-Kamiokande's $(L/E)_{Data}/(L/E)_{MC}$ spectrum	30
2.11	The CHOOZ and Palo Verde 90% excluded regions	32
2.12	The normal mass hierarchy	33
3.1	A schematic of the Soudan Underground Laboratory	36
3.2	The MINOS local coordinate system	38
3.3	A view of one plane of the MINOS detector	39
3.4	MINOS module types	41
3.5	Veto shield construction	43
3.6	Inside a mux box	45
3.7	Phototube cookie diagram	46

3.8	A module mapper “map”	56
3.9	Muon energy loss curve in copper	58
3.10	Atmospheric neutrino flight path diagram	63
3.11	Neutrino zenith angle resolution (from MC)	64
3.12	Neutrino energy resolution (from MC)	65
3.13	The relative error on the reconstructed neutrino energy (from MC)	66
3.14	The absolute error on the reconstructed neutrino zenith angle (from MC)	67
3.15	The relative error on the reconstructed neutrino $\log(L/E)$ (from MC)	68
4.1	Zenith angle distribution of cosmic ray muons (no cuts)	74
4.2	Ratio of data and MC zenith angle distributions for all tracks	75
4.3	Zenith angle distribution of stopping muons	77
4.4	Ratio of data and MC zenith angle distributions for stopping muons	78
5.1	Track vertex direction cosine y and z	83
5.2	Fiducial regions	85
5.3	Neural net schematic	86
5.4	Vertex z trace diagram	88
5.5	Number of strips in an event (from MC)	89
5.6	Track vertex y (from MC)	89
5.7	Track vertex direction cosine y (from MC)	90
5.8	Track vertex direction cosine z (from MC)	90
5.9	Track end direction cosine z (from MC)	91
5.10	Track vertex z trace (from MC)	91
5.11	Total vertex charge (from MC)	92
5.12	The event charge per plane (from MC)	92
5.13	The track charge fraction (from MC)	93
5.14	Atmospheric and cosmic MC neural net output	94

5.15	Neural net figure of merit	95
5.16	Beam timing distributions from data	96
5.17	The distribution of shield strip times (from data)	98
5.18	The spatial vertex distribution of selected events	102
5.19	The summed number of selected events month by month	103
5.20	The neutrino y distribution for the final data set and MC	104
5.21	The neutrino zenith angle and energy for the final data set and MC	105
5.22	The distribution of $\log(L/E)$ for the final data set	106
6.1	Monte Carlo oscillated and unoscillated $\log(L/E)$ distribution	108
6.2	Showering neutrino background zenith angle distribution (from MC)	112
6.3	Showering background E distribution (from MC)	113
6.4	Showering background $\log(L/E)$ distribution (from MC)	114
6.5	Total cosmic ray background zenith angle distribution (from data)	115
6.6	Total cosmic ray background E distribution (from data)	115
6.7	Cosmic ray background $\log(L/E)$ distribution (from data)	116
6.8	The final $\log(L/E)$ distribution	117
6.9	The final, minimum-subtracted, likelihood surface	119
6.10	The final, minimum-subtracted, likelihood contour	120
6.11	The one dimensional, minimum subtracted, likelihood contour in Δm_{23}^2	121
6.12	The one dimensional, minimum subtracted, likelihood contour in $\sin^2(2\theta_{23})$	122
6.13	The 68%, 90%, and 99% sensitivity	124
6.14	The 68% likelihood and sensitivity contours	125
6.15	The χ^2 distribution for an ensemble of MC experiments	127
6.16	The χ^2 distribution for no oscillations	129
7.1	Comparison of flux measurements to the GSHL	132
7.2	Pion and kaon production uncertainty as a function of x_{lab}	133

7.3	68% contours for varying the neutrino <i>DIS</i> cross-section	135
7.4	68% contours for varying the neutrino <i>QE</i> cross-section	136
7.5	68% contours for varying the neutrino <i>RES</i> cross-section	137
7.6	68% contours for varying neutral current and cosmic ray backgrounds .	138
7.7	68% contours for varying the neural net cut	139
7.8	68% contours for varying the reconstructed muon momentum	141
7.9	68% contours for varying the reconstructed hadronic shower energy . . .	142
7.10	The cumulative systematic contour and the original 68% contour	143
7.11	The unphysical likelihood surface	145
7.12	The unphysical likelihood contours	146
8.1	The 68% confidence limits for the Cambridge analysis and this analysis	149
8.2	The 90% confidence limits for the Cambridge analysis and this analysis	150
8.3	The 99% confidence limits for the Cambridge analysis and this analysis	151
8.4	The resolution on $\log(L/E)$ split in kinematic y	153
A.1	The spatial vertex distribution of low $\log(L/E)$ data events	157
A.2	The zenith angle and energy distributions of low $\log(L/E)$ data events .	158
A.3	The track direction cosine x and z distributions for low $\log(L/E)$ data events	159
A.4	The cumulative number of low $\log(L/E)$ data events	160

Chapter 1

Introduction

The last two decades were an exciting time for neutrino physics, particularly in the area of neutrino oscillation. Through the diligent and insightful work of a handful of experiments, physicists have illuminated the machinery at work underneath the naive Standard Model assumption of three massless, uncoupled neutrino flavors. The lepton sector, as it turns out, shares intriguing similarities with the quark sector, some of which have been explored through three different routes: reactor neutrinos, solar neutrinos, and atmospheric neutrinos. The first demonstrated that the coupling between the ν_e and ν_τ , because of the now suspected flavor composition of the ν_1 and ν_3 mass states, must be very small; the second showed the interplay between the medium size coupling of the ν_e and ν_μ and the matter enhancement of electron neutrino oscillations; and the third displayed the nearly maximal mixing between the ν_μ and ν_τ .

This thesis investigates oscillations in the atmospheric neutrino sector by looking for a deficit of muon-type atmospheric neutrino event as a function of the event's $\log(L/E)$, where L is the length of the neutrino's flight path from production to interaction (in km), and E is the neutrino's energy (in GeV). This combination of variables is chosen because, as described in Section 2.3.1, neutrino oscillation is sensitive to this ratio.

These atmospheric neutrinos originate in meson decay, and these mesons are produced when cosmic rays interact with Earth's upper atmosphere. The flux of these neutrinos has been modeled by, among others, the Bartol group [17]. The analysis presented in this thesis is based on a comparison of the observed muon-type atmospheric neutrino $\log(L/E)$ distribution with the predicted distribution.

The subsequent chapters follow this general path:

- Chapter 2 discusses the historical emergence of the neutrino, lays some theoretical groundwork, and summarizes the current experimental situation.
- Chapter 3 describes the MINOS hardware in detail.
- Chapter 4 lays out the Monte Carlo framework of the MINOS experiment, an important piece of any high energy physics endeavor.
- Chapter 5 describes the difficult task of sifting through the large data set to extricate signal events.
- Chapter 6 details how the signal events are analyzed once they are selected.
- Chapter 7 enumerates the relevant systematic errors inherent to the analysis.
- Chapter 8 summarizes all relevant details and presents the projected sensitivity of the MINOS experiment to atmospheric neutrino oscillations.

Chapter 2

Tour of Neutrino Physics and Neutrino Oscillations

In a sentence, the goal of this thesis is to present a measurement of the $\nu_\mu \rightarrow \nu_\tau$ oscillation parameters using atmospheric neutrinos. That succinct statement belies many details which conceal many devils, and so vocabulary must first be established, as well as historical and theoretical motivation for this measurement. Section 2.1 provides a brief summary of the history of the neutrino, Sections 2.2 and 2.3 expound on the physics of neutrino interactions and oscillations, Sections 2.4 and 2.5 lay out the relevant experimental background, and Section 2.6 highlights some of the unanswered questions about neutrinos.

2.1 Neutrino History Timeline

The historical emergence of the neutrino has been described extensively in places too numerous to count. Here a timeline is provided as a summary. Noteworthy oscillation milestones are discussed later on. For a more complete discussion of most of the timeline,

see [55].

1914 The β spectrum from nuclei undergoing radioactive decay is observed to be continuous, violating conservation of energy in what was assumed to be a two body decay [29].

1930 Pauli postulates the existence of a spin 1/2, electrically neutral particle that is emitted in nuclear β decay [62].

1934 Fermi utilizes Pauli's neutral particle in a theoretical description of a four fermion interaction which successfully predicts the continuous β spectrum and decay rate of nuclear decays. Fermi also coins the name "neutrino" [37].

1946 Pontecorvo proposes direct neutrino observation through inverse β decay, $p + \bar{\nu} \rightarrow n + e^+$ [67].

1953 Reines and Cowan publish results on the first direct observation of neutrinos [70], employing Pontecorvo's idea of inverse β decay.

1956 Lee and Yang realize that spatial parity may not be respected in β decay [54].

1957 Wu, *et al.* provide experimental verification that, indeed, parity is not respected in β decay [75].

1958 Goldhaber establishes via photon polarization observation from nuclear electron capture de-excitation that the neutrino field has its momentum and spin anti-parallel (left-handed) [45].

1959 Pontecorvo proposes that two flavors of neutrino, electron and muon, may exist [66].

1962 The existence of the muon neutrino is experimentally verified at Brookhaven [32].

1975 The τ lepton is discovered by Perl, *et al.* [63]. The existence of this third charged lepton implied the existence of a third neutrino, ν_τ .

2001 The DONUT collaboration makes the first experimental observation of ν_τ [53].

2.2 The Standard Model

Many people have extensively described the Standard Model and, specific to this thesis, neutrinos and the weak interaction (see, for example, [25] and [7]), but it would be negligent to not provide a brief overview.

Currently, the Standard Model includes three species of neutrino, which are partnered with the charged leptons e , μ , and τ to give six total leptons in three families.

$$\begin{pmatrix} e \\ \nu_e \end{pmatrix} \begin{pmatrix} \mu \\ \nu_\mu \end{pmatrix} \begin{pmatrix} \tau \\ \nu_\tau \end{pmatrix}$$

The leptons are one of three groups of particles which comprise the physical basis of the Standard Model. The second group contains the gauge bosons, listed in Table 2.2, which mediate the four fundamental forces. The higgs boson, associated with the higgs field, and essential to explaining particle masses within the Standard Model, is also in this group. The W^\pm and Z° bosons are quite massive (80.4 and 91.2 GeV, respectively) when compared to the massless gluon, photon, and graviton. Consequently, the range of the weak interaction is extremely short, so much so that it was initially treated as a point interaction. The other three forces have an effectively infinite range, though the strong force never has the opportunity to act at long distances because of color confinement.

The third and final group of particles in the Standard Model is the quarks, which make up hadrons. As with the leptons, there are six quarks grouped in pairs based on charge, rather than flavor as with the leptons. The “upper” quarks all have charge $\frac{2}{3}e$,

Field	Strong	Electromagnetic	Weak	Gravity
Particle	gluon	photon	W^\pm, Z^0	graviton
Coupling Constant	α_s	α	G_F	G_N
Scale	10^{-1}	10^{-2}	10^{-5} GeV^{-2}	$10^{-11} \frac{\text{m}^3}{\text{kg}\cdot\text{s}^2}$

Table 2.1: Fields and their associated bosons and coupling constants

and the “lower” quarks have charge $-\frac{1}{3}e$.

$$\begin{pmatrix} u \\ d \end{pmatrix} \begin{pmatrix} c \\ s \end{pmatrix} \begin{pmatrix} t \\ b \end{pmatrix}$$

The quarks and leptons have further similarities. Each family is more massive than the previous one (so that the (τ, ν_τ) and (t, b) families are heaviest), and both the leptons and quarks mix in weak interactions. For the former, the flavor couplings are encompassed by the MSN matrix, as discussed in Section 2.3, and the CKM matrix encompasses the couplings for the latter [48].

The weak interaction can be viewed as a “current-current” interaction, with two categories of currents participating, neutral currents and charged currents. The Z^0 couples neutral currents, such as ee and $\nu_e\nu_e$; this represents neutrino-electron scattering. A particular charged current pair, and one of fundamental importance to this thesis, is $\nu_\mu\mu^-$ and ud (see Figure 2.2). The weak interaction violates numerous conservation laws, including charge conjugation parity (C), space parity (P), combined CP parity, charm, and strangeness [61], all of which are conserved by the strong and electromagnetic interactions.

The neutrinos range from at least two to five orders of magnitude lighter than their charged leptonic counterparts; the Standard Model has no explanation for this wild range of lepton mass, nor for why the neutrino masses should be close to or identically zero. Several interesting physics questions are intertwined with the specter of neutrino mass, and as such much experimental effort has gone toward measuring the neutrino

masses. Direct kinematic measurements, such as the beta decay spectrum of tritium, have yielded only upper limits for the masses. Oscillation experiments can provide only mass-squared differences.

2.2.1 Neutrino Production and Interactions

The leptons do not interact strongly, and the neutrinos are electrically neutral and so do not interact electromagnetically. Thus the weak interaction is responsible for all known neutrino physics. Gravity is extremely feeble, even when compared to the “weak” interaction, and can be neglected.

Neutrinos are born in a multitude of weak decays, such as beta decay,

$$n \rightarrow p + e^- + (\bar{\nu}_e) \quad (2.1)$$

muon decay,

$$\mu^\pm \rightarrow e^\pm + \nu_e(\bar{\nu}_e) + \bar{\nu}_\mu(\nu_\mu) \quad (2.2)$$

and the decays of the π and K mesons.

$$\pi^\pm \rightarrow \mu^\pm + \nu_\mu(\bar{\nu}_\mu) \quad (2.3)$$

$$K^\pm \rightarrow \mu^\pm + \nu_\mu(\bar{\nu}_\mu) \quad (2.4)$$

$$K_L^0 \rightarrow \pi^\pm + \mu^\mp + \bar{\nu}_\mu(\nu_\mu) \quad (2.5)$$

$$K_L^0 \rightarrow \pi^\pm + e^\mp + \bar{\nu}_e(\nu_e) \quad (2.6)$$

It is possible within the Standard Model to calculate these decay widths. For example,

$$\Gamma(\mu^+ \rightarrow e^+ \nu_e \bar{\nu}_\mu) = \frac{G_F^2 m_\mu^5}{192\pi^3} \quad (2.7)$$

and

$$\Gamma(\pi^+ \rightarrow \mu^+ \nu_\mu) = \frac{G_F^2 |V_{ud}|^2}{8\pi} f_\pi^2 m_\mu^2 m_\pi \left(1 - \frac{m_\mu^2}{m_\pi^2}\right)^2 \quad (2.8)$$

where V_{ud} is the up-down CKM matrix element and f_π is the pion weak decay constant, which has a value of 131 MeV [48].

Once a neutrino is produced, it may interact with other particles it encounters. Such interactions can fall into two broad categories, neutral current or charged current. In the former, the neutrino scatters off of another particle by exchanging the neutral Z^0 and continues on its way, as shown in Figure 2.1. In the latter, the neutrino exchanges the charged W^\pm with another particle, thus coupling to its partner lepton, demonstrated in Figure 2.2. In both cases, the struck particle is altered, either gaining momentum or transforming into another particle. Further, if the neutrino interacts with a quark, other hadrons may be produced.

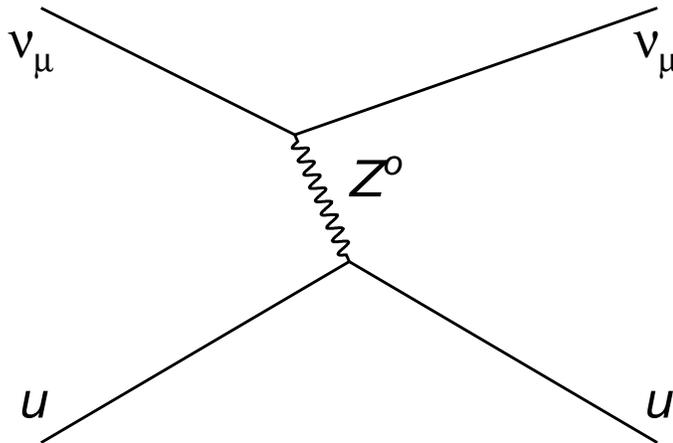


Figure 2.1: Neutral current ν interaction with an up quark

The charged current interaction is actually subdivided into three categories, quasi-elastic (QE), resonance production (RES), and deep inelastic scattering (DIS), based on the hadronic production. In the first, the outgoing lepton takes essentially all of the neutrino energy and the struck nucleon is left with very little kinetic energy. In the second, the hadronic portion of the interaction gets just the right amount of energy to

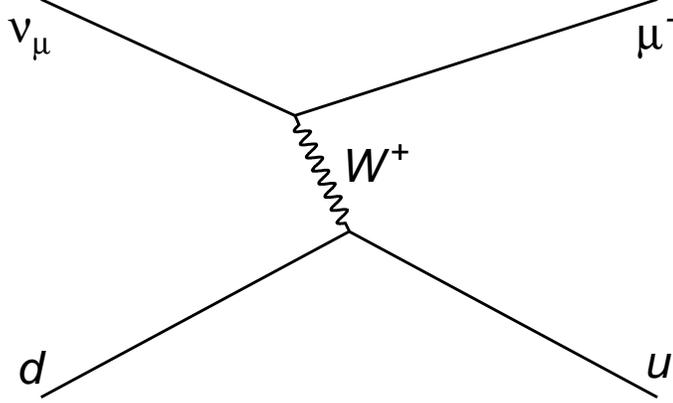


Figure 2.2: Charged current ν_μ interaction

produce, for example, a Δ resonance which would yield a single pion in addition to a nucleon in the final state. In the last, a hadronic shower is produced which carries away much of the incoming neutrino energy. These three different regimes of charged current interaction have different dependencies on the neutrino energy, so that below ~ 500 MeV, the QE portion is totally dominant. Above this point, the RES cross section begins to rise and the QE cross section starts to fall. At 1 GeV, the DIS cross section begins to rise, and becomes totally dominant by 10 GeV, where the other two have dropped off significantly. In between 1 and 10 GeV, all three have non-negligible contributions to the total cross section. Figure 2.3 shows these cross sections plotted against neutrino energy, along with selected experimental data.

All neutrino interactions, neutral or charged current, are extremely unlikely, as the neutrino interaction cross section is so fantastically small. The total neutrino interaction cross section around 1 GeV is $O(10^{-42})$ m²; juxtapose that value with the cross-sectional area of a typical nucleus, $O(10^{-30})$ m².

Neutrino interactions are, ignoring the internal structure of any subsequent hadronic

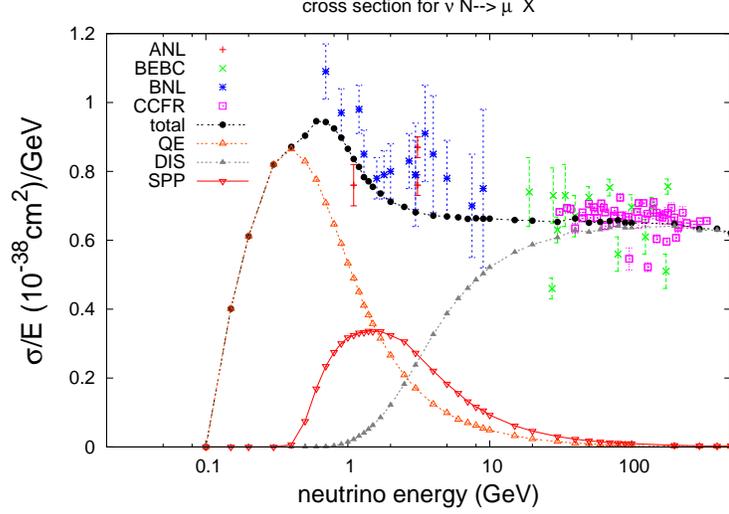


Figure 2.3: Total neutrino charged current cross section, along with the QE , RES (Single Pion Production), and DIS components [60].

shower, two-body to two-body processes. Such processes are described in the Mandelstam variables:

$$s \equiv (p_1 + p_2)^2 = (p_3 + p_4)^2 \quad (2.9)$$

$$t \equiv (p_1 - p_3)^2 = (p_2 - p_4)^2 \quad (2.10)$$

$$u \equiv (p_1 - p_4)^2 = (p_2 - p_3)^2 \quad (2.11)$$

p_1 and p_2 are the four momenta of the incoming particles, such as the ν_μ and the d quark in Figure 2.2, and p_3 and p_4 are the four momenta of the outgoing particles, such as the μ and the u quark in Figure 2.2. The Mandelstam variables are quite convenient to describe two-body to two-body processes because, not only are they Lorentz invariant (and thus can be evaluated in any reference frame), but they fully encompass all available, independent kinematic information.

Two other kinematical variables often used in neutrino interaction calculations are y_{kin} and the Bjorken x_B . To expound on these variables, it is necessary to first introduce three more useful quantities, Q^2 , W^2 , and ν_{kin} . First, Q^2 is the square of the four

momentum of the mediator boson in the interaction,

$$Q^2 = -(p_1 - p_3)^2 = -t^2 \quad (2.12)$$

where the notation defined for the Mandelstam variables (above) has been used. Second, W is simply the invariant mass of the hadronic products in the interaction. Third, ν_{kin} is defined as

$$\nu_{kin} = \frac{W^2 + Q^2 - M^2}{2M} \quad (2.13)$$

where M is the proton or neutron mass. ν_{kin} , evaluated in the nucleon rest frame, is equal to the energy lost by the leptonic part of the interaction, $E_1 - E_3$ (the incoming neutrino energy minus the outgoing lepton energy). This statement ignores nuclear effects and the initial nucleon state.

With those variables defined, y_{kin} and x_B can now be tackled. y_{kin} is the fraction of the total neutrino energy taken by the hadronic portion of the interaction products, defined as

$$y_{kin} = \frac{E_H}{E_\nu} \quad (2.14)$$

E_H is illustrated graphically in Figure 2.4, and is the total energy taken away from the interaction by the hadronic products. y_{kin} is easily calculated in the lab frame as ν_{kin}/E_ν . y_{kin} is interesting because the neutrino-quark interaction cross section is flat in y_{kin} , but the neutrino-antiquark cross section goes like $(1 - y_{kin})^2$. Similarly, the antineutrino-quark cross section goes like $(1 - y_{kin})^2$ but the antineutrino-antiquark cross section is flat in y_{kin} . These dependencies are beautiful manifestations of the $V - A$ charged current [48]. The Bjorken x_B is defined as

$$x_B = \frac{Q^2}{2M\nu_{kin}} \quad (2.15)$$

The structure functions used to parametrize the hadronic production in lepton-nucleon scattering are, for a given Q^2 , only weakly dependent on x_B and insensitive to an overall

change in scale (as in the masses of the particles). This behavior is the essence of Bjorken scaling.

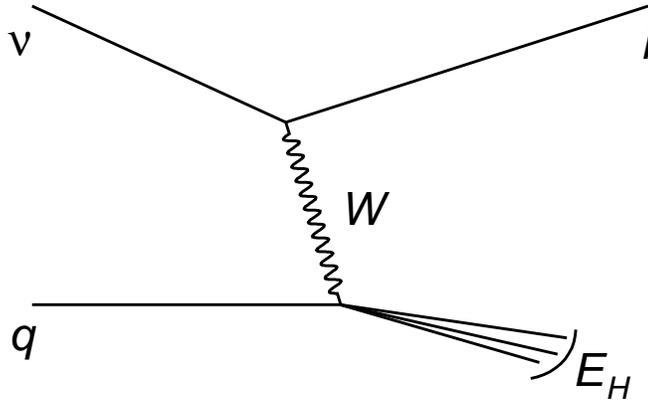


Figure 2.4: A charged current neutrino interaction showing the hadronic shower. q represents a quark (u or d), l represents the lepton (e , μ , or τ) which results from the neutrino interaction, and E_H represents the total energy taken away by the hadronic products.

2.3 Oscillation Physics

Forty-two years ago, Raymond Davis and John Bahcall endeavored to observe solar electron neutrinos in an effort to confirm the contemporary theoretical model describing energy generation in main-sequence stars. The ensuing twenty years of experimental observation served to contradict the predictions of this and all later models. Specifically, all fusion models predicted a higher flux of neutrinos than was actually observed.

One proposed solution to this conundrum was that the neutrinos were somehow changing during transit from the Sun to the Earth. It was suspected that all neutrinos originating from the Sun were electron neutrinos (ν_e). If, by some mechanism, the ν_e changed flavor, they would not interact to produce their charged partner because their

low energy would leave them below the kinematic limits to produce a muon or tau, and hence a deficit would be observed. Thus the idea of neutrino oscillation was forwarded to explain the solar neutrino problem.

The idea of neutrino oscillation was not new, however. In his seminal paper [65], Bruno Pontecorvo theoretically described neutrino oscillation.

Twenty years after Davis and Bahcall, proton decay experiments began operating all over the world. One of the secondary capabilities of these experiments was atmospheric neutrino observation. Specifically, the proton decay experiments measured the ratio of atmospheric ν_μ to ν_e , and most experiments observed a lower ratio than expected. The lower ratio translated into a deficit of neutrinos, as in the solar case, but this time of ν_μ . Once again, oscillation was forwarded as a possible explanation.

The physics of neutrino oscillation is simply quantum mechanical mixing of weak interaction eigenstates. To begin, consider the usual three flavor scheme, where U is a complex and unitary matrix.

$$\begin{pmatrix} \nu_e \\ \nu_\mu \\ \nu_\tau \end{pmatrix} = \begin{pmatrix} U_{e1} & U_{e2} & U_{e3} \\ U_{\mu1} & U_{\mu2} & U_{\mu3} \\ U_{\tau1} & U_{\tau2} & U_{\tau3} \end{pmatrix} \begin{pmatrix} \nu_1 \\ \nu_2 \\ \nu_3 \end{pmatrix} \quad (2.16)$$

The flavor eigenstates ν_e , ν_μ , and ν_τ are expressed as a linear combination of the three mass eigenstates ν_1 , ν_2 , and ν_3 . These mass eigenstates are the eigenstates of the free Hamiltonian, and thus as a neutrino propagates through space-time, it is the time evolution of these states which is specified directly. Each flavor eigenstate is a different admixture of the three mass eigenstates, governed by the U matrix, and each component in the admixture evolves differently. As such, a neutrino created in one flavor eigenstate can oscillate into another flavor eigenstate at a later time.

The mixing matrix can be represented in many ways. As a consequence of the

experimental knowledge acquired about neutrinos and their couplings, a particular parametrization of the mixing matrix which factorizes according to the various types of neutrino oscillations has become popular.

$$\begin{pmatrix} \nu_e \\ \nu_\mu \\ \nu_\tau \end{pmatrix} = \begin{pmatrix} 1 & 0 & 0 \\ 0 & c_{23} & s_{23} \\ 0 & -s_{23} & c_{23} \end{pmatrix} \begin{pmatrix} c_{13} & 0 & s_{13}e^{-i\delta} \\ 0 & 1 & 0 \\ -s_{13}e^{-i\delta} & 0 & c_{13} \end{pmatrix} \times \\ \begin{pmatrix} c_{12} & s_{12} & 0 \\ -s_{12} & c_{12} & 0 \\ 0 & 0 & 1 \end{pmatrix} \begin{pmatrix} e^{i\alpha_1/2} & 0 & 0 \\ 0 & e^{i\alpha_2/2} & 0 \\ 0 & 0 & 1 \end{pmatrix} \begin{pmatrix} \nu_1 \\ \nu_2 \\ \nu_3 \end{pmatrix} \quad (2.17)$$

Here, c_{ij} and s_{ij} respectively denote $\cos\theta_{ij}$ and $\sin\theta_{ij}$, the mixing angles between eigenstates and δ , α_1 , and α_2 are CP violating phases. The latter two are Majorana phases and only have a physical consequence should neutrinos be Majorana particles. Additional phases are irrelevant and have been subsumed in the definitions of the lepton fields. The factorization of the mixing matrix is illustrative in the following way; the first piece demonstrates mixing between the second and third mass eigenstates (atmospheric neutrino oscillations), the third piece demonstrates mixing between the first and second mass eigenstates (solar neutrino oscillations), the second piece encapsulates cross mixing, and the fourth piece accounts for the possibility of neutrinos being their own antiparticles (neutrinoless double beta decay, for example).

2.3.1 Two Flavor Oscillation Phenomenology

In most experimental situations, a neutrino is produced somewhere, be it in a nuclear reactor [31], the Sun [49], or an anthropogenic neutrino beam (MINOS), and it is detected somewhere else via its charged current (CC) or neutral current (NC) interaction. To obtain the probability of observing a neutrino of type b (e, μ, τ) when it was produced

as a neutrino of type a (e, μ, τ), consider the initial state

$$|\nu_a\rangle = U_{a1}|\nu_1\rangle + U_{a2}|\nu_2\rangle + U_{a3}|\nu_3\rangle \quad (2.18)$$

Because each state $|\nu_1\rangle, |\nu_2\rangle$, and $|\nu_3\rangle$ are eigenstates of the free Hamiltonian, their time dependence is encompassed by the usual phase factor $e^{iE_{1,2,3}t}$, where we have adopted natural units ($\hbar = c = 1$). Thus,

$$|\nu_a(t)\rangle = U_{a1}e^{iE_1t}|\nu_1\rangle + U_{a2}e^{iE_2t}|\nu_2\rangle + U_{a3}e^{iE_3t}|\nu_3\rangle \quad (2.19)$$

The final state in question is

$$|\nu_b\rangle = U_{b1}|\nu_1\rangle + U_{b2}|\nu_2\rangle + U_{b3}|\nu_3\rangle \quad (2.20)$$

The probability of oscillation is the absolute square of the transition amplitude, found by projecting the initial state onto the final state.

$$P(\nu_a \rightarrow \nu_b) = |A(\nu_a \rightarrow \nu_b)|^2 = |\langle \nu_b | \nu_a(t) \rangle|^2 \quad (2.21)$$

and

$$\begin{aligned} |\langle \nu_b | \nu_a(t) \rangle|^2 = & \left| \left(U_{a1}^* U_{b1} e^{-iE_1t} + U_{a2}^* U_{b2} e^{-iE_2t} + U_{a3}^* U_{b3} e^{-iE_3t} \right) \right|^2 \times \\ & \left| \left(U_{a1} U_{b1}^* e^{iE_1t} + U_{a2} U_{b2}^* e^{iE_2t} + U_{a3} U_{b3}^* e^{iE_3t} \right) \right|^2 \end{aligned} \quad (2.22)$$

Usually, physicists consider a two flavor oscillation scheme. The mathematics is much simpler, and because of the small coupling between mass eigenstates 1 and 3, a two flavor scheme is a very good approximation for either $\nu_e \rightarrow \nu_\mu$ or $\nu_\mu \rightarrow \nu_\tau$. Within this constrained model, there is only one mixing angle for each of these oscillation modes, typically called θ . As before, consider two arbitrary flavors a and b .

$$\begin{pmatrix} \nu_a \\ \nu_b \end{pmatrix} = \begin{pmatrix} \cos \theta_{ab} & \sin \theta_{ab} \\ -\sin \theta_{ab} & \cos \theta_{ab} \end{pmatrix} \begin{pmatrix} \nu_1 \\ \nu_2 \end{pmatrix}$$

Again, consider a neutrino produced as flavor a . The initial state is given by

$$|\nu_a\rangle = \cos \theta_{ab} |\nu_1\rangle + \sin \theta_{ab} |\nu_2\rangle \quad (2.23)$$

The time dependent expression for this state is

$$|\nu_a(t)\rangle = \cos \theta_{ab} e^{iE_1 t} |\nu_1\rangle + \sin \theta_{ab} e^{iE_2 t} |\nu_2\rangle \quad (2.24)$$

The final state, $|\nu_b\rangle$, is expressed as

$$|\nu_b\rangle = -\sin \theta_{ab} |\nu_1\rangle + \cos \theta_{ab} |\nu_2\rangle \quad (2.25)$$

The probability of oscillation from flavor a to flavor b is again found by projecting the initial state onto the final state and taking the absolute square.

$$P(\nu_a \rightarrow \nu_b) = |A(\nu_a \rightarrow \nu_b)|^2 = |\langle \nu_b | \nu_a(t) \rangle|^2 \quad (2.26)$$

Explicitly,

$$|\langle \nu_b | \nu_a(t) \rangle|^2 = |-\sin \theta_{ab} \cos \theta_{ab} e^{iE_1 t} + \sin \theta_{ab} \cos \theta_{ab} e^{iE_2 t}|^2 \quad (2.27)$$

Taking the absolute square and simplifying,

$$P(\nu_a \rightarrow \nu_b) = \sin^2 \theta_{ab} \cos^2 \theta_{ab} \left(2 - e^{i(E_1 - E_2)t} - e^{i(E_1 - E_2)t} \right) \quad (2.28)$$

Expanding the exponentials and collecting terms,

$$P(\nu_a \rightarrow \nu_b) = \sin^2(2\theta_{ab}) \sin^2\left(\frac{(E_1 - E_2)t}{2}\right) \quad (2.29)$$

At this point, some approximations can be made. Note that

$$E_1 = \sqrt{p_1^2 + m_1^2} \quad (2.30)$$

Because of the smallness of the neutrino mass, the square root can be expanded about $m_1^2/p_1^2 = 0$ with a high degree of accuracy.

$$p_1 \sqrt{1 + \frac{m_1^2}{p_1^2}} \approx p_1 \left(1 + \frac{m_1^2}{2p_1^2} + \dots \right) \quad (2.31)$$

Substituting this approximation in for the two energy terms yields

$$P(\nu_a \rightarrow \nu_b) = \sin^2(2\theta_{ab}) \sin^2\left(\frac{t}{2} \left[p_1 + \frac{m_1^2}{2p_1} - p_2 - \frac{m_2^2}{2p_2} \right]\right) \quad (2.32)$$

To further the calculation, it must now be assumed that the different mass components of a neutrino produced as a flavor eigenstate have either the same momentum or the same energy. Because the mass of each mass eigenstate is different, and because their momenta and energies cannot both be equal (see Equation 2.30), a choice must be made. As it turns out for relativistic neutrinos, either choice will yield the correct result. In truth, neither assumption is wholly correct. The neutrino production process is localized in both space and time and, because of quantum mechanical uncertainty, neither the momentum nor the energy are precisely defined [7] [25].

Subtleties aside, assume that the momenta are equal, or $p_1 = p_2 = p$. This assumption yields

$$P(\nu_a \rightarrow \nu_b) = \sin^2(2\theta_{ab}) \sin^2\left(\frac{\Delta m_{12}^2 t}{4p}\right) \quad (2.33)$$

where $\Delta m_{12}^2 = m_1^2 - m_2^2$. The next step is to make the approximation $E \approx p$, which again is very accurate for relativistic neutrinos. Note that, earlier, the approximation $E \approx p + m^2/2p$ was made. Dropping the second term in the sum is, for a neutrino with $p = 1 \text{ GeV}$ and $m = 1 \text{ eV}$ (around the current direct upper limit on the electron neutrino mass), sacrificing accuracy at the level of 10^{-19} . Another consequence of the extreme relativistic nature of the neutrino is that the time it travels, t , is equal to the distance it travels, typically called L . Setting $t = L$ gives

$$P(\nu_a \rightarrow \nu_b) = \sin^2(2\theta_{ab}) \sin^2\left(\frac{\Delta m_{12}^2 L}{4E}\right) \quad (2.34)$$

Finally, the argument of the second sin term, which as shown above has units of *Energy*·*Length*, must be made dimensionless by insertion of the appropriate factor of $\hbar c$.

$$\frac{\Delta m_{12}^2 (\text{eV}^2) L (\text{km})}{4E (\text{GeV}) \hbar c} \cdot \frac{\frac{1 \text{ MeV}^2}{10^{12} \text{ eV}^2} \cdot \frac{10^{18} \text{ fm}}{1 \text{ km}}}{\frac{10^3 \text{ MeV}}{1 \text{ GeV}} 197.327 \text{ MeV} \cdot \text{fm}} = 1.27 \Delta m_{12}^2 \frac{L}{E} \quad (2.35)$$

This final step yields the well known formula

$$P(\nu_a \rightarrow \nu_b) = \sin^2(2\theta_{ab}) \sin^2\left(1.27\Delta m_{ab}^2 \frac{L}{E}\right) \quad (2.36)$$

where P is the probability of oscillation, θ_{ab} is the mixing angle, Δm_{ab}^2 is the mass squared difference between neutrino types (in eV^2), L is the distance traveled (in km), and E is the neutrino energy (in GeV).

2.4 Atmospheric Neutrino Summary

In the experimental particle physics community, three different types of neutrinos, distinguished by their source, have received much attention in recent decades. Solar neutrinos are neutrinos generated in the sun by its various thermonuclear processes. An observed deficit of solar neutrinos was physicists' first indication that neutrinos could be oscillating. Anthropogenic neutrinos are copiously produced at nuclear power plants, and these were the first neutrinos to be directly detected. More recently, man-made neutrinos at LSND and K2K have yielded weak evidence for neutrino oscillation. Atmospheric neutrinos, which are neutrinos produced by the interaction of cosmic rays with the Earth's atmosphere, have provided possibly the most compelling evidence that neutrinos oscillate between flavors. These neutrinos will be the focus of the rest of this section.

As mentioned, atmospheric neutrinos are generated when cosmic rays interact with the terrestrial atmosphere. These cosmic rays are composed of many nuclei and charged particles, but are predominantly hydrogen and helium. As shown in Figure 2.5, the momentum spectrum of primary cosmic rays roughly obeys a power law spectrum which falls with energy. Also, because of the magnetic field generated by the Earth, cosmic rays must have a certain minimum energy to reach the atmosphere; this minimum energy increases as latitude decreases, so that, for example, cosmic rays with energies down to

$\sim 3 \text{ GeV}$ can reach the atmosphere above Soudan, but only cosmic rays with energies over $\sim 10 \text{ GeV}$ can reach the atmosphere above Kamiokande [17]. When these cosmic rays strike the atmosphere, they produce hadronic showers. The hadrons thus produced then decay, producing leptons. The dominant processes which produce atmospheric neutrinos are shown in Equations 2.3, 2.4, 2.5, and 2.6.

The $\mu \nu_\mu$ decay channel is dominant for both the pion and kaon, having a branching ratio of roughly 100% for the former and 63.5% for the latter. For the K_L^0 , the $\pi\mu\nu$ decay channel has a branching ratio of 27.2% and the $\pi e\nu$ channel has a ratio of 38.8%. Further, the muons produced in these reactions decay via Equation 2.2, producing another muon and electron neutrino. This muon decay mode has a branching ratio of effectively 100%. The only other decays include a radiated virtual or real photon in addition to the $e\nu_e\bar{\nu}_\mu$.

These neutrinos have, like the cosmic ray flux, a power law momentum spectrum that falls with energy. See Figure 2.6.

Spatially, the atmospheric neutrino flux is roughly uniform across the sky. This makes sense because the flux of primary cosmic rays is also isotropic over all solid angles. Because of the Earth's magnetic field, however, there is an east-west asymmetry. See [40].

2.5 The Atmospheric Neutrino Problem

The production of atmospheric neutrinos is dominated by the processes denoted in Equations 2.3, 2.4, and 2.2. Upon inspection, one easily sees that muon and electron neutrinos should, to first order, come in a ratio of $2 :: 1$. Several experiments have found that this is not the case.

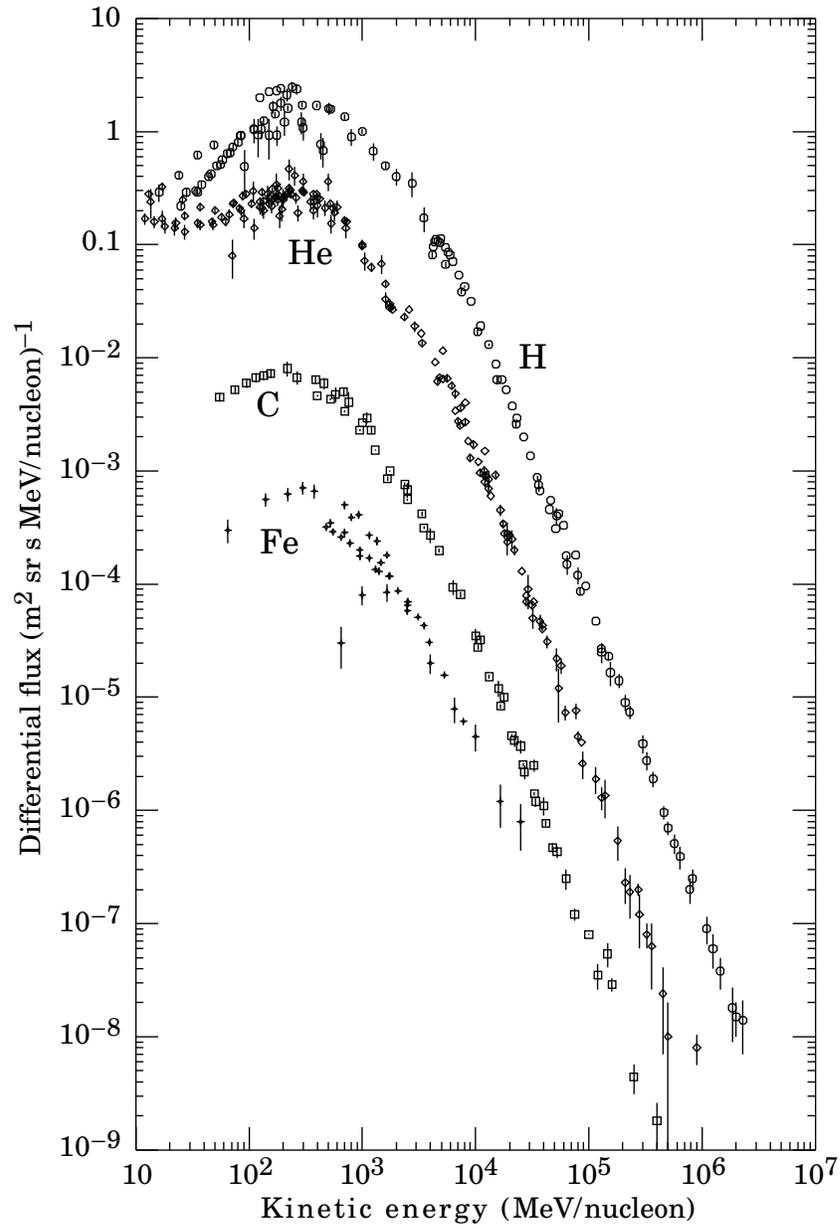


Figure 2.5: Flux of primary cosmic radiation from various balloon experiments, and the Alpha Magnetic Spectrometer [34].

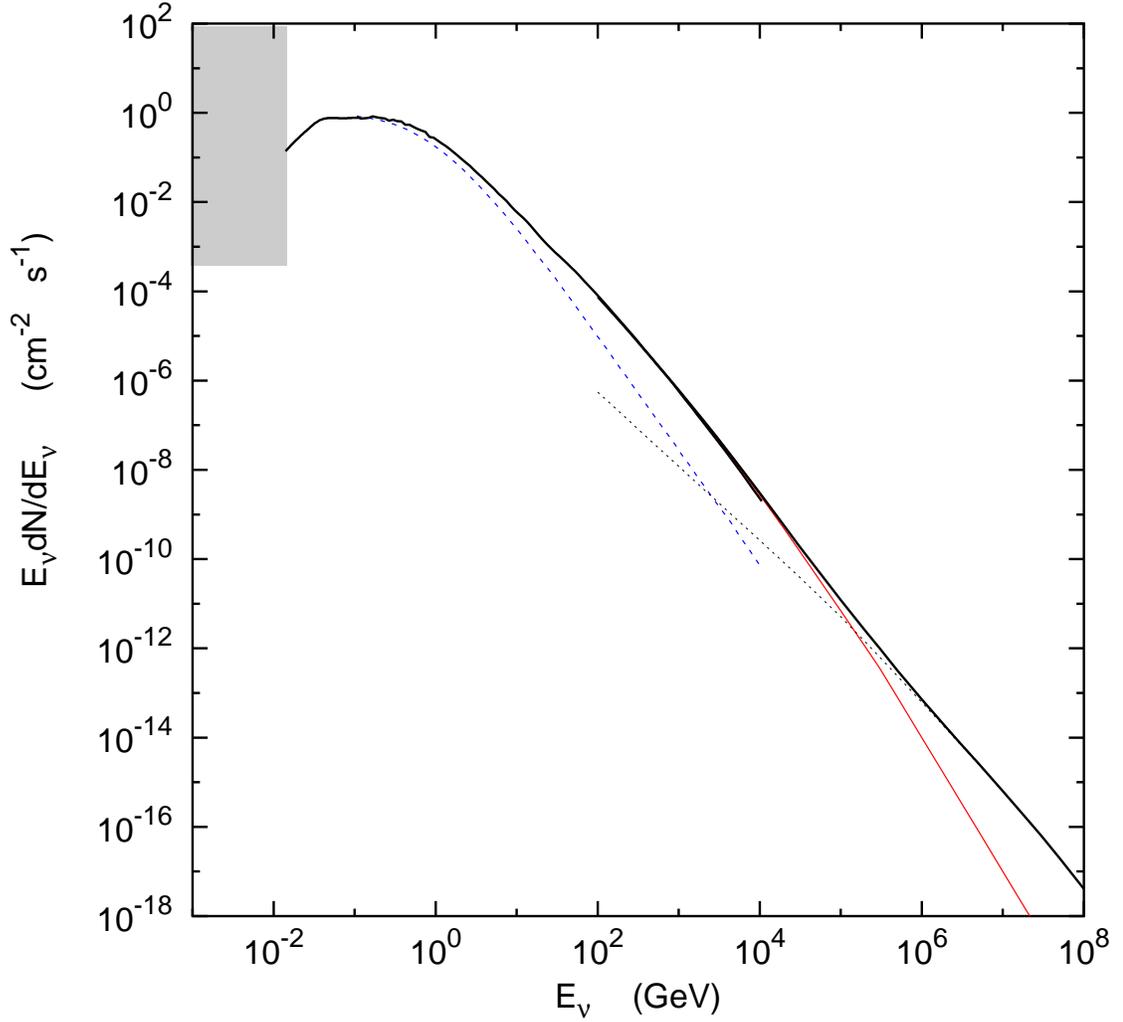


Figure 2.6: Total atmospheric neutrino flux. The solid black line shows $\nu_\mu + \bar{\nu}_\mu$, the dashed blue line shows $\nu_e + \bar{\nu}_e$, the gray dotted line shows the contribution of neutrinos from charmed meson decay, and the red solid line shows $\nu_\mu + \bar{\nu}_\mu$ from kaon and pion (i.e. uncharged meson) decay. At very low energies (inside the gray box), solar neutrinos dominate [40].

One quantity that most experiments have measured is R , which is a ratio of ratios.

$$R = \frac{\nu_\mu/\nu_e |_{\text{measured}}}{\nu_\mu/\nu_e |_{\text{Monte Carlo}}} \quad (2.37)$$

This rather complicated representation of the relative number of muon and electron neutrinos is used because the absolute atmospheric neutrino flux is not well understood and using the ratio of the number of muon to electron neutrinos eliminates most of the uncertainty that this flux brings into calculations [42].

2.5.1 Kamiokande

Kamiokande was a 4.5 kt (*kiloton*) water Čerenkov detector located at a depth of 2700 m.w.e. (meters water equivalent) in the Kamioka Mine, Japan. Essentially, it was a cylindrical tank of ultrapure water lined with photomultiplier tubes which were used to collect Čerenkov radiation from any relativistic particles traveling in the detector. Specifically, neutrinos were mainly detected via the interaction

$$\nu_l + N \rightarrow l + X \quad (2.38)$$

In both processes, the outgoing lepton (e or μ) emits a single ring of Čerenkov radiation and, through a pattern recognition algorithm, the lepton’s flavor is determined. The pattern recognition algorithm accomplishes this flavor tagging by looking at the characteristics of the Čerenkov cone. A muon’s cone has well-defined edges because it typically does not produce other particles in traversing the detector, while an electron’s cone has “fuzzy” edges due to the electromagnetic showering it induces; these are the differences that the algorithm seeks out. Through PMT timing, the lepton’s direction could be found, and by counting photoelectrons, experimenters could also measure its momentum.

The Kamiokande experiment performed several different atmospheric neutrino event analyses. First, they calculated R directly from all single ring events, using only the

flavor count obtained from the pattern recognition algorithm, as a function of zenith angle. Second, they calculated the fraction of muon single ring events by looking for muon decay. This fraction of muon events was then used to calculate R . In both cases, Kamiokande found a value of R which was significantly below 1; $R = 0.57^{+0.08}_{-0.07} \pm 0.07$ in the first analysis [39] and $R = 0.70 \pm 0.07$ in the second [47]. Further, R exhibited a zenith angle dependence, showing that there was a deficit of muon neutrinos coming through the Earth, but not from directly above. Finally, in [47], the Kamiokande collaboration made an attempt to interpret their results in the context of neutrino oscillations.

2.5.2 Fréjus

The Fréjus experiment is noteworthy because it is the only proton decay experiment which saw no evidence for neutrino oscillations in the atmospheric neutrinos. The Fréjus detector was a 900 ton iron calorimeter located in the French Alps under roughly 1400 m of rock overburden. In their atmospheric analysis, the Fréjus collaboration used the e/μ flavor ratio to divine their confidence limits in $(\sin^2(2\theta_{23}), \Delta m_{23}^2)$. They excluded the region $\sin^2(2\theta_{23}) > 0.6$, $\Delta m_{23}^2 > 3.5 \times 10^{-3} \text{ eV}^2$, and so while the Fréjus data did not suggest oscillation, it also did not exclude any of the Super-Kamiokande 99% confidence allowed region, save for a small slice at the very highest Δm_{23}^2 values (see Figure 2.9) [22]. In a modern context, the Fréjus analysis is inconclusive.

2.5.3 IMB

The IMB detector was a 6.9 kt total mass rectangular water Čerenkov detector located at 1570 m.w.e. in a salt mine near Cleveland, Ohio. Apart from its shape, IMB was similar to Kamiokande. And like Kamiokande, IMB carried out single ring and muon decay analyses to calculate R . Based on the single ring analysis, IMB found

$R = 0.54 \pm 0.05 \pm 0.12$. Using the muon decay analysis, IMB obtained $R = 0.67 \pm 0.06$ [20].

IMB went further by also focusing on upward going muons. Any muons that are upward going must have resulted from neutrino interactions in the rock around the detector since no muons are energetic enough to travel all the way through the Earth. These upward going muons were divided into stopping and through-going, and then a ratio of upward going to through-going muons was calculated. This ratio was used to constrain the $(\sin^2(2\theta_{23}), \Delta m_{23}^2)$ parameter space [21].

2.5.4 Soudan 2

The Soudan 2 detector was an iron tracking calorimeter located in the Soudan Underground Laboratory. 224 separate, identical modules made up the Soudan 2 detector. Each module was filled with 1.6 mm thick steel plates and drift tubes made of Hytrel, a slightly resistive plastic. Each drift tube was wrapped with thin copper strips spaced evenly along its length. The center strip was maintained at -9 kV, and the voltage on each subsequent strip going outward stepped down so as to create a fairly uniform electric field inside the tube. Thus, when a charged particle ionized gas inside the drift tube, ionization electrons drifted to the closest end, where they were collected by anode wires [42].

Soudan 2, like IMB and Kamiokande, measured R , finding $R = 0.64 \pm 0.11 \pm 0.06$ [8]. Soudan 2 has also performed a full oscillation analysis based on their observed $\log(L/E)$ spectrum for both the full data set, which yielded a best fit point of $\sin^2(2\theta_{23}) = 0.97$, $\Delta m_{23}^2 = 1.7 \times 10^{-3} \text{ eV}^2$ [71], and the upward-stopping muon subset, which yielded a best fit point of $\sin^2(2\theta_{23}) = 0.97$, $\Delta m_{23}^2 = 5.2 \times 10^{-3} \text{ eV}^2$ [9].

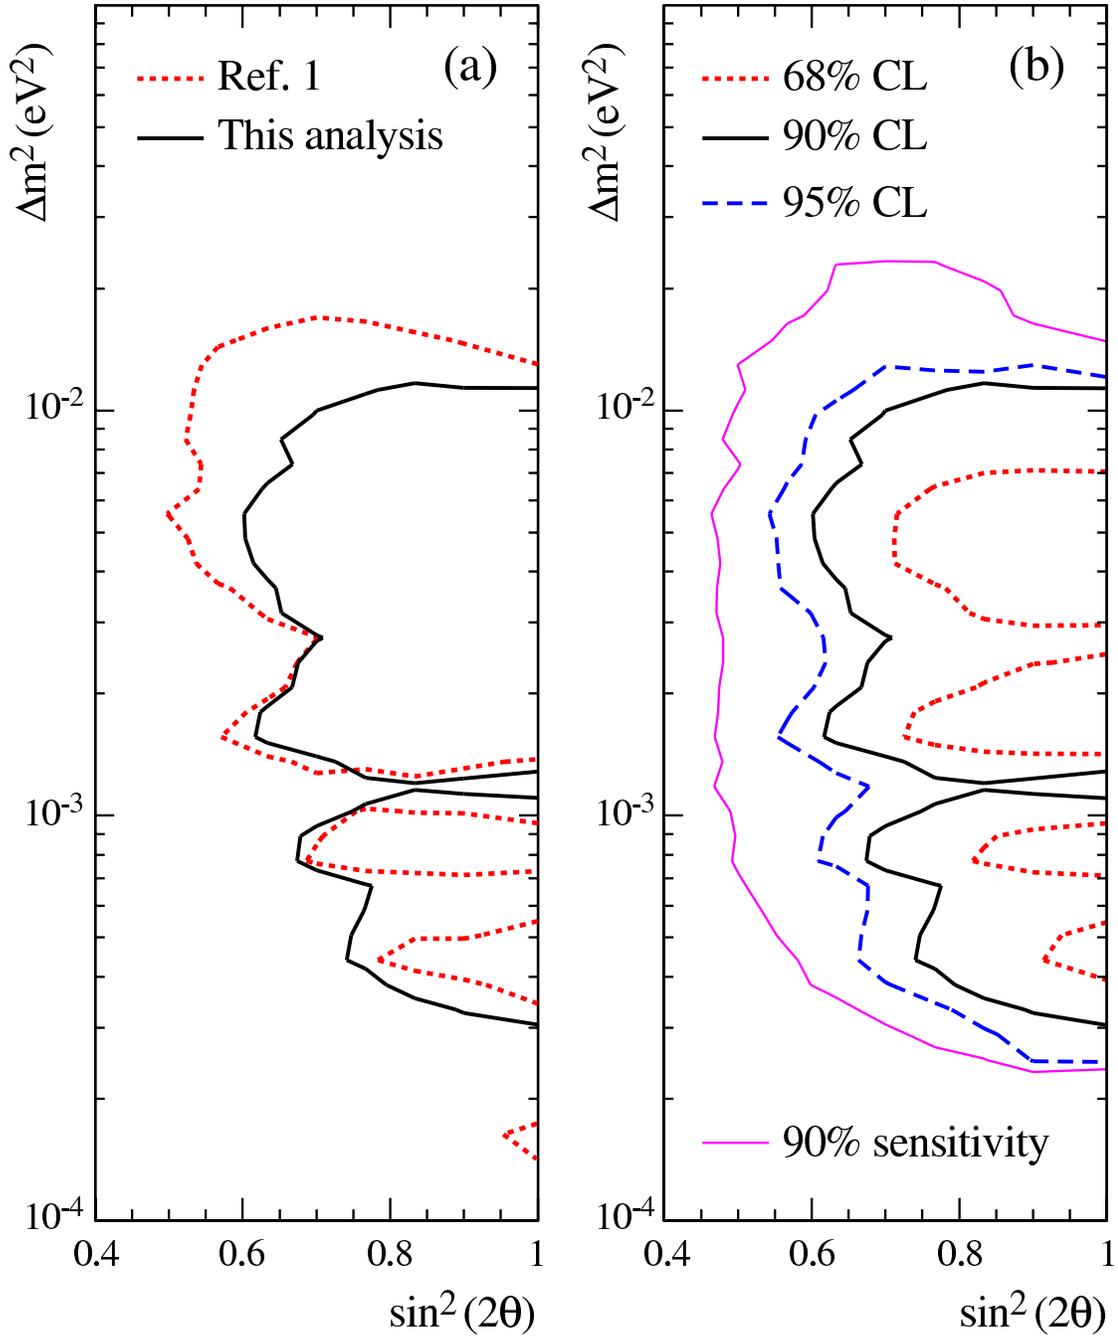


Figure 2.7: The confidence limits for the final Soudan 2 data set, based on their two separate analyses. In panel (a), the red, dashed line shows the 90% confidence allowed region for the analysis of the full data set [71], and the solid black line shows the same for the upward-stopping muon analysis [9], which is the analysis that “This analysis” refers to in the figure. In panel (b), the 68%, 90%, and 95% confidence levels for the upward-stopping muon analysis are respectively shown as a red, dotted line, a black, solid line, and a blue, dashed line. The 90% sensitivity is the outermost line, shown in purple.

2.5.5 MACRO

MACRO, the Monopole, Astrophysics and Cosmic Ray Observatory, was located in the Gran Sasso Laboratory, and it took the shape of a rectangular box. It had two detector components, streamer tubes and planes of long rectangular boxes containing liquid scintillator. MACRO also used the upward muon flux to constrain the oscillation parameter space, and found that their data favored maximal mixing with $\Delta m_{23}^2 \simeq 2.3 \times 10^{-3} \text{ eV}^2$ [44].

2.5.6 Super-Kamiokande

Super-Kamiokande is a second-generation version of the original Kamiokande detector. It is much larger, having a total mass of 50 kt and a fiducial volume of 22.5 kt, but has the same geometry and detection method. Super-K has divided its sample of fully-contained neutrino events according to energy, measuring $R = 0.658 \pm 0.016 \pm 0.035$ for “sub-GeV” neutrinos and $R = 0.702_{-0.030}^{+0.032} \pm 0.101$ for “multi-GeV” neutrinos [16].

Further, the number of muon neutrino events observed shows a strong dependence on zenith angle, manifesting itself as a growing deficiency in events as the distance traveled through the earth increases ($\cos \Theta \rightarrow -1$). The electron neutrino events, however, agree quite well with the expected flux. See Figure 2.8. This, coupled with the muon neutrino deficiency, strongly indicates that the dominant effect is $\nu_\mu \rightarrow \nu_\tau$. Super-K has also, like IMB, studied upward stopping and through-going muons produced by neutrino interactions in the rock below the detector.

The data contained in these various subsets has been used to constrain the $(\sin^2(2\theta_{23}), \Delta m_{23}^2)$ parameter space, shown in Figure 2.9.

Super-K has, by far, the highest statistics of any atmospheric neutrino experiment. Consequently, it has become the benchmark in atmospheric neutrino studies and it is worthwhile to more closely consider their analyses and results. In their “standard”

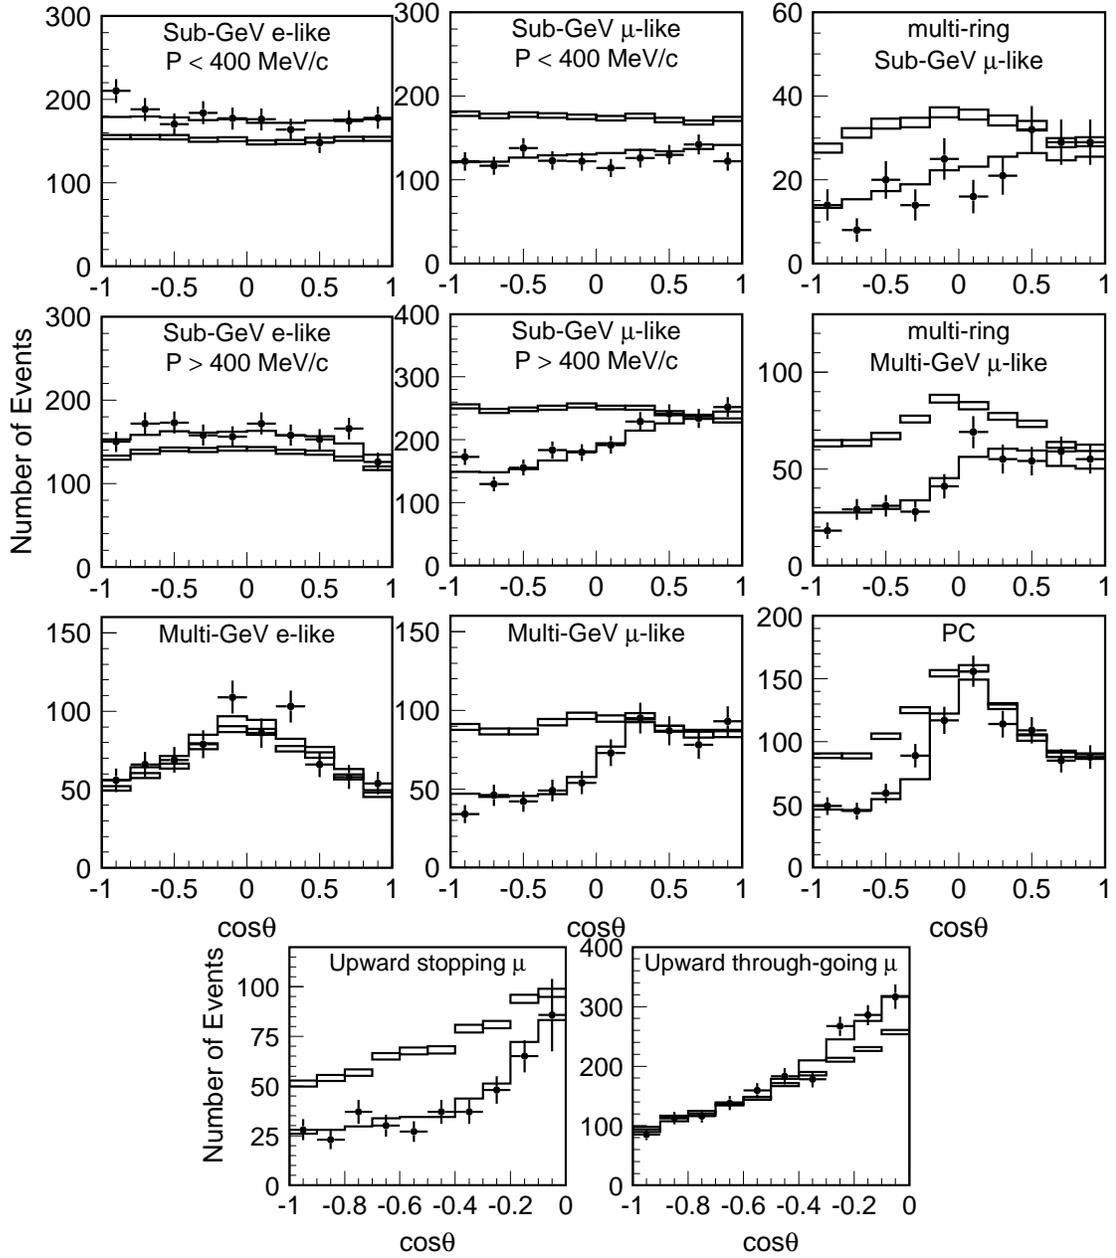


Figure 2.8: Super-Kamiokande's zenith angle distribution for fully-contained 1-ring events, multi-ring events, partially-contained events, and upward muons. The points show the data, box histograms show the non-oscillated Monte Carlo events, and the lines show the best-fit expectations for $\nu_\mu \rightarrow \nu_\tau$ oscillations with $\sin^2(2\theta_{23}) = 1.00$ and $\Delta m_{23}^2 = 2.1 \times 10^{-3} \text{ eV}^2$. The height of the boxes shows the statistical error of the Monte Carlo [16].

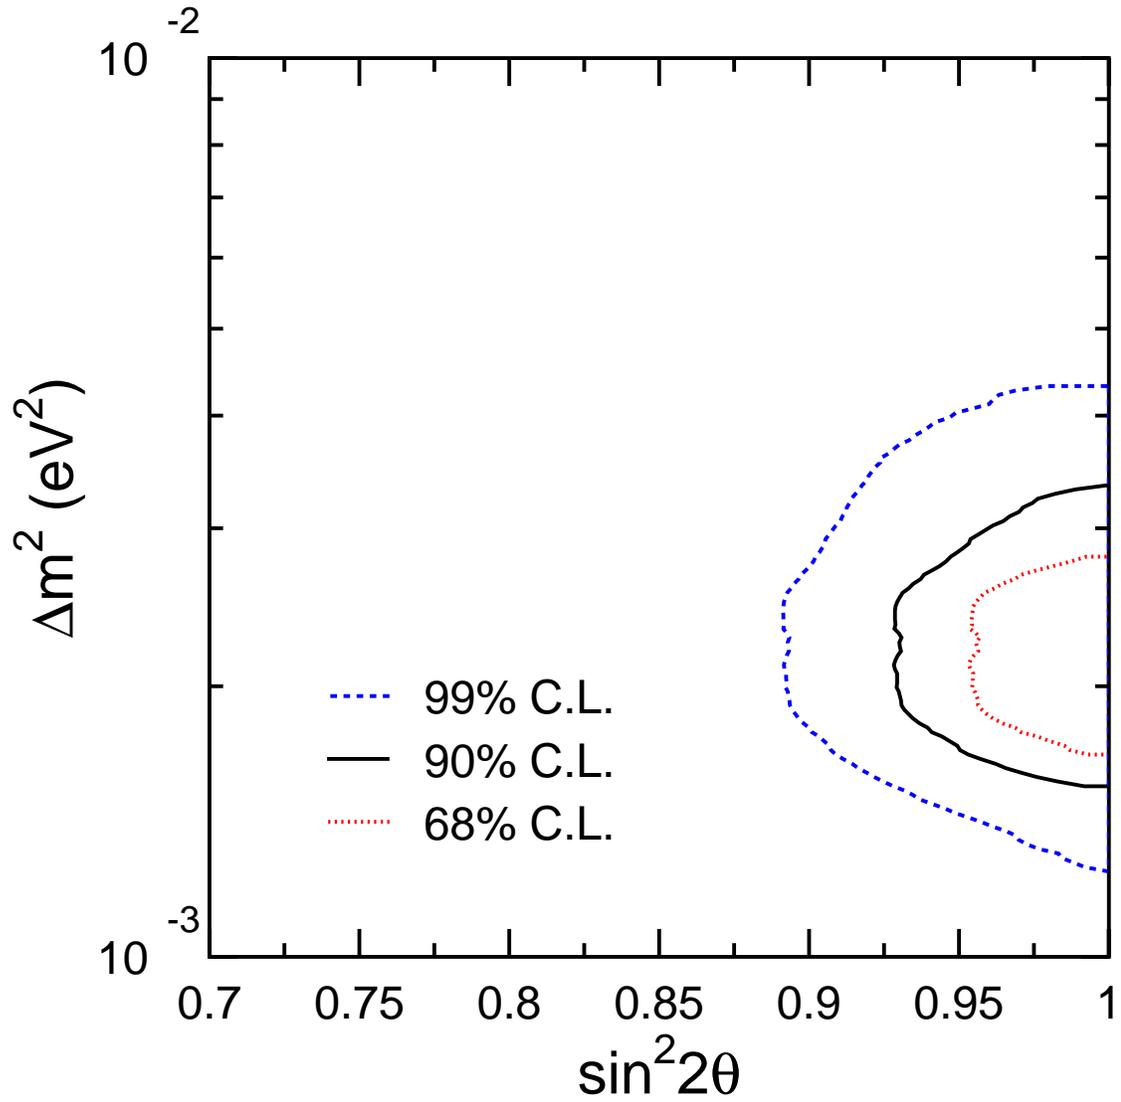


Figure 2.9: Super-Kamiokande's allowed oscillation parameters for $\nu_\mu \rightarrow \nu_\tau$ oscillations. Three contours correspond to the 68% (dotted line), 90% (solid line), and 99% (dashed line) C.L. allowed regions [16].

atmospheric neutrino oscillation analysis, they fit the zenith angle distributions of the various classes of events they separate out (fully contained, partially contained, single ring, multi-ring, upward stopping muons, etc.) by forming a χ^2 statistic and searching for the χ^2_{min} minimum on a $(\sin^2(2\theta_{23}), \Delta m^2_{23})$ grid. They have also calculated the ratio of ratios, shown in Equation 2.37. Both of these are nearly universal among the existing atmospheric neutrino oscillation analyses.

Recently, Super-K has published some novel results [15]. Specifically, their high statistics have allowed them to somewhat convincingly demonstrate the “dip” in the $\log(L/E)$ spectrum, a global first for atmospheric neutrinos. In the ratio $(L/E)_{Data}/(L/E)_{MC}$ as a function of $\log(L/E)$ for unoscillated MC, it is apparent that a dip should manifest itself at whatever $\log(L/E)$ maximizes the term $\sin^2(1.27\Delta m^2 L/E)$ in the oscillation probability. As shown in Figure 2.10, Super-K’s data exhibits (perhaps with an imaginative eye) a dip around ~ 500 km/GeV. The location of the dip is driven exclusively by Δm^2_{23} , and the best fit parameters from this analysis are almost identical to the best fit parameters from their zenith angle fit.

2.5.7 Beam Experiments

Currently NuMI/MINOS is running, and K2K recently finished; both were designed to investigate neutrino oscillation at the same Δm^2 relevant to atmospheric neutrinos using a neutrino beam. Both also employ a similar strategy. In both cases, a very pure muon neutrino beam is produced and passed through a “near” detector. This near detector is very close to the origin of the beam and monitors the beam’s flavor content and energy spectrum. Both experiments then use a “far” detector which is also in the beam’s path but at a distance of ~ 250 km for K2K and ~ 735 km for NuMI/MINOS. For K2K, the far detector is the Super-Kamiokande detector; for NuMI/MINOS, it is the MINOS detector. A deficit of muon neutrino events and a distortion of the energy

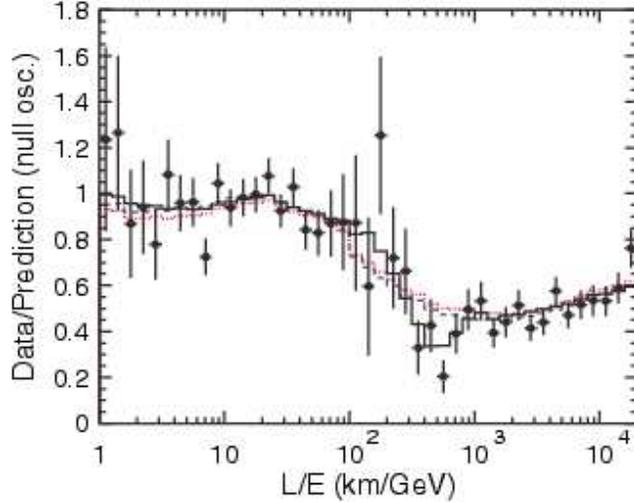


Figure 2.10: Super-Kamiokande’s $(L/E)_{Data}/(L/E)_{MC}$ spectrum, shown as points with errors. The solid line shows the best fit oscillation expectation; the red, dotted line shows the best-fit expectation assuming neutrino decoherence; and the blue, dashed line shows the best-fit expectation for neutrino decay [15].

spectrum from near detector to far is the hallmark of neutrino oscillations.

K2K amassed 1.049×10^{20} protons on target (POT) and found the results $\sin^2(2\theta_{23}) = 1.0$ and $\Delta m_{23}^2 = 2.8 \times 10^{-3} \text{ eV}^2$ [6]. For NuMI/MINOS, the most recent published results are based on 1.27×10^{20} POT [56], and as of Fall 2007, preliminary results on a 2.5×10^{20} POT data set have been released. Both are similar; the latter analysis yields $\sin^2(2\theta_{23}) > 0.84$ and $\Delta m_{23}^2 = 2.38_{-0.16}^{+0.20} \times 10^{-3} \text{ eV}^2$ at 90% confidence [1].

2.6 Outstanding Issues

Currently, there is very compelling, essentially universally accepted evidence that neutrinos change flavor and consequently have mass. But this “answer” leads to several very intriguing questions.

2.6.1 Number of Neutrino Species

Two famous measurements of the number of neutrino types unambiguously find three [34]. The Liquid Scintillator Neutrino Detector (LSND) data, on the other hand, demands at least four, the usual three (ν_e, ν_μ, ν_τ) and at least one sterile species which interacts only gravitationally. MiniBoone was designed to verify or disprove LSND and recently finished [19]. Their first results exclude LSND at 90% confidence [5].

2.6.2 θ_{13}

Ignoring LSND, all neutrino oscillation data are well described by a three neutrino basis and $\nu_e \rightarrow \nu_\mu$ (solar) or $\nu_\mu \rightarrow \nu_\tau$ (atmospheric) oscillations, giving $\theta_{12} \approx 34^\circ$ and $\theta_{23} \approx 45^\circ$. Thus far, the world has nothing but an upper limit on θ_{13} .

The Palo Verde [24] and CHOOZ [13] experiments were designed to measure θ_{13} using electron antineutrinos from nuclear power plants. The experiments were quite similar in most respects, results notwithstanding. They both found no evidence of oscillation and, for $\Delta m_{13}^2 > 2 \times 10^{-3}$, their data demand $\theta_{13} < 10^\circ$; see Figure 2.11. These experiments completely excluded the then allowed $\nu_\mu \rightarrow \nu_e$ region for Kamiokande, which strongly suggested that atmospheric neutrino oscillations were due to $\nu_\mu \rightarrow \nu_\tau$.

Unfortunately, no experiment has yet had the sensitivity to probe θ_{13} via, for example, $\nu_\mu \rightarrow \nu_e$ subdominant oscillations. The *NO ν A* experiment, currently in the final stages of approval, is designed to observe exactly this subdominant oscillation using a 30 kt detector with good spatial resolution located off axis in the NuMI beam. *T2K*, a Japanese experiment utilizing the Super-Kamiokande detector and an off-axis beam, is very similar to *NO ν A* and due to start operating in 2009.

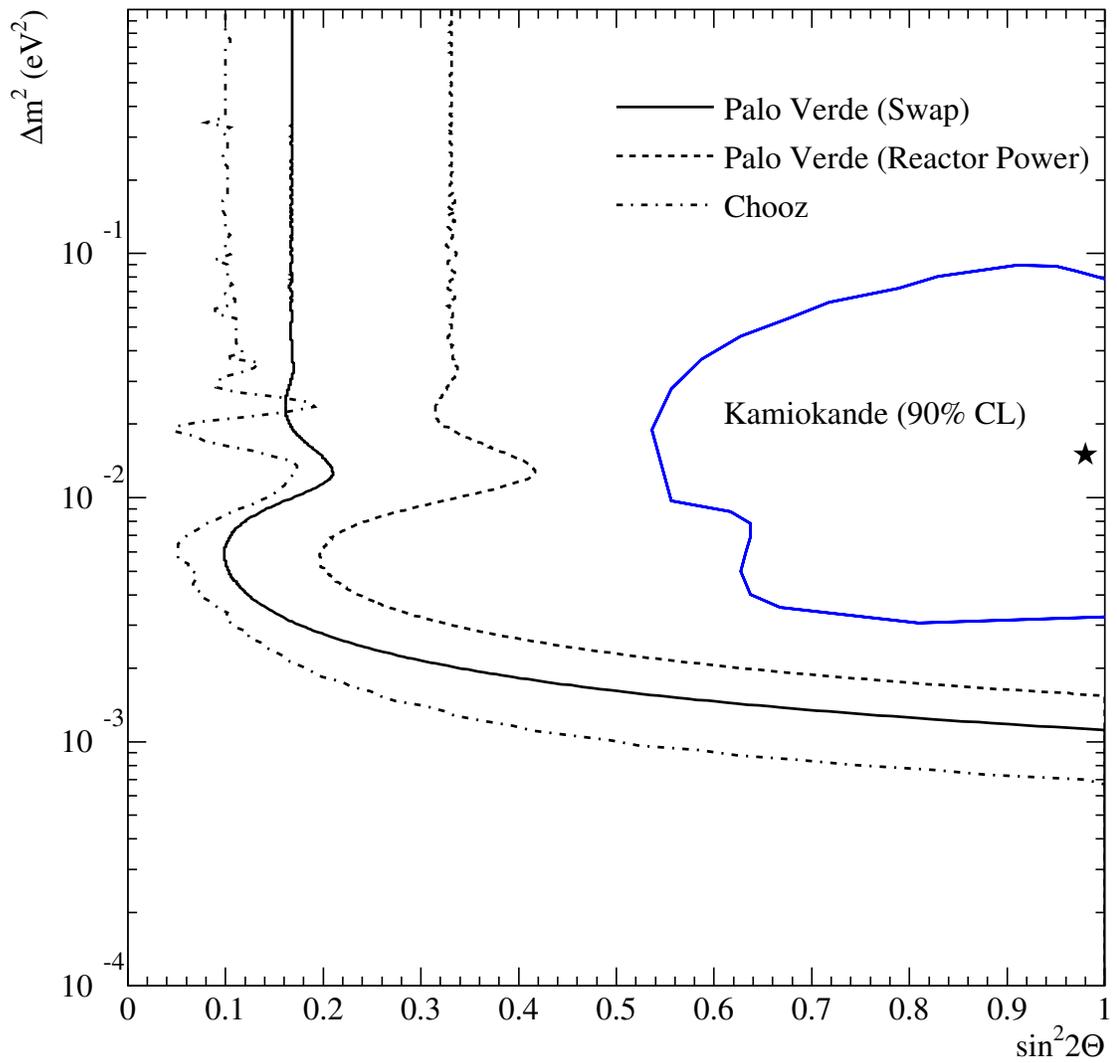


Figure 2.11: The 90% excluded regions for the CHOOZ and Palo Verde $\bar{\nu}_e$ reactor experiments. The CHOOZ line is the outermost. Also shown is the original Kamiokande $\nu_\mu \rightarrow \nu_e$ allowed region (blue) [24].

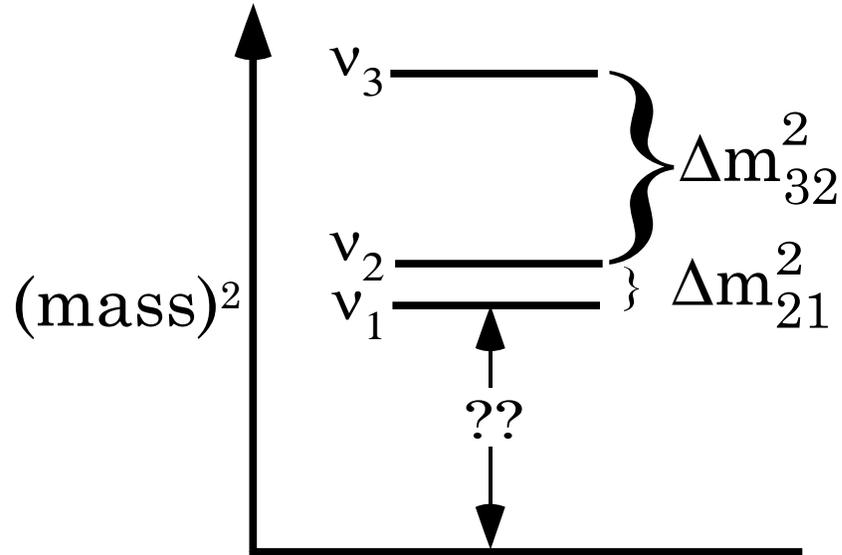


Figure 2.12: The normal mass hierarchy. The “??” indicates that the absolute position of the mass eigenstates is unknown [52].

2.6.3 Mass Hierarchy

The experiments which have endeavored to directly measure the neutrino masses have only provided loose upper bounds. The oscillation experiments are only sensitive to mass squared differences. As such, the absolute mass scale of the neutrinos is still unknown, as shown in Figure 2.12. The “??” shows that the absolute position of the 2 and 1 structure is unknown. Further, the arrangement of the mass states is unknown. If nature is consistent and neutrino mixing is analogous to quark mixing, then the two lightest neutrinos will be very close in mass and the third neutrino will be much heavier, known as the normal hierarchy. We have no evidence to exclude an inverted hierarchy, where there are two heavy neutrinos and one lighter neutrino.

CP violation is another issue intertwined with the mass hierarchy. For neutrino beam experiments, the beam must travel through the earth to the detector. In transit, the neutrinos feel an effective potential from interacting with the earth, and this effective

potential differs in sign for ν_e and $\bar{\nu}_e$. Thus, if an experiment were to use a $\nu_\mu/\bar{\nu}_\mu$ beam and combination of $\log(L/E)$ sensitive to the atmospheric mass splitting, that experiment could probe CP violation by comparing $\nu_\mu \rightarrow \nu_e$ and $\bar{\nu}_\mu \rightarrow \bar{\nu}_e$ [52].

Chapter 3

The MINOS Detector and Infrastructure

Pursuing neutrino measurements requires a very massive detector, or an extremely high neutrino flux, or some combination of the two. The atmospheric neutrino flux is not under human control, and so detectors must be made as large as feasibly possible. The MINOS detector is among the most massive underground detectors ever built, and is technologically distinct from the most common, water Čerenkov type detectors. MINOS has the potential to make an atmospheric neutrino oscillation measurement second in sensitivity only to Super-Kamiokande. Section 3.1 paints a general picture of the MINOS detector and detector components, Section 3.2 describes the identical layers (planes) that make up the MINOS detector, Section 3.3 explains the necessity, design, operation and construction of the veto shield, Section 3.4 walks through the MINOS data acquisition and electronics chain from energy deposition in the detector to the final digitized data, Section 3.5 enumerates the various measures the collaboration has taken to calibrate the detector, Section 3.6 describes how events are extracted from the data, and finally Section 3.7 discusses what kinematic quantities are of interest in a neutrino

interaction and how well they are reconstructed.

3.1 Detector Panorama

The MINOS (Main Injector Neutrino Oscillation Search) detector is a 5.4 kt (kiloton) magnetized iron calorimeter with tracking and charge separation capabilities [2]. It is located in the Soudan Underground Laboratory in Soudan, Minnesota, at a depth of 2,341 ft or 710 m.w.e. (meters of water equivalent).



Figure 3.1: A schematic of the Soudan Underground Laboratory. The Soudan2 cavern is on top right, the MINOS cavern is on the bottom left, and the connecting tunnel runs between. The elevator shaft is the vertical line. For a sense of scale, the area of the MINOS cavern floor is comparable to an American football field.

Geometrically, the MINOS detector has an 8 m octagonal cross-section and is 31 m long. The detector is divided into two supermodules. The first supermodule contains 248 instrumented planes; the second, 236.

There are essentially two components in the MINOS detector, steel and solid scintillator. Steel provides the mass necessary to have neutrino interactions in the detector, and scintillator is the active detector. More specifically, the detector is comprised of alternating layers of steel and scintillator planes having respective thicknesses of 2.54 cm

and 1 cm. Each of these layers is further subdivided into eight modules.

Blanketing the top and sides of the detector is a veto shield constructed of scintillator modules identical to the modules that make up the MINOS detector. The veto shield, while not critical for beam analyses, is absolutely essential for atmospheric neutrino studies.

As mentioned, the MINOS detector has charge separation capability. Each super-module is equipped with a 15 kA-turn coil to provide a toroidal magnetic field, which has an average value of 1.3 T (Tesla).

The coordinate system employed must be described, as it is referred to repeatedly later in this thesis. The far detector local coordinate system was chosen so that

- the y direction was vertical
- the z direction was along the beamline and $z = 0$ was at the point closest to Fermilab (so that all beam neutrino events increase in z)
- the x direction completed a standard right-handed coordinate system

Figure 3.2 illustrates these features.

To date, MINOS has been used to study the upward going muon flux and to make a measurement of the $\nu_\mu \rightarrow \nu_\tau$ oscillation parameters in atmospheric neutrinos [3] and beam neutrinos [56], described below.

The MINOS detector is one part of the long baseline experiment NuMI/MINOS. A neutrino beam generated at the Fermi National Accelerator Laboratory (FNAL) is aimed at the MINOS detector, and a near detector very similar to MINOS monitors the beam's content. The distance between FNAL and the MINOS detector is 735 km, and with that in mind, the beam and detectors have been designed to operate around the values of Δm_{23}^2 and E which maximize the oscillation probability given in Equation 2.36. Specifically, the beam generates neutrinos with energies peaked around 3 GeV

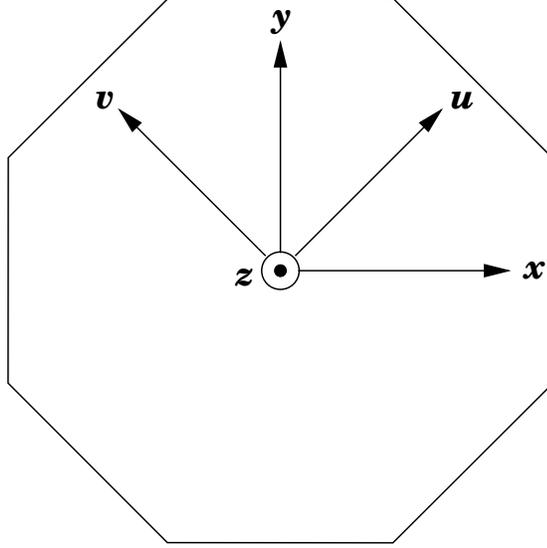


Figure 3.2: The MINOS local coordinate system as seen looking at the northern face of the detector. z is out of the page, and y is opposite gravity.

because of the current evidence that Δm_{23}^2 is somewhere between 0.001 and 0.010 eV². The oscillation parameters are measured by searching for distortions in the measured energy spectrum of the neutrino beam from the near detector to the far detector.

3.2 Plane Design

A MINOS plane is comprised of 192 strips of scintillator attached to an octagon of steel measuring 8 m across and 2.54 cm thick, as shown in Figure 3.3. The MINOS collaboration adopted this planar geometry for ease of assembly and maximal resolution for neutrino events coming from the direction of FNAL (the planes are oriented perpendicular to the line connecting Soudan and FNAL). Both of the materials that comprise a plane, steel and scintillator, arrive at Soudan in pieces which will fit down the 2 m by 2 m elevator shaft leading from the surface to the Laboratory. The steel is fabricated in sheets which are 1.27 cm thick and are either rectangular or trapezoidal, corresponding

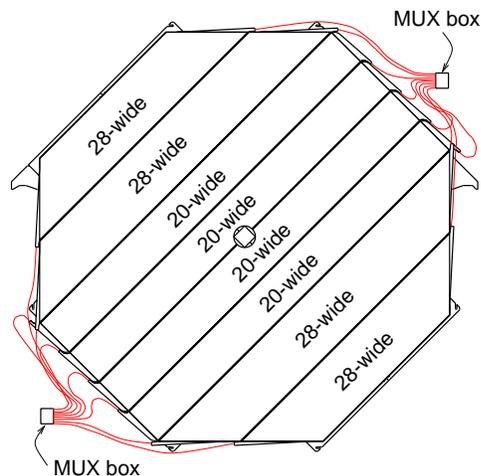


Figure 3.3: A view of one plane of the MINOS detector. The 20 or 28 wide refers to the number of strips in the individual module, and a MUX box is discussed in Section 3.4.1.

to the middle and outer portions of an octagon.

Four of these sheets are laid down to form an octagon on a support structure known as a “strongback,” a structure of steel beams used to support the otherwise flimsy steel planes during assembly and installation. On top of the first four sheets of steel, four more sheets are laid to form another octagon which is rotated 90° from the first. Thus each sheet in the lower octagon crosses all four sheets in the upper octagon and vice versa. Strategically placed plug weld holes in the sheets allow for assembling eight individual sheets into one solid, octagonal plane.

3.2.1 Scintillator

Scintillator is the active detector for MINOS, and the University of Minnesota played the primary role in the design, fabrication, and quality assurance of the MINOS scintillator modules. Caltech and Argonne also contributed to the scintillator module effort.

Scintillator is a material which emits scintillation photons when traversed by a

charged particle. Particular to MINOS, the scintillator is polystyrene doped with the fluors PPO, roughly 1%, and POPOP roughly 0.03% and co-extruded with titanium dioxide, which provides a reflective outer surface to trap emitted photons. Because scintillator is the active detector component of MINOS, its location and shape is what gives the position and spatial resolution of a hit within the detector. The final physical form of the scintillator, a strip 4.1 cm wide, 1 cm high, and cut to length to span the face of the plane, was chosen based on an optimization balancing cost and physics capability.

To collect the scintillation light, MINOS employs wavelength-shifting optical fibers. The cost of a phototube is approximately proportional to the area of its photocathode and, as such, an obvious way to minimize cost is to use the least photocathode possible. Optical fiber allows the light from a comparatively large scintillator strip to be focused on a tiny area of photocathode and thus save significant money. Another benefit of optical fiber is that its attenuation length is longer than scintillator's attenuation length.

Strips of scintillator arrived at the MINOS module factory with a length of 8 m and were cut to length so as to conform to the shape of the module they were included in. Workers would first epoxy the strips to the bottom aluminum skin.

There is a clear groove in the extrusion of each scintillator strip which facilitates collection of these photons by a wavelength-shifting fiber, glued into the groove during assembly. All of the fibers in the module run out of the light tight skin through a plastic manifold that guides the fibers into an optical connector. From here, the light is transmitted to a phototube, discussed later, via a clear optical fiber. The phototube converts the light into an electrical signal.

During assembly, the fibers are epoxied into an optical connector and the face of the connector and fibers is fly cut and polished to maximize optical transmission. Further, when in the detector, the modules are connected to phototubes via bundles of fiberoptic cable whose ends are mirror images of the optical connectors on the modules. Optical

grease was used at all of these connections to maximize light transmission.

The scintillator, much like the steel, must be modular to facilitate shipping and access to the Laboratory. Coincidentally, the scintillator comprising one plane of active detector is divided into eight “modules.” These modules take the shapes shown in Figure 3.4 and are comprised of either twenty or twenty-eight individual scintillator strips epoxied into an aluminum, light tight skin.

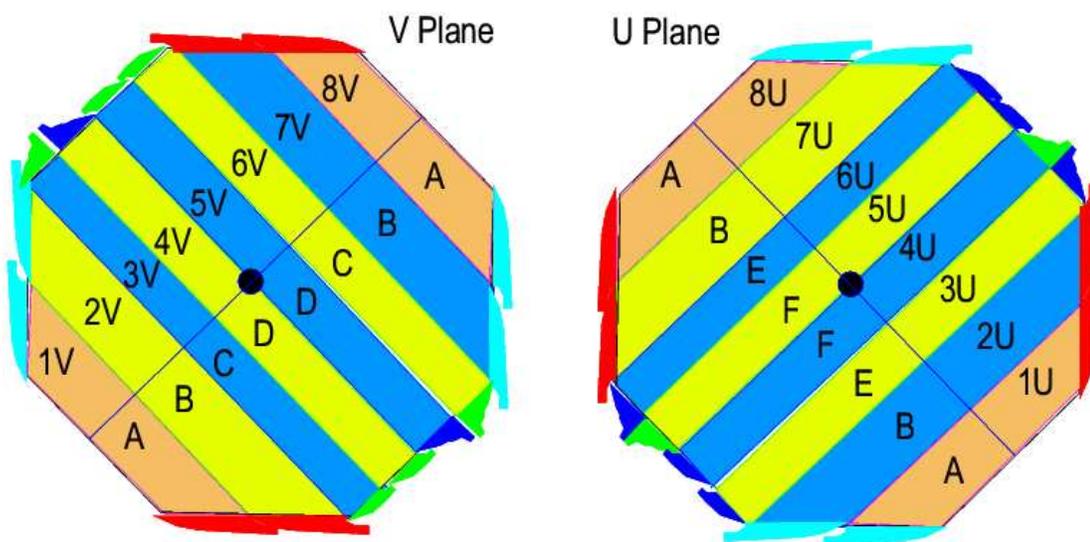


Figure 3.4: The various module types (*A* through *F*) for the two distinct types of planes (*U* or *V*) used in MINOS.

3.3 Shield Design

The MINOS veto shield was not included in the original plans for the MINOS hall and detector and, as such, has a rather unorthodox structure. Rather than having specialized components optimized for the spatial constraints of the cavern and the physics needs of MINOS, the veto shield is made almost entirely of spare detector materials. Being constrained to a tiny budget made designing the shield a very tricky problem, but the

logistical quandry came with an upside; readout of the shield could easily be handled by the existing electronics and data acquisition system and analysis of the shield data could be done in the same software framework as the detector data.

The primary motivation for a veto shield for MINOS came from early investigations of atmospheric and cosmic ray data collected before the entire detector was complete. The measured cosmic ray rate was consistent with expectations, but steep cosmic ray muons created a large unanticipated background for atmospheric neutrino interactions. Without a shield, MINOS could do no atmospheric neutrino analysis with a reasonable signal to background ratio. Thus the collaboration scrambled to piece together a shield. The result is a testament to physicist ingenuity. The shield contains four essentially identical sections, two for each supermodule. At 15 m long, a supermodule requires two scintillator modules laid end to end to span its length. As shown in Figure 3.5, the shield is merely scintillator modules laid lengthwise along the detector so that no muon coming from above can get into the detector without going through a shield module. All told, the shield has roughly 708 m^2 of active surface area, 3360 strips, and two layers over most of the detector.

The lateral spatial resolution of the shield is roughly that due to the width of a scintillator strip, $\frac{4.1}{\sqrt{12}}$ cm. The longitudinal resolution, using timing as a constraint, is $\sim 2 \text{ m}$.

The veto shield achieves approximately 96% efficiency at tagging cosmic ray muons.

3.4 Data Acquisition

The MINOS data acquisition system is an assembly of well-tested technology brought together to translate a pulse of light into a number stored in digital form.

The data acquisition begins at the mux box, an apparatus designed to house the photomultiplier tubes and achieve the optical summing, described below, employed to



Figure 3.5: Veto shield construction. The detector lies under the modules to the right. The vertical modules to the left tag muons coming in at an angle.

minimize readout cost. The internal configuration of the mux box is shown in Figure 3.6.

3.4.1 Optical Summing

As mentioned, the scintillator light output of each strip is collected by an optical fiber and subsequently read out by a phototube. Phototubes are particularly expensive components and, because of the multitude of strips in the MINOS detector, the collaboration decided to employ multiplexed optical readout (technically optical summing) to save costs.

Each electronics channel reads out eight scintillator strips, and the strips which share a readout channel are chosen so that, while all in the same plane, are separated by 1 m along the plane. Also, because both sides of a scintillator strip are read out, the optical summing scheme was designed such that if a hit strip produces a signal on both sides of its plane, and there is only one hit strip in the plane, that combination of electronics channels (one east, one west) is unique and maps back to only that strip. See [68] for a detailed description of the scheme. The mux box, described below, is the component which implements the optical summing.

Traveling with a photon from inside the detector, one first encounters the connectors shown near the top of Figure 3.6. Each MINOS plane has eight modules and a corresponding optical fiber bundle connecting it to the mux box, and each mux box has sixteen connectors. Thus each mux box reads out one side of two planes. From any one connector, there are either twenty or twenty-eight fibers to be read out, corresponding to the two broad classes of scintillator modules. All told, there are 192 fibers coming from each side of each plane and, as such, 384 fibers coming into one mux box. Each phototube, shown in the middle of Figure 3.6 and described below, has sixteen pixels. The three phototubes together have 48 pixels and, consequently, each pixel reads out

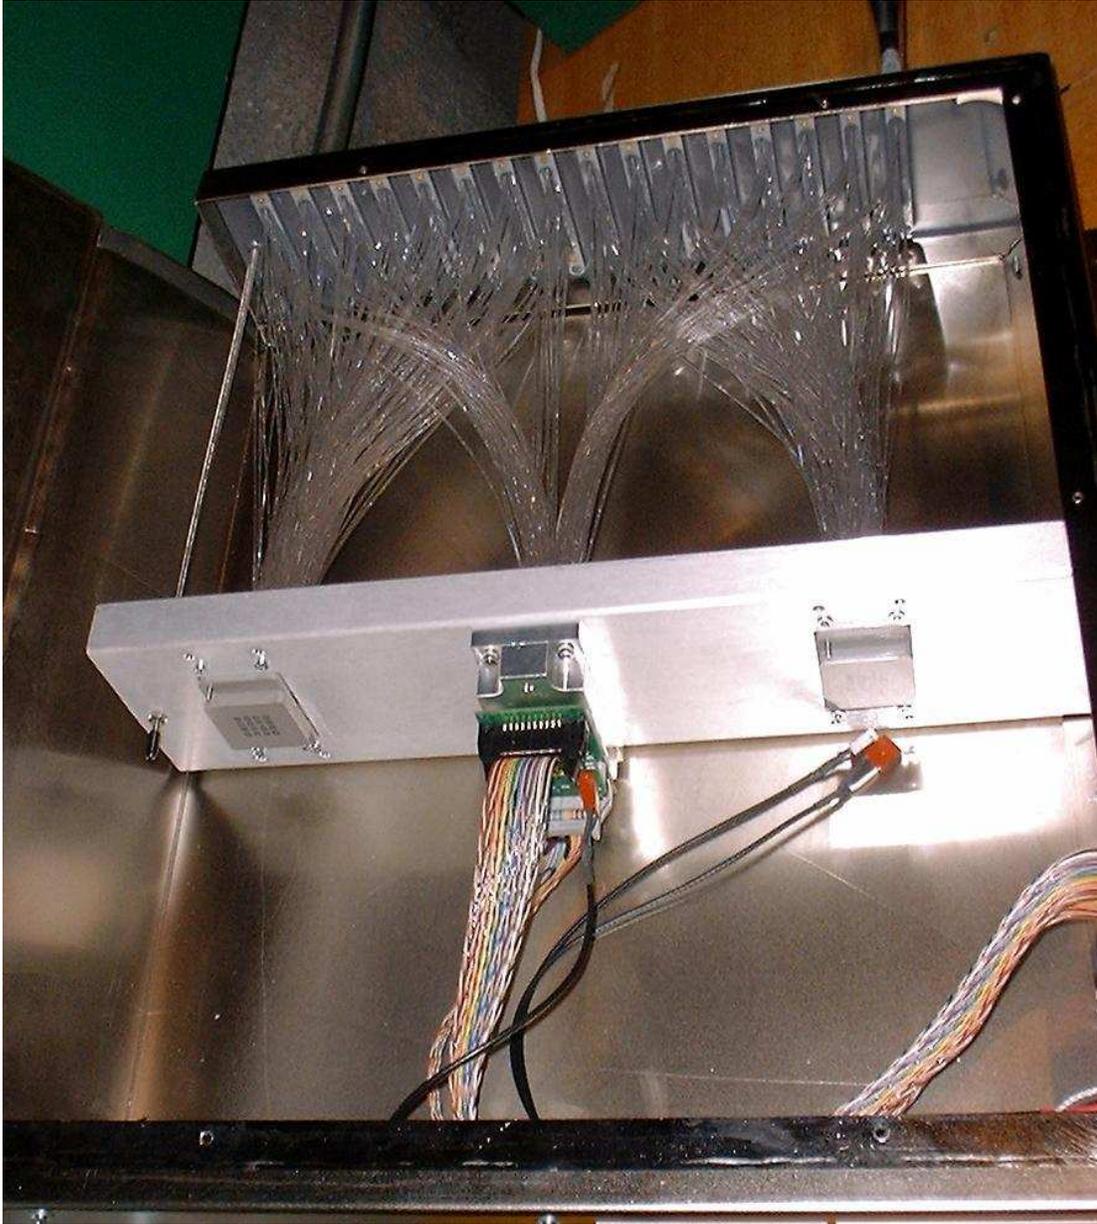


Figure 3.6: Inside a mux box. The fibers from the detector attach at the top, and the clear fibers run into the “cookie” (described in the next section), and the phototube mounts to the other side of the cookie. In the picture, the colorful ribbon cable attaches to the phototube base.

eight fibers. The fibers are held in place against the face of the phototube by an implement known as the “cookie” in MINOS jargon, shown in Figure 3.7. The face of the cookie is visible in the outer two phototube spots in Figure 3.6, and each small spot has eight fibers coming into it from the connectors.

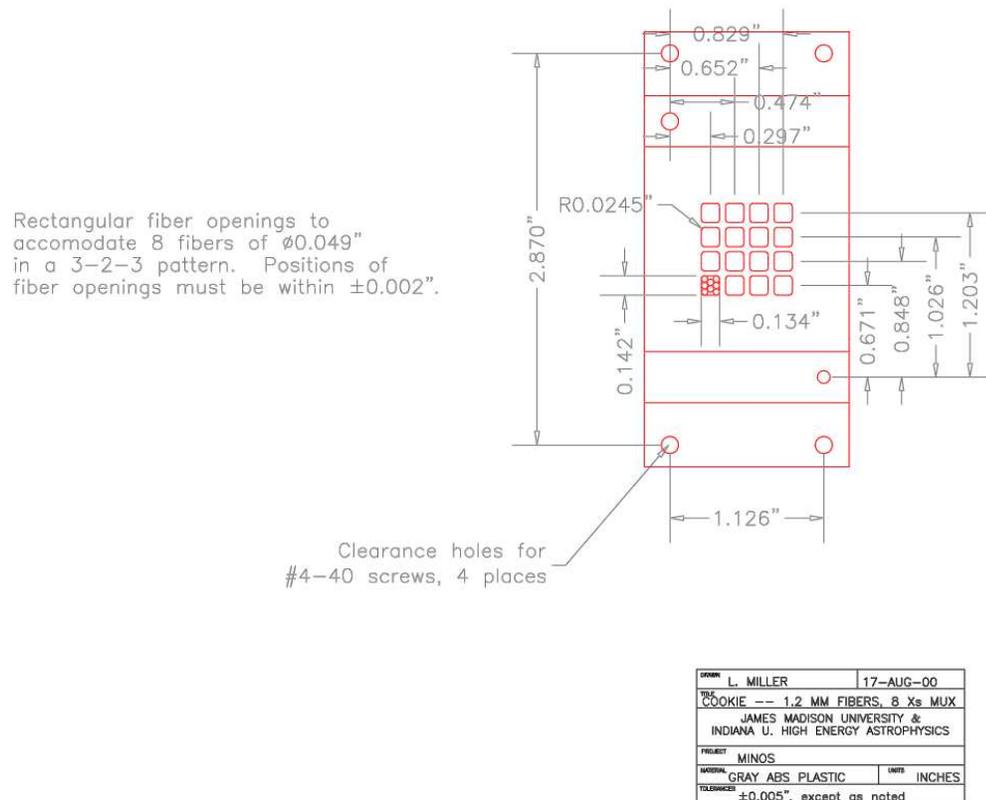


Figure 3.7: Diagram of a cookie, which holds multiple optical fibers against one PMT channel inside the mux box. Note that the cookie has sixteen holes, one for each of the PMT channels. Each hole in the cookie holds 8 fibers.

Next in the photon’s path is the phototube itself, a Hamamatsu R5900-00-M16 16-anode photomultiplier. A well-established phototube, the M16 has been used by several other experiments, including CDF and DESY’s HERA-B detector [2], and represented the optimal photodetector choice for MINOS at the time MINOS was being designed.

The M16 has sixteen separate photocathodes and corresponding anodes. The phototube is sensitive to light with wavelengths between 300 and 650 nm, and the peak sensitivity wavelength is 420 nm, at which the quantum efficiency of the photocathode is roughly 15%. The emission spectrum of the MINOS scintillator and WLS is around 520 nm, and the M16 has at this wavelength a quantum efficiency of roughly 13%. The sixteen channels share the twelve stage dynode chain, which employs electrostatic focusing to keep the electrons within one channel of the tube. This dynode chain provides a maximum gain of 3.3×10^6 . In practice, the phototubes in MINOS are run at a voltage which provides an approximate gain of 1×10^6 . After traversing the dynode section, the electrons are collected on the appropriate anode and read out as an electrical current.

In generic terms, the job of a phototube is to convert light into an electrical signal; this small pulse of charge is what constitutes the next stage of MINOS data acquisition. MINOS employs a slew of electronics to digitize and catalog the phototube signal; VA chips (made by the IDE Corporation) handle the initial analog signal processing, and they are read out by VA Readout Controllers (VARCs) in VersaModule Eurocard (VME) crates. On the VARCs are VARC Mezzanine Modules (VMMs), which handle all the digitization tasks and act as the analog-digital interface between the VA chips and the VARCs.

3.4.2 VA Chip

The MINOS VA32_HDR11 chip, produced by IDEAS of Norway, is a customized version of a standard product they produce, the VA32_HDR2. Both models are 32 channel analog CMOS chips with digital control logic. Each channel has a preamp, a shaper, a track and hold stage, and an output switch. The standard VA chips were intended for use with silicon detectors, and as such low noise was a primary design goal. MINOS had, at one time, considered using hybrid photodiodes (HPDs) as its photodetectors,

and the standard VA chip would have been almost ideal off the shelf. Phototubes are fairly quiet compared to silicon detectors, and they also have a much larger dynamic range, and this is where the MINOS VA chips diverge from the VA32_HDR2. The main modification to this chip was a modified preamp capable of handling up to a 30 pC signal with solid linearity.

3.4.3 VA Frontend Board

As mentioned, each mux box contains three phototubes, each of which requires a VA chip for readout. Rather than have totally separate power distribution and support circuitry for each chip, a VA Frontend Board (VFB) hosts the three chips and attaches directly to the mux box. The VA support circuitry on the VFB includes a multi-channel Amp-Shaper-Discriminator (ASD) chip and a serial slow control interface for biasing, regulating, and monitoring. The VFB is, strictly speaking, a slave device; all action on the VFB is actually controlled by the VARC to which it is attached. When a dynode trigger is received by the ASD, a signal is sent to the VARC which timestamps the trigger and, if it has coincident triggers from other VFBs, the VARC initiates readout of the appropriate VA chip. This process is described further in the VARC section. Two cables connect each VFB to a VARC, a fifty conductor twisted pair control cable and a twisted pair analog signal cable.

3.4.4 VMM

The VARC Mezzanine Module (VMM) is where the digitization takes place in the MINOS readout chain and is what the analog signal cable from a VFB connects to. The VMM resides on the VARC on an elevated platform (hence the name Mezzanine), and a VARC can accommodate six VMMs.

The heart of a VMM is a high speed, 14 bit ADC. The maximum digitization,

approximately 10 MHz, is so high compared to the MINOS event rate, roughly 4 MHz, that one VMM handles the digitization of two VFBS with no increase in dead time.

3.4.5 VARC

MINOS is a relatively simple detector, judged on an absolute scale. Still, MINOS has upwards of 180,000 strip ends, and the readout and handling of such an ensemble of detector elements is a formidable task. Elements must be ganged together and managed in an automated way, noise must be removed from the data stream, and hits must be digitized and timestamped. The electronics components described thus far have primarily been slave devices with one function. The VA Readout Controller (VARC) is the first element in the data acquisition chain which exercises control over other components. The VARC has three primary functions: timestamp signals, sparsify signals, and interface with the next level of signal processing, the Readout Processor (ROP), via VME.

A FPGA chip known as the Event Timestamp Controller (ETC) is the gatekeeper of the readout process. One ETC services a VMM, and there are six of these pairs in total on each VARC. A dynode trigger from a VFB initiates the readout process. The ETC receives this trigger, timestamps it, and queries the VA chip the trigger originated from to find out whether or not the chip is waiting to be read out. If not, the ETC puts the VA chip into “sample and hold” mode, in which the VA chip waits to acquire the anode signal from the phototube. Once this signal is acquired, the ETC causes the VA chip to dump the signal to the VMM, where it is digitized. This digitized signal is then sent to a FIFO to await sparsification, to be described in the next paragraph. It should be noted that if a VA chip is waiting to be read out, any dynode triggers received from that chip are ignored.

The VMM returns, for an input charge, some integer ADC value; because of electronics noise, each channel will have some ADC output even for zero initial charge. This noise is referred to as the pedestal, and all channels have their respective pedestal value subtracted from their initial ADC value whenever they are read out. Not only are the channels pedestal subtracted, they also are subject to a minimum criteria known as the sparsification threshold. If, after pedestal subtraction, a channel's ADC value is less than its sparsification threshold, that channel's signal is ignored. This sparsification threshold is set to $\frac{1}{5}$ pe. These operations are carried out by the Sparsifier. Each VARC has only one Sparsifier, which reads from all six FIFOs.

Should a signal survive sparsification, the Sparsifier then puts it into one of two 32 kb RAMs which are read out cyclically via VME (see Table 3.1). Finally, readout of the RAM is initiated if and only if the VARC detects signals from two of its 36 VA chips above sparsification threshold which lie within a coincidence window of 500 ns.

Bits:	0-13	14-15	16-20	21-22	23	24-26	27-28	29	30	31
Use:	ADC Data	HD0,1	Mux [4:0]	VA0,1	SEL	F0-F2	VID0,1	DC0	P	1
Bits:	0-15			16-29			30	31		
Use:	TS [15:0] Time Stamp Lower			TS [16:29] Time Stamp Upper			P	0		

31st bit - Data Identifier: 0=timestamp, 1= ADC address

P - Parity (including ID)

HD[1:0] - Header code during Time-of-Flight(TOF) event

F[2:0] - 110 signifies a TOF event

SEL - unused

VA[1:0] - unused

MUX[4:0] - unused

ADC[13:0] - unused

VID[1:0] - Identify VARC in crate

TS[29:0] - Time-stamp

DC0 - data code (unused)

Table 3.1: The upper and lower 32 bit data words for a digit [36].

3.4.6 ROPs, BRPs, TPs, Oh My

Beyond the VARC, the data acquisition and processing is handled by a series of computers in a hierarchical configuration. Within each VME crate is a Read Out Processor (ROP) which, through VME, pulls the data out of the two RAM buffers on each of the three VARCs in the VME crate. The ROP pieces together the data over a *timeframe*, currently one second in length, and buffers timeframes until they are read out by the next stage of the DAQ.

Both supermodules of the MINOS detector have four VME crates per side, and each of these groups of four crates form a *branch*; the four ROPs in a given branch are read out by one Branch Read Out Processor (BRP). The BRPs cobble together the timeframes from their ROPs, and then pass the data on to the trigger farm. This readout and transfer is controlled by a Master Branch Read Out Processor (MBRP), which coordinates the four BRPs and selects which Trigger Processor (TP) in the trigger farm receives the timeframe under consideration from the BRPs.

The TP goes beyond simple data shuffling; it orders all the data temporally within a VME crate and implements the four out of five plane trigger, which requires that four planes out of any five contiguous planes in the detector contain at least one hit each. If this trigger is satisfied, the TP forms a *snarl*, which is the final form of the MINOS raw data.

The last stage of the DAQ is the Data Collection Process (DCP), a final ordering and quality assurance stage. While the data coming from a specific TP are time ordered, the data coming from the various TPs in the trigger farm are not, and so one important task of the DCP is to carry out the final time ordering. The DCP buffers the timeframes it receives from the trigger farm and orders them temporally. Afterwards, the DCP tests checksum words in the data stream as a sentinel against corrupted data and does some timeframe overlap correction. Lastly, it writes the data to file [36].

3.5 Calibration

A detector by itself is not of much use. One must know how to interpret its response, and this is the essence of calibration. MINOS has two primary readout quantities which must be calibrated, charge and time.

3.5.1 Charge Calibration

The fundamental goal of charge calibration is to provide the constant of proportionality relating ADC count and energy deposited in a particular scintillator strip by a charged particle. To this end, numerous factors must be calibrated out. First, unavoidable channel to channel variations in the Front End Electronics (FEE) cause different channels to produce different ADC counts for an identical amount of charge collected. To remedy this, the FEE are equipped with a charge injection system. The charge injection system inserts a known amount of charge in place of an anode signal, and the FEE read out this charge as usual. The initial charge is varied to cover the dynamic range of the ADC, and thus any nonlinearity in a channel's response can be calibrated out.

Second, phototubes also have a response which is not only nonlinear with incident light, but varies according to ambient conditions. Thus MINOS requires a system to monitor these time variations, as well as linearize the response of the phototubes. The light injection system accomplishes these goals, along with evaluating the PMT gain and confirming that the optical portion of the readout chain is intact.

The heart of the light injection system is the pulser box, which contains blue LEDs, optical connectors on the rear of the box to interface with fiberoptic cable running to the scintillator modules, and clear fiber to route the LED light from the LED to the connectors. The basic idea is simple: direct light from an LED onto a WLS fiber coming from a strip of scintillator and read out the light which enters the fiber and propagates to the phototube. The details, of course, are more complex. The fiberoptic cables running

between the pulser box and the modules plug into, on the module side, a highly reflective chamber in the manifold which holds all the WLS fibers coming out of the module. One fiber from the pulser box illuminates numerous fibers in the module. The optical connections inside the pulser box are arranged to ensure that no two WLS fibers which are read out by the same PMT pixel are illuminated simultaneously. The LED pulse length and amplitude are controlled automatically and, in a Light Injection (LI) run, the amplitude is varied to map out the gain curve of the phototube. During this LI run, the light output of the LEDs must be independently monitored to account for any variations in the LED light output, which can change by a few percent. Thus the light from every pulser box is also routed to two Positive-Intrinsic-Negative (PIN) Photodiodes, whose response is both linear and stable within the requirements of MINOS. These components, in concert, allow MINOS to correct for seasonal changes, intrinsic PMT nonlinearity, and hardware swaps.

3.5.2 Timing

The heart of the MINOS timing system is a GPS unit located in the Laboratory and connected to the surface via an optical antenna. The rather long antenna induces a $5.1 \mu\text{s}$ offset from UTC but doesn't affect the relative accuracy of the GPS. The timing hardware within the FEE is connected to the GPS via the Timing Central Unit (TCU). This TCU receives one pulse per second from the GPS, as well as a 10 MHz clock signal. Using these two signals, the TCU assembles timing signals which it sends to the FEE. Here, the 10 MHz signal is multiplied up to a 640 MHz signal to achieve the 1.5 ns timestamping accuracy required for dynode triggers.

3.5.3 CalDet

Ultimately, the NuMI/MINOS experiment will culminate in a comparison of the neutrino energy spectrums in the near and far detectors, and consequently an accurate absolute and relative energy calibration is imperative. CalDet addresses the former.

Absolute calibration of a calorimeter can only be done by exposing the calorimeter to a particle source of known type and energy. This approach is possible for neither the near nor far detector. Thus was born the Calibration Detector (CalDet). CalDet was a miniature version of the two MINOS detectors placed in various particle beams at CERN. The beams were of known content and energy and were employed to measure the detector response. See [58] for a detailed description.

3.5.4 MINOS Module Mapper

Scintillator module fabrication took place in three separate places, and each factory was equipped with a module mapper. A mapper was a device used to immediately test scintillator modules for fabrication errors and to provide a metric with which the collaboration could detect individual module damage during shipping as well as perform calibration tasks. A mapper consisted of a long table to support the scintillator modules, a radioactive source (Cs137), a lead housing called “the pig” which collimated the photons emitted by the source, a motorized cage which controlled the position of the pig along the table, two mux boxes for module readout, and an automated data acquisition and motor control system.

Starting at one end of the module, the pig was moved across the module perpendicular to the direction of the strips. The pig was then moved down the module a few centimeters, and the pig was again moved across the module perpendicular to the direction of the strips. This process was repeated until the pig had traversed the entire length of the module. In this way, the response of the module to a constant source along

a fine grid of points on the module was mapped out.

Knowledge of the response as a function of the position along the module allowed for calculation of the attenuation length of the optical fibers within the module. Also, the collimation was pyramidal, and knowledge of the shape allowed for calculation of the strip positions in the module. See Figure 3.8 for a typical map of some scintillator strips.

With the mapper information, and full track reconstruction (described in Section 3.6.6) to give position along the hit strip, correction for attenuation is possible.

3.5.5 Strip-to-Strip Calibration

Three main factors exist which light injection and mapping calibration cannot compensate for. Scintillator light yield, defects in optical fiber, and variations in the optical connectors all contribute to uncertainty in the detector response. A strip-to-strip calibration of the detector using cosmic ray muons ameliorates this problem at the 2% level [59].

3.5.6 MEU Calibration

With the previous calibrations complete, a digitized pulse height has some meaning relative to the other channels and strips in the detector, but has no meaning in terms of real energy. Thus one final calibration must be done, to translate this almost-fully-calibrated pulse height into a specific amount of energy deposition, known as the Muon Energy Unit, or *MEU*. The *MEU* is, ignoring small strip-to-strip variation, the detector response to a 1 GeV muon traversing 1 plane of scintillator at normal incidence [46].

The calculation of the *MEU* is done by exploiting knowledge of how particles lose energy when they traverse various materials, scintillator in this case. Figure 3.9 shows this energy loss, dE/dx , as a function of the particle momentum. First, note that the

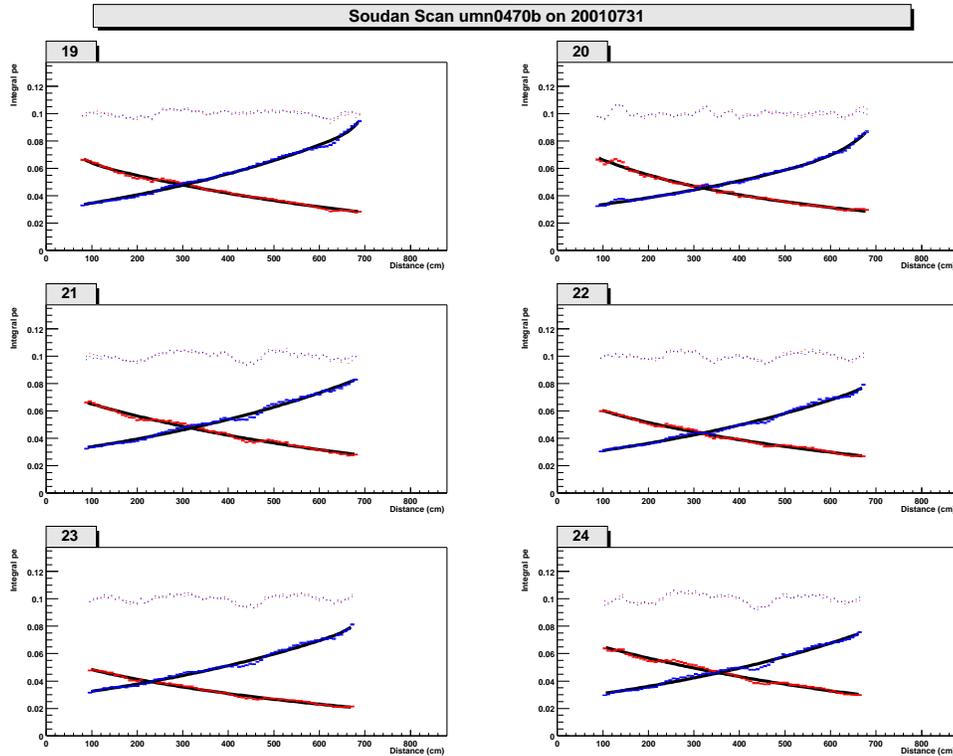


Figure 3.8: A typical “map” for six strips in module 470b. The red (decreasing from left to right) and blue (increasing from left to right) points show the pe response from the two ends of the strip when the source of the mapper was at the z -position along the strip shown on the horizontal axis. The black lines are double exponential fits to said data, and the “floating” points show the ratio of the real response to the fit. It should be noted that the fit follows the data points quite well for this particular module, so in many places the fit and the data points are on top of one another.

vertical scale is in stopping power, which can be converted into energy loss (MeV/cm) by simply multiplying by the density of copper in g/cm^3 . Second, note that in the region around minimum ionization, the energy loss is only weakly dependent on the actual particle momentum. This feature is particularly useful for one type of event, stopping cosmic ray muons. For stopping muons, it is a given that, very near the track end, these muons have fallen below the minimum ionizing momentum. Further, knowing the detector composition and the location of the track end allows for easy selection of a specific muon momentum by simply projecting back the energy loss along the track. In this way, the *MEU* calibration algorithm selects, for any stopping muon track, a region of the track which has a momentum of about 1 GeV, and, consequently, an accurately known energy loss. A small window around this point is selected, and the average detector response per path length, shown below, is calculated in this window.

$$\left\langle \frac{dE}{dx} \right\rangle = \frac{1}{N_p} \sum_{i=1}^{N_p} \frac{S_i}{L_i} \quad (3.1)$$

Here, N_p is the number of planes in the track window, S_i is the total detector response in the i^{th} plane of the track window, and L_i is the path length through the i^{th} plane (the plane thickness divided by cosine of the angle of incidence). A large ensemble of these averages is acquired, and the median of that ensemble is taken to be one *MEU* [46].

3.6 Event Reconstruction

As mentioned, the essence of a bit of MINOS data, known as a *digit*, is an ADC count, a time, and an electronics address. Deciding what sort of particle or particles produced the *digits* in an event is the goal of event reconstruction. In MINOS, the reconstruction happens in sequential stages, each of which will now be succinctly described.

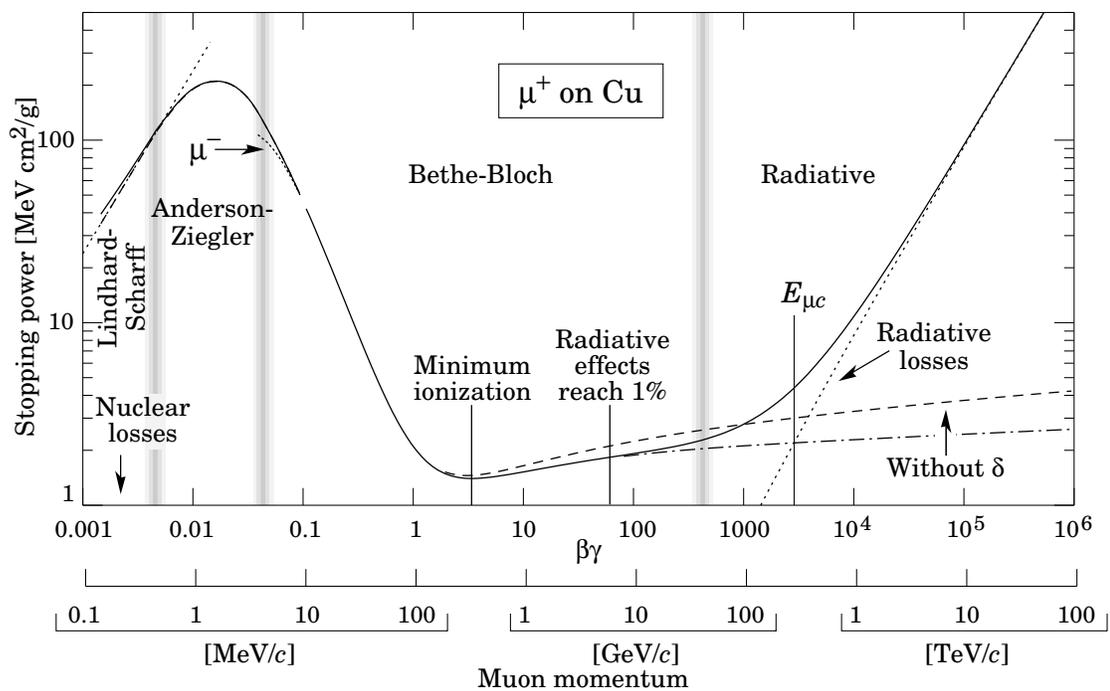


Figure 3.9: The energy loss curve for positive muons incident on copper across a wide range of muon momentum [34].

3.6.1 Noise Filtering

Despite the implementation of an online plane trigger condition, as described in Section 3.4.6, most of the *snarls* (see Section 3.4.6) written to disk are junk due to random noise coincidence. As such, the first stage in the reconstruction chain is a simple noise filter, which only passes *snarls* which satisfy very loose pulse height and event length requirements.

3.6.2 Demultiplexing

Section 3.4.1 described the optical summing, also referred to as multiplexing, the collaboration employed as a cost saving measure. This savings incurred cost in another place. With each electronics channel (alternately stated as each PMT pixel) now reading out eight possible scintillator strips, the reconstruction software acquired the additional task of disambiguating which strip produced the light read out by the PMT.

The demultiplexer the collaboration currently uses is described in [74]. The demultiplexer used to reconstruct the data analyzed in this thesis is described in [69]. As detailed previously, the optical summing scheme was designed so that, in an ideal world free of cross talk and random noise, a hit strip will emit light from both ends and produce a unique pair of *digits* (assuming that said hit strip is the *only* hit strip in the plane). Both demultiplexers begin by searching for these unambiguous *digit* pairs. Unfortunately, the demultiplexers do have to filter out cross talk and stray hits, and both employ all available information (pulse height, timing, strip locations, etc.) to untangle the myriad of strip combinations.

Both demultiplexers, on average, correctly assign over 97% of the total event pulse height to the correct strips, as discussed in [74].

3.6.3 Slicing

In the near detector, the neutrino flux is so high during beam spills that there are, quite often, several neutrino interactions resulting from the same spill. In such a high rate environment, *digits* from the same physical event must be grouped together into *slices* before higher level reconstruction is executed. This grouping prevents, for example, the track finder from creating one track from two distinct tracks born from two different physical events. The slicer performs the grouping based on timing and spatial information.

The far detector has such a low event rate, around 0.5 Hz, that the slicer takes no action when reconstructing far detector events.

3.6.4 Clustering

The next step in the reconstruction is clustering. A *cluster* on its own isn't of much physical significance; it is merely a group of hit strips which are close to one another in space and time. Such proximity suggests that these strips are part of the same track or shower, and so clustering aids in the later reconstruction. It is worth noting that a single strip which has no neighbors in its plane can be its own *cluster*, and such a cluster is designated as *tracklike*. Any other *cluster* is designated as *showerlike*.

Realistically, the granularity of the MINOS detector allows for only two gross event topologies, showers and tracks. Shower and track finding are the next two stages in the reconstruction chain.

3.6.5 Shower Finding

The shower finder searches through the *showerlike* clusters, looking for pairs of showers in opposite views (i.e. one u and one v cluster or vice versa) which overlap in z and time, and are consistent in pulse height. If such a pair is found, they are deemed a 3D

shower.

3.6.6 Tracking

If one subscribes to the philosophy that physics is fundamentally an empirical science, then almost all of high energy is, in essence, tracks. From bubble chambers and spark chambers to the incredibly sophisticated detectors of today, physicists have “tracked down” the myriad of particles we now know by observing their paths of energy deposition.

The MINOS track reconstruction starts by considering the demultiplexed strips in the two independent views, u and v . In each view, the tracker uses a combination of timing, linear fitting, and the Hough transform [33], to assemble the strips into a $2D$ track and to find the track’s vertex and end. After the tracker completes the u vs z and v vs z tracks, it combines them into one $3D$ track and calculates the track’s range, momentum from range, and other parameters of interest.

3.6.7 Track Fitting

The MINOS detector is equipped with a magnetic field, and so it can measure momentum from curvature in the field as well as identify the charge of the particle producing the track. The magnetic field throughout the detector is parametrized in a “map” available to the reconstruction software.

To measure the curvature of a track, the track fitter employs the Kalman Filter algorithm [50] to iteratively fit the track in the forward and reverse directions (i.e. vertex to end and end to vertex) until the measured momentum converges. The fitter allows for the addition and removal of strips to and from the original track if such action improves the fit.

3.6.8 Event Assembly

The final stage of event reconstruction is event assembly. This last stage consists mostly of bookkeeping, as the event finder collects associated tracks and showers and places them into an *event*. The event finder also performs some secondary reconstruction tasks, such as tagging the primary shower (should multiple showers be reconstructed), adjusting downward any shower energy pierced by a track, and investigating stray *strips* and *clusters* for possible *tracks* or *showers* the initial reconstruction missed.

3.7 Neutrino Variables and Resolutions

The immediate question for any reconstruction software is how well it extracts from the data the true kinematic characteristics of the particle responsible for the data. This question is impossible to answer with real data, where the true nature of the event is unknown. But with Monte Carlo, reconstructed quantities may be compared with the known true quantities to estimate the accuracy of the reconstruction. The MINOS Monte Carlo is discussed in the next chapter; suffice it to say for now that a thorough Monte Carlo exists, with a large ensemble of Monte Carlo atmospheric neutrino events.

3.7.1 The Neutrino Variables L and E

In anticipation of later chapters, the two single most important bits of information to be extracted by the reconstruction, L and E , are discussed here. First, definitions must be laid out. L is the path length the neutrino traverses, and it enters the expression for the oscillation probability (Equation 2.36). Figure 3.10 shows a schematic of the Earth and the neutrino flight path. The angle θ in the figure is the zenith angle, and using the law of cosines, one can express L :

$$L = \sqrt{R^2 \cos^2 \theta + 2Rh + h^2} - R \cos \theta \quad (3.2)$$

where R is the Earth's radius and h is the neutrino production height in the atmosphere (assumed to be 20 km for this analysis). The detector is assumed to be on the surface of the Earth because the actual depth, 737 m, is completely negligible for the calculation of L .

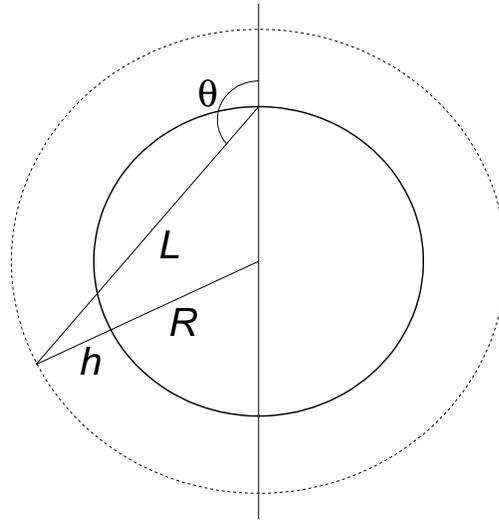


Figure 3.10: The neutrino path length L , atmospheric production height h , zenith angle θ , and Earth's radius R . The atmosphere is shown as a dashed line. (Drawing not to scale)

There remains one ambiguity in this calculation of L , namely where the zenith angle actually comes from. With a perfect detector, its lucky operator could observe all products in a neutrino interaction and unambiguously assign the incoming neutrino direction using simple kinematics. As MINOS is not a perfect detector, a choice must be made. It has been decided that the outgoing lepton track provides the best indication of the incoming neutrino direction, based on a preliminary study that was carried out to investigate the utility of including shower kinematic information while determining the neutrino direction. More explicitly, an effective shower mass was used to calculate how much momentum the shower carried off based on the total shower energy. Including this shower momentum actually makes reconstruction of the neutrino zenith angle slightly

less accurate, as Figure 3.11 shows. The decrease in accuracy results from the poor directional resolution on the hadronic shower.

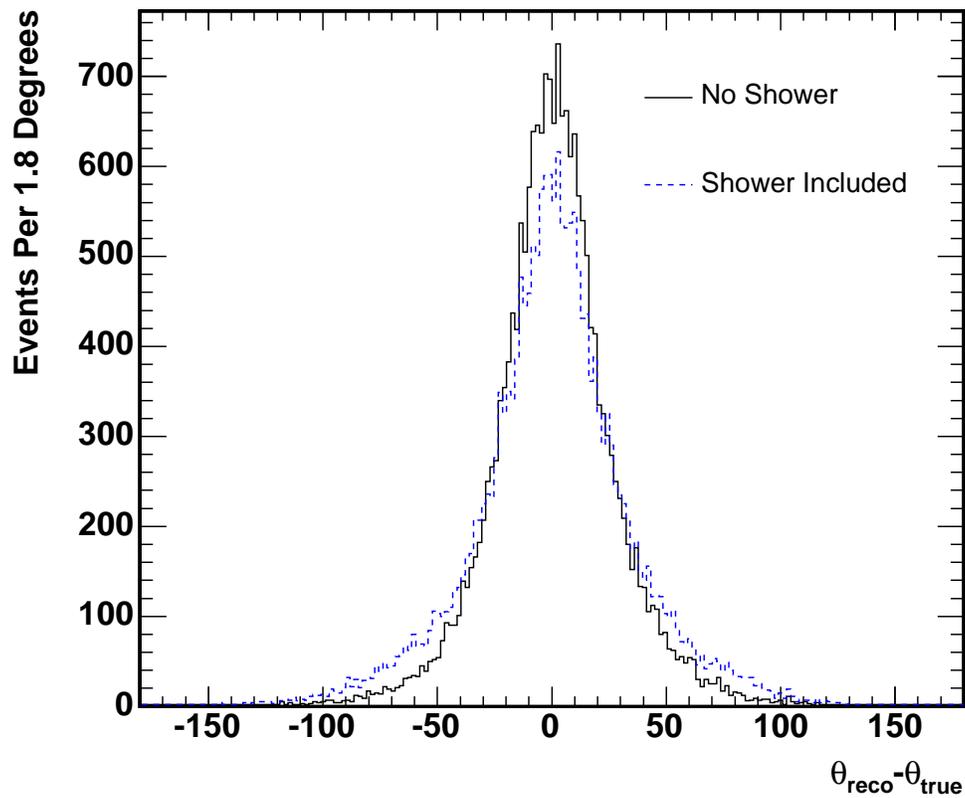


Figure 3.11: The neutrino zenith angle resolution from Monte Carlo. Black (solid) shows the resolution when using the lepton direction as the neutrino direction, and blue (dashed) shows the resolution when shower information is included in calculating the neutrino direction.

E is the total neutrino energy, which also enters the oscillation probability (Equation 2.36). In analogy with L , the utility of shower information for the calculation of E has been investigated. Here, including the shower energy significantly improves the accuracy of the neutrino energy estimation, as shown in Figure 3.12.

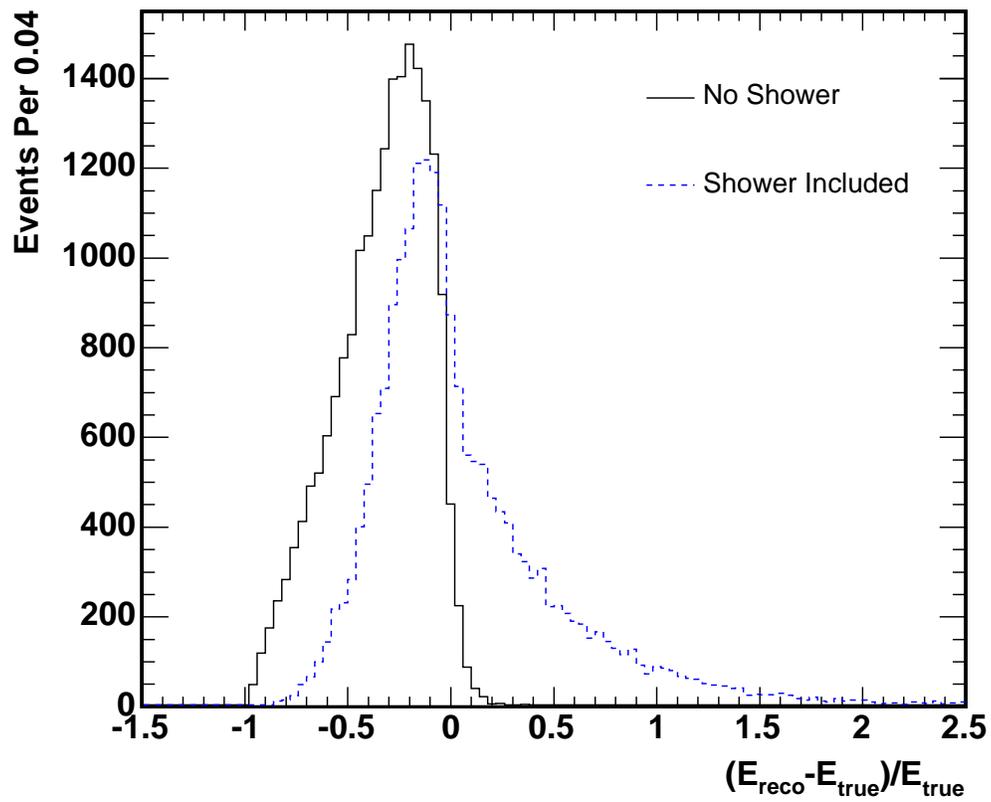


Figure 3.12: The neutrino energy resolution from Monte Carlo. Black (solid) shows the resolution when ignoring any shower energy, and blue (dashed) shows the resolution when shower information is included in calculating the neutrino energy.

3.7.2 Resolutions

Now that the methods of calculating these quantities are established, a quantitative comparison of their true and reconstructed values is in order.

Neutrino Energy

Figure 3.13 shows the relative error

$$\frac{E_{reco} - E_{true}}{E_{true}}$$

on the neutrino energy for charged current muon neutrino interactions which satisfy our event selection criteria up to and including the fiducial cut (event selection is described in Chapter 5).

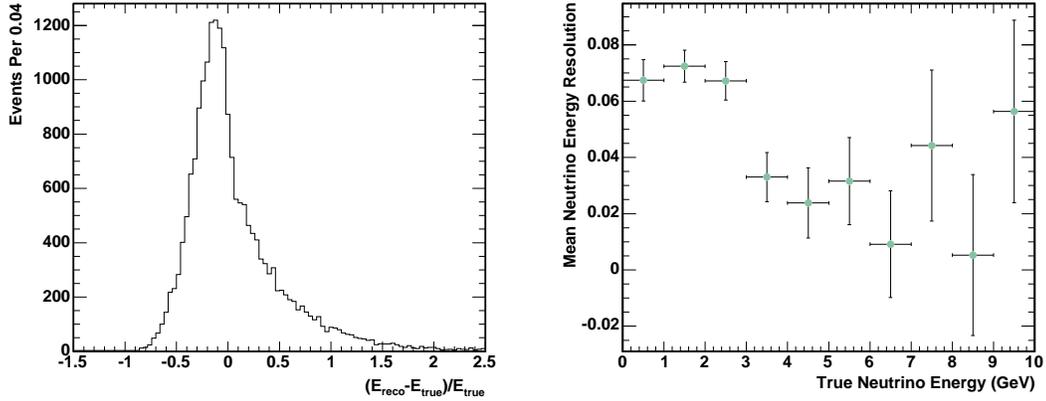


Figure 3.13: The left panel shows the relative error on the reconstructed neutrino energy, calculated from Monte Carlo. The mean is 0.06 and the rms is 0.47. The right panel shows the mean and rms (as an error bar) of the distribution on the left for 1 GeV segments of true neutrino energy.

Neutrino Zenith Angle

Figure 3.14 shows the absolute error

$$\theta_{true} - \theta_{reco}$$

on the reconstructed zenith angle for charged current muon neutrino interactions which satisfy our event selection criteria up to and including the fiducial cut (event selection is described in Chapter 5).

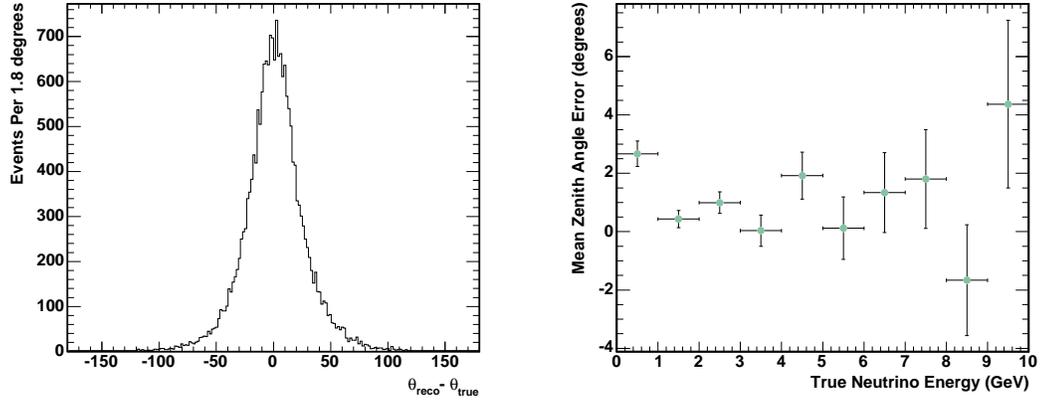


Figure 3.14: The left panel shows the absolute error on the reconstructed neutrino zenith angle, calculated from Monte Carlo. The mean is 1.16° and the rms is 26.7° . The right panel shows the mean and rms (as an error bar) of the distribution on the left for 1 GeV segments of true neutrino energy.

Neutrino $\log(L/E)$

Figure 3.15 shows the relative error on the reconstructed $\log(L/E)$ for charged current muon neutrino interactions which satisfy our event selection criteria up to and including the fiducial cut (event selection is described in Chapter 5). This particular quantity is considered for reasons described in Section 5.3.

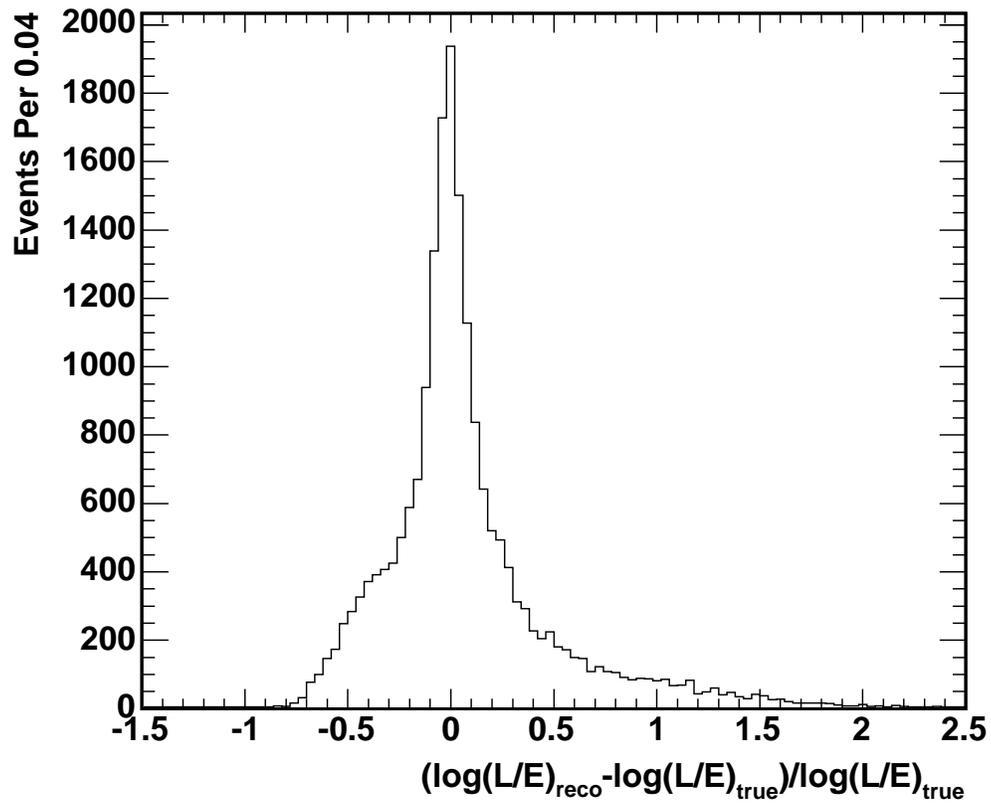


Figure 3.15: The relative error on the reconstructed neutrino $\log(L/E)$, calculated from Monte Carlo. The mean is 0.08 and the rms is 0.43.

Chapter 4

Monte Carlo

Computer simulation has become the bedrock of practically every particle physics analysis, the current one notwithstanding. Tuning cuts, exploring detector efficiency and resolution, estimating sensitivity, and a myriad of other tasks are all done via Monte Carlo (MC). This analysis uses MC to simulate the heavily dominant cosmic ray background to atmospheric neutrino interactions, to discern how various cuts affect the two samples, and to estimate the veto shield efficiency, among several other tasks.

4.1 Monte Carlo Components

The MINOS MC is a full detector simulation based on GEANT3 [27], a software package out of CERN. GEANT is interfaced with a representation of the MINOS geometry and composition of the detector, and it handles the generation and propagation of particles through the detector, handles the secondaries these particles produce, and keeps account of the energy deposited by these particles. Once the energy deposited in the active detector components is known, a special bit of MINOS software known as DetSim/PhotonTransport simulates the propagation of scintillation light to the PMT

photocathode and the MINOS PMT and FEE response.

The bulk of the MC generation process was done using pre-existing code written in Fortran, and was carried out at Cambridge University. Once the Fortran processing was complete, the output was converted to a format compatible with the MINOS C++ software framework.

4.1.1 Atmospheric Neutrino and Cosmic Ray Muon Fluxes

All atmospheric neutrino MC was generated using the Barr solar maximum flux [17]. An ensemble of neutrinos which obeys the Barr flux is generated at the detector and is then handed off to the next stage of MC processing.

The observed flux of cosmic ray muons for a given detector will depend on the depth, overburden profile, and to a small extent, the latitude of that detector. Consequently, it is very difficult to make an absolute prediction of the flux and, rather than attempt such a prediction, MINOS has taken an alternative approach. The Monte Carlo designers have parametrized the above impacts on the flux, generated a large ensemble of cosmic ray muons based on those parametrizations, and then an end user can normalize that ensemble to his or her data sample.

The MC begins by randomly choosing an arrival direction for each muon from a pdf flat in zenith and azimuth. It then picks a muon energy from the distribution shown in Equation 4.1 [68]. In this equation, θ represents the zenith angle and E_μ represents the muon energy (in GeV).

$$\frac{dN_\mu}{dE_\mu} \approx \frac{0.14 E_\mu^{-2.7}}{\text{cm}^2 \text{ sec sr GeV}} \left(\frac{1}{1 + \frac{1.1 E_\mu \cos \theta}{115 \text{ GeV}}} + \frac{0.054}{1 + \frac{1.1 E_\mu \cos \theta}{850 \text{ GeV}}} \right) \quad (4.1)$$

Once a muon has been assigned an energy and direction, the MC checks it against the known Soudan overburden profile [51] to see whether or not the muon would make it underground. This energy cutoff is parametrized by Equation 4.2, which gives the

minimum muon energy necessary to traverse a slant depth X ($\epsilon \approx 500 \text{ GeV}$ and $\xi \approx 2.5 \times 10^5 \text{ g/cm}^2$) [68].

$$E_{\mu}^{min} = \epsilon(e^{X/\xi} - 1) \quad (4.2)$$

As with the atmospheric neutrinos, once a muon is generated, it is handed off to the next MC stage.

4.1.2 Neutrino Event Generation

All atmospheric neutrino MC was generated using NEUGEN3, a Fortran package maintained by Gallagher [43]. The primary purpose of NEUGEN is to simulate a neutrino interaction with a nucleon; to this end, NEUGEN comes with a cross section library valid over an energy range of 100 MeV to 100 GeV. Once NEUGEN conjures an interaction and the resultant particles, be they hadrons from a neutral current scattering or a muon from a ν_{μ} , these particles and their momenta are handed off to GEANT.

4.1.3 Particle Propagation

As mentioned, GEANT lords over the domain of detector composition, geometry, and the propagation of particles with given initial conditions through this virtual detector. GEANT provides methods to describe a real detector by knowing material parameters for many substances, such as steel and scintillator, used to construct particle detectors. GEANT also implements geometry description, so that a realistic form for the detector can be used. It can then transport particles, documenting the particle trajectories, energy deposition, and detector response along the way.

4.1.4 Detector Response

Once a simulated particle has been propagated through the detector and its energy deposition is known, the work of the Fortran legacy code is complete. From here, PhotonTransport and DetSim take over. These C++ packages take a given energy deposition in scintillator, simulate the light produced, propagate that light through the scintillator and optical fiber to the phototube, model the phototube response, and produce an ADC response which can then be processed identically to real data.

4.2 Monte Carlo Samples

For the current analysis, there are four MC samples, one atmospheric ν and one cosmic ray muon sample for tuning cuts and training the neural network (see Chapter 4), referred to as the *tuning* samples, and one atmospheric ν and one cosmic ray muon sample for evaluation of cuts and the neural network, called the *eval* samples. Both of the *eval* samples are scaled to the actual livetime of the data sample. Table 4.1 shows the sizes of the four samples before scaling.

	Atmos Tuning	Cosmic Tuning	Atmos Eval	Cosmic Eval
Kiloton Years	11644.0	2.6	4254.0	1.0

Table 4.1: MC Samples

The scaling of the *eval* samples is achieved in two different ways. To scale the cosmic ray *eval* sample, all throughgoing tracks in both the real data and cosmic ray *eval* samples are counted. A throughgoing track, in this case, is a track which has its vertex higher than $y = 2$ m, its end lower than $y = -2$ m, and its point of closest approach to the center of the detector at a radius smaller than 3 m. All numbers from the cosmic ray *eval* sample are then scaled by the ratio of the number of throughgoing tracks in the data to the number of throughgoing tracks in the cosmic ray *eval* sample.

The reader should note that any subset of the two samples, such as stopping muons, could also be used to find the scale factor.

The atmospheric ν *eval* sample is scaled using the known exposure of the real data. Because the atmospheric MC is generated with the Barr flux, the exposure that a given amount of MC represents is known. Thus all that must be done is to find the exposure, in kt-yr, of the real data sample, and then scale the atmospheric ν *eval* sample by the ratio of the data exposure to the MC exposure. The real data exposure is discussed in Section 5.1.

4.3 Comparison to Data

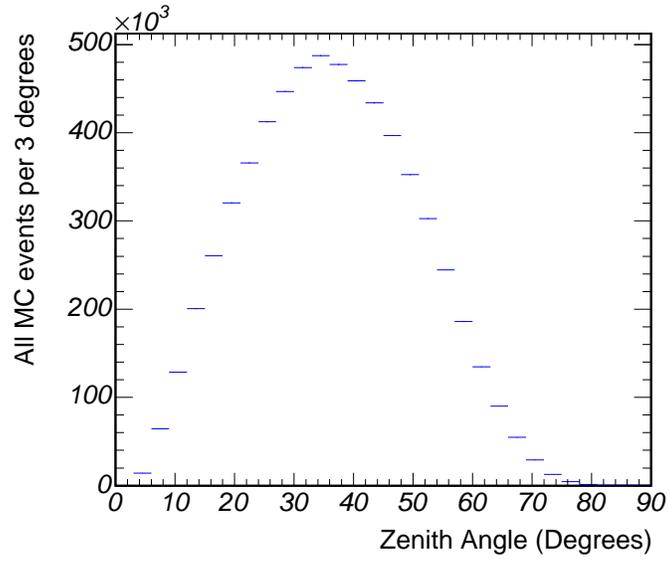
4.3.1 Cosmic Ray Muons

The accuracy with which the cosmic ray Monte Carlo reproduces the collected data must be investigated. Another very thorough study has been done [57], and this study showed that the cosmic ray MC reproduces the data well in most respects. Here, attention will be confined to one quantity of particular import for the current analysis, the zenith angle distribution of reconstructed tracks.

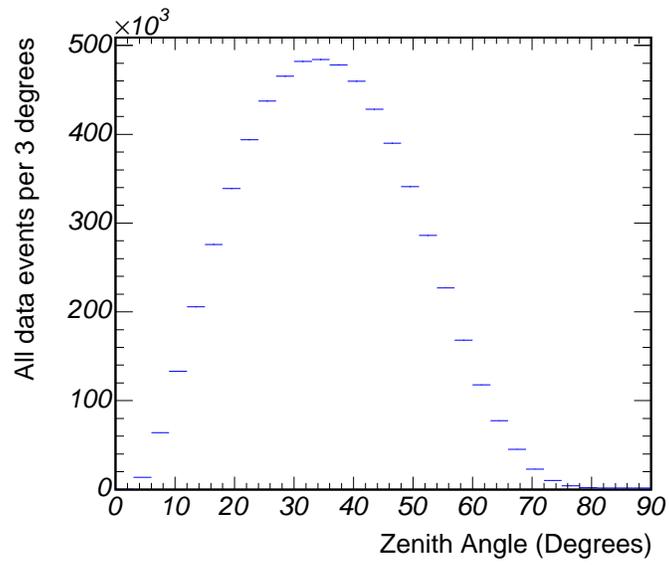
As will be discussed in the next chapter, the cosmic ray muon background is highly asymmetric in zenith angle ($0^\circ \leq \theta_{zen} \leq 180^\circ$), while neutrino-induced muons are isotropic in solid angle. Thus event selection based on the zenith angle of reconstructed tracks is very powerful; see Section 5.2.2. Tuning such cuts requires accurate MC.

A data sample and cosmic ray MC sample, with roughly the same number of reconstructed tracks, have been selected for the purpose of comparing directly their zenith angle distributions. Figure 4.1 shows the zenith angle distribution for cosmic ray muons with no cuts applied.

More elucidating, perhaps, is the ratio of these two distributions, shown in Figure



(a) MC



(b) Data

Figure 4.1: Zenith angle distribution of cosmic ray muons (no cuts).

4.2. At first blush, this ratio does not look very encouraging, as it deviates significantly from a straight line across the entire range $0^\circ \leq \theta_{zen} \leq 180^\circ$. It is believed that the somewhat naive method of generating cosmic ray muons (see Section 4.1.1), coupled with possible errors in the rock map of the overburden at Soudan, are responsible for the discrepancy between the cosmic ray MC and data.

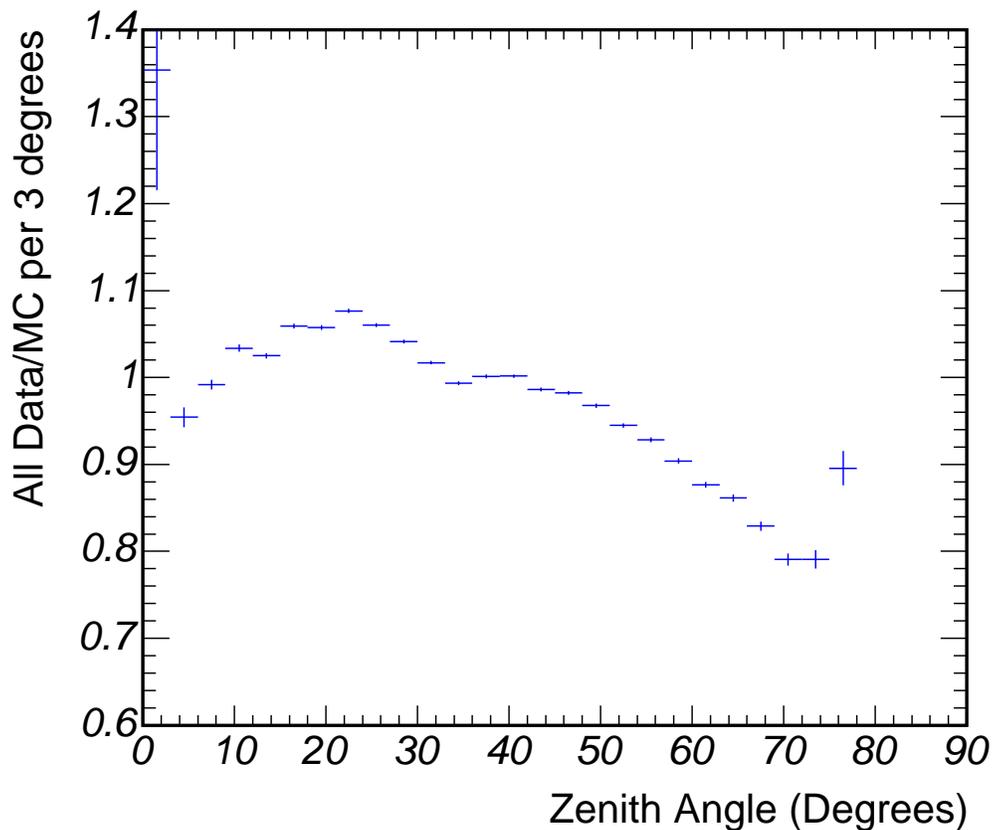


Figure 4.2: Ratio of data and MC zenith angle distributions for all tracks.

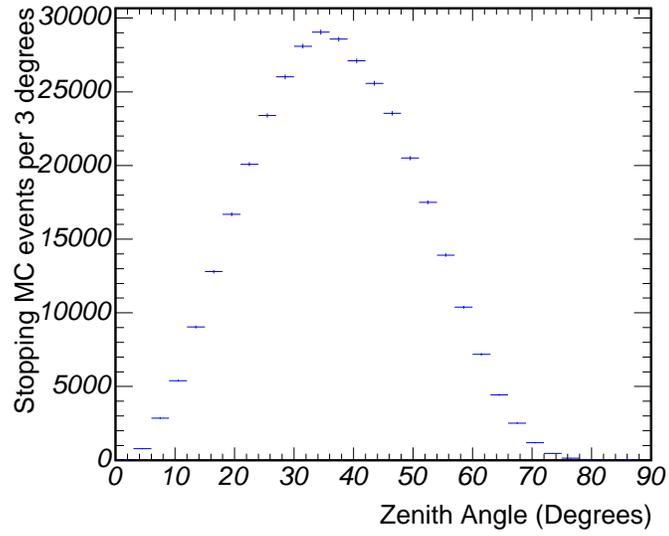
Fortunately, this analysis is based on only fully contained events. As such, the throughgoing muons, a large fraction of this total sample, are of little interest. Consequently, a sample of stopping muons, muons which stop in the fiducial volume of the detector, is more representative of the relevant background confronted. Things are

slightly better for the stopping muons. Still, $\sim 20\%$ accuracy in the MC is far from ideal. Fortunately, all is not lost. First, as will be discussed in Section 5.2.2, one of the early event selection stages is a hard cut on the vertex track y direction cosine that translates into requiring the track zenith angle to be larger than $\sim 37^\circ$. As such, any disagreement between data and MC in the region $\theta_{zenith} < 37^\circ$ is, from a pragmatic point of view, irrelevant. Second, in the region $\theta_{zenith} > 37^\circ$, the MC actually *overestimates* the background by as much as 20%. Consequently, any cuts which are tuned on this MC sample will be, if anything, too aggressive and produce a *signal : background* ratio which is worse than what it is in reality.

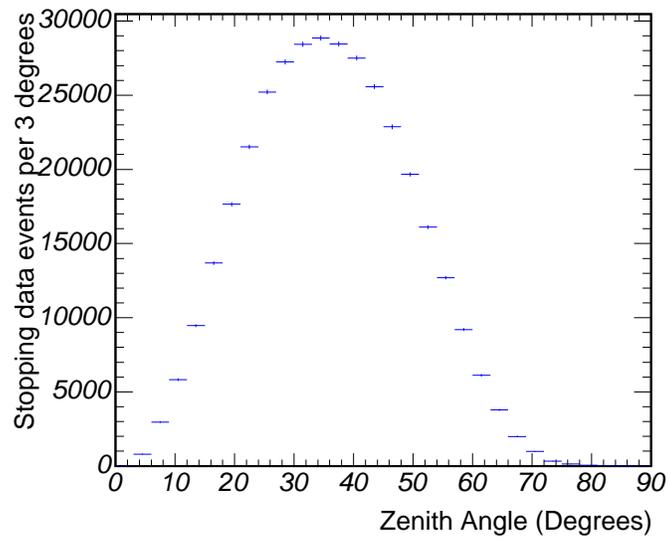
Figure 4.3 shows the zenith angle distribution of stopping muons in data and MC. By eye, it is easy to see that selecting stopping muons alters slightly the shape of the zenith angle distribution; Figures 4.1 and 4.3 show small differences in shape for zenith angles less than 35° . The data to MC ratio, however, is somewhat flatter for stopping muons, as shown in Figure 4.4.

4.3.2 Atmospheric Neutrinos

No direct comparison to data is possible for the atmospheric neutrino MC, as it is impossible in real data to say on an event-by-event basis whether a muon is neutrino-induced or from cosmic rays. Uncertainties in the atmospheric neutrino MC are discussed in Chapter 6.



(a) MC



(b) Data

Figure 4.3: Zenith angle distribution of stopping muons.

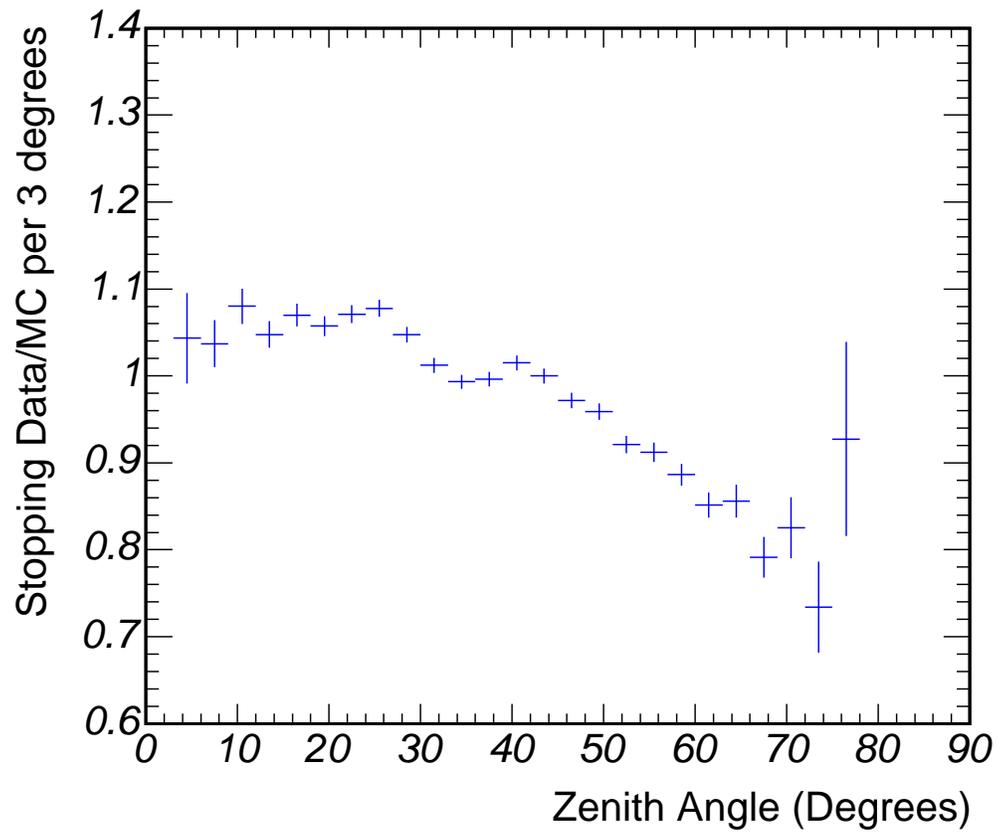


Figure 4.4: Ratio of data and MC zenith angle distributions for stopping muons.

Chapter 5

Data Samples and Event Selection

The subset of neutrino events in the entire MINOS data sample can be analogized to a message in a bottle adrift in the ocean. The neutrinos carry desired information, but are completely lost in the enormity of the background. This analysis uses a combination of techniques to divide signal and background, and this chapter describes the various attacks employed. Sections 5.1 and 5.3 detail the MINOS data set before and after event selection, and Section 5.2 is divided into subsections dedicated to each stage of the event selection.

5.1 Initial Data Sample

This analysis is based on data collected by the MINOS Far Detector from 2003.08 to 2006.04. Construction of the MINOS detector took place from roughly 2001.07 to 2003.08. As new planes were added, they were commissioned and activated in groups of four. Thus there exists 5 kt-yr of data not included in this analysis. There are several

good reasons for excluding this data. First, each magnetic coil can magnetize an entire supermodule (SM) only; consequently all of the data before SM1 was complete and all of the SM2 data before SM2 was complete were taken with the field in absentia. Second, the veto shield, critical for this analysis, took on several configurations during the construction period before its present, steady-state configuration.

2006.04 was not the end of the experiment. The MINOS detectors, and the NuMI beam, are all currently running and are scheduled to run for another four years. Data collected after 2006.04 is not included in this analysis because the calibration lags data collection and was not complete at the time of this writing.

Another consideration is the quality of data being collected. It is possible for a subsystem of the detector to fail without stopping data acquisition. For example, parts of the detector at various times have lost high voltage to the phototubes, leaving a completely dead section. Also, diagnostic runs must be removed from further consideration. The Cambridge group has designed a very thorough data quality check to flag unacceptable runs [23]. We use their run list in this analysis.

The number of seconds spanned by each run has been counted, and these individual run lengths have been summed. All told, our initial data sample is 7.63×10^7 sec long, which amounts to 883 live days, and 13.1 kt-yr.

5.2 Event Selection

Extraction of atmospheric ν_μ events from the overall MINOS data sample is an extremely challenging problem. The MINOS detector records approximately one cosmic ray muon every two seconds; it records one neutrino induced muon per day. Separating signal from background when the initial signal to background ratio is of the order $10^5 : 1$ requires some ingenuity, or at the very least a stubborn resolve.

5.2.1 Quality Control

No reconstruction is perfect, and many events are low energy and yield little useful information. As such, the first step in event selection is to make quality control cuts which ensure that the reconstruction has accurately reconstructed the event and that the event has nominal physics potential.

This first stage of event selection requires

- Each event must contain one reconstructed track.
- The track must contain at least six planes total.
- The track must have at least three strips in each view, u and v .
- The sign of the difference between the vertex and end y positions must be consistent with what the y vs T (time) slope dictates
- $S_s < 3T_s$ if $S_s > 2$ m. Here, S_s is the transverse span of the strips, $S_s = (stpmax - stpmin)$, in one view (u or v), and T_s is the transverse span of the reconstructed track, $T_s = (trkmax - trkmin)$, in the same view. This requirement is imposed on both views.
 - $stpmax(stpmin)$ is the highest(lowest) transverse position among all strips in a view.
 - $trkmax(trkmin)$ is the highest(lowest) transverse position on the track in a view.
 - Requiring an absolute strip span of at least 2 m guards against cutting out very flat tracks.
 - **This cut eliminates poorly reconstructed events.**

- $\frac{T_r}{N_T} < 1.4$, where T_r is the range (in plane number) of the track, $T_r = |N_l - N_f| + 1$, and N_T is the total number of planes the track contains. In MINOS-specific jargon,
 - The numerator is the number of planes spanned by the track, obtained by $N_f = trk \rightarrow plane.beg$ and $N_l = trk \rightarrow plane.end$ in the MINOS standard reconstruction ntuple.
 - The denominator is the number of planes with digits on the track, obtained from $N_T = trk \rightarrow plane.n$ in the MINOS standard reconstruction ntuple.
 - **This cut eliminates events with large gaps.**

5.2.2 Kinematic Cuts

The cosmic ray background is highly asymmetric, totally downgoing, and consequently has some very distinct kinematic characteristics. Thus it is safe to cut on the track vertex direction cosine y ($\cos \theta_y$) and z ($\cos \theta_z$) to significantly reduce the background while minimally affecting the signal. The specific cuts are

- $\cos \theta_y > -0.8$
- $|\cos \theta_z| > 0.2$

These cuts reduce the remaining background by 57% and the remaining signal by 7% and are illustrated in Figure 5.1.

5.2.3 Fiducial Cuts

The main difference between cosmic ray muons and muons resulting from ν_μ interactions is most cosmic ray muons enter the top of the detector and leave the bottom. Thus, a fiducial cut is a common and very effective step in eliminating the cosmic ray background. Further, this analysis employs an asymmetric containment cut because of

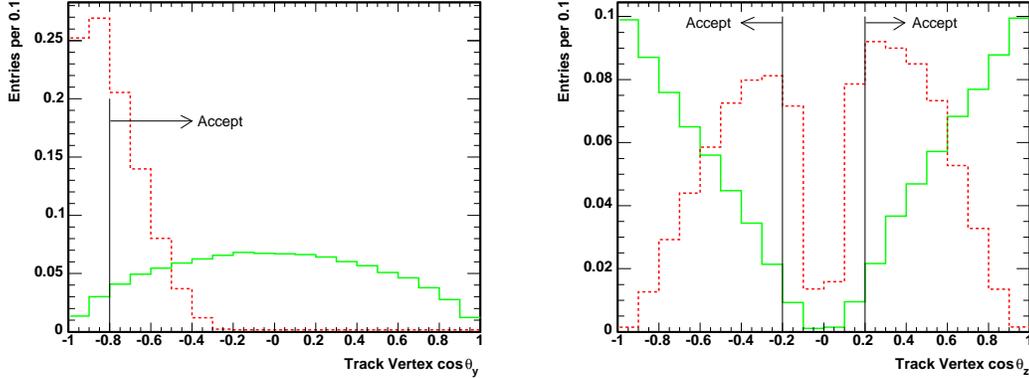


Figure 5.1: Track vertex direction cosine y (left) and z (right) before cuts. Red (dashed) is cosmic MC, green (solid) is atmos MC.

the inherent asymmetry in the background, which is overwhelmingly downward going, and the signal, which is isotropic.

The first step in estimating containment is to calculate the y vs t slope of the muon track. The MINOS coordinate system has the positive y axis pointing straight up (opposite gravity). As such, a track with negative slope is downward going and one with positive slope is upward going. For the former, an octagonal fiducial region 3.35 m from center to edge is assigned; for the latter, 3.75 m is used as the center to edge distance. In both cases, both the track vertex and end are required to be within the fiducial region. The first and last five planes of each supermodule are also excluded, as is a cylindrical region of 38 cm at the center of the detector. These two areas are excluded because many cosmic ray muons enter the detector through its ends and the coil hole, neither of which are covered by the veto shield (see Section 5.2.6). The coil hole is particularly problematic because, should a muon enter it roughly parallel to the z axis of the detector, the muon can penetrate quite far into detector without producing any hits. The 38 cm excluded volume at the center of the detector ensures none of these “coil muons” produce false fully contained neutrino events.

Rather than require simply the track vertex and end lie within the fiducial region, a more complicated containment metric is used because of the idiosyncrasies of the MINOS detector and reconstruction software. Many events have a track which is fully contained but strips not included in the track which lie outside the fiducial region. This event pathology often arises from very steep cosmic ray muons which penetrate deep into the detector before interacting and then sharply scatter into the detector. Cutting only on the track vertex and end positions does not remove these events, which is not the desired decision. Thus the strip positions must also be considered. To this end, a least squares fit is performed on the (u, z) and (v, z) positions of all strips with summed charge greater than 1.5 pe. The $v(u)$ fit is then used to produce a $v(u)$ coordinate for a $u(v)$ strip at its Z position. A complete (u, v) point for every strip in the detector is thus obtained and all points are required to lie within an octagonal fiducial region 3.9 m from center to edge. Figure 5.2 shows a diagram of these various fiducial regions.

The final containment decision is the *AND* of the vertex/end containment and the strip level containment.

As mentioned, containment is a powerful discriminator between signal and background. The various fiducial distances have been optimized on MC, and these optimal values yield a 99.89% reduction of remaining background. Regrettably, 50.4% of the remaining signal is sacrificed.

5.2.4 Neural Network

Unfortunately, containment alone is not enough to separate signal from background. Further effort must be expended on extracting signal events from the sea of background events.

Many of the cosmic ray muons which pass the containment requirement look, to first order, the same as neutrino induced muons. So making simple cuts on kinematic and

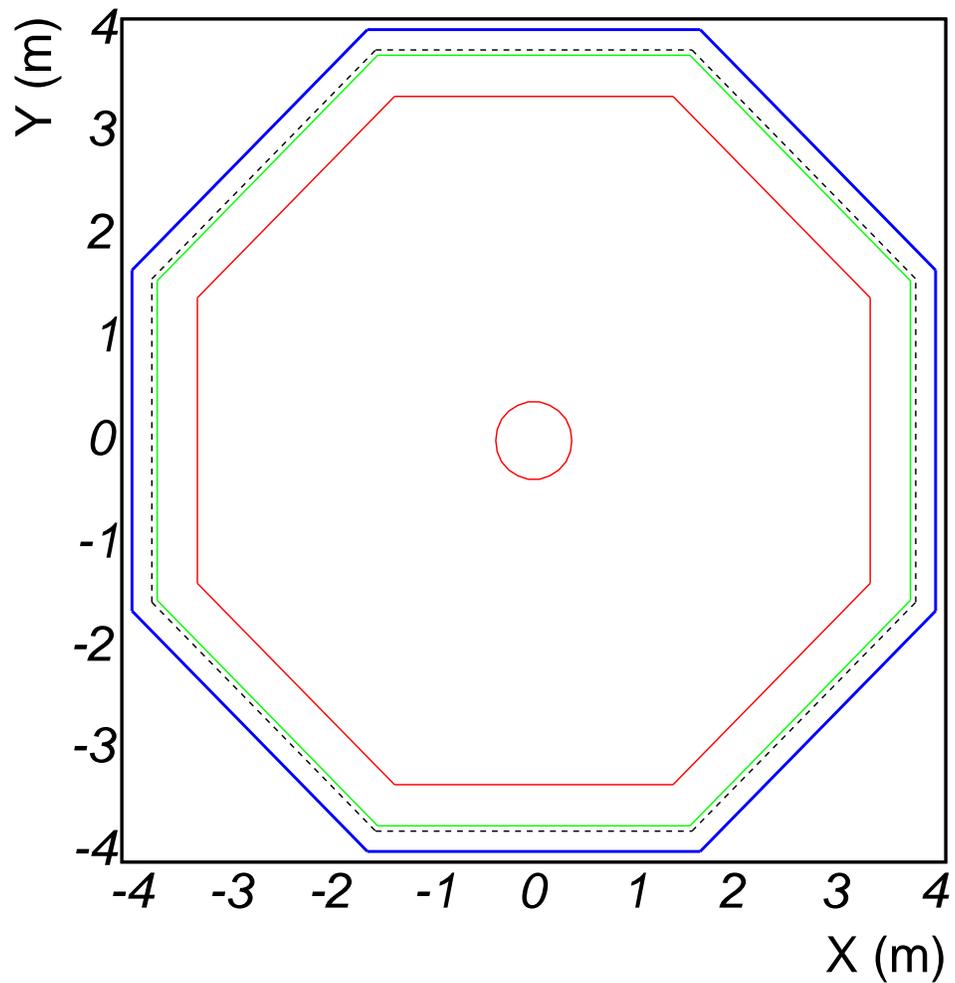


Figure 5.2: The fiducial regions used in the containment decision. From the outside in, the lines show the detector outline (blue), the region for the strip level containment (dashed black), the fiducial region for upward going tracks (green), and the region for downward going tracks (red). At the center is the excluded region around the magnetic coil.

topological variables is not efficient enough in rejecting background while simultaneously preserving a useable number of signal events. As such, a neural network (NN) has been employed to carry out event classification in this analysis.

The NN is a standard Multi-Layer Perceptron (MLP) backpropagation network [30] with one hidden layer between the input layer and output node. Figure 5.3 illustrates such a NN (the one this analysis actually employs has nine inputs and nine hidden nodes, but is otherwise the same as the simplified diagram). One can imagine, in the diagram, information flowing along the lines from left to right.

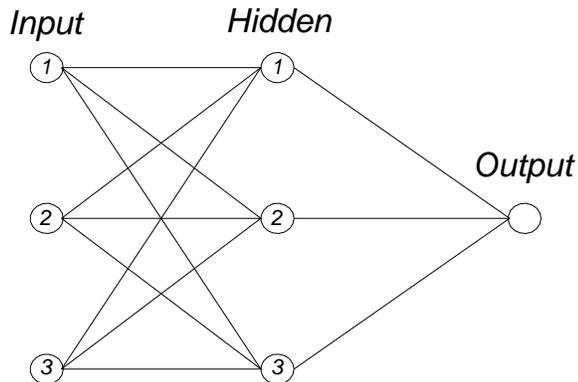


Figure 5.3: Schematic of a neural net with three inputs, one hidden layer with three nodes, and one output node.

The input nodes receive, as discussed below, values which come from the event reconstruction and which serve to characterize the track and event. These values are scaled so that they always lie between -1 and 1 , and they are next fed into the hidden layer nodes. Each hidden layer node receives as input a weighted sum of the values from the input nodes, and each connection has a unique weight. For example, the input of the first hidden node in Figure 5.3 is $W_{11}I_1 + W_{12}I_2 + W_{13}I_3$, where W_{1i} is the weight between the first hidden node and the i^{th} input node and I_i is the scaled input node

value described above. The output of a hidden layer node is this sum run through the node’s activation function. The activation function for each hidden node (and the output node) is the usual sigmoid function.

$$f(t) = \frac{1}{1 + e^{-t}} \quad (5.1)$$

The output node receives as input a weighted sum of the hidden node outputs, in complete analogy with the dynamic between the hidden and input layers. The final output of the output node is again analogous to the previous stage; the weighted sum the output node receives as input is run through the activation function. The result represents the NN’s final decision on what sort of event it has been presented with.

Suitable input to a NN are variables which have some discriminating power between cosmic ray muons and neutrino induced muons. For example, the y position of the track vertex statistically looks different for signal and for background; even with containment applied, the cosmic ray vertex y distribution is peaked toward higher y values while for neutrino induced muons, the vertex y distribution is fairly flat. Nine such quantities have been identified: the number of strips in an event (Figure 5.5), the track vertex y (Figure 5.6), the track vertex direction cosine y (Figure 5.7), the track vertex direction cosine z (Figure 5.8), the track end direction cosine z (Figure 5.9), the track vertex z trace (Figures 5.4 and 5.10), the total charge within ± 3 planes of the track vertex (Figure 5.11), the event charge per plane (Figure 5.12), and the fraction of the event charge contained in the track (Figure 5.13). With these quantities in hand, the NN is then trained using the training MC sample of cosmic ray muons and atmospheric neutrinos.

Training is the process of iteratively dragging the internode weights to their “correct” values, values which result in distinct net output when the NN is presented with a cosmic ray muon and a neutrino induced muon. The NN will put out a number between 0 and 1, and the output has been (arbitrarily) chosen so that 0 is the target for cosmic

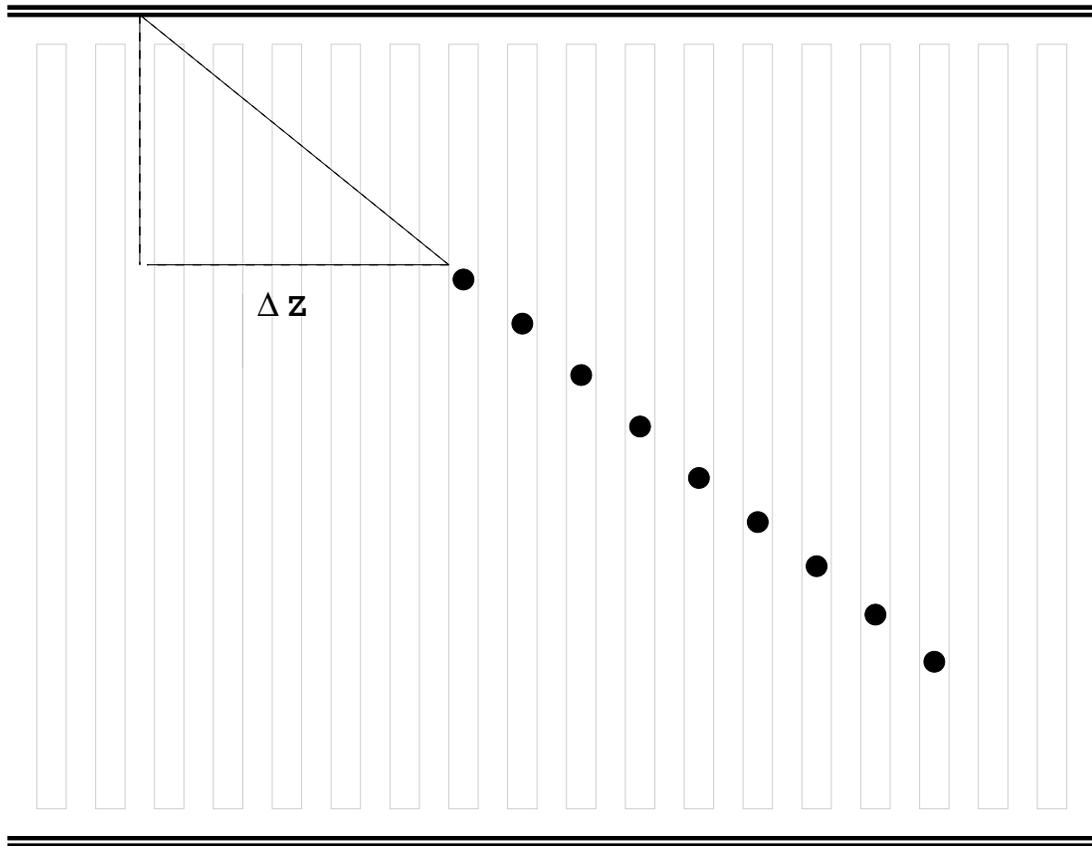


Figure 5.4: A diagram of the vertex z trace. The light gray boxes represent planes of the detector, the black bars are the detector edges, and the dots show track hits. The track is projected back to the edge of the detector, and the z distance covered by the projection between the track vertex and the detector edge (Δz) is the trace.

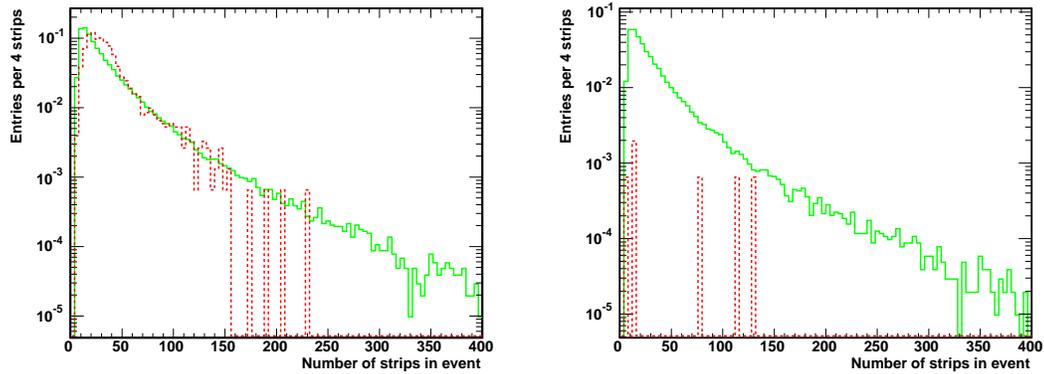


Figure 5.5: Number of strips contained in an event before (left) and after (right) NN cut. Red (dashed) is cosmic MC, green (solid) is atmos MC.

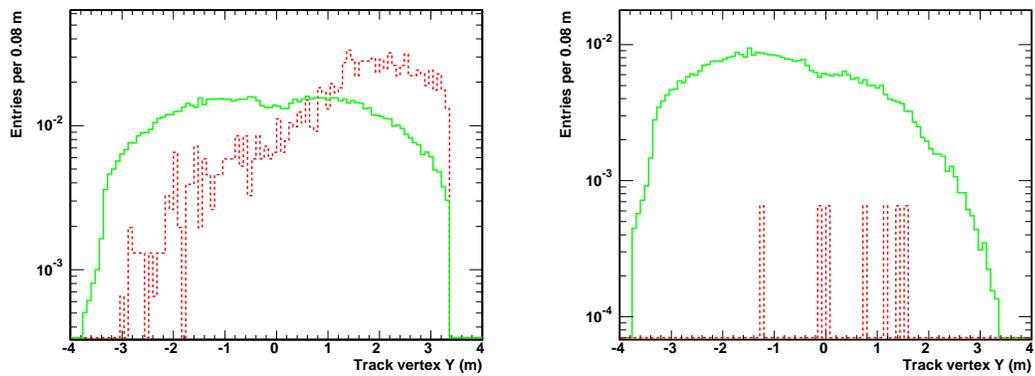


Figure 5.6: Track vertex y (m) before (left) and after (right) NN cut. Red (dashed) is cosmic MC, green (solid) is atmos MC.

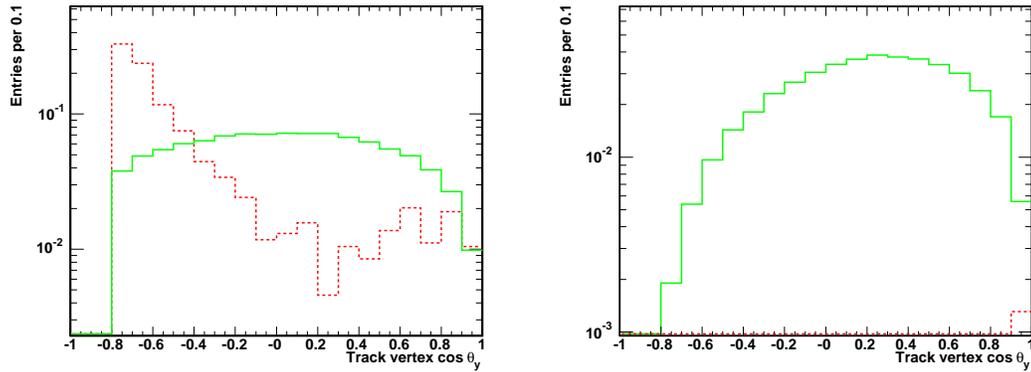


Figure 5.7: Track vertex y direction cosine before (left) and after (right) NN cut. Red (dashed) is cosmic MC, green (solid) is atmos MC.

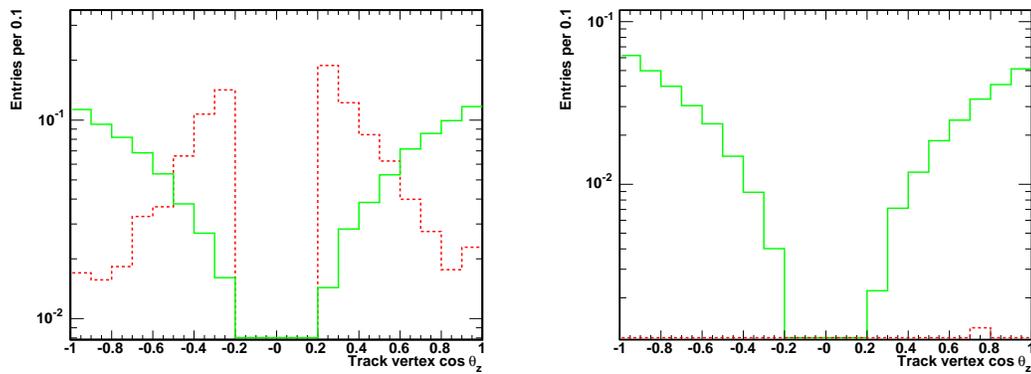


Figure 5.8: Track vertex z direction cosine before (left) and after (right) NN cut. Red (dashed) is cosmic MC, green (solid) is atmos MC.

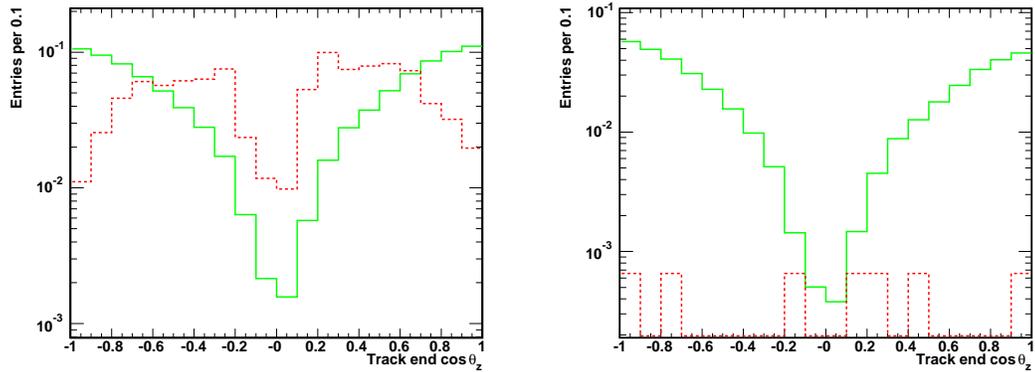


Figure 5.9: Track end z direction cosine event before (left) and after (right) NN cut. Red (dashed) is cosmic MC, green (solid) is atmos MC.

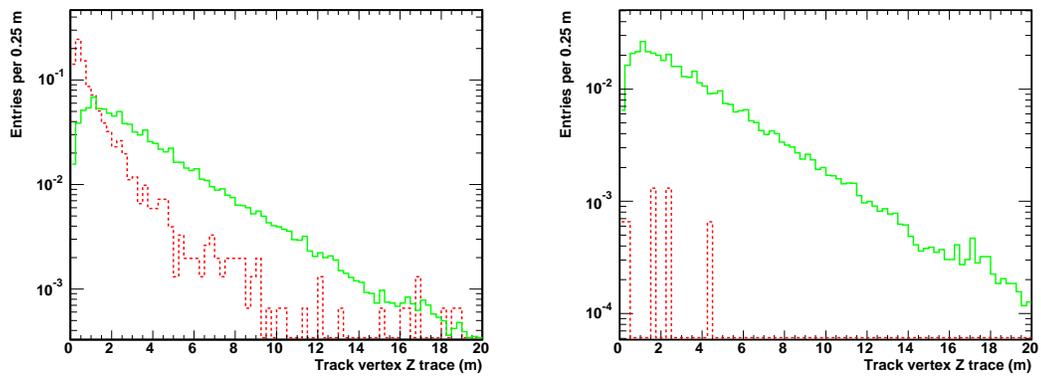


Figure 5.10: Track vertex z trace (m) before (left) and after (right) NN cut. Red (dashed) is cosmic MC, green (solid) is atmos MC.

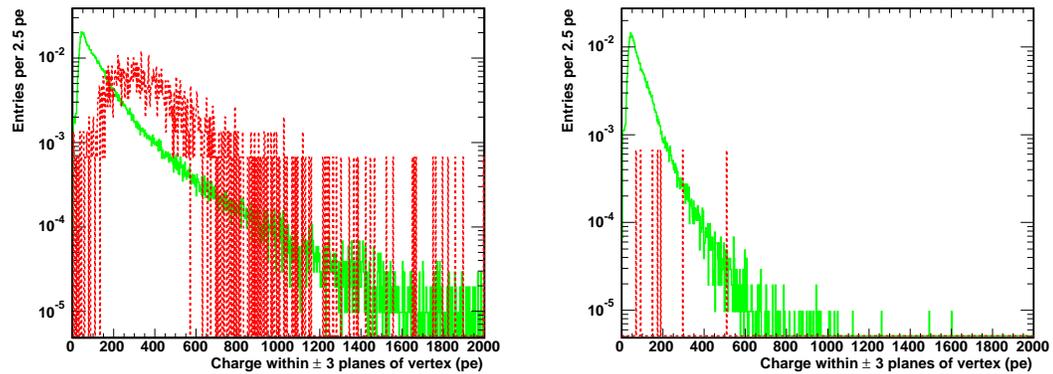


Figure 5.11: Total charge (pe) within ± 3 planes of vertex before (left) and after (right) NN cut. Red (dashed) is cosmic MC, green (solid) is atmos MC.

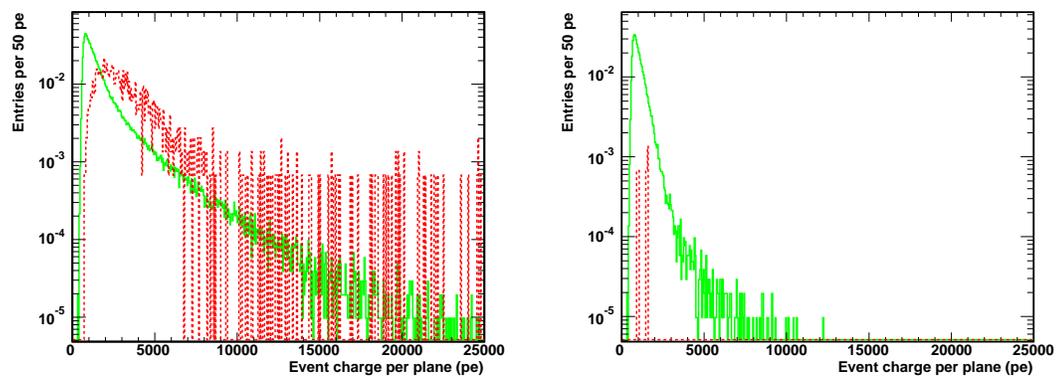


Figure 5.12: The charge per plane (pe) before (left) and after (right) NN cut. Red (dashed) is cosmic MC, green (solid) is atmos MC.

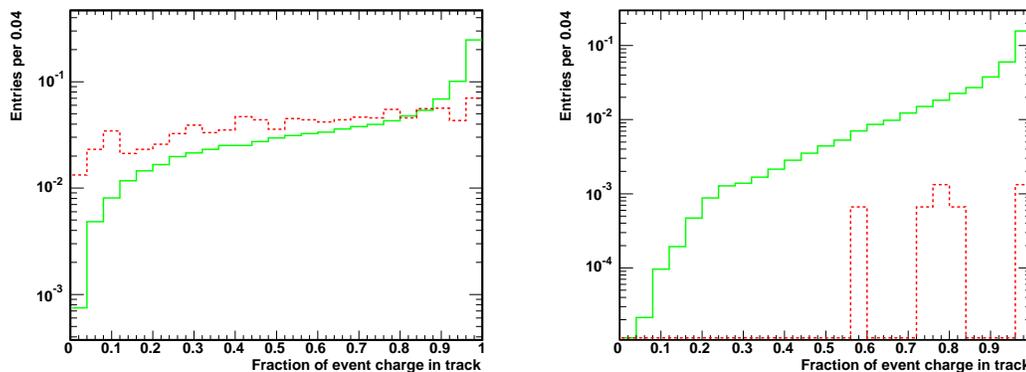


Figure 5.13: The fraction of the event charge the track contains before (left) and after (right) NN cut. Red (dashed) is cosmic MC, green (solid) is atmos MC.

ray muons and 1 is the target for neutrino induced muons. To train the NN, all the internode weights are initialized to a small, arbitrary value (0.1 in this case). The the NN is presented with a known event type. Once the final output is calculated, the weights are adjusted by applying a correction, unique to each weight, which is proportional to the difference between this output and the target output (e.g. 0 for cosmic ray muons). See [30] for more detail on the derivation of these corrections.

A crudely adaptive learning rate has been employed; the learning rate tends to zero as iterations (known as *epochs*) in training the net are completed. The adaptive learning rate, along with direct observation of the net behavior as a function of the number of training epochs, ensure that the net is not overtrained.

Once training is done, the MC evaluation sample mentioned earlier is then analyzed to ascertain the effectiveness of the NN. Each event in both the cosmic ray and atmospheric neutrino samples is presented to the NN, and its output is logged. Ideally, for a cosmic ray muon, the NN should output 0.; for a neutrino induced muon, 1. In reality, there is a range of output, and by optimizing on the Figure Of Merit (FOM) shown in Equation 5.2, the optimal cut on NN output is found to be to accept everything above

0.92. See Figures 5.14 and 5.15. As one goes to higher net output, the number of cosmic ray muons left decreases much more rapidly than the number of atmospheric neutrino events because almost all of the latter events are clustered toward 1. As such, the FOM steadily increases as the net cut varies from 0.0 to 0.9.

$$FOM = \frac{signal^2}{signal + background} \quad (5.2)$$

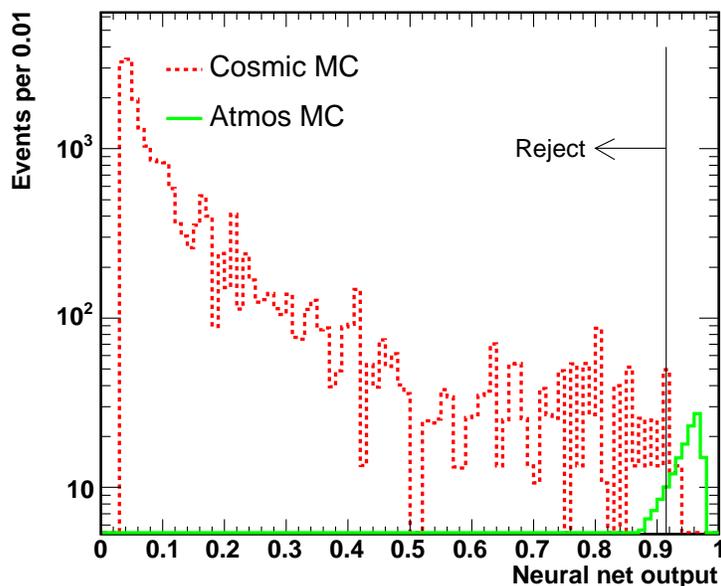


Figure 5.14: NN output distribution for atmospheric and cosmic MC. Red (dashed) is cosmic MC, green (solid) is atmos MC.

5.2.5 Beam Timing

As mentioned in Chapter 2, MINOS is only half of the NuMI/MINOS experiment, and its primary goal is to observe neutrinos generated in Batavia, Illinois, not in the atmosphere. This beam of neutrinos turned on in May of 2005, and the data set for this

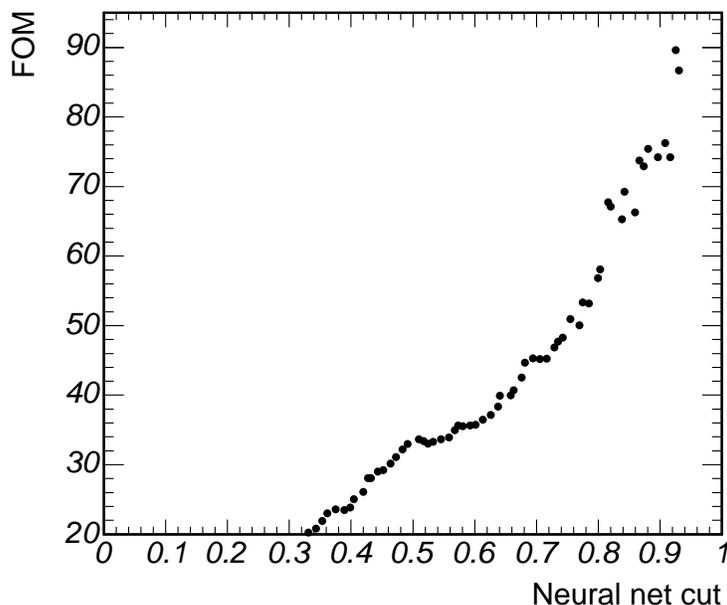


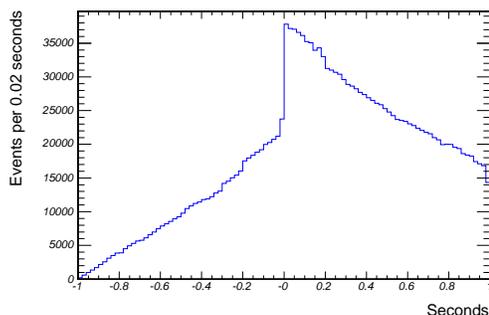
Figure 5.15: Figure of merit as a function of neural net cut (from MC).

thesis runs through April of 2006. Consequently, care must be taken to remove beam neutrino events from this last year of data.

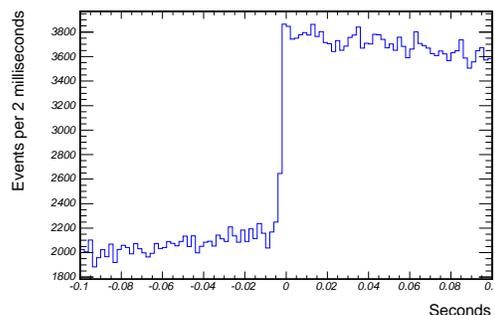
The time of each beam spill is stored in the MINOS database. A simple comparison between any event trigger time (T_{trig}) in the far detector and the nearest beam spill time (T_{spill}) to that trigger time allows complete removal of beam events with little penalty in the form of downtime.

Two months of data, 2005.11 and 2006.03, have been analyzed to evaluate the utility of the beam timing information. Figure 5.16 shows the distribution of snarl trigger times relative to the time of the nearest beam spill for these two months. As the progression of successively narrower time intervals shows, the beam events are well isolated on the ten microsecond scale. Further, previous work [64] has shown that the amount of time when the beam is live but there is no record of the beam spills, because of timing system or beam monitoring failures, is less than 1% of the total beam live time. Consequently,

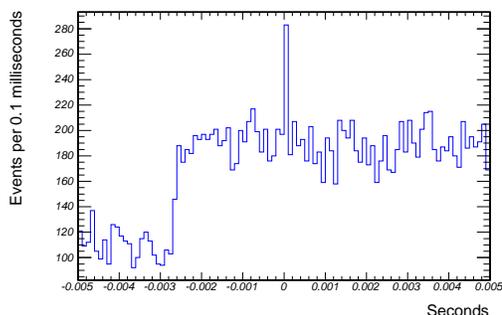
it may be safely required that $-2. \times 10^{-5} \text{ sec} < (T_{trig} - T_{spill}) < 3. \times 10^{-5} \text{ sec}$.



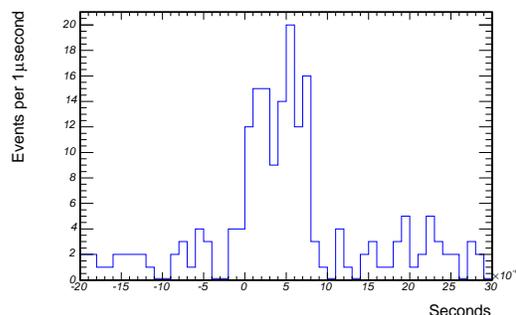
(a) All time differences (-1 to 1)



(b) Time differences from -0.1 to 0.1



(c) Time differences from -0.005 to 0.005



(d) Time differences from $-2. \times 10^{-5}$ to $3. \times 10^{-5}$

Figure 5.16: Beam timing distributions from data collected during 2005.11 and 2006.03. All times are in seconds.

5.2.6 Shield

In the conquest of background, there is one final weapon, the veto shield. As mentioned in Chapter 2, the veto shield achieves nearly complete coverage of the top and sides of the MINOS detector and is read out in an identical fashion to the detector.

To make the most efficient use of the veto shield, the timing of shield hits relative to the time of the event in the detector must be considered. Doing so significantly

reduces accidental tagging with minimal impact on the rejection efficiency for cosmic rays. Further, timing is more precise than attempting to correlate spatially the track projection and the shield hits.

From Figure 5.17, one can make a reasonable guess about the optimal time window inside which to look for shield hits. The specific requirements are that

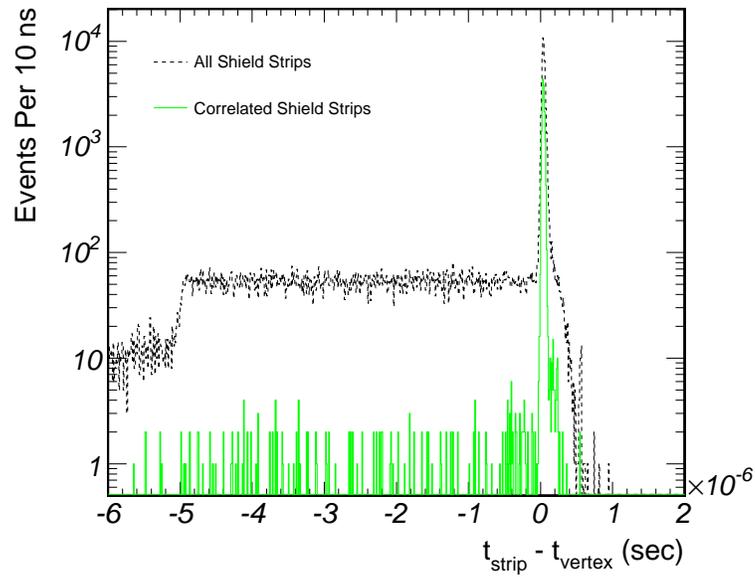
- the time of the strip is between -50 and 150 ns relative to the track vertex time
- the track vertex must lie between the ends of the shield strip in z

Should any shield strip satisfy these requirements, the event is considered vetoed.

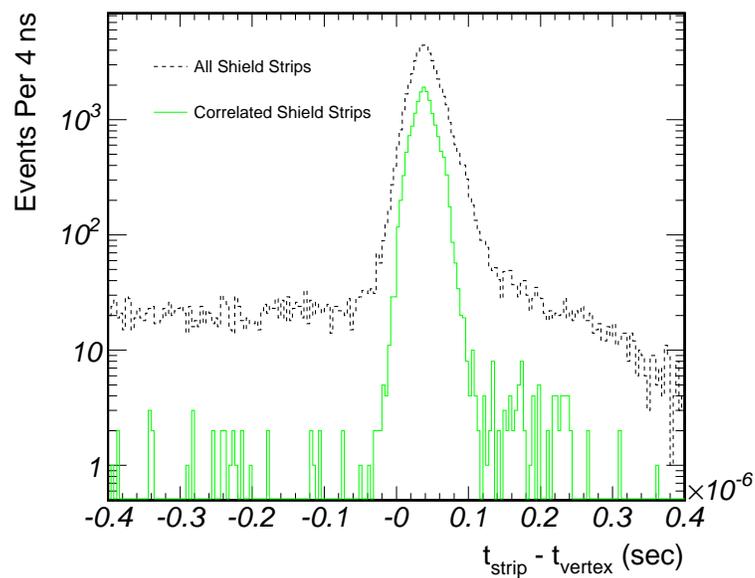
The background rejection and signal acceptance of these cuts has been measured using real data and atmospheric MC. To measure the background rejection, a subsample of cosmic ray muons which do not pass the NN cut has been selected. These events are independent of the events actually accepted. As these events do pass the event selections leading up to the NN, they are an accurate representation of the real background confronted. Specifically, events are considered which have a NN output between 0.05 and 0.8. For the entire data sample, there are 14092 such events, and 13473 of these are tagged by the veto shield. Thus the background rejection efficiency is $95.6 \pm 0.8\%$

An attempt has been made to quantify the systematic error associated with this background rejection efficiency. In words, three separate intervals of neural net output between 0.05 and 0.8 have been considered independently, and any differences between the three intervals should indicate the size of the systematic error in question. Table 5.2.6 contains the results of this study. The largest deviation from the overall efficiency of 95.6% is 1.3%, and this number is taken as the systematic error.

The ideal way to measure the signal acceptance of the shield is to implement a veto shield simulation in the MC. Unfortunately, this portion of the MC is not yet complete. For a neutrino event, where no particle actually deposits energy in the shield, the essence



(a) Full time scale



(b) Zoomed

Figure 5.17: The distribution of shield strip times relative to the track vertex time (from data).

	Net Range	Total/Tagged	Efficiency	Deviation
Interval 1	$0.05 < NN < 0.1$	4437/4269	96.2%	+0.6%
Interval 2	$0.1 < NN < 0.23$	4740/4571	96.4%	+0.8%
Interval 3	$0.23 < NN < 0.8$	4915/4633	94.3%	-1.3%

Table 5.1: The breakdown of shield efficiency for three similarly sized ranges of net output. NN denotes the net output value.

of a shield MC would be the production of shield hits which, statistically, reproduce the known spatial and temporal distributions of the random shield noise in the real data. These hits would then serve as shield readout for a MC event. As such, a shield readout from a data snarl where there is no cosmic ray muon present would be suitable “fake” shield readout for the atmospheric MC. To select such snarls, a small sample of real data was processed using the reconstruction noise filter, described in Section 3.6.1, in reverse mode. Using the noise filter “flipped” ensures that any event which passes has minimal activity in the detector and thus has totally random shield hits, not hits caused by a real cosmic ray muon. All told, more than 10^5 of these random shield readouts have been collected, which is large enough to ensure that, in processing the atmospheric MC, no shield readout is used more than once.

When processing the atmospheric MC, a shield snarl from the ensemble is randomly selected, and then the whole event (including the shield hits) is subjected to the same cuts previously described to decide whether or not the event is tagged by the shield. Using this method, we find that the shield has a $2.8 \pm 0.5\%$ accidental signal rejection.

5.3 Final Data Set

All told, 105 events are selected. The effect of each cut is shown in Table 5.2.

Table 5.2 requires further explanation. The “Cosmics” column is derived from the cosmic ray *eval* sample, described in Section 4.2. The “Shield” entry in this column is the

	Cosmics	ν BG	ν CC	Data
Total	$(3761.2 \pm 2.2) \times 10^4$	730.9 ± 1.5	1077.5 ± 1.8	4.1×10^7
Quality	$(3196.7 \pm 2.1) \times 10^4$	114.8 ± 0.6	589.7 ± 1.3	3.1×10^7
Kinematic	$(1375.0 \pm 1.4) \times 10^4$	111.8 ± 0.6	549.4 ± 1.3	1.3×10^7
Fiducial	$(204.4 \pm 5.2) \times 10^2$	66.0 ± 0.5	251.8 ± 0.9	20190
Neural Net	53.8 ± 26.9	5.9 ± 0.1	115.7 ± 0.6	283
Beam Timing	53.8 ± 26.9	5.9 ± 0.1	115.7 ± 0.6	177
Shield	2.4 ± 1.2	5.7 ± 0.1	112.4 ± 0.6	105

Table 5.2: The number of data and MC events left after each cut.

product of the shield background rejection efficiency listed in the previous section and the number of events left after the neural net cut. Interestingly, one can also compare this number, 2.4 events, with the result obtained by counting the number of real data events which pass the neural net but *fail* the shield cut. This total number, 72 ± 8.5 , scaled by the shield background rejection inefficiency, 4.4% (discussed in Section 5.2.6), yields 3.2 ± 0.4 . Reassuringly, the two values are consistent.

The “ ν BG” “ ν CC” columns are derived from the atmospheric ν *eval* sample, also described in Section 4.2. In the background category, both neutral current events for all neutrino flavors and charged current ν_e and ν_τ (should there be any above the threshold energy) events have been grouped together. In all cases, these interactions produce showering events which should be rejected. The “ ν CC” column contains only ν_μ charged current events. The “Shield” entry for both the background and charged current columns is found in the same manner as just described for the “Cosmics” column, except that the relevant efficiency is now the signal acceptance, which is $97.2 \pm 0.5\%$, as described in the previous section.

Lastly, the beam timing cut is taken to have no impact on the cosmic ray background, the neutrino backgrounds, nor the neutrino signal. The punchline of Table 5.2 is that 120.5 ± 1.3 events total are expected, and 105 are observed.

These events are evenly distributed spatially through the detector, as Figure 5.18 shows. The events also “arrive” at a constant rate in time, as demonstrated by Figure 5.19. The neutrino kinematic y distributions from data and MC are shown in Figure 5.20, and the zenith angle and energy distributions are shown in Figure 5.21.

Oscillations most clearly manifest themselves in the $\log(L/E)$ distribution of the final data set, where L is the length (in km, see Equation 3.2) the parent neutrino traveled from production to interaction, and E is its energy (in GeV). The $\log(L/E)$ distribution of the final data set is shown in Figure 5.22.

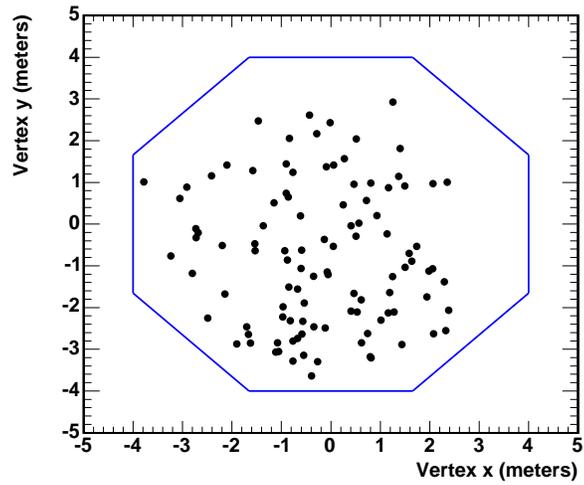
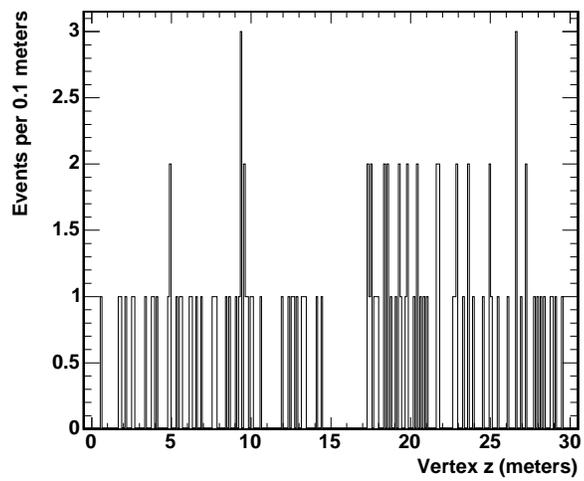
(a) Vertex y vs x (b) Vertex z

Figure 5.18: The spatial vertex distribution of selected events.

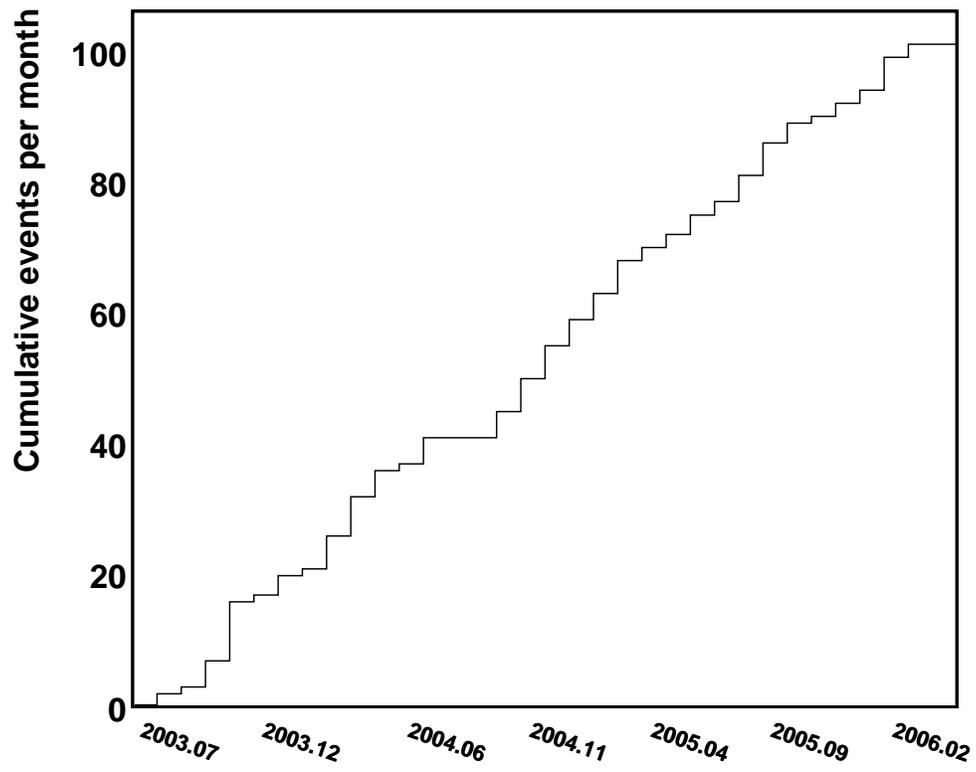


Figure 5.19: The summed number of selected events month by month for the entire livetime.

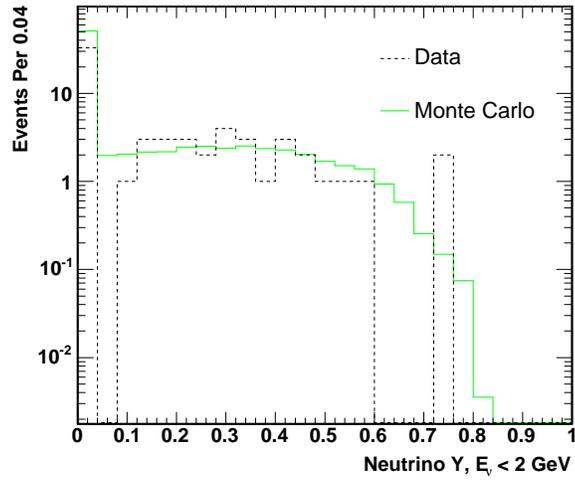
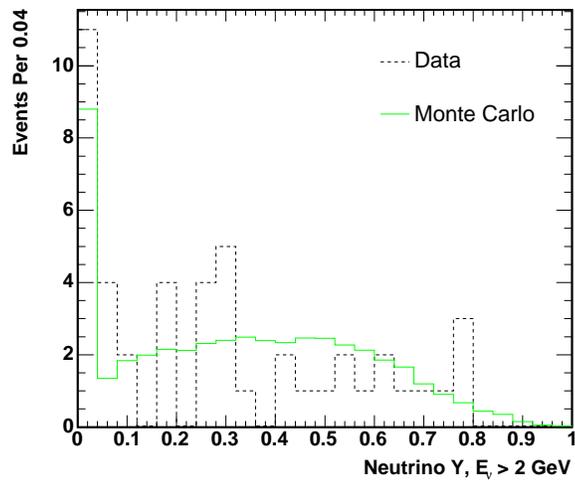
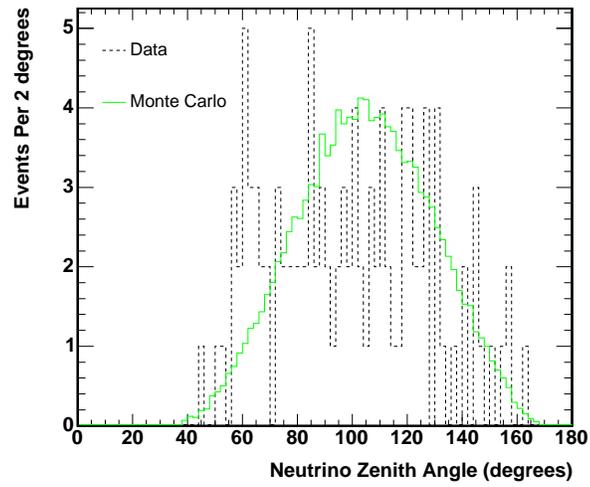
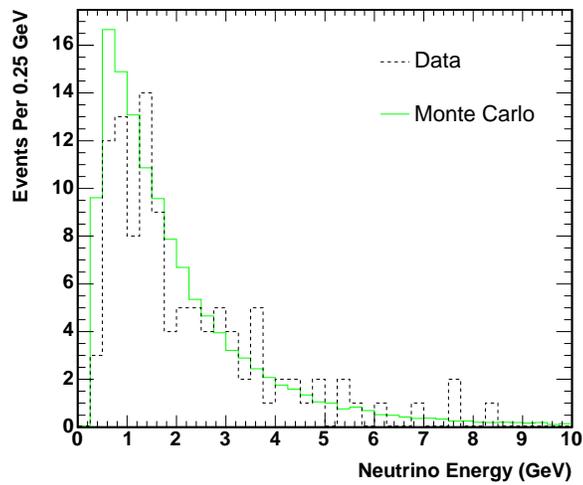
(a) Neutrino y for $E_\nu < 2$ GeV(b) Neutrino y for $E_\nu > 2$ GeV

Figure 5.20: The neutrino y distribution for the final data set and MC. The MC is unoscillated, includes all backgrounds described in Table 5.2, and represents the Bartol flux normalized to the data exposure.



(a) Neutrino zenith angle



(b) Neutrino energy

Figure 5.21: The neutrino zenith angle and energy for the final data set and MC. The MC is unoscillated, includes all backgrounds described in Table 5.2, and represents the Bartol flux normalized to the data exposure.

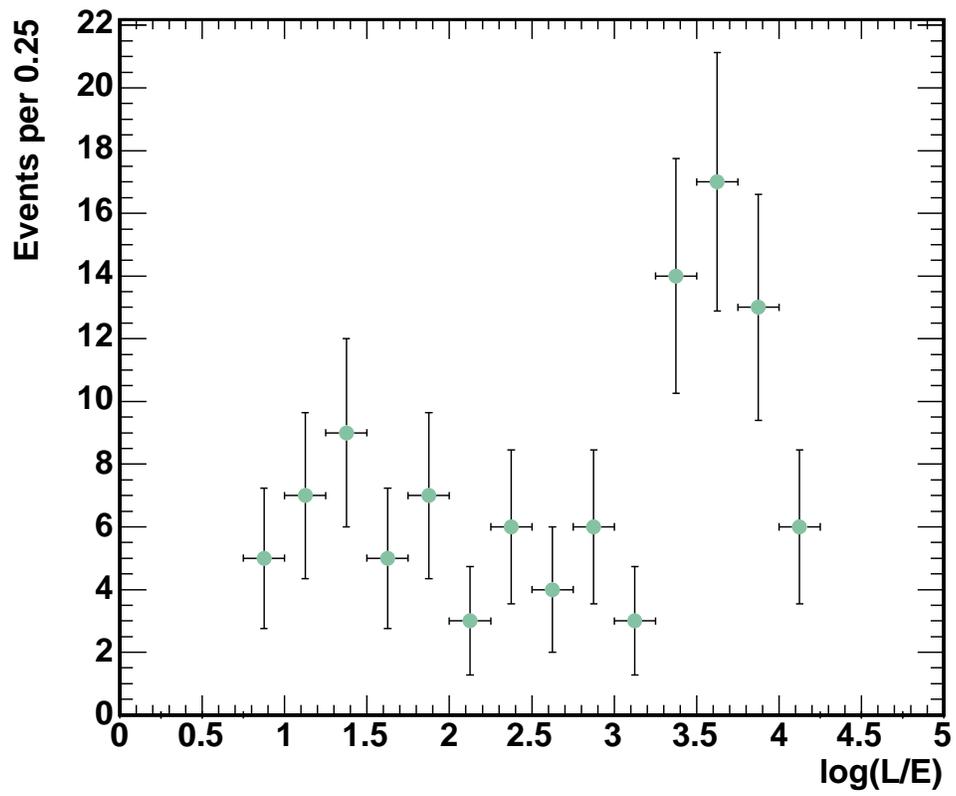


Figure 5.22: The distribution of $\log(L/E)$ for the final data set.

Chapter 6

Oscillation Fit

With the event selection complete, the compatibility of the selected data and the hypothesis of neutrino oscillation must be evaluated. To accomplish this task, some distribution which is sensitive to oscillations is considered. As described at the end of Chapter 5 and in Section 3.7, and as used in many other experiments, the $\log(L/E)$ distribution is an obvious choice.

Broadly speaking, neutrino oscillations will produce an overall deficit of events compared with expectations. Ignoring detector resolution, one should expect a pronounced deficit around the first oscillation maximum, and a roughly 50% decrease in event yield in the high $\log(L/E)$ region, where oscillations are rapid. Lastly, in the very low $\log(L/E)$ region, little difference from the no oscillation prediction should be observed, because the oscillation probability is small. Resolution tends to “soften” such features, as L and E are not reconstructed perfectly. Figure 6.1 illustrates these features given the MINOS detector resolution, along with the oscillation probability itself as a function of $\log(L/E)$.

This thesis employs a maximum likelihood method to calculate what oscillation parameters the data suggest. The data $\log(L/E)$ distribution is compared with the

$\log(L/E)$ distribution of the MC under various oscillation scenarios, and the most probable scenario is selected.

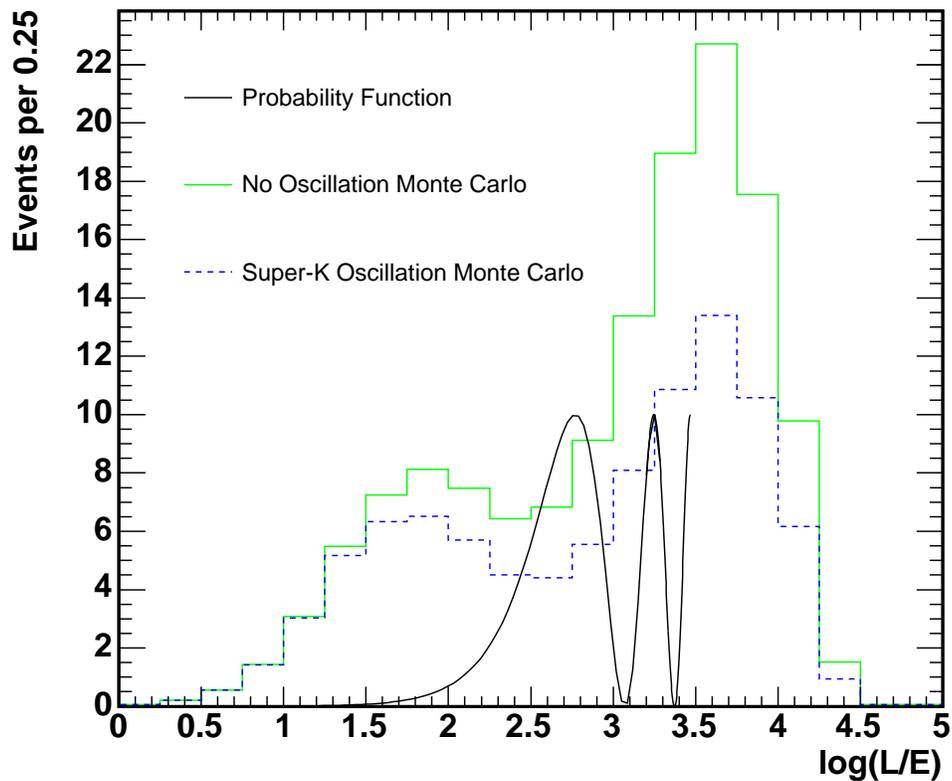


Figure 6.1: The reconstructed $\log(L/E)$ distribution for unoscillated MC (green) and MC oscillated with Super-Kamiokande’s best fit oscillation parameters (blue). Also shown in gray is the oscillation probability (with Super-K’s best fit parameters) as a function of true $\log(L/E)$, multiplied by 10 so as to be visible on the vertical scale. The probability function has been truncated at $\log(L/E) \approx 3.5$ because, around and above this value, the oscillations are too rapid to resolve.

Section 6.1 lays the mathematical foundation for the fitting procedure, Section 6.2 describes how the background is handled in the fitting, Section 6.3 details the actual results of the fitting, and Section 6.4 discusses how well the “best-fit” results describe the data.

6.1 Bayesian Interlude

Bayes originally formulated the basic principles that Laplace later cast into what is now known as Bayes' Theorem [73]:

$$P(X | Y, I) = \frac{P(Y | X, I) \times P(X | I)}{P(Y | I)} \quad (6.1)$$

where X and Y are random events or propositions, I is all relevant background information and initial conditions, and

- $P(X | Y, I)$ is the posterior probability
- $P(Y | X, I)$ is the likelihood
- $P(X | I)$ is the prior probability
- $P(Y | I)$ is the evidence

The utility of this theorem is immediately evident “hypothesis” and “data” are recognized as possible meanings of X and Y .

$$P(\text{hypothesis} | \text{data}, I) = \frac{P(\text{data} | \text{hypothesis}, I) \times P(\text{hypothesis} | I)}{P(\text{data} | I)} \quad (6.2)$$

Bayes Theorem provides the connection between what is sought, the probability that the hypothesis is true, to something that can easily be calculated, the probability of observing the data assuming that the hypothesis is true [73].

The evidence $P(\text{data} | I)$ serves as a normalization, and in analyses concerned with parameter estimation, it can be ignored. Further, the prior probability, which encompasses previous knowledge about the hypothesis being tested, is often taken to be flat in whatever parameters are being estimated. In principle, the prior can take any form, and its variations in the fit parameters would reflect how likely or unlikely the hypothesis is thought to be. For this oscillation analysis, one can think of the prior as identically

zero outside the physical region and flat within the physical region. Really, these details are academic because, as described later, only logarithmic likelihood differences (i.e. the probability relative to the most probable set of parameters) are used, and so the prior and evidence both subtract off. Thus, finding the maximum value of the posterior probability is entirely equivalent to maximizing the likelihood.

In the atmospheric arena, the likelihood of the i^{th} neutrino event is

$$\mathcal{L}_i(\vec{\theta} | E_i^R \pm \sigma_i^E, L_i^R \pm \sigma_i^L) = \mathcal{P}(\vec{\theta} | E_i^R \pm \sigma_i^E, L_i^R \pm \sigma_i^L) + \mathcal{L}^{bg} \quad (6.3)$$

where $\vec{\theta}$ denotes the oscillation parameters ($\vec{\theta} = \Delta m_{23}^2, \sin^2(2\theta_{23})$), $E_i^R \pm \sigma_i^E$ is the reconstructed energy, and $L_i^R \pm \sigma_i^L$ is the reconstructed length of flight. Contained in the first term are the flux, cross-section, detector acceptance, and oscillation survival probability. Specifically,

$$\mathcal{P}(\vec{\theta} | E_i^R \pm \sigma_i^E, L_i^R \pm \sigma_i^L) = \frac{P(\vec{\theta} | E_i^R \pm \sigma_i^E, L_i^R \pm \sigma_i^L)}{\iint P(\vec{\theta} | E_i^R \pm \sigma_i^E, L_i^R \pm \sigma_i^L) dE^R dL^R} \quad (6.4)$$

where

$$P(\vec{\theta} | E_i^R \pm \sigma_i^E, L_i^R \pm \sigma_i^L) = \int W(E_i^T, L_i^T) \times \eta(E_i^T, L_i^T) \times \frac{1}{N} \frac{dN(E^T)}{dE^T} \times G_E(E_i^R - E_i^T, \sigma_i^E) \times G_L(L_i^R - L_i^T, \sigma_i^L) \times dE^T \times dL^T \quad (6.5)$$

Here, W is the neutrino survival probability, η is the detector acceptance, $\frac{1}{N} \frac{dN(E^T)}{dE^T}$ is the normalized flux times cross-section, and $G_{E,L}$ are the resolutions in E and L .

Rather than attempt to parametrize all terms in the probability for observing the i^{th} neutrino event, physicists usually circumnavigate the problem by allowing the MC to numerically handle these various factors.

Algorithmically, the first step is to generate the expected $\log(L/E)$ spectrum given the acceptance, flux, and cross-section. The $(\sin^2(2\theta_{23}), \Delta m_{23}^2)$ parameter space is then divided into a grid and, at each grid point, the nominal $\log(L/E)$ spectrum is oscillated

according to the survival probability. Once the MC spectrum is oscillated, the background $\log(L/E)$ distribution is added on, which yields the full expected spectrum. The likelihood of the data distribution given the oscillated distribution is then calculated. The likelihood is defined as

$$\mathcal{L} = \prod_i^{N_{bins}} \left(\frac{N_i^{mc} N_i^d e^{-N_i^{mc}}}{N_i^d!} \right) \times \exp\left(\frac{(\alpha_B - \alpha_N)^2}{2\sigma_\alpha^2}\right) \quad (6.6)$$

where N_{bins} is the number of bins in $\log(L/E)$, N_i^{mc} is the number of Monte Carlo events in the i th $\log(L/E)$ bin, N_i^d is the number of data events in the i th $\log(L/E)$ bin, α_B is the normalization for the best fit, α_N is the nominal normalization, and σ_α is the uncertainty on the normalization, taken to be 15%.

Customarily, the negative (natural) log likelihood is used as it is easier to work with and is a monotonically increasing function of the likelihood.

$$-\ln \mathcal{L} = \sum_i^{N_{bins}} \left(N_i^{mc} + \ln(N_i^d!) - N_i^d \ln(N_i^{mc}) \right) + \frac{(\alpha_B - \alpha_N)^2}{2\sigma_\alpha^2} \quad (6.7)$$

The sum in Equation 6.7 is simply related to the Poisson probability of the data distribution given the oscillated MC distribution. The second term is a penalty for the optimal renormalization of the MC spectrum. At each grid point, the normalization is allowed to float from 75% to 125%.

6.2 Background

The dominant background in the data is showering neutrino interactions. To better quantify this background, showering events from the atmospheric neutrino MC sample are subjected to the same event selection described in Chapter 4. Here, “showering” includes both neutral current events for all neutrino flavors and charged current ν_e and ν_τ (which are assumed to not oscillate) events, as described in Section 5.3. The latter produce showers which, in the MINOS detector, are often flavor ambiguous and difficult to distinguish from neutral current showers.

The showering neutrino events which pass the event selection are used to generate the showering background $\log(L/E)$ distribution. This background is then scaled the same way as the Monte Carlo signal distribution, as they both originate from the same flux (Barr solar max [17]). The showering background zenith angle, energy and $\log(L/E)$ distributions are shown in Figures 6.2, 6.3, and 6.4. The integral of all three of these distributions is 5.7 events for the nominal flux prediction and 6.5 events for the best fit flux value.

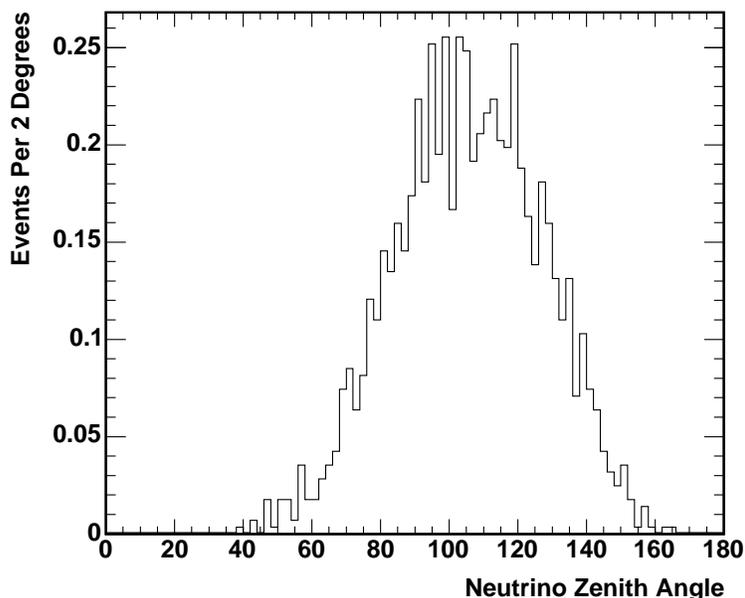


Figure 6.2: Showering neutrino background zenith angle distribution (from MC).

The second largest background component is still cosmic ray muons. The cosmic ray background is modeled by collecting events which pass the neural net cut but are tagged by the veto shield. These events are used to generate the cosmic ray background kinematic distributions, shown in Figures 6.5, 6.6, and 6.7. The integral of each of these distributions is 3.2 events.

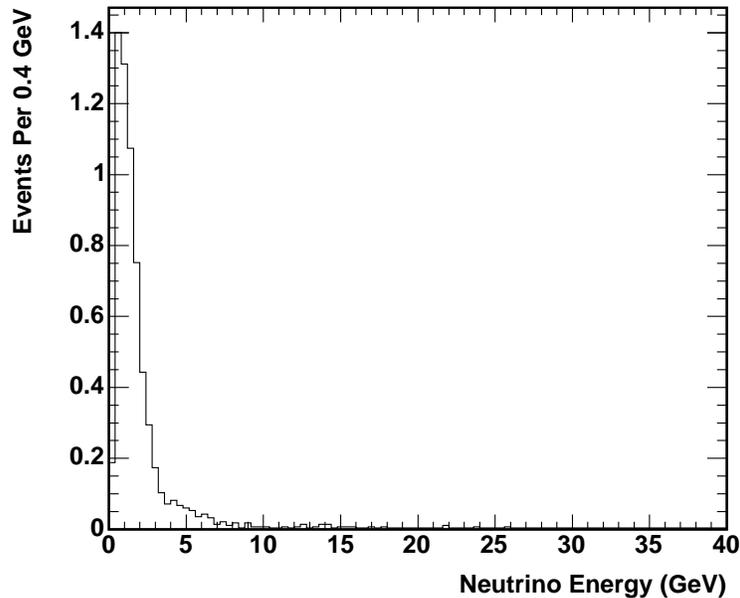


Figure 6.3: Showering background E distribution (from MC).

6.3 Fit Results

The data, unoscillated, and best fit oscillation $\log(L/E)$ distributions are shown in Figure 6.8. Some features in Figure 6.8 further explanation. First, the overall shape of this $\log(L/E)$ distribution may look somewhat odd to those familiar with the atmospheric neutrino oscillation world. Ordinarily, the heights of the two “humps” in the $\log(L/E)$ distribution are similar, with the hump at higher $\log(L/E)$ shorter due to the rapid oscillations at those $\log(L/E)$ values. For the current analysis, acceptance for downward going events, which have small L values, has been deliberately sacrificed to maintain acceptance for upward going events, as described in Section 5.2. As the value of L drives the value of $\log(L/E)$, the $\log(L/E)$ distribution has fewer events in the low $\log(L/E)$ region.

Second, in the low $\log(L/E)$ region, the data appear systematically higher than the

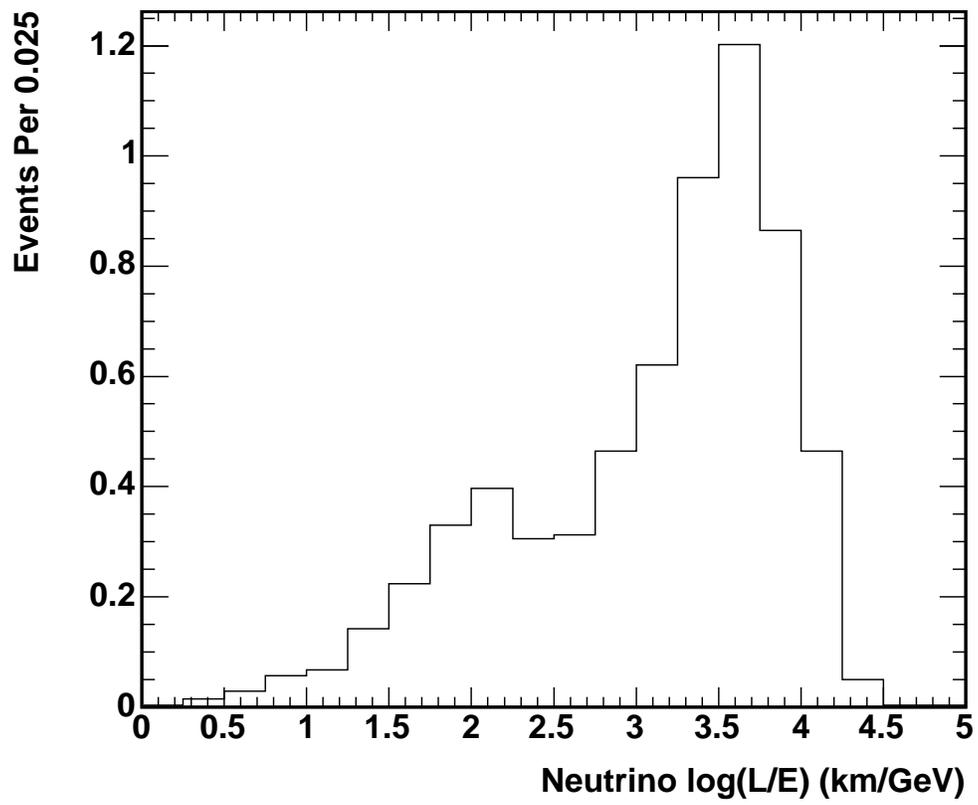


Figure 6.4: Showering background $\log(L/E)$ distribution (from MC).

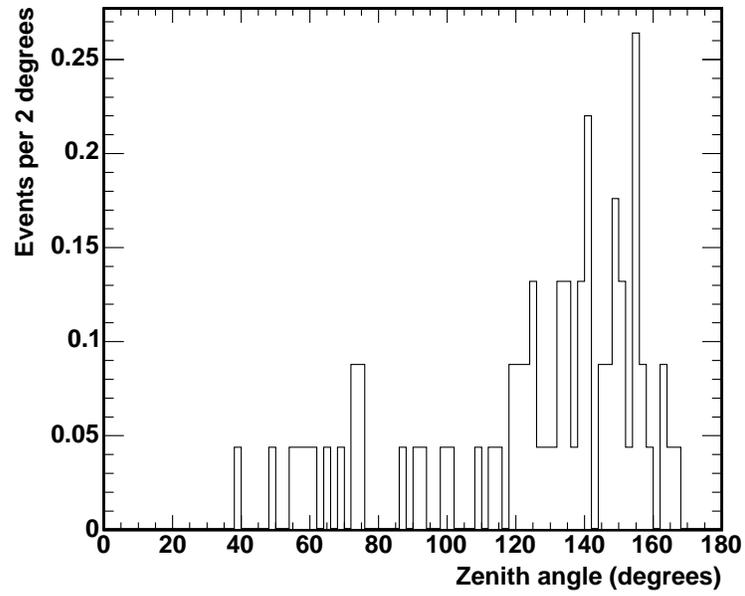


Figure 6.5: Total cosmic ray background zenith angle distribution (from data).

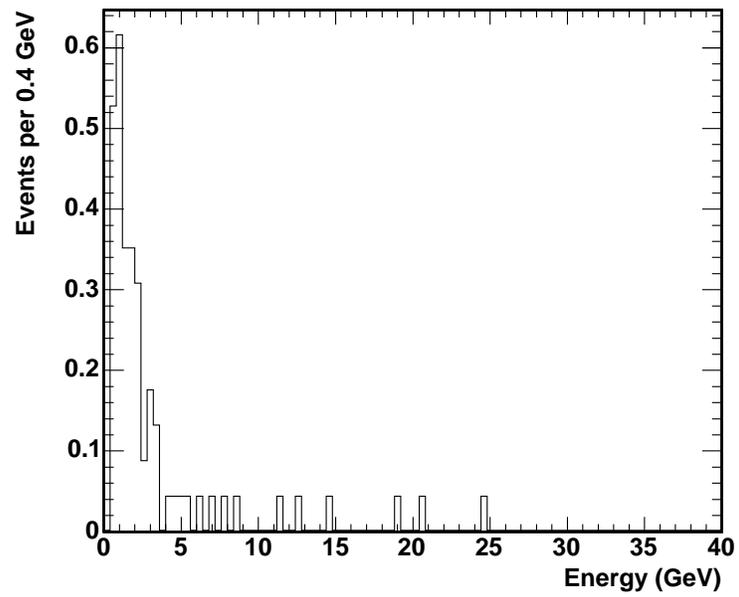


Figure 6.6: Total cosmic ray background E distribution (from data).

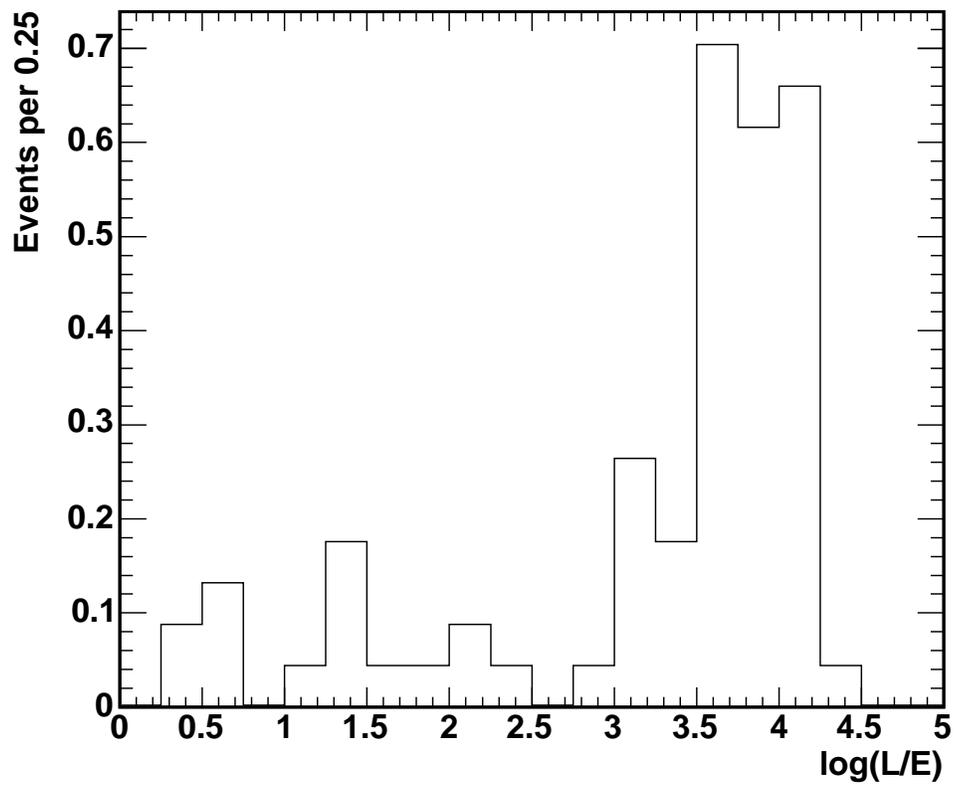


Figure 6.7: Cosmic ray background $\log(L/E)$ distribution (from data).

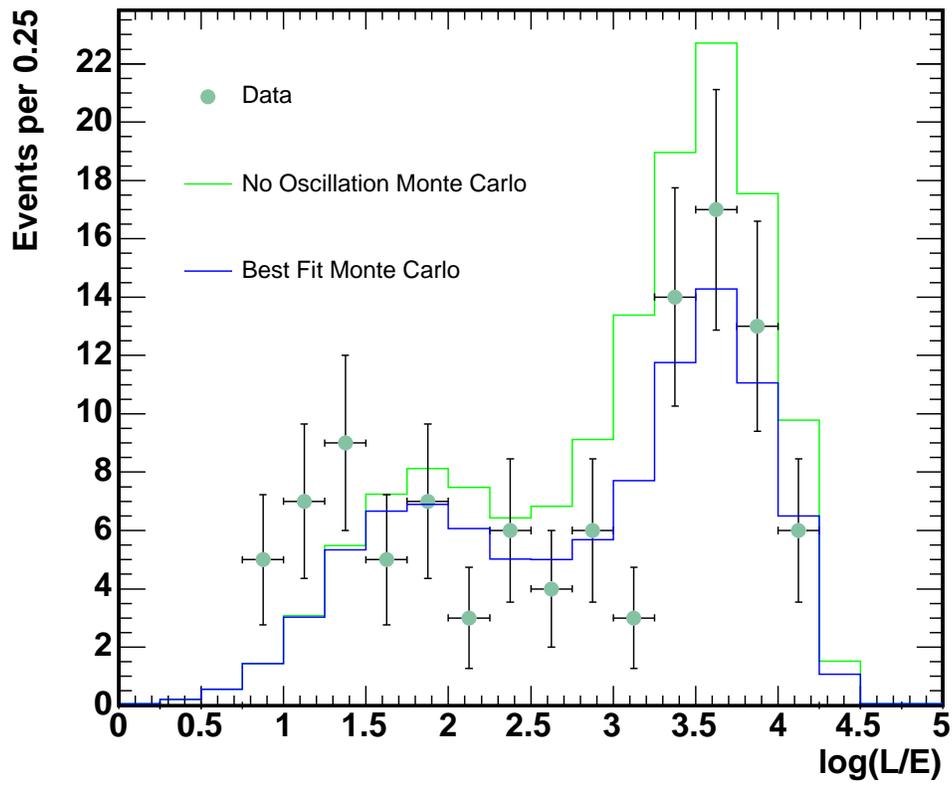


Figure 6.8: The final $\log(L/E)$ distribution. Unoscillated MC is green, data are points with error bars, and best fit MC is blue. The MC includes all backgrounds described in Table 5.2.

MC prediction. With such a small number of expected events, statistical fluctuation is not out of the realm of possibility. A pessimist would immediately cry “background!” about these events. Either way, they warrant a closer look, to which Appendix A is dedicated. Suffice it to say here that no distribution indicates these events to be cosmic ray (or any other) background.

The likelihood results are shown in Figures 6.9 and 6.10. In Figure 6.9, the peak at $\sin^2(2\theta_{23}) = 1$ and $\log(\Delta m_{23}^2) = -1$ indicates that this region of the oscillation grid is very strongly excluded. Consider that, if one calculates what $\log(L/E)$ maximizes the oscillation probability under the assumption that $\Delta m_{23}^2 \approx 0.1$, one finds that $\log(L/E) \approx 1$, precisely the $\log(L/E)$ region in which there appears to be an excess of events, as described in the previous paragraph.

The one dimensional maximum likelihood fit, in $\Delta m_{23}^2 m^2$ space only, has also been executed. In this fit, Δm_{23}^2 is fixed at some value and the $(\sin^2(2\theta_{23}), \text{normalization})$ parameter space is searched for the lowest $-\ln(L)$ value, which is then taken as the value of the likelihood for the value Δm_{23}^2 in question. A new value of Δm_{23}^2 is selected, and the process is repeated. Figure 6.11 shows this likelihood contour. This distribution gives the asymmetric statistical error on Δm_{23}^2 ; the two points away from the minimum where the likelihood rise first reaches 0.5 are the 68% limits, analogous to $\pm 1\sigma$ in the Gaussian distribution.

An analogous one dimensional fit in $\sin^2(2\theta_{23})$ space has also been performed, and this fit yields a one-sided 68% limit on the best fit value. The result is shown in Figure 6.12.

The best fit point resides at $\Delta m_{23}^2 = 0.93_{-0.44}^{+3.94} \times 10^{-3} \text{ eV}^2$ and $\sin^2 2\theta_{23} = 0.95_{-0.32}$.

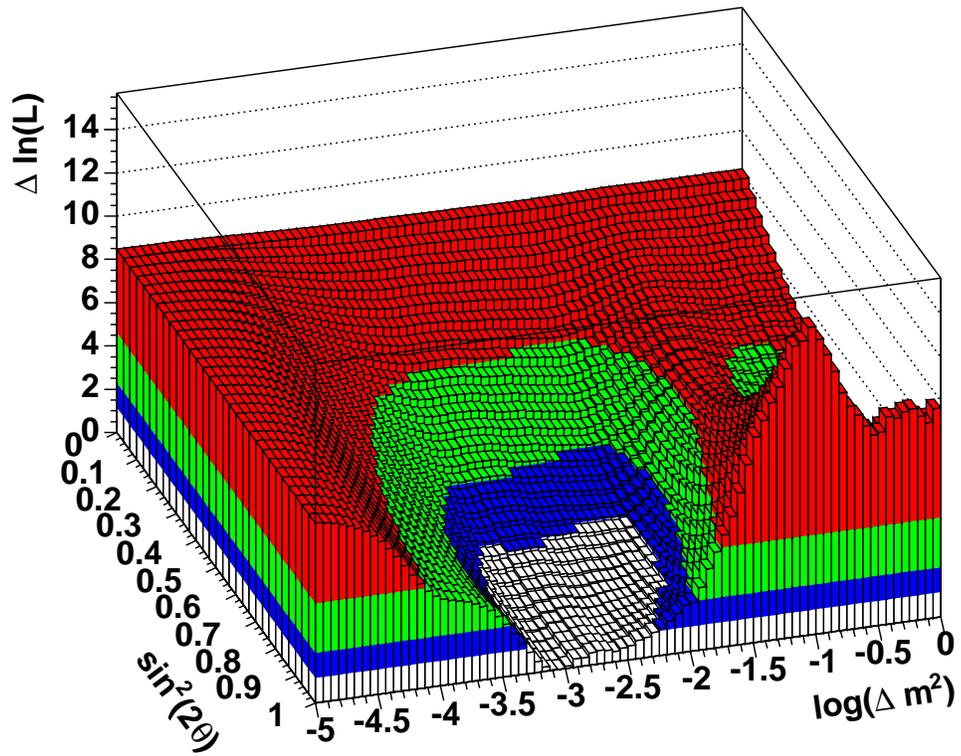


Figure 6.9: The final, minimum-subtracted, likelihood surface. The levels are 68% (blue), 90% (green), and 99% (red) confidence.

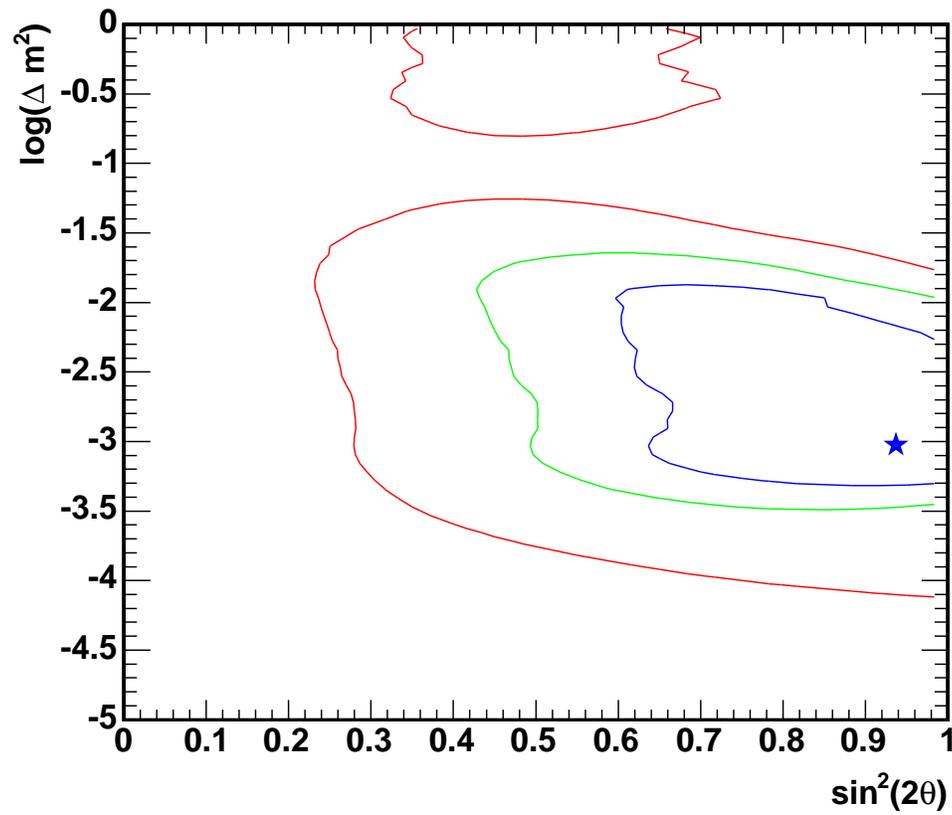


Figure 6.10: The final, minimum-subtracted, likelihood contour. The levels are, moving out from the innermost contour, 68% (blue), 90% (green), and 99% (red) confidence. The best fit point is shown as a blue star.

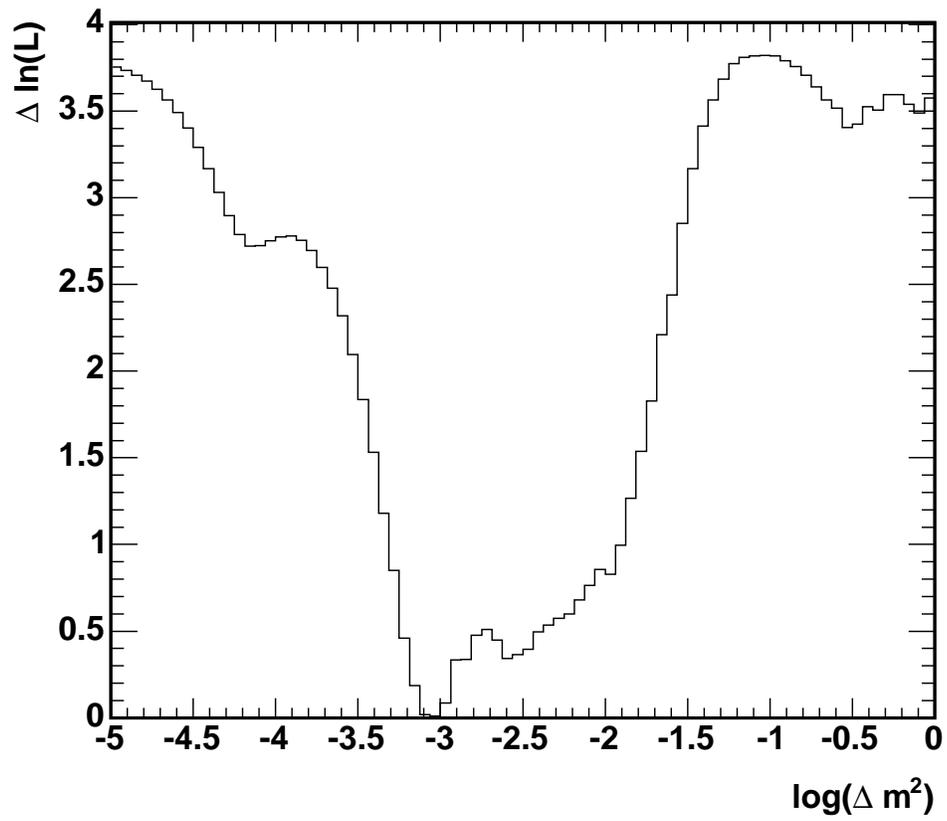


Figure 6.11: The one dimensional, minimum subtracted, likelihood contour as a function of $\log(\Delta m_{23}^2)$.

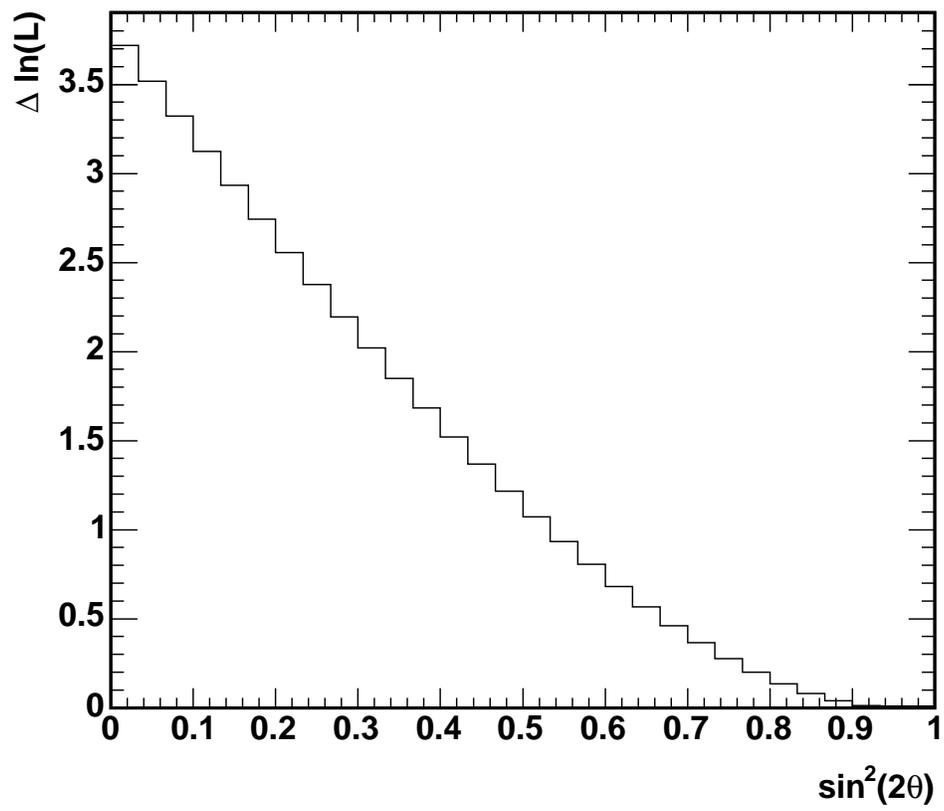


Figure 6.12: The one dimensional, minimum subtracted, likelihood contour as a function of $\sin^2(2\theta_{23})$.

6.3.1 Sensitivity

The sensitivity contour is shown in Figure 6.13. The sensitivity is obtained by fitting the best fit oscillated Monte Carlo $\log(L/E)$ distribution as though it were real data, against the same oscillation grid used to fit the data $\log(L/E)$ distribution. The sensitivity is one measure of the statistical significance of the results. The “average” $\log(L/E)$ spectrum from an ensemble of experiments should converge toward the best fit $\log(L/E)$ distribution and so, if the sensitivity contour appears significantly different from the data likelihood contour, the data $\log(L/E)$ distribution is the result of a large fluctuation from the average, at least in the eyes of the MC.

The sensitivity surface is much the same as the data likelihood surface, except in the region $\log(\Delta m_{23}^2) > -1.5$ across all values of $\sin^2(2\theta_{23})$. Again, in the data, the higher than predicted number of events in the low $\log(L/E)$ region drives this band in parameter space to be more strongly excluded. Since the excess isn't present in the best fit MC $\log(L/E)$ spectrum, the likelihood difference is consequently smaller, and the 90% and 99% confidence limits fail to close on the high Δm^2 side in the sensitivity as they do in the likelihood. Features, such as the saddle point around $\sin^2(2\theta_{23}) = 0.5$, $\log(\Delta m_{23}^2) = -1$, are still present.

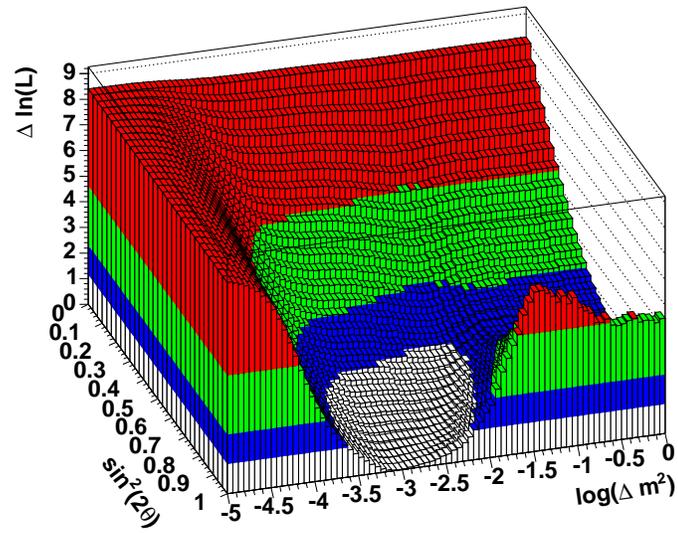
Figure 6.14 shows the 68% contours from the likelihood and the sensitivity. The sensitivity 68% contour is 35% larger than the data likelihood contour.

6.4 Goodness of Fit

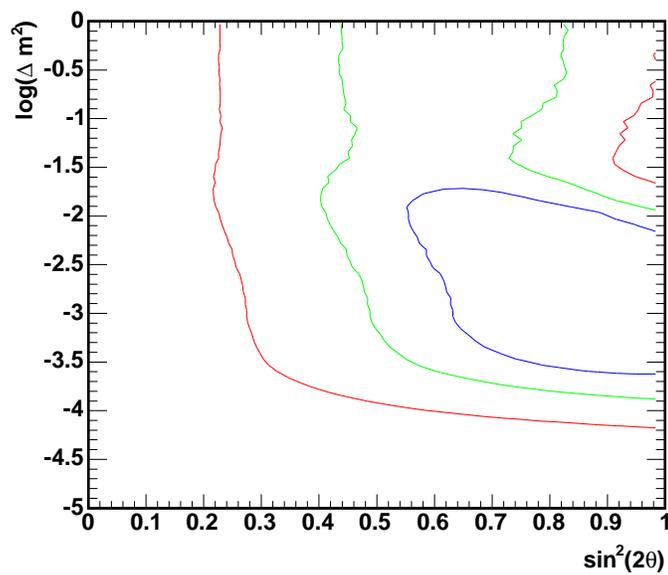
The modified χ^2 of the fit has been calculated,

$$\chi^2 = \left(\sum_{i=1}^n 2(N_{exp} - N_{obs}) + 2N_{obs} \log \left(\frac{N_{obs}}{N_{exp}} \right) \right) + \frac{(\alpha_N - \alpha_{best})^2}{(2\sigma_\alpha^2)} \quad (6.8)$$

For the best fit parameters, $\chi^2 = 23.3$ per 17 DOF.



(a) Sensitivity Surface



(b) Sensitivity Contour

Figure 6.13: The 68% (blue), 90% (green), and 99% (red) sensitivity.

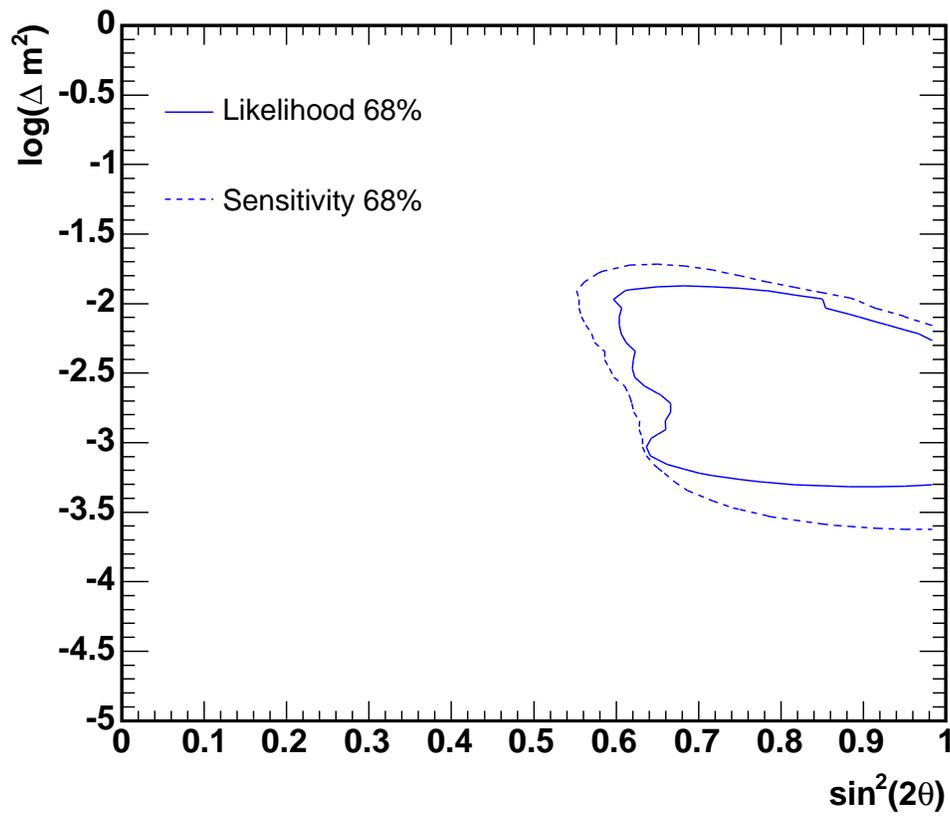


Figure 6.14: The 68% likelihood and sensitivity contours. The sensitivity contour is 35% larger than the likelihood contour.

With the initial fitting done, how well the best fit actually describes the data must be evaluated. To this end, a series of MC experiments has been carried out with the same fitting procedure used to analyze the real data. The various χ^2 values are then collected.

First, approximately half of the atmospheric neutrino MC sample is designated as “fake data”, events which will be used to generate the data $\log(L/E)$ distributions in the ensemble of fake experiments. The other half of the Monte Carlo sample is used to generate the grid of $\log(L/E)$ distributions against which the data distributions are fit, as described in Section 6.1.

To simulate an experiment, the number of desired neutrino events must be assigned, based on the number of neutrino events observed in the real data analysis. In this case, 105 total events, less 3 expected cosmic ray background events, gives 102 real neutrino data events (see Section 5.3). For each MC experiment, the target number of events is selected from a gaussian distribution centered on 102 and having a width equal to the atmospheric neutrino flux uncertainty, 15%. Thus the number of neutrino events in the MC experiments ranges roughly from 87 to 118. Once the number of neutrino events is set for a given experiment, that many events (both ν_μ charged current and showering events) are selected from the “fake data” MC ensemble to fill the experiment’s neutrino $\log(L/E)$ distribution.

Next, the size of the cosmic ray background, acquired from data, is varied in a similar manner to the number of neutrino events. Here though, the relevant width is the statistical error on the number cosmic ray background events, which is 12%. The rescaled cosmic ray background $\log(L/E)$ distribution is added to neutrino $\log(L/E)$ distribution to complete the experiment’s $\log(L/E)$ distribution. We then fit it in exactly the same manner that we fit the real data $\log(L/E)$ distribution.

Finally, this process is repeated for numerous experiments. Given that around 102

events are desired in each experiment, there is enough MC for 500 independent experiments. This ensemble yields $\chi^2_{avg} = 15.7 \pm 6.0$ (RMS). The distribution of χ^2 values from these 500 MC experiments is shown in Figure 6.15. The data fit χ^2 is 1.3σ from the average value of 15.7, which means the data χ^2 has approximately 20% coverage, a very satisfactory and reassuring result.

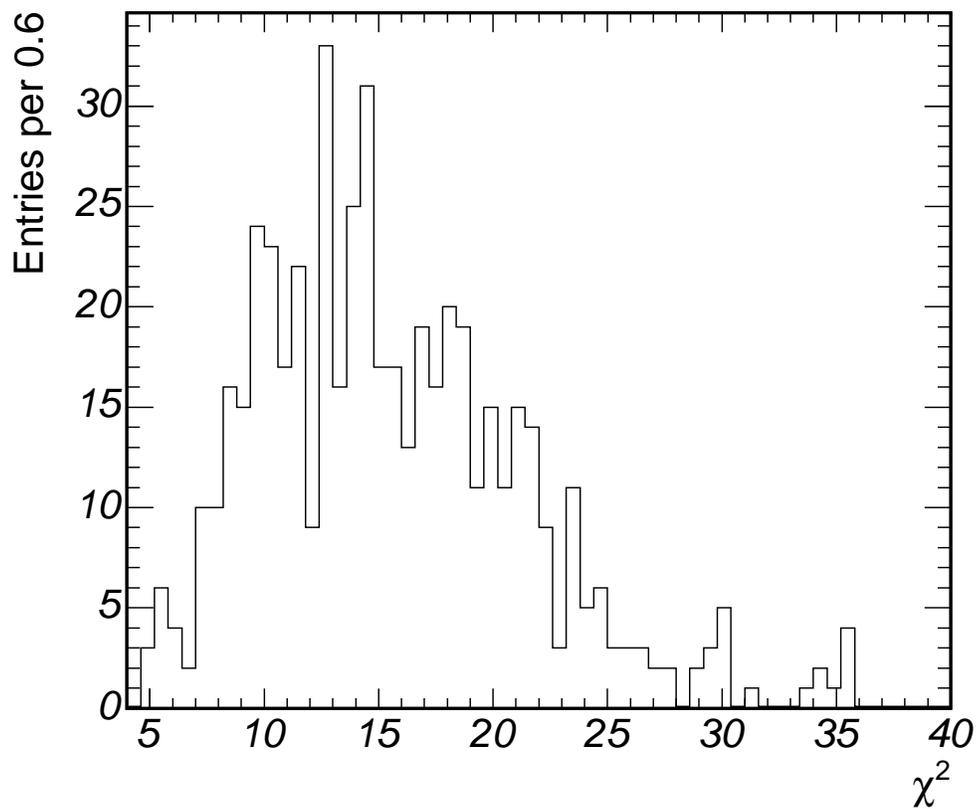


Figure 6.15: The χ^2 distribution for an ensemble of 500 MC experiments.

6.5 The Null Oscillation Scenario

Most of this chapter addresses a two-part question: “What oscillation scenario do the data suggest, and how well does that scenario describe the data?” An equally important question is “How incompatible with the idea of null oscillations are the data?” This subject has been investigated in a manner similar to the Goodness of Fit test, described above (Section 6.4). An ensemble of 500 experiments, each with 102 unoscillated neutrino events, was generated. These $\log(L/E)$ distributions were then fit against the null oscillation $\log(L/E)$ distribution, with only the normalization allowed to vary, and the χ^2 values of those fits were calculated.

Figure 6.16 shows the distribution of these 500 χ^2 values. This ensemble yields $\chi^2_{avg} = 18.0 \pm 6.4$ (RMS). The real data yields $\chi^2 = 31.5$ when fit against the null oscillation hypothesis with varying normalization (at 89% of the nominal flux). The data χ^2 is 2.1σ away from the mean. In other words, less than 3.5% of the time would an experiment observing neutrinos that do not oscillate acquire a data set which yields a χ^2 that is as incompatible with the null oscillation hypothesis as is the current real data set.

It should be noted that this result excludes null oscillations about 3% less strongly than does the likelihood surface (Figure 6.9) obtained from the standard fit. This difference, while not fully understood, likely arises from the same idiosyncrasies in the data that drive the difference between the sensitivity and the likelihood contour, as discussed in Appendix A.

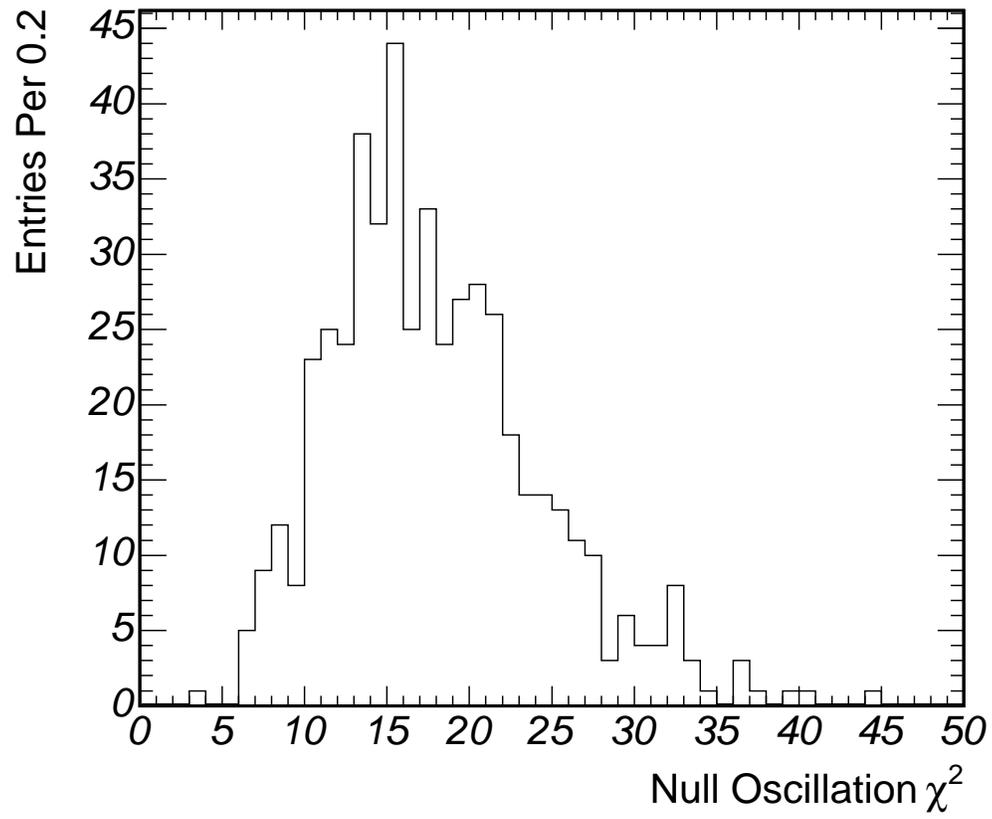


Figure 6.16: The χ^2 distribution for an ensemble of MC experiments with no oscillations applied.

Chapter 7

Systematic Error Analysis

A measurement is nothing without knowledge of how accurate it is. In this chapter, the error associated with the measurement of the oscillation parameters will be investigated.

Oscillation measurements on atmospheric neutrinos often are based on small samples, have numerous uncertainties, and consequently have broad, non-gaussian likelihood or χ^2 minima. As is evident from Figure 6.9, this measurement has a fairly flat likelihood surface around the minimum. In such cases, attempting to specify error in the usual vocabulary (*bestfit*) \pm (*error*) is inaccurate, as the notion of an error on a parameter ceases to have meaning when the likelihood surface is so flat in the region of the minimum. As such, the effect of errors on the likelihood surface itself will be shown in this Chapter.

Section 7.1 discusses the unavoidable uncertainties in the simulation of atmospheric neutrino production, Section 7.2 details the current uncertainty on the various neutrino cross-sections, Section 7.3 deals with the impact of the uncertainty on the size of the backgrounds, Section 7.4 demonstrates what bias the event selection itself introduces, Section 7.5 shows how systematic errors on the reconstructed muon momentum and hadronic energy affect the final results, and Section 7.6 collects all of these various

errors to demonstrate their overall effect. Finally, Section 7.7 addresses the important question of how fitting near a physical boundary influences the fit results.

7.1 Uncertainties in Neutrino Production

In chapter 4, the MC framework was described, but nothing was said of any inherent error in the MC study. The single largest error in the measurement of the atmospheric neutrino oscillation parameters is the absolute neutrino flux, carrying approximately 15% [18] uncertainty in the neutrino energy range of interest to most detectors, save perhaps ICECUBE. This value for the uncertainty is the culmination of a very involved evaluation of the impact of the various simulation uncertainties in generating the neutrino flux. The two least certain components of the simulation are uncertainties in the primary cosmic ray flux itself and in hadron production.

7.1.1 Primary Flux

Measurements of the composition and magnitude of the cosmic ray flux have been conducted using small detectors lofted with balloons or attached to spacecraft, the most statistically significant of which are AMS (actually mounted on the International Space Station) [4], CAPRICE [26], and BESS [72]. Unfortunately, these detectors are capable of measuring particle energy up to only roughly 200 GeV. To observe cosmic rays above this energy, similarly lofted but less accurate emulsion-calorimetry techniques must be employed, as done with RUNJOB [12] and JACEE [14]. Together, these experiments have measured the flux of primary cosmic rays over a large range of energy.

Gaisser, Stanev, Honda, and Lipari [41] have collected most of the available flux measurements and parametrized the primary fluxes. According to Barr, et al., this GSHL parametrization carries approximately 6% uncertainty in the energy range 1 to 100 GeV and, above 100 GeV, the uncertainty increases to around 30% at 1000 GeV

[18]. Figure 7.1 illustrates this uncertainty for a large range of primary kinetic energy for protons and helium, the two dominant components of the cosmic ray flux.

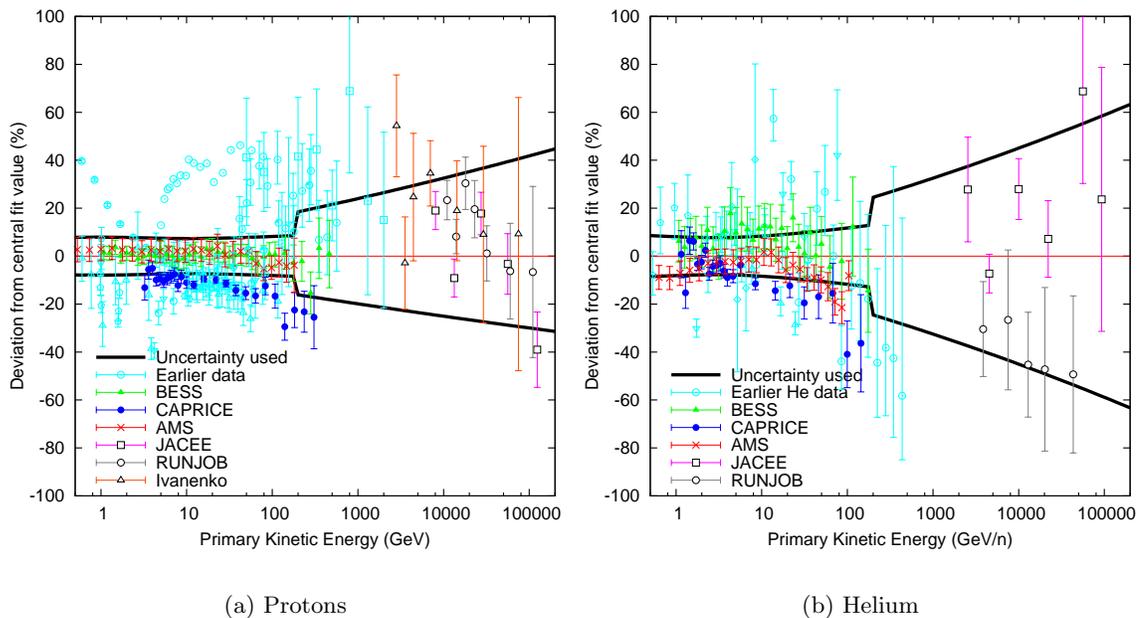


Figure 7.1: Comparison of flux measurements to the GSHL [18].

7.1.2 Hadron Production

The origin of atmospheric neutrinos lies in the collisions of primary cosmic rays with the earth's upper atmosphere (essentially N_2 and O_2). Obviously a thorough understanding of hadron production, particularly pion and kaon production, is crucial for an accurate prediction of the atmospheric neutrino flux.

Hadron production in nucleon-nucleus collisions is a complicated problem, both theoretically (delving into QCD) and experimentally. Early experimental exploration has been kinematically limited; many production experiments (for example, [10] and [11]) have probed only a small portion of scattering angle, incoming particle energy (because

of beam constraints), x_{lab} , etc. Currently, MIPP [38] and HARP [28] are running at Fermilab and CERN, respectively, and will use detectors with nearly 4π sr angular coverage to measure the full pion and kaon production cross sections, a critical bit of information for the overall normalization of the atmospheric neutrino flux [35]. Unfortunately, only preliminary results from HARP and no MIPP results are currently available.

Figure 7.2 shows the uncertainty in the production rates of pions and kaons used in the generation of the MC.

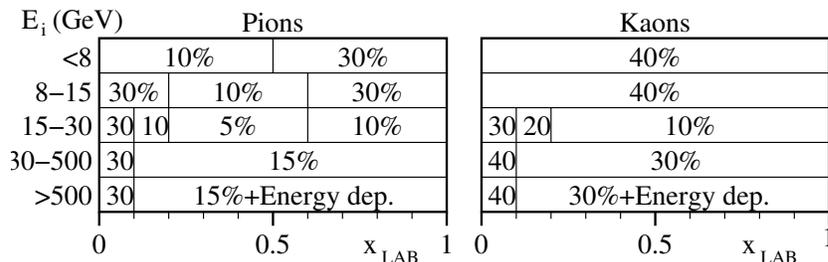


Figure 7.2: Pion and kaon production uncertainty as a function of x_{lab} for various energy ranges of incoming proton energy E_i [18].

7.2 Neutrino Cross-Section

Currently, the overall neutrino interaction cross-section is known to about 10% percent, though the actual uncertainty varies with energy. Deep inelastic scattering (*DIS*), quasi-elastic scattering (*QE*), and resonance production (*RES*) sum to produce the total cross-section, but the individual sizes are known less accurately.

The impact of the neutrino cross-section uncertainty has been estimated with a MC simulation. In this simulation, the uncertainties are

- $\sigma_{DIS} = \pm 5\%$
- $\sigma_{QE} = \pm 15\%$

- $\sigma_{RES} = \pm 15\%$

Since the total cross-section is known more accurately than any of its components, a coupled variation of the three separate cross-sections yields the most optimistic and, perhaps, most correct estimate of the effect of the cross-section uncertainty. In this analysis, a more pessimistic approach has been taken, and all three cross-sections have been varied simultaneously up or down. In even this most pessimistic method, the best fit point is not pulled out of the original 68% contour, as shown in Figures 7.3, 7.5, and 7.4.

7.3 Background Estimation

As mentioned in Chapter 5, the two primary backgrounds are real “showering” (all flavor neutral current and ν_e and ν_τ charged current) neutrino interactions and cosmic ray muons. The former is modeled with MC, and the latter is taken from real data as described in the previous chapter. As the neutral current background estimation is based off of MC, it has negligible statistical error. The cosmic ray background, on the other hand, has a fairly large statistical error, around 12%. As such, both backgrounds have been systematically varied by $\pm 12\%$. The result is shown in Figure 7.6, which demonstrates that the measurement is almost completely insensitive to the background effects.

7.4 Neural Net

The best cut on neural net output is found by optimizing the *FOM* shown in Equation 5.2. As Figure 5.14 shows, the relative mixture of signal and background varies quickly as a function of this cut. As such, any systematic bias this cut introduces must be accounted for when considering systematic errors. To this end, the effect of changing

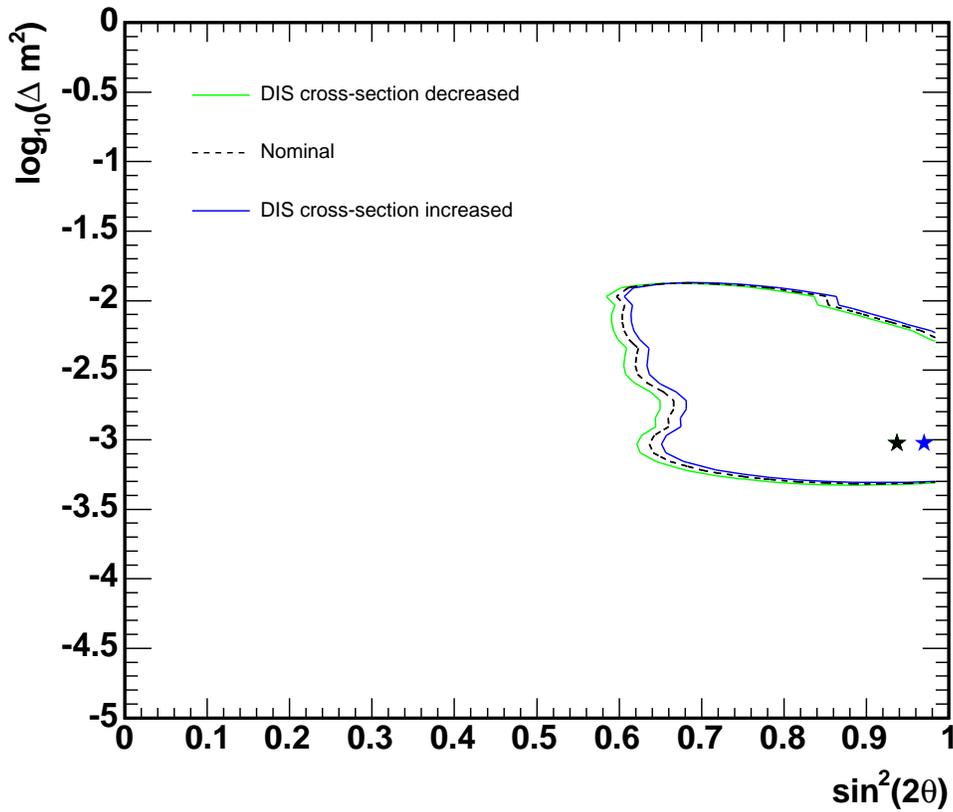


Figure 7.3: 68% contours for varying the neutrino *DIS* cross-section. Black (dashed) shows the original data fit, green (light solid) shows the results of reducing the *DIS* cross-sections by 5%, and blue (dark solid) shows the result of increasing the *DIS* cross-section by 5%. Stars show best-fit points. Note that varying the *DIS* cross-section did nothing to the best fit value of Δm_{23}^2 , and that lowering it did nothing to the best-fit point at all.

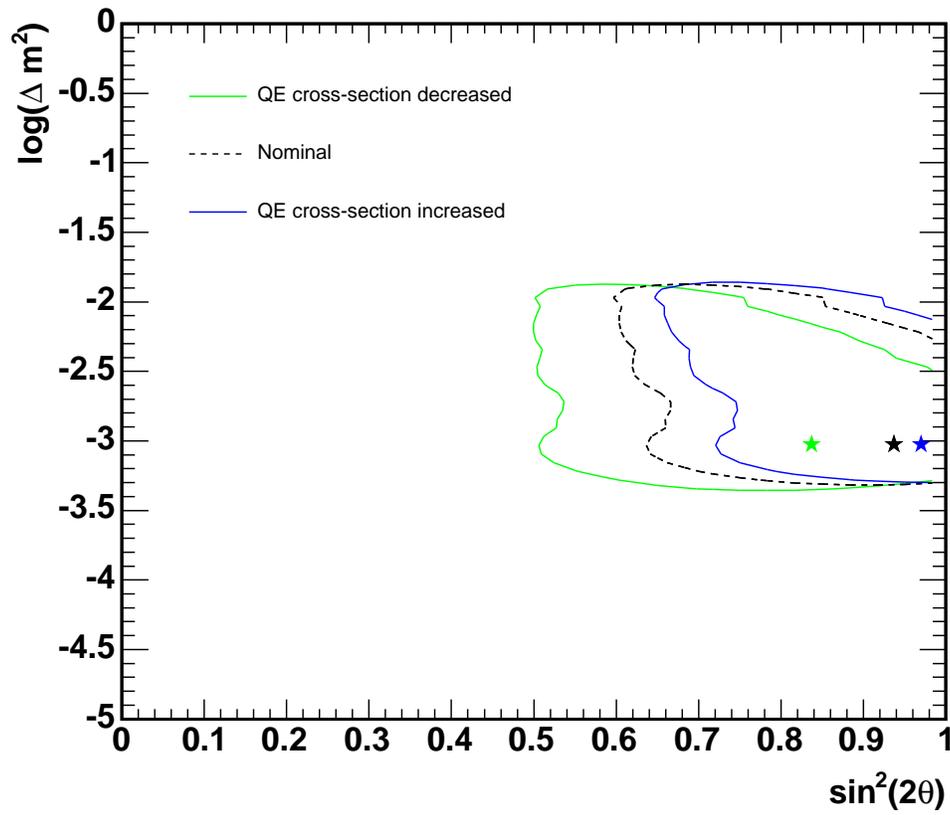


Figure 7.4: 68% contours for varying the neutrino QE cross-section. Black (dashed) shows the original data fit, green (light solid) shows the results of reducing the QE cross-sections by 15%, and blue (dark solid) shows the result of increasing the QE cross-section by 15%. Stars show best-fit points. Note that varying the QE cross-section did nothing to the best fit value of Δm_{23}^2 .

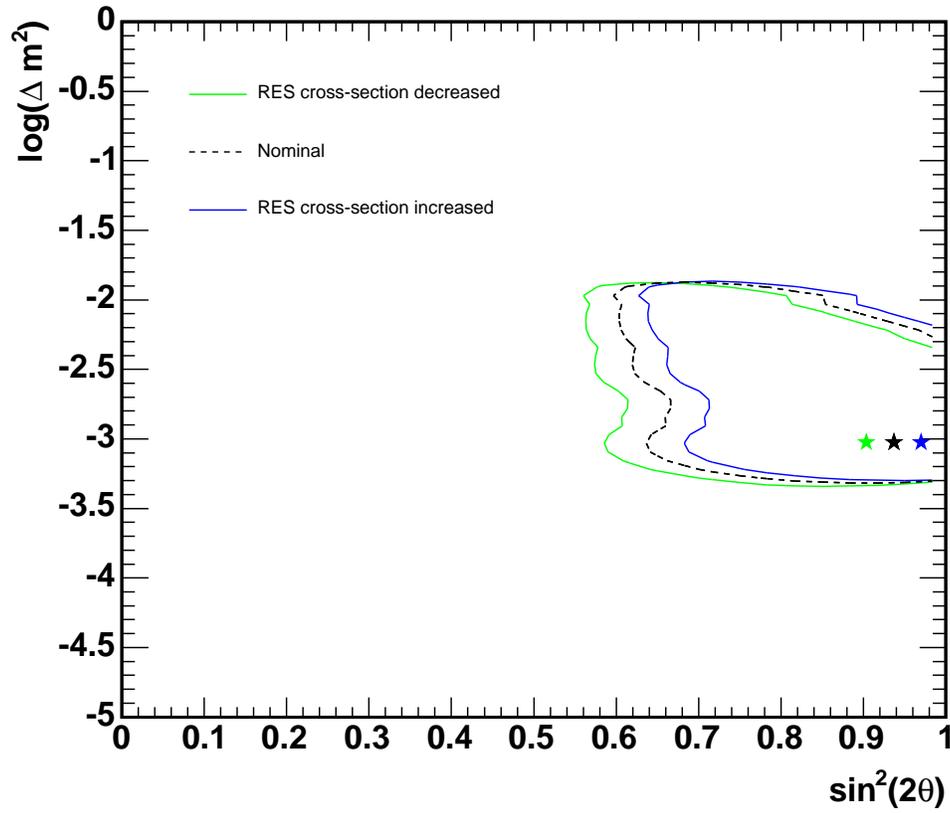


Figure 7.5: 68% contours for varying the neutrino *RES* cross-section. Black (dashed) shows the original data fit, green (light solid) shows the results of reducing the *RES* cross-sections by 15%, and blue (dark solid) shows the result of increasing the *RES* cross-section by 15%. Stars show best-fit points. Note that varying the *RES* cross-section did nothing to the best fit value of Δm_{23}^2 .

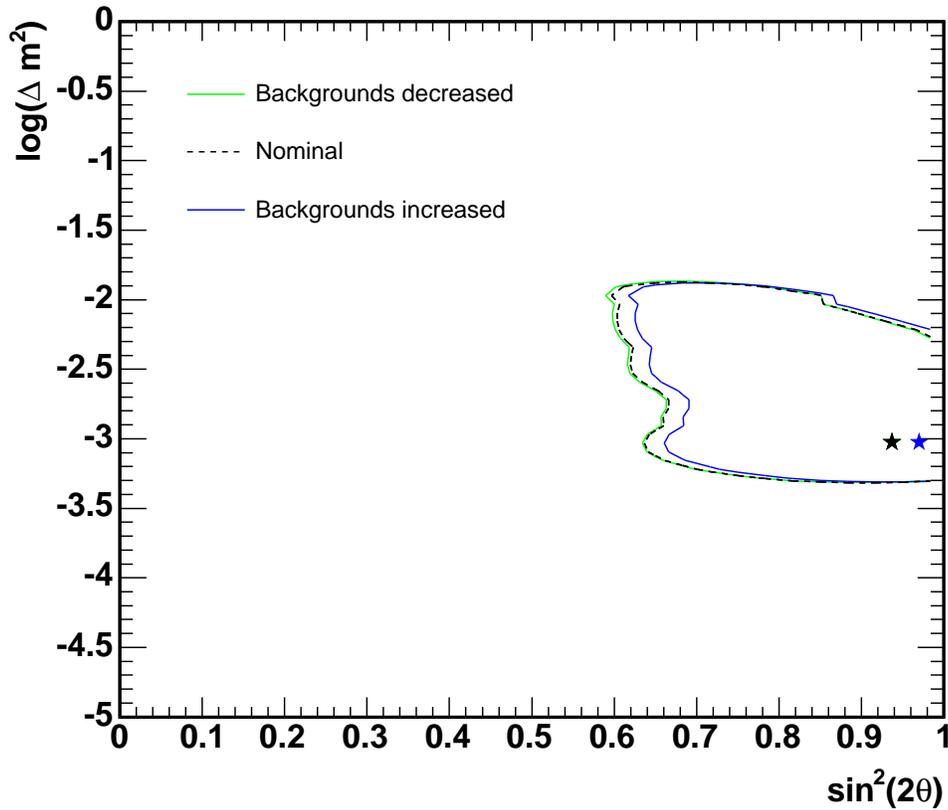


Figure 7.6: 68% contours for varying neutral current and cosmic ray backgrounds. Black (dashed) shows the original data fit, green (light solid) shows the result of reducing the backgrounds by their specified errors, and blue (dark solid) shows the result of increasing the backgrounds by their specified errors. Stars show best-fit points. Note that reducing the background amounts did nothing to the best fit point, and increasing the background amounts did nothing to the best fit value of Δm_{23}^2 .

the neural net cut by ± 0.01 , which changes the FOM by $\sim 10\%$, has been investigated.

The results are shown in Figure 7.7.

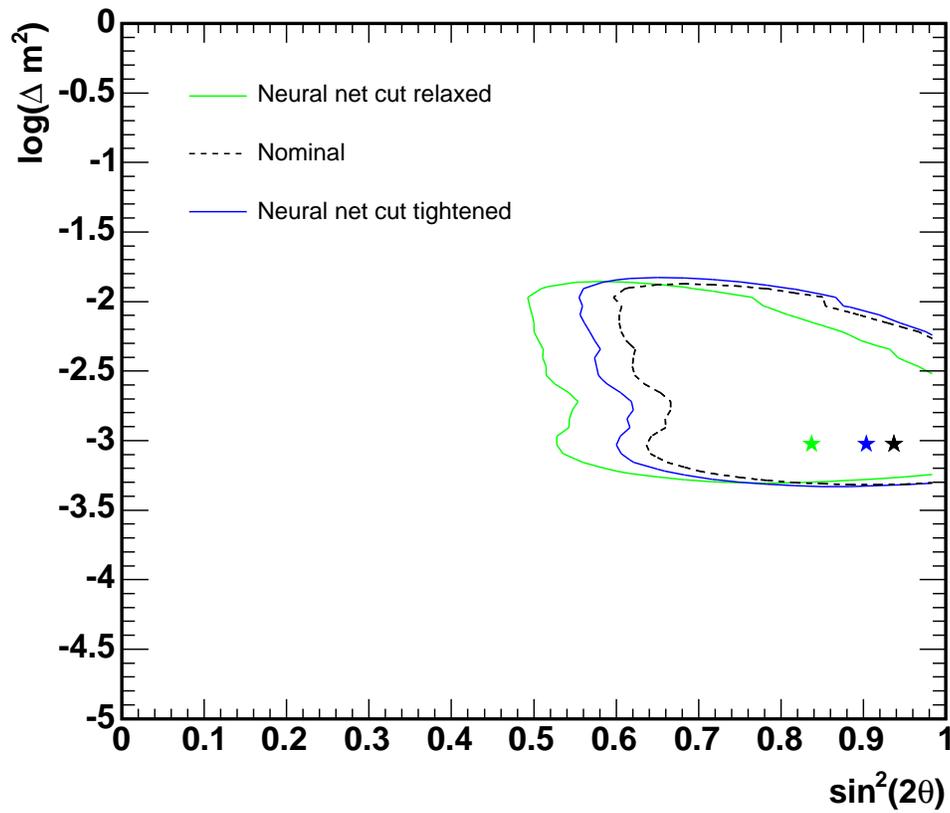


Figure 7.7: 68% contours for varying the neural net cut. Black (dashed) shows the original data fit, green (light solid) shows the result of reducing the neural net cut by 0.01, and blue (dark solid) shows the result of increasing the neural net cut by 0.01. Stars show best-fit points. Note that varying the neural net cut did nothing to the best fit value of Δm_{23}^2 .

7.5 Muon Momentum and Hadronic Energy Reconstruction

Reference [3] discusses possible systematic effects in the muon momentum reconstruction, and the largest effect is estimated to be $\pm 3\%$. Similarly, in [3], the largest systematic effect in hadronic energy reconstruction is estimated to be $\pm 5\%$. Figures 7.8 and 7.9, respectively, show the impact on the final measurement of these uncertainties.

7.6 Cumulative Systematic Effect

Perhaps the best estimation of the total systematic effect is the size of the “envelope” contour enclosing all the various 68% contours which result from the previously described error estimates. In simple words, this envelope is the outline formed by drawing all the contours from all the systematic variations on the same plot and tracing the outermost border that they all form. Figure 7.10 shows this envelope, along with the original data fit contour.

This contour is 47% larger than the original 68% contour, and so in some sense the systematic error for this analysis is 47%, though that statement applies to the 68% confidence limit and not the quoted best fit parameters. The systematic error result is verified by calculating the size difference between each individual error contour (the *DIS* low contour, the *DIS* hi contour, the *QE* low contour, etc.) and the data contour in the same way as was described above for the envelope contour. The quadratic sum of these fourteen individual errors is 46%, a reassuring result.

Lastly, it is worthwhile to emphasize that the best fit value of Δm_{23}^2 is completely insensitive to all systematic error investigated, save for the reconstructed muon momentum.

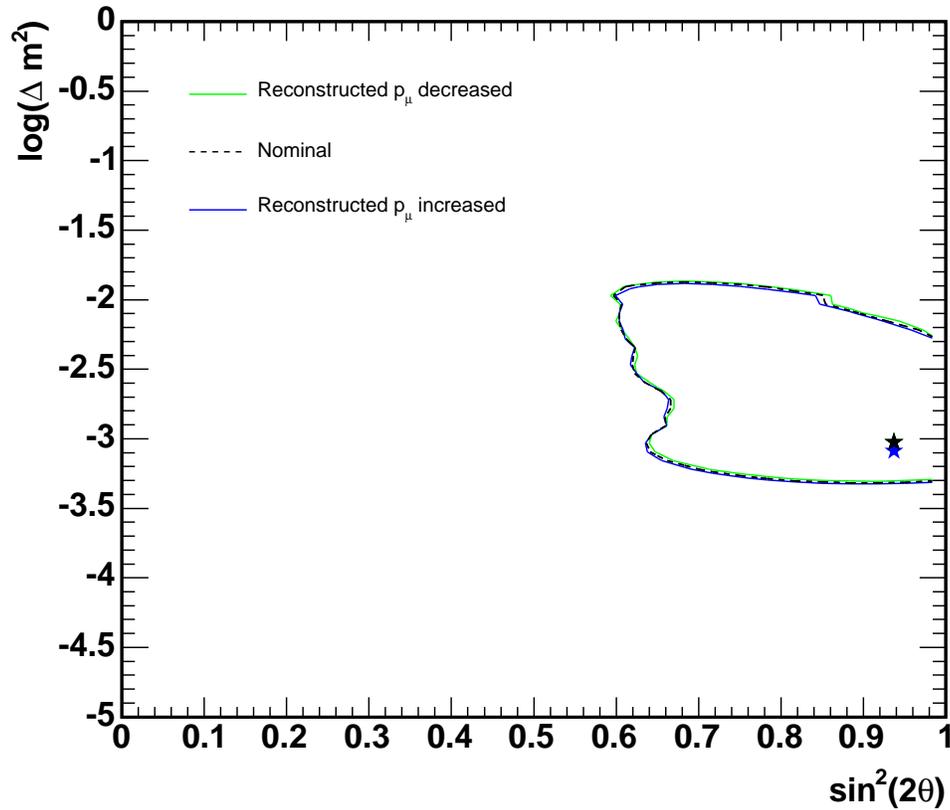


Figure 7.8: 68% contours for varying the reconstructed muon momentum. Black (dashed) shows the original data fit, green (light solid) shows the result of reducing all momenta by their specified error, and blue (dark solid) shows the result of increasing all momenta by their specified error. Stars show best-fit points. Note that decreasing the momentum did nothing to the best fit point, and that increasing the momentum did nothing to the best fit value of $\sin^2(2\theta_{23})$.

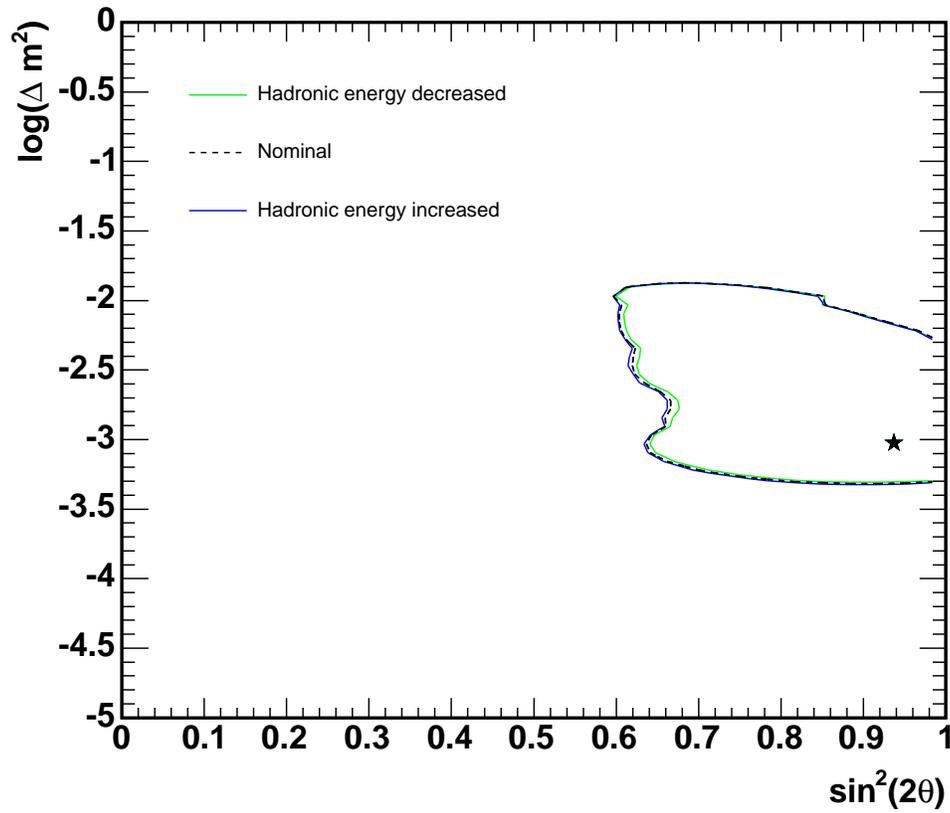


Figure 7.9: 68% contours for varying the reconstructed hadronic shower energy. Black (dashed) shows the original data fit, green (light solid) shows the result of reducing all shower energies by their specified error, and blue (dark solid) shows the result of increasing all shower energies by their specified error. Stars show best-fit points. Note that varying the hadronic energy did nothing to the best fit point.

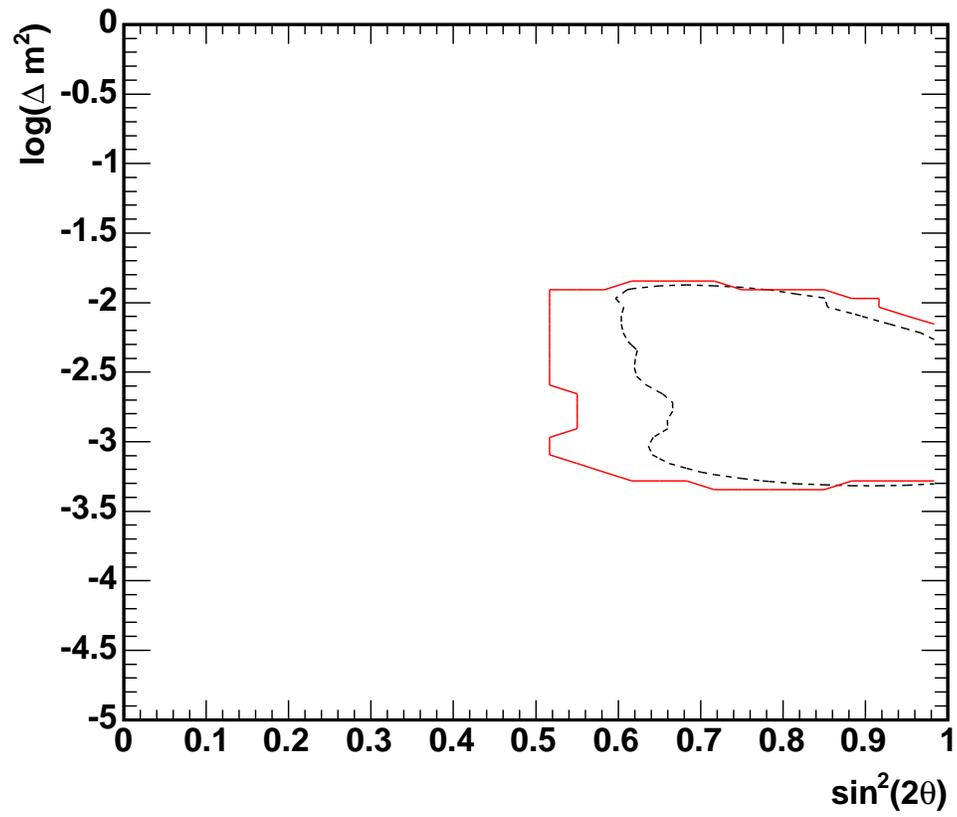


Figure 7.10: The cumulative systematic contour (red) and the original 68% contour (dashed black).

Neutrino Cross Section Error			
	Variation	Low/Hi $\sin^2 2\theta$	Low/Hi Δm_{23}^2
<i>DIS</i> Cross Section	$\pm 5\%$	0.95 / 0.98	0.93 / 0.93
<i>QE</i> Cross Section	$\pm 15\%$	0.85 / 0.98	0.93 / 0.93
<i>RES</i> Cross Section	$\pm 15\%$	0.92 / 0.98	0.93 / 0.93
Background Error			
	Variation	Low/Hi $\sin^2(2\theta_{23})$	Low/Hi Δm_{23}^2
Showering ν and CR μ	$\pm 12\%$	0.95 / 0.98	0.93 / 0.93
Reconstruction Error			
	Variation	Low/Hi $\sin^2(2\theta_{23})$	Low/Hi Δm_{23}^2
Neural Net	$\pm 10\%$ (in FOM)	0.85 / 0.92	0.93 / 0.93
Muon Momentum	$\pm 3\%$	0.95 / 0.95	0.93 / 0.81
Hadronic Energy	$\pm 5\%$	0.95 / 0.95	0.93 / 0.93
Nominal		$\sin^2 2\theta_{23} = 0.95$	$\Delta m_{23}^2 = 0.93 \times 10^{-3} \text{ eV}^2$

Table 7.1: A summary of the effects of all systematic error sources. The Δm_{23}^2 values should be understood as only the significant figures of the nominal value, listed at the bottom. For the Background section, *both* the neutrino background and the cosmic ray background were increased or decreased by 12%, and the results quoted are the cumulative effect of both variations.

7.7 Physical Boundary

Physically, a negative oscillation probability is meaningless, and so the parameter space is constrained to $\sin^2(2\theta_{23}) < 1$. Mathematically, nothing precludes the region $\sin^2(2\theta_{23}) > 1$, and artificially cutting off the fitting at this physical boundary can affect the results. As such, care must be taken to ensure these effects do not catastrophically alter the fit results.

This potential problem of the physical boundary has been directly investigated by repeating the likelihood fit over an extended grid which covers the same Δm_{23}^2 range, but goes out to $\sin^2(2\theta_{23}) = 2$. The result is shown in Figures 7.11 and 7.12. Allowing the fit to go into the unphysical region does not change the best fit point at all. It is reassuring to see that the best fit point is insensitive to the proximity of the physical boundary.

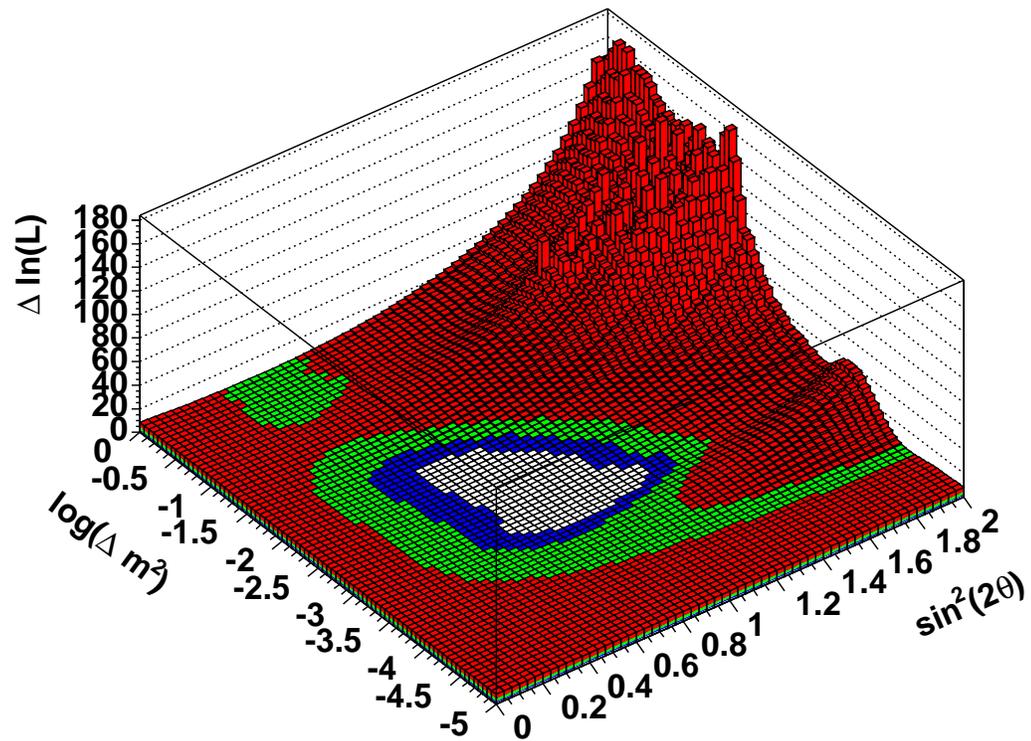


Figure 7.11: The likelihood surface obtained when the fit is allowed to go into the unphysical region. As usual, blue shows the 68% contour, green shows 90%, and red shows 99%. Obviously the region around $\sin^2(2\theta_{23}) = 2$ is very strongly excluded.

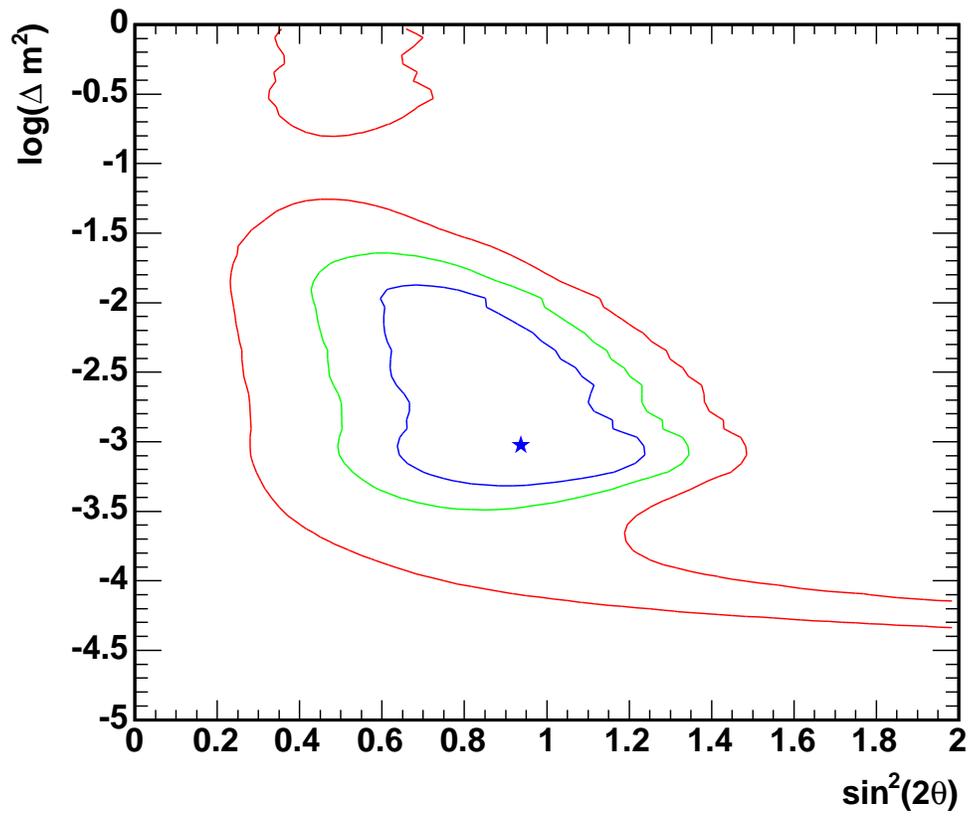


Figure 7.12: The likelihood contours obtained when the fit is allowed to go into the unphysical region. As usual, blue shows the 68% contour, green shows 90%, and red shows 99%.

Chapter 8

Conclusions

Science does not happen in a vacuum, and in most cases an analysis has peers. In this final chapter, other apropos results will be compared with the current results, some suggestions for refining the analysis will be put forth, and the analysis will be summarized.

8.1 Comparisons

The current measurement is not the first estimation of $\sin^2(2\theta_{23})$ and Δm_{23}^2 , so an immediate question is how this measurement compares to previous ones. First, as mentioned in Section 2.5.6, the Super-Kamiokande measurement is globally dominant, with some 12000 fully contained neutrino events over ~ 1500 live days, as of June 2005 [16]. They have also amassed around 1000 partially contained neutrino events, and are still collecting data as of this writing. The Super-Kamiokande constraints on the parameter space are quite tight, with their 99% contour falling well within this analysis' 68% contour, which is reassuring in the sense that the Super-K best fit point is compatible with these results.

Soudan 2 has also recently published their final atmospheric neutrino analysis [9] (the Soudan 2 detector is now disassembled). Figure 2.7 shows these results, along with their earlier results [71]. Ignoring the small local maxima which divide the Soudan 2 contours into lobes (these local maxima are largely irrelevant anyway, because the likelihood differences these maxima represent are very small), their contours are quite similar to the contours found in this analysis, covering roughly the same area in parameter space.

Two noteworthy beam experiments have published results. As mentioned in Section 2.5.7, K2K has analyzed 1.049×10^{20} protons on target (POT) and found $\sin^2(2\theta_{23}) = 1.0$ and $\Delta m_{23}^2 = 2.8 \times 10^{-3} \text{ eV}^2$ [6]. NuMI/MINOS has also recently released results based on 2.5×10^{20} POT: $\sin^2(2\theta_{23}) > 0.84$ and $\Delta m_{23}^2 = 2.38_{-0.16}^{+0.2} \times 10^{-3} \text{ eV}^2$ at 90% confidence [1]. Both of these beam result best fit points lie within the current analysis' 68% confidence limit.

Cambridge has a very active MINOS atmospheric neutrino group, and this group spearheaded the analysis upon which the first published MINOS analysis was based [3]. As this analysis was performed on some of the same data as Cambridge analysis (though with different reconstruction), it is worthwhile to compare and contrast the two.

This first Cambridge analysis was based on 107 total fully and partially contained muon neutrino and antineutrino interactions collected during the time interval August 2003 through February 2005. Figures 8.1, 8.2, and 8.3 show the 68%, 90%, and 99% contours for both the Cambridge analysis and this one. For what it's worth, the 68% contour for this analysis is 35% smaller than the corresponding Cambridge contour.

It is curious that the confidence limits presented in this thesis are tighter than the Cambridge confidence limits even though the total numbers of selected events for the two analyses are almost the same. This apparent contradiction should not cause alarm, as the two analyses accept very different classes of events. Specifically, Cambridge accepts partially contained events, whereas this analysis does not. Also, the Cambridge event

list and the event list for this analysis have only about a 10% overlap. Thus the two event sets are largely independent, and differences in the final likelihood surfaces should not come as a shock.

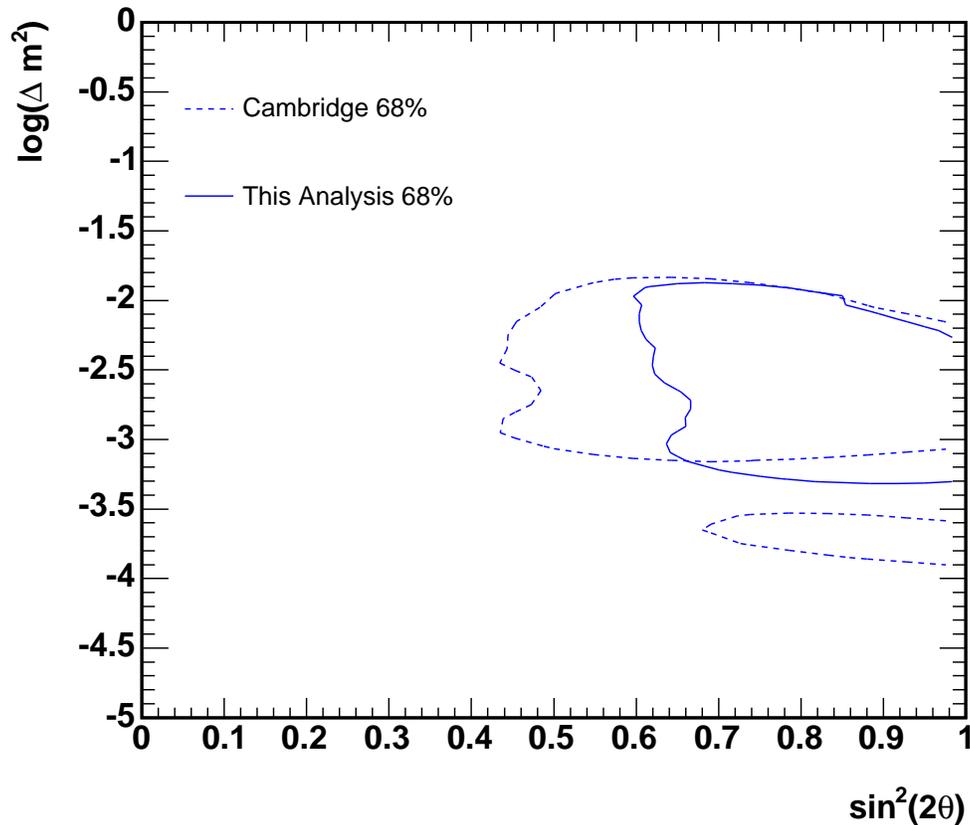


Figure 8.1: The 68% confidence limits for the Cambridge analysis and the present one. The allowed region for both contours is to the higher $\sin^2(2\theta_{23})$ side.

Cambridge has a second generation analysis nearly complete, which includes data collected through early 2007. With more than twice the data, as well as further refinement of the analysis itself, the forthcoming Cambridge results will no doubt be significantly more restrictive on the oscillation parameter space.

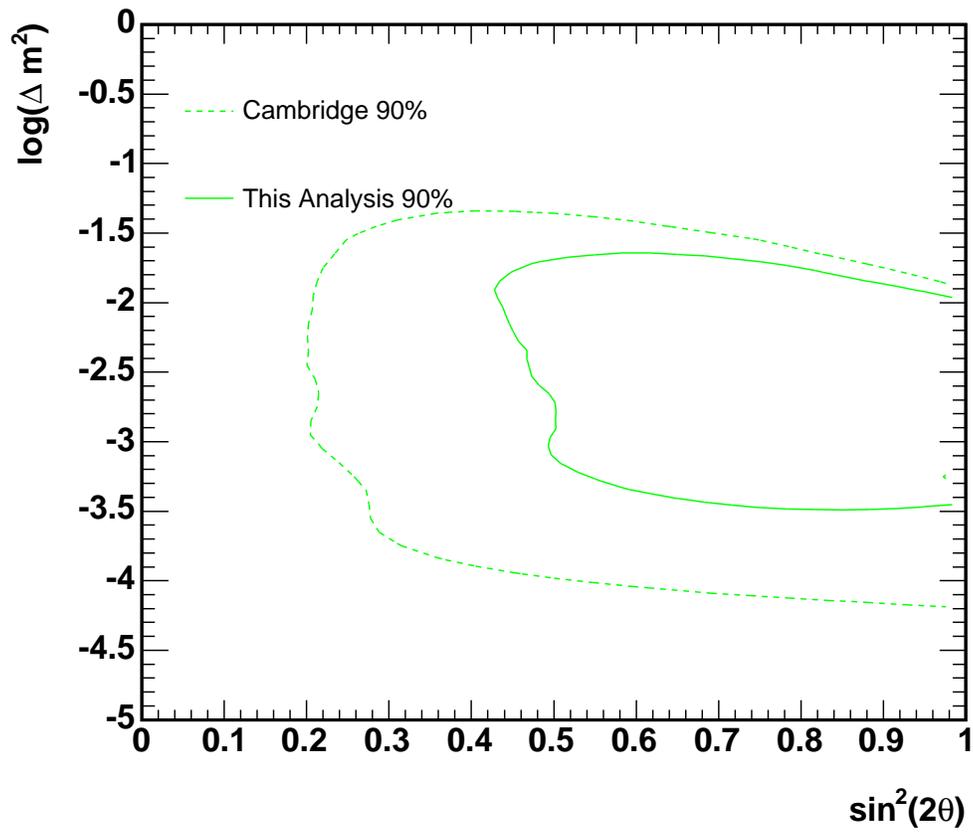


Figure 8.2: The 90% confidence limits for the Cambridge analysis and the present one. The allowed region for both contours is to the higher $\sin^2(2\theta_{23})$ side.

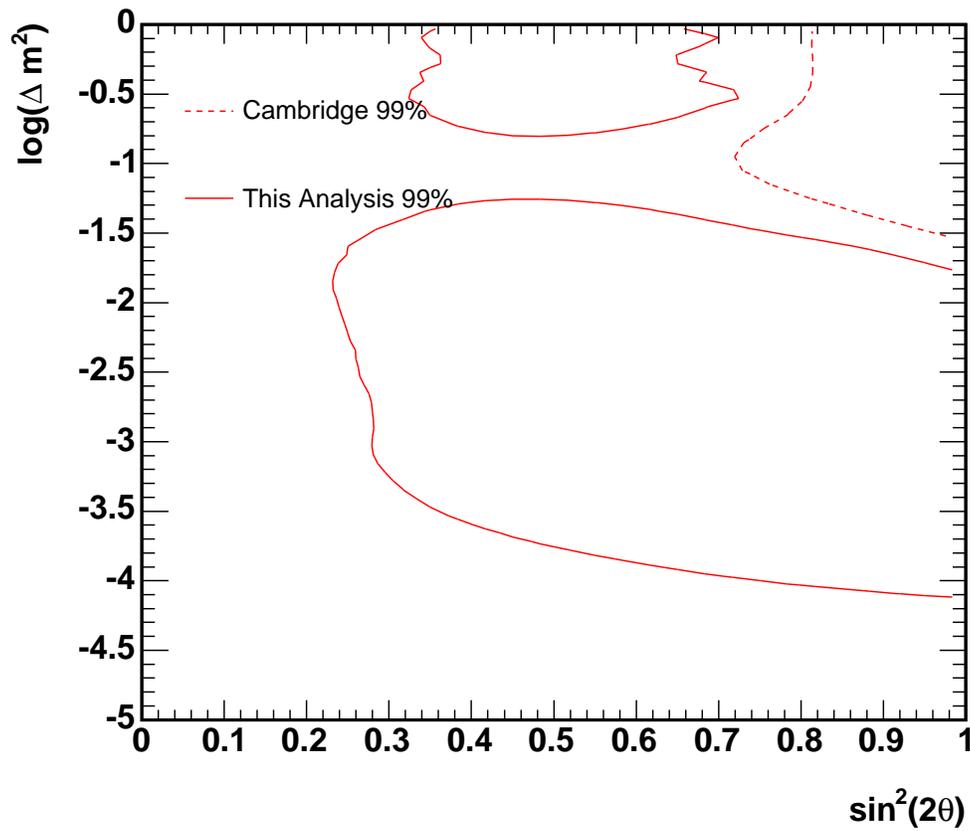


Figure 8.3: The 99% confidence limits for the Cambridge analysis and the present one. The allowed region for both contours is to the higher $\sin^2(2\theta_{23})$ side.

8.2 The Future

To paraphrase a popular bit of physics wisdom, an analysis is never really complete, its author just gives up. While that is certainly hyperbole, an analysis can always be refined further.

8.2.1 Sample Splitting

The ability of a detector to track an event and the reconstruction software to subsequently characterize that event depends on the event's physical properties. This effect is universal, as all detectors are inherently better at observing certain classes of events than others. As such, physicists often take advantage of this variable accuracy by dividing a total data set into subsets based on resolution. In principle, events which are more accurately reconstructed provide better discrimination for whatever hypothesis is being tested, and so can be weighted more heavily in the fitting procedure being carried out.

A preliminary study was carried out to investigate the utility of splitting the MINOS data into three subsamples based on the kinematic y , defined in Equation 2.14, of the event. In words, y is the fraction of the neutrino energy taken away by the hadronic portion of the interaction products, as shown in Figure 2.4. The lower y is, the more energy the lepton takes away and, particularly in the case of muon neutrinos, the more accurately MINOS should be able to reconstruct the neutrino direction and energy. The three ranges of y under consideration are

- Low y : $y < 0.2$
- Mid y : $0.2 < y < 0.5$
- High y : $y > 0.5$

Figure 8.4, which shows the resolution on $\log(L/E)$ for the three samples, should demonstrate physical motivation for splitting the sample. Surprisingly, the resolution is not significantly better for the low y sample than it is for the high y sample.

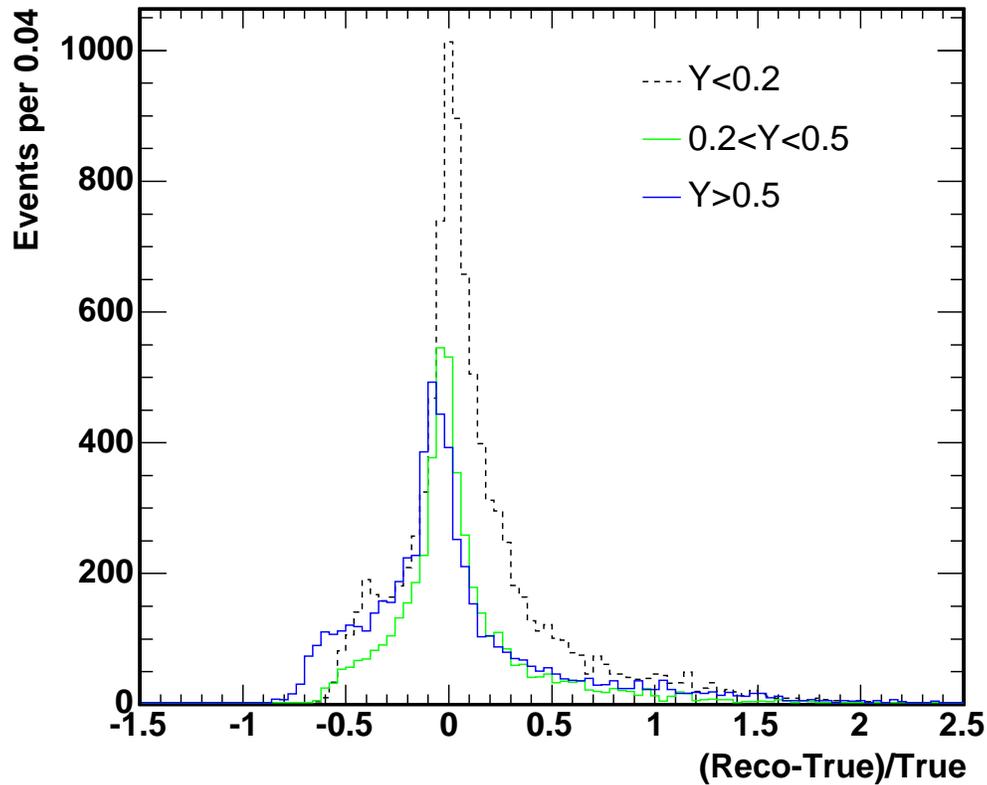


Figure 8.4: The resolution on $\log(L/E)$ for low y events (dashed black), mid y events (green), and hi y events (blue).

Still, there probably exists some other quantity on which to base the splitting that will yield an advantageous division of the data, and this subject could be one which yields an immediate improvement for this analysis.

8.3 Summary

883 live days of MINOS far detector data has been processed with a multiprong event selection and subsequently analyzed using a binned maximum likelihood method. The best fit oscillation parameters are $\sin^2 2\theta_{23} = 0.95_{-0.32}$ and $\Delta m_{23}^2 = 0.93_{-0.44}^{+3.94} \times 10^{-3} \text{ eV}^2$, and the 68% confidence limit includes the entire Super-Kamiokande 99% contour.

Appendix A

Low $\log(L/E)$ Data Events

As previously described, the overwhelming background for this measurement is cosmic ray muons. These muons can reach a detector only from above, not below, and consequently will have small reconstructed pathlength (L) values. As such, cosmic ray muon events tend to populate the low $\log(L/E)$ region, which is why the excess of events at low $\log(L/E)$ in Figure 6.8 demands special consideration.

First, consider the spatial distribution of these events through the detector. Figure A.1 shows the y vs x and z distributions of the vertices for events in the region $\log(L/E) < 1.5$. The vertices are not concentrated near the top of the detector ($y > 2$ m) as one would expect if the events were due to cosmic ray muons. Also, the events are not grouped near the ends of the detector ($z < 1$ m or $z > 29$ m), which would be another obvious signature of cosmic ray background slipping through the event selection.

The zenith angle distribution of these events, shown in Figure A.2, is consistent with the range of L values these events must have to fall in the range $\log(L/E) < 1.5$, and the range of event energies is uniform between approximately 1 and 8 GeV (also shown in Figure A.2). Together, these plots demonstrate that there was not some catastrophic reconstruction failure which artificially brought, for example, a medium L event with a

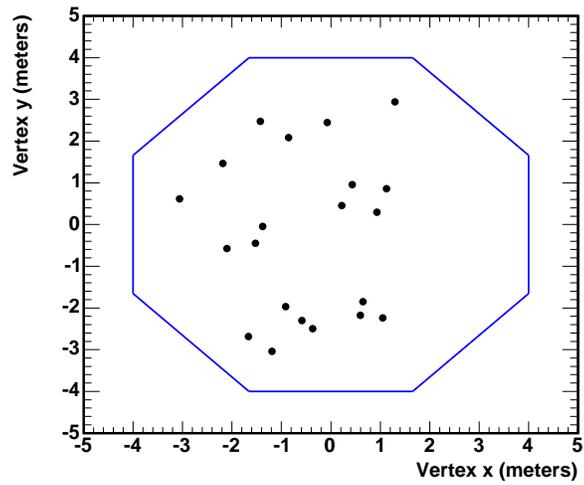
highly overestimated E into this $\log(L/E)$ region.

Figure A.3 shows the track direction cosine x and z distributions of these events; the x direction cosine is quite uniformly distributed. The z direction cosine is, perhaps, suspicious since it has 14 events with $\cos\theta_z < 0$ and 7 events with $\cos\theta_z > 0$. This distribution is $\sim 6\%$ probable according to the binomial distribution, treating the sign of $\cos\theta_z$ as a coin flip. As mentioned in Section 3.1, the z axis in MINOS runs down the long axis of the detector, pointing away from Fermilab. As such, if beam events were missed by the beam timing cut (Section 5.2.5), they would have *positive* $\cos\theta_z$ and so, while the large fraction of low $\log(L/E)$ events with $\cos\theta_z < 0$ is curious, it does not indicate any problem with the removal of beam events. Because of the detector's planar geometry, it is unsurprising that there are no events with $\cos\theta_z \approx 0$; the acceptance for such events is very poor.

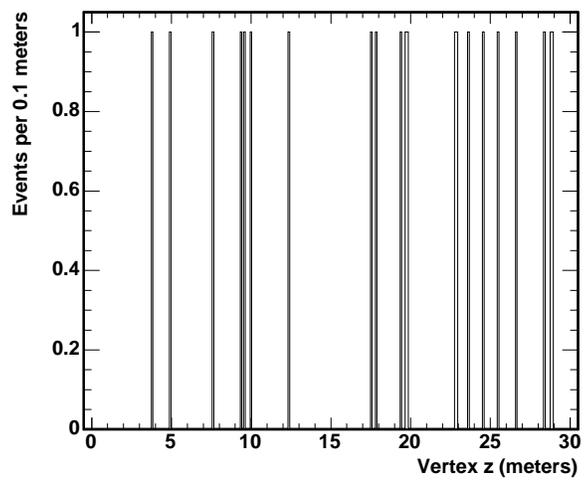
The last distribution for these events is shown in Figure A.4, which demonstrates that these events are evenly distributed temporally through the total data set and not due to some particular hardware failure or other (highly pathological) scenario localized in time.

Finally, these events have been hand scanned, and only one event (run 33750, subrun 2, snarl 36268) is obviously a cosmic ray muon. For this particular event, the veto shield had in-time activity but the track vertex lay just outside the spatial range of the hit shield plank, and thus the event was not tagged. All other events in this low $\log(L/E)$ region visually pass for neutrino induced muons.

Having considered these various characterizations of the low $\log(L/E)$ data events, it seems reasonable to assume that the excess of events with $\log(L/E) < 1.5$ is merely statistical fluctuation.

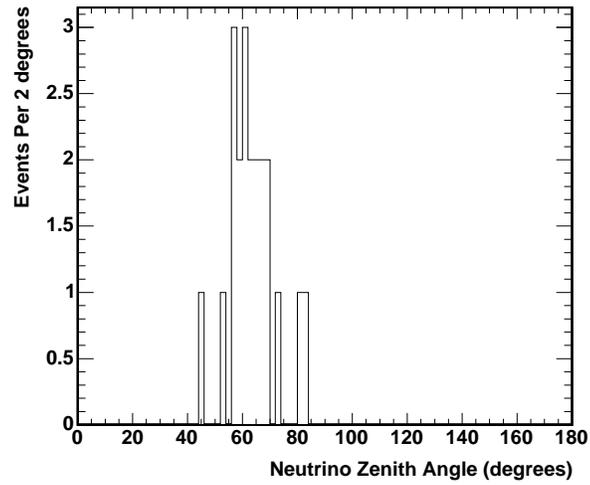


(a) Vertex Y vs X

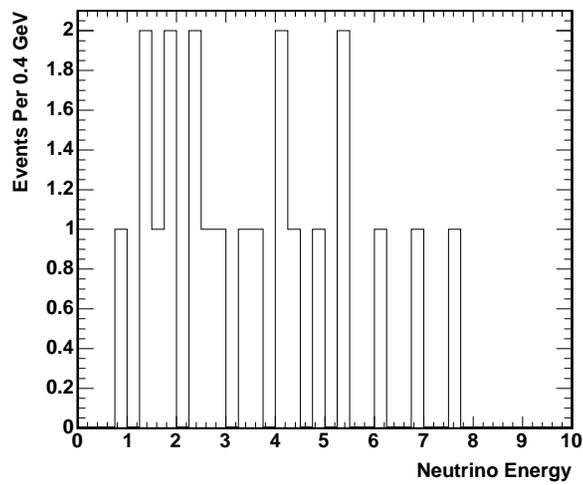


(b) Vertex Z

Figure A.1: The spatial vertex distribution of data events with $\log(L/E) < 1.5$.



(a) Zenith Angle



(b) Energy

Figure A.2: The zenith angle and energy distributions of data events with $\log(L/E) < 1.5$.

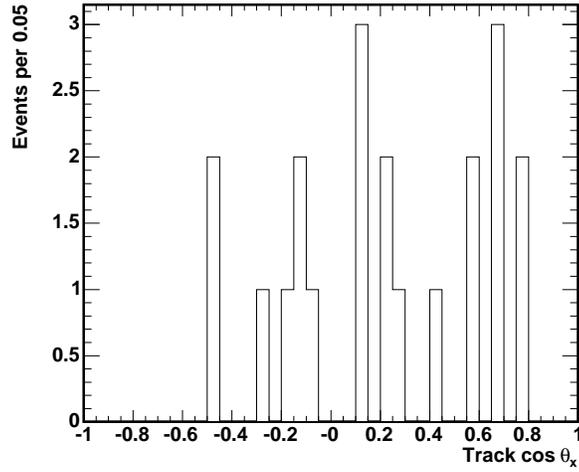
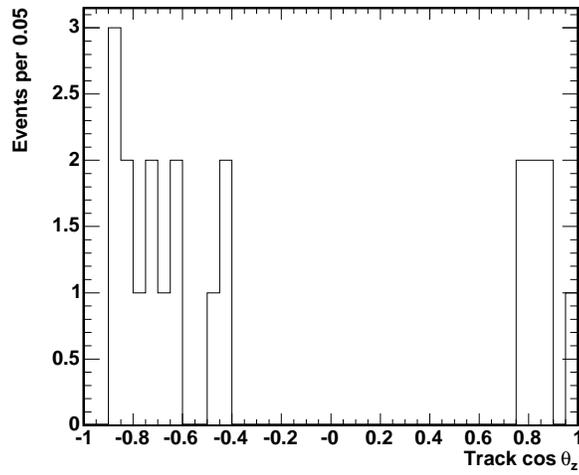
(a) Track direction cosine x (b) Track direction cosine z

Figure A.3: The track direction cosine x and z distributions for data events with $\log(L/E) < 1.5$.

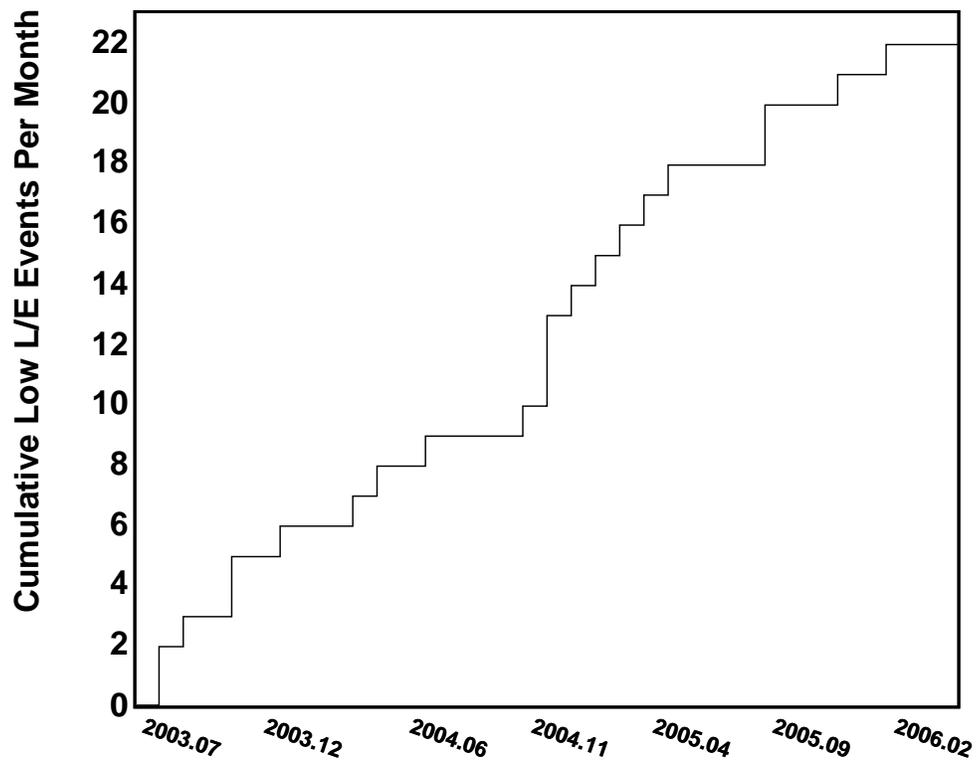


Figure A.4: The cumulative number of data events with $\log(L/E) < 1.5$.

References

- [1] Preliminary results from minos on muon neutrino disappearance based on an exposure of 2.5×10^{20} 120- gev protons on the numi target. 2007.
- [2] P. Adamson et al. The minos detectors technical design report. Technical report, MINOS Collaboration, 1998.
- [3] P. Adamson et al. First observations of separated atmospheric ν_μ and $\bar{\nu}_\mu$ events in the minos detector. *Phys. Rev.*, D73:072002, 2006.
- [4] M. Aguilar et al. The alpha magnetic spectrometer (ams) on the international space station. i: Results from the test flight on the space shuttle. *Phys. Rept.*, 366:331–405, 2002.
- [5] A. A. Aguilar-Arevalo et al. A search for electron neutrino appearance at the delta $m^2 \sim 1 \text{ eV}^2$ scale. 2007.
- [6] M. H. Ahn et al. Measurement of neutrino oscillation by the k2k experiment. *Phys. Rev.*, D74:072003, 2006.
- [7] E. Kh. Akhmedov. Neutrino physics. *arXiv:hep-ph/001264v2*, 2000.
- [8] W. W. M. Allison et al. The atmospheric neutrino flavor ratio from a 3.9 fiducial kiloton-year exposure of soudan 2. *Phys. Lett.*, B449:137–144, 1999.

- [9] W. W. M. Allison et al. Neutrino oscillation effects in soudan-2 upward-stopping muons. *Phys. Rev.*, D72:052005, 2005.
- [10] C. Alt et al. Inclusive production of charged pions in p + c collisions at 158 GeV/c beam momentum. 2006.
- [11] G. Ambrosini et al. Measurement of charged particle production from 450 GeV/c protons on beryllium. *Eur. Phys. J.*, C10:605–627, 1999.
- [12] A. V. Apanasenko et al. Study of high-energy primary cosmic rays in the runjob russian-japanese balloon experiment. *Bull. Russ. Acad. Sci. Phys.*, 65:476–479, 2001.
- [13] M. Apollonio et al. Limits on neutrino oscillations from the chooz experiment. *Phys. Lett.*, B466:415–430, 1999.
- [14] K. Asakimori et al. Cosmic-ray proton and helium spectra: results from the jacee experiment. *Astrophysical Journal*, 502:278–283, 1998.
- [15] Y. Ashie et al. Evidence for an oscillatory signature in atmospheric neutrino oscillation. *Phys. Rev. Lett.*, 93:101801, 2004.
- [16] Y. Ashie et al. A measurement of atmospheric neutrino oscillation parameters by super-kamiokande i. *Phys. Rev.*, D71:112005, 2005.
- [17] G. D. Barr, T. K. Gaisser, P. Lipari, Simon Robbins, and T. Stanev. A three-dimensional calculation of atmospheric neutrinos. *Phys. Rev.*, D70:023006, 2004.
- [18] G. D. Barr, T. K. Gaisser, S. Robbins, and Todor Stanev. Uncertainties in atmospheric neutrino fluxes. *Phys. Rev.*, D74:094009, 2006.
- [19] Andrew Bazarko. Miniboone: Status of the booster neutrino experiment. *Nucl. Phys. Proc. Suppl.*, 91:210–215, 2001.

- [20] R. Becker-Szendy et al. The electron-neutrino and muon-neutrino content of the atmospheric flux. *Phys. Rev.*, D46:3720–3724, 1992.
- [21] R. Becker-Szendy et al. Neutrino measurements with the imb detector. *Nucl. Phys. Proc. Suppl.*, 38:331–336, 1995.
- [22] Christoph Berger et al. A study of atmospheric neutrino oscillations in the frejus experiment. *Phys. Lett.*, B245:305–310, 1990.
- [23] A. Blake. *A Study of Atmospheric Neutrino Oscillations in the MINOS Far Detector*. PhD thesis, University of Cambridge, 2005.
- [24] F. Boehm et al. Final results from the palo verde neutrino oscillation experiment. *Phys. Rev.*, D64:112001, 2001.
- [25] F. Boehm and P. Vogel. *Physics of Massive Neutrinos*. Cambridge University Press, 1987.
- [26] M. Boezio et al. The cosmic-ray proton and helium spectra measured with the caprice98 balloon experiment. *Astropart. Phys.*, 19:583–604, 2003.
- [27] R. Brun, R. Hagelberg, M. Hansroul, and J. C. Lassalle. Geant: Simulation program for particle physics experiments. user guide and reference manual. 1978. CERN-DD-78-2-REV.
- [28] M. G. Catanesi et al. Measurement of the production cross-section of positive pions in p al collisions at 12.9 GeV/c. *Nucl. Phys.*, B732:1–45, 2006.
- [29] J. Chadwick. The intensity distribution in magnetic spectrum of β rays of radium b + c. *Verhandl. Dtsch. phys. Ges.*, 16:383, 1914.
- [30] V. Cherkassky and F. Mulier. *Learning From Data : Concepts, Theory, and Methods*. Wiley, 1998.

- [31] C. L. Cowan, F. Reines, F. B. Harrison, H. W. Kruse, and A. D. McGuire. Detection of the free neutrino: A confirmation. *Science*, 124:103–104, 1956.
- [32] G. Danby et al. Observation of high-energy neutrino reactions and the existence of two kinds of neutrinos. *Phys. Rev. Lett.*, 9:36–44, 1962.
- [33] Richard O. Duda and Peter E. Hart. Use of the hough transformation to detect lines and curves in pictures. *Commun. ACM*, 15(1):11–15, 1972.
- [34] S. Eidelman et al. Review of particle physics. *Phys. Lett.*, B592:1, 2004.
- [35] Ralph Engel, T. K. Gaisser, and Todor Stanev. Pion production in proton collisions with light nuclei: Implications for atmospheric neutrinos. *Phys. Lett.*, B472:113–118, 2000.
- [36] N. Felt et al. Minos far detector electronics user’s manual. Technical report, Harvard University and others, 2001.
- [37] E. Fermi. An attempt of a theory of beta radiation. 1. *Z. Phys.*, 88:161–177, 1934.
- [38] Y. Fisyak et al. The mipp experiment, proposal 907. www.fnal.gov, 2000.
- [39] Y. Fukuda et al. Atmospheric muon-neutrino / electron-neutrino ratio in the multi-gev energy range. *Phys. Lett.*, B335:237–245, 1994.
- [40] T. K. Gaisser and M. Honda. Flux of atmospheric neutrinos. *Ann. Rev. Nucl. Part. Sci.*, 52:153, 2002.
- [41] T. K. Gaisser, T. Stanev, M. Honda, and P. Lipari. Primary spectrum to 1 TeV and beyond. 2001. Prepared for 27th International Cosmic Ray Conference (ICRC 2001), Hamburg, Germany, 7-15 Aug 2001.

- [42] H. Gallagher. *Neutrino Oscillation Searches With the Soudan 2 Detector*. PhD thesis, University of Minnesota, 1996.
- [43] H. Gallagher. The neugen neutrino event generator. *Nucl. Phys. Proc. Suppl.*, 112:188–194, 2002.
- [44] G. Giacomelli and A. Margiotta. Macro results on atmospheric neutrinos. *Nucl. Phys. Proc. Suppl.*, 145:116–119, 2005.
- [45] M. Goldhaber, L. Grodzins, and A. W. Sunyar. Helicity of neutrinos. *Phys. Rev.*, 109:1015–1017, 1958.
- [46] J. Hartnell. *Measurement of the Calorimetric Energy Scale in MINOS*. PhD thesis, University of Oxford, 2005.
- [47] K. S. Hirata et al. Observation of a small atmospheric ν_μ/ν_e ratio in kamiokande. *Phys. Lett.*, B280:146–152, 1992.
- [48] Quang Ho-Kim and Xuan-Yem Pham. *Elementary particles and their interactions*. Springer-Verlag. Berlin, Germany, Springer-Verlag, 1998. 661p.
- [49] Raymond Davis Jr., Don S. Harmer, and Kenneth C. Hoffman. Search for neutrinos from the sun. *Phys. Rev. Lett.*, 20:1205–1209, 1968.
- [50] Rudolph Emil Kalman. A new approach to linear filtering and prediction problems. *Transactions of the ASME—Journal of Basic Engineering*, 82(Series D):35–45, 1960.
- [51] S. Kasahara. *A Study of Cosmic Ray Composition in the Knee Region using Multiple Muon Events in the Soudan 2 Detector*. PhD thesis, University of Minnesota, 1997.
- [52] Boris Kayser. Neutrino physics. *ECONF*, C040802:L004, 2004.

- [53] K. Kodama et al. Observation of tau-neutrino interactions. *Phys. Lett.*, B504:218–224, 2001.
- [54] T. D. Lee and Chen-Ning Yang. Question of parity conservation in weak interactions. *Phys. Rev.*, 104:254–258, 1956.
- [55] M. Messier. *Evidence for Neutrino Mass From Observations of Atmospheric Neutrinos With Super-Kamiokande*. PhD thesis, Boston University, 1999.
- [56] D. G. Michael et al. Observation of muon neutrino disappearance with the minos detectors and the numi neutrino beam. *Phys. Rev. Lett.*, 97:191801, 2006.
- [57] S.L. Mufson and B.J. Rebel. Cosmic muon physics with the minos far detector. Technical report, Indiana University, 2005.
- [58] R. Nichol. *Calibration of the MINOS Detectors*. PhD thesis, University College London, 2003.
- [59] R. Nichol and C. Smith. Strip-to-strip calibration at the calibration detector. Technical report, University College London and California Institute of Technology, 2003.
- [60] Jaroslaw A. Nowak and Jan T. Sobczyk. Modelling cc neutrino cross sections in the few gev energy region. *Acta Phys. Polon.*, B37:1955–1966, 2006.
- [61] L.B. Okun. *Leptons and Quarks*. Elsevier Science B.V., 1982.
- [62] W. Pauli. Open letter to radioactive ladies and gentlemen (translation into english). *Physics Today*, 31:27, 1978.
- [63] Martin L. Perl et al. Evidence for anomalous lepton production in $e^+ e^-$ annihilation. *Phys. Rev. Lett.*, 35:1489–1492, 1975.

- [64] D. Petyt. Far detector event selection and data quality cuts for the 1e20 analysis. Technical report, University of Minnesota, 2006.
- [65] B. Pontecorvo. Inverse beta processes and nonconservation of lepton charge. *Sov. Phys. JETP*, 7:172–173, 1958.
- [66] B. Pontecorvo. Electron and muon neutrinos. *Sov. Phys. JETP*, 10:1236–1240, 1960.
- [67] B. Pontecorvo. Inverse beta process. *Cambridge Monogr. Part. Phys. Nucl. Phys. Cosmol.*, 1:25–31, 1991.
- [68] B. Rebel. *Neutrino-Induced Muons in the MINOS Far Detector*. PhD thesis, Indiana University, 2004.
- [69] B. Rebel and S. Mufson. A demultiplexing algorithm for cosmic rays in the minos far detector. Technical report, Indiana University, 2002.
- [70] F. Reines and C. L. Cowan. Detection of the free neutrino. *Phys. Rev.*, 92:830–831, 1953.
- [71] Mayly C. Sanchez et al. Observation of atmospheric neutrino oscillations in soudan 2. *Phys. Rev.*, D68:113004, 2003.
- [72] T. Sanuki et al. Precise measurement of cosmic-ray proton and helium spectra with the bess spectrometer. *Astrophys. J.*, 545:1135, 2000.
- [73] D.S. Sivia. *Data Analysis A Bayesian Tutorial*. Oxford University Press, 1996.
- [74] Mark Thomson. Alternative demuxer. Talk given at collaboration meeting, March 2005.

- [75] C. S. Wu, E. Ambler, R. W. Hayward, D. D. Hoppes, and R. P. Hudson. Experimental test of parity conservation in beta decay. *Phys. Rev.*, 105:1413–1414, 1957.