

Measuring ν_μ Disappearance with the MINOS Experiment

Jessica Mitchell

of

Murray Edwards College, University of Cambridge

A dissertation submitted to the University of Cambridge
for the degree of Doctor of Philosophy

Abstract

The MINOS Experiment consists of two steel-scintillator calorimeters, sampling the long baseline NuMI muon neutrino beam. It was designed to make a precise measurement of the ‘atmospheric’ neutrino mixing parameters, $\Delta m_{\text{atm.}}^2$ and $\sin^2(2\theta_{\text{atm.}})$. The Near Detector measures the initial spectrum of the neutrino beam 1 km from the production target, and the Far Detector, at a distance of 735 km, measures the impact of oscillations in the neutrino energy spectrum. Work performed to validate the quality of the data collected by the Near Detector is presented as part of this thesis.

This thesis primarily details the results of a ν_μ disappearance analysis, and presents a new sophisticated fitting software framework, which employs a maximum likelihood method to extract the best fit oscillation parameters. The software is entirely decoupled from the extrapolation procedure between the detectors, and is capable of fitting multiple event samples (defined by the selections applied) in parallel, and any combination of energy dependent and independent sources of systematic error.

Two techniques to improve the sensitivity of the oscillation measurement were also developed. The inclusion of information on the energy resolution of the neutrino events results in a significant improvement in the allowed region for the oscillation parameters. The degree to which $\sin^2(2\theta) = 1.0$ could be disfavoured with the exposure of the current dataset if the true mixing angle was non-maximal, was also investigated, with an improved neutrino energy reconstruction for very low energy events.

The best fit oscillation parameters, obtained by the fitting software and incorporating resolution information were: $|\Delta m^2| = 2.32_{-0.08}^{+0.12} \times 10^{-3} \text{ eV}^2$ and $\sin^2(2\theta) > 0.90(90\% \text{ C.L.})$. The analysis provides the current world best measurement of the atmospheric neutrino mass splitting Δm^2 . The alternative models of neutrino decay and decoherence are disfavoured by 7.8σ and 9.7σ respectively.

Declaration

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

Jessica Mitchell

Acknowledgements

This thesis would never have been possible without the help and support of a great many people, to whom I am exceedingly grateful.

Firstly, I would like to thank my supervisor, Mark Thomson, for his support, guidance and excellent physics and statistics advice. It was always delivered with a smile, and no doubt will stand me in good stead for the future. Many thanks also to Andy Blake, my office mate, sounding board and stalwart travelling companion, who saw me through volcanic ash clouds and writing my first class in C++. It has been a great pleasure to be part of the Cambridge HEP group, long may Thursday pub and curry continue. I would like to thank everyone for making the group a fun place to be, with a special mention for Ruth Toner, an excellent office mate and friend.

Thank you to my MINOS colleagues for providing the support and opportunities which made this thesis happen. I would like to express my gratitude to the MINOS Charged Current analysis group, in particular Justin Evans, whose patience and guidance was very much appreciated. To those who suffered alongside me in the CC box opening of 2010; I salute you. Also, to all the members of Young MINOS I knew during my time on the experiment, you made collaboration meetings fun, and expanded my knowledge of good American beer by several orders of magnitude.

Last but not least, thank you to my friends and family, Eva and Stephen Mitchell, and my partner, David Valinsky, for their help and support over the past four years; for giving me the confidence to continue and for making endless cups of tea. I wouldn't have made it without you.

Finally, I would like to dedicate this thesis to my father, John Mitchell. His memory is a constant source of inspiration to me.

Contents

1	Introduction	1
2	The Physics of Neutrino Oscillations	4
2.1	Neutrinos in the Standard Model	4
2.2	Neutrino Oscillation Theory	6
2.2.1	Two-Flavour Neutrino Case	10
2.2.2	Neutrino Interactions	13
2.2.3	Alternative Models of Neutrino Disappearance	14
2.3	Neutrino Oscillation Experiments	15
2.3.1	Long Range Oscillations	16
2.3.2	Short Range Oscillations	21
2.3.3	Sub-Dominant Short Range Oscillations	29
2.4	Summary and Current Status	33
3	The MINOS Experiment	34
3.1	The NuMI Beam	35
3.2	The MINOS Detectors	39
3.2.1	Steel and Scintillator	40
3.2.2	Far Detector	42

3.2.3	Calibration Detector	43
3.3	Calibration, Simulation and Data Reconstruction	44
3.3.1	Detector Calibration	44
3.3.2	Monte Carlo Simulation	45
3.3.3	Event Reconstruction	48
3.4	Analysis Dataset	51
4	The Near Detector and Data Quality	54
4.1	Detector Overview	54
4.1.1	Front-End Electronics	56
4.1.2	Data Acquisition and Triggers	60
4.2	Near Detector Data Quality	61
4.2.1	Summary of Quality Cuts	61
4.2.2	Details of Failure Modes	62
4.3	Implementation	70
4.4	Rejected Data	71
5	The Charged Current Analysis	74
5.1	Overview of an Oscillation Analysis	74
5.2	Event Selection	75
5.2.1	Pre-Selection	75
5.2.2	Fiducial Volume	76
5.2.3	Particle Identification (PID)	77
5.3	k NN Shower Energy	80
5.4	Extrapolation	86
5.5	Rock and Anti-Fiducial Events	89

6	Fitting the Far Detector Data	93
6.1	Fit Mechanism and Implementation	94
6.1.1	Energy Binning	96
6.2	Systematic Uncertainties	101
6.2.1	Incorporating Systematics into the Framework	104
6.2.2	Correlation of Systematics	108
6.3	Fit Validation	111
6.3.1	Simple MC Tests	112
6.3.2	Mock Data Challenges	114
6.4	Sensitivities	116
6.5	Alternative Models	119
6.6	Summary	119
7	Event Resolution	121
7.1	Resolution Calculation	121
7.1.1	Comparison to Calibration Detector Energy Resolution	125
7.1.2	σ/E Distributions	129
7.1.3	Oscillation Fitting in Bins of Resolution	131
7.2	Performance Studies	132
7.2.1	Sensitivity Studies	132
7.2.2	Detector Effects	137
7.2.3	Alternative Model Discrimination	144
7.3	Summary	145
8	Sensitivity to Non-Maximal Mixing	146
8.1	Motivation	146

8.2	Improving the Energy Reconstruction	149
8.2.1	Kinematics of Neutrino-Nucleon Scattering	149
8.2.2	Event Selection in True Neutrino Energy	150
8.2.3	Potential Sensitivity Gain	151
8.2.4	Correcting the k NN Shower Energy	152
8.2.5	Neutrino Energy Calculation	156
8.2.6	Muon Angle Resolution	157
8.3	Results for Events Selected with $E_{\text{true}} < 2 \text{ GeV}$	159
8.4	Event Selection in Reconstructed Quantities	163
8.5	Summary	165
9	Results	168
9.1	Previous MINOS Oscillation Results	168
9.2	Selected Events	170
9.3	Cross Checks	172
9.3.1	Reproducing the 2008 Result	172
9.3.2	Fitting 2008 Data with Analysis Improvements	174
9.3.3	Comparison of Fitting Software Results	175
9.4	Results	177
9.4.1	Cross Checks Performed Post-Unblinding	182
9.4.2	Alternative Models	185
10	Conclusion	188
10.1	Future Developments	190
	Bibliography	1

List of Figures	9
List of Tables	17

“... Your cleverest poets will not give you leave to be just and good, even if you want to. For here was Minos, a man who exceeded all men in cruelty, and who enslaved with his navies the inhabitants of continent and islands alike, and yet they honour him by placing in his hand a sceptre of justice and give him a throne in Hades to be umpire of spirits ...”

— Flavius Philostratus, *Life of Apollonius of Tyana* 3.25

Chapter 1

Introduction

The existence of the neutrino was first postulated in 1930 by Wolfgang Pauli [1], a ‘desperate remedy’ in order to preserve the conservation of energy and explain the electron energy spectrum observed from β -decay. Originally named the ‘neutron’, it was required to be electrically neutral, have a mass much less than that of the electron, be spin- $\frac{1}{2}$ and interact weakly with matter. It would be 26 years before its existence was confirmed by experiment, with neutrinos from a fission reactor detected by the Savannah River experiment [2].

After the confirmation of the existence of the neutrino, a number of experimental results followed in the late 1950s and 1960s, in parallel with the theoretical work which developed the nature of the weak interaction and electroweak unification by Glashow, Weinberg and Salam [3–5]. The second generation of neutrino, in addition to the electron-neutrino postulated by Pauli, was identified in 1962 by Lederman *et al.* [6]. The discovery of the τ lepton in 1975 implied the existence of a corresponding third neutrino flavour, which was discovered by the DONUT collaboration [7] in 2000. The LEP experiments [8] on the resonance peak of the Z^0 boson provided strong evidence that only three active generations of neutrinos exist in the Standard Model.

The first evidence that neutrinos were massive was seen in deficits of neutrinos produced by the sun, the “solar neutrino problem” and from cosmic rays, the “atmospheric neutrino anomaly”. There is now a large body of compelling evidence that these deficits are the result of the phenomenon of neutrino oscillations, where neutrinos undergo periodic transformations between flavour states as they propagate, with the flavour states as mixtures of the neutrino mass eigenstates. Chapter 2 describes the theory of neutrino

oscillations, and gives a summary of the historical developments from the observation of the solar neutrino problem in the late 1960s to the present day, providing the context for the MINOS Experiment.

The MINOS Experiment, with its two detectors sampling the long baseline NuMI muon neutrino beam was designed to make a precise measurement of the ‘atmospheric’ neutrino mixing parameters, $\Delta m_{\text{atm.}}^2$ and $\sin^2(2\theta_{\text{atm.}})$. These parameters govern the transition between $\nu_\mu \leftrightarrow \nu_\tau$ in the region where $\Delta m_{23}^2 L/E \sim \mathcal{O}(1)$. The Near Detector measures the initial spectrum of the neutrino beam 1 km from the production target, and the Far Detector, at a distance of 735 km, measures the impact of oscillations in the neutrino energy spectrum.

Chapter 3 describes in detail the NuMI beam, the MINOS Far Detector, the MINOS Monte Carlo simulation and the event reconstruction software. The Near Detector is described in Chapter 4, with details of software, developed for this thesis, which validates the quality of the data collected at this detector.

Chapter 5 describes the procedures used in previous MINOS analyses; differentiating a Charge Current ν_μ sample from the Neutral Current, $\bar{\nu}_\mu$ and ν_τ backgrounds; and how extrapolating information from the Near Detector provides a more accurate Far Detector prediction. Two analysis refinements for the most recent published result, each the subject of a thesis in its own right, are also summarised; a new technique to reconstruct the hadronic shower energy, and the selection of an optimised event sample including events which interact outside the fiducial volume of the Far Detector. An overview of the analysis procedure is also provided.

Developed for this thesis, two further analysis refinements are described in detail in Chapters 6 and 7. A new fitting method to extract the best fit oscillation parameters was developed and is presented in Chapter 6, decoupled from the extrapolation procedure and capable of fitting multiple event samples (defined by the selections applied) in parallel, and any combination of energy dependent and independent systematics across these samples. A significant improvement in sensitivity is achieved with the use of muon track and hadronic shower energy resolution information as described in Chapter 7, in conjunction with the improved fitting method.

MINOS has previously published results where $\sin^2(2\theta)$ is unphysical or very close to maximal mixing. Chapter 8 summarises a second sensitivity study in addition to the inclusion of resolution information; investigating how sensitive the detectors would be to a non-maximal mixing angle, and the degree to which $\sin^2(2\theta) = 1.0$ could be

disfavoured with the exposure of the current dataset.

Chapter 9 combines the analysis components summarised in Chapter 5, and the improvements described in Chapters 6 and 7 to extract an optimised measurement of the oscillation parameters. The analysis provides the world best measurement of the atmospheric neutrino mass splitting Δm^2 . Also presented are the results of fitting neutrino decay and decoherence as alternative models of neutrino disappearance.

Finally, the conclusion in Chapter 10 gives a summary of the work presented in this thesis, discussing the unknowns that remain in the sector, and how future results from MINOS and related experiments aim to further expand our knowledge of the neutrino physics sector.

Chapter 2

The Physics of Neutrino Oscillations

2.1 Neutrinos in the Standard Model

The Standard Model (SM) of Particle Physics, which describes the interactions of the 17 fundamental particles, incorporates massless neutrinos which interact via the unified electroweak force. The SM originated as a theory of massless fields with the local gauge symmetry $SU(3) \times SU_L(2) \times U(1)$, unifying the strong force (interactions described by an $SU(3)$ symmetry) with the electroweak force ($SU_L(2) \times U(1)$). Particle masses are introduced via spontaneous symmetry breaking and the Higgs mechanism.

The Standard Model places pairs of leptons in weak doublets, where each flavour of massless neutrino is associated with a charged lepton: the electron neutrino (ν_e) with the electron; the muon neutrino (ν_μ) with the muon and finally the tau neutrino (ν_τ) with the tau lepton. The Standard Model dictates that interactions of leptons are confined within these doublets, so the total number of any generation is conserved.

Neutrinos couple only to the weak force, and so undergo two types of interaction, mediated by the massive W^\pm and Z^0 bosons. Charged current (CC) interactions occur where the neutrino is converted into its partner lepton via mediation by the W^\pm boson with a quark or lepton; the flavour of the neutrino can be determined by observation of its leptonic partner. Neutral current (NC) interactions occur where the neutrino exchanges a Z^0 boson with a quark and scatters; the flavour of the incoming neutrino cannot be



Figure 2.1: *Left:* A neutral current interaction where a neutrino of flavour l scatters off a nucleon, exchanging a neutral Z^0 boson. *Right:* A charged current interaction, where a neutrino of flavour l interacts with a nucleon, exchanging a charged W boson and producing a lepton of type l .

determined. The SM Lagrangian terms for these two interactions are as follows:

$$\mathcal{L}_{CC} = -\frac{g}{2} \sum_j (\bar{\varphi}_{jL} \gamma^\mu W_\mu^- \nu_{jL} + \varphi_{jL} \gamma^\mu W_\mu^+ \bar{\nu}_{jL}), \text{ and} \quad (2.1)$$

$$\mathcal{L}_{NC} = -\frac{g}{2 \cos \theta_W} \sum_j \bar{\nu}_{jL} \gamma^\mu Z_\mu^0 \nu_{jL}, \quad (2.2)$$

where the sum over j implies the sum over all lepton flavours: $\varphi_j = (e, \mu, \tau)$ and $\nu_j = (\nu_e, \nu_\mu, \nu_\tau)$. The label L indicates only left-handed fermions interact with the weak force; neutrinos maximally violate parity (\mathcal{P}) and charge conjugation (\mathcal{C}). Feynman diagrams for the CC and NC processes are given in Figure 2.1. At low energies, the large masses of these mediating gauge bosons give rise to the exceedingly small cross sections for neutrino interactions with matter.

An increasing number of experiments have been constructed to observe neutrino interactions and reconstruct their kinematic properties. Such observations have provided compelling evidence neutrinos undergo a quantum mechanical process known as “neutrino oscillations”, where neutrino flavour changes with periodic variations. This implies a non-zero mass for the neutrino, required for such leptonic flavour mixing to occur.

The remainder of this Chapter provides an overview of the theory of neutrino oscillations, with a discussion of additional alternative models of neutrino propagation. The current status of the experimental evidence for neutrino oscillations is reviewed to provide context for the MINOS Experiment, which aims to confirm the existence of neutrino oscillations and measure the governing parameters to high precision.

2.2 Neutrino Oscillation Theory

Non-zero neutrino masses are not included in the GSW electroweak theory, and so must be incorporated by adding mass terms to the Standard Model Lagrangian. The mass term has the following form, assuming the neutrino is a Dirac spinor and not a Majorana particle:

$$\mathcal{L}_{mass} = - \sum_{\alpha, \alpha'} \bar{\nu}_{\alpha L} M_{\alpha\alpha'} \nu_{\alpha' R} + \text{h.c.} \quad (2.3)$$

where $\nu_{\alpha L, R}$ ($\alpha = e, \mu, \tau$) are the neutrino ‘flavour’ states which couple to the weak force, and M is a 3×3 complex matrix. If neutrinos are massive particles, they also exist in a spectrum of mass eigenstates $|\nu_i\rangle$ ($i = 1, 2, 3$) which give the vacuum free particle solutions to the wave equation describing the propagation of neutrinos in space and time. They can be expressed as linear combinations of the flavour eigenstates $|\nu_{\alpha L, R}\rangle$ ($\alpha = e, \mu, \tau$):

$$|\nu_{\alpha L}\rangle = \sum_{i=1,2,3} U_{\alpha i} |\nu_{iL}\rangle \quad |\nu_{\alpha R}\rangle = \sum_{i=1,2,3} V_{\alpha i R} |\nu_{iR}\rangle \quad (2.4)$$

where $U_{\alpha i}$ is the lepton mixing matrix for the left handed neutrino, satisfying the unitary condition $UU^\dagger = \mathbf{1}$. A similar matrix exists for the right handed neutrino given by $V_{\alpha i}$. The fields written in this mass basis diagonalise the matrix $M = U^\dagger m_i \delta_{ij} V$ in Equation 2.3, rewriting the Lagrangian mass term in the standard form:

$$\mathcal{L}_{mass} = - \sum_{i=1}^3 m_i \bar{\nu}_i \nu_i \quad (2.5)$$

where $m_i > 0$ and represents the non-zero neutrino masses.

The contributions of Maki, Nakagawa and Sakata to the theory of the mixing of massive neutrinos, and Pontecorvo to the discussion of the phenomenology of neutrino mixing in the 1960s, have caused the unitary matrix U to be commonly known as the *Pontecorvo-Maki-Nakagawa-Sakata* (PMNS) matrix [9, 10]. The PMNS matrix components are shown here in the most common parameterisation, as the product of three rotation matrices based on mixing angles θ_{ij} between the i th and j th mass eigenstates

$(\theta_{12}, \theta_{13}, \theta_{23})$, and a phase factor δ :

$$\begin{aligned}
 U &= \begin{pmatrix} U_{e1} & U_{e2} & U_{e3} \\ U_{\mu1} & U_{\mu2} & U_{\mu3} \\ U_{\tau1} & U_{\tau2} & U_{\tau3} \end{pmatrix} \\
 &= \begin{pmatrix} 1 & 0 & 0 \\ 0 & c_{23} & s_{23} \\ 0 & -s_{23} & c_{23} \end{pmatrix} \begin{pmatrix} c_{13} & 0 & s_{13}e^{-i\delta} \\ 0 & 1 & 0 \\ -s_{13}e^{i\delta} & 0 & c_{13} \end{pmatrix} \begin{pmatrix} c_{12} & s_{12} & 0 \\ -s_{12} & c_{12} & 0 \\ 0 & 0 & 1 \end{pmatrix} \begin{pmatrix} e^{i\alpha} & 0 & 0 \\ 0 & e^{i\beta} & 0 \\ 0 & 0 & 1 \end{pmatrix} \quad (2.6)
 \end{aligned}$$

where $c_{ij} \equiv \cos \theta_{ij}$ and $s_{ij} \equiv \sin \theta_{ij}$. The phase factor δ is known as the ‘Dirac phase’, and if non-zero is the source of \mathcal{CP} violation in the neutrino sector. The phases α and β are required if neutrinos are Majorana particles, where the neutrino is identical to its anti-particle. These Majorana elements do not affect the observation of neutrino oscillations and are thus neglected for the purposes of this thesis.

Assuming neutrinos are produced in a vacuum, they propagate as free particle mass eigenstates with position four-vector $\mathbf{X} = (t, \mathbf{x})$ and momentum four vector $\mathbf{p} = (E, \mathbf{p})$, such that state $|\nu_i\rangle$ propagates as:

$$|\nu_i(\mathbf{x})\rangle = e^{-i\mathbf{p}_i \cdot \mathbf{X}} |\nu_i\rangle \quad (2.7)$$

Therefore a neutrino produced in flavour state $|\nu_\alpha\rangle$ evolves as:

$$|\nu_\alpha(\mathbf{x})\rangle = \sum_{i=1,2,3} e^{-i\mathbf{p}_i \cdot \mathbf{X}} U_{\alpha i} |\nu_i\rangle \quad (2.8)$$

Inverting the PMNS matrix in Equation 2.4 using the principles of unitarity, the evolution of the flavour state can be rewritten as a combination of all flavour states:

$$|\nu_\alpha(\mathbf{x})\rangle = \sum_{\beta=e,\mu,\tau} \sum_{i=1,2,3} e^{-i\mathbf{p}_i \cdot \mathbf{X}} U_{\alpha i} U_{\beta i}^* |\nu_\beta\rangle \quad (2.9)$$

Assuming the neutrino masses m_i are small compared to the energy at which they are produced and working in natural units, the three-momentum of mass state i can be approximated as follows. In the case of relativistic neutrinos, where $v_i t \approx ct = L$ and L is the distance travelled from the point of production, the phase of state i can be

expressed as:

$$p_i = \sqrt{E_i^2 - m_i^2} \simeq E_i - \frac{m_i^2}{2E_i}$$

$$\therefore \mathbf{p}_i \cdot \mathbf{X} = E_i t - p_i L \simeq \frac{m_i^2}{2E_i} L$$

Substituting for this phase in Equation 2.9, and assuming all mass eigenstates have the same energy E gives the following relationship between phases:

$$|\nu_\alpha(\mathbf{x})\rangle = \sum_{\beta=e,\mu,\tau} \sum_{i=1,2,3} e^{-im_i^2 L/2E} U_{\alpha i} U_{\beta i}^* |\nu_\beta\rangle \quad (2.10)$$

Equation 2.10 demonstrates that if the neutrino masses m_i are different, the phases of these eigenstates will evolve at different rates. If the PMNS matrix exhibits non-zero off-diagonal components, an initially pure neutrino flavour state will, therefore, with time and distance develop contributions from other flavours as its mass eigenstate propagates. Therefore a neutrino produced in one flavour state may be detected as a different flavour at some distance L after time t .

After the neutrino, initially produced in state $|\nu_\alpha\rangle$ undergoes a weak interaction, its wavefunction collapses into state $\langle\nu_\beta|$:

$$\langle\nu_\beta|\nu_\alpha(\mathbf{x})\rangle = \sum_i U_{\alpha i} U_{\beta i}^* e^{-im_i^2 L/2E} \quad (2.11)$$

Substituting in from Equation 2.9, the probability of detecting the neutrino in this state at a distance \mathbf{x} is given by:

$$\begin{aligned} P(\nu_\alpha \rightarrow \nu_\beta) &= |\langle\nu_\beta|\nu_\alpha(\mathbf{x})\rangle|^2 \\ &= \left(\sum_i U_{\alpha i}^* U_{\beta i} e^{im_i^2 L/2E} \right) \left(\sum_j U_{\alpha j} U_{\beta j}^* e^{-im_j^2 L/2E} \right) \\ &= \sum_i \sum_j U_{\alpha i}^* U_{\beta i} U_{\alpha j} U_{\beta j}^* e^{-i\Delta m_{ij}^2 L/2E} \\ &\quad + \left\{ \sum_i \sum_j U_{\alpha i}^* U_{\beta i} U_{\alpha j} U_{\beta j}^* - \sum_i \sum_j U_{\alpha i}^* U_{\beta i} U_{\alpha j} U_{\beta j}^* \right\} \end{aligned}$$

where $\Delta m_{ij}^2 = m_i^2 - m_j^2$, the squared mass splitting between the i th and j th eigenstates.

The last two terms are identical and can be added to give:

$$P(\nu_\alpha \rightarrow \nu_\beta) = \sum_i \sum_j \underbrace{U_{\alpha i}^* U_{\beta i} U_{\alpha j} U_{\beta j}^*}_{A_{ij}} (e^{-i\Delta m_{ij}^2 L/2E} - 1) + \underbrace{\sum_i \sum_j U_{\alpha i}^* U_{\beta i} U_{\alpha j} U_{\beta j}^*}_B$$

A_{ij} can be rewritten using the relationship $z + z^* = 2\Re[z]$, as $A_{ii} = 0$ and $A_{ij} = A_{ji}^*$. Using the unitarity of the PMNS matrix, the second term, B , can be rewritten as:

$$\begin{aligned} B &= \sum_i \sum_j U_{\alpha i}^* U_{\beta i} U_{\alpha j} U_{\beta j}^* \\ &= \left(\sum_i U_{\alpha i}^* U_{\beta i} \right) \left(\sum_j U_{\alpha j} U_{\beta j}^* \right) \\ &= \left(\sum_i U_{\beta i} U_{i\alpha}^\dagger \right) \left(\sum_j U_{\alpha j} U_{j\beta}^\dagger \right) \\ &= \delta_{\alpha\beta} \end{aligned}$$

where $\delta_{\alpha\beta}$ is the Kronecker delta.

Hence the oscillation probability can be written in the form:

$$P(\nu_\alpha \rightarrow \nu_\beta) = \delta_{\alpha\beta} + 2 \sum_{i(>j)} \sum_j \Re \left[U_{\alpha i}^* U_{\beta i} U_{\alpha j} U_{\beta j}^* (e^{-i\Delta m_{ij}^2 L/2E} - 1) \right]$$

Expanding the complex exponential and using

$$\cos\left(\frac{\Delta m_{ij}^2 L}{2E}\right) = 1 - 2 \sin^2\left(\frac{\Delta m_{ij}^2 L}{4E}\right)$$

gives the common form of the oscillation probability:

$$\begin{aligned} P(\nu_\alpha \rightarrow \nu_\beta) &= \delta_{\alpha\beta} \\ &\quad - 4 \sum_{i(>j)} \sum_j \Re \left[U_{\alpha i}^* U_{\beta i} U_{\alpha j} U_{\beta j}^* \right] \sin^2\left(\frac{\Delta m_{ij}^2 L}{4E}\right) \\ &\quad + 2 \sum_{i(>j)} \sum_j \Im \left[U_{\alpha i}^* U_{\beta i} U_{\alpha j} U_{\beta j}^* \right] \sin\left(\frac{\Delta m_{ij}^2 L}{2E}\right) \end{aligned} \quad (2.12)$$

The transition probability in Equation 2.12 is periodic with distance from the point of origin of the neutrino. If the mass eigenstates possess different masses and the neutrino and flavour states are not coincident, the neutrino will oscillate between flavours as it

propagates.

A similar expression for anti-neutrinos can be defined as that for the case of neutrinos. Assuming \mathcal{CPT} invariance, it can be shown that:

$$P(\nu_\alpha \rightarrow \nu_\beta) = P(\bar{\nu}_\beta \rightarrow \bar{\nu}_\alpha)$$

where $\bar{\nu}_\alpha$ denotes the anti-particle. Substituting $\alpha \leftrightarrow \beta$ in Equation 2.12, the probability as a function of neutrino type ($\nu_{\alpha,\beta}$) and mixing matrix (U) becomes:

$$P(\nu_\alpha \rightarrow \nu_\beta, U) = P(\bar{\nu}_\alpha \rightarrow \bar{\nu}_\beta, U^*) \quad (2.13)$$

Hence by replacing the unitary matrix with its complex conjugate, the anti-neutrino oscillation probabilities can be extracted from the neutrino case. Clearly if the PMNS matrix, as shown in Equation 2.6, has a complex component, the above equality will no longer hold, and result in \mathcal{CP} violation in the neutrino sector.

2.2.1 Two-Flavour Neutrino Case

Current experimental data (discussed fully in Section 2.3) suggests the neutrino mass splittings conform to a hierarchy where $|\Delta m_{23}^2| \gg |\Delta m_{12}^2|$, and where the angles θ_{12} and θ_{23} are large, with θ_{13} being small. In most experiments, including MINOS, this leads the three-flavour oscillation probability to effectively decouple into three sets of two-flavour oscillation models, each associated with one of the rotation matrices comprising the PMNS matrix.

From Equation 2.12, extracting the muon neutrino survival probability of interest to MINOS:

$$P(\nu_\mu \rightarrow \nu_\mu) = 1 - 4 \sum_{i(>j)} \sum_j |U_{\mu i}|^2 |U_{\mu j}|^2 \sin^2 \left(\frac{\Delta m_{ij}^2 L}{4E} \right) \quad (2.14)$$

where the \mathfrak{Im} component disappears for $\alpha = \beta$. For the above to simplify to a two-neutrino case, two approximations must hold: the oscillations driven by the solar mass splitting (Δm_{21}^2) must be small, and only the mixing angle θ_{23} can be relevant.

The first approximation is a good one, as experimentally $\Delta m_{12}^2 \sim \mathcal{O}(7 \times 10^{-5} \text{ eV}^2)$ and so the $\sin^2 \left(\frac{\Delta m_{12}^2 L}{4E} \right)$ term tends to zero for the MINOS energy $E \sim 1 \text{ GeV}$ and base-

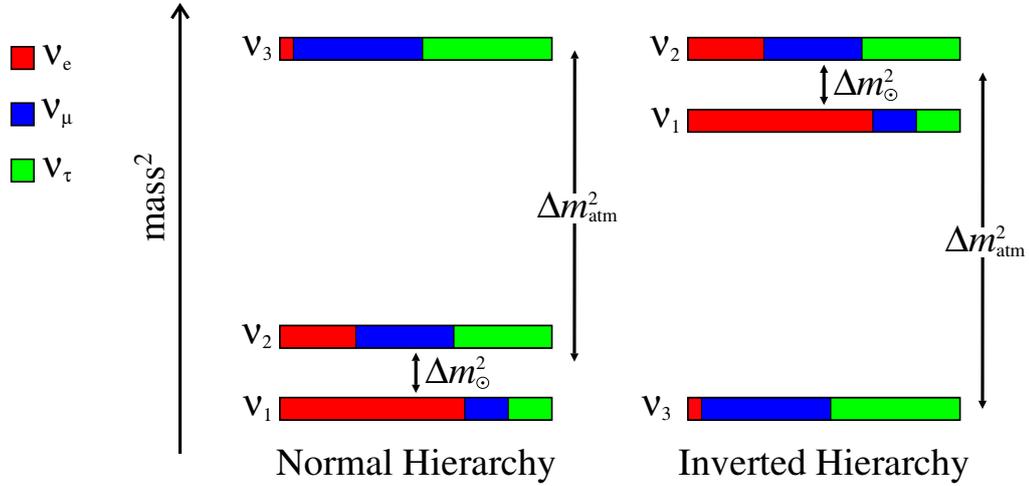


Figure 2.2: Schematic of the permitted mass orderings in the neutrino sector. Δm_{atm}^2 denotes the atmospheric mass splitting and Δm_{\odot}^2 the solar mass splitting (also referred to as Δm_{21}^2). The diagram on the left denotes the “normal” configuration, while the diagram on the right shows the inverted configuration. The colours give the approximate flavour compositions of each mass eigenstate. Taken from [11].

line $L = 735$ km. Hence, from $\Delta m_{13}^2 - \Delta m_{12}^2 = \Delta m_{23}^2$ we can substitute $\Delta m_{13}^2 \approx \Delta m_{23}^2$, used interchangeably with the label Δm_{atm}^2 . The second approximation holds regardless of the size of the solar mixing angle θ_{12} , and as shown later in this chapter experiment has found $\sin(\theta_{13})$ to be small, permitting the approximations $\sin(\theta_{13}) = 0$ and $\cos(\theta_{13}) = 1$. Figure 2.2 shows the two possible orderings of the neutrino masses for the three neutrino case, the “normal” and “inverted” hierarchies.

Substituting the PMNS matrix components into Equation 2.14 and applying the approximations above, the ν_{μ} survival probability in the two neutrino approximation reduces to:

$$P(\nu_{\mu} \rightarrow \nu_{\mu}) = 1 - \sin^2(2\theta_{23}) \sin^2\left(\frac{1.27\Delta m_{atm}^2 L}{E}\right) \quad (2.15)$$

with the units shown in Equation 2.16. Figure 2.3 shows this two-neutrino probability as a function of energy for the MINOS baseline ($L = 735$ km).

$$\frac{\Delta m_{ij}^2 L}{4E} \cong 1.27 \frac{\Delta m_{ij}^2 (\text{eV}^2) L (\text{km})}{E (\text{GeV})} \quad (2.16)$$

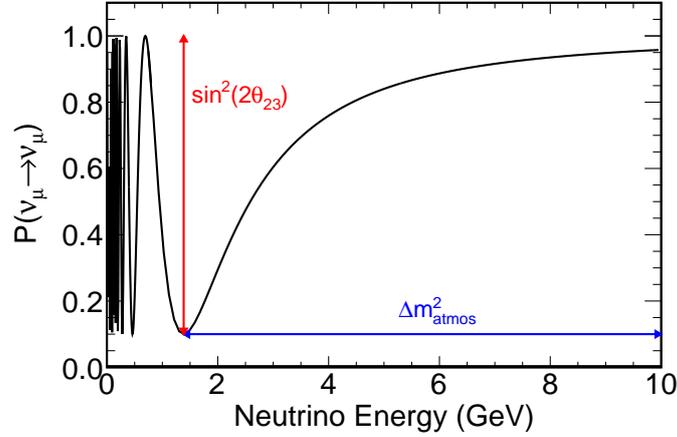


Figure 2.3: The $\nu_\mu \rightarrow \nu_\mu$ survival probability in the two neutrino flavour case, as in Equation 2.15 for the MINOS baseline $L = 735$ km. The depth of the oscillation minimum is determined by the mixing angle $\sin^2(2\theta_{23})$, while the position of the dip in neutrino energy is determined by the mass splitting $\Delta m_{atm.}^2$. At low energies the rapid variation in the probability cannot be resolved by experiment and averages to $\frac{1}{2}$. Taken from [12].

The full set of decoupled two-flavour models are defined as follows:

- **Short Range Oscillations:** Defined as the range where $\Delta m_{23}^2 L/E \sim \mathcal{O}(1)$, the dominant oscillation mode is between $\nu_\mu \leftrightarrow \nu_\tau$; governed by Δm_{23}^2 and $\sin^2 2\theta_{23}$ as derived above. Oscillations of this type are investigated by atmospheric neutrino detection experiments such as Super-Kamiokande, and long-baseline beam experiments such as MINOS, with oscillation probability:

$$P(\nu_\mu \leftrightarrow \nu_\tau) \simeq \sin^2(2\theta_{23}) \sin^2\left(\frac{1.27\Delta m_{23}^2 L}{E}\right) \quad (2.17)$$

- **Long Range Oscillations:** Over longer ranges where $\Delta m_{12}^2 L/E \sim \mathcal{O}(1)$, the transition between $\nu_e \leftrightarrow \nu_\tau/\nu_\mu$ dominates, and is governed by θ_{12} and $|\Delta m_{12}^2|$. These are observed in the disappearance of ν_e solar neutrinos and long baseline reactor experiments. The oscillation probability is given by:

$$P(\nu_e \leftrightarrow \nu_\tau/\nu_\mu) \simeq \sin^2(2\theta_{12}) \sin^2\left(\frac{1.27\Delta m_{12}^2 L}{E}\right) \quad (2.18)$$

- **Sub-Dominant Short Range Oscillations:** In the third sector of the neutrino mixing matrix, the mode $\nu_e \leftrightarrow \nu_\tau/\nu_\mu$ is governed by the small mixing angle θ_{13} . Generally observed by short baseline reactor experiments, it is also the focus of a MINOS analysis in addition to the short range oscillation mode detailed above. The oscillation probability is given by:

$$P(\nu_e \leftrightarrow \nu_\tau/\nu_\mu) \simeq \sin^2(2\theta_{13}) \sin^2\left(\frac{1.27\Delta m_{23}^2 L}{E}\right) \quad (2.19)$$

A variety of neutrino oscillation experiments have been constructed in order to probe different L/E values, and determine limits on different regions of oscillation parameter space. In order for the oscillations to be observable, a characteristic $L/E \simeq \pi/\Delta m^2$ is required; the neutrino flux is then sampled in the region with maximum oscillation probability. In the case where $L/E \ll \pi/\Delta m^2$, the oscillation probability is too small to be observable; where $L/E \gg \pi/\Delta m^2$, the oscillation probability averages to $\frac{1}{2} \sin^2 \theta$ and all information regarding the mass splittings is lost.

2.2.2 Neutrino Interactions

The theory of the weak interaction permits two types of neutrino interaction: Neutral Current (NC) and Charged Current (CC) as shown in Figure 2.1. The detection of the final state lepton after the exchange of a W^\pm with the nucleon allows the determination of the incoming neutrino flavour. The CC interaction can be further subdivided into three categories, each possessing distinct characteristics:

1. Quasi-Elastic (QE):

$$\nu_l + n \rightarrow l + p \quad (2.20)$$

Where the neutrino scatters off the nucleus. This form of interaction dominates for low energy (< 1 GeV) neutrinos.

2. Resonance (RES):

$$\nu_l + N \rightarrow N' + l \quad (2.21)$$

Where N' is an excited baryonic state such as Δ^+ , which decays to a pion in the final state.

3. Deep Inelastic Scattering (DIS):

$$\nu_l + N \rightarrow l + X \quad (2.22)$$

Where the neutrino interacts with a quark and the target nucleus breaks up, with X denoting the resultant hadronic shower. This form of interaction dominates at higher energies.

2.2.3 Alternative Models of Neutrino Disappearance

While the focus here has been on the theory of oscillations to explain changes in neutrino flavour composition, alternative models have been proposed to account for these phenomena. Different expressions for the $|\nu_\mu\rangle$ survival probabilities are constructed, and can be tested alongside pure oscillations to place limits in parameter space and determine which model best describes the data. The MINOS Experiment has considered two such alternative models, neutrino decoherence and neutrino decay.

Neutrino Decoherence

Neutrino decoherence [13] has been proposed as a method for neutrino flavour changes, affecting the quantum mechanical wavefunction of the neutrino. Flavour-dependent coherence lengths would eventually destroy the phase relationship between the mass states. Neutrinos travelling over very long baselines, such as from supernovae would be expected to experience this phenomenon. New physics beyond the Standard Model is required for this effect to be observable over small distances, such as quantum gravity in the context of black hole thermodynamics [14].

With a pure $|\nu_\mu\rangle$ neutrino source, the survival probability for a combined decoherence and oscillations model can be expressed as follows, with the introduction of a single new parameter μ^2 :

$$P(\nu_\mu \rightarrow \nu_\mu) = 1 - \frac{1}{2} \sin^2(2\theta_{23}) \left[1 - e^{-\frac{\mu^2 L}{2E}} \cos\left(\frac{\Delta m_{32}^2 L}{2E}\right) \right] \quad (2.23)$$

where the phenomenological parameter μ^2 determines the degree of decoherence. For the limiting case where $\mu^2 \rightarrow 0$, the survival probability reverts to the standard pure oscillation case of Equation 2.15. In the absence of oscillations, setting $\Delta m_{32}^2 = 0$ obtains

the pure decoherence survival probability:

$$P(\nu_\mu \rightarrow \nu_\mu) = 1 - \frac{1}{2} \sin^2(2\theta_{23}) \left[1 - e^{-\frac{\mu^2 L}{2E}} \right] \quad (2.24)$$

Neutrino Decay

Neutrino decay [15] permits one mass phase $|\nu_i\rangle$ to be unstable, and decay with some defined lifetime τ_i . Decay could then occur alongside neutrino oscillations, and if the squared mass splitting is sufficiently small to average to zero the decay mode could dominate experimental observations. The survival probability for $|\nu_\mu\rangle$ in a combined decay and oscillation model, permitting mass state $|\nu_2\rangle$ to decay is as follows:

$$P(\nu_\mu \rightarrow \nu_\mu) = \sin^4(\theta_{23}) + \cos^4(\theta_{23}) e^{-\frac{\alpha L}{E}} + 2 \sin^2(\theta_{23}) \cos^2(\theta_{23}) e^{-\frac{\alpha L}{2E}} \cos\left(\frac{\Delta m_{32}^2 L}{2E}\right) \quad (2.25)$$

where $\alpha = m_2/\tau_2$. If $|\nu_2\rangle$ decays into a sterile (non-observable) neutrino state and the squared mass splitting is sufficiently small, the pure decay mode has the resulting survival probability:

$$P(\nu_\mu \rightarrow \nu_\mu) = \left(\sin^2(\theta_{23}) + \cos^2(\theta_{23}) e^{-\frac{\alpha L}{2E}} \right)^2 \quad (2.26)$$

2.3 Neutrino Oscillation Experiments

Neutrino detection is based upon observing and distinguishing charged current and neutral current neutrino interactions. Such experiments generally require intense neutrino beams and detectors with large fiducial masses to provide a useful event rate. Charged current events are observed and determine the flavour composition, while Neutral current events are required to determine the overall flux of the neutrino source. Detectors are also generally located underground to provide shielding from cosmic ray induced background.

Experiments can be grouped broadly into two categories; ‘disappearance’ measurements, where a deficit of the flavour of interest is searched for relative to the initial neutrino flux, and ‘appearance’ measurements, where an excess in a sub-dominant flavour is searched for relative to the incident neutrino flux. The following section details the cur-

rent status of the experimental neutrino physics field to provide context for the MINOS Experiment.

2.3.1 Long Range Oscillations

The first evidence for neutrino oscillations was observed in solar neutrinos, an intense flux of ν_e produced by the nuclear fusion reactions within the Sun. The current best understanding of the Sun is given by the Standard Solar Model (SSM) [16], which can be used to determine the expected neutrino flux from the rates of the solar fusion reactions.

The primary method of energy generation within the sun results from the pp chain, the dominant chain of fusion reactions which overall combines four hydrogen nuclei (protons) into a helium nucleus, and two electron neutrinos:



where the full set of fusion interactions are summarised in Figure 2.4 and Table 2.1, as predicted by the SSM. Greater than 90% of solar neutrinos are emitted in the above process, yet due to their low maximum energy of 0.425 MeV they are more difficult to detect. Decays of the by-products of the fusion process emit higher energy neutrinos, as shown in Figure 2.4. Electron neutrinos resulting from ${}^8\text{B} \rightarrow {}^8\text{Be} + e^+ + \nu_e$ decay are produced at energies up to 14.1 MeV, and so become ideal experimental candidates.

Radio-Chemical Experiments

The first experiment to measure the solar neutrino flux was the Homestake Experiment [18], a pioneering radio-chemical detector which made its first measurements in the late 1960s and ran until 1995. The Homestake experiment detected solar neutrinos from the conversion of chlorine into argon, via the inverse beta decay process: ${}^{37}\text{Cl} + \nu_e \rightarrow {}^{37}\text{Ar} + e^-$ in a large tank of C_2Cl_4 . The detector was primarily sensitive to ${}^8\text{B}$ neutrinos, as the threshold for this interaction is 0.841 MeV. The ${}^{37}\text{Ar}$ was permitted to accumulate for several months prior to extraction from the tank, and the number of atoms counted by monitoring their decays; a process which could be performed with an efficiency of $\geq 90\%$.

The final measurement of the solar neutrino flux was $2.56 \pm 0.16(\text{stat.}) \pm 0.16(\text{syst.})$ SNU (Solar Neutrino Units, where $1 \text{ SNU} \equiv 1$ interaction per 10^{36} atoms per second.). This is a clear deficit compared to the value predicted by the SSM at Homestake, which at

Reaction	Abbr.	Flux ($\text{cm}^{-2} \text{s}^{-1}$)
$pp \rightarrow d e^+ \nu$	pp	$5.97(1 \pm 0.006) \times 10^{10}$
$pe^- p \rightarrow d \nu$	pep	$1.41(1 \pm 0.011) \times 10^8$
${}^3\text{He} p \rightarrow {}^4\text{He} e^+ \nu$	hep	$7.90(1 \pm 0.15) \times 10^3$
${}^7\text{Be} e^- \rightarrow {}^7\text{Li} \nu + (\gamma)$	${}^7\text{Be}$	$5.07(1 \pm 0.06) \times 10^9$
${}^8\text{B} \rightarrow {}^8\text{Be}^* e^+ \nu$	${}^8\text{B}$	$5.94(1 \pm 0.11) \times 10^6$
${}^{13}\text{N} \rightarrow {}^{13}\text{C} e^+ \nu$	${}^{13}\text{N}$	$2.88(1 \pm 0.15) \times 10^8$
${}^{15}\text{O} \rightarrow {}^{15}\text{N} e^+ \nu$	${}^{15}\text{O}$	$2.15(1^{+0.17}_{-0.16}) \times 10^8$
${}^{17}\text{F} \rightarrow {}^{17}\text{O} e^+ \nu$	${}^{17}\text{F}$	$5.82(1^{+0.19}_{-0.17}) \times 10^6$

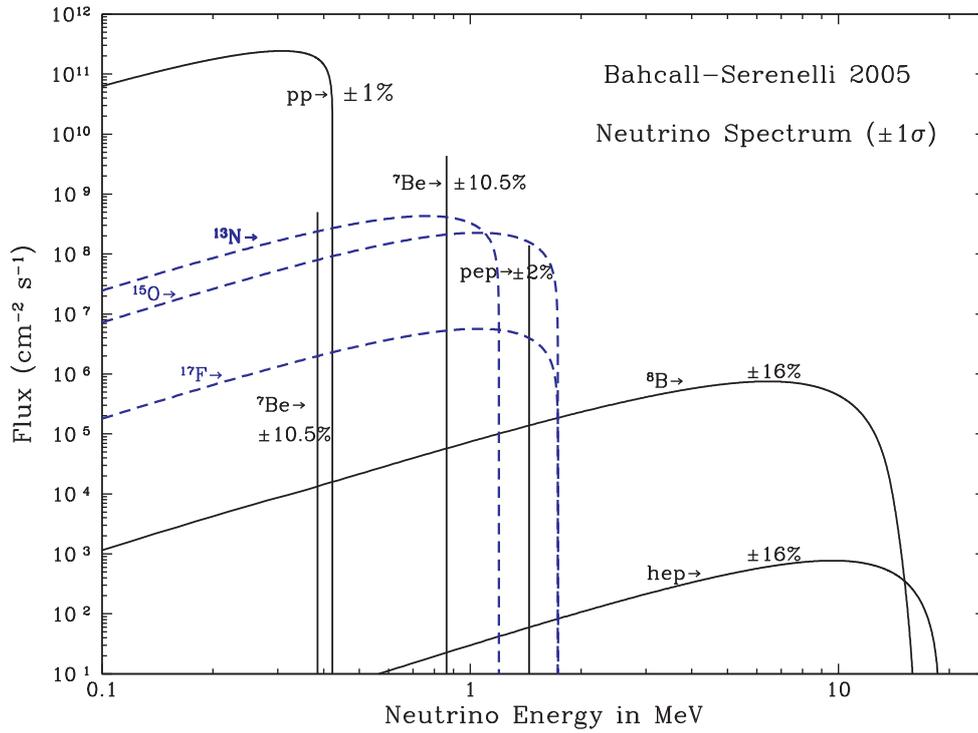


Table 2.1 & Figure 2.4: Table 2.1: Solar neutrino production reactions as predicted by the SSM. The first five rows are branches of the pp chain. Figure 2.4: Solar neutrino fluxes and energy spectra as predicted by the SSM. Images taken from [17].

this time was 8.5 ± 1.8 SNU [19].

In the 1980s and 1990s a series of next generation experiments were constructed: SAGE [20], GALLEX [21] and GNO [22], which used gallium to detect solar neutrinos at an energy threshold of 0.23 MeV through the reaction ${}^{71}\text{Ga} + \nu_e \rightarrow {}^{71}\text{Ge} + e^-$; solar neutrinos sensitive to the first stage of the pp chain. These experiments confirmed the deficit observed by Homestake, measuring the flux to be $70.8^{+5.3}_{-5.2}$ (stat.) $^{+3.7}_{-3.2}$ (syst.) SNU and 69.3 ± 5.5 (stat. + syst.) SNU (combined GALLEX+GNO) respectively, compared to an SSM prediction of 129^{+8}_{-6} SNU [23] for the gallium observation channel.

This deficit became known as the ‘solar neutrino anomaly’. It was postulated the discrepancy was due to substantially different temperatures and pressures within the Sun than were predicted by the SSM. However, advances in helioseismology [24] (understanding of wave propagation within the Sun) inferred an internal temperature for the Sun in agreement with the SSM and so showed these claims to be incorrect. Neutrino oscillations offered an alternative solution, with the transition from $\nu_e \leftrightarrow \nu_\mu/\nu_\tau$ resulting in the observed deficit. Further experiments were required to confirm this mechanism, as radio-chemical methods are unable to observe the corresponding ν_μ or ν_τ appearance signal. Experiments such as the Sudbury Neutrino Observatory (SNO) and the Kamioka Liquid Scintillator Anti-Neutrino Detector (KamLAND) were able to confirm the oscillation signal, and in the case of the former provide sensitivity to all three neutrino flavours.

Heavy Water and Liquid Scintillator Experiments

The Sudbury Neutrino Observatory [25] was a heavy water (D_2O) Cerenkov detector, located in Canada which took data from 1999 to 2006. The neutrino target consisted of 1000 tonnes of heavy water contained in a 6 m radius transparent acrylic vessel. The Cerenkov light resulting from interactions in the medium were detected by an array of 9456 photomultiplier tubes (PMTs).

The motivation for SNO was to provide direct evidence for neutrino flavour transitions, where the use of D_2O provided a unique sensitivity to the total solar neutrino flux. The 5 MeV interaction threshold made SNO sensitive to primarily ${}^8\text{B}$ solar interactions,

where three different rates could be measured:

$$\text{CC} : \nu_e + d \rightarrow p + p + e^-$$

$$\text{ES} : \nu_x + e^- \rightarrow \nu_x + e^-$$

$$\text{NC} : \nu_x + d \rightarrow p + n + \nu_x$$

where ν_x refers to any active neutrino flavour. The elastic scattering (ES) channel is sensitive to all neutrino flavours through the exchange of the Z boson, though the cross section is largest for electron neutrinos as they can also interact via the W boson. The NC channel is equally sensitive to all neutrino types, and so can determine the total solar neutrino flux in the presence of flavour changes. By combining measurements of the three different interaction types, it is possible to extract the ν_e flux and ν_μ/ν_τ fluxes independently.

The CC and ES interactions are detected directly through the production of Cerenkov light from the recoil electron. For the NC interaction in pure D_2O , the capture of the neutron by another deuteron produces a 6 MeV γ -ray, which scatters electrons to produce detectable Cerenkov light. This mode of data taking with pure heavy water was designated as Phase I. Over the lifetime of SNO, data taking was performed in three distinct phases, each enhancing the NC channel sensitivity. For Phase II, 2 tons of salt (NaCl) were added to the heavy water to improve the neutron capture efficiency, where neutron capture on ^{35}Cl has a higher cross section and improves γ -ray multiplicity. In Phase III, the salt was removed and an array of ^3He proportional counters were installed to measure the neutrons directly.

In Phase II, the following ^8B fluxes were measured [26]:

$$\phi_{\text{CC}} = 1.68_{-0.06}^{+0.06}(\text{stat.})_{-0.09}^{+0.08}(\text{syst.}) \times 10^6 \text{ cm}^{-2}\text{s}^{-1}$$

$$\phi_{\text{ES}} = 2.35_{-0.22}^{+0.22}(\text{stat.})_{-0.15}^{+0.15}(\text{syst.}) \times 10^6 \text{ cm}^{-2}\text{s}^{-1}$$

$$\phi_{\text{NC}} = 4.94_{-0.21}^{+0.21}(\text{stat.})_{-0.34}^{+0.38}(\text{syst.}) \times 10^6 \text{ cm}^{-2}\text{s}^{-1}$$

Figure 2.5 illustrates the results of this second phase. The non-electron flavour component of the flux ($\phi_{\mu\tau}$) is plotted against the electron neutrino flux (ϕ_e), with possible bands for the CC, ES and NC interactions plotted along with the SSM prediction and the Super-Kamiokande experiment prediction. All measurements are consistent with each other and the SSM, and meet at a common non-zero $\phi_{\mu\tau}$ value. This represents strong evidence for neutrino flavour transformations, where the non- ν_e flux component

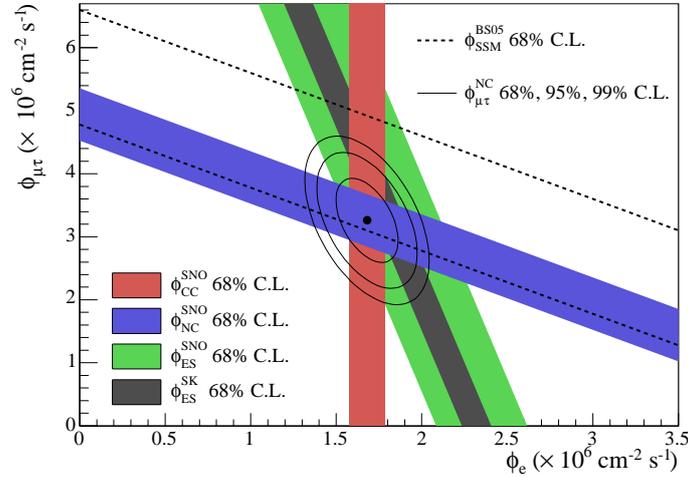


Figure 2.5: The results from Phase II of the SNO experiment. The flux of $\mu+\tau$ neutrinos versus the flux of electron neutrinos is plotted, with CC, ES and NC interaction bands from SNO and SKK. The intercepts of these bands with the axes represent the $\pm 1\sigma$ uncertainties. The flux predicted by the SSM is shown by the dotted lines. Taken from [26].

exists and when combined with the electron neutrino flux agrees with the overall flux predicted by the SSM.

To extract the best fit oscillation parameters in this case, the impact of neutrino interactions in the matter of the Sun must be considered, the so-called Mikheyev-Smirnov-Wolfenstein (MSW) effect [27, 28]. As neutrinos propagate through a medium rich in electrons, they undergo forward scattering. All flavours scatter via NC interactions, exchanging a Z boson. Additionally, electron neutrinos experience an additional potential from CC interactions with the solar electrons, which alters the transition probabilities between the neutrino flavours.

For Phase II of the SNO experiment, the overall best fit point from a two-flavour oscillation analysis is $\Delta m_{21}^2 = 5.0_{-1.8}^{+6.2} \times 10^{-5} \text{ eV}^2$, $\tan^2 2\theta_{12} = 0.45_{-0.10}^{+0.11}$ [26]. The Phase III flux results taken using ^3He proportional counters [29] are in agreement with previous results and the SSM.

The solar neutrino oscillation signal was confirmed by the Kamioka Liquid scintillator Anti-Neutrino Detector (KamLAND) [30]; located in the Kamioka Observatory, Japan where data-taking began in 2002. It utilises 1 kton of highly purified liquid scintillator suspended in a 13 m diameter nylon-based balloon. The scintillator is surrounded by an

array of 1879 PMTs measuring the flux of electron anti-neutrinos from the 55 surrounding nuclear reactors, each an isotropic neutrino source. These reactors are distributed at distances ranging from $L \simeq 100 - 700$ km from the KamLAND detector, and their operational information is used to derive the predicted anti-neutrino flux in the absence of oscillations.

The $\bar{\nu}_e$ are detected via the inverse β -decay channel: $\bar{\nu}_e + p \rightarrow e^+ + n$, which has a 1.8 MeV energy threshold. The prompt scintillation light of the positron (including kinetic and annihilation energy) provides an estimate of the neutrino energy, and when combined with the subsequent 2.2 MeV γ -ray resulting from neutron capture (approximately 200 μ s later) on a free proton tags the interaction and significantly reduces the background. Figures 2.6 to 2.7 summarise the current published KamLAND neutrino oscillation results [30], based on data collected between 2002 and 2007. A significant deficit in the anti-neutrino flux is observed, with $2179 \pm 89(\text{syst.})$ events expected and 1609 observed.

Figure 2.7 shows the ratio of the KamLAND background and geo-neutrino subtracted spectrum to the unoscillated prediction, plotted as a function of L_0/E , where $L_0 = 180$ km (the flux weighted effective baseline over the surrounding reactors). The oscillatory signal is clearly confirmed in the plot structure, with two cycles of the periodic variation visible.

A maximum likelihood fit to a two-flavour neutrino oscillation model ($\theta_{13} = 0$) was performed, accounting for time variations in reactor flux and simultaneously fitting the geo-neutrino contribution and incorporating matter effects within the Earth. The allowed contours for both KamLAND data alone, and combined KamLAND and solar neutrino experiment data are shown in Figure 2.6, with the latter dominated by the SNO contribution. The best fit oscillation parameters are $\Delta m_{21}^2 = 7.58_{-0.13}^{+0.14}(\text{stat.})_{-0.15}^{+0.15}(\text{syst.}) \times 10^{-5} \text{ eV}^2$ and $\tan^2 2\theta_{12} = 0.56_{-0.07}^{+0.10}(\text{stat.})_{-0.06}^{+0.10}(\text{syst.})$; the currently most precise measurement of Δm_{21}^2 [17].

2.3.2 Short Range Oscillations

Atmospheric neutrinos are produced as the result of collisions between high energy cosmic rays with nuclei in the upper atmosphere, resulting in a cascade of secondary pions and kaons, which decay to produce electrons and muons and their associated neutrino flavours. The typical energies of these interactions result in a wide range of

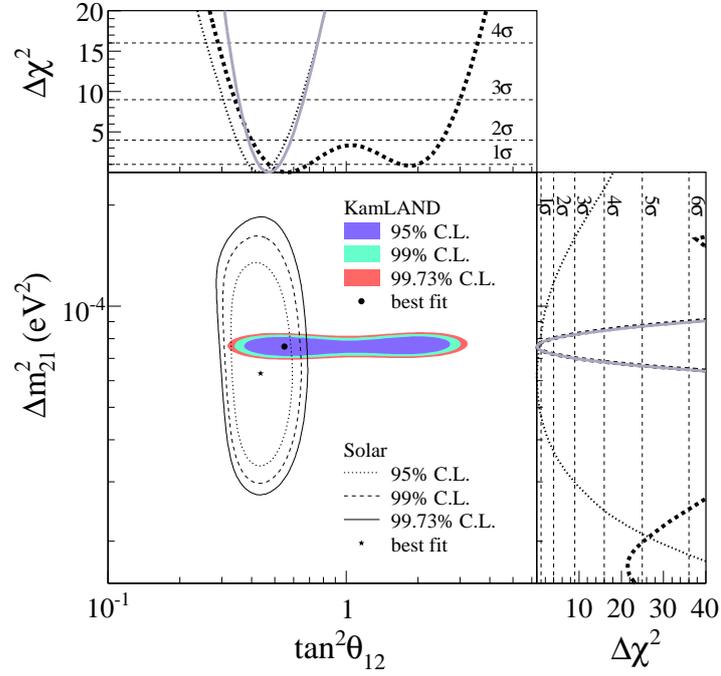


Figure 2.6: Allowed region for neutrino oscillation parameters, for a KamLAND only analysis and a combined analysis of KamLAND and solar data. The $\Delta \chi^2$ profiles are shown in the sidepanels, for KamLAND (dashed), solar (dotted) and combined (solid). Taken from [30].

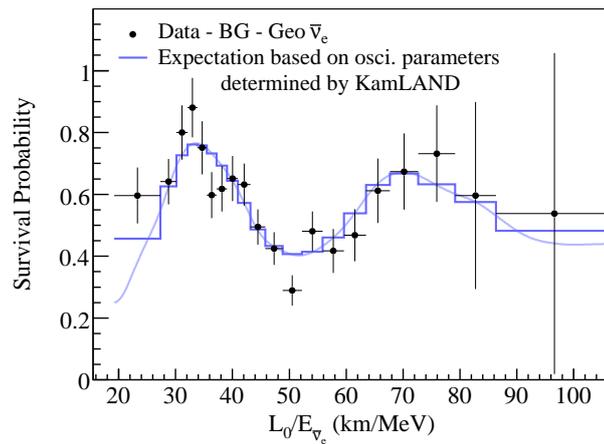


Figure 2.7: Ratio of the KamLAND background and geo-neutrino subtracted spectrum to the unoscillated prediction, as a function of L_0/E ; where L_0 is the flux weighted effective baseline ($L_0 = 180$ km). The error bars are statistical only. Taken from [30].

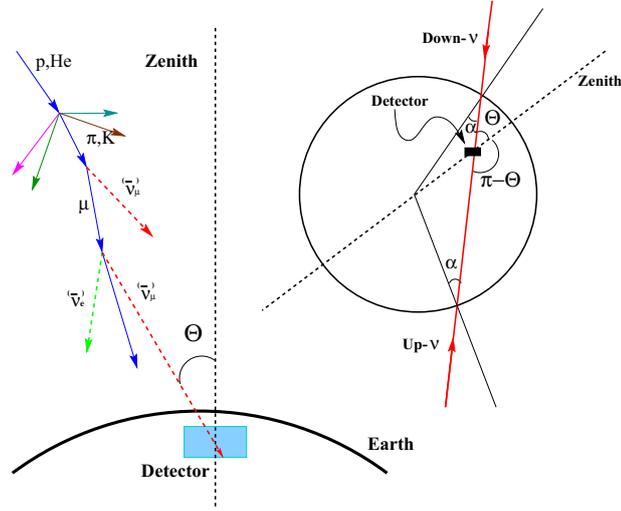
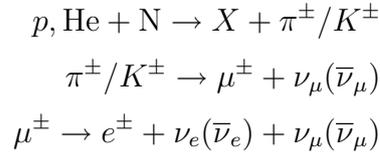


Figure 2.8: Interactions of cosmic rays with nuclei in the upper atmosphere to produce atmospheric neutrinos. The definition of the zenith angle Θ is shown. Taken from [32].

neutrino energies, with $E \simeq 0.1 - 100 \text{ GeV}$, with most neutrinos produced at lower energy. The processes are summarised in Figure 2.8 and as follows:



Considering the ratio of muon neutrinos to electron neutrinos in the decay chains above, the predicted ratio of fluxes is $R = N(\nu_\mu + \bar{\nu}_\mu) / N(\nu_e + \bar{\nu}_e) \approx 2$, with an uncertainty of $\sim 5\%$ [31]. This ratio is energy dependent, and should exhibit a gradual rise above 5 GeV due to the increasing likelihood the intermediate stage muons will reach the Earth's surface before decaying.

Experimentally, measurements of the atmospheric neutrino fluxes are usually presented as a “ratio of ratios”, $R' \equiv R_{Data} / R_{MC}$ between the observed and predicted neutrino flavour rates, using the charged lepton resulting from CC interactions in the atmosphere to tag the flavour of the incident neutrino. The first measurements of the atmospheric neutrino flux were performed in the 1980s, by the water Cerenkov experiments IMB [33] and Kamiokande [34]. The charged leptons emitted in CC interactions were detected via their Cerenkov rings, with flavour separation possible due to ring topology. Muons produce clean, sharp rings, where the electromagnetic showers

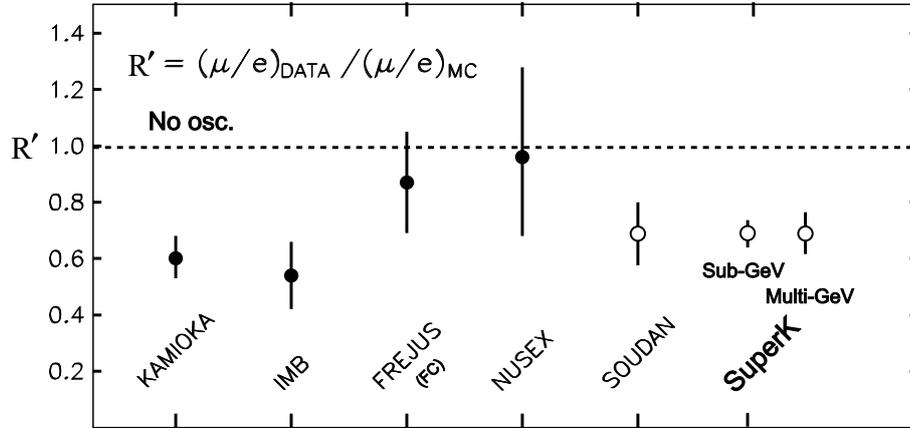


Figure 2.9: Summary of measurements of the double ratio (R') from a range of atmospheric neutrino experiments. Taken from [35].

induced by the less massive electrons create ‘fuzzy’ edged rings. Both experiments observed a deficit in the ν_μ flux relative to the expectation, measuring double ratios of $R' = 0.54^{+0.05}_{-0.05}(\text{stat.})^{+0.11}_{-0.11}(\text{syst.})$ and $R' = 0.60^{+0.07}_{-0.06}(\text{stat.})^{+0.5}_{-0.5}(\text{syst.})$ respectively. A summary of R' values measured by a range of experiments is given in Figure 2.9; the deficit shown became known as the ‘atmospheric neutrino anomaly’.

Neutrino oscillations provide an explanation for the ν_μ neutrino deficit, permitting transitions from ν_μ to other neutrino flavours. A method of detecting an oscillatory signal is to measure the distribution of the muon neutrino flux by zenith angle Θ , defined in Figure 2.8. Neutrinos from all parts of the atmosphere can be detected due to the very low level of attenuation of neutrinos within the Earth, giving a range of possible path lengths L from 15 – 500 km for downward going neutrinos, and from 500 – 13,000 km for upward going neutrinos. If neutrino oscillations occur over any of these length scales, deficits in the flux will be observed at angles where the path length corresponds to a high transition probability.

Super-Kamiokande

Super-Kamiokande (SK) [36], a 50 kton cylindrical water Cerenkov detector, began data taking in 1996 and provided precise measurements of the atmospheric neutrino fluxes. It is instrumented with 11146 PMTs facing the inner detector; an 18 kton vessel of highly purified water. The outer veto volume is instrumented by 1885 outward facing PMTs, which serve to identify muons entering the detector and shield the inner fiducial volume from γ -rays and neutrons from the surrounding rock.

SK events are sorted into a number of classes: Fully Contained (FC) if the neutrino interaction vertex and all resulting particles are confined to the inner detector, Partially Contained (PC) if the vertex is within the fiducial volume but one or more emitted particles exit, or as neutrino-induced Up-Going Muons (UPMU) where neutrinos interact in the rock beneath the detector and the resultant muons are detected. While muons travelling in the downward direction cannot be distinguished from cosmic rays, muons travelling in an upward direction are neutrino induced. FC events with large amounts of observed Cerenkov light are further classified as ‘multi-GeV’, where the emitted lepton closely follows the path of the incident neutrino, so the neutrino propagation distance can be accurately determined.

An analysis of atmospheric neutrino data collected between April 1996 and November 2001 was published [37] in 2005. Figure 2.10 shows the zenith angle distributions for different classes of selected neutrino events at different event energies. Also plotted are the expected distributions in the absence of oscillations, and the best fit distributions for $\nu_\mu \leftrightarrow \nu_\tau$ oscillations. For muon-like neutrinos there is a clear zenith-angle dependent deficit, while electron-like samples are consistent with the expectations. The interpretation of these results is that muon neutrinos are undergoing flavour transitions to tau neutrinos with a large mixing angle, as the multi-GeV flux for muon-like neutrinos drops to almost half the non-oscillatory expectation. Figure 2.11 shows the best fit and confidence limits extracted by fitting to these distributions. The best fit point is: $\Delta m_{23}^2 = 2.1 \times 10^{-3} \text{ eV}^2$, $\sin^2 2\theta_{23} = 1.0$. The 90% confidence limit on $\sin^2 2\theta_{23} > 0.92$ is currently the world’s best limit on this parameter.

K2K

The atmospheric neutrino $\nu_\mu \leftrightarrow \nu_\tau$ oscillation signal was confirmed by the KEK-to-Kamioka (K2K) experiment [38], which published its final oscillation result in 2006. K2K utilised an accelerator produced neutrino beam to probe the atmospheric neutrino energy region. The beam was produced by firing 12 GeV protons produced in the KEK synchrotron onto an aluminium target, producing a spray of pions and kaons. The positively charged particles were focused by a pair of magnetic horns into a 200 m decay pipe, where they decayed to produce a beam of 98% ν_μ with a mean energy of 1.3 GeV.

The initial spectrum was measured 300 m from the production point by a 1 kton water Cerenkov ‘near’ detector, and subsequently travelled 250 km and was measured by the Super-Kamiokande ‘far’ detector. The expected energy spectrum for different models of neutrino propagation were produced by extrapolating the neutrino flux measured at the

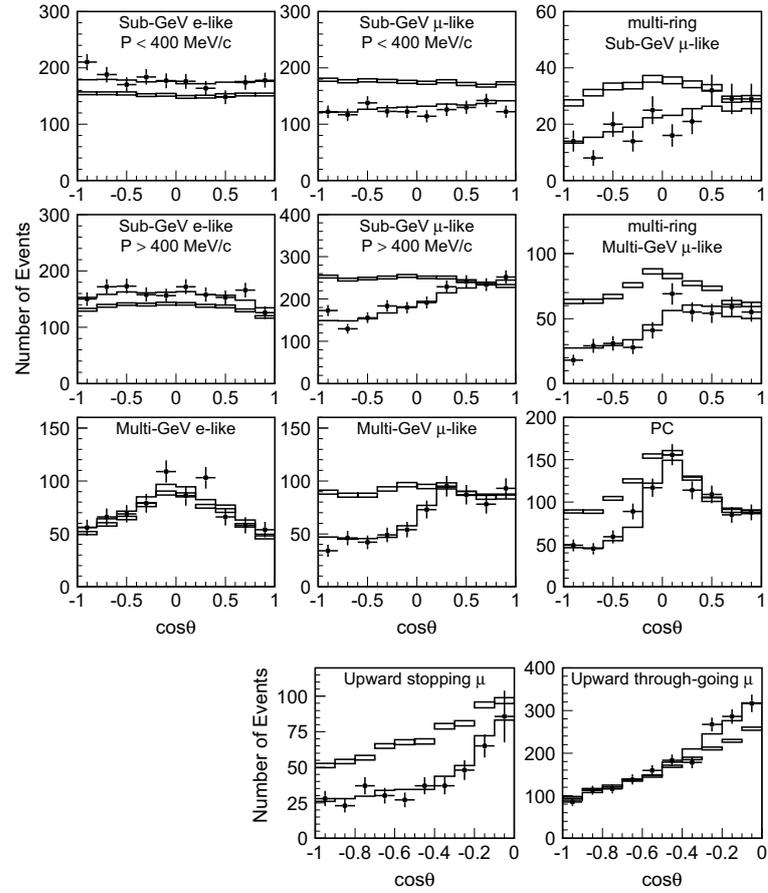


Figure 2.10: Zenith angle distributions for different classes of Super-Kamiokande events. Data are shown by points, unoscillated predictions by boxes and the oscillated predictions by lines, for $\nu_\mu \leftrightarrow \nu_\tau$ oscillations with the best fit values of $\sin^2 2\theta = 1.0$ and $\Delta m^2 = 2.1 \times 10^{-3} \text{ eV}^2$. Taken from [37].

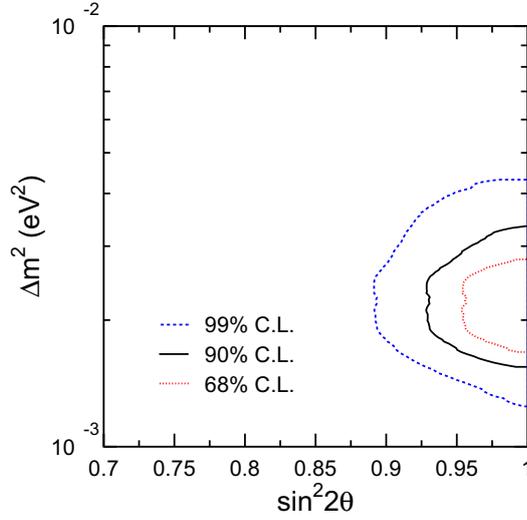


Figure 2.11: Allowed oscillation parameters for $\nu_\mu \leftrightarrow \nu_\tau$ oscillations from Super-Kamiokande data. Three contours correspond to the 99% (dashed), 90% (solid) and 68% (dotted) confidence limits. Taken from [37].

near detector; a reduction in systematic uncertainties is achieved by using two detectors with similar constructions and extrapolating between them. The expected signal for neutrino disappearance, in addition to a deficit of muon neutrinos is a distorted energy spectrum from which best fit parameters can be extracted.

K2K published their final oscillation results from a dataset collected between 1999 and 2004 [38]. A deficit of muon neutrinos was measured at Super-Kamiokande, with an unoscillated prediction of $158_{-8.6}^{+9.2}$ events predicted in the fiducial volume and 112 events observed. Spectral distortion was also observed consistent with neutrino oscillations, for the 58 events for which a neutrino energy was reconstructed. These 58 events were selected as quasi-elastic CC interactions, with single ring muon-like properties. Figure 2.12 shows the observed far detector spectrum, with the prediction in the absence of oscillations and the best fit spectrum. A clear spectral distortion is observed, and interpreted as the effect of neutrino oscillations. Figure 2.13 shows the confidence limit contours obtained in the two-flavour oscillation analysis, which are consistent with the previous Super-Kamiokande analysis. The best fit point (constrained to the physical region) is: $\Delta m_{23}^2 = 2.8 \times 10^{-3} \text{eV}^2$, $\sin^2 2\theta_{23} = 1.0$.

The MINOS Experiment was primarily designed to identify neutrino oscillations at the atmospheric Δm^2 scale, and is described in detail in the following Chapters. The latest world-best measurements of the mass splitting and mixing angle were published

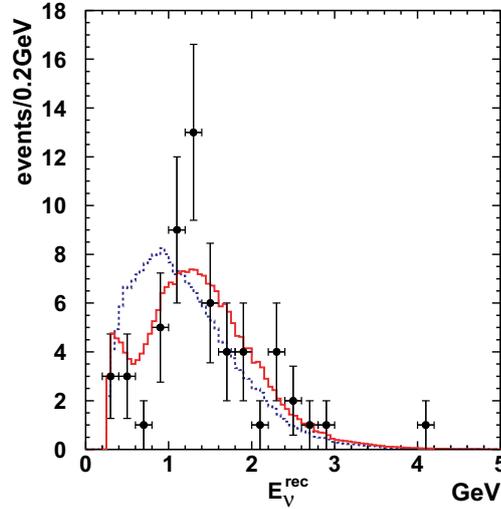


Figure 2.12: The reconstructed neutrino energy spectrum for the muon-like 1-ring sample. The points with error bars are data, the solid line the best fit spectrum with oscillations at the best fit point of $\Delta m_{23}^2 = 2.8 \times 10^{-3} \text{ eV}^2$, $\sin^2 2\theta_{23} = 1.0$ and the dashed line the prediction for no disappearance. Taken from [38].

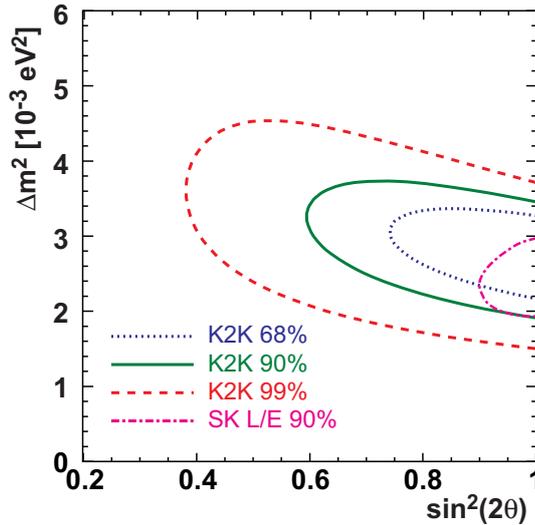


Figure 2.13: Allowed oscillation parameters for $\nu_\mu \leftrightarrow \nu_\tau$ oscillations from the K2K two-flavour oscillation analysis. Results from a Super-Kamiokande atmospheric analysis are also displayed. Taken from [38].

in 2010 [39], with $\Delta m_{23}^2 = 2.32_{-0.08}^{+0.12} \times 10^{-3} \text{eV}^2$ and $\sin^2(2\theta_{23}) > 0.90$ to 90% confidence. Recently starting and future off-axis beam experiments such as T2K [40] (a successor to K2K) and NO ν A [41] will provide tighter constraints on the atmospheric neutrino oscillation parameters. T2K presented preliminary results in July 2011 [42], with $2.1 \times 10^{-3} < \Delta m_{23}^2 < 3.1 \times 10^{-3} \text{eV}^2$ and $\sin^2(2\theta_{23}) > 0.85$ to 90% confidence for 31 selected events. Data taking is expected to resume in December 2011.

2.3.3 Sub-Dominant Short Range Oscillations

In addition to solar and atmospheric neutrino oscillations, the flavour change $\nu_e \rightarrow \nu_\mu/\nu_\tau$ controlled by the mixing angle θ_{13} is also permitted, coupling the two dominant oscillation sectors. Measurements can be made searching for the flavour change over distances associated with atmospheric neutrino oscillations. The current generation of experiments have only provided an upper limit on this parameter, due to its small value compared to the other mixing angles. The current best limit until very recently was set by the CHOOZ experiment [43], with new limits published in Summer 2011 by MINOS [44] and T2K [45].

The CHOOZ experiment searched for the subdominant mode $\bar{\nu}_e \leftrightarrow \bar{\nu}_\mu/\bar{\nu}_\tau$ in reactor produced anti-neutrinos from the CHOOZ power plant, located in Northern France. The isotropic flux of electron anti-neutrinos were produced by two high-pressure water reactors, with a mean energy of 3 MeV and an intensity known to better than 2%. The detection medium consisted of 5 tons of liquid scintillator, located 1 km from the neutrino source which detected neutrinos via the inverse beta decay reaction $\bar{\nu}_e + p \rightarrow e^+ + n$ (as in KamLAND). The signature consisted of the prompt e^+ signal, followed by photon emissions from neutron capture. The scintillator was loaded with Gadolinium to enhance the efficiency of neutron capture and the energy of the resultant photons.

The signal for neutrino oscillations published by CHOOZ [43, 46] was a distorted $\bar{\nu}_e$ spectrum and a deficit of events. Figure 2.14 shows the prompt positron spectrum, where no deficit of events or spectral distortion is observed. Figure 2.15 shows the confidence limits on Δm_{23}^2 and θ_{13} for a two-flavour oscillation analysis; subdominant oscillations were excluded at 90% confidence for $\Delta m_{23}^2 \geq 8 \times 10^{-4} \text{eV}^2$ and $\sin^2 2\theta_{13} \geq 0.10$ for large Δm_{23}^2 values.

Though not designed for the identification of ν_e interactions, MINOS also has some

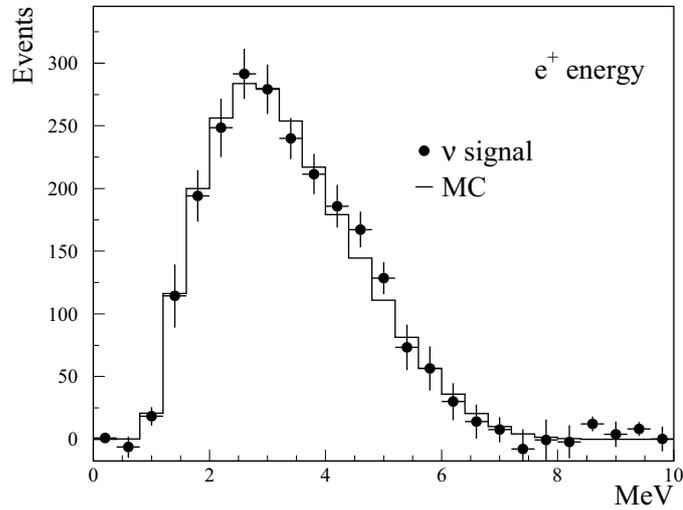


Figure 2.14: CHOOZ prompt e^+ spectrum. The points with error bars are data and the solid line the prediction for no disappearance. Taken from [43].

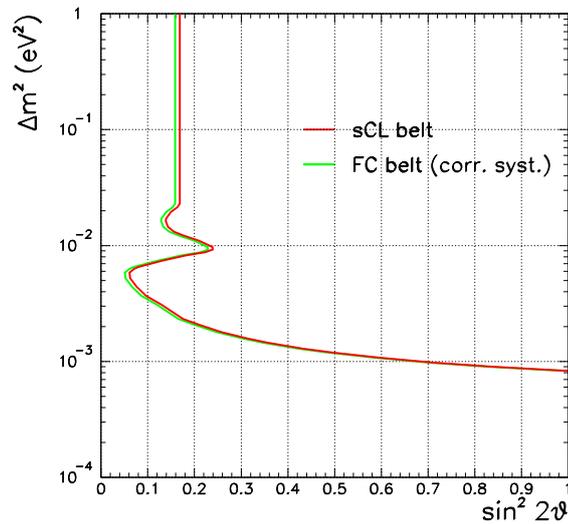


Figure 2.15: Allowed parameter region for θ_{13} and Δm_{23}^2 from CHOOZ data. The area to the right of the line is excluded to 90% confidence, and contours from both Feldman-Cousins and strong Confidence Limit (sCL) methods are shown. Taken from [43].

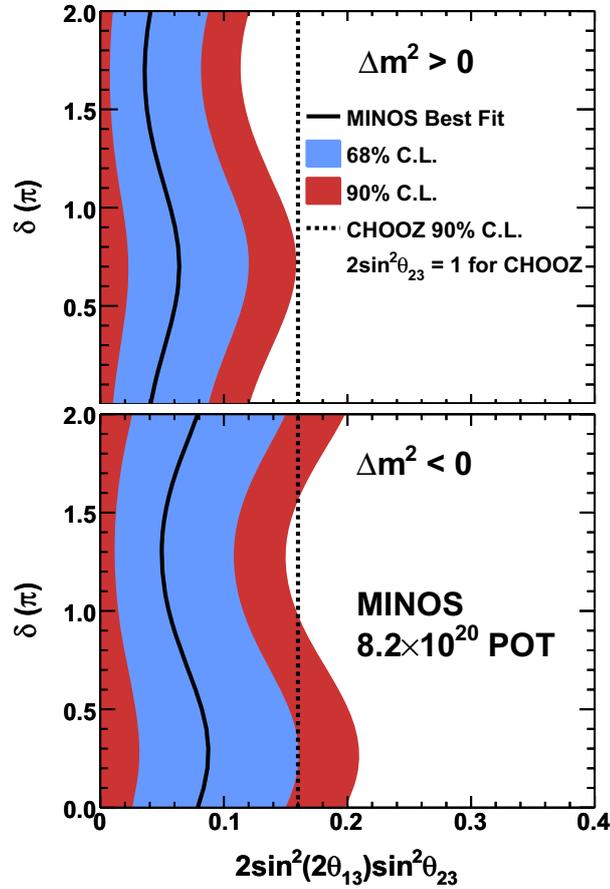


Figure 2.16: Confidence limits for $2 \sin^2(\theta_{23}) \sin^2(2\theta_{13})$ as a function of δ , for the normal (upper panel) and inverted (lower panel) mass hierarchies. The solid line indicates the MINOS best fit values. The dashed line indicates the CHOOZ upper 90% confidence limit evaluated at $\theta_{23} = \pi/4$ and $\Delta m_{23}^2 = 2.32 \times 10^{-3} \text{eV}^2$.

sensitivity to θ_{13} . The oscillation probability in MINOS to first order is as follows:

$$P(\nu_\mu \rightarrow \nu_e) \approx \sin^2(\theta_{23}) \sin^2(2\theta_{13}) \sin^2(1.27 \Delta m_{32}^2 L/E) \quad (2.28)$$

Unlike in CHOOZ, the oscillation probability is also a function of δ and the mass hierarchy, due to matter effects and possible leptonic \mathcal{CP} violation. The results published most recently in [44] find that $2 \sin^2(\theta_{23}) \sin^2(2\theta_{13}) < 0.12(0.20)$ to 90% confidence, for the normal (inverted) mass hierarchy and $\delta = 0$. The $\theta_{13} = 0$ hypothesis is disfavoured at 89% confidence. The allowed ranges and best fits are shown in Figure 2.16, evaluated at the current atmospheric best fit parameters of $\theta_{23} = \pi/4$ and $\Delta m_{23}^2 = 2.32 \times 10^{-3} \text{eV}^2$.

The T2K Experiment [40] utilises a muon neutrino beam produced at the JPARC

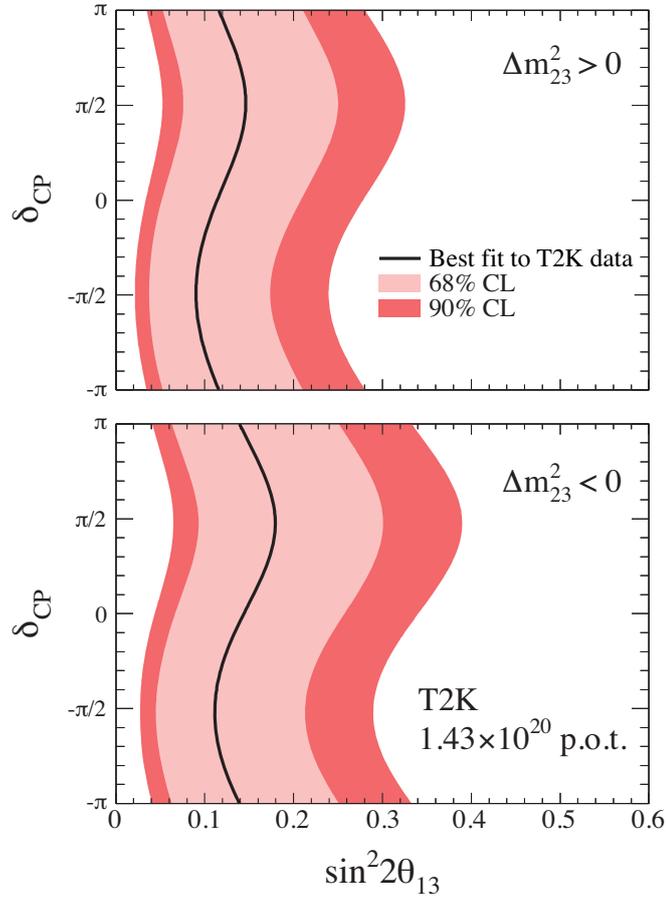


Figure 2.17: Confidence limits for $\sin^2(2\theta_{13})$ as a function of δ , for the normal (upper panel) and inverted (lower panel) mass hierarchies.

facility, which currently fires 30 GeV protons at a graphite target to produce pions and kaons which are sign selected and focused by three magnetic horns into a 96 m decay tunnel. The beam is directed 2.5° off-axis to the Super Kamiokande water Cerenkov Far Detector at a distance of $L = 295$ km. The Near Detector complex is sited 280 m downstream of the target, and consists of an on-axis detector used for beam monitoring, and an off-axis detector to measure the initial neutrino spectrum and extrapolate to the Far Detector, as was done in K2K (and is done in MINOS).

The first electron neutrino appearance results were published in [45], and observed an excess of six events over an expectation of $1.5 \pm 0.3(\text{syst.})$, corresponding to 2.5σ significance. The limits for a three-flavour analysis are shown in Figure 2.17, and are consistent with $0.03(0.04) < \sin^2(2\theta_{13}) < 0.28(0.34)$ for the \mathcal{CP} violating phase $\delta = 0$ and normal (inverted) hierarchy.

Future reactor experiments, such as Double-CHOOZ (the successor to CHOOZ) [47], Reno [48] and Daya Bay [49] aim to measure or constrain further the value of θ_{13} via the oscillation of $\bar{\nu}_e$. Double-CHOOZ aims to achieve a sensitivity of the order of $\sin^2 2\theta_{13} \leq 0.03$ after approximately 3 years of running with both near and far detectors operational. The NO ν A experiment [41] (a successor to MINOS) is an off-axis two detector experiment designed to observe electron neutrino appearance in the Fermilab NuMI beam, and is currently under construction.

2.4 Summary and Current Status

Over the last 50 years, clear experimental evidence has been observed for mixing between neutrino flavours. The solar neutrino anomaly has been addressed by results from experiments such as SNO and KamLAND. Super-Kamiokande and K2K have strongly suggested the atmospheric neutrino anomaly can be explained by $\nu_\mu \leftrightarrow \nu_\tau$ oscillations. The solar and atmospheric regimes are coupled by the mixing angle θ_{13} . Current best fits and limits on the neutrino mixing parameters with the corresponding experiments follow, taken from [17, 39, 44, 45]:

- $\sin^2(2\theta_{12}) = 0.87 \pm 0.03$, combined solar and KamLAND data.
- $\Delta m_{12}^2 = 7.59_{-0.21}^{+0.19} \times 10^{-5} \text{eV}^2$, combined solar and KamLAND data.
- $\sin^2(2\theta_{23}) > 0.92$ to 90% confidence, Super-Kamiokande.
- $\Delta m_{23}^2 = 2.32_{-0.08}^{+0.12} \times 10^{-3} \text{eV}^2$, MINOS (work presented in this thesis).
- $2 \sin^2(\theta_{23}) \sin^2(2\theta_{13}) < 0.12$ to 90% confidence, MINOS.
 $0.03 < \sin^2(2\theta_{13})$ to 90% confidence, T2K.

The MINOS experiment, described in detail in Chapter 3, aims to provide high precision measurements of Δm_{23}^2 and $\sin^2(2\theta_{23})$ in the region indicated by Super-Kamiokande. This thesis describes an analysis to provide updated and high precision values of the atmospheric neutrino mixing parameters, with a dataset taken between May 2005 and June 2009.

Chapter 3

The MINOS Experiment

The Main Injector Neutrino Oscillation Search (MINOS) is a two detector, long baseline neutrino oscillation experiment. Its aim is to confirm the existence of neutrino oscillations and to provide a precision measurement of the mixing parameters in the region indicated by previous atmospheric neutrino experiments.

The neutrino beam is generated at the Fermi National Accelerator Laboratory (Fermilab) in Batavia, Illinois by the Neutrinos at the Main Injector (NuMI) facility. The neutrino beam is measured at two different distances from production by two functionally identical detectors, as illustrated in Figure 3.1. The Near Detector is situated close to the beam source to sample the beam before neutrino oscillations are visible, while the Far Detector is located 735 km along the beamline in the Soudan Underground laboratory, Minnesota.

The neutrino energy spectrum observed at the Far Detector is compared to that predicted in the absence of oscillations, to search for the energy dependent neutrino deficit associated with an oscillation signal. To reduce the impact of systematic effects such as detector efficiencies, uncertainties in the neutrino flux and cross sections the expected Far Detector prediction is generated by extrapolating the measured Near Detector spectrum to predict the Far Detector spectrum.

In this Chapter details of the design and operation of the neutrino beam and the Far Detector are given. The methods of calibration and the calibration detector are also discussed. Finally, the simulation of the beam line and event reconstruction methods are summarised, with details of the dataset analysed in this thesis. The Near Detector is described in Chapter 4, in the context of the work performed as part of this thesis to



Figure 3.1: The MINOS neutrino beamline. Neutrinos are produced by the NuMI facility in Fermilab, Illinois and travel 735 km through the Earth to the Soudan Underground Laboratory in Soudan, Minnesota.

establish the quality of the dataset used for the analysis.

3.1 The NuMI Beam

A diagram of the NuMI beamline [51] is shown in Figure 3.2. To produce the neutrino beam, protons are accelerated to 120 GeV in the Main Injector (MI) accelerator, and

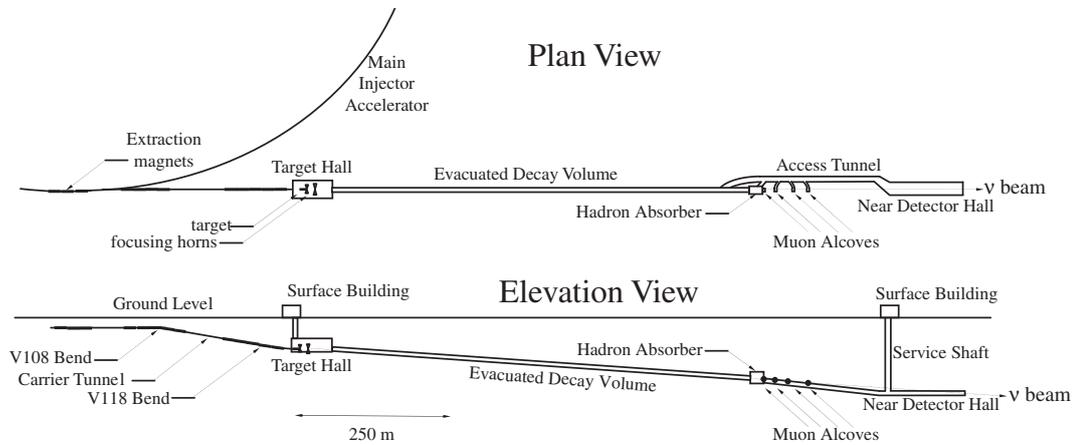


Figure 3.2: Plan and elevation of the NuMI beam facility, Fermilab. A beam of protons is incident on a graphite target, producing pions and kaons that are focused by two magnetic horns into the decay volume. The subsequent kaon and pion decays produce the neutrino beam. The evacuated decay pipe was filled with helium in November 2007, in order to prevent single point failure at a weakness in the decay volume. Taken from [50].

extracted in beam spills of approximately $10 \mu\text{s}$ in duration, by three kicker magnets. Spills are spaced approximately 2 s apart, and typically contain 2.1×10^{13} protons. The proton beam is bent downwards by 58 mrad to account for the curvature of the Earth to point in the direction of the Far Detector, and strikes a graphite target 350 m from the point of extraction.

The target is a rectangular graphite rod, constructed in 47 longitudinal segments with dimensions of 6.4 mm in width, 15 mm in height and 940 mm in length. The total thickness is designed to correspond to 2.4 incident proton interaction lengths. A schematic of the target is given in Figure 3.3. Equipment is stationed upstream of the target to protect the target assembly and monitor the beam. The baffle, a hollow graphite cylinder located as shown in Figure 3.4 protects the target and horns if the proton beam is misaligned. A toroid which measures the current of protons passing through it by magnetic induction provides a measurement of the beam intensity in units of ‘Protons-on-Target’ (PoT).

The incident protons interact within the target to produce pions and kaons with a wide range of longitudinal and transverse momenta. The narrow and long shape of the target allows the resulting particles to easily escape, and minimises interactions of the pions and kaons in the target. A pair of magnetic horns are then used to focus the

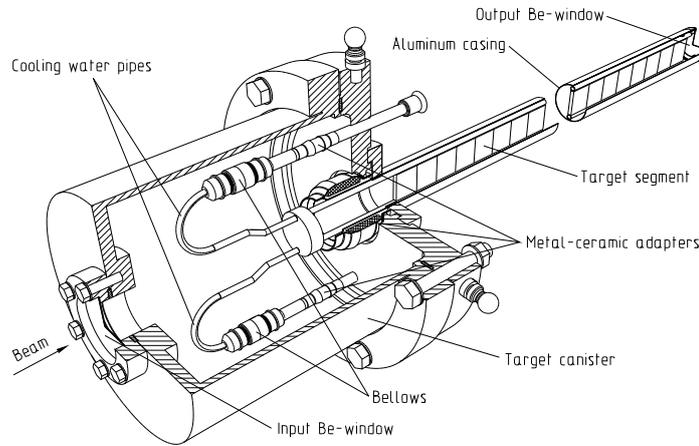


Figure 3.3: Schematic of the NuMI target. The target is continuously water cooled by pipes running through the assembly, and consists of 47 graphite segments. Taken from [52].

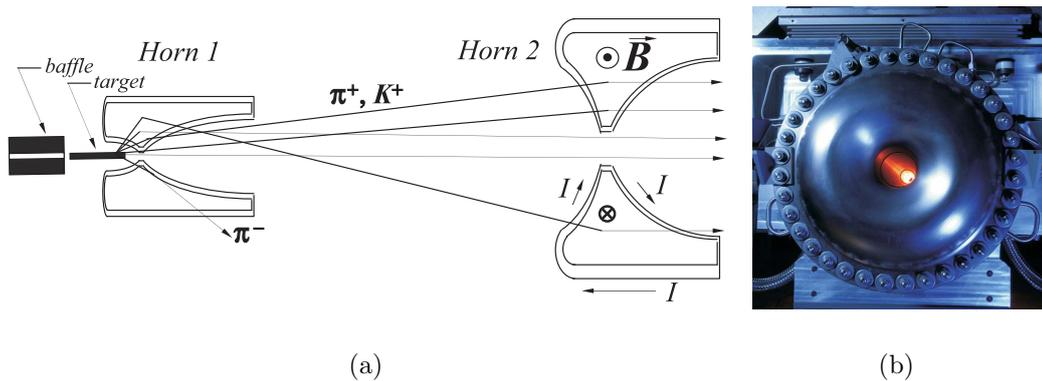


Figure 3.4: (a) Hadrons produced by the NuMI target are focused by a pair of magnetic horns spaced 10 m apart, capable of focusing neutrinos or anti-neutrinos. The baffle protects the horns from exposure to misdirected proton pulses. Taken from [50]. (b) View through one of the focusing horns. Taken from [53].

secondary particles (predominantly pions and kaons) into a beam.

Figure 3.4 shows the relative target horn positioning, and the focusing of the hadrons into a collimated beam. The horns are sited 10 m apart, and consist of parabolic-shaped inner conductors with cylindrical shaped outer conductors. During a beam spill, a maximum pulsed current of 200 kA produces a toroidal magnetic field which acts as a lens. The magnetic field polarity can be changed to focus π^+ and K^+ (for ν_μ) and

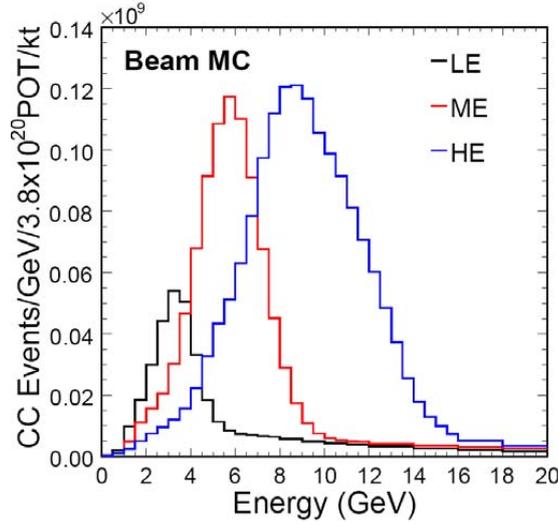


Figure 3.5: Expected rate of ν_μ interactions in the Near Detector for three different target/horn configurations; LE: ‘Low Energy’ (LE10/185kA), target 10 cm upstream of the first horn, 185 kA horn current; ME: ‘Medium Energy’ (LE100/200kA); HE: ‘High Energy’ (LE250/200kA). Taken from [54].

defocus π^- and K^- in the “forward” configuration, or focus π^- and K^- (for $\bar{\nu}_\mu$) and defocus π^+ and K^+ in the “reverse” configuration. For the analysis presented in this thesis, the forward horn current is used. In this configuration, some π^- and K^- will pass directly through the centre of the focussing horns and contribute a $\bar{\nu}_\mu$ background component in the ν_μ beam.

The effective focal length of the horns is proportional to the incident particle momentum; by adjusting the relative positions of the target and horns the momentum of the pions and kaons can be controlled, and hence the neutrino energy spectrum can be modified. By moving the target closer to the magnetic horns the secondary particles and the neutrino spectrum are both shifted to lower energies. The target has a maximum longitudinal range of 2.5 m, permitting a variety of neutrino energy spectra to be obtained. Figure 3.5 shows the expected ν_μ CC interaction rates and energies for three beam configurations.

The majority of the data used in this thesis was taken in the Low Energy (LE) configuration, where the target is placed 10 cm upstream from the horns, which are pulsed at 185 kA. This results in a neutrino energy spectrum peaked around 3 GeV. Results from Super-Kamiokande suggest increasing the number of neutrino interactions at lower energies will give a greater sensitivity where the oscillation signal is maximal.

After focusing, the hadron beam passes into a decay pipe, initially evacuated. It was subsequently filled with helium in November 2007 to prevent the failure of the decay volume due to corrosion of the aluminium window at the upstream end. The decay pipe is 2 m in diameter and 675 m in length, the approximate decay length of a 10 GeV pion. The pions and kaons decay into neutrinos and muons as follows:

$$\begin{aligned}\pi^\pm &\rightarrow \mu^\pm + \nu_\mu (\bar{\nu}_\mu) \\ K^\pm &\rightarrow \mu^\pm + \nu_\mu (\bar{\nu}_\mu)\end{aligned}\tag{3.1}$$

At the end of the decay pipe, any remaining hadrons are removed by an absorber, constructed of concrete and steel blocks faced by a water-cooled steel and aluminium core. After the absorber, the beam passes through 240 m of rock to remove decay pipe muons before reaching the Near Detector.

The beam composition at the Near Detector is simulated using FLUKA05 [55] Monte Carlo and a GEANT4 [56] simulation of the NuMI beamline; Section 3.3.2 discusses this in more detail. The predicted spectrum for the LE10/185kA beam configuration (see Figure 3.5) is composed of 92.9% ν_μ , 5.8% $\bar{\nu}_\mu$ and 1.3% ν_e [50]. The electron and anti-neutrino components are mainly due to $K^+ \rightarrow \pi^0 e^+ \nu_e$ and $\mu^+ \rightarrow e^+ \bar{\nu}_\mu \nu_e$ decays, in addition to decays of any negatively charged mesons passing through the focusing horns.

3.2 The MINOS Detectors

The two MINOS detectors are designed to be as similar as possible in terms of their responses to neutrino interactions. This enables a reduction in the impact of the systematic uncertainties in the neutrino flux, cross sections and detector acceptance on oscillation measurements. The detectors are magnetised steel-scintillator sampling calorimeters, constructed of alternate planes of plastic scintillator strips and steel plates, capable of observing muon and electron neutrino interactions with a visible energy larger than about 500 MeV. The detectors are magnetised to contain the muons, and reconstruct the muon momentum and charge, thus distinguishing between neutrino and anti-neutrino interactions.

The definitive description of the detectors, their construction and readout systems is given in [57]. In this section, the design and performance of the Far Detector is detailed. Additionally, the third ‘calibration’ detector, used to test detector response, is briefly

reviewed. The Near Detector is discussed in Chapter 4.

3.2.1 Steel and Scintillator

The active medium in the MINOS detectors is comprised of strips of plastic scintillator, made of extruded polystyrene 4.1 cm thick and 1 cm wide. The scintillator records the passage of ionising particles, and is used for both tracking and calorimetry purposes. The scintillator strips are grouped, placed side by side and sandwiched between aluminium sheets to form light-tight modules. Eight modules are then combined to form a single scintillator plane, attached to 2.54 cm thick planes of steel.

The steel planes are used as nuclear targets for neutrino interactions, passively absorbing the products of these interactions and providing the structure of the detectors. The steel was made in a series of batches, where subsets of each batch were used in each detector, to ensure similar magnetic and density properties. The steel density was measured to be $7.85 \pm 0.03 \text{ g/cm}^3$, with no systematic differences between the two detectors.

The detectors consist of a series of these scintillator/steel planes, aligned vertically with a 5.94 cm spacing. Three-dimensional reconstruction of events is rendered possible due to the fact the strips in successive planes are oriented orthogonally to one another, alternating between $\pm 45^\circ$ to the vertical. These are referred to as the ‘*u*’ and ‘*v*’ views, as shown in Figure 3.6.

A cross section through a strip is shown in Figure 3.7. The strips contain a $2.0 \times 2.0 \text{ mm}$ groove into which a 1.2 mm diameter wavelength-shifting (WLS) fibre is glued. The groove is sealed with reflective tape, while the remainder of the facing of the strips is coated with a TiO_2 doped polystyrene reflective layer. When a charged particle passes through the detector, excitations of the scintillator material produce light. The light is collected by the WLS fibre and transported with minimal re-absorption to the end of the strip, transferred to optically transparent fibres and then to multi-anode photomultipliers (PMTs). This system is illustrated in Figure 3.8.

The PMTs convert the scintillator light into photoelectrons, which are multiplied by cascading secondary emission resulting from the dynode chain. The Data Acquisition software (DAQ) controls the digitisation of the PMT output by the readout electronics in each detector.

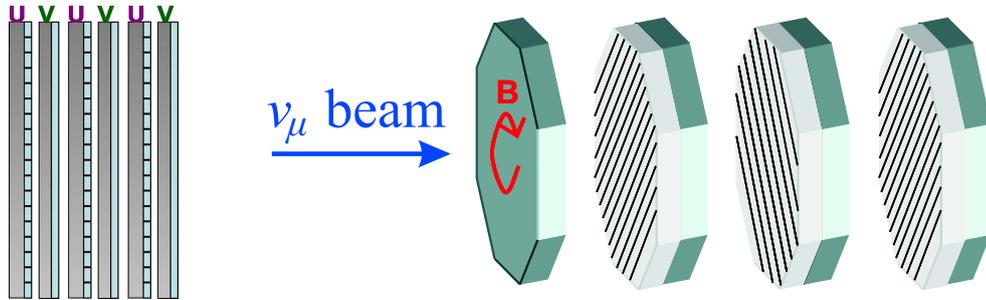


Figure 3.6: Arrangement of the scintillator and steel planes in the detectors, illustrating the two orthogonal strip orientations relative to the beam direction and plane spacing. Adapted from [12] and [54].

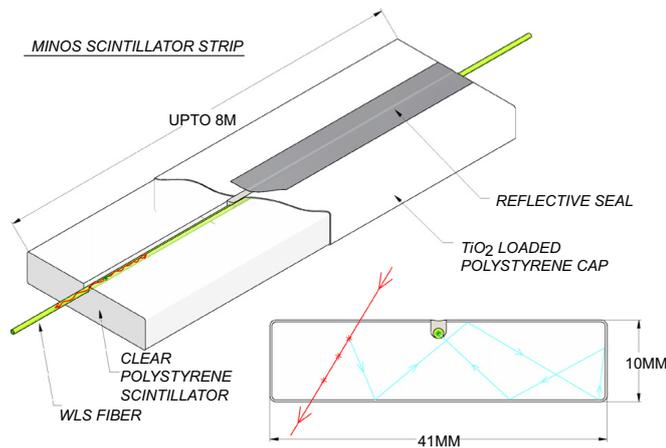


Figure 3.7: Cross section through a scintillator strip. Light from an ionising particle is reflected inside the strip, absorbed and re-emitted isotropically by a wavelength shifting (WLS) fibre. The photons are transported without significant loss to clear fibres and routed to the PMTs. Taken from [57].

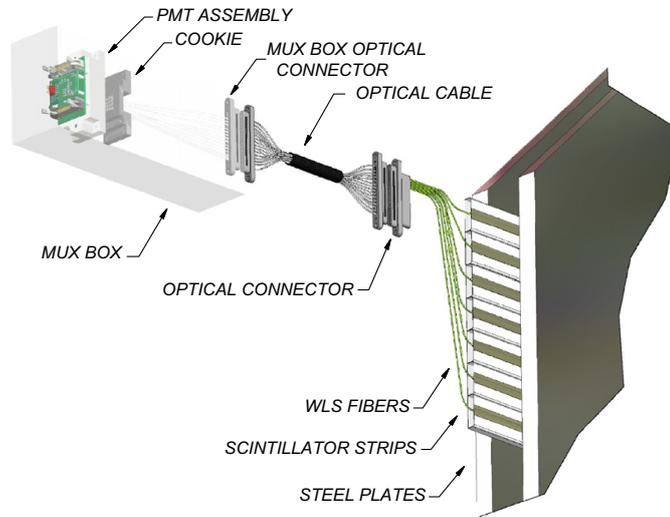


Figure 3.8: Optical readout for the detectors. Scintillation light captured by the WLS fibres is transported to photomultipliers via clear optical fibres. Taken from [57].

3.2.2 Far Detector

The Far Detector, shown in Figure 3.9, is located in the Soudan Underground Laboratory in Minnesota 735 km from the NuMI target, at a depth of 705 m. The detector consists of 486 octagonal steel-scintillator planes with an edge-to-edge dimension of 8 m. Each steel plane consists of eight 1.27 cm-thick plates welded together to produce a single 2.54 cm-thick plane, as the size of the mine shaft at the Soudan mine constrained the size of sheet which could be winched into the laboratory.

The 5.4 kton detector is divided into two ‘supermodules’ with a 1.1 m air gap between them, composed of 249 planes and 237 planes respectively, where the first plane of each module is uninstrumented. The total length of the detector is approximately 30 m.

Each supermodule is independently magnetised by a 15.2 kA coil running through the centre of each plane, inducing a magnetic field of 1.27 T in the steel. The direction of the coil current is generally chosen to focus negatively charged particles, which increases the fraction of ν_μ interaction products which are contained within the detector volume. The detector coordinate systems are defined in Figure 3.10, demonstrating the relationship between the x - y view and the u - v view.

The eight scintillator modules comprising a plane are composed of 192 scintillator strips, which are read out at both ends by Hamamatsu 16-anode PMTs. The signals



Figure 3.9: The MINOS Far Detector.

from eight strip ends (located in the same plane and spatially separated by about 1 m) are summed into a single photomultiplier channel, with different summing patterns on the east and west sides of the detector. This permits the demultiplexing of the signal to reconstruct a unique event signature (see Section 3.3.3).

A veto shield is installed covering the top and sides of the Far Detector, to reduce backgrounds to the atmospheric neutrino measurement from cosmic rays. It is constructed from two layers of the same scintillator modules as those used in the main detector, and Monte Carlo studies [58] have shown it successfully tags 99.9% of cosmic ray muons which penetrate eight or more detector planes, excepting those which enter through the front or back faces. Details of the Far Detector front end electronics and Data Acquisition system (DAQ) are given in [57].

3.2.3 Calibration Detector

A calibration detector (CALDET) was used in a CERN test beam to measure the response to beams of π , e and p beams of varying energy, ranging from 0.6 GeV to 10 GeV. It consisted of 60 unmagnetised planes one metre square, with the same alternating steel and scintillator construction as the Near and Far Detectors. One end of each scintillator strip was read out using the Far Detector electronics setup, and the other end the setup in the Near Detector, enabling a comparison to be made between the different readout systems. The responses of the electronics were found to be identical to within 0.6% [59].

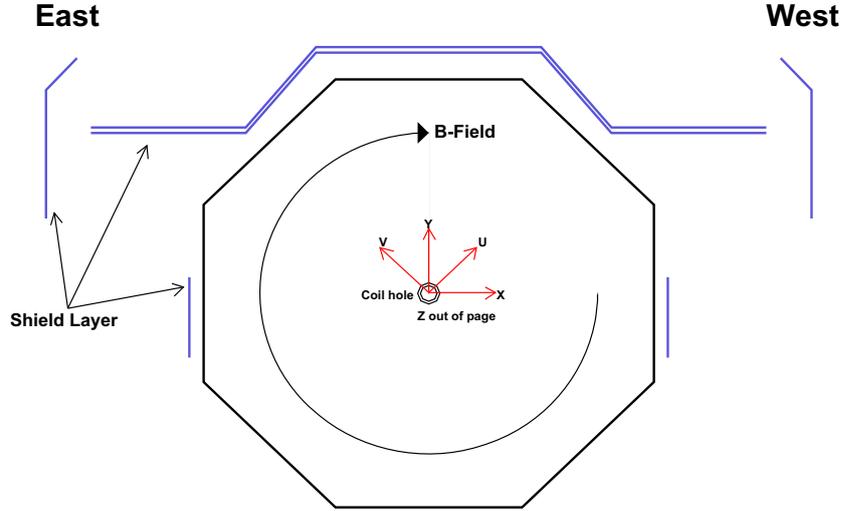


Figure 3.10: The MINOS Far Detector coordinate system, showing the relationship between the x - y view and the u - v view. The orientation of the neutrino beam is $+3^\circ$ to the z direction. Taken from [58].

3.3 Calibration, Simulation and Data Reconstruction

3.3.1 Detector Calibration

To perform an oscillation analysis, the MINOS detectors are required to provide a measurement of the absolute muon energy and absolute hadronic energy deposited after a neutrino interaction. The energy reconstruction must take into account the fact that changes in the detector response with detector region and over time may significantly affect the accuracy of the reconstruction. It is vital to correct for these effects in analyses which require the comparison of energy spectra between the Near and Far Detectors.

A calibration process is required to remove these variations and convert a raw photomultiplier signal $Q_{\text{raw}}(i, x, t)$, measured at time t in channel i for an energy deposit in position x into a true energy deposit. The process takes the form of a series of multiplicative calibration constants, calculated independently for each detector (Near, Far and Calibration):

$$Q_{\text{corr}} = Q_{\text{raw}}(i, x, t) \times D(t) \times L(i) \times S(i, t) \times A(i, x) \times M \quad (3.2)$$

with the multiplicative factors defined as follows:

- i. **Drift correction, $D(t)$** : This coefficient corrects for changes in the detector response with time; due to ageing of the scintillator and PMTs in the longer term, and temperature or replacement of electronics components in the short term.
- ii. **Linearity calibration, $L(i)$** : Ensuring the linearity of the PMTs and electronics is important for correctly measuring large energy deposits. The Light Injection (LI) system maps the linearity of the PMTs with pulse height on a strip-to-strip basis once a month to provide the linearity correction term in Equation 3.2. The LI also pulses multiple times an hour in both detectors to provide short term gains used to correct the Monte Carlo.
- iii. **Strip-to-strip calibration, $S(i, t)$** : The total response of each strip in each detector depends on differences in efficiencies between the WLS fibres and the PMTs. The correction variable $S(i, t)$ is obtained by monitoring the strip response to through-going cosmic muons, and is a function of strip and time.
- iv. **Attenuation, $A(i, x)$** : This correction accounts for the variation in light as a function of distance along the strip, as light which travels further along the WLS fibre is subjected to greater attenuation.
- v. **Absolute energy scale (inter-detector) calibration, M** : A relative calibration between the two detectors is required for signal comparison, and to assign an absolute energy value to the signals observed in both detectors, and in the calibration detector. The constant M converts the measured pulse-height into the same energy unit (Muon Energy Unit, MEU) comparable between detectors.

A summary of the calibration procedure is given in Figure 3.11. The procedure is not perfect, and differences in calorimetric response between the two detectors to the same energy scale are incorporated as detector-specific systematic uncertainties (see Chapter 6).

3.3.2 Monte Carlo Simulation

Simulated data are used extensively in MINOS oscillation analyses; the Far Detector prediction is compared with the data collected to observe evidence of the energy dependent spectral distortion resulting from ν_μ disappearance. The Monte Carlo data are also

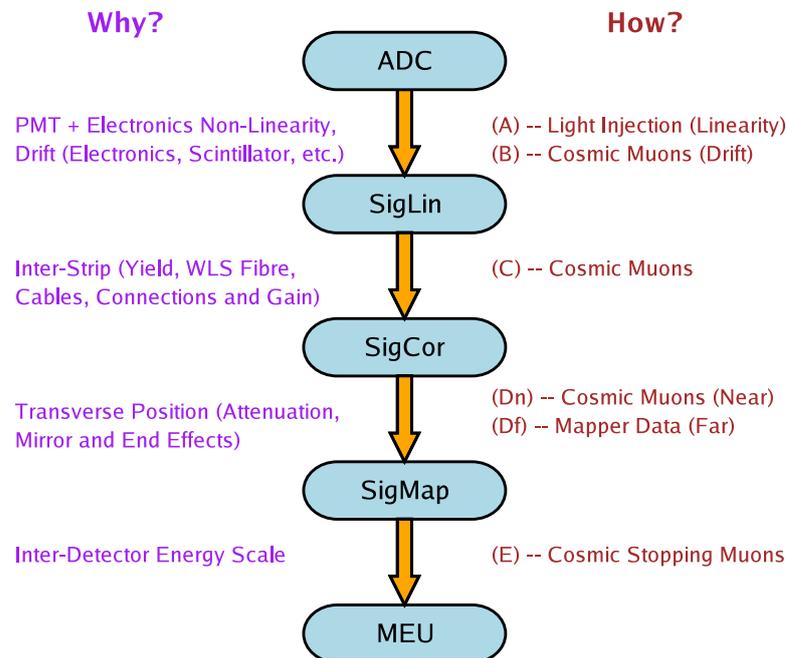


Figure 3.11: A summary of the MINOS calibration procedure. Each step corresponds to a correction factor in Equation 3.2. The correction from raw pulse height (ADC) to SigLin applies the linearity calibration, L and drift correction D ; from SigLin to SigCor the strip-to-strip calibration S ; from SigCor to SigMap the attenuation correction A and finally the absolute energy scale calibration M to convert the raw signal into the standardised energy unit MEU. Adapted from [60].

utilised to validate event selections, the reconstruction software and investigate oscillation sensitivities and systematic uncertainties. The simulation of the beamline, from the secondary particles produced in the target to detailed models of detector response to the final state particles is implemented using Monte Carlo methods.

The FLUKA [55] package, a general purpose package for modelling particle transport and interactions in matter, calculates the production of secondary mesons in the target. The FLUGG [61] package records and swarms the resulting hadrons through the horns and decay pipe, interfacing with the GEANT4 [56] simulation of the beamline. Documentation on the specifics of the latest NuMI beam simulation are given in [62]. All decays producing a neutrino are stored, given a trajectory passing through the Near or Far Detectors, a neutrino energy based on the decay kinematics, and a probability for the particular meson decay; providing an overall prediction for the beam neutrino flux. The neutrinos can then be input into event generation software.

The detector simulation samples neutrino events from the stored outputs, propagating them through the Near and Far Detector halls. Neutrino interaction events are generated within the detectors and the surrounding halls and rock by interfacing with the neutrino event generator NEUGEN [63]. NEUGEN acts as an event generator and cross-section library, permitting quasi-elastic, resonance and deep inelastic scattering interactions between 100 MeV and 100 GeV. A GEANT3 [64] simulation propagates the particles produced in the neutrino interaction through the detector representation, generating a series of ‘truth hits’. GCALOR [65] simulates the final interactions of the hadronic particles arising from the neutrino interactions with the detector mass.

The MINOS software packages `DetSim` and `PhotonTransport` are then utilised to convert hits in the detector into realistic detector readout. `PhotonTransport` simulates the scintillator response to an energy deposit, and takes the scintillation light through the wavelength shifting fibres to the PMTs. `DetSim` models the amplification of the photoelectron signal in the PMTs, the front end electronics and DAQ. It outputs the simulated data in the same format as the real data, a `RawDigitDataBlock` which can then be read into the reconstruction software. The package also models detector effects such as noise and PMT cross-talk (where light from one PMT can also register hits in one adjacent to it).

3.3.3 Event Reconstruction

The event reconstruction software uses low level objects such as hit topology and timing information contained in real or simulated raw data to reconstruct high level objects, such as muon tracks and hadronic showers. This enables an estimation to be made of the visible neutrino energy for muon and electron neutrino CC and NC events, in addition to defining a set of quantities to aid discrimination between these interaction processes.

The first stage of the reconstruction process is to take in the raw data blocks output by the readout systems and form strip objects, each consisting of an energy deposit in a single scintillator strip. In the Far Detector, an additional demultiplexing step must also be performed; in this case, each strip is read out at both ends resulting in two possible channels, and with the multiplexing of cables into each pixel each digit has eight candidate strip ends associated with it.

To identify which pair of strip ends truly corresponds to each pixel, a demultiplexing algorithm [66] compares pairs of strip ends in each plane, and for those pixels with unambiguous solutions the timing information is used to constrain the region the event could have occurred in and the likely event type; this template is then used to demultiplex the remaining strips. The pixels are multiplexed in such a way as to make this possible, for strips have unique combinations of pixels on the east and west sides of the detector, and those sharing a common pixel are separated by more than 23 strips in that plane. Once demultiplexing is complete in the Far Detector the digits are formed into the strips determined by the demultiplexing, read out at both ends. In the Near Detector the strips are read out at one end, and groups of digits occurring in the strip at a similar time form the strip object.

The high flux intensity of the NuMI beam means multiple neutrino events per beam spill are expected in the Near Detector. Hence, the next stage in the reconstruction procedure is to slice the digits in the strips by time and proximity into likely single neutrino interactions. The comparatively low event rate in the Far Detector generally results in a single slice per strip, as most beam spills contain no interactions and rarely contain more than one event. Tracks and showers are formed within slices to fully reconstruct the events.

3.3.3.1 Muon Tracks

The defining characteristic of a ν_μ CC neutrino interaction is a muon track with associated shower. Muons as heavy, relativistic particles lose energy in matter by ionising the surrounding medium and exciting atomic nuclei. The mean rate of energy loss is described by the Bethe-Bloch equation (taken from [17]):

$$-\left\langle \frac{dE}{dx} \right\rangle = Kz^2 \frac{Z}{A} \frac{1}{\beta^2} \left[\frac{1}{2} \ln \frac{2m_e c^2 \beta^2 \gamma^2 T_{\max}}{I^2} - \beta^2 - \frac{\delta(\beta\gamma)}{2} \right] \quad (3.3)$$

where z is the charge of the incident particle in units of electron charge, Z and A are the atomic mass and atomic number of the medium, m_e is the mass of the electron, T_{\max} is the maximum kinetic energy which can be imparted to a free electron in a single collision and I is the mean excitation energy in the medium. K is a combination of constants ($K = 4\pi N_A r_e^2 m_e c^2$), where N_A is Avogadro's number and r_e is the charge radius of the electron. β and γ have the usual relativistic kinematic definitions. In MINOS, the theorised mean rate of energy loss is calculated from the tables given in [67], which incorporates modifications to the classic Bethe-Bloch theory.

The energy loss as a function of $\beta\gamma$ falls rapidly, reaching a minimum between $2 < \beta\gamma < 4$ and plateauing at higher energies. This shape is shown in Figure 3.12, which shows the theoretical Bethe-Bloch stopping power for muons in the Far Detector as a function of energy, in addition to the stopping power derived from cosmic ray data and Monte Carlo. Aligning the minima of these curves gives an absolute energy scale calibration, with a 2% uncertainty in the dip position resulting in a 0.2% uncertainty on the absolute energy scale [57].

The muons resulting from CC interactions in the MINOS detector are generally produced with $\beta\gamma$ values near the minimum of the stopping power curve, and so are defined as minimally-ionising particles (MIPs). As the muons travel through the detector they lose energy, and when their total energy is below the minimally-ionising region they deposit more energy per distance travelled.

Muons in the MINOS detectors form tracks from ionising particles within the scintillator. Track finding and fitting algorithms [32] have been developed to clearly identify and reconstruct the kinematics of the muons. The track finder identifies 'seed tracks', the basic structures which define candidate muon tracks. These are then used as inputs to a Kalman filter algorithm, which estimates the path of the muons as they move helically in the magnetic field. The muon momentum can then be reconstructed from range

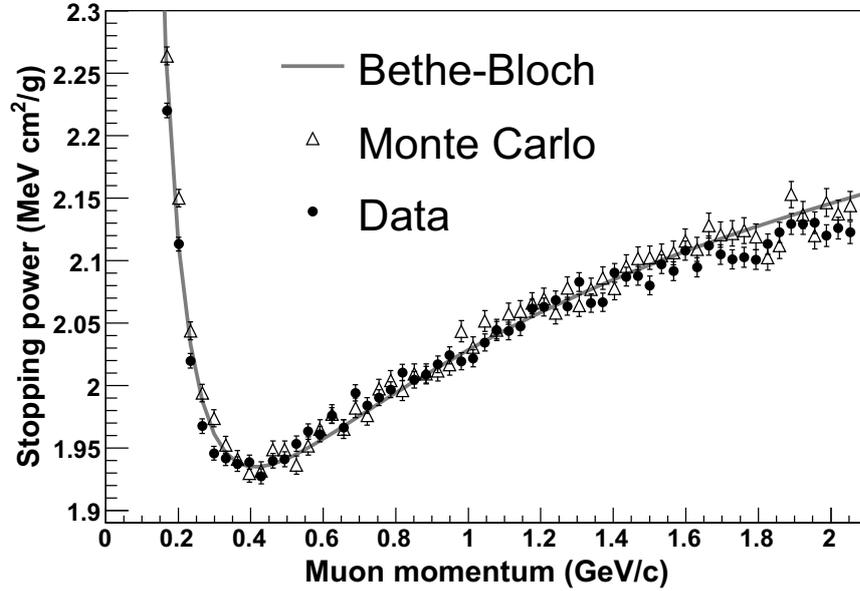


Figure 3.12: Stopping power for muons in the Far Detector, from data, Monte Carlo and the theoretical Bethe-Bloch calculation. The minimum ionisation for muons in the detector is found to be 0.4 GeV/c. Taken from [57].

if the track is fully contained, and additionally from the curvature of the track which also reconstructs the muon charge sign.

3.3.3.2 Showers

Strips which are clustered together in space and time are used to construct showers, though they are not resolved on an individual particle level. Instead, the shower energy is calculated from the summed pulse height of the individual hits, with those hits associated with the track on strips that are shared subtracted. Multiple showers can be reconstructed within each slice, where it is vital to correctly identify the hadronic shower arising directly from the neutrino interaction vertex.

Detector response to the beams differs if a hadronic shower or electromagnetic shower is produced. After a CC or NC ν_μ interaction, the charged secondary particles lose energy through the Bethe-Bloch process and interact via the strong force with nuclei producing low energy hadrons. The energy deposits from these process make up the visible hadronic shower energy.

A CC ν_e interaction can also occur, from $\nu_\mu \rightarrow \nu_e$ oscillations or beam ν_e con-

tamination. The resulting electrons at MINOS energies can radiate photons through Bremsstrahlung which produce e^+e^- pairs which undergo further radiation, resulting in a cascading electromagnetic shower. These ν_e events are relatively rare in the MINOS detector and are the subject of a recently published analysis [68].

The final event object is therefore the combination of the vertex shower and the most energetic muon track. The optimal combination of objects is identified by considering the temporal and spatial distribution within each slice. For muon neutrino interactions, the neutrino energy is estimated to be the sum of the energies of the longest track and the vertex shower. The presence of a track with a clearly defined vertex is a condition applied when selecting a ν_μ sample, in addition to cuts on other quantities used to remove NC events where short tracks are mistakenly reconstructed within showers (see Chapter 5).

3.4 Analysis Dataset

Figure 3.13 shows a summary of the Protons-on-Target accumulated between May 2005 and June 2009. Four distinct data taking periods can be identified and are used to divide the analysis sample:

- **Run I:** Run I commenced in May 2005 and ended with the accelerator complex scheduled shutdown in February 2006, with the NuMI target in the LE-10 position.
- **Run I-pHE:** Immediately following the shutdown there was a two month period of high energy beam running, with the target 250 cm upstream of the focusing horns in the ‘pseudo-High Energy’ configuration (pHE). These data are considered a subset of that taken previously with the same target and so is referred to as Run I-pHE. During this time problems with the motion system of the target became apparent, and the target was replaced and installed in the LE-10 position in late August/early September 2006.
- **Run II:** Run II commenced in September 2006 and ran until July 2007. For previous MINOS publications and conference results, the dataset has previously been divided into Run IIa (September 2006 - March 2007) and Run IIb (April 2007 - July 2007). The combined exposure at this point formed the dataset analysed for the 2008 ν_μ disappearance result [69]. Runs I and II are analysed separately

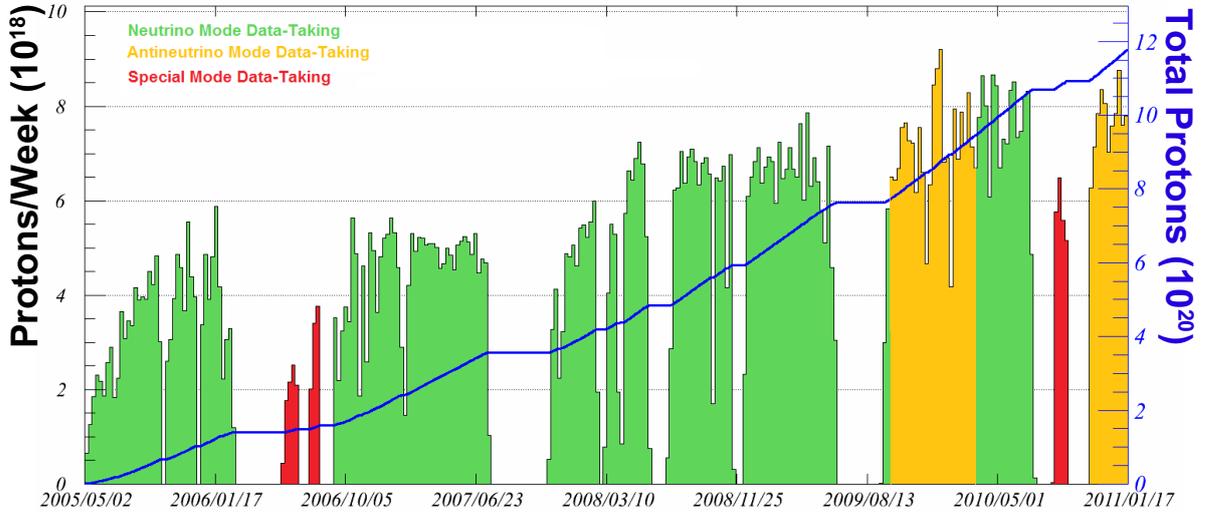


Figure 3.13: Protons-on Target (PoT) as a function of time from the start of MINOS data taking; per week and the integrated total. The data used in this thesis was recorded as shown between May 2005 and June 2009 with a total of 7.25×10^{20} PoT accumulated. Taken from [70].

due to a ~ 1 cm longitudinal shift in the LE-10 position after the swapping of the targets, which required some modifications to the Monte Carlo.

- **Run III:** Run III began taking data in the LE-10 configuration in November 2007 after an accelerator long shutdown, during which time the beampipe was filled with helium. The Run III data more than doubled the exposure from the previously published analyses to a total of 7.25×10^{20} PoT used for this thesis.

Table 3.1 summarises the start and end dates of the Run periods, the Run/Subrun numbers corresponding to these endpoints as well as the total PoT accumulated at the Far Detector during this time. The term ‘Run’ has two meanings in the context of MINOS data; the first mainly used here refers to the long periods spanning months of data taking when hardware and configuration changes determine the boundaries of each Run. The second refers to a single period of 24 hours of data taking, which is divided into ‘subruns’, each corresponding to an hour of data taking.

	Start Date	End Date	Far Det. (Run/Subrun)		Far Det. PoT / 10^{20}
			Start	End	
Run I	20/05/2005	26/02/2006	31720/0	33794/8	1.27
Run I-pHE	11/06/2006	13/08/2006	35765/16	36145/22	0.15
Run II	12/09/2006	17/07/2007	36570/0	38449/3	1.94
Run III	18/11/2007	13/06/2009	39965/10	43639/7	3.89

Table 3.1: Summary of the Far Detector analysis dataset.

Chapter 4

The Near Detector and Data Quality

It is essential that the MINOS physics analyses utilise data taken during periods of normal running of the detectors; with anomalies in Near and Far Detectors identified and the appropriate subruns removed from the analysis datasets. This Chapter discusses in detail the operation and electronics of the Near Detector, with Section 4.2 dedicated to details of the implementation of the Near Detector data quality validation software developed as part of this thesis.

4.1 Detector Overview

The Near Detector, shown in Figure 4.1, is located 1040 m from the NuMI target and 100 m underground at Fermilab. The 0.98 kton detector consists of a single module of 282 steel-scintillator planes; each plane has a characteristic “squashed-octagon” shape, 6.2 m in width and 3.8 m in height. The detector is magnetised by a 40 kA coil offset by 1.49 m from the horizontal centre of the detector, generating a 1.17 T field at the intersection point with the neutrino beam. The detector and magnetic field geometries were designed to provide a close approximation at the beam spot to the magnetic field and uniformity of instrumentation present in the Far Detector, and contain the neutrino interactions in the region surrounding the intersection. Away from the beam spot financial constraints affected the level of instrumentation installed.

Two different types of plane are present in the Near Detector; fully and partially



Figure 4.1: The MINOS Near Detector.

instrumented. The former consist of 96 scintillator strips, covering an area of 13.2 m^2 of the steel plane. The latter consist of only 64 scintillator strips, covering a reduced area of 6.0 m^2 . Every fifth plane along the length of the detector is fully instrumented, where four of every five planes of the first 121 are partially instrumented. The different configurations are shown in Figure 4.2. The z direction at the Near Detector is defined perpendicular to the planes, with the beam angled at -3° relative to the z axis. Scintillator strips are read out at one end by Hamamatsu 64-anode PMTs; due to the shorter strip lengths in the Near Detector the effect of attenuation is reduced. There are 9248 live PMT pixels (channels) in both sections of the detector.

The detector can be subdivided into two sections, where the 121 plane front section forms the calorimeter, and the remaining 161 planes downstream form the muon spectrometer. The calorimeter is designed for neutrino interaction vertex identification and hadronic shower containment, in addition to accurately reconstructing the early segments of muon tracks. There is no multiplexing of strip ends in the calorimeter section, and all planes are fully or partially instrumented.

The muon spectrometer is utilised to track high energy muons resulting from interactions within the calorimeter. Only one in five planes is instrumented, and each PMT pixel reads out signals from four strip ends. The track segment observed in the calorimeter is extrapolated into the spectrometer in order to resolve the resulting four-fold ambiguity.

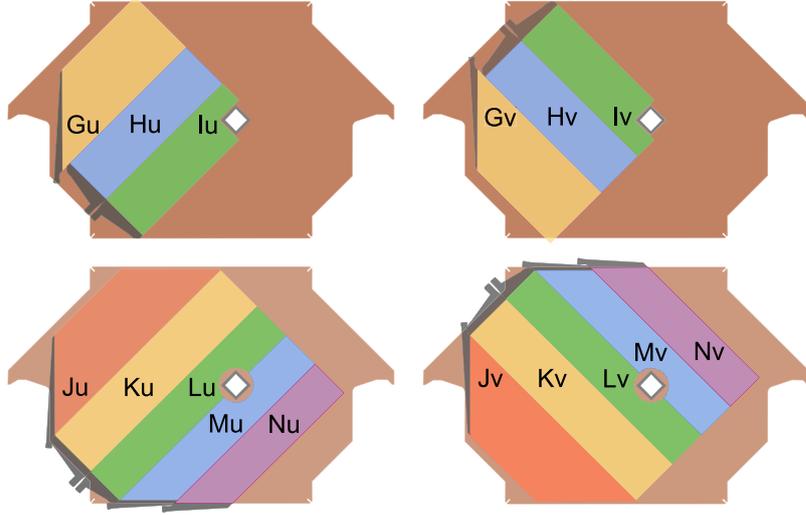


Figure 4.2: Configurations of planes in the Near Detector. Those on the left are u view planes and those on the right are v views; partially instrumented in the upper two figures and fully instrumented in the lower. The labels G to N denote the different lengths and shapes of the scintillator modules. The z direction is defined perpendicular to the planes, with the beam angled at -3° relative to the z axis.

The typical beam intensity at the Near Detector is 2.2×10^{13} PoT/spill with a beam diameter of approximately 50 cm, inducing an average of 16 neutrino interactions per beam spill during LE running. This differs significantly from the Far Detector, which averages significantly less than one ($\sim 10^{-4}$) interaction per beam spill. About half of the Near Detector neutrino interactions will occur in the calorimeter region and hence will be fully reconstructed. Due to the much higher rate of interactions in the Near Detector and the occurrence of multiple events per beam spill, the detector uses specially designed front-end electronics [71], capable of processing continuously throughout a spill with no downtime.

4.1.1 Front-End Electronics

The fundamental requirement of the Near Detector electronics is to be able to accurately separate signals during the $10 \mu\text{s}$ beam spill by digitising continuously at a sufficient frequency, and to do so with no downtime during beam spills to record and measure all neutrino interactions. The digitisation frequency is identical to the Main Injector RF clock and equal to 53 MHz, which results in an intrinsic timing resolution of 18.8 ns.

The PMTs induce a cascade of photo-electrons (p.e.) with a gain of 1×10^6 , so a signal of one initial p.e. is equivalent to an output integrated charge of 160 fC. The electronics are required to be able to resolve a signal range from one to multiple hundreds of photo-electrons, with a separation time of the order of 100 ns. The digitisation frequency is designed to provide this level of signal discrimination.

There are four main elements to the Near Detector electronics: the MENUs (MINOS Electronics for Neutrinos) which interface directly with the output from the PMTs; the MINDERs (MINOS Near Detector Electronics Readout) which power and read out the MENU components to the MASTER (MINOS Acquisition, Sparsifier and Time-stamper for Event Readout), which interfaces between the front-end systems and the Data Acquisition (DAQ) system, the processors and the trigger farm. The final component is the clock system, which controls all time signals across the detector.

MENU (MINOS Electronics for Neutrinos)

The MENU comprises the fundamental channel unit of the Near Detector, hosting the Charge Integrator and Encoder (QIE) electronics and Analogue to Digital Converter (ADC) which digitise the PMT readout, and a data storage object known as a FIFO (*First In First Out*). The QIE has excellent noise and linearity characteristics and a large dynamic range, integrating the input current from the anode of the PMT every 18.8 ns, with additional charge from the detector being added cumulatively during this time. It acts as a current divider, splitting the input current into eight binary weighted ranges: 1/2, 1/4, 1/8 and so on to 1/256. This is performed by grouping sets of transistors, forward biased by a current which divides in the same way as the current from the detector, resulting in a series of binary weighted voltages on storage capacitors.

The operation is split into four phases: *integration*, *comparison*, *output* and *reset*; each phase lasting for a QIE clock cycle. Hence to ensure the electronics do not evince 75% deadtime in recording data, for each current range four sets of eight capacitors are required. The integration phase is described in the previous paragraph. The voltage on the capacitors is then compared (*comparison period*) to preset threshold voltages; the bias current and current-splitting ratios are designed so only one range will fall into the selection band to be digitised by the ADC. This digitisation period corresponds to the *output* cycle. The capacitors are then reset.

The digitisation is performed by an 8 bit flash ADC, running on the same 53 MHz clock as the QIE. An idealised transfer function between the current measured by the

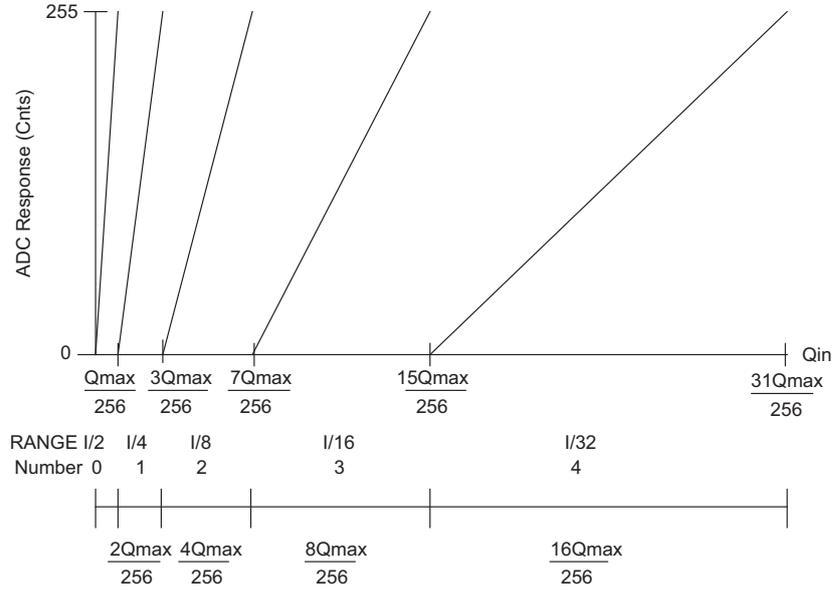


Figure 4.3: Response of the QIE electronics, for the first five of eight current ranges. The x -axis shows the amount of charge injected into a chip, while the y -axis shows the ADC response for different current ranges. These idealised transfer functions are used for calibration to reconstruct the input charge from the ADC count. Taken from [57].

QIE and the ADC count assigned is shown in Figure 4.3, for the first five of the eight ranges. These curves are constructed by applying a known calibration current to the QIE. Range 0 is defined as spanning an input charge of $Q_{max}/256$, and each subsequent range spans double that of the previous, with half of the charge resolution. The resolution for range 0 is 2.6 fC per count, and the maximum of the total dynamic range is 85 pC.

After digitisation, the data recorded have a 13 bit format; 8 bits storing the ADC count, 3 RANGE bits which label which current range was selected for digitisation, and 2 CAPID bits which indicate which of the four sets of capacitors read in the datum word. A datum is produced every 18.8 ns, the intrinsic timing unit which is generally expressed with the unit of a ‘bucket’ or ‘TDC’ for the purposes of data monitoring and quality plots shown in Section 4.2. Each data word is then stored in the FIFO, which during beam spills holds up to 530 data words.

MINDER (MINOS Near Detector Electronics Readout)

The MINDER is the motherboard of the MENUs, with 16 MENUs to each of the 578 MINDERS. It provides power to the MENU components, reads out the data stored in

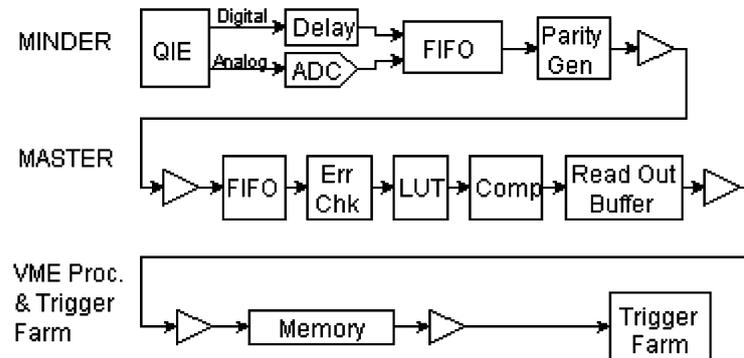


Figure 4.4: Block diagram of data flow through the front end electronics of the Near Detector. The progression of data from the QIE circuit at the MENU level to the VME crate and trigger processing farm is visible, including error checking and the data linearisation through the use of the Look-Up-Table (LUT). Taken from [71].

the FIFOs, timestamps it and pushes it up the electronics chain to the MASTER module, as well as supporting calibration and other functions of the MENU module. At the end of each beam spill gate, the MINDER sends a start of datastream header and initial timestamp to the MASTER, before reading out the first word from each MENU. The timestamp is then incremented and the second word from each MENU is read out and so on until the end of the data stream, where a footer is appended. Readout is sufficiently fast to occur in the 1.4s gap between beam spills, and so doesn't occur simultaneously with data collection. First level trigger information is also recorded; triggers are discussed in the following section.

MASTER (MINOS Acquisition, Sparsifier and Timestamper for Event Readout)

The MASTER services up to eight MINDERS, and acts as the interface between the front-end electronics and the Data Acquisition system (DAQ). Up to 12 MASTERS are hosted by an individual electronics 'Crate'. Each MASTER linearises the QIE data using a large Look-Up-Table (LUT) containing output values resulting from charge-injection calibration for every possible QIE code for all channels. This converts the datum into a 16-bit number tagged as a 'digit'.

It also performs checks on data integrity; a PARITY bit is attached to the data in the MINDER and checked for any corruption of the data word, and header and footer bits are checked for the length of the data stream, where a problem assigns a MISCOUNT

error to the record. If the digit passes the pedestal threshold (a value above the non-zero ADC count for no input) the channel ID and timestamp is appended and data are then sent to memory and to the trigger farm for further processing. The data flow through the system from QIE to trigger farm is summarised in Figure 4.4.

4.1.2 Data Acquisition and Triggers

The DAQ systems for the Near and Far Detectors are functionally identical, with the purpose of transferring data to a PC farm where events are built, and those that pass a selection of triggers are retained for physics analysis. The DAQ also performs monitoring and calibration tasks. Data from adjacent time periods from each MASTER is concatenated in the memory buffers into units (called timeframes) of 1 s in duration; the smallest time unit the DAQ interfaces with. Each timeframe retains a header with all information from the linearised digits from the MASTER, and is sent to the trigger processing (TP) farm.

At this stage, the DAQ writes out a series of data blocks known as `TpSingles`, mapping all hits onto their channel IDs, after the first level triggers have been applied and prior to the second level triggers being considered. First level triggers define the main modes the front-end electronics run in, and determine at the MENU level which pixels are read out and digitised, depending on the presence or not of a beam spill. Second level triggers are applied later by the trigger processing (TP) farm, determining which data is written for analysis. The triggers generally applied for physics analyses are detailed as follows. After the triggers have been applied, the data streams are written as files containing `RawDigit` data blocks, formatted to be input into reconstruction and data quality software.

- **First Level Triggers:**

- *Spill Mode:* Continuous processing of all pixels during the spill gate (SGATE); a $13\ \mu\text{s}$ window surrounding the beam spill, opened by a signal from the main injector to the MINDERS $1.5\ \mu\text{s}$ before the spill and closed $1.5\ \mu\text{s}$ after the spill. Intended to record every hit with redundancy time built in.
- *Dynode Mode:* Designed to record out-of-spill cosmic ray data, continuous digitisation for 150 ns is initiated by each PMT independently when the dynode signal exceeds a set threshold.

- **Second Level Triggers:**

- *Plane trigger*: A customisable trigger, where M detector planes in any set of N contiguous planes contain at least one hit. Generally $M = 4, N = 5$ for physics analyses.
- *Activity trigger*: Require N planes to register activity across the entire detector. Nominally $N = 20$.

4.2 Near Detector Data Quality

Subjecting the data collected at the Near Detector to quality checks is imperative, to ensure physics analyses are based on time periods of normal detector functioning. Previously data validation was performed manually on a run by run basis [72], based on the MINOS Online Monitoring system. A new validation scheme has been developed¹, making direct use of the raw data in determining the status of the Near Detector. Raw data blocks (`TpSingles` and `RawDigit`) are used directly to calculate data quality variables for each subrun, measuring the fraction of readout channels that are recording normally and identifying known anomalies in the functioning of the detector systems.

4.2.1 Summary of Quality Cuts

The data quality cuts derived in this thesis were based on data accumulated between March 2005 and October 2009, comprising 33441 total subruns in the Near Detector. The Near Detector dataset has been subjected to the full set of checks detailed below, analogous to the corresponding quality controls already in place at the Far Detector [73].

If any of the following criteria are satisfied, the subrun is flagged as “bad” in the Near Detector run quality database. The cuts are applied at the subrun level, discarding an hour of data taking and preventing a transient problem resulting in an entire run (equivalent to approximately 24 hours of running) being discarded. Of the 33441 subruns analysed, 3.0% are rejected.

¹The work presented here was performed by the author in collaboration with Dr A. Blake and Dr G. Lefeuvre.

- **TP Singles Quality Cuts:**

- i. Any number of Crates missing in readout
- ii. More than 4 cold MINDERS in either the calorimeter or spectrometer
- iii. More than 96 cold MENUs in the calorimeter
- iv. More than 96 hot MENUs in the calorimeter

- **Raw Digits Quality Cuts:**

- v. A Light Injection (LI) rate during a spill (SGATE) trigger greater than 1 Hz
- vi. More than 5 channels (spectrometer) or 15 channels (calorimeter) recording a CAPID error bit greater than the threshold rate of 100 Hz
- vii. More than 5 channels (spectrometer) or 10 channels (calorimeter) recording a PARITY or MISCOUNT error bit greater than the threshold rate of 10 Hz
- viii. More than 25 channels (spectrometer) or 50 channels (calorimeter) recording any other error bit greater than the threshold rate of 10 Hz (30 Hz for CAPIDSameMenu errors)
- ix. Anomalously high ADC counts in either the calorimeter or the spectrometer, more than 200 channels recording an average signal of high ADC counts (>100) per spill trigger greater than 1000
- x. Acquisition window not well synchronised with the spill

- **Additional Quality Cuts:**

- x. Subrun duration less than 400 s
- xi. Any subrun not tagged as ‘Physics’, ‘Modified’ is allowed
- xii. Any subrun tagged as ‘Test’
- xiii. Anomalously high median or maximum spill rate; greater than 50 Hz or 500 Hz respectively
- xiv. Anomalously low plane and spill triggers; fewer than 400 combined

4.2.2 Details of Failure Modes

The range of criteria on which selection cuts are placed were identified from the information available in the raw data files produced before and after the second level triggers

are applied. A training sample of approximately 2000 subruns² was used to identify where to place the selection cuts on a range of variables.

Readout Electronics

Readout problems are the most common failure modes seen in the Near Detector, typified by localised regions where the readout is either dead or persistently reading out very high values. The readout holes and spikes are chiefly caused by malfunctioning electronics which affect only a small number of MENUs, but can be caused by more widespread problems which affect entire MINDERS and readout Crates. Selection cuts at these three granularities remove $\sim 54\%$ of the subruns flagged as bad from the full dataset.

Utilising data from the training sample, within the calorimeter section of the detector each MENU with a singles rate (number of hits in one channel per timeframe) of 0 Hz is tagged as ‘cold’, and with a rate greater than 1000 Hz as ‘hot’; compared to the average rate of around 15 Hz. In the spectrometer, the boundaries of the definitions of cold and hot are set at 20 Hz and 1000 Hz respectively; compared to an average of approximately 70 Hz. The boundary is set at a non zero value in the spectrometer due to two particular malfunctioning MINDERS evidencing persistently low rates, resulting in a secondary peak in the rate below 20 Hz. A MINDER with all channels identified as having an anomalously low rate is labelled as cold.

To define the characteristic number of cold MENUs and MINDERS for each subrun, the distribution of the maximum number of cold MENUs or MINDERS in each timeframe is considered. The highest 30 entries are removed, truncating the tail of the distribution and minimising the effect of very high values indicating transitory problems which rapidly return to a lower level. The highest remaining value is defined as the characteristic or maximum value for the subrun. This is shown in Figure 4.5 on the MENU scale, with the characteristic number of cold channels plotted for all analysed subruns from 2005 to 2009, and the final cuts placed on this variable. Similarly, Figure 4.6 shows the cut on the number of cold MINDERS for both the calorimeter and the spectrometer. Figure 4.7 shows the distribution of crates failing to read out data correctly. In the case of the latter, any value greater than zero indicates a large readout hole in the detector and the subrun is removed.

²The training sample comprises 1797 subruns; including all runs identified as bad in the preceding data quality study [72] and a random selection of good runs from the same time period.

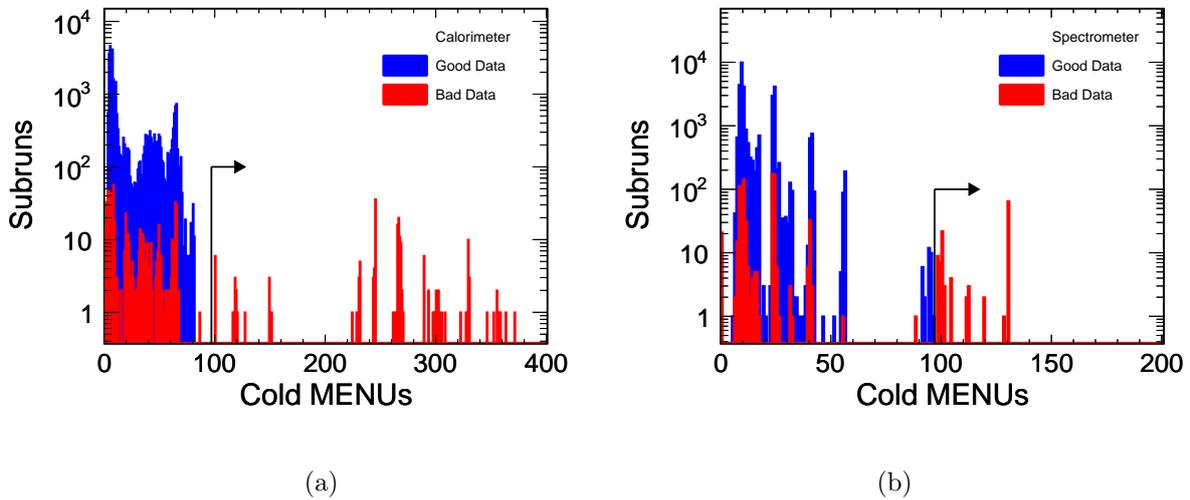


Figure 4.5: Characteristic number of cold MENUUs observed in the calorimeter (a) and the spectrometer (b) for the full dataset by subrun. Those classified as bad by the full selection are indicated in red, and good subruns in blue. The selection cut for this variable is also shown.

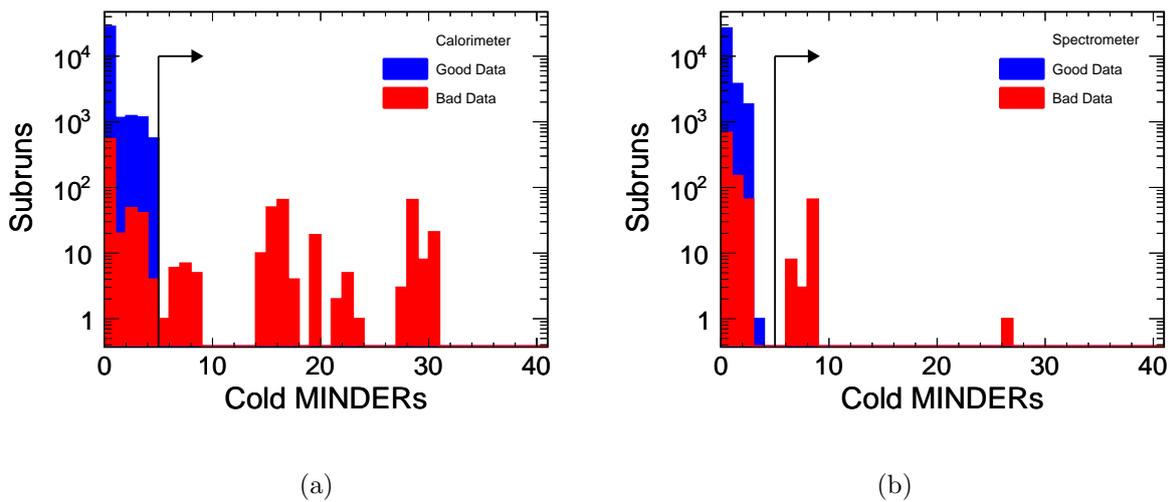


Figure 4.6: Characteristic number of cold MINDERS observed in the calorimeter (a) and the spectrometer (b) by subrun. Those classified as bad by the full selection are indicated in red, and good subruns in blue. The selection cut for this variable is also shown.

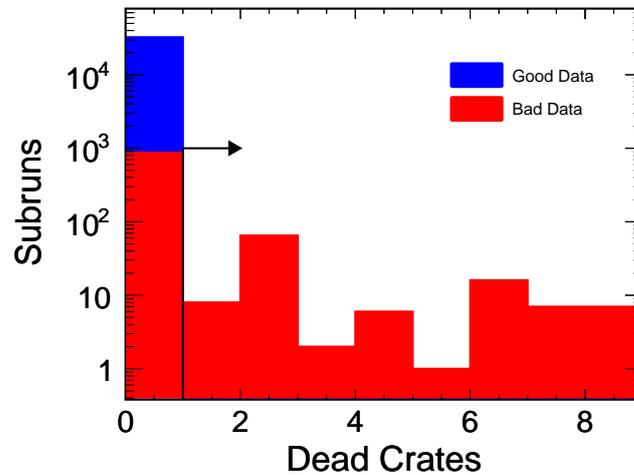


Figure 4.7: Distribution of electronics crates not reading out for all or part of each subrun. Those shown in red are classified as bad by the full selection summarised in Section 4.2.1, with good subruns shown in blue. The arrow denotes the position of the selection cut on this variable.

Light Injection

The Light Injection (LI) system maps the linearity of the PMTs and electronics and monitors the integrity of the optical readout systems, on both short and long time scales. The LI system uses UV LEDs to inject light pulses into the WLS fibres, mimicking scintillation light signals. During a period of data taking, the LI system pulses every Far Detector channel about 300 times per hour, and each Near Detector channel 1000 times per hour. These data provide corrections for transient instabilities in the electronics and PMT gain, and are shown to be highly correlated with short term changes in the detector hall environment.

The LI system is hardware inhibited, and the pulses recorded for calibration purposes should never coincide with an SGATE trigger. All subruns with an average rate of LI hits greater than 1 Hz are flagged as bad, to distinguish between those few subruns with very low rates of leakage and those with high rates that occur throughout the subrun.

Error bits

Three possible error bits can be set during data readout and acquisition, with an additional four bits indicating a hit in a suspicious channel. A CAPID error occurs when, as described previously, the wrong set of capacitors was selected for digitisation

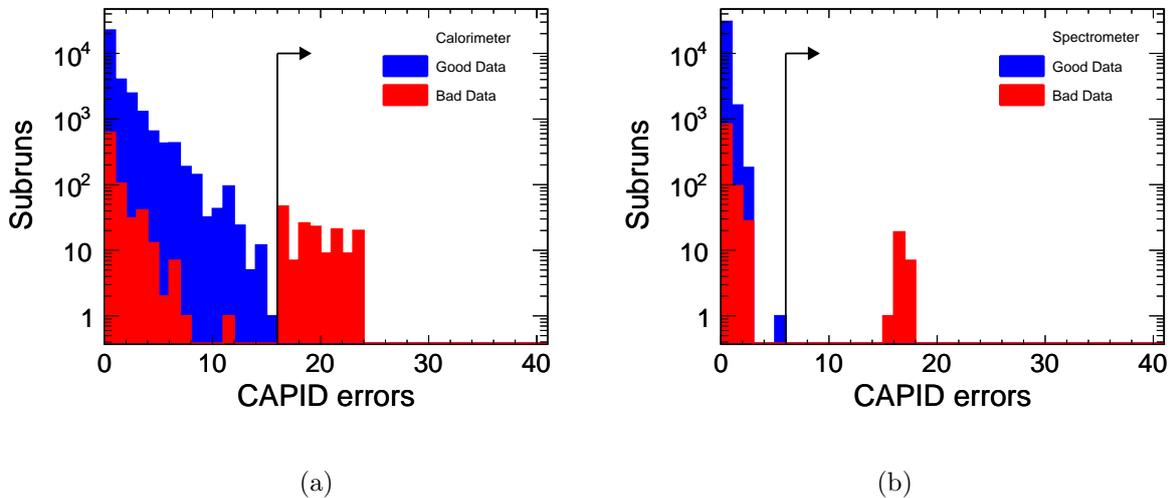


Figure 4.8: Maximum number of channels with high CAPID error rates by subrun, observed in the calorimeter (a) and the spectrometer (b). Those classified as bad by the full selection are indicated in red, and good subruns in blue. The selection cuts for this variable is also shown.

in the MENU. A PARITY error indicates a problem in the transition of the data word through the MINDER, and a MISCOUNT error is indicative of a problem in the length of the data word, the incorrect transfer of headers and trailers, or a buffer overflow.

A CAPID error rate in one channel greater than 100 Hz is taken to define an unacceptable rate; Figure 4.8 shows the distributions of the maximum number of channels with a rate greater than this threshold during each subrun, for the calorimeter and spectrometer. The cut on the maximum permitted number of malfunctioning MENUs is also indicated. A PARITY or MISCOUNT error rate of greater than 10 Hz defines the unacceptable threshold for these variables; Figures 4.9 and 4.10 illustrate the distributions and cuts placed in the calorimeter and spectrometer.

In addition to the error bits, four warning flags indicate either an interruption of the readout stream by higher priority data (PriorityTruncate) or that the channel shares a MENU or MINDER with one indicating a CAPID, PARITY or MISCOUNT error (CapidSharedMENU, CapidSharedMinder and TransferSharedMinder). Unacceptable rate thresholds are also set on these variables, at 30 Hz for the CapidSharedMENU bit and 10 Hz for the remaining bits, to take into account more widespread detector problems. Thresholds on the maximum permitted number of such channels are set higher

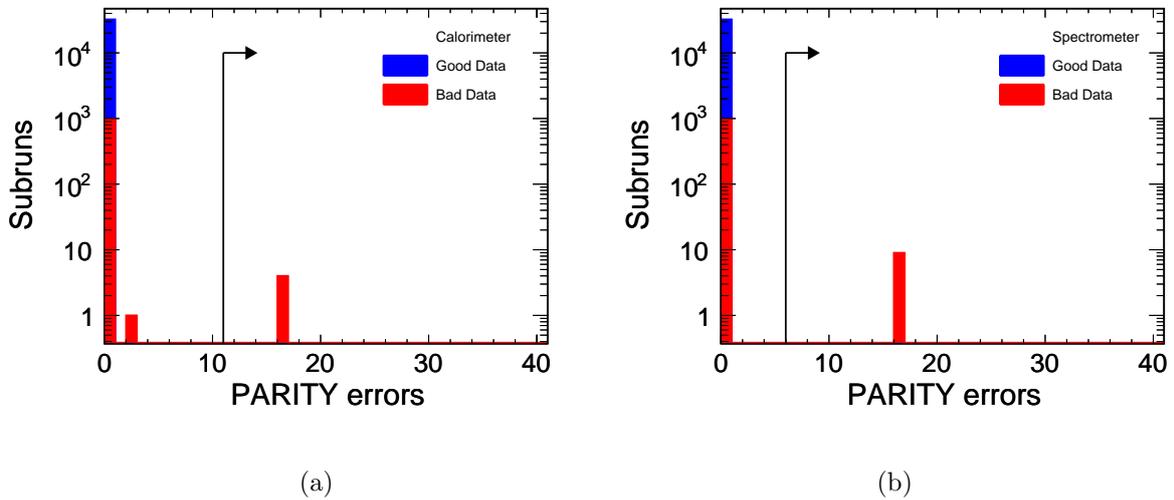


Figure 4.9: Maximum number of channels with high PARITY error rates by subrun, observed in the calorimeter (a) and the spectrometer (b). Those classified as bad by the full selection are indicated in red, and good subruns in blue. The selection cuts for this variable is also shown.

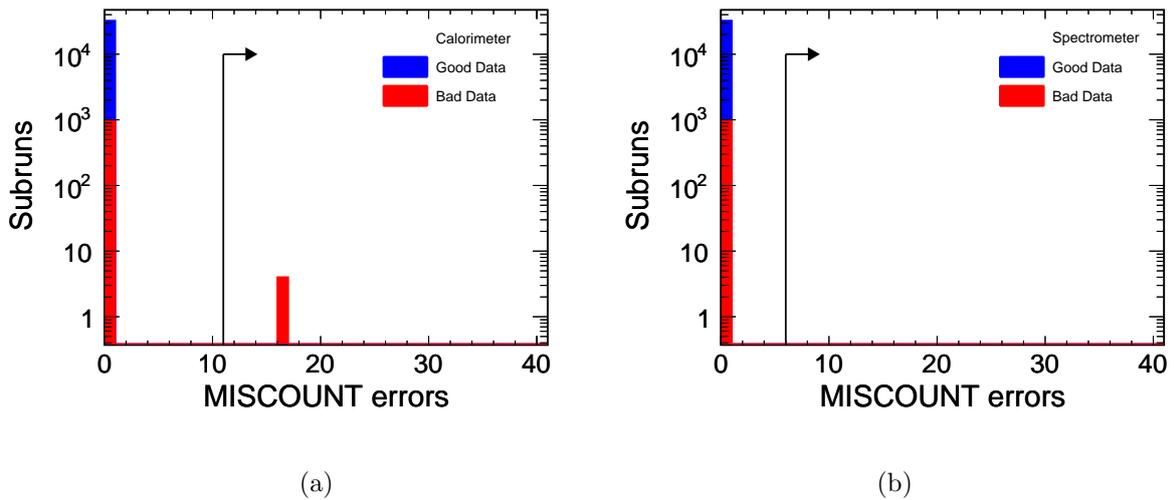


Figure 4.10: Maximum number of channels with high MISCOUNT error rates by subrun, observed in the calorimeter (a) and the spectrometer (b). Those classified as bad by the full selection are indicated in red, and good subruns in blue. The selection cuts for this variable is also shown.

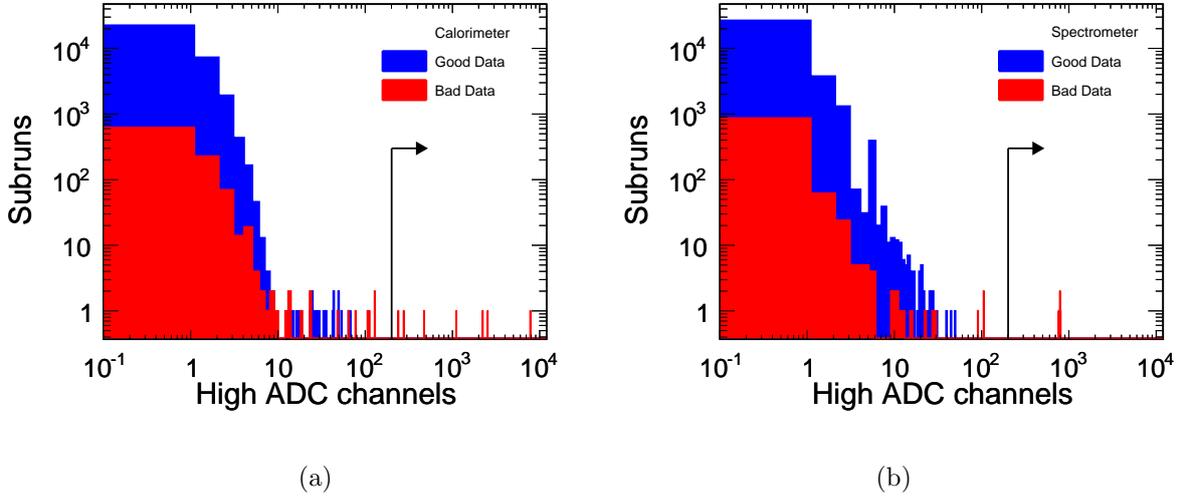


Figure 4.11: Distribution of channels with high ADC counts averages over the number of beam spills per subrun in the calorimeter (a) and spectrometer (b). Those classified as bad by the full selection are indicated in red, and good subruns in blue. The selection cuts for this variable is also shown.

than those for the true error bits, with 25 per subrun in the spectrometer and 50 per subrun in the calorimeter.

ADC Values

A large number of channels recording high ADC values for prolonged periods through the subrun can be indicative of electronics problems on the MENU scale. This failure mode is characterised by averaging the number of ADC counts above 100 (compared to the maximum ADC count of 255) over the number of spill triggers for each channel. The channel is flagged if the average exceeds 1000, and the subrun is discarded if the number of channels with high ADC averages exceeds 200, as shown in Figure 4.11.

SGATE Timing

As stated in the previous section, the SGATE trigger opens the window for data acquisition about $2 \mu\text{s}$ before the beam spill arrives in the detector, and it remains open for up to $20 \mu\text{s}$. The five or six ‘batches’ which make up a spill should then be well centred in the timing window, as indicated in Figure 4.12(a). If the trigger signal is not received, portions of the beam spill can be lost as shown in Figure 4.12(b), which if the

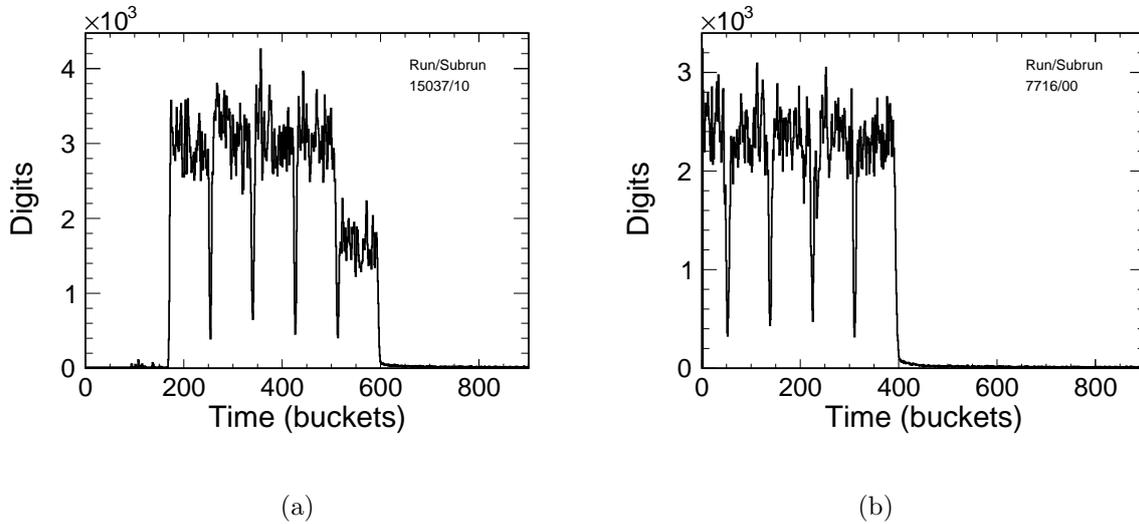


Figure 4.12: Beam spill timing distributions in the Near Detector in units of buckets (18.8 ns); for a spill well centred in the SGATE timing window (a) and one showing poor synchronisation with the trigger (b).

subrun were to be included would result in bad PoT counting for the physics analyses.

The shape of the TDC spectrum is utilised to determine if the beam spill is well synchronised with the timing window. As can be seen in Figure 4.12, the spectrum can be broken down into three regions; the 520-bucket SGATE and the plateaux either side of it, labelled as the ‘Low’ and ‘High’ regions. The average number of entries is calculated for each of the three regions, by counting the number of entries and averaging over the width of each region.

In the Low region no hits earlier than bucket 15 are used, to prevent spikes at zero which occur when a significant number of cold channels are observed (such as a whole crate malfunction) artificially inflating the average. This cut is applied to all data taken after May 2005, as prior to this date the window was only opened 20 buckets before the arrival of the spill. In the High region, the width is defined as 100 buckets after the end of the spill, as the true end of the spectrum is not easily discerned.

Under normal running conditions the average in the SGATE region should be well above zero, and above the average in both the low and high regions. Hence for the subrun to pass the timing cut, the average in the gate must be greater than 100 entries, and the ratio of the gate average to the Low and the High averages must be greater than 1.05 in both cases.

4.3 Implementation

The flagging of each Near Detector subrun as good or bad has been fully automated on the dedicated MINOS online machine hosted at Fermilab. The data quality variables defined above for each subrun stored in the offline database table *DbuNearRunQuality* in the *RunQuality* package. The *TpSingles* and *RawDigits* data streams output by the Data Acquisition system (DAQ) are automatically processed once an hour, and the data quality variables are then written to the database by *FillNearRunQuality* in the *CandMorgue* package.

The database table can be queried by the access tool *RunQualityFinder*, which takes a run and subrun number, or detector type (Near or Far) and timestamp. The MySQL database tools in the Minos offline software (*minosoft*) access the appropriate database table row, and returns a *RunStatus* bitmap encoding the problem which causes the subrun to fail the data quality selection, and an overall run status flag. The *RunStatus* enumeration is defined as below:

```
namespace RunStatus {
    typedef enum ERunStatus {
        kUnknown          = 0x00,    // unknown (no run)
        kOK                = 0x01,    // good subrun
        kBad               = 0x02,    // bad subrun
        kBadRunType        = 0x10,    // incorrect run type
        kBadReadout        = 0x20,    // incorrect readout
        kShort              = 0x100,   // too short
        kLowRate           = 0x200,   // too low event rate
        kHighRate          = 0x400,   // too high event rate
        kLIRate            = 0x800,   // too high LI rate
        kColdReadout       = 0x1000,  // cold readout
        kHotReadout        = 0x2000,  // hot readout
        kBadReadoutErrors  = 0x4000,  // bad readout errors
        kBadSpillWindow    = 0x10000, // bad timing window
        kHighADC           = 0x20000  // high ADC values
    } RunStatus_t;
}
```

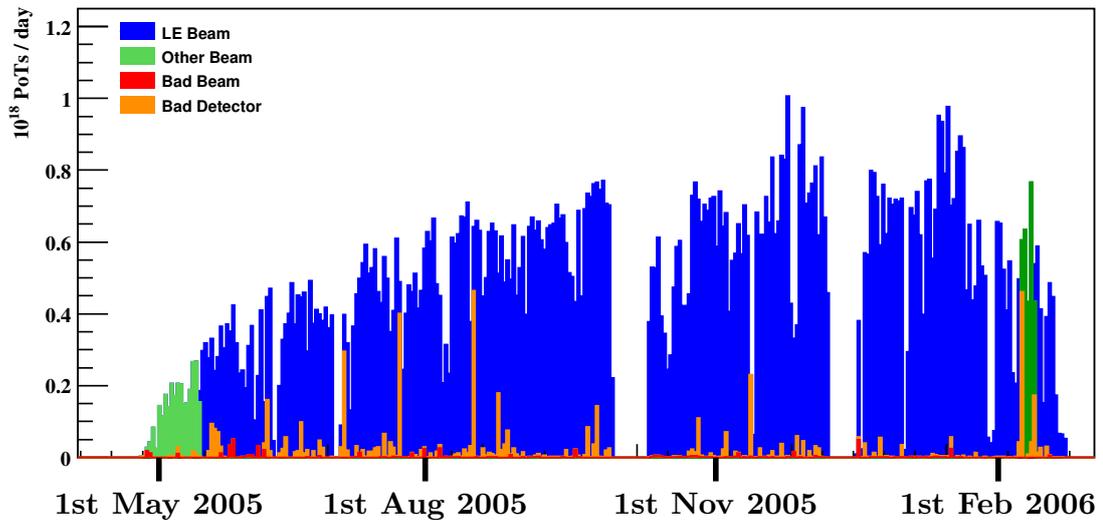


Figure 4.13: Distribution of PoT taken between March 2005 to February 2006, including the data period defined as Run I. Good data taken in the Low Energy configuration (LE) are shown in blue, and good data in other beam configurations in green. Data rejected due to bad beam operating conditions are shown in red, and due to bad detector conditions are shown in orange.

4.4 Rejected Data

Figures 4.13 to 4.16 show the distributions of PoT collected for all data collected between March 2005 and October 2009. The data flagged as bad due to beam operating conditions are shown in red, with data rejected due to the selection cuts detailed in this Chapter shown in orange. Overall 3.0% of the 33411 total subruns are considered unsuitable for inclusion in a physics analysis. The data quality procedure was applied to all data used in the analysis described in the following chapters.

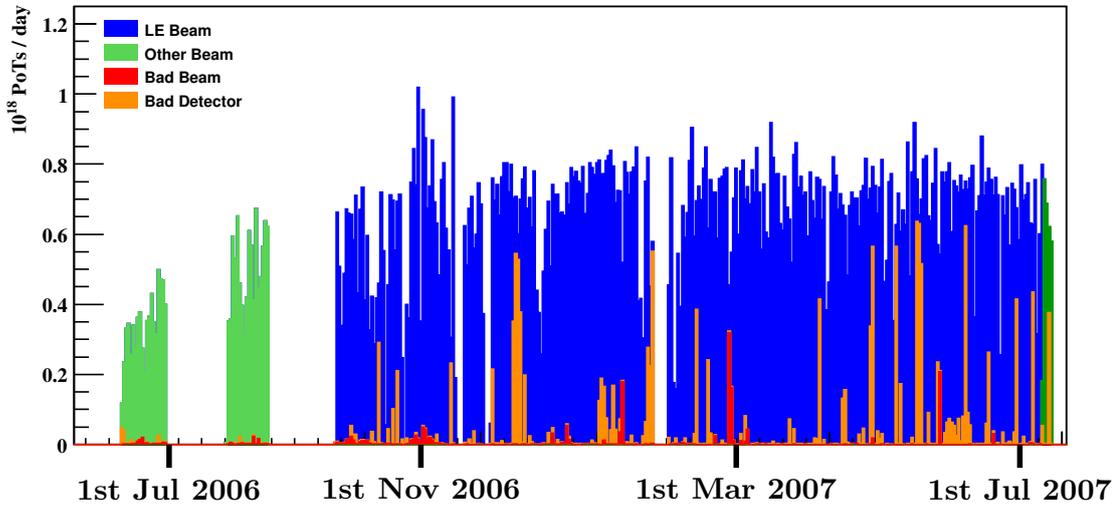


Figure 4.14: Distribution of PoT taken between June 2006 to July 2007, including the data periods defined as Run I-pHE and Run II. Good data taken in the Low Energy configuration (LE) are shown in blue, and good data in other beam configurations in green. Data rejected due to bad beam operating conditions are shown in red, and due to bad detector conditions are shown in orange.

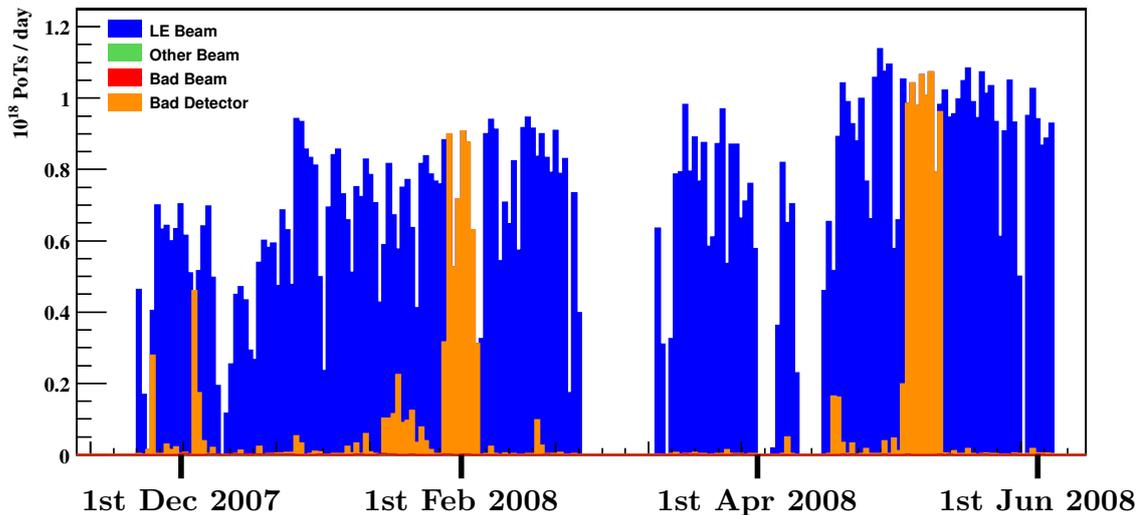


Figure 4.15: Distribution of PoT taken between November 2007 to June 2007, including part of the data period defined as Run III. Good data taken in the Low Energy configuration (LE) are shown in blue, and good data in other beam configurations in green. Data rejected due to bad beam operating conditions are shown in red, and due to bad detector conditions are shown in orange.

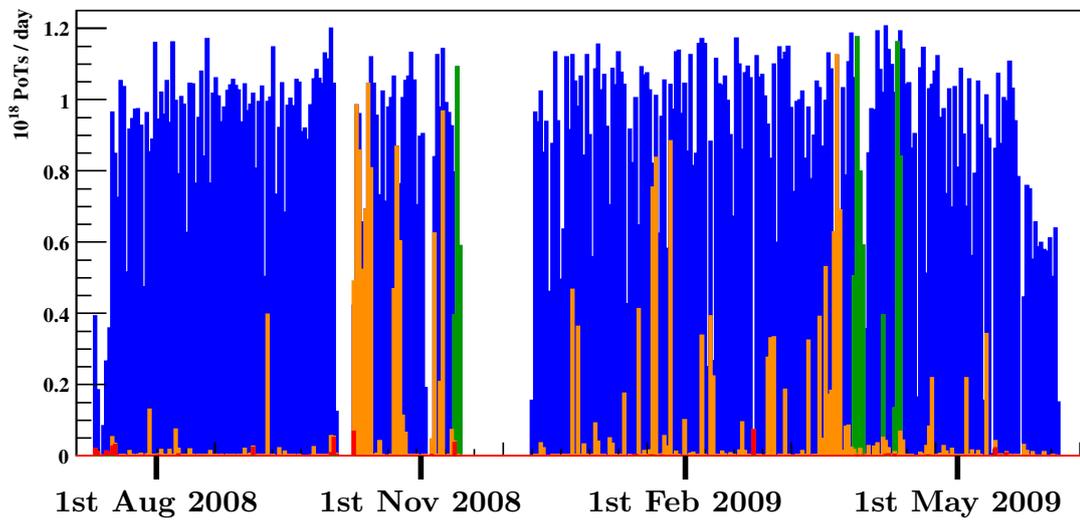


Figure 4.16: Distribution of PoT taken between July 2008 to June 2009, including the part of the data period defined as Run III. Good data taken in the Low Energy configuration (LE) are shown in blue, and good data in other beam configurations in green. Data rejected due to bad beam operating conditions are shown in red, and due to bad detector conditions are shown in orange.

Chapter 5

The Charged Current Analysis

5.1 Overview of an Oscillation Analysis

A series of steps must be performed in order to complete a MINOS neutrino disappearance analysis; the selection of a sample of appropriate events in the Near and Far Detectors and defining the neutrino energy for these events. A method is then required to utilise information from the Near Detector to obtain a Far Detector prediction, and finally to extract the best fit oscillation parameters and confidence limits, taking into account systematic errors.

The classes of events induced by the beam in the detectors are as follows: Charged Current (CC) ν_μ , $\bar{\nu}_\mu$, ν_e and $\bar{\nu}_e$, where electron neutrino events are both inherent in the beam, and a possible appearance signal; CC ν_τ and $\bar{\nu}_\tau$ appearance events in the Far Detector only, and Neutral Current (NC) events. Selection criteria summarised in Section 5.2 for this analysis aim to obtain pure separate samples of CC ν_μ and $\bar{\nu}_\mu$ events, where the main backgrounds are NC, ν_τ , and $\bar{\nu}_\mu$ or ν_μ events respectively.

The event energy is determined by summing the energy of the track, taken from the range of the muon or the track curvature, discussed in Section 3.3.3, and the hadronic shower energy, a new method for the calculation of which is detailed in Section 5.3 and [74]. The distributions of these event energies form the Near and Far Detector neutrino energy spectra. These spectra are not identical due to meson decay kinematics, the geometry of the beamline and the relative positioning of the Near and Far Detectors to where the neutrino beam is produced. Hence, an extrapolation method is utilised, to convert the measured Near Detector spectrum into a predicted Far Detector spectrum.

The two detector design has the advantage in minimising systematic uncertainties that affect the two detectors in the same way. Details of the extrapolation procedure are given in Section 5.4 and [12].

The Far Detector prediction is produced with a variety of oscillation parameters, which are varied across a grid in $(\sin^2(2\theta), \Delta m^2)$ space to obtain the parameters which minimise the likelihood of the observed data in the Far Detector given the prediction. A new fitting procedure has been developed for this thesis, and is detailed in Chapter 6. Additional analysis methods designed to improve the sensitivity of the result have also been implemented. A procedure to divide events into bins of energy resolution was developed for this thesis and is detailed in Chapter 7, and the inclusion of rock and events outside the fiducial volume to maximise statistics is described in Section 5.5 and [75].

5.2 Event Selection

A series of pre-selection cuts is applied to Near and Far Detector Monte Carlo and data in order to extract relatively pure samples of ν_μ and $\bar{\nu}_\mu$ CC interactions, removing cosmic data and some Neutral Current events prior to the main selection. The pre-selection also applies fiducial volume cuts in both detectors to mitigate geometric effects near the detector edges, and a series of data only constraints on good beam operating conditions and data quality. The main selection is then applied to define the fiducial samples; CC ν_μ interacting to produce muon tracks with negative curvature, and CC $\bar{\nu}_\mu$ comprising muon anti-neutrinos with positive curvature tracks. The pre-selection and selection cuts are detailed here, and summarised in Table 5.1.

5.2.1 Pre-Selection

- **Require a track:** Require all events to possess at least one reconstructed track. This cut removes a large number of potential NC background events where no track is reconstructed.
- **Track fit Successful:** Ensure the track fitter summarised in Section 3.3.3 has successfully converged. Tracks typically fail the filter when their hits result in an unphysical noise covariance matrix, which accounts for uncertainties in the muon

state vector at each fit step due to scattering and ionisation energy loss between the planes under consideration.

The remaining variables are applied to data only:

- **Data Quality:** Data quality selectors are applied in the Near and Far Detectors, to ensure good operating conditions when the data was taken. Data quality cuts for the Near Detector are detailed in Chapter 4, with Far Detector quality cuts summarised in [73].
- **Good Beam:** The beam is required to have fallen within certain parameters for normal operation for the data to be utilised, as summarised in [76]. Checks include the size of the beam spot, typically the RMS width and mean position to be centred on the target end, and the status of the focussing horns, with cuts on acceptable horn current values.
- **Light Injection:** Ensuring light injection calibration pulses did not occur during data taking, mimicking high energy deposits in the strips.
- **Muon direction:** A cut on the direction cosine of the reconstructed muon track, relative to the event vertex and the beam direction. If $\cos(\theta) < 0.6$, the track is likely to have resulted from a cosmic neutrino and the event is discarded.
- **Spill Window:** To minimise the cosmic ray background, the detector only records data during a window with limits of $2 \mu\text{s}$ either side of the $10 \mu\text{s}$ beam spill duration. The centring of the beam spills in this window is a check carried out as part of the data quality procedure.

5.2.2 Fiducial Volume

The boundaries of the fiducial volume in each detector are designed to ensure that if the primary event vertex lies within the fiducial limits, and the hadronic shower is fully contained. The cuts differ between the two detectors due to detector geometry and event acceptance requirements. In the Far Detector, the volume is set to be as large as possible to maximise the data statistics; the higher event rate in the Near Detector permits a smaller and more conservative definition.

In the Near Detector, the fiducial volume is a cylinder, angled downwards by 3° to run parallel to the beam. The event vertex is required to be within 80 cm of the beam

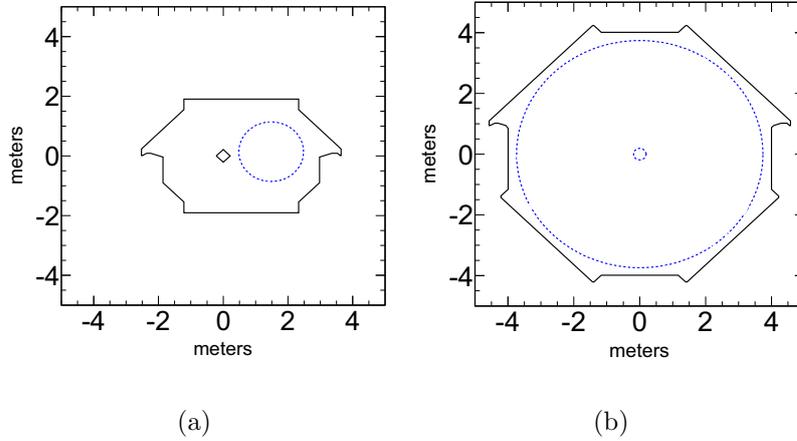


Figure 5.1: Fiducial volumes in the (a) Near and (b) Far Detectors. Events with their primary vertices contained within these volumes are considered for the primary analysis. The Rock and Anti-Fiducial analysis utilises events outside these limits.

centre, and between the 14th and 68th planes. This results in the containment of showers within the calorimeter section of the detector.

In the Far Detector, the event vertex must lie within an annular cylinder; between 40 cm and $\sqrt{14}$ m from the centre of the beam line. This ensures no events less than 35 cm from the edge of the detector are included in the main analysis sample. The vertex must also not lie within the first four planes of either supermodule, the last eight planes of supermodule one or the last 20 planes of supermodule two. This discards events which originate outside the detector and enter through the front face, and those which lose a proportion of their hadronic energy out of the back face.

Diagrams of the fiducial volumes are given in Figure 5.1. The Rock and Anti-Fiducial (RAF) analysis described in Section 5.5 utilises the events which lie outside these volumes.

5.2.3 Particle Identification (PID)

In order to perform a disappearance analysis, it is vital to identify a highly pure CC sample, removing NC events without sacrificing efficiency. The ν_μ and $\bar{\nu}_\mu$ samples are treated as one category in the PID, with the same selector applied to all CC events.

The unique topology of a CC event is a muon track extending beyond a hadronic shower. In cases of low inelasticity (y -value) these events are clearly distinguishable to the reconstruction. However, NC events which consist entirely of hadronic energy deposits can be assigned falsely reconstructed short tracks, which are difficult to distinguish from low energy or high inelasticity CC events. The purpose of the Particle-Identification (PID) algorithm is to evaluate how likely it is candidate tracks are truly muons or are due to mis-identified hadronic energy deposits. As shown in Table 5.1, the first pre-selection cut is to require all events possess at least one reconstructed track of at least 10 planes in length before the PID is applied.

For the analysis presented in this thesis, two PID algorithms were utilised. The first is the selector used in the previous published analysis [69], described in detail in [77]. The second was developed for the most recent published analysis [39], and aimed to improve selection efficiency at low energies ($E < 5$ GeV). It maximises sensitivity to alternative disappearance models and is described in detail in [78]. Both algorithms utilise a k -Nearest-Neighbours (k NN) method to discriminate between CC and NC events.

The k NN algorithm takes a training sample of Monte Carlo, constructing an n -dimensional space for the n reconstructed variables used to discriminate between event types. The space is then populated with these training events. A simulated event is placed into the space, and compared with its k nearest neighbours. The proportion of the k neighbours that are CC events in truth (for example 1/5th of events) is defined as the PID variable x for that test event (in this case, $x = 0.20$). Therefore the parameter is valued between zero and one for each candidate event, with NC-like events assigned a low x value and CC-like events a high x value.

The primary selector R was tuned so the input variables and values of k and n maximised oscillation sensitivity. A Monte Carlo study was performed to determine the cut position, with fake data generated at specific oscillation parameters. The optimised cut position was defined to be that which gave the smallest errors on the mixing parameters when an oscillation fit was performed [79]. The optimum values were found to be $k = 80$, $n = 4$, for the input variables defined as follows:

- **Number of planes:** CC interactions produce long, curving muon tracks as they deposit energy in the detector. Falsely reconstructed tracks within hadronic showers typically have a smaller range in the z -direction.
- **Mean energy deposited per strip:** Defined as the total pulse height in the

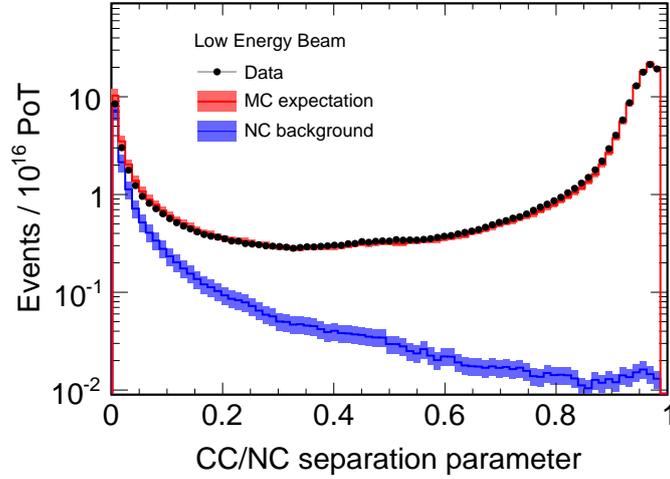


Figure 5.2: Distribution of the optimised PID selection variable for data and MC in the Near Detector, for the primary selector R .

event, divided by the number of strips recording hits, after the first 30% of planes after the track vertex are discarded. Muons are minimally-ionising (MIPs) and so deposit an almost constant amount of charge in each plane. The multiple particles comprising the shower travel more slowly and deposit larger amounts of energy in each plane.

- **Signal fluctuation parameter:** The first 30% of planes nearest the track vertex are discarded, and the remaining planes are sorted by pulse height. The parameter is defined as the ratio of the mean of the lowest 50% of pulse heights to the mean of the upper 50%. Hadronic showers identified as tracks exhibit larger fluctuations in their energy deposits than muons, and tend to a lower value of the fluctuation parameter. A larger values implies more uniform energy deposits.
- **Transverse profile parameter:** This variable quantifies the separation of the track from the rest of the event, defined as the fraction of event energy within a 4-strip envelope around the reconstructed track, excluding the 50% of planes closest to the vertex.

Figure 5.2 shows the distribution of the primary selector output, R , for Near Detector Monte Carlo and data. When optimised, the cut position for CC events is placed at $R > 0.30$ when this selector is used in isolation.

The second selector, J , is applied to those events which fail the primary case, and

aims to recover events with energies below 5 GeV. It is constructed in the same fashion as the PID above, with four variables input into a k NN. It maximises the $\Delta \ln \mathcal{L}$ of pure decay and decoherence models with respect to oscillations for high statistics fake data, constrained by the requirement that the sensitivity to oscillations is not worsened. Due to the low energy nature of the events a track is still required, but the cut on the minimum number of planes it traverses is removed.

- **Number of planes:** Defined as in the primary selector R .
- **Pulse height in the last 5 planes of the track:** A value larger than that of a minimally-ionising particle is expected if a hadronic track terminates in a nuclear interaction.
- **Track scattering:** Muon-like hadronic tracks undergo nuclear interactions within the detector, and hence exhibit more scattering than muon tracks. Additionally, tracks reconstructed from unrelated hits are likely to be less smooth than tracks from a single particle. Two variables are constructed from Pearson correlation coefficients, in the $u - z$ and $v - z$ views to quantify the degree of scattering.

Figure 5.3 shows the distribution of the secondary selection, J , in the Near Detector.

The final PID is constructed as a logical disjunction between the two selectors, if an event passes either R or J it is included in the CC samples. The cut positions on both R and J are re-optimised when working in conjunction, which relaxes the cut on R :

$$R > 0.25 \vee J > 0.5 \quad (5.1)$$

The overall efficiency and level of NC contamination in the Far Detector as a result of applied the combined PID is shown in Figure 5.4, compared to the single primary selection applied in [69]. With the low energy selector the purity decreases and the efficiency increases; as a result an improvement in alternative model discrimination is seen but the oscillation sensitivity is unchanged.

5.3 k NN Shower Energy

The calorimetric shower energy measurement used in previous MINOS analyses [69] permits large potential gains in oscillation sensitivity with a hypothetical improvement

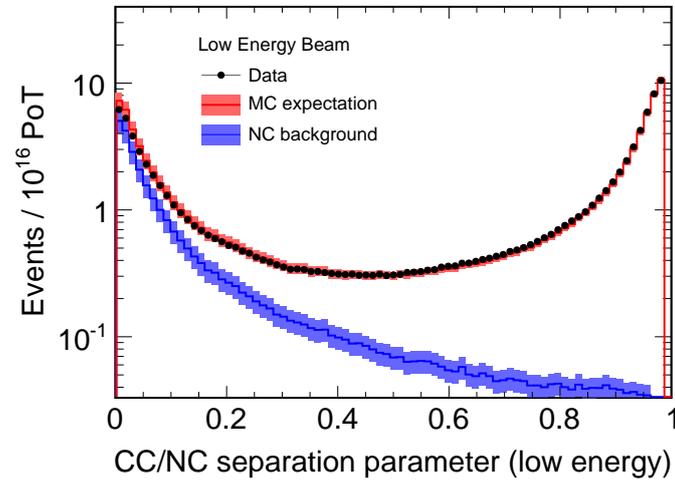


Figure 5.3: Distribution of the low energy optimised PID selection variable J , for data and MC in the Near Detector.

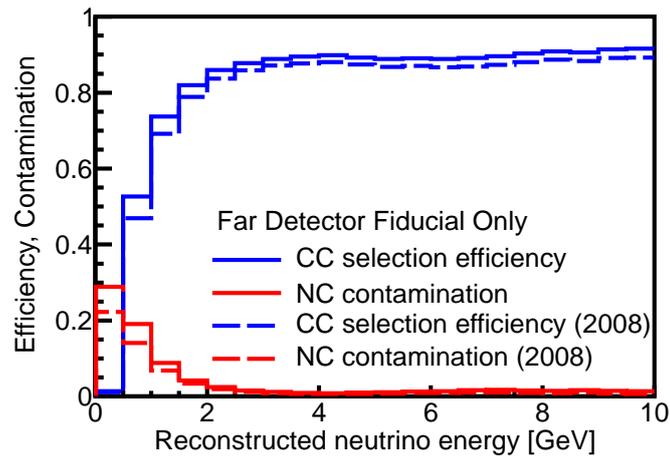


Figure 5.4: Efficiency and contamination of the k NN selection in the Far Detector. The dashed line indicates the single primary selector used in the 2008 published analysis, with the solid line the combined selector applied for the analysis described here.

Pre-Selection	Main Selection
Monte Carlo and Data	
Require a track <code>ntrk > 0</code>	Cut on PID variable $R > 0.25 \vee J > 0.5$
Track fit Successful <code>trkfitpass==true</code>	Cut on charge-sign ν_μ sample: <code>charge== -1</code> $\bar{\nu}_\mu$ sample: <code>charge== +1</code>
Fiducial Volume	
Data Only	
Data Quality <code>isGoodDataQuality==true</code>	
Good Beam <code>goodBeamToUse==true</code>	
Light Injection <code>isLI==0</code>	
Muon direction <code>dirCosNu > 0.6</code>	
Spill Window $-2 \mu s < \text{GoodTimeToNearestSpill}$ $\wedge \text{GoodTimeToNearestSpill} < 12 \mu s$	

Table 5.1: Summary of pre-selection and selection cuts applied to Monte Carlo and data samples, prior to an analysis being performed.

in energy resolution. To quantify the potential of this hypothetical improvement, the reconstructed neutrino energy is adjusted as follows, to reduce the difference between the true and reconstructed energies:

$$E' = E_{reco} + \frac{E_{reco} - E_{true}}{f} \quad (5.2)$$

where $f = 1.1$ and 1.5 for a 10% and 50% improvement respectively. Figure 5.5 shows the gains in sensitivity resulting from these improvements, compared to the neutrino energy calculated using the calorimetric estimator, and an oscillation fit performed in true neutrino energy. By reconstructing a shower energy that results in a more clearly defined oscillation dip, the sensitivity to $\sin^2(2\theta)$ and Δm^2 is improved; the range of values in parameter space matching the position and depth of the dip is reduced. For

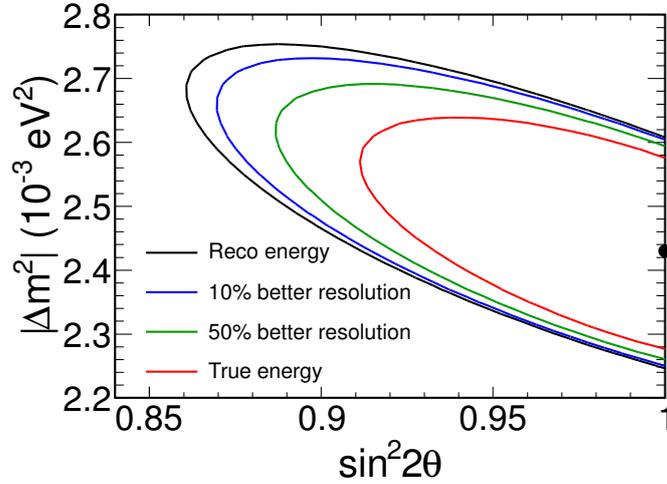


Figure 5.5: Comparison of oscillation sensitivities for an oscillation fit performed using the calorimetric shower energy reconstruction (black), for 10% (blue) and 50% (green) improvements in neutrino energy resolution, and true neutrino energy (red). The Monte Carlo exposure was scaled to 7.2×10^{20} PoT, and oscillated at $\sin^2(2\theta) = 1.0$, $\Delta m^2 = 2.43 \times 10^{-3} \text{ eV}^2$.

this analysis, a k NN approach was taken to determine the shower energy, the details of which are given in full in [74].

The general k NN method has already been described in Section 5.2.3; in this case three variables were chosen which closely correlate with shower energy, and the optimised value of k was found to be $k = 400$. Seven variables were considered, and the resulting oscillation sensitivity of Δm^2 at 1σ was taken as the figure of merit. The three variables that were selected are:

- **Number of planes:** The number of planes in the primary shower.
- **Calorimetric shower energy:** The combined calorimetric shower energy of the first two reconstructed showers, if a second was present in the event.
- **Deweighted shower energy:** The total deweighted shower energy within 1 m of the track vertex. The deweight is an energy dependent function that alters the relative importance given in the reconstruction to the total number of photoelectrons recorded and the number of strips hit [80].

The k NN procedure produces a spectrum closer to the true distribution in Monte Carlo, as shown by Figure 5.6. The gain is particularly significant below 1 GeV, where

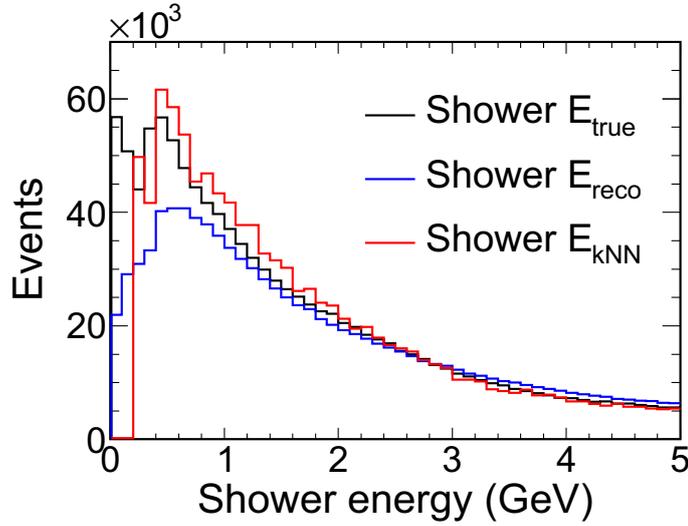


Figure 5.6: Shower energy distributions for different methods of determination. The black line denotes the true shower energy from the MC events, the blue the calorimetric shower energy used previously, and the red the k NN shower energy. The k NN matches the true distribution more closely, especially at energies below 1 GeV.

the calorimetric energy is a poor estimator of the true distribution for the simulated events. The improvement in resolution in 0.5 GeV bins of true neutrino energy is shown in Figure 5.7. The k NN energy does however exhibit a hard cutoff below 0.25 GeV, where events cannot be assigned an energy greater than zero and less than 0.25 GeV. This is investigated further in Chapter 8.

A series of energy corrections are applied before the final result is determined, due to a bias introduced by the k NN procedure at energies below 5 GeV. The Monte Carlo training sample is not populated with events uniformly distributed in energy, it is produced with the majority of events around the peak at 3 GeV as would be expected. As a result, the 400 nearest neighbours for a very low E_{shw} event are not distributed isotropically, biasing the mean shower energy of these events to a higher value. The converse is true for events above the beam peak, which results in the ratio of $E_{k\text{NN}}/E_{\text{true}}$ shown in Figure 5.8.

The polynomial fit shown in Figure 5.8 is applied as a weighting function in $\log E_{k\text{NN}}$ to attempt to flatten the distribution, so $\left\langle \frac{E_{k\text{NN}}}{E_{\text{true}}} \right\rangle \simeq 1$. Three iterations of the polynomial correction are applied to obtain the final shower estimation, where the deviations from a mean valued at one are below the level of one percent. Below 500 MeV, the final energy resolution obtained is improved by 50% from the calorimetric estimator [74].

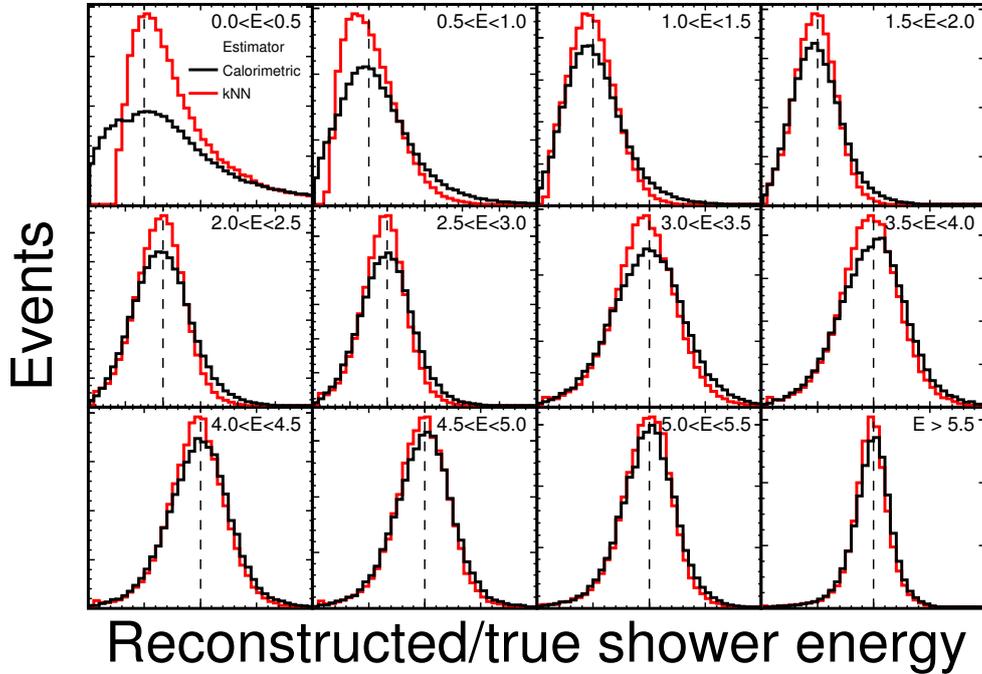


Figure 5.7: Distributions of calorimetric (black) and k NN shower energy (red) resolutions in bins of 500 MeV. The dashed line indicates the position of $E_{reco}/E_{true} = 1$ for the calorimetric energy. The misalignment in the lowest two energy bins is clearly visible, and the improvement resulting from the k NN method.

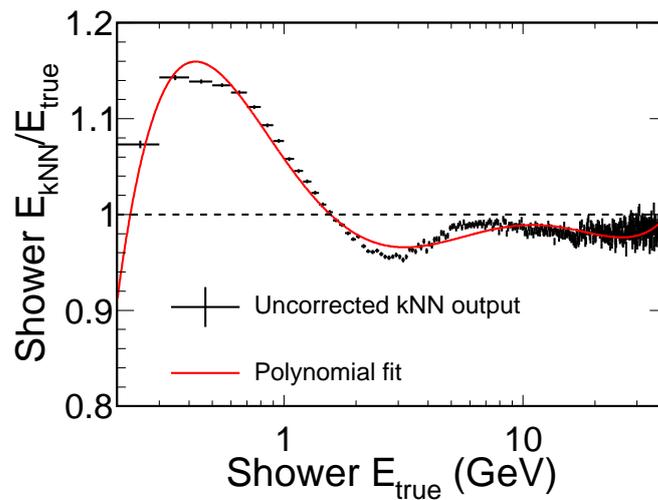


Figure 5.8: The bias from the k NN shower energy estimator, and polynomial correction applied to remove this bias. Taken from [74].

5.4 Extrapolation

The two detector design of MINOS was adopted in order to use information gathered at the Near Detector to minimise systematic uncertainties in the energy spectrum of the Far Detector. Certain sources of uncertainty such as neutrino flux, interaction cross-sections and the measurement of the hadronic shower energy affect the data at both detectors, and extrapolating the spectrum measured in the Near Detector to the Far Detector mitigates the impact of these uncertainties.

The extrapolation used here, the “beam matrix” method, is identical to that used in the published MINOS analyses [39, 69]; with the addition in the latter case of resolution information, discussed in Chapter 7. A full and detailed description of the beam matrix method and its derivation can be found in [12].

The neutrino flux experienced at the Near and Far Detectors is not identical, due to the geometry of pion and kaon decay in the NuMI beam pipe, as illustrated by Figure 5.9. Neutrinos are emitted isotropically in the pion’s rest mass frame. However, the neutrino flux in the lab frame is a function of the angle, θ , between the parent’s direction of flight and that of the neutrino:

$$\frac{dN}{d \cos \theta} = \frac{1}{2\gamma_p^2 (1 - \beta_p \cos \theta)^2} \quad (5.3)$$

where β_p is the velocity of the parent and γ_p is its Lorentz factor. The Near Detector subtends a greater angle than the Far Detector due to its proximity to the decay volume, thus neutrinos from lower energy pions decaying at larger angles will give rise to interactions in the Near Detector but miss the Far Detector. This leads to the differences in the energy spectra shown in Figure 5.10, where a set of parent particles with the same decay vertices and momenta (shaded areas) contribute sets of neutrinos with different energies in the Near and Far Detectors. The peak of the energy spectrum is shifted higher in energy in the Far Detector, as more neutrinos emitted from higher energy pions are produced at smaller angles.

In order to use the Near Detector spectrum to correctly predict the Far Detector spectrum, the Monte Carlo is used to construct a beam matrix, which contains the relevant kinematics and geometric information to transfer between them. The format of this matrix is shown in Figure 5.11, where for each true neutrino energy recorded at the Near Detector on the x axis, each column contains the true energy distribution

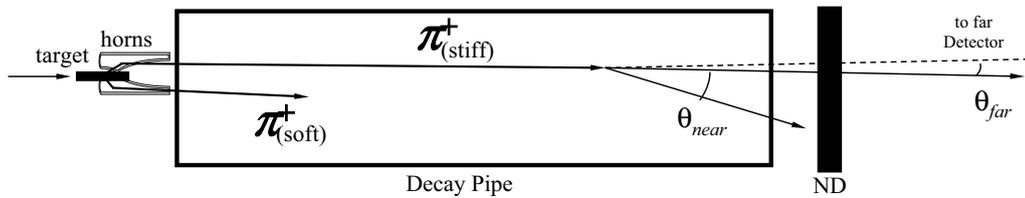


Figure 5.9: Illustration of neutrino parent decays in the NuMI beam pipe. Each parent will typically produce neutrinos at a wider range of angles that intersect the Near Detector than will intersect the Far Detector.

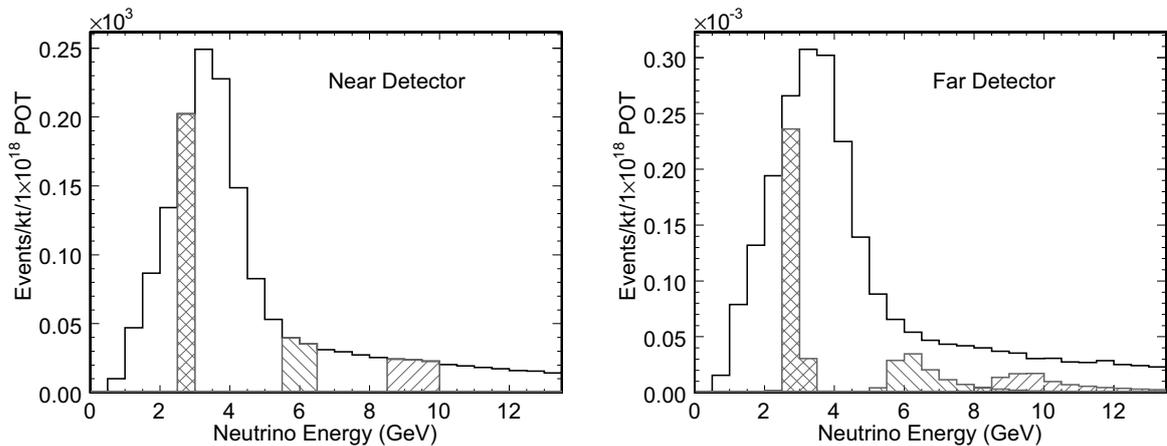


Figure 5.10: The relationship between neutrino energy spectra at the Near and Far Detectors. The shaded regions in each plot indicate neutrinos resulting from parents with the same decay vertices and momenta in the beam pipe. Taken from [74].

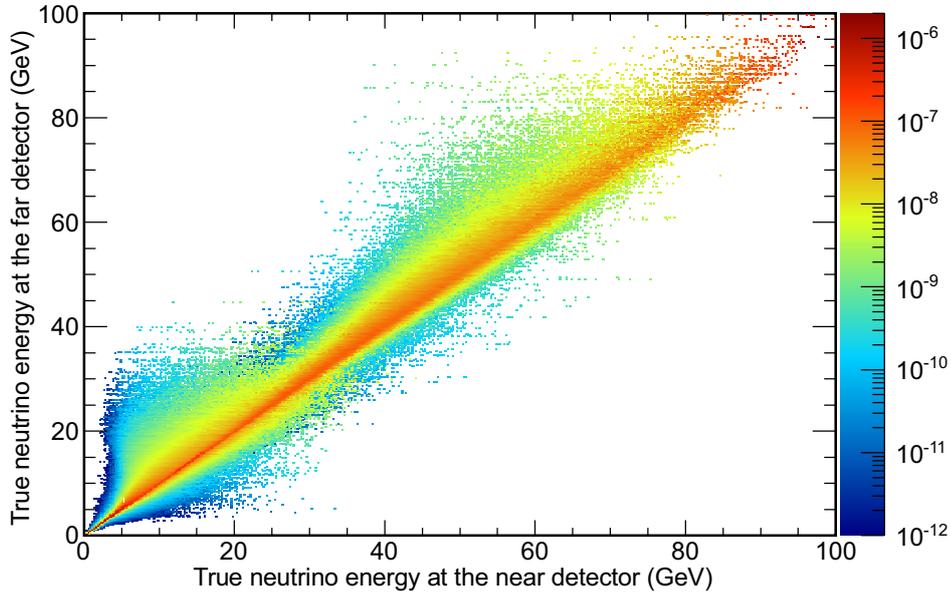


Figure 5.11: The beam matrix for neutrinos from Run III Monte Carlo, which multiplies the measured Near Detector flux to give a Far Detector flux prediction. Each cell represents the mean number of neutrino events expected at the Far Detector for one neutrino event at the Near Detector.

from the same parent particles at the Far Detector, normalised to one Near Detector neutrino. The ND flux can then be multiplied directly by the matrix to obtain the Far flux prediction.

The matrix is populated by considering a large number of neutrino parents and matching them to a randomly selected interaction vertex within the Near Detector, which fixes the resulting neutrino energy observed at the ND. The same parent is then decayed towards the Far Detector, and a matrix cell filled at these energies, weighted by the probability that these decay directions would occur and the neutrino would interact in each detector. A matrix is produced for each Run period separately to account for differences in the beam state, and for neutrinos and anti-neutrinos separately¹.

The beam matrix is designed to relate the measured fluxes between the two detectors. A series of additional corrections must be applied as part of the extrapolation procedure, to convert the Near Detector reconstructed neutrino energy spectrum to a measured flux, and the predicted Far Detector flux into a reconstructed energy spectrum. The corrections are all determined from the Monte Carlo simulations at the appropriate

¹Matrices are also split at this stage into bins of resolution, the concept of which is detailed in Chapter 7.

detector.

Efficiency and purity corrections (see Equation 8.11) are carried out on a bin-by-bin basis, and determined separately for each detector in order to remove background contamination from ν_τ and wrong sign events, and account for reconstruction efficiency. The energy spectra are converted from reconstructed to true energy (and vice versa) by multiplying or dividing by a Monte Carlo populated reconstructed vs true energy matrix. The true spectrum is divided by the CC ν_μ cross-section², the mass of the detector fiducial region and the total number of Protons-on-Target for a true flux prediction per PoT.

After the beam matrix has been applied, the same corrections are applied in reverse order to output the Far Detector prediction in reconstructed energy. At the true energy stage, oscillations or alternative disappearance models can be applied to compare directly to data, but this is not the method used to obtain the best fit oscillation parameters for this analysis (see Chapter 6). The full extrapolation procedure is summarised in Figure 5.12.

5.5 Rock and Anti-Fiducial Events

Previous analyses have disregarded the classes of events which have their true vertices either in the rock surrounding the Far Detector, or in the regions outside the fiducial volume, as shown in Figure 5.13. The muon resulting from rock interactions can enter the detector, but little or none of the hadronic energy is recorded and beam events are difficult to distinguish from those originating in the rock. For anti-fiducial events, poor containment of the shower results in energy loss from the edges or faces of the detector planes.

These events are included as an additional class to the CC ν_μ and $\bar{\nu}_\mu$ samples for this analysis, providing an exposure of comparable size to that of the fiducial events and providing a measurable improvement in oscillation sensitivity (as shown in Figure 6.17 in Chapter 6). For a complete discussion of the development of the treatment of RAF events, see [75].

Due to the lack of shower containment in the RAF sample, the visible track energy

²The model within NEUGEN is used here, and the cross-section is the sum of quasi-elastic, resonance and deep inelastic scattering cases.

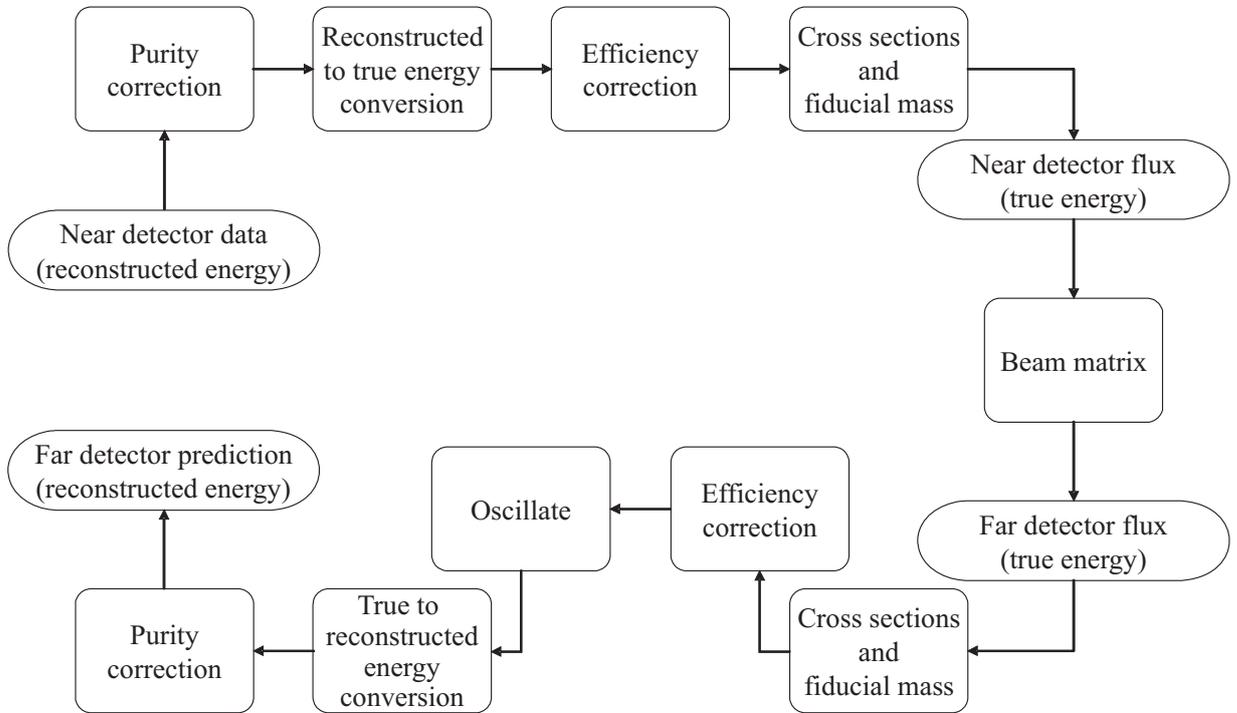


Figure 5.12: Flowchart summarising the beam matrix extrapolation procedure, which converts a measured Near Detector neutrino spectrum into a predicted Far Detector spectrum. For details of each step, see the main text.

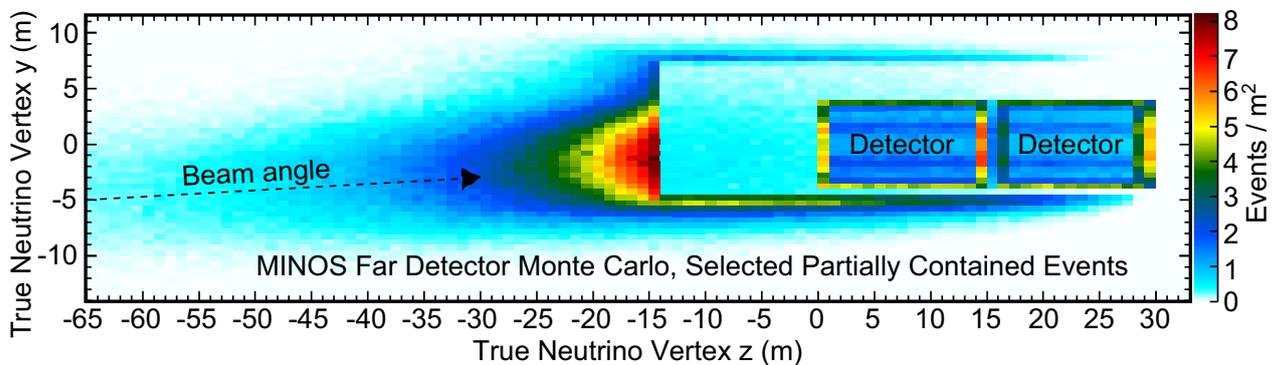


Figure 5.13: Distribution of the true vertices of events in the rock and anti-fiducial sample. The majority of events occur in the rock upstream of the cavern wall, which is clearly visible.

provides the maximum amount of information about the neutrino energy and so the reconstructed shower energy is not utilised. Similar pre-selection cuts are applied as summarised in Table 5.1, inverting the fiducial volume requirement, discarding anti-fiducial events near the coil hole and relaxing the need for the track fit to have converged. No charge sign cut is applied, so the backgrounds considered for this sample are NC and ν_τ events only.

To identify muon-like tracks from hadronic showers penetrating the anti-fiducial regions, cuts are made on the same PID variables R and J . However, the cuts are varied depending on which one of six “regions” in the Far Detector the first track hit of the event is recorded. Events with their true vertices in the anti-fiducial region will exhibit better energy resolution to those originating in the rock, and so dividing by region is analogous to dividing the fiducial sample into energy resolution bins as detailed in Chapter 7. The six regions are defined as follows, and shown schematically in Figure 5.14.

- **Front face:** Predominantly rock events, with hits in the first four planes of supermodule one.
- **Outside radial edge region, rock or detector-like:** Events which lie outside $\sqrt{14}$ m from the beam centre are classified as rock or detector-like, based on where the first hit of the track originates. If recorded in the outermost strip the true event vertex is more likely to be within the rock, and if recorded several strips into the detector volume the event is assumed to have originated in the detector.
- **Supermodule gap and gap edge:** Events which occur in the last eight planes of the first supermodule or the first four planes of the second. Events in this region but which also lie outside the radial cut are defined as a separate region.
- **Back face:** Events with vertices lying in the last 20 planes of the second supermodule, likely to have poorly contained transverse showers.

The RAF analysis also utilises information from the Near Detector in order to correct the neutrino flux, but does not use rock and anti-fiducial events recorded in the Near Detector directly. The differing cavern geometries, poor modelling of the Near Detector mass in the cavern and differences in rock density and composition make a direct extrapolation technically demanding, so the Near Detector fiducial events are used. Data and Monte Carlo spectra are run through the extrapolation procedure summarised in Figure 5.12, as far as obtaining the Far Detector flux prediction. The ratio of the data

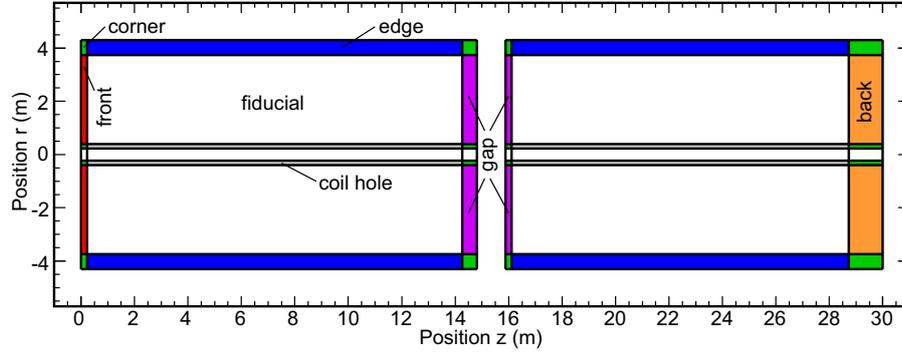


Figure 5.14: Diagram of the anti-fiducial regions used for the RAF analysis. Events with hits in the coloured regions are assigned to the categories detailed in the main text.

flux to the Monte Carlo flux is then used to correct the Far Detector RAF Monte Carlo in true energy. This procedure is considered to be analogous to the fiducial beam matrix method, the corrections not separated into separate steps but implied in the use of the Monte Carlo flux.

The improvements in oscillation sensitivity resulting from the PID, k NN shower energy and inclusion of RAF events discussed here are given in the following Chapter, which details a new fitting method developed for this thesis. When the new fitting procedure is implemented with the analysis methods detailed here and in Chapter 7, it provides the ν_μ disappearance analysis results given in Chapter 9.

Chapter 6

Fitting the Far Detector Data

The purpose of an oscillation fit is to compare Far Detector predicted neutrino energy distributions with data, and extract a measurement of the best fit parameters. The comprehensive oscillation fitting framework `GhostFitter`, developed for this thesis, is designed to fit in parallel any number of *samples*, where a sample is defined by the event selection which is applied, and which is affected by its own specific backgrounds. Systematic errors are incorporated both uniquely for each sample and in common across some or all samples where appropriate.

The samples that are considered in this thesis are Charged Current Negative Curvature events, composing the main sample of muon neutrino interactions (ν_μ); Charged Current Positive Curvature events ($\bar{\nu}_\mu$) which comprises muon anti-neutrinos; and the Rock and Anti-Fiducial sample (RAF), including all muon neutrino events which occur outside the standard MINOS detector fiducial volume.

A separate sample of Neutral Current events was also considered, but it was found that while including this sample did marginally improve the sensitivity to Δm^2 ($< 3\%$), electron neutrino appearance in the NC sample added a level of complexity that was not justified for the small improvement in the sensitivity contour [81].

Within each sample the fitting software contains a number of independent *vessels*. A vessel is specifically defined in this context as a particular (unoscillated) subset of events in Monte Carlo predictions and in data; each sample contains different vessels which are summarised by each row in Table 6.1. As an example, a vessel in the ν_μ sample could be defined as a set of histograms containing all the events with the best 20% neutrino energy resolution (Chapter 7) from the Run I data taking period.

To perform a measurement of the disappearance parameters, the contents of each prediction vessel are oscillated as appropriate at varying sets of oscillation parameters, have systematics errors applied if required and are finally summed to produce a single prediction histogram per vessel. The output histogram is then compared to the histogram residing in the corresponding data vessel, by defining a log-likelihood between them. The measured parameters are those for which the predicted vessels and the data vessels most closely agree.

ν_μ Fiducial	$\bar{\nu}_\mu$ Fiducial	ν_μ RAF
Run I $\times N_{Res.}$	Run I	Run I rock detector
Run II $\times N_{Res.}$	Run II	Run II rock detector
Run III $\times N_{Res.}$	Run III	Run III rock detector
Run I-pHE $\times N_{Res.}$		Run I-pHE rock detector
Backgrounds		
NC events	NC events	NC events
ν_τ events	ν_τ events	ν_τ events
$\bar{\nu}_\mu$ events	ν_μ events	

Table 6.1: Details of vessels included in the selected event samples used for this analysis, where each sample is defined as described in Chapter 5, and afforded a column above. The ν_μ fiducial events are also split by resolution information in N_{res} bins as detailed in Chapter 7. Events with their vertices reconstructed in the rock or in the detector anti-fiducial region are extrapolated as two separate categories and summed at the Far Detector to give one overall prediction for the RAF sample. The backgrounds considered for each sample are also shown.

6.1 Fit Mechanism and Implementation

In order to perform the comparison with data, the expected spectra at each point in parameter space must be generated, including all backgrounds and with any combination of systematic shifts. The framework has been developed to utilise a series of

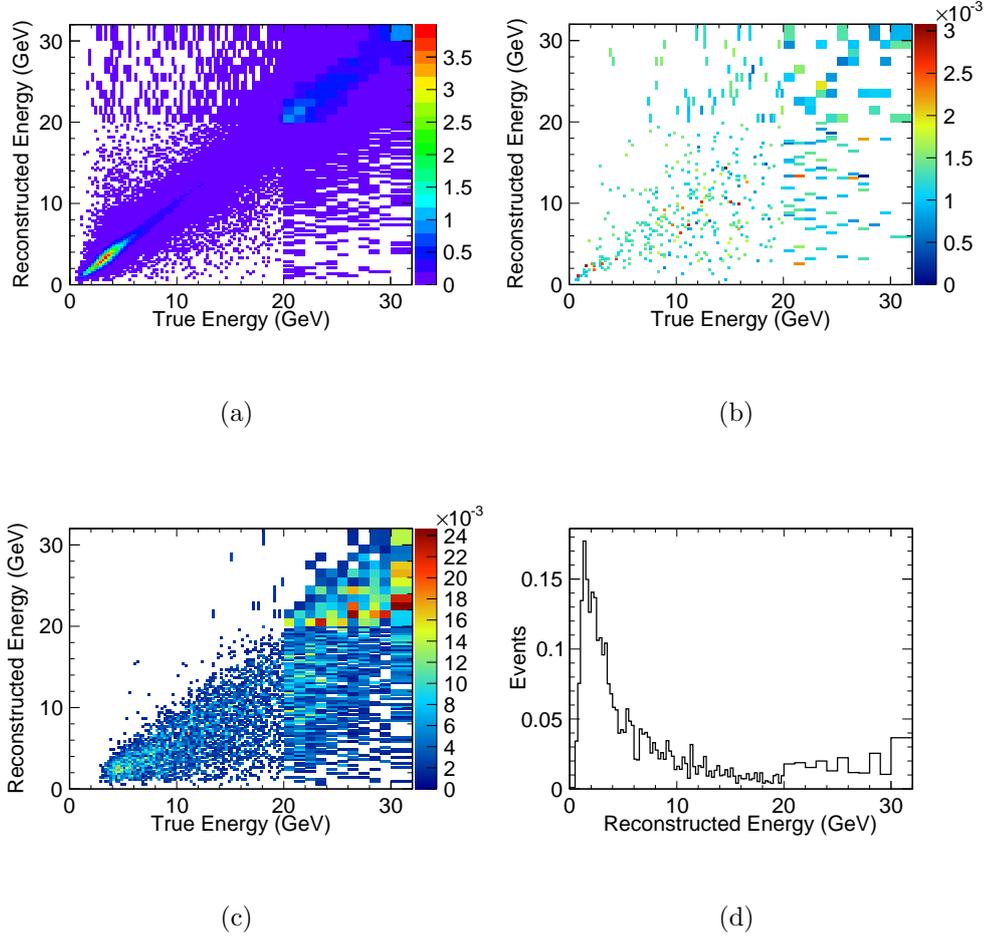


Figure 6.1: Example of the Monte Carlo ‘template’ histograms for the ν_μ sample and the vessel containing all Run I ν_μ selected events; for the true ν_μ events (a), true $\bar{\nu}_\mu$ events (b), true ν_τ events (c) and the NC background events (d). The input histograms correspond to the first column of Table 6.1.

two-dimensional histogram ‘templates’ to perform this task, which are stored in the vessels and samples defined previously. These consist of true neutrino energy against the appropriate energy variable for the sample as shown in Figure 6.1, binned as detailed in 6.1.1, and extrapolated from the Near to the Far Detector using the Beam Matrix Method outlined in Chapter 5. This enables the expected spectrum at any point in parameter space to be uniquely recovered by weighting by survival probability using the true neutrino energy.

A schematic of the 4-dimensional C++ Standard Template Library (STL) vector structure which holds the templates is shown in Figure 6.2; one sample is shown con-

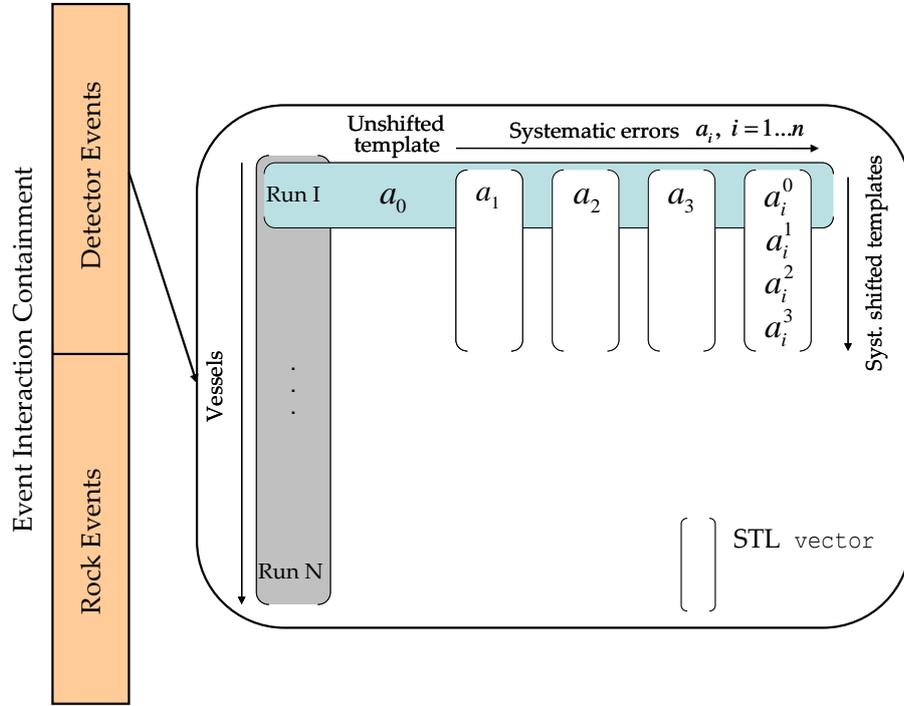


Figure 6.2: Schematic of the C++ sample and vessels container within the `GhostFitter` software, including systematic errors. Events are initially split by their fiducial containment, with the rock vector structure empty for the non-RAF samples; an instance of the 3-dimensional structure shown is stored for each of the event containment types. Each vessel contains an unshifted template, in addition to a series of systematically shifted templates for each source of systematic error. Each background contribution (NC, ν_τ , $\bar{\nu}_\mu$) to each sample has an identical dedicated structure.

taining a series of vessels, including the storage of systematic errors templates. Each background contribution (such as Neutral Currents, ν_τ , $\bar{\nu}_\mu$) to each sample has an identical dedicated `vector` structure. For details on the systematic errors see Section 6.2.

6.1.1 Energy Binning

The 2008 MINOS analysis [69] used bin widths of 1 GeV to cover the peak of the neutrino energy distribution, and increasingly wider bins to cover the tail region. However, the neutrino energy resolution (see Chapter 7) in the 1 - 2 GeV region where the oscillation dip is located has values typically distributed between 0.2 - 0.4 GeV. This indicates the 2008 analysis used a relatively coarse binning scheme, and a sensitivity improvement

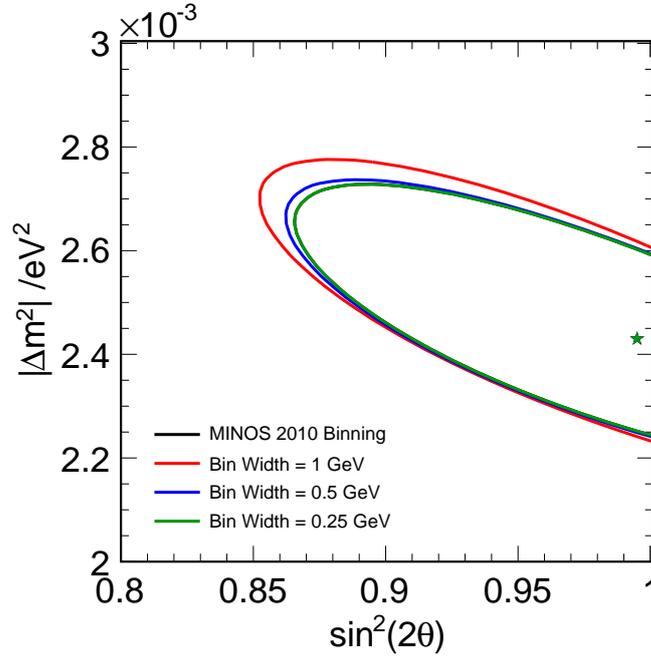


Figure 6.3: The 90% sensitivity confidence intervals for different energy spectra binning in the oscillation fit, for fiducial event samples (CC ν_μ and CC $\bar{\nu}_\mu$) summed over all Runs. The official MINOS energy binning used for the 2010 analysis is indistinguishable from the 0.25 GeV case (green). Contours are generated from Far Detector MC scaled to the data exposure of 7.2×10^{20} PoT, oscillated at $\sin^2(2\theta) = 1.0$ and $\Delta m^2 = 2.43 \times 10^{-3} \text{ eV}^2$.

could potentially be made by reducing the bin width of the neutrino energy spectrum.

Coarse binning schemes can conceal information regarding the position and depth of the oscillation dip, and make it more difficult to distinguish between predictions for different sets of oscillation parameters. However, using a bin width narrower than the typical resolution of the detector in the region of interest may modify the energy spectrum in a way unrelated to the underlying physics, and introduce unmodelled systematic errors. It also must be noted that the CPU required for the oscillation fit scales with the number of bins in the neutrino energy spectra.

Figure 6.3 shows sensitivity contours obtained using a range of binning configurations. The new configurations used bin widths of 1 GeV, 0.5 GeV and 0.25 GeV up to 30 GeV, with an overflow bin above this energy. The 1 GeV case is approximately equivalent to the 2008 MINOS analysis. Successive improvements are seen with the reduction in bin

width as expected, suggesting the reduction in bin width leads to an improvement in the distinction between predictions with different oscillation parameters. The final binning scheme used for the MINOS 2010 analysis for fiducial events (CC ν_μ and CC $\bar{\nu}_\mu$ samples) is summarised below, with 0.25 GeV bins covering the peak region and successively wider bins in the tail. This case is also plotted in Figure 6.3, and is indistinguishable from the 0.25 GeV case with one overflow bin.

- **ν_μ and $\bar{\nu}_\mu$ samples:**

- i. One bin from 0 GeV to 0.5 GeV.
- ii. From 0.5 GeV to 20 GeV, 79 bins of width 0.25 GeV.
- iii. From 20 GeV to 30 GeV, 10 bins of width 1 GeV.
- iv. From 30 GeV to 50 GeV, 10 bins of width 2 GeV.
- v. From 50 GeV to 200 GeV, one bin of width 150 GeV.
- vi. All events at a higher energy are placed in the overflow bin.

- **RAF sample:** The fit is performed in 11 bins of global bin number, where events are subdivided into six detector regions as detailed in [75], and binned within each region in reconstructed track energy¹. The binning scheme used in the standalone analysis in [75] also utilises track angle, however this increase in the number of bins led to problems with low statistics when systematically shifting events for the fitter [82] and a simpler energy only binning was developed.

The framework requires a significant number of templates to be stored, and the CPU required increases rapidly with the addition of a significant number of systematics. To limit the number of histograms which are required, when performing a grid search the $\sin^2(2\theta)$ term in the survival probability can be transformed into a scaling factor across all bins between maximally oscillated ($\sin^2(2\theta) = 1.0$) and unoscillated spectra:

$$\mathcal{H}\left(1.0 - \sin^2(2\theta) \sin^2 \frac{1.27L\Delta m^2}{E}\right) = \mathcal{H}\left(1.0 - \sin^2 \frac{1.27L\Delta m^2}{E}\right) \times \sin^2(2\theta) - \mathcal{H}(1) \times (\sin^2(2\theta) - 1) \quad (6.1)$$

¹For the standalone analysis detailed in [75], a binning scheme optimised for sensitivity in track energy and muon angle is used. The angular dependence was discarded in the combined analysis due to low statistics in the histograms required to incorporate sources of systematic error into the oscillation fit.

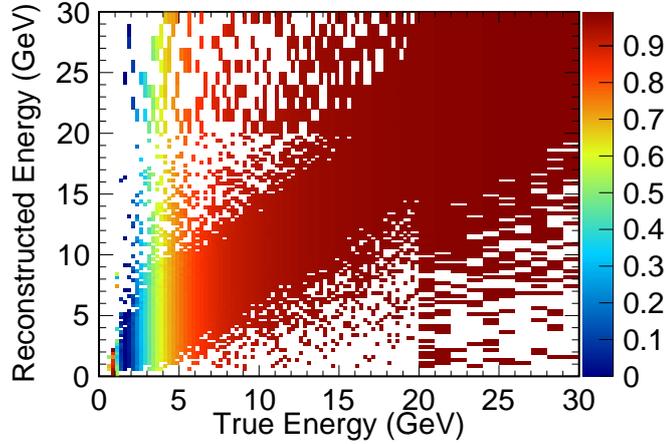


Figure 6.4: Example histogram showing the ratio of oscillated to unoscillated MC events in true vs reconstructed energy, oscillated at $\Delta m^2 = 2.78 \times 10^{-3}, \text{eV}^2$ and $\sin^2(2\theta) = 1.0$. For bins where no events are present in the DST ratio distribution, the bin centre is used to oscillate events in the prediction.

where $\mathcal{H}(1.0 - \sin^2(2\theta) \sin^2 1.27L\Delta m^2/E)$ represents the standard 2D template histogram in reconstructed and true energy, shown as a function of the oscillation probability which scales the contents of each bin. $\mathcal{H}(1.0 - \sin^2 1.27L\Delta m^2/E)$ represents a 2D histogram scaled by the oscillation probability at a specific Δm^2 value with $\sin^2(2\theta) = 1.0$, and $\mathcal{H}(1)$ the unoscillated template histogram. The same formula is applied to templates for the CC background events in each sample. An equivalent expression exists for the small effect of tau neutrino appearance:

$$\mathcal{H}\left(\sin^2(2\theta) \sin^2 \frac{1.27L\Delta m^2}{E}\right) = \mathcal{H}\left(\sin^2 \frac{1.27L\Delta m^2}{E}\right) \times \sin^2(2\theta) \quad (6.2)$$

and a set of templates is made for tau neutrinos accordingly. To further reduce the CPU time required to run the fitting software, all ROOT histograms were replaced by a standalone class which serves as a wrapper to an array, greatly reducing the complexity of the histogram object and the time required to access the entries. After implementing both of these modifications to the fitting software, the time required to perform a statistical fit was decreased by $\sim 1/2$.

By specifying a particular grid to be run over in oscillation space, templates can be pre-oscillated for each point in Δm^2 . This enables the same set of templates to be used for a variety of grids in $\sin^2(2\theta)$, and for different combinations of systematics. To

apply the oscillations, a set of histograms identical in binning to those output by the extrapolation are filled with the ratio of oscillated to unoscillated Monte Carlo events, where the oscillations are applied for the Δm^2 value in question. An example of the ratio histogram is shown in Figure 6.4. Each bin in the input template is weighted by the ratio of oscillated to unoscillated Monte Carlo. This is repeated for every grid point in Δm^2 prior to running the fit.

However, the pre-oscillated templates cannot be utilised when the CERN minimisation program MINUIT [83] is used instead of a grid search method. The MINUIT program requires the function to be minimised to be continuous and double-differentiable. When MINUIT steps in parameter space and calculates the log-likelihood, if using pre-oscillated templates on a rigidly defined grid it will see discontinuities in the likelihood function when it steps within a single bin width on the grid compared to stepping into an adjacent bin. It is unable to converge on a solution if these discontinuities are present, and so must apply the oscillations at each point in parameter space it visits. In order to minimise the CPU and time required to run the fitting software, the grid search method is thus preferred and used in all instances in this thesis.

A maximum likelihood analysis is performed to obtain the best fit oscillation parameters. Equation 6.3 expresses the standard form [17] of the binned negative log-likelihood for Poisson distributed data used in the fitting package, which performs a shape and normalisation fit to every vessel:

$$-2 \ln \mathcal{L} = 2 \sum_i^{N_{Dis}} \sum_j^{N_{Bins}} \left(N_{ij}^{MC} - N_{ij}^{Data} + N_{ij}^{Data} \ln \frac{N_{ij}^{Data}}{N_{ij}^{MC}} \right) \quad (6.3)$$

where N_{ij}^{MC} and N_{ij}^{Data} are the number of expected and observed events respectively, for neutrino energy bin j in vessel i . The likelihood is then summed over all energy bins (N_{Bins}), all vessels (N_{Dis}) and all samples to calculate the final likelihood value.

Confidence limits are defined by comparison of the likelihood at each point in $(\sin^2(2\theta), \Delta m^2)$ parameter space and the best fit value, with the x number of standard deviations interval defined as in Equation 6.4 [17]:

$$\frac{x^2}{2} = \Delta \ln \mathcal{L} = \ln \mathcal{L}_{\text{bestfit}} - \ln \mathcal{L} \quad (6.4)$$

For the two dimensional parameter fit to $(\sin^2(2\theta), \Delta m^2)$ resulting in two degrees of freedom, the 68% and 90% confidence intervals are described by those points for which

$\Delta \ln \mathcal{L} = 2.279/2$ and $\Delta \ln \mathcal{L} = 4.605/2$ respectively. When quoting 1σ confidence limits on the individual parameters, the errors are calculated by marginalising over the other parameter. For example, for each value of Δm^2 all $\sin^2(2\theta)$ values are scanned and the lowest likelihood value corresponds to the marginalised likelihood for that Δm^2 value. The minimum of the curve is at the best fit point by construction, and the 68% and 90% limits are defined by the 1D $\Delta 2 \ln \mathcal{L}$ values from the minimum of 1.0 and 2.71 respectively.

The best fit parameters and contours can be determined by either MINUIT, or alternatively a standard grid search method. For the fitting framework the latter method is implemented across $(\sin^2(2\theta), \Delta m^2)$ space in order to utilise the factorisation of $\sin^2(2\theta)$ to minimise the time and CPU required. MINUIT is utilised at each point in parameter space to minimise over the systematic errors and associated nuisance parameters, as described in the following section.

6.2 Systematic Uncertainties

As the statistical error in the MINOS oscillation result is decreased, potential systematic uncertainties become of greater significance. These can result in errors in the predicted Far Detector spectrum due to shifts in event energies, errors in the estimated level of background or the mistaken tagging of events as signal when they should not pass the selection cuts.

There are 20 systematic errors which are expected to contribute to the result presented in this thesis. The impact of the different uncertainties is estimated by determining the magnitude of the offset in the oscillation parameter if an unmodelled $\pm 1\sigma$ shift in a particular systematic is present in high-statistics Monte Carlo “fake data”. For systematics affecting both the Near and Far Detectors, both were shifted accordingly. This fake data is then input and fit to unmodified Monte Carlo, and the best fit point compared to the best fit when the fake data has no systematic shifts applied. The offset is taken to be representative of the bias due to this systematic on the final result.

Figure 6.5 shows the “star plot” for the nine most significant systematic errors identified in [84], extrapolated using the Beam Matrix Method and fit using the 2008 MINOS analysis fitter (called `NuSystFitter`) [85]. The centre represents the best fit to unshifted Monte Carlo, and the arms the deviations from the true best fit observed when $\pm 1\sigma$

shifts are applied for each systematic.

The nine systematics that were considered are summarised below (full details are given in [84]):

- **Absolute hadronic energy scale:** Comprising uncertainties common to both detectors in the shower energy measurement, with a scale given by $6.6\% + (3.5\%) \times e^{E_{shw}/1.44}$. Details of the calculation of this uncertainty are given in [74].
- **Absolute track energy scale:** Comprising uncertainties common to both detectors. For events with momentum measured from range, a systematic uncertainty of 2% is given, and for events where curvature is utilised an error of 3% is used. The errors in both cases are based on known detector simulation, particle propagation and energy loss unknowns, with the curvature measurement including a comparison between range and curvature measurements for specific tracks [86].
- **Relative normalisation:** An uncertainty in the number of neutrino events expected due to factors not present in both detectors. A relative normalisation uncertainty of 1.54% is used, arising from detector fiducial mass uncertainties, PoT counting and selection efficiencies between the Near and Far Detectors [87].
- **NC background:** Addressing uncertainty in the number of true NC events in the CC selected samples, incorporating NC cross sections. A systematic error of 20% is applied in both detectors [78, 88].
- **Cross sections:** An error incorporating uncertainties in a variety of neutrino and anti-neutrino interaction cross sections. The contributions are summed in quadrature for the overall error shown in Figure 6.5. For details on the individual cross sections and their contributions, see [84].
- **Relative hadronic energy scale:** Incorporates differences in detector calibrations and responses in hadronic energy measurement between detectors. A 1.1% systematic uncertainty on the energy scale in the Far Detector is applied, or a 1.9% uncertainty on the energy scale in the Near Detector [84].
- **Beam:** Significant disagreement exists in the high energy edge of the peak in the neutrino energy distribution between Near Detector Monte Carlo and data, due to poor modelling in the simulation of the pions and kaons which decay to produce

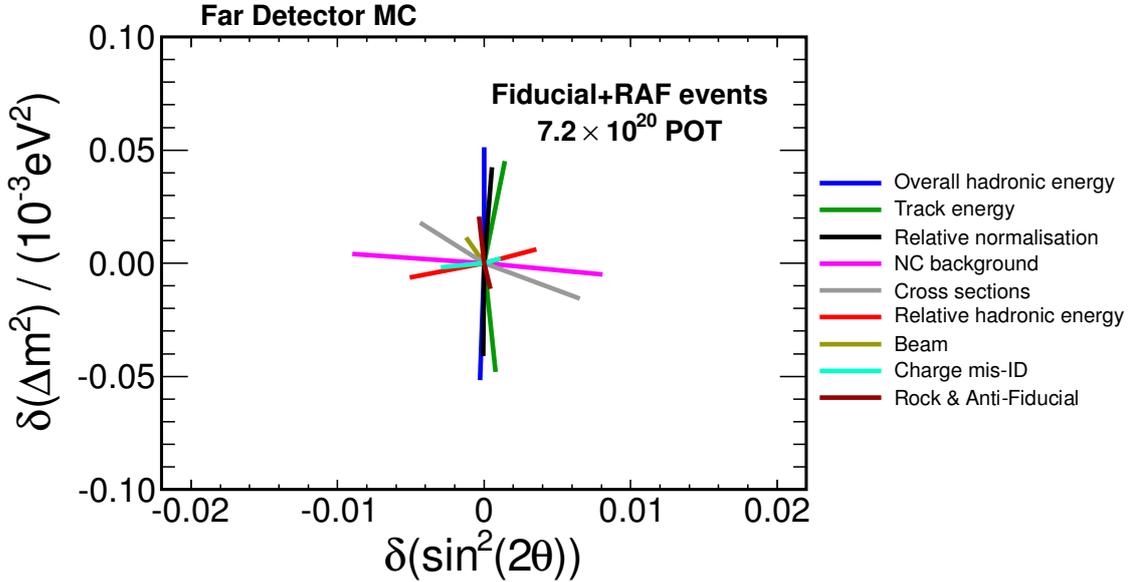


Figure 6.5: Shifts in the best fit point induced by imposing a $\pm 1\sigma$ shift in each main category of systematic error [85]. The centre point represents the input oscillation parameters to the scaled fake data, and each line represents the shift when an error is applied to the fake data and fitted to unaltered Monte Carlo predictions.

the neutrino beam. A beam fitting tuning process described in [89] constrains the simulation, with the 1σ error on the remaining beam uncertainties taken as the systematic uncertainty.

- **Charge mis-identification:** A systematic error of 40% on the relative size of the sample of true $\bar{\nu}_\mu$ events in the ν_μ sample, applied across all energies [84].
- **Rock and Anti-Fiducial:** A combination of errors specific to the RAF sample, summed in quadrature for Figure 6.5. These errors include the uncertainty in separating true rock and true detector originating events along the detector edge, a 0.9% normalisation uncertainty on events originating in the rock and a 1.1% error on the level of non-DIS cross-sections of rock events [75, 84].

Table 6.2 summarises the sizes of the shifts in the best fit parameters, corresponding to ± 1 standard deviations in the four most dominant systematics. The four systematics inducing the largest deviations are included as nuisance parameters in the fit, as shown in Equation 6.5.

Systematic		$\delta (\Delta m^2)$	$\delta (\sin^2 (2\theta))$
Absolute Shower Energy	+1 σ	-0.052	-0.0003
	-1 σ	+0.051	+0.0000
Track Energy	+1 σ	-0.048	-0.0008
	-1 σ	+0.045	+0.0014
Normalisation	+1 σ	+0.041	-0.0001
	-1 σ	-0.042	+0.0001
NC Background	+1 σ	-0.005	+0.0081
	-1 σ	+0.004	-0.0090

Table 6.2: Shifts in the best fit point due to the four most significant systematic errors, fitting to all Runs (I, II, III and pHE) and samples (ν_μ , $\bar{\nu}_\mu$ and RAF).

With the inclusion of pseudo-High Energy (pHE) beam configuration data the dataset considered in this thesis comprises multiple beam operating conditions. It was investigated whether the normalisation systematic error should be applied across all data regardless of the beam conditions, or whether a second independent normalisation systematic should be applied to Runs taken in the pHE configuration. Figure 6.6 shows the difference in sensitivity by considering two normalisation systematics compared to a common single systematic shift. With the minimal change in sensitivity achieved, a normalisation error common to all Runs was utilised.

6.2.1 Incorporating Systematics into the Framework

To include the systematics in the oscillation fit, the likelihood function is allowed freedom to adjust the parameters at each point in $(\sin^2 (2\theta), \Delta m^2)$ using MINUIT until the minimum $\ln \mathcal{L}$ is found. The likelihood formula therefore becomes:

$$-2 \ln \mathcal{L} = 2 \sum_i^{N_{Dis}} \sum_j^{N_{Bins}} \left(N_{ij}^{MC} - N_{ij}^{Data} + N_{ij}^{Data} \ln \frac{N_{ij}^{Data}}{N_{ij}^{MC}} \right) + \sum_k^{N_{Syst}} \frac{a_k^2}{2\sigma_k^2} \quad (6.5)$$

where parameter a_k is defined as the magnitude of the systematic shift, and acts as a penalty term which penalises the fit as the parameter is permitted to vary. The error σ_k is the estimated uncertainty in each systematic given above.

In the absence of any systematic shift in the data, the fit preferentially sets the

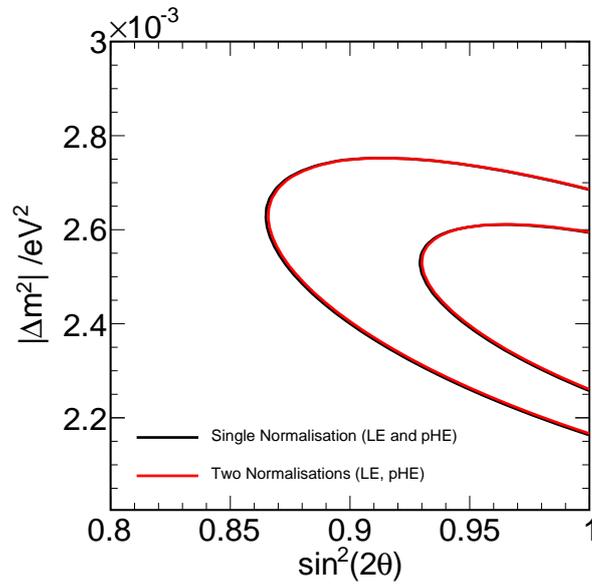


Figure 6.6: The 68% and 90% sensitivity confidence intervals for individual normalisation systematic errors (red), for data taken in the Low Energy beam configuration (Runs I and II), and pseudo-High Energy beam configuration (Run I-pHE). The sensitivity obtained using a common normalisation error for all beam configurations (LE and pHE) is shown in black. Contours are generated from Far Detector MC in the ν_μ sample, for Runs I, II and pHE scaled to the data exposure of 4.7×10^{20} PoT and oscillated at $\sin^2(2\theta) = 1.0$ and $\Delta m^2 = 2.32 \times 10^{-3} \text{ eV}^2$.

nuisance parameters to zero. With systematic shifts present in the data (as detailed in the following section), the $\ln \mathcal{L}$ value improves as the fit to the data improves, and the sensitivity contours increase in size as the height of the $\Delta 2 \ln \mathcal{L}$ surface decreases. As expected, this reflects the additional uncertainty in the best fit with the inclusion of these errors.

6.2.1.1 Implementation

In order to make the fitting process with multiple systematics computationally feasible, an alternative approach to brute force minimisation over four systematic parameters is required. Template histograms cannot sensibly be regenerated for every value of systematic shift at each point in $(\sin^2(2\theta), \Delta m^2)$ parameter space; including the shower and track energy systematics, (as both are applied in both detectors) would require adjusting the Monte Carlo energies for all events and re-extrapolating histograms through the Beam Matrix for every value of each systematic shift required by the fitter.

In order to reduce the number of templates required, the approach taken is to interpolate between templates at certain set values of the systematic errors to reconstruct a continuous set of systematic shifts. At each point in $(\sin^2(2\theta), \Delta m^2)$ parameter space, a instance of MINUIT is created which, for each systematic being fit, provides an initial shift to a function which interpolates between the systematically shifted templates. The log likelihood for this grid point is minimised by MINUIT, which provides a range of nuisance parameter values to the interpolation function.

A variety of methods for interpolating between systematically shifted templates at the integer values of $\pm 1, 2\sigma$ were investigated. The times required to evaluate the log-likelihood per grid point in $(\sin^2(2\theta), \Delta m^2)$ parameter space are summarised in Table 6.3, for oscillation fits including three and five systematics, compared to the time for a statistical only oscillation fit. A simple linear interpolation between neighbouring templates would cause MINUIT to exhibit unwanted behaviour at the boundaries between templates, as the function would not be differentiable across the entire $-2\sigma < \eta < +2\sigma$ range. Thus a cubic interpolation between the shifted templates and the nominal Monte Carlo, applied on a bin-by-bin basis and utilising MINUIT was considered, in addition to a bin-by-bin cubic spline function.

However, as Table 6.3 shows, the application of the above interpolation methods took a large amount of CPU time and so a weighted interpolation scheme independent

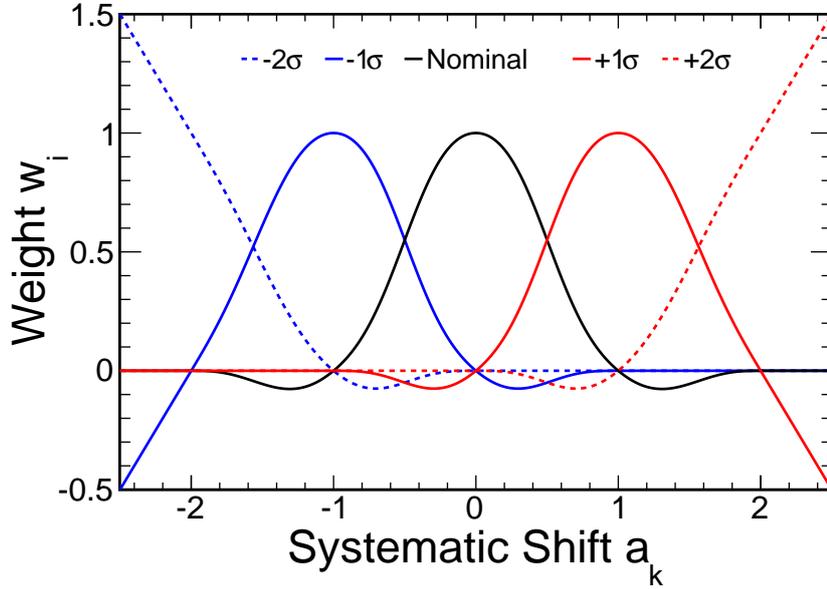


Figure 6.7: Weights applied to the five input histograms, systematically shifted by $0, \pm 1\sigma$ and $\pm 2\sigma$ to interpolate between them to generate a template with any given shift applied. The curves are maximal and all other weights fall to zero at the integer values of σ to return the known shifted spectra.

of neutrino energy was also developed². To determine the expected spectrum for a given shift η in the range $-2\sigma < \eta < +2\sigma$, the function generates a weight which is applied to the entire template, a more efficient approach than a bin-by-bin method. The weights are generated from an average of linear interpolations between neighbouring templates, with values shown in Figure 6.7. At an integer value of σ the function returns an input template with weight one and all other templates with weight zero; at intermediate values the function assigns more significant weights to neighbouring contributions.

Equation 6.6 then provides the expected spectrum by summing the residuals from the nominal for all systematics to be fit to obtain the final prediction:

$$n(\eta) = \omega_{-2}(\eta)n_{-2} + \omega_{-1}(\eta)n_{-1} + \omega_0(\eta)n_0 + \omega_{+1}(\eta)n_{+1} + \omega_{+2}(\eta)n_{+2} \quad (6.6)$$

The use of the weighted interpolation scheme decreased the CPU time required from ~ 2 days to ~ 8 hours for a grid search with 100×100 points, and hence is the approach adopted for use by the `GhostFitter` software.

²Developed by Dr A. Blake in the Cambridge group.

	Average Time/ Grid Point (<i>s</i>)		
	Cubic Interpolation	Cubic Spline	Weighted Template
Statistical	1.8	1.8	1.8
3 Systematics	5.4	5.1	2.2
5 Systematics	16.0	12.7	3.0

Table 6.3: Average time required per grid point to run the fitting software, by interpolation method between systematics templates. The three systematics considered were the normalisation, NC background and absolute shower energy; absolute track energy and relative shower energy were added for the five systematics case. The time taken to perform a statistical fit is also included.

The four dominant systematics considered here apply across multiple samples; for example, the normalisation uncertainty is applied to all CC events (ν_μ , $\bar{\nu}_\mu$ and ν_τ) which occur within the detector. The same value of the nuisance parameter must be applied simultaneously to the CC ν_μ and $\bar{\nu}_\mu$ samples, in addition to the detector component of the RAF sample. Thus the relative normalisation is applied as a simple flat scaling of the contents of each reconstructed energy bin, to all CC events, after the remaining systematic shifts have been included. In the case of the RAF sample, the scaling factor is applied to the detector component only, and the rock and detector spectra are then summed for the final prediction.

To include the remaining sources of systematic error, the `vector` structure described in Figure 6.2 holds the four shifted templates required for each systematic as indicated. If a systematic is included which does not apply to that sample (such as charge misidentification which is not considered in the RAF sample), a null pointer is stored in the entry corresponding to that systematic and is skipped accordingly when the nuisance parameters are minimised. In this fashion any number of systematics which apply to any combination of subsets of samples can be handled by the fitting software.

6.2.2 Correlation of Systematics

Summing the residuals to obtain the overall prediction assumes the systematics are all uncorrelated, as does the use of nuisance parameters in the likelihood function, by not using the full covariance matrix method to incorporate the systematics. In order to justify this supposition, Figure 6.8 shows the ratios of Monte Carlo predicted neutrino energy spectra, generated by changing the systematics simultaneously and by changing

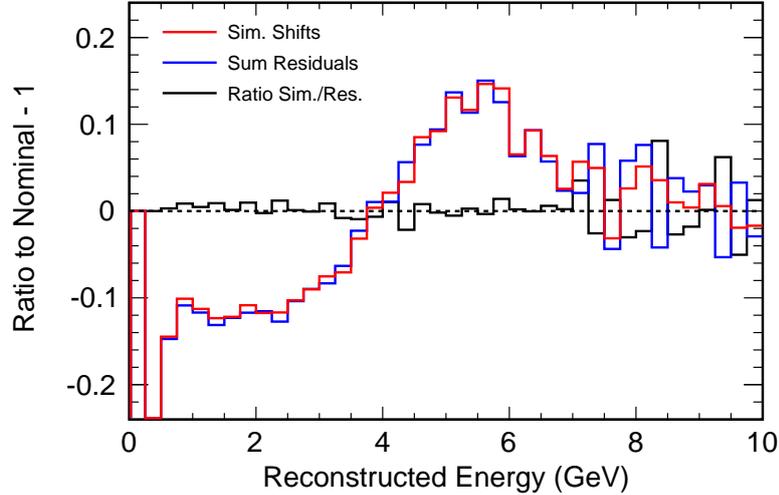


Figure 6.8: Ratio to the nominal spectrum with a $+1\sigma$ shift in each of the four dominant systematics included in the oscillation fit, both applied simultaneously and applied separately and summing the residual from the nominal case. In the energy region of interest, the disagreement is of the order of less than 2% between the two cases.

them independently and summing the residuals from the nominal case. The degree of correlation is less than 2% in the energy region where oscillations are predominant.

When multiple systematics were incorporated into the fitting software assuming they are essentially uncorrelated, it was found degeneracies occurred when combining more than two sources of systematic errors. Fitting more than two systematic errors resulted in discontinuities appearing in the nuisance parameter surfaces in oscillation space as shown in Figure 6.9(a), visible as peaks and troughs between $0.75 < \sin^2(2\theta) < 0.85$ and $2.85 \times 10^{-3} < \Delta m^2 < 3.0 \times 10^{-3}$ for the track energy systematic shift. Figure 6.9(b) shows the log-likelihood distribution at $\sin^2(2\theta) = 0.825$ and $\Delta m^2 = 2.85 \times 10^{-3} \text{ eV}^2$ as a function of the track energy systematic error, where the fitting software is run on Monte Carlo data including the normalisation, neutral current and track energy systematics as nuisance parameters. The track energy systematic error is fixed at specific values consecutively, with the normalisation and neutral current systematic errors permitted to roam freely. The resulting global and local minima are clearly visible.

The MINUIT fit to any one of the nuisance parameters can become trapped in the local minimum (shown by the best fit marker in Figure 6.9(b) at $a_k = -0.5\sigma$), and is

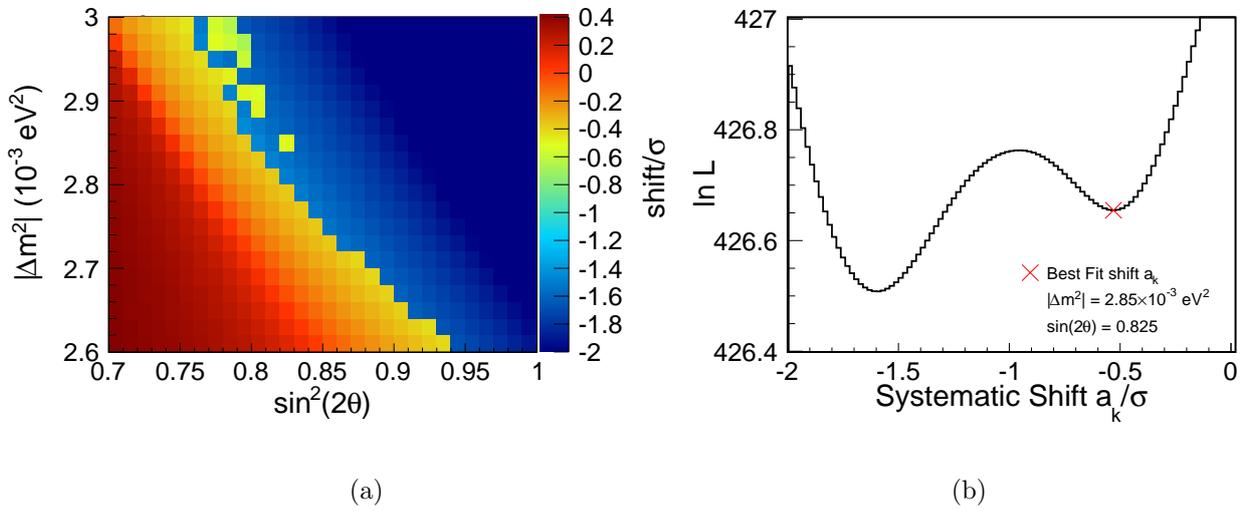


Figure 6.9: (a) The best fit nuisance parameter surface for the track energy systematic, when performing an oscillation fit to the normalisation, neutral current and track energy systematics. (b) The distribution of log-likelihood against track energy systematic error, at $\sin^2(2\theta) = 0.825$ and $\Delta m^2 = 2.85 \times 10^{-3} \text{ eV}^2$. The normalisation and neutral current systematics are allowed to roam freely at each fixed value of the track systematic error. The figures are generated from Far Detector MC scaled to the data exposure of $7.2 \times 10^{20} \text{ PoT}$, oscillated at $\sin^2(2\theta) = 1.0$ and $\Delta m^2 = 2.43 \times 10^{-3} \text{ eV}^2$.

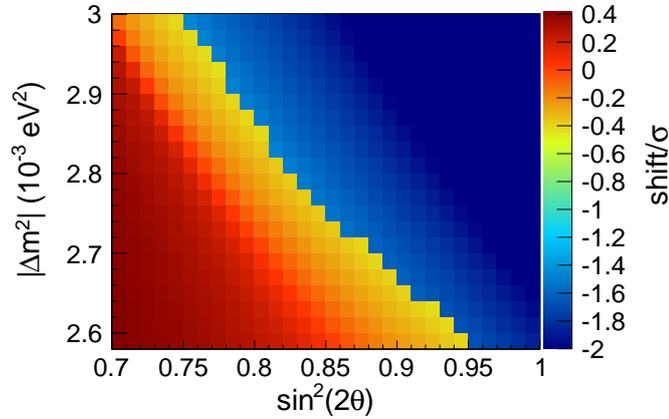


Figure 6.10: The best fit nuisance parameter surface for the track energy systematic, when performing an oscillation fit to the normalisation, neutral current and track energy systematics. Initial values for the oscillation fit to all systematics simultaneously are found from best fit nuisance parameters for each systematic individually.

unable to converge to the global minimum with the lowest likelihood value. To prevent this from occurring, at each point in parameter space the fitting software minimises the likelihood for each systematic individually, fixing all other nuisance parameters. The results of each of the individual fits are fed into the instance of MINUIT which minimises all systematic errors simultaneously, as the initial values for each systematic error, instead of giving all systematic errors an initial value of zero. Figure 6.10 shows the likelihood distribution as a function of the track energy systematic and the best fit nuisance parameter surface in oscillation space for the oscillation fit discussed above, with the `MinuitInitialValues()` method implemented. The discontinuities visible in Figure 6.9(a) are no longer observed. This study was repeated for all combinations of systematics to ensure no discontinuities in the best fit nuisance parameter surfaces remained.

6.3 Fit Validation

Sensitivity contours and predictions of the parameter errors are constructed by using fake data, weighting the full MC sample by a set of oscillation parameters and scaling down to an equivalent exposure to that of the data collected. A basic test of the fitting framework is to use this spectrum in place of data and ensure the correct oscillation parameters

are obtained, and compare the best fit spectra and sensitivity contours to those output from the fitting software used for 2008 published MINOS analysis (`NuSystFitter`). This test is then performed in a more sophisticated fashion using the Mock Data Challenge detailed in Section 6.3.2. Sensitivity contours and best fits are detailed in Section 6.4.

6.3.1 Simple MC Tests

Preliminary testing on the fitting framework involved inputting pre-oscillated spectra at specified disappearance parameters into the software in place of data, with a range of applied systematic error shifts in combinations of parameters. By removing the nuisance parameter term from the likelihood function, an identical spectrum to that input as data should be output as the best fit, with the correct best fit systematic shifts. After passing these basic tests, the method used to oscillate the predicted spectra was investigated by comparison with an oscillated spectrum at the same parameters from the `NuSystFitter` software. All tests were carried out on the ν_μ sample, Far Detector Run III MC; oscillated at $\sin^2(2\theta) = 1.0$ and $\Delta m^2 = 2.43 \times 10^{-3} \text{ eV}^2$.

It was found the oscillated spectra from the different software differed by $\sim 3\%$ in the region of the peak of the reconstructed neutrino energy distribution, between 0.5 - 2 GeV. Figure 6.11 shows the ratios of oscillated true CC ν_μ events in Run III to unoscillated as output by the fitting software in black, compared to the true neutrino disappearance probability function in red; for the `GhostFitter` software and the `NuSystFitter` software. In the the case of the latter, the default method employed is to apply oscillations at the neutrino energy at the bin centre, clearly visible where the function intersects the lowest two bins in energy. This is not the case for the `GhostFitter` method (Section 6.1), where the function intersects the distribution at a higher energy than the bin centre above the peak and lower energy below it.

The impact of using different methods of applying oscillations to the sensitivity in parameter space is shown in Figure 6.12. The Far Detector MC fake data was oscillated in true neutrino energy on an event by event basis, and the software was permitted to move into the unphysical region $\sin^2(2\theta) > 1.0$. The `GhostFitter` software provides the least biased determination of the true best fit point, compared to the four alternative methods in `NuSystFitter`, which defaults to the `BinCentres` method. Both methods are shifted relative to the oscillated data by less than the largest systematic, with the star plot superimposed at the best fit point on Figure 6.12 for comparison.

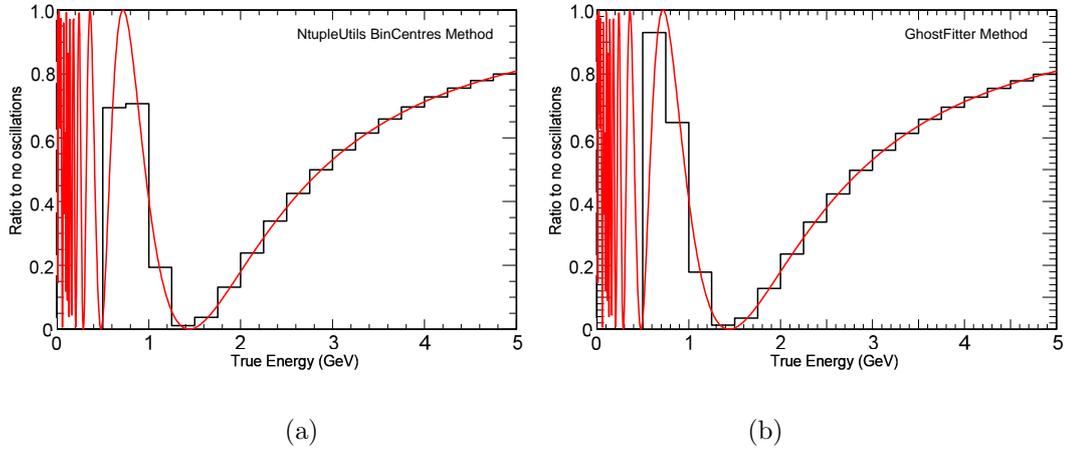


Figure 6.11: Ratios of oscillated true CC ν_μ events in Run III Far Detector Monte Carlo to the unoscillated true energy spectrum (black), compared to the true neutrino disappearance probability function (red); as output by the official MINOS 2008 software `NuSystFitter` (a), and the `GhostFitter` software (b).

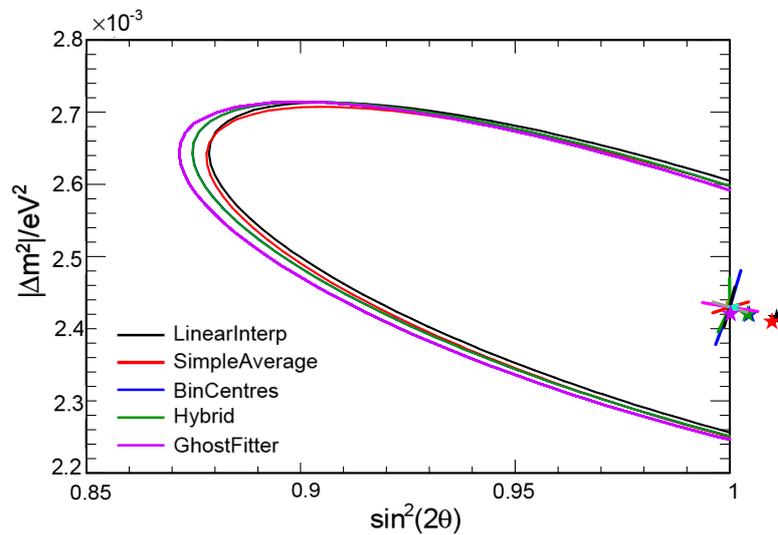


Figure 6.12: The 90% sensitivity confidence interval for different methods of incorporating neutrino oscillations, for fake data oscillated on an event-by-event basis. The `GhostFitter` software best fit and contour (purple) are compared to alternative methods employed by the `NuSystFitter` software (black-green). The systematics star plot (Figure 6.5) is overlaid.

6.3.2 Mock Data Challenges

The standard test within the MINOS Collaboration used to validate extrapolation and fitting methods is to perform a Mock Data Challenge (MDC). A large dataset of fake data is produced, with oscillation parameters applied that are unknown to the analyser prior to performing the fit. The extrapolation and fitting code is deemed to be validated if the result lies within 1σ of the applied parameters.

The fake dataset is produced with 100 times the statistics of the total real data exposure, and is split into 100 fake “experiments”. These experiments are fit individually and simultaneously, summing the 100 likelihood surfaces to provide an overall prediction and sensitivity contour. A distribution of best fit points is expected as statistical fluctuations are applied to each energy bin within each experiment, which test the determination of the 68% and 90% confidence limits by observing the distribution of the best fit points in relation to these regions of parameter space. The resulting best fit points from the MDC are shown in Figure 6.13.

The true oscillation parameters were revealed to be $\sin^2(2\theta) = 0.9756$ and $\Delta m^2 = 2.1704 \times 10^{-3} \text{ eV}^2$; the best fit parameters from the `GhostFitter` for the different fits are summarised in Table 6.4. The MDC was also performed with the `NuSystFitter` as an additional cross check, and the best fit parameters from this package were found to be $\sin^2(2\theta) = 0.966$, $\Delta m^2 = 2.17 \times 10^{-3} \text{ eV}^2$.

Samples Fit	$\sin^2(2\theta)$	Δm^2
Truth	0.9756	2.1704
Fiducial, no Res. Binning	0.9625	2.1675
Fiducial, with Res. Binning	0.9725	2.1525

Table 6.4: Mock Data Challenge best fit points for the fiducial sample, with no resolution binning and with the data split into five resolution bins (resolution binning is discussed in Chapter 7). The true input oscillation parameters are also shown. These points are compared to the 68% and 90% confidence limits in Figure 6.14, with and without resolution binning.

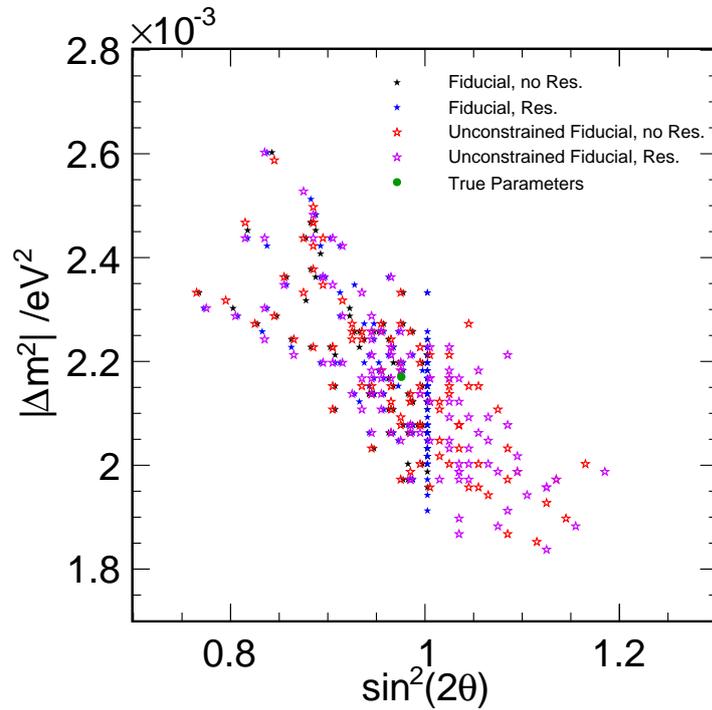


Figure 6.13: The distribution of 100 fake experiments performed as part of the MDC. Fits were performed for fiducial events with and without resolution binning.

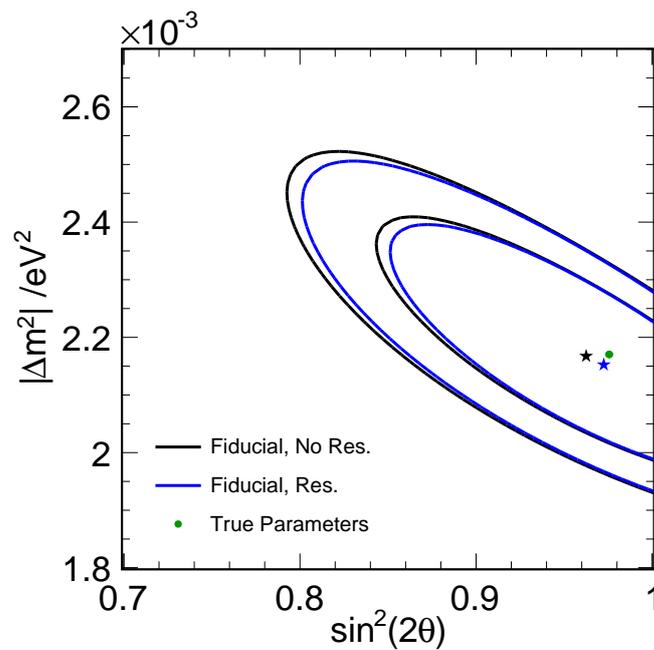


Figure 6.14: The 90% and 68% sensitivity confidence intervals for the MDC performed prior to fitting the real data, showing the fits performed for fiducial events with and without resolution binning.

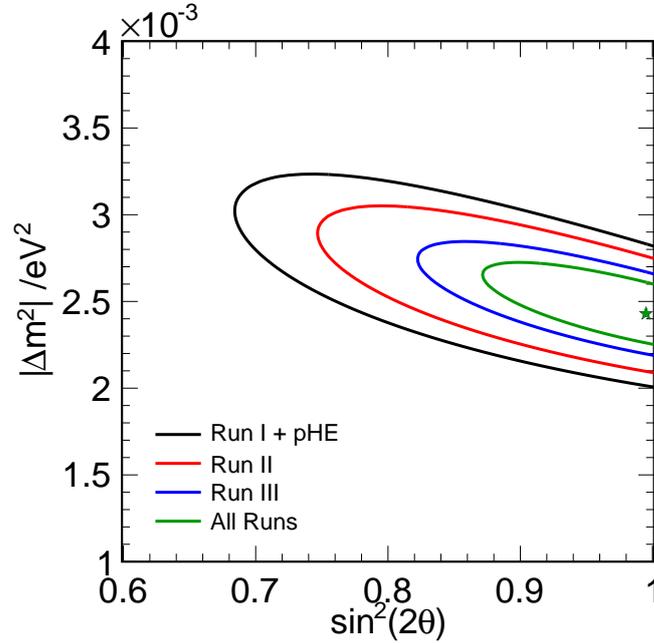


Figure 6.15: The 90% sensitivity confidence intervals for individual Runs and all Runs combined, for fiducial event samples (CC ν_μ and CC $\bar{\nu}_\mu$). Contours are generated from Far Detector MC scaled to the data exposure of 7.2×10^{20} PoT, oscillated at $\sin^2(2\theta) = 1.0$ and $\Delta m^2 = 2.43 \times 10^{-3}$ eV².

6.4 Sensitivities

Figures 6.15 to 6.18 show a variety of sensitivity contours, with and without systematics for the various improvements in the fitting method (designed for this thesis) from the NuSystFitter method. Figures 6.15 and 6.16 exhibit the improvements in sensitivity obtained by adding Runs (the discrete data taking periods defined in Table 3.1) for the fiducial and RAF samples, with the overall result shown in green for both plots. Figure 6.17 shows the change in the contour by adding each of these samples independently for all Runs; with the RAF sample having a more significant impact than the addition of the CC $\bar{\nu}_\mu$ events.

The effect of fitting for the four major systematics is shown in Figure 6.18, where the expected broadening of the contour from the flattening of the $\ln \mathcal{L}$ surface is clearly visible. Figure 6.19 exhibits the marginalisation of these surfaces over each of the oscillation parameters.

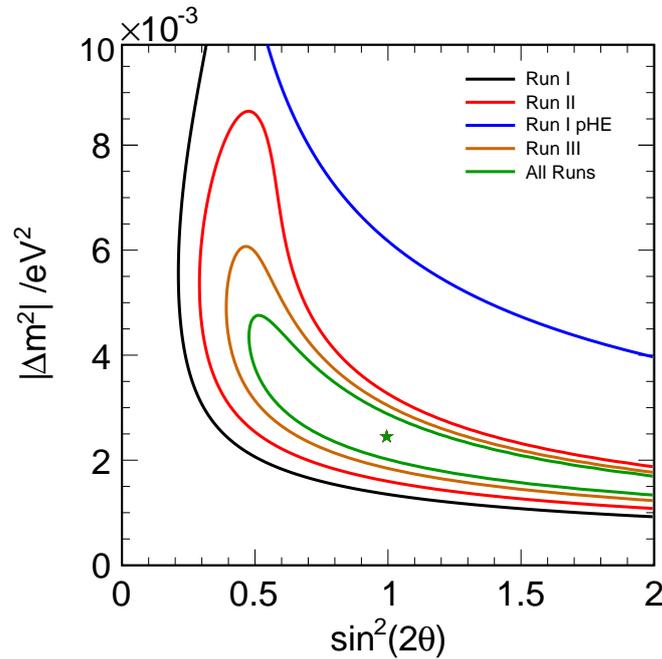


Figure 6.16: The 90% sensitivity confidence intervals for individual Runs and all Runs combined, for the RAF event sample. Contours are generated from Far Detector MC scaled to the data exposure of 7.2×10^{20} PoT, oscillated at $\sin^2(2\theta) = 1.0$ and $\Delta m^2 = 2.43 \times 10^{-3} \text{ eV}^2$.

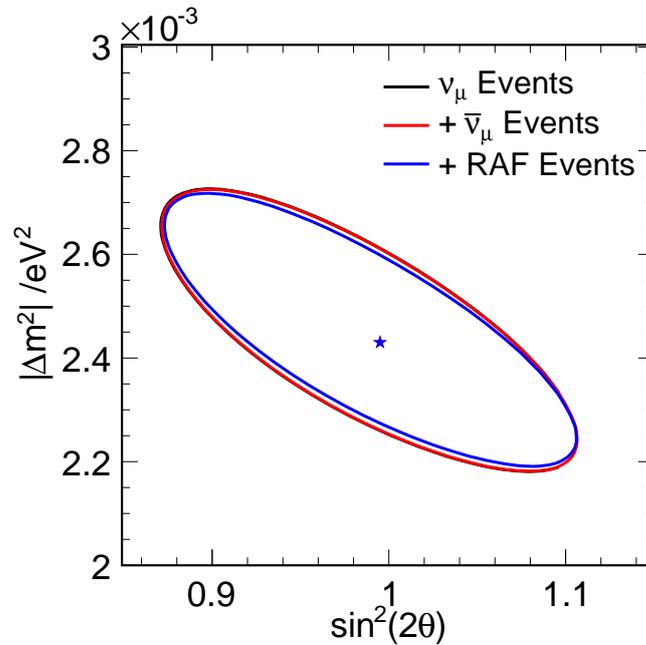


Figure 6.17: The 90% sensitivity confidence intervals for the ν_μ sample, adding $\bar{\nu}_\mu$ events and RAF events. Contours are generated from Far Detector MC scaled to the data exposure of 7.2×10^{20} PoT, oscillated at $\sin^2(2\theta) = 1.0$ and $\Delta m^2 = 2.43 \times 10^{-3} \text{ eV}^2$.

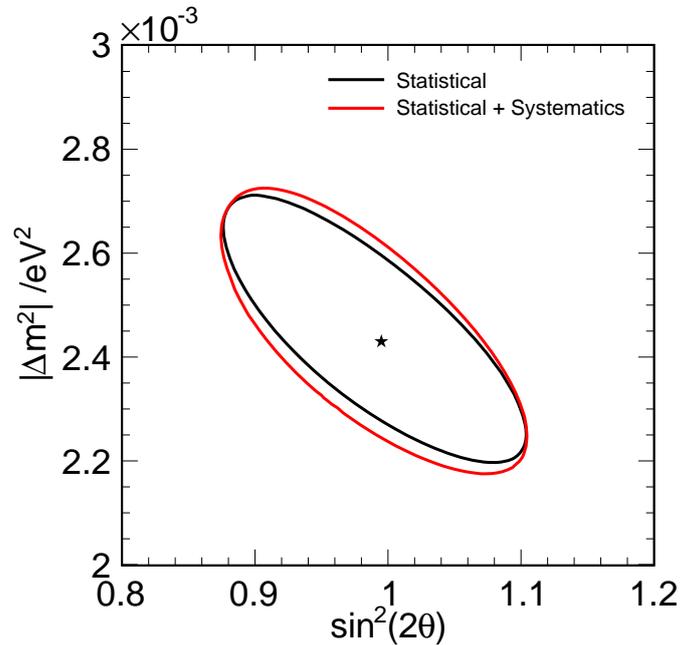


Figure 6.18: The 90% sensitivity confidence intervals for the statistical fit and fit including the four major systematic errors, to MC scaled to the data exposure of 7.2×10^{20} PoT, oscillated at $\sin^2(2\theta) = 1.0$ and $\Delta m^2 = 2.43 \times 10^{-3} \text{ eV}^2$. The fits include all three event samples and all Runs. The expected broadening of the contour when systematic errors are included is observed.

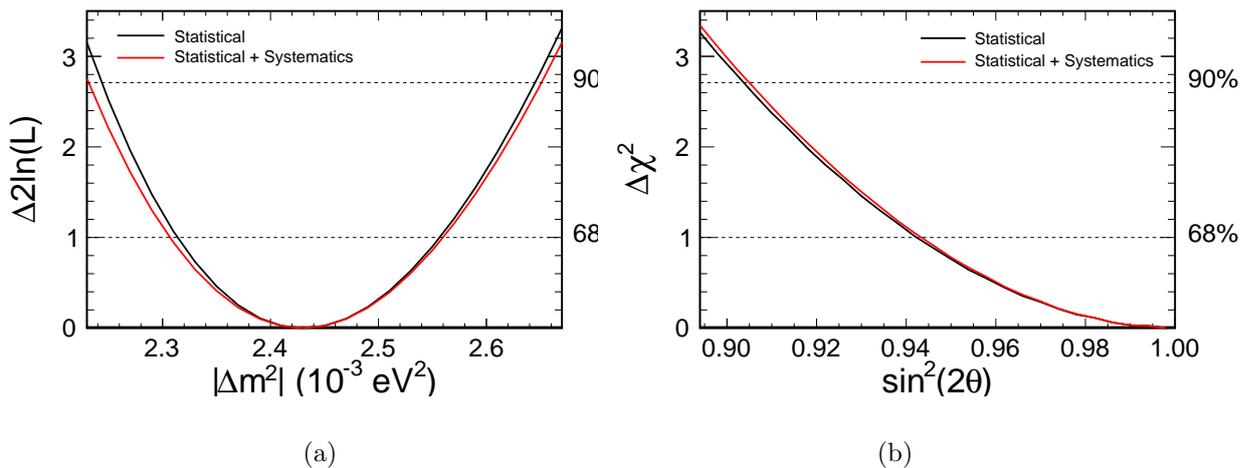


Figure 6.19: Marginalisation of the log-likelihood surface onto the Δm^2 (a) and $\sin^2(2\theta)$ (b) axes, for the statistical fit (black) and fit including the four major systematic errors (red). The fits include all three event samples and all Runs.

6.5 Alternative Models

The fitting framework also implements the pure neutrino decay and decoherence models, in addition to the oscillation analysis. For decoherence, the mixing angle component of the survival probability factorises as for the oscillation case (see Equation 6.1), and so two pre-decohered templates are generated in a similar fashion for the muon neutrino survival probability; one with no disappearance (template $\mathcal{H}(1)$) and one disappeared for each value of μ^2 with $\sin^2(2\theta) = 1.0$ (template $\mathcal{H}\left(1 - \frac{1}{2}\left(1 - e^{-\frac{\mu^2 L}{2E}}\right)\right)$). They are then scaled between at each point on the grid as follows:

$$\mathcal{H}\left(1 - \frac{1}{2}\sin^2(2\theta)\left(1 - e^{-\frac{\mu^2 L}{2E}}\right)\right) = \mathcal{H}\left(1 - \frac{1}{2}\left(1 - e^{-\frac{\mu^2 L}{2E}}\right)\right) \times \sin^2(2\theta) - \mathcal{H}(1) \times (\sin^2(2\theta) - 1) \quad (6.7)$$

In the case of neutrino decay, the form of the survival probability does not factorise as readily, and three pre-generated templates are required: one with no disappearance (template $\mathcal{H}(1)$) and two with the appropriate exponential energy dependent functions applied (templates $\mathcal{H}\left(e^{-\frac{\alpha L}{2E}}\right)$ and $\mathcal{H}\left(e^{-\frac{\alpha L}{E}}\right)$.) They are then scaled between as in Equation 6.8:

$$\begin{aligned} \mathcal{H}\left(\left(\sin^2(\theta) + \cos^2(\theta)e^{-\frac{\alpha L}{2E}}\right)^2\right) &= \mathcal{H}\left(e^{-\frac{\alpha L}{2E}}\right) \times 2\sin^2(\theta)(1 - \sin^2(\theta)) \\ &\quad + \mathcal{H}\left(e^{-\frac{\alpha L}{E}}\right) \times (1 - \sin^2(\theta))^2 \\ &\quad + \mathcal{H}(1) \times \sin^4(\theta) \end{aligned} \quad (6.8)$$

Figure 6.20 shows the 90% confidence intervals for statistical fits to these models with PoT scaled Far Detector Monte Carlo.

6.6 Summary

A new standalone framework for oscillation fitting has been developed and validated, for this thesis and for the MINOS result published in [39] to augment the significant gain in statistics from previously published results. The following chapter discusses an additional improvement to the result sensitivity, by utilising energy resolution information of the events to subdivide the sample vessels.

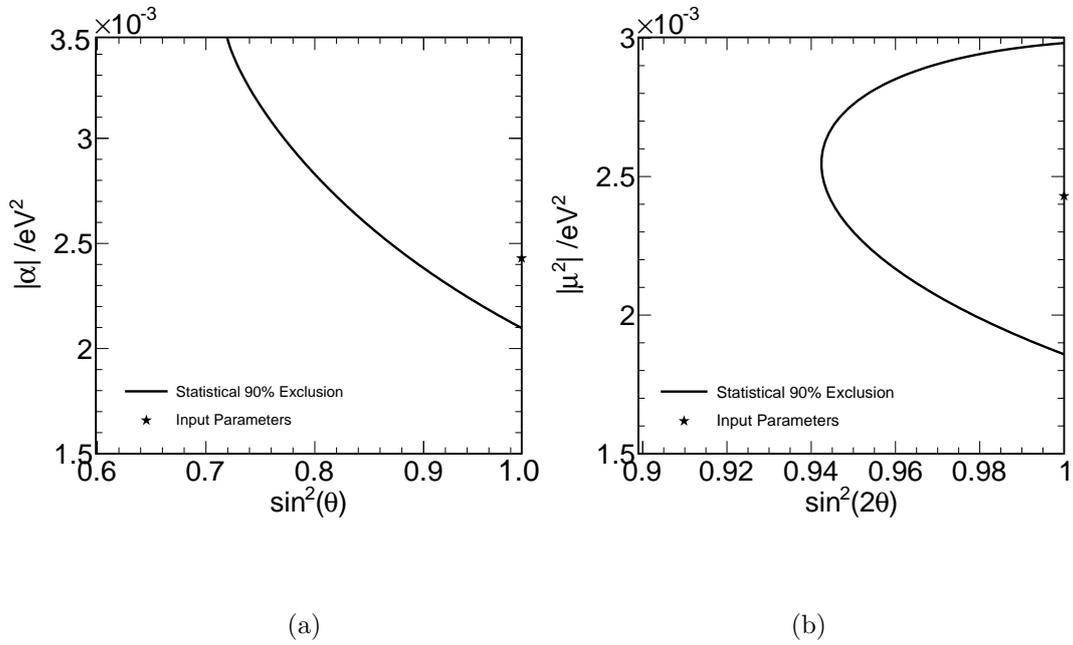


Figure 6.20: The 90% sensitivity confidence intervals for pure neutrino decay (a) and pure neutrino decoherence (b). Contours are generated from Far Detector MC scaled to the data exposure of 7.2×10^{20} PoT, disappeared at $\sin^2(2\theta) = 1.0$ and $\Delta m^2 = 2.43 \times 10^{-3} eV^2$.

Chapter 7

Event Resolution

The energy resolution of the calorimeter is the limiting factor to the precision with which the energy of the neutrino can be determined. By utilising the hadronic shower energy and track energy resolutions in the Far Detector to quantify the degree to which an event has been accurately reconstructed, and incorporating this information into the oscillation fit, an improvement in sensitivity to the oscillation parameters may be achieved.

7.1 Resolution Calculation

The reconstructed neutrino energy is the sum of the shower and track energies, the latter of which can be determined from either the range of the muon or the curvature of the track. Tracks which are fully contained have their momentum estimated from range as described in [90], where an algorithm swims from the final state vector at the end of the track back to the track vertex, calculating the energy loss at each step to determine the muon momentum. Exiting muons have their momentum taken from the curvature of the track as determined by the Kalman track fitter [32].

Separate parameterisations of the resolution are obtained for each of these momentum measurements from high statistics Monte Carlo Far Detector data and summed in quadrature:

$$\sigma_{E_\nu}^2 = \sigma_{E_{\text{trk}}}^2 \oplus \sigma_{E_{\text{shw}}}^2 \quad (7.1)$$

to define the overall resolution of an event. These resolution parameterisations are then

used to divide the extrapolated Far Detector predictions in the ν_μ sample and Far Detector data into bins of resolution. In the context of the code structure described in Figure 6.2, each bin of resolution in each Run period defines a single vessel, which have their log-likelihoods calculated independently and summed.

The resolution of a calorimeter can be expressed as a function of energy with terms incorporating differing energy dependences, summed in quadrature. The resolution parameterisation function takes the form $\sigma_x = A\sqrt{x} \oplus Bx \oplus C$; where x is a label indicating the shower energy or muon momentum. The term where $\sigma_x/x \propto 1/\sqrt{x}$ corresponds to the Poisson uncertainty in counting the number of photoelectrons output by the PMTs; the term $\sigma_x/x \propto 1/x$ incorporates contributions from electronic noise in the readout chain, and the term $\sigma_x/x \propto 1$ corresponds to energy independent contributions such as calibration errors [91].

For showers and contained tracks, the parameter x is taken to be the reconstructed energy (showers) or momentum (contained tracks). In the case of uncontained tracks, the parameterisation is defined in terms of $p^2\sigma_{q/p}$:

$$\begin{aligned}
 p &= \frac{q}{q/p} \\
 \therefore dp &= \frac{d}{d(q/p)} \left(\frac{q}{q/p} \right) d(q/p) \\
 dp &= -\frac{p^2}{q} \sigma_{q/p} \\
 \Rightarrow \sigma_p &\propto p^2 \sigma_{q/p}
 \end{aligned} \tag{7.2}$$

where $\sigma_{q/p}$ is an estimate of the error in q/p (reconstructed muon charge/momentum) output by the track fitting algorithm. When multiplied by p^2 (as derived above) it provides an estimate of the error in p .

To determine whether to use $p^2\sigma_{q/p}$ instead of the reconstructed muon momentum (as is done for the momentum from range events), Figure 7.1 shows the best fit lines obtained between the Gaussian widths of the $\Delta E = (E_{\text{reco}} - E_{\text{true}})$ distributions in bins of σ_{E_ν} . σ_{E_ν} is calculated for each event using the two possible measures of the momentum from curvature error: $p^2\sigma_{q/p}$ and the reconstructed muon momentum. A gradient of one would signify the parameterisation perfectly reconstructs the resolution of the muon momentum, and is clearly better modelled by the $p^2\sigma_{q/p}$ parameterisation with a gradient of $m = 0.95$, than using muon momentum with a gradient of $m = 0.88$.

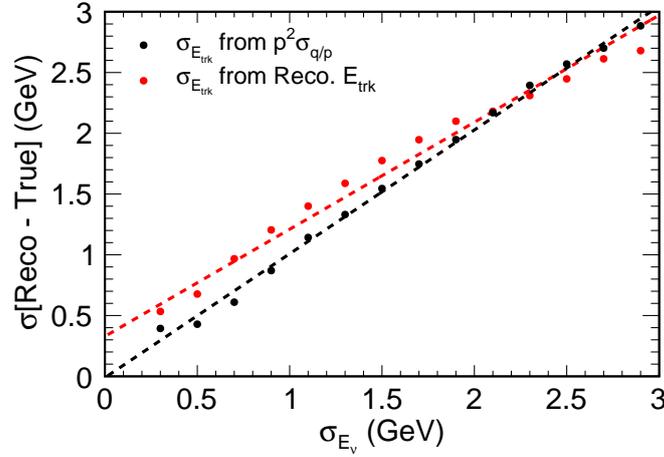


Figure 7.1: Distributions of Gaussian widths of $E_{\text{reco}} - E_{\text{true}}$ distributions in bins of σ_{E_ν} , with resolution calculated using $p^2\sigma_{q/p}$ (black) and the reconstructed muon energy (red). The dashed lines are linear fits between the two parameters, with gradients of $m = 0.88$ for the muon momentum case and $m = 0.95$ for the $p^2\sigma_{q/p}$ case.

The widths of Gaussian functions, fit to $\Delta E = (E_{\text{reco}} - E_{\text{true}})$ distributions in bins of reconstructed energy or $p^2\sigma_{q/p}$ were extracted, and plotted as functions of these reconstructed parameters. While the standard approach would be to fit to $\Delta E/E$, the values obtained diverge rapidly at low energy using this variable so ΔE is considered. The functional forms were then fit using MINUIT, as shown in Figure 7.2 for the shower and muon energy from range cases, and Figure 7.3 for the muon energy from curvature case. The resulting parameterisations of the resolutions (in the standard form) are as follows:

Shower energy

$$\frac{\sigma_{E_{\text{shw}}}}{E_{\text{shw}}} = \frac{40.4\%}{\sqrt{E}} \oplus 8.6\% \oplus \frac{257 \text{ MeV}}{E} \quad (7.3)$$

Muon momentum from range

$$\frac{\sigma_{E_\mu}}{E_\mu} = \frac{5.1\%}{\sqrt{E}} \oplus 6.9\% \quad (7.4)$$

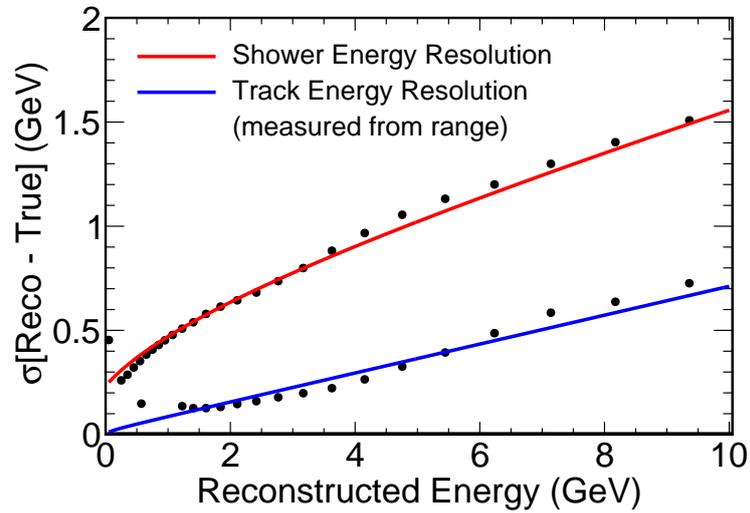


Figure 7.2: Energy resolution parameterisations for shower energy and muon momentum determined from range, as a function of reconstructed shower and muon momentum respectively.

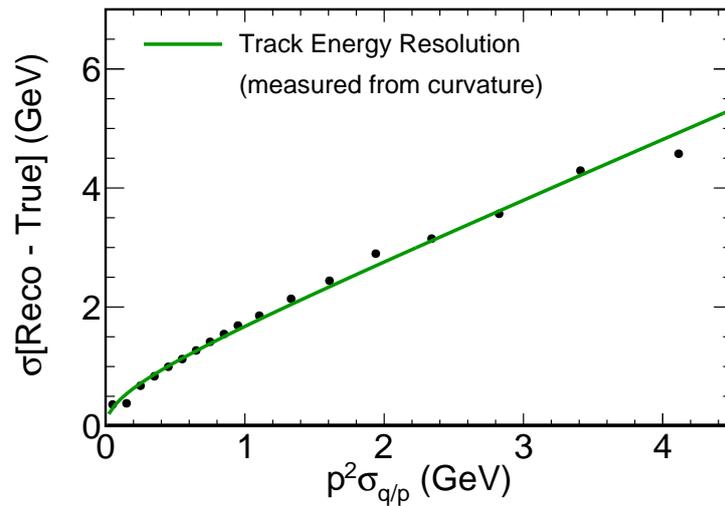


Figure 7.3: Energy resolution parameterisation for muon momentum determined from curvature, as a function of $p^2\sigma_{q/p}$.

Muon momentum from curvature

$$\sigma_{E_\mu} = p^2 \sigma_{q/p} \oplus 1.34 \sqrt{p^2 \sigma_{q/p}} \quad (7.5)$$

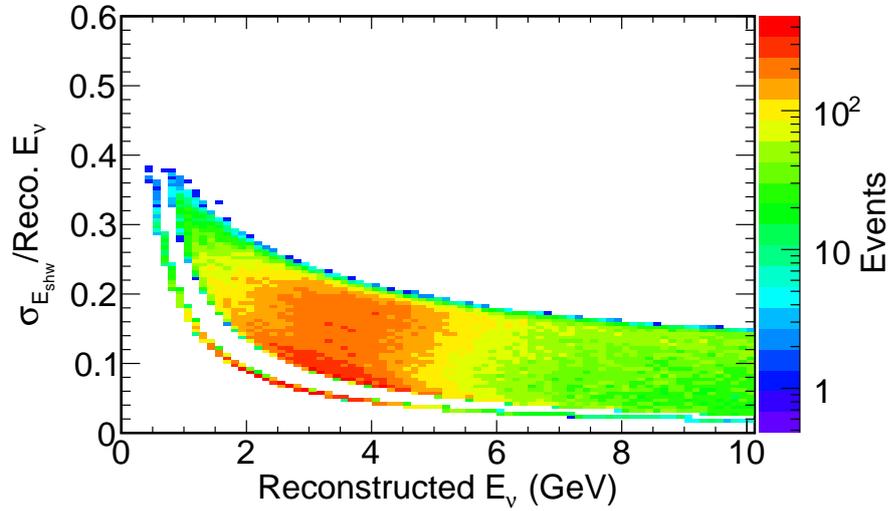
For events with zero shower energy, the energy resolution has a minimum value of 257 MeV, an artifact of the kNN-based shower energy reconstruction summarised in Chapter 5 and discussed further in Chapter 8. Figure 7.4(a) shows the fractional shower energy resolution as a function of reconstructed neutrino energy, where the isolated strip of events are those reconstructed with zero shower; demonstrating the non-zero $\sigma_{\text{shw}}^{\text{minimum}} = 257 \text{ MeV}$.

The fractional track energy resolution in Figure 7.4(b) shows two distinct regions, with stopping tracks with low σ_{trk} measured from range and a wider spread of curvature resolutions. The form of the parameterisation of the muon momentum from curvature has, as expected from Equation 7.2, the constant multiplying the $p^2 \sigma_{q/p}$ term rounding to unity. The term proportional to $\sqrt{p^2 \sigma_{q/p}}$ therefore acts as a correction factor, compensating for any inaccuracies in estimating the error in q/p from the Kalman filter.

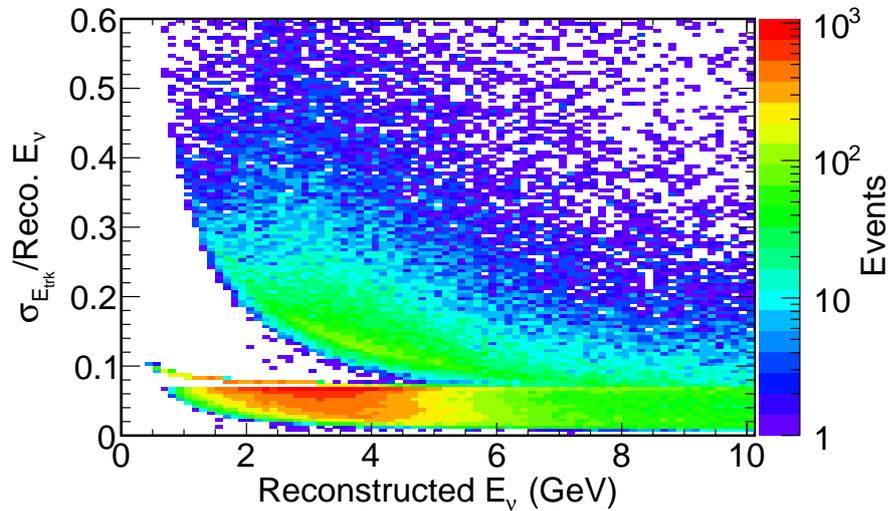
Each event in the Far Detector (in both data and Monte Carlo) is assigned a σ_{E_ν} , and all events are divided into vessels based on the fractional error on the reconstructed neutrino energy E_ν , σ_{E_ν}/E_ν . Each bin in reconstructed neutrino energy is expected to contain events with a spread of resolutions, and hence the division is applied by considering the distribution of σ_{E_ν}/E_ν in each bin *independently*, and dividing the events into N number of quantiles. As these cuts are considered on a bin-by-bin basis, the positions of the quantile boundaries are energy dependent, and determined from a high statistics Monte Carlo sample for each Run. Figure 7.5 shows this division for Run I MC and five quantiles, and a slice of the fractional resolution at 1 GeV. If the data are divided into five samples, the 20% “best” measured events can be considered to be those in the lowest σ_{E_ν}/E_ν quantile, and the 20% “worst” in the uppermost quantile.

7.1.1 Comparison to Calibration Detector Energy Resolution

The calibration detector (CalDet, as described in Section 3.2.3) was used in a CERN test beam to determine the response of the steel-scintillator and electronics to beams of known momentum. The data from the calibration detector were compared to Monte Carlo, which parameterised the energy resolution for hadronic showers as $56\%/\sqrt{E} \oplus 2\%$ [57] from the original calorimetric shower energy estimator.

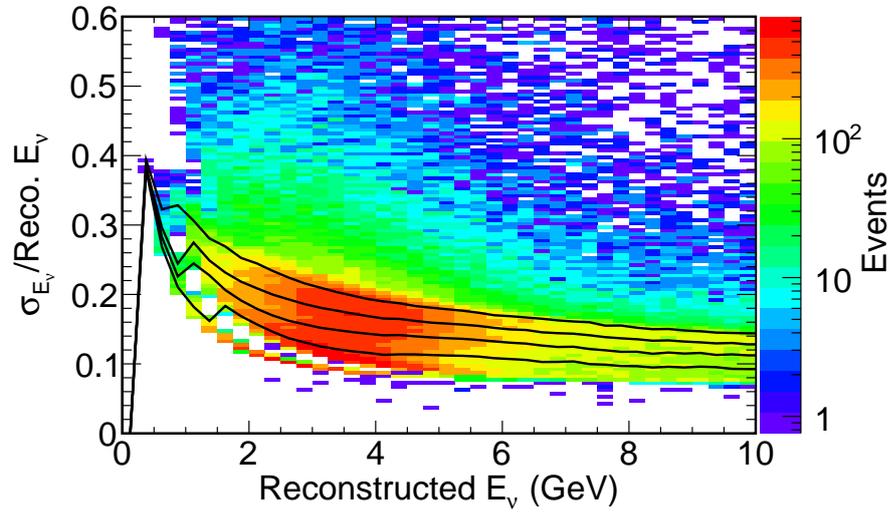


(a)

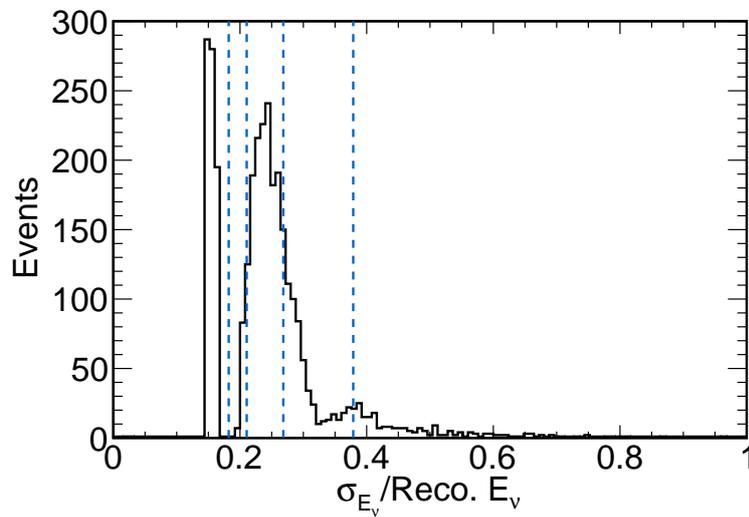


(b)

Figure 7.4: The fractional resolution shower energy contribution $\sigma_{\text{shw}}/E_\nu$ (a) and track energy from range and curvature contribution $\sigma_{\text{trk}}/E_\nu$ (b), for high statistics Run I MC, for all events which pass the MINOS 2010 CC selection. Quasi-elastic events with zero shower energy appear as a band of events at the low edge of the main distribution in (a), and a thin strip of higher fractional resolution events above the cluster of low $\sigma_{E_{\text{trk}}}$ stopping events in (b).



(a)



(b)

Figure 7.5: (a) The fractional resolution σ_{E_ν}/E_ν as a function of reconstructed energy for Run I high statistics Monte Carlo. Lines show the division of the Monte Carlo dataset into five quantiles. A slice of the fractional resolution error for the 1.0 - 1.25 GeV bin with quantile edges is shown in (b).

Table 7.1 compares resolution parameterisations using the most recent Monte Carlo and reconstruction, for different forms of the parameterised function with and without the lowest shower energy bin included in calculating the constants, to the CalDet result. The lowest shower energy bin is higher in value than would be expected due to an excess of events with $E_{shw} < 0.25$ GeV, an artifact of the k NN shower energy summarised in Chapter 5; where events reconstructed with their energy below 0.25 GeV have a reconstructed shower energy equal to zero assigned to them. By removing the zero shower events, the constants defining three-term functional form of the parameterisation move closer in value to the original CalDet values, indicating the lowest shower energy bin does act to pull up the B and C terms and pull down the A term. By reducing the function to the two-term form used in [57], the values move closer again to those obtained by CalDet, with the remaining differences attributed to changes in the Monte Carlo and shower energy reconstruction between 2004 and the present analysis.

	Parameterised Constants		
	A	B	C
CalDet	56%	2%	–
$\sigma_E/E = A/\sqrt{E} \oplus B \oplus C/E$ including $E_{shw} = 0$ GeV events	40.4%	8.6%	257 MeV
$\sigma_E/E = A/\sqrt{E} \oplus B \oplus C/E$ $E_{shw} = 0$ GeV events removed	44.5%	5.7%	211 MeV
$\sigma_E/E = A/\sqrt{E} \oplus B$ including $E_{shw} = 0$ GeV events	48.7%	1.0%	–
$\sigma_E/E = A/\sqrt{E} \oplus B$ $E_{shw} = 0$ GeV events removed	48.7%	1.0%	–

Table 7.1: Shower energy resolution parameterisations for the functional forms $\sigma_E/E = A/\sqrt{E} \oplus B \oplus C/E$ and $\sigma_E/E = A/\sqrt{E} \oplus B$, with and without the inclusion of $E_{shw} = 0$ GeV events. These events cause the anomalous first point in the shower resolution parameterisation in Figure 7.2.

To ensure the different forms of the shower energy parameterisation do not affect the final sensitivity, each of the four functions listed in Table 7.1 are used to determine the resolution value σ_ν for Far Detector Monte Carlo and the fitting software is run

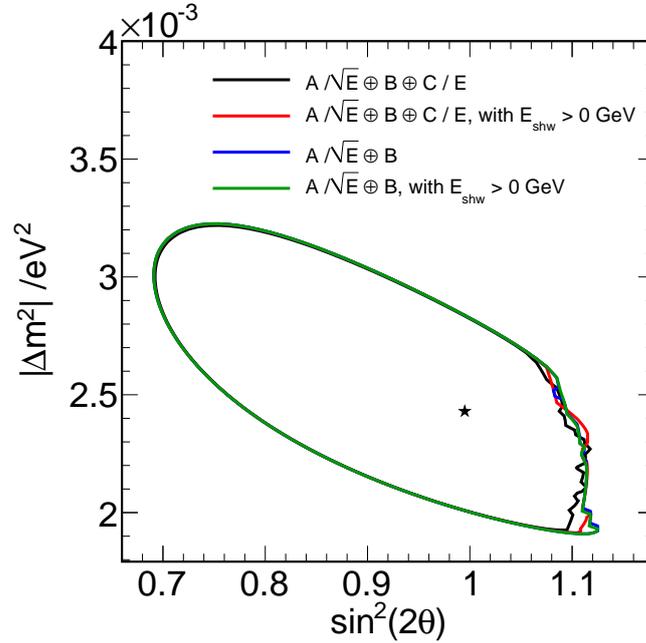


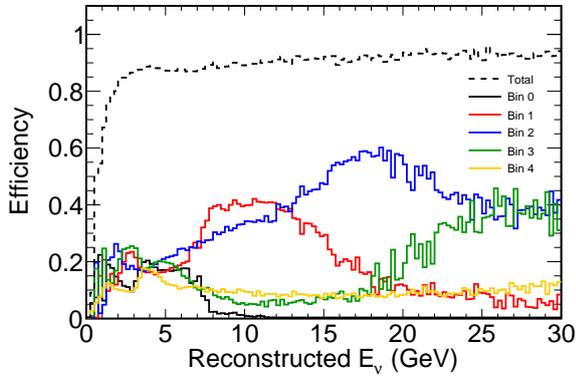
Figure 7.6: The 90% sensitivity confidence intervals for an oscillation fit performed for different parameterisations of shower energy resolution; using the functional form $\sigma_E/E = A/\sqrt{E} \oplus B \oplus C/E$ with (black) and without (red) including zero shower energy events, and the functional form $\sigma_E/E = A/\sqrt{E} \oplus B$ with (blue) and without (green) including zero shower energy events in the parameterisation. Zero shower energy events are included in the oscillation fit for all cases. The contours are generated from Run I Far Detector MC scaled to the data exposure of 1.27×10^{20} PoT, oscillated at $\sin^2(2\theta) = 1.0$ and $\Delta m^2 = 2.43 \times 10^{-3} \text{ eV}^2$.

to obtain the 90% sensitivity contour in each case. The resulting confidence limits are shown in Figure 7.6, where in the physical region the differences are less than 0.01% in Δm^2 . Hence the parameterisation of the shower energy shown in Figure 7.2 is utilised.

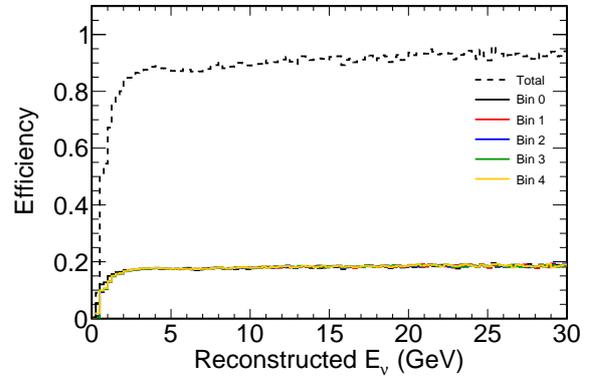
7.1.2 σ/E Distributions

The σ/E distributions used to define the quantiles are divided using the ROOT `GetQuantiles()` method, which splits the histogram by identifying x values which bound $1/N$ of the integral. If this bound lies between two bin edges, the method defaults to the lower bin edge. When determining the quantile divisions using high

Res. Bin	ν_μ selected events			
	100 bins	200 bins	1000 bins	100k bins
Bin 0	42.76	54.16	66.74	66.52
Bin 1	79.18	74.68	66.10	66.40
Bin 2	99.02	75.16	66.63	66.40
Bin 3	67.26	74.16	67.49	66.40
Bin 4	44.17	54.24	65.44	66.67



(a)



(b)

Table 7.2 & Figure 7.7: Table 7.2: Event counts for Run I MC, scaled to the Run I data exposure of 1.3×10^{20} PoT in each bin of resolution for four binning schemes for the σ/E distributions. Figure 7.7: Efficiency for identifying ν_μ CC events split into five bins of resolution, for 100 bins in energy Figure 7.7(a) and 100k bins Figure 7.7(b). The highly uneven event counts in Table 7.2 for the coarsest binning clearly correlate with the uneven efficiencies in Figure 7.7(a).

statistics Monte Carlo, this can cause large numbers of events to move between bins based on how coarsely binned the initial distributions are; leading to a bias which can produce quantiles containing uneven numbers of events.

Table 7.2 details the event counts in each of the quantiles, for distributions with 100 bins, 200 bins, 1000 bins and 100k bins in energy. Event counts are more even with a finer binning, with minimal differences between the integrals for each quantile in the 100k bins case, and clearly uneven event counts for the distributions with fewer than 1000 bins.

The effect of uneven quantile divisions due to coarse binning are also visible in Figure 7.7, which shows the efficiency, ϵ , of selected CC ν_μ events in each of the five bins of resolution, for distributions with 100 bins and 100k bins. The efficiency ϵ is defined as:

$$\epsilon = \frac{\# \text{ True CC Selected Events}}{\# \text{ Total Selected Events}} \quad (7.6)$$

where the selection applied consists of the selection cuts summarised in Table 5.1. As would be expected, with a finer binning the efficiencies in each resolution bin are equalised at $\simeq 0.17$ relative to an overall selection efficiency with no resolution binning of $\simeq 0.85$.

7.1.3 Oscillation Fitting in Bins of Resolution

Resolution information is included in the oscillation fit by treating each of the resolution bins in each Run as an independent vessel (as defined in Chapter 6). As shown in Figure 7.7, the mechanism can be thought of as fitting five Runs, each with an efficiency of $\simeq 0.17$, in place of one Run with an efficiency of $\simeq 0.85$.

In the high statistics limit, by construction the unoscillated neutrino spectra in reconstructed energy are identical in each quantile. However, oscillations are applied as a function of true energy, so the oscillated neutrino energy spectra are no longer identical. The best resolution quantile should exhibit the most diagonal distribution in reconstructed vs true energy, and the worst resolution quantile the most spread reconstructed vs true distribution, as shown in Figure 7.8.

The fit treats the sample in question (CC ν_μ) as comprising $N_{\text{Runs}} \times N_{\text{Res. Bins}}$ separate experiments, which are each fed into the likelihood function in Equation 6.5 as independent distributions. Alternatively, it is possible to include resolution information by carrying out separate shape-only likelihood fits for each resolution bin, and an overall normalisation fit for the Run:

$$-2 \ln \mathcal{L} = 2 \left(\sum_i^{N_{Dis}} N_i^{MC} - N_i^{Data} \ln N_i^{MC} - \sum_i^{N_{Dis}} \sum_j^{N_{ResBins}} \sum_m^{N_{Bins}} \mu_{ijm}^{MC} - \mu_{ijm}^{Data} \ln \mu_{ijm}^{MC} \right) \quad (7.7)$$

where N_{ijm}^{MC} , N_{ijm}^{Data} , μ_{ijm}^{Data} and μ_{ijm}^{MC} are the number of expected events, data events and the relative data and Monte Carlo probabilities respectively, for reconstructed energy bin m and resolution bin j in distribution i . However, the use of this form of the likelihood

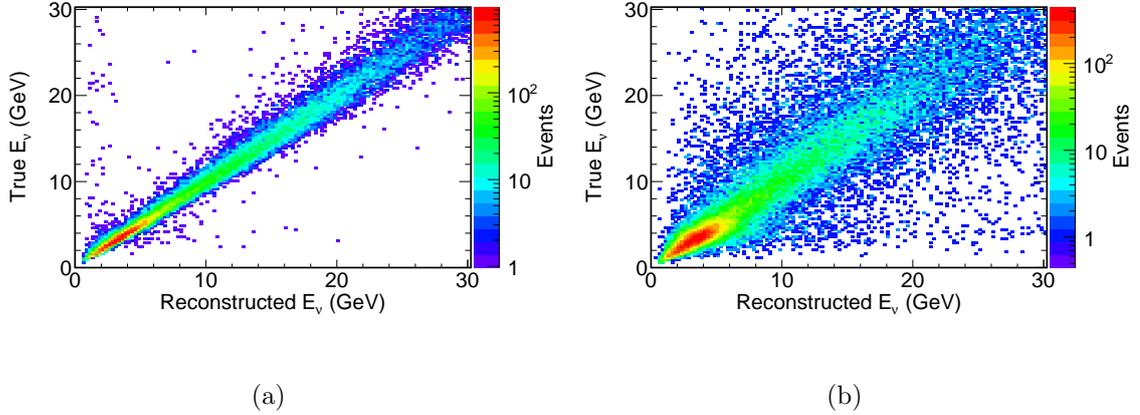


Figure 7.8: Reconstructed vs true neutrino energy for Run I MC for the events with the best 20% resolution (a) and the worst 20% resolution (b).

function exhibited $< 0.5\%$ improvement in sensitivity to Δm^2 , with indistinguishable sensitivity contours to those generated using Equation 6.5.

Only the CC ν_μ sample is fit in bins of resolution, with the positive curvature $\bar{\nu}_\mu$ sample undivided; leading to the fit of fiducial events being designated as “ $N_{\nu_\mu} \oplus 1$ ”, where N_{ν_μ} is the number of resolution bins. Studies were performed to determine if additional sensitivity improvements could be obtained by parameterising and fitting with the $\bar{\nu}_\mu$ sample event resolution, the “ $N_{\nu_\mu} \oplus N_{\bar{\nu}_\mu}$ ” case. While a 1% improvement in sensitivity was observed for the case of $N_{\nu_\mu} = N_{\bar{\nu}_\mu} = 5$, the small number of events in some resolution bins and run periods led to highly uneven unoscillated Far Detector predictions[92]. Coupled with the relatively small size of the improvement this approach was not adopted for use in this thesis. The binning scheme in energy and region utilised for the RAF sample is roughly analogous to splitting by resolution.

7.2 Performance Studies

7.2.1 Sensitivity Studies

Figure 7.9 shows the expected confidence limits on the oscillation parameters using the GhostFitter software, obtained by dividing the CC ν_μ event sample into one, two, five and ten quantiles; confirming the expectation that a greater number of bins leads to a

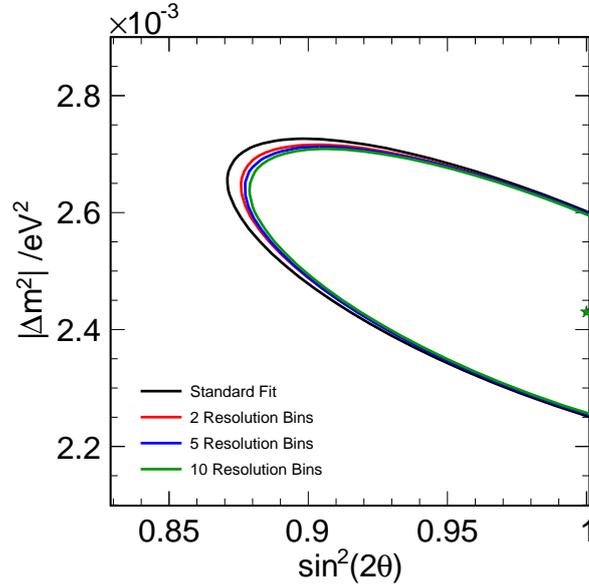


Figure 7.9: The 90% sensitivity confidence intervals for the standard fit, and fits performed in two, five and ten quantiles of resolution. The contours are generated from Far Detector MC scaled to the data exposure of 7.2×10^{20} PoT, oscillated at $\sin^2(2\theta) = 1.0$ and $\Delta m^2 = 2.43 \times 10^{-3} \text{ eV}^2$.

greater improvement in sensitivity. The projections onto $\sin^2(2\theta)$ and Δm^2 over the $\Delta 2 \ln \mathcal{L}$ surface are given in Figure 7.10, representing a 6% improvement in sensitivity in $\sin^2(2\theta)$ and an 8% improvement in Δm^2 for five bins of resolution. The improvement values are taken from the positions of the sensitivity curves at $\Delta 2 \ln \mathcal{L} = 1.0$. The gain in sensitivity from one to five bins is significantly greater than from five to ten, noting that the CPU time and memory required for the fit scale linearly with the number of quantiles utilised. Hence the use of five divisions was chosen.

Ratios of Far Detector Monte Carlo data, oscillated at $\sin^2(2\theta) = 1.0$ and $\Delta m^2 = 2.43 \times 10^{-3} \text{ eV}^2$ to the unoscillated prediction, for the best and worst quantiles in the case of five resolution bins are shown in Figure 7.11. The oscillation fit was performed with all bins of resolution simultaneously, for all Runs scaled to the data exposure of 7.2×10^{20} PoT. As expected, the events in the best resolution quantile exhibit a sharper and deeper oscillation minimum than the worst events, closer to the shape and position of the true oscillation probability.

The events in each quantile of resolution were input and used to perform an oscil-

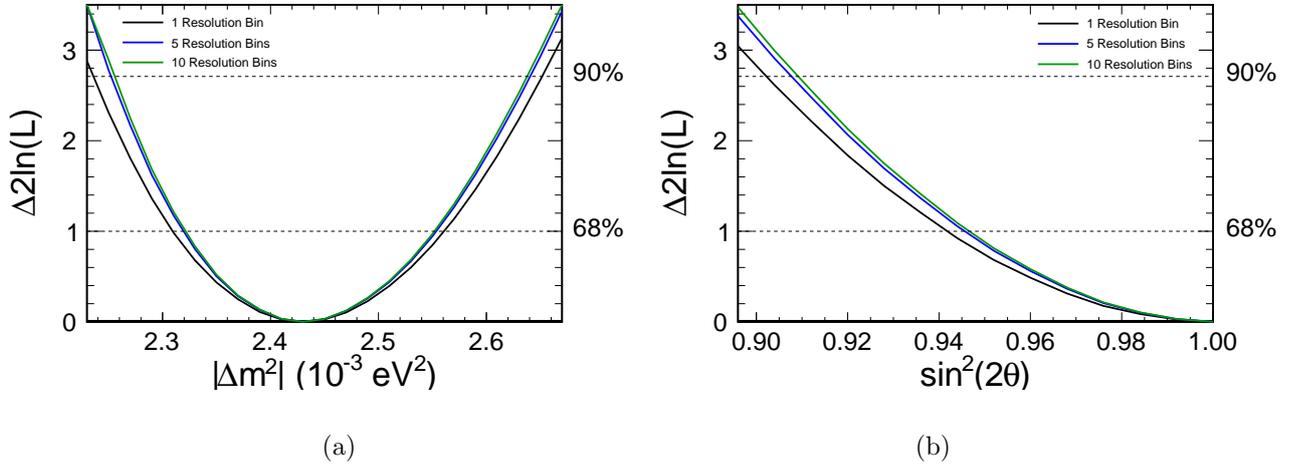


Figure 7.10: Projections of the log-likelihood surface onto the (a) Δm^2 and (b) $\sin^2(2\theta)$ axes, for the statistical fit with one (black), two (blue) and ten (green) bins of resolution. The contours are generated from Far Detector MC scaled to the data exposure of 7.2×10^{20} PoT, oscillated at $\sin^2(2\theta) = 1.0$ and $\Delta m^2 = 2.43 \times 10^{-3} \text{ eV}^2$.

lation fit individually using the fitting software; somewhat counter-intuitively, the best resolution bin does not exhibit the best sensitivity as shown in Figure 7.13(a). The remaining bins exhibit the expected relationship, with the events with the worst resolution providing the worst sensitivity to the oscillation parameters. The division into resolution bins is performed in reconstructed neutrino energy, yet the oscillations are applied in true neutrino energy; thus the true energy distribution in each resolution bin gives it its power. The true neutrino energy distribution for the best 20% of events is shifted upwards in true energy compared to the adjacent resolution bin, depleting the region of the oscillation dip and worsening the sensitivity in the best quantile relative to its neighbouring quantiles.

To explain the sensitivity in the lowest resolution bin, a variety of event variables were investigated. The k NN shower energy exhibits a hard cutoff below 0.25 GeV. By pathologically assigning events a reconstructed energy lower than their true energy, this event set is assigned a better resolution value due to this missing energy. The highest zero shower energy event proportion is present in the best resolution bin, as shown in Figure 7.12. This has the effect of biasing the true energy distribution upwards and so depleting the number of events with true energy in the oscillation region in this bin.

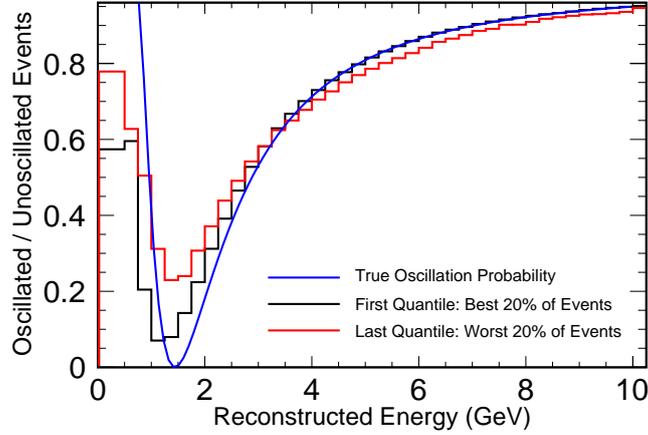


Figure 7.11: Ratios of oscillated and unoscillated neutrino spectra, for the best and worst resolution bins for five quantiles oscillated at $\sin^2(2\theta) = 1.0$ and $\Delta m^2 = 2.43 \times 10^{-3} \text{ eV}^2$. The true oscillation probability function is also shown in the region of interest (the true oscillation probability function is not drawn below 0.7 GeV).

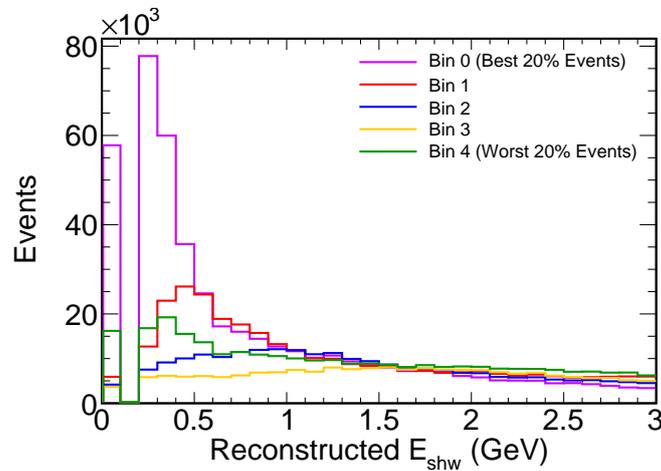
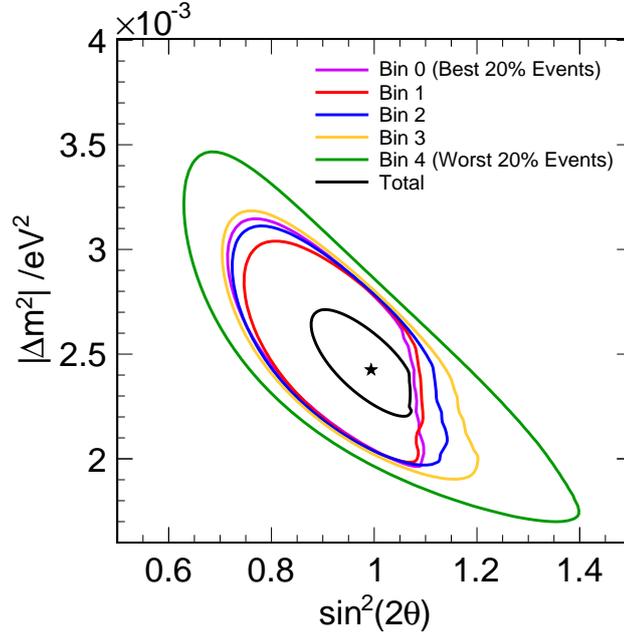
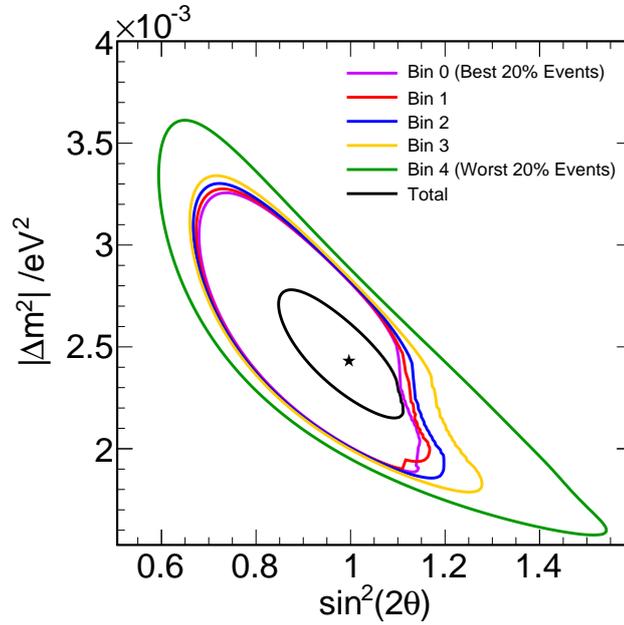


Figure 7.12: Reconstructed shower energy distributions by resolution bin, for high statistics Far Detector Monte Carlo over all Runs.

The sensitivity contours obtained by performing an oscillation fit with an additional selection applied of $E_{\text{shw}}^{\text{reco}} > 0 \text{ GeV}$ are shown in Figure 7.13(b), and exhibit the expected relationship between bins of resolution, with the best resolution bin providing the best sensitivity.



(a)



(b)

Figure 7.13: 90% confidence limits for each individual bin of resolution compared to the total fit, for (a) all events passing the selection, and (b) cutting out events with reconstructed $E_{\text{shw}} = 0$ GeV. The contours are generated from Far Detector MC scaled to the data exposure of 7.2×10^{20} PoT, oscillated at $\sin^2(2\theta) = 1.0$ and $\Delta m^2 = 2.43 \times 10^{-3}$ eV². Discontinuities in the unphysical region are due to low statistics for individual resolution bins, resulting in negative bin content predictions which are ignored by the $2 \ln \mathcal{L}$ function.

7.2.2 Detector Effects

The inclusion of resolution information splits the CC ν_μ sample into five separate bins, each of which contains events with different y-distributions, numbers of QE, RES and DIS events and Neutral Current background events. As a result, systematic mis-modelling of certain parameters in the Monte Carlo may affect each resolution bin to a different degree, and thus using the resolution information could potentially increase the sensitivity of the oscillation fit results to potential biases.

Three possible sources of bias are considered here; a smeared or biased shower or track energy; mis-modelling of the event kinematics; or a significantly underestimated Neutral Current background. To test the impact of these effects, Far Detector fake data spectra are generated from Monte Carlo events with a biased shower or track energy, adjusted event kinematics or a scaled up neutral current component. The fitting software then compares predicted spectra made from the nominal, unmodified Far Detector Monte Carlo to the biased data spectra. The resulting movements of the best fit point are compared, for an oscillation fit performed in one and five bins of resolution, to see if an additional sensitivity to each bias results from the use of the resolution information.

Biased Energy Scales Both the shower and track energy could be affected by a systematic offset, resulting from unknown calibration or reconstruction problems. Two different forms of this error are considered; firstly a reconstructed energy smeared from the true energy:

$$E'_{\text{reco}} = E_{\text{true}} + (E_{\text{reco}} - E_{\text{true}}) \times (1 + \alpha) \quad (7.8)$$

where α is a value picked at random from a Gaussian distribution of width w , and secondly a fractional offset of reconstructed from true energy:

$$E'_{\text{reco}} = E_{\text{reco}} \times (1 + \beta) \quad (7.9)$$

where β represents a fractional offset of the reconstructed energy. To select values of w and β for testing, the events with $3 \text{ GeV} < E_\nu < 5 \text{ GeV}$ were considered, as they inhabit the peak of the neutrino energy spectrum. For these events, the distribution of $\sigma_{\text{shw}}/E_\nu$ peaks at $\sigma_{\text{shw}} \approx 0.3/4$, and the $\sigma_{\text{trk}}/E_\nu$ distribution peaks at $\sigma_{\text{shw}} \approx 0.1/4$. Hence, $w_{\text{shw}} = 0.3$, $w_{\text{trk}} = 0.1$ and $\beta_{\text{shw}} = \pm 0.3$, $\beta_{\text{trk}} = \pm 0.1$ have been selected, as approximately four times greater than the most typical measurement error on shower and track energies, to ensure any possible sensitivity to mis-modelling is not underestimated.

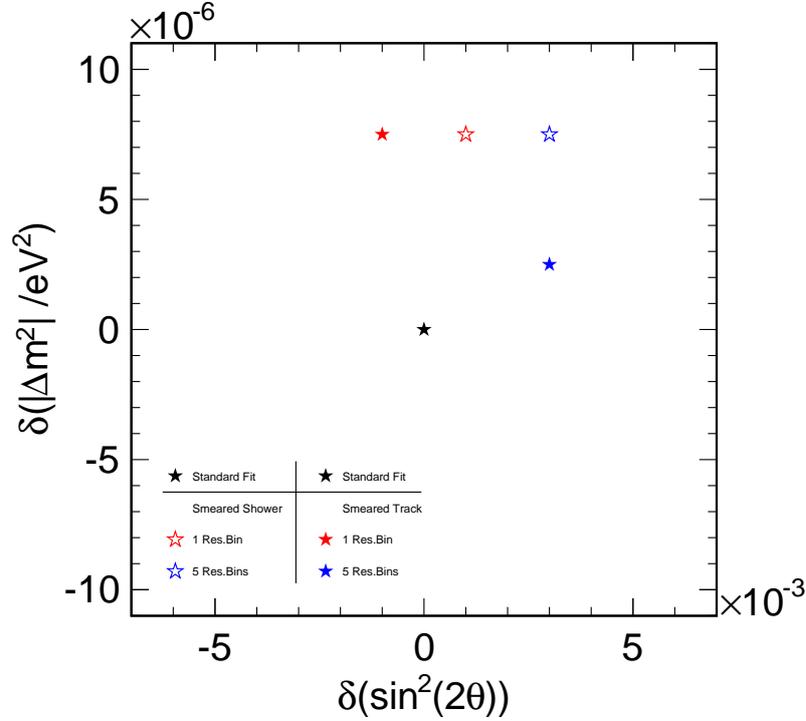


Figure 7.14: Best fit points for fake data with a smeared reconstructed shower energy (empty red and blue stars) or track energy (filled red and blue stars), using an unbiased prediction Far Detector MC scaled to 7.2×10^{20} PoT.

By smearing the shower or track energies, the change in the best fit $\sin^2(2\theta)$ becomes larger with the move from one to five resolution bins, as shown in Figure 7.14. In the case of an energy offset by a factor β , the fitted value of $\sin^2(2\theta)$ is biased less with an increasing number of resolution bins as seen in Figure 7.15. The fitted value of Δm^2 however becomes more biased with an increasing number of quantiles for smeared and offset track energy, and remains constant with any change in shower energy. In all cases, the effect of the biasing in Δm^2 is significantly less than the precision to which the best fit parameter is quoted.

Event kinematics When events are divided by resolution, the peak of the distribution of the reconstructed kinematic variable y : $y = E_\nu - E_\mu/E_\nu$ is expected to shift upwards moving from the best to worst quantiles, as a greater proportion of energy is found in the shower and less in the track. Thus it would be expected that a greater number of quasi-elastic events would be found in the better quantiles, as they tend to be characterised

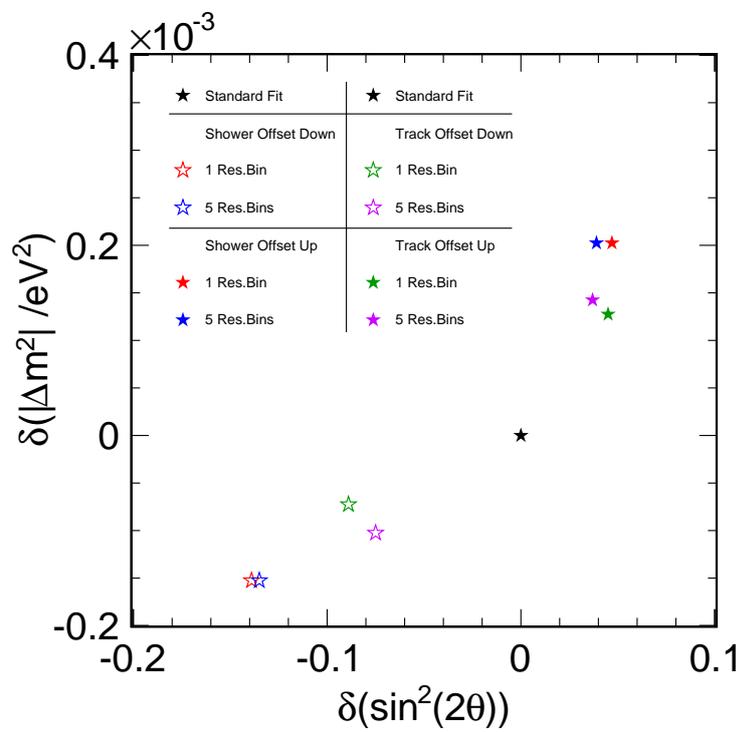


Figure 7.15: Best fit points for fake data with a positive and negative fractional offset applied to shower energy (red and blue stars) or track energy (green and violet stars), using an unbiased Far Detector MC prediction scaled to 7.2×10^{20} PoT.

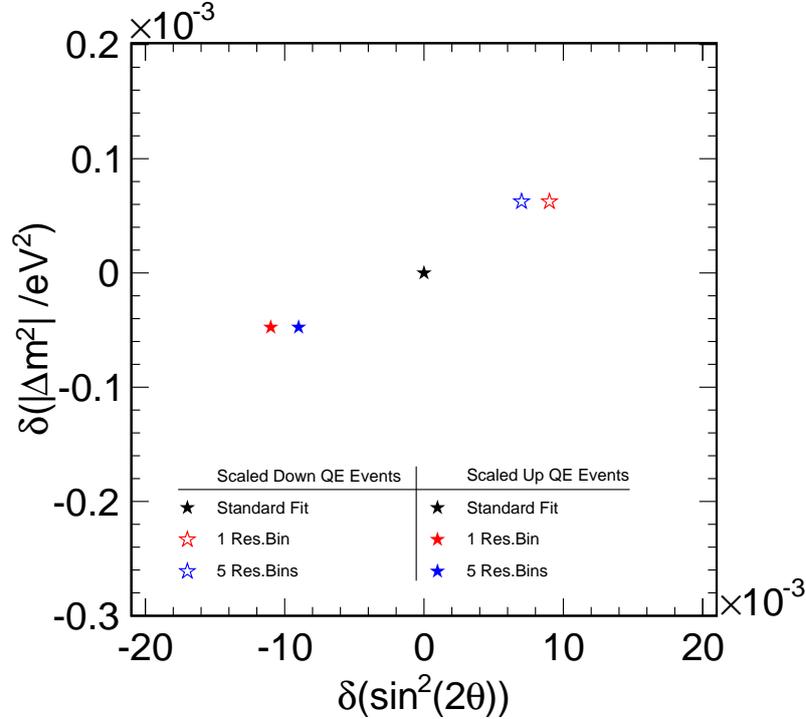


Figure 7.16: Best fit points for fake data with a 20% increase (open star) and 20% decrease (full star) in proportion of quasi-elastic events, using an unbiased Far Detector MC prediction scaled to 7.2×10^{20} PoT.

by small showers. By scaling up and down the proportion of quasi-elastic events in the sample by 20%, this has the effect of increasing the number of low- y events and decreasing the number of high- y events (and vice versa), to investigate the effect of poorly reconstructing the y value. Figure 7.16 shows by increasing the number of quantiles, the sensitivity to the shift in the proportion of QE events improves.

Enhanced Neutral Current Background The topology of an NC event is typically a low energy hadronic shower, with a falsely reconstructed very low energy track stub which permits it to pass the CC selection criteria. These characteristics mean that these events tend to be assigned poor resolution values compared to CC events at the same energy, as low energy CC events are dominated by a single track which can be more accurately measured than a low energy shower. The trend of assigning a greater proportion of the neutral current background events to the “worse” quantiles is shown in Table 7.3, with 62.3% of neutral current events in the Monte Carlo sample assigned

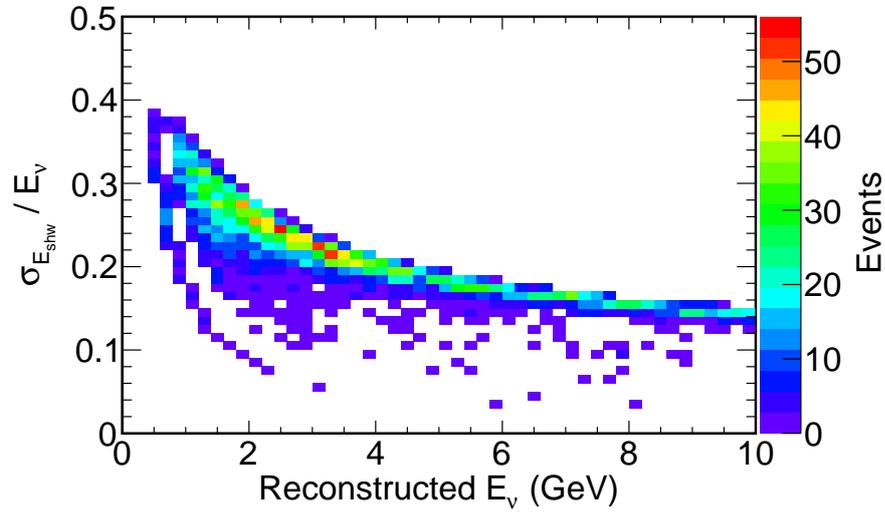
to the worst resolution bin, and 94.5% in the upper two bins.

The distribution of the fractional resolution σ_{E_ν}/E_ν for NC events peaks at a higher value than for CC events (Figure 7.17(b)), as expected from the distribution of the NC events amongst the higher resolution quantiles. The double peak in this distribution is explained by the rapid fall off in the shower fractional resolution with energy (Figure 7.17(a)); the events in the tail ‘pile up’ and a deficit is observed in the space between this region and where the events cluster at 2 – 4 GeV in reconstructed energy.

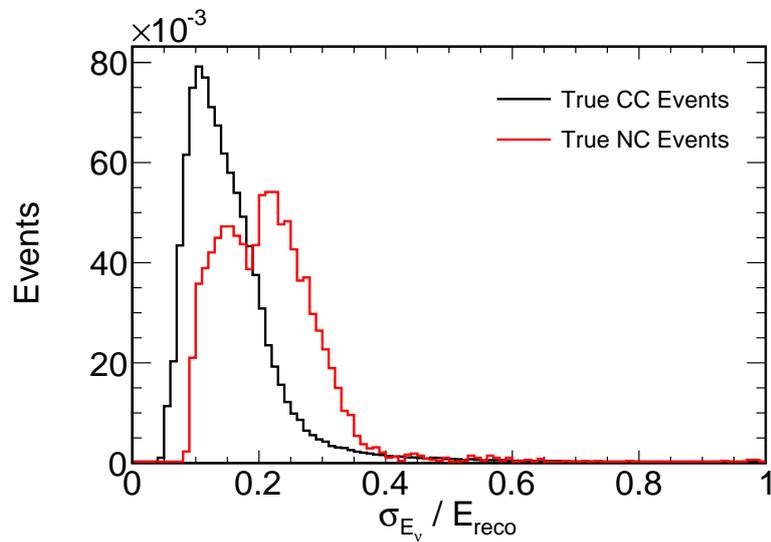
	NC Event Count	
	MC Exposure (7.5×10^{23})	Data Exposure (7.2×10^{20} PoT)
Bin 0	95	0.084
Bin 1	259	0.222
Bin 2	684	0.547
Bin 3	7042	5.11
Bin 4	13930	9.91
Total	22010	15.9

Table 7.3: Neutral current event counts in each bin of resolution for Far Detector MC, scaled to the data exposure of 7.2×10^{20} PoT and unscaled. 62.3% of neutral current events are assigned to the worst resolution bin, with 94.5% in the upper two bins.

If the level of the NC background is significantly mis-modelled, the highly non-uniform division of NC events between bins could result in an increase in sensitivity to the mis-modelling. Figure 7.18 shows the effect of moving from one to five resolution bins for a fourfold increase or decrease in the background level. The offset of the best fit point from the unbiased fit value decreases in both $\sin^2(2\theta)$ and Δm^2 when moving from one resolution bin to five resolution bins, indicating that by using resolution information the sensitivity to a mis-modelled NC background would decrease.



(a)



(b)

Figure 7.17: (a) The distribution of the shower resolution parameterisation $\sigma_{E_{\text{shw}}}/E_{\nu}$ as a function of reconstructed energy for true NC events in Run I high statistics Monte Carlo. When summed over energy, the cluster of events at low energy and the rapid fall off in the tail result in the double peak seen in (b).

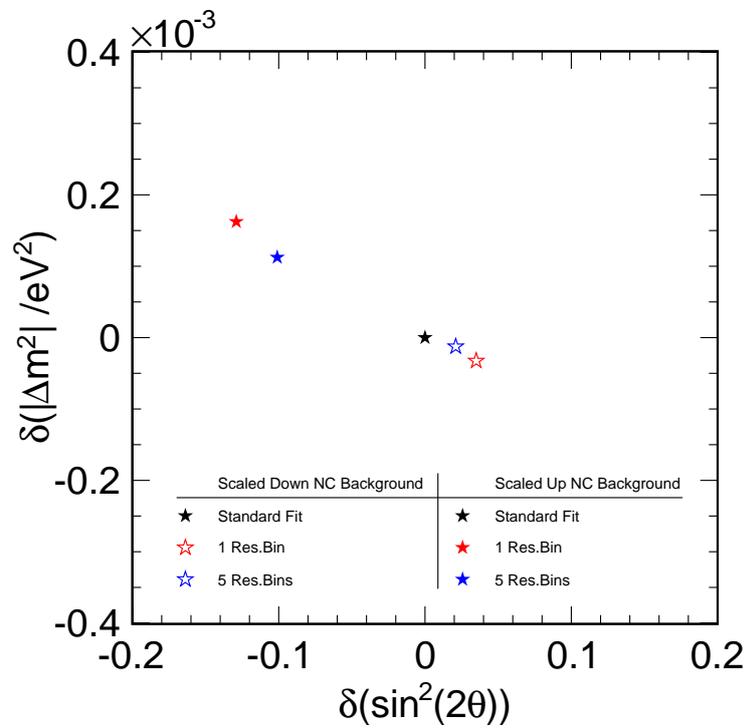


Figure 7.18: Best fit points for fake data with the neutral current background scaled up (open star) and down (full star) by a factor of four, using an unbiased Far Detector MC prediction scaled to 7.2×10^{20} PoT.

7.2.3 Alternative Model Discrimination

As noted in Chapter 6, the confidence limits on the measurement of a parameter using a log-likelihood method are defined by specific values of $2\Delta \ln \mathcal{L} = 2(\ln \mathcal{L} - \ln \mathcal{L}_{\text{bestfit}})$, where the contours of constant $\Delta 2 \ln \mathcal{L}$ include a specific percentage of the true measurements; thus the 68% confidence limit should contain 68% of a large number of repeated experiments, the ‘‘coverage probability’’. In the case of a measurement of parameters distributed with Gaussian errors, the 68% limit can also be expressed as ‘‘ 1σ ’’, where σ denotes the standard deviation of the Gaussian probability distribution.

To quantify the degree to which a model is disfavoured compared to another, the best fit point for oscillations is considered. The value of $2 \ln \mathcal{L}_{\text{bestfit}}$ is taken to equal the oscillation likelihood at this point, and is subtracted from the likelihood of the alternative model being tested. The resulting $2\Delta \ln \mathcal{L}$ corresponds to a coverage probability for single parameter estimation, and as it is considered to be a single parameter problem the following relationship applies:

$$n\sigma = \sqrt{2 \ln \mathcal{L}_{\text{model}} - 2 \ln \mathcal{L}_{\text{oscillations}}} \quad (7.10)$$

where σ is defined as above: the standard deviation of a Gaussian distribution.

The impact of fitting in bins of resolution on the alternative disappearance models of pure neutrino decoherence and decay is shown in Table 7.4; improving the degree of exclusion of decoherence by 0.3σ and decay by 0.2σ .

	Decoherence		Decay	
	$N_{Res} = 1$	$N_{Res} = 5$	$N_{Res} = 1$	$N_{Res} = 5$
$\mu^2 \alpha (10^{-3} eV^2)$	2.36	2.37	2.44	2.52
$\sin^2(2\theta)$	1.0	1.0	1.0	1.0
$\Delta 2 \ln \mathcal{L}$	68.12	72.81	33.74	38.46
Exclusion	8.25σ	8.53σ	5.81σ	6.20σ

Table 7.4: Improvements in alternative model discrimination with the addition of resolution information. Calculated from Far Detector MC scaled to the data exposure of 7.2×10^{20} PoT, oscillated at $\sin^2(2\theta) = 1.0$, $\Delta m^2 = 2.43 \times 10^{-3} eV^2$. The exclusion of each model is calculated as $\sqrt{2\Delta \ln \mathcal{L}}$, where $2\Delta \ln \mathcal{L}$ is the difference with reference to the oscillation best fit $2 \ln \mathcal{L}$.

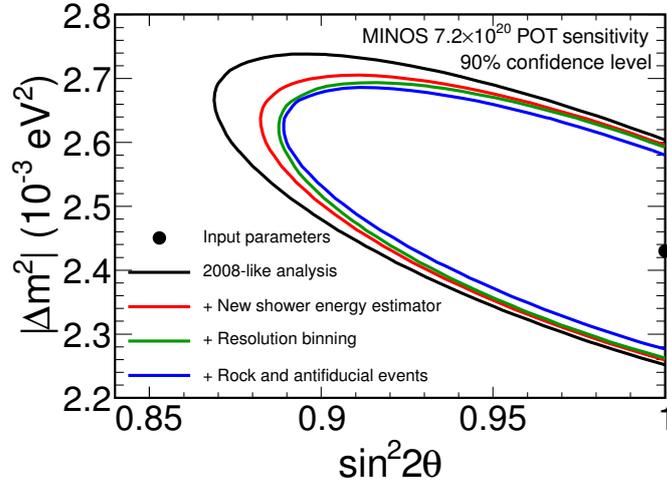


Figure 7.19: The MINOS 90% sensitivity to oscillation parameters with successive analysis improvements, for MC scaled to 7.2×10^{20} PoT of data. The red curve shows the improvement from using the kNN shower energy, the green when five resolution bins are used and the blue when rock and anti-fiducial events are added.

7.3 Summary

The inclusion of resolution information in conjunction with the `GhostFitter` as performed for this thesis, when combined with the other analysis modifications for the 2010 MINOS published analysis described in Chapter 5, results in a substantial improvement in oscillation sensitivity. Figure 7.19 shows the cumulative improvements as these modifications are included.

The following chapter discusses sensitivity studies performed to attempt to improve the exclusion of maximal mixing for an assumed non-maximal value of $\sin^2(2\theta)$. Chapter 9 then summarises the fits performed to the real data, including the results published in [39].

Chapter 8

Sensitivity to Non-Maximal Mixing

MINOS has previously published results where the atmospheric mixing angle $\sin^2(2\theta)$ is consistent with maximal mixing. The aim of the study presented in this chapter is to determine the level of accuracy at which MINOS could disfavour maximal mixing, if the true mixing angle $\sin^2(2\theta) \neq 1.0$. All sensitivities in this Chapter are based on Far Detector Monte Carlo, scaled to a data exposure of 7.25×10^{20} PoT, and assume $\Delta m^2 = 2.32 \times 10^{-3} \text{ eV}^2$, the best fit mass splitting from the following chapter. Only fiducial CC events ($\nu_\mu + \bar{\nu}_\mu$) are used.

8.1 Motivation

Figure 8.1 shows the projections of statistical log-likelihood surfaces onto $\sin^2(2\theta)$, for true oscillation parameters of $\Delta m^2 = 2.32 \times 10^{-3} \text{ eV}^2$ and $\sin^2(2\theta) = 0.9$. The oscillation fits were performed with the neutrino energy spectra binned in true neutrino energy and reconstructed neutrino energy; with and without resolution binning implemented. The “exclusion of maximal mixing” is defined for the purpose of this study as the $\Delta 2 \ln \mathcal{L}$ value at $\sin^2(2\theta) = 1.0$, when the true $\sin^2(2\theta)$ oscillation parameter is defined to be $\sin^2(2\theta) = 0.9$ by the fitting software. The fit performed in true neutrino energy excludes $\sin^2(2\theta) = 1.0$ at $\Delta 2 \ln \mathcal{L} = 9$, which corresponds to 3σ as defined by Equation 7.10. When using the k NN shower energy with and without resolution binning when running the fitting software, an exclusion of $\approx 1\sigma$ is achieved.

The sensitivity of a MINOS analysis to the oscillation parameters is governed by how well the depth and position of the oscillation dip can be resolved, with the depth

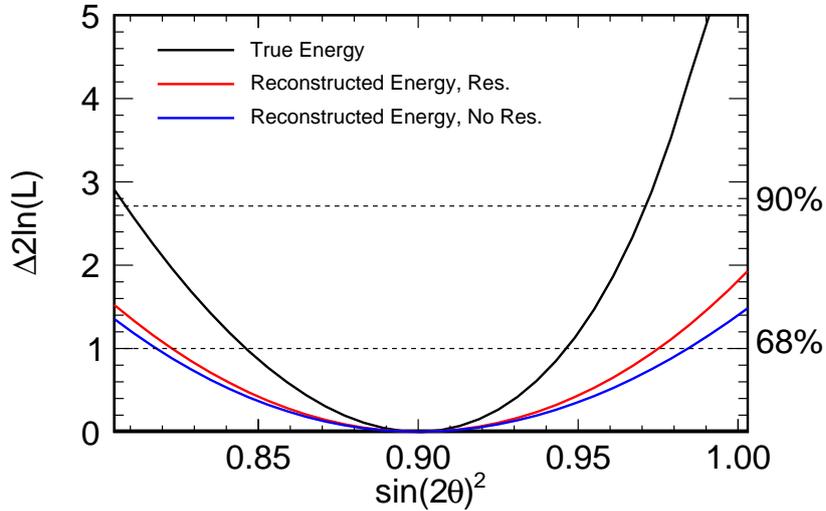


Figure 8.1: Projections of the statistical log-likelihood surfaces onto $\sin^2(2\theta)$ for oscillation fits performed with spectra binned in true (black) and reconstructed energy with (red) and without (blue) resolution binning. The underlying oscillation parameters were set to be $\sin^2(2\theta) = 0.9$ and $\Delta m^2 = 2.32 \times 10^{-3} \text{ eV}^2$. The $\Delta 2 \ln \mathcal{L}$ values corresponding to 68% confidence (1σ) and 90% confidence are shown by the dashed lines.

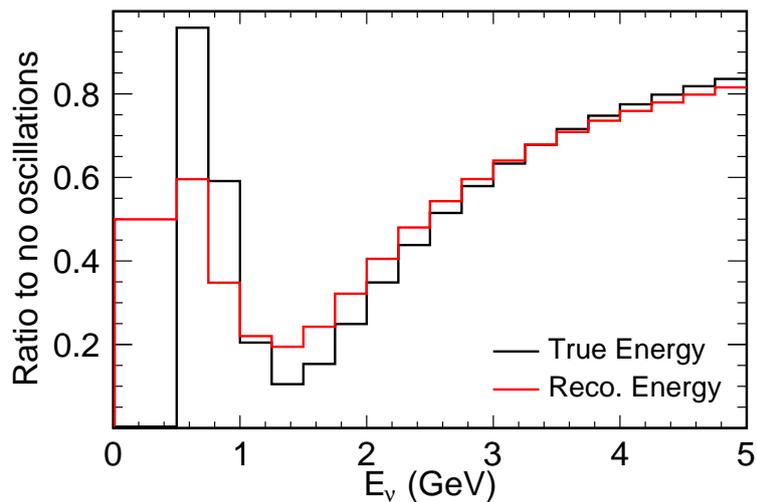


Figure 8.2: The ratio of the best fit oscillated spectrum to the unoscillated prediction is shown for fits to spectra binned in true (black) and reconstructed (red) energy. The underlying oscillation parameters were set to be $\sin^2(2\theta) = 0.9$ and $\Delta m^2 = 2.32 \times 10^{-3} \text{ eV}^2$.

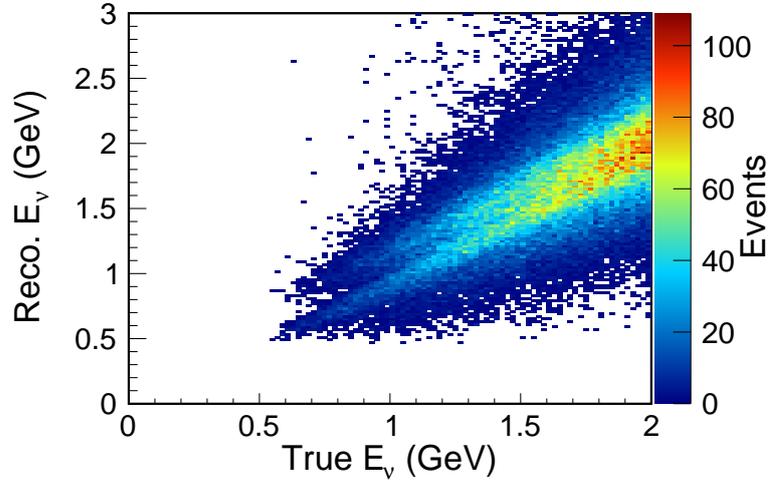


Figure 8.3: The reconstructed vs true neutrino energy distribution, from Far Detector Monte Carlo summed over all Runs.

a function of the $\sin^2(2\theta)$ parameter. As shown in Figure 8.2, there is a clear “filling in” of the oscillation dip when comparing oscillation fits in true and reconstructed neutrino energy. The events in this region are either poorly reconstructed very low energy events, or events with higher energies than the oscillation dip that are systematically shifted into the 1 – 2 GeV region. Poorly reconstructed low energy events may benefit from an improved reconstruction method which utilises the relativistic kinematics of the interaction to obtain the neutrino energy.

As has been discussed in previous Chapters, the k NN shower energy reconstruction sets the shower energy of events below 0.25 GeV to zero. These events have neutrino energies that are systematically underestimated, as is shown in the double peak effect in Figure 8.3. Events with their reconstructed shower energy set to zero will lie in the peak with lower neutrino energy, where events with shower energies set to 0.25 GeV forming a second peak with higher reconstructed neutrino energy. These events have their neutrino energies corrected as described in Section 8.2.4, and may also benefit from a kinematics-based reconstruction method.

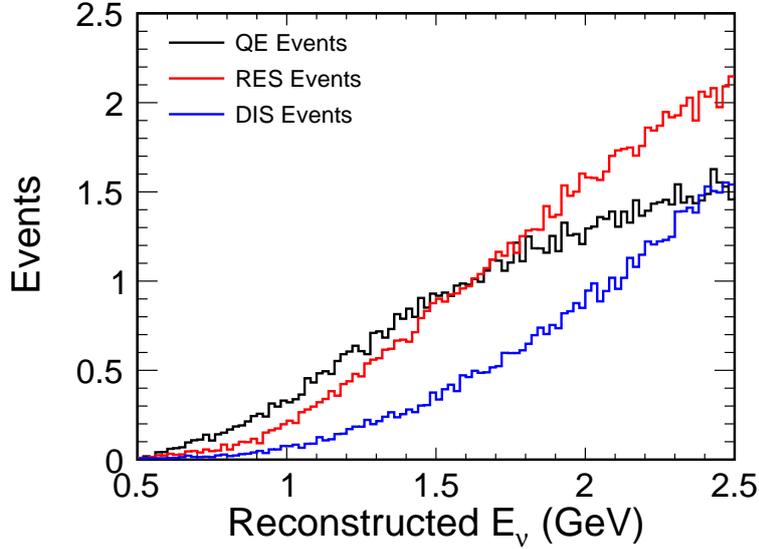


Figure 8.4: Distributions of reconstructed neutrino energy, separately for Quasi-Elastic (QE), Resonance (RES) and Deep Inelastic Scattering (DIS) events in the region of the oscillation dip for $\sin^2(2\theta) = 0.9$ and $\Delta m^2 = 2.32 \times 10^{-3} \text{ eV}^2$.

8.2 Improving the Energy Reconstruction

As described in Chapter 2, there are three types of CC neutrino interaction; Quasi-Elastic (QE), Resonance production (RES) and Deep Inelastic Scattering (DIS). Figure 8.4 shows the distributions of the three CC event types in the region of the oscillation dip, for $\sin^2(2\theta) = 0.9$ and $\Delta m^2 = 2.32 \times 10^{-3} \text{ eV}^2$. The dominance of QE and RES events over DIS implies that an improvement in the shower energy reconstruction using the kinematics of neutrino-nucleon scattering may be possible for these low energy events, and thus improve the sensitivity to $\sin^2(2\theta)$.

8.2.1 Kinematics of Neutrino-Nucleon Scattering

By conserving four-momentum and assuming a stationary neutron target, an estimate of the neutrino energy can be easily derived from the event kinematics:

$$E_\nu = \frac{2m_n E_\mu - m_\mu^2 - m_n^2 + W^2}{2m_n - 2E_\mu + 2E_\mu \cos(\theta)} \quad (8.1)$$

where m_n and m_μ are the neutron and muon masses respectively, θ is the angle of the muon relative to the incoming neutrino and W the invariant mass of the recoiling hadronic system.

In a Quasi-Elastic (QE) event ($\nu_\mu + n \rightarrow l + p$), the neutrino scatters and modifies the target nucleus, with the resulting hadronic system constituting a single proton: $W^2 = m_p^2$. The hadronic system may have very little energy, and so often no shower is measured by the detector. The nature and kinematics of QE events are explored in detail in [93]. Resonance production events (RES) in MINOS ($\nu_\mu + N \rightarrow N + \mu$) typically produce a $\Delta(1232)$: $W^2 = m_{\Delta(1232)}^2$ (though a range of other states are possible), which decays to a pion, and can look very similar to QE events with a muon track and low energy shower. Deep inelastic scattering events (DIS) break up the nucleus and produce large hadronic energy deposits, corresponding to a broad distribution in energy of possible values of W^2 . For this reason correcting the QE and RES events will be the focus in attempting to improve the neutrino energy reconstruction.

8.2.2 Event Selection in True Neutrino Energy

An event selection is required to identify the events for which a corrected shower energy would improve their reconstruction. To determine if a significant improvement in sensitivity is possible, the analysis was first performed selecting events in true neutrino energy. If an improvement was seen, a selection in reconstructed quantities would then be developed, as discussed in Section 8.4. To identify the region of poorly reconstructed events in true neutrino energy, a series of oscillation fits were performed using the `GhostFitter` software, where all events with a true neutrino energy below $E_{\text{true}} < x$ GeV (the “threshold” energy) were binned in true energy, and all events above the threshold were binned in reconstructed neutrino energy. The value of x was varied between 0 – 4 GeV.

Figure 8.5(a) shows the values of $\Delta 2 \ln \mathcal{L}$ at $\sin^2(2\theta) = 1.0$ (as the exclusion of maximal mixing was defined previously as the value of $\Delta 2 \ln \mathcal{L}$ at maximal mixing), with the energy threshold for using true neutrino energy binning plotted on the x -axis. The rapid rise in $\Delta 2 \ln \mathcal{L}$ between $x = 1$ GeV and $x = 2$ GeV indicates those events with $E_{\text{true}} < 2$ GeV would benefit the sensitivity most if their energy reconstruction was improved. This point is reinforced by the projection of the log-likelihood surfaces onto $\sin^2(2\theta)$ with different true energy thresholds shown in Figure 8.5(b), where a true energy cut of 2 GeV provides a sensitivity almost equal to using true neutrino energy for all events in the oscillation fit. Selecting events with a true energy less than 2 GeV

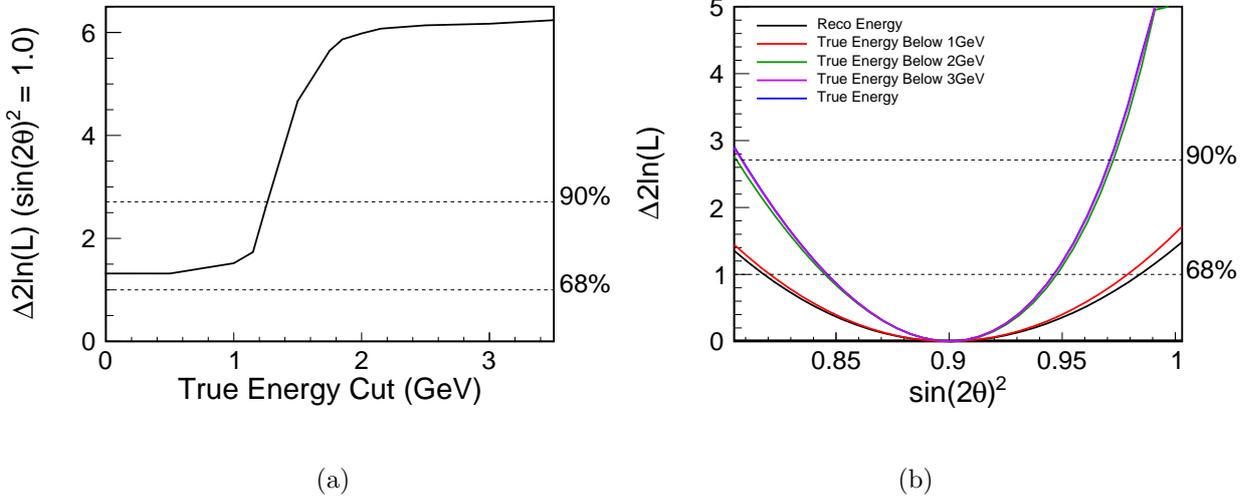


Figure 8.5: (a) $\Delta 2 \ln \mathcal{L}$ at $\sin^2(2\theta) = 1.0$ against the threshold energy variable x GeV, where all events with true neutrino energy below x GeV had their true energy used in the oscillation fit, and all events above the threshold had their reconstructed energy used by the fitting software. (b) Projections onto $\sin^2(2\theta)$ over the statistical log-likelihood surfaces for oscillation fits performed with values of $x = 1, 2$ and 3 GeV, all events binned in true energy and all events binned in reconstructed energy.

with a high purity and efficiency should maximise the improvement in sensitivity to non-maximal mixing.

8.2.3 Potential Sensitivity Gain

The possible gains in sensitivity to be made by correcting the energies of the truly QE and RES events are shown in Figure 8.6, for selected events with a true neutrino energy $E_{\text{true}} < 2$ GeV. Correcting truly quasi-elastic events provides most of the improvement, with the cumulative effect of correcting both interaction types shown in blue. The fractional improvement in the exclusion of maximal mixing, I , is defined for this Chapter by the following expression:

$$I = (\Delta \ln \mathcal{L}_{\text{New Reco.}} - \Delta \ln \mathcal{L}_{\text{Std. Reco.}}) / \Delta \ln \mathcal{L}_{\text{Std. Reco.}} \quad (8.2)$$

where $\Delta \ln \mathcal{L}$ is defined as the value at maximal mixing, for an oscillation fit with true $\sin^2(2\theta) = 0.9$. Recalculating the neutrino energy for the two event types results in a

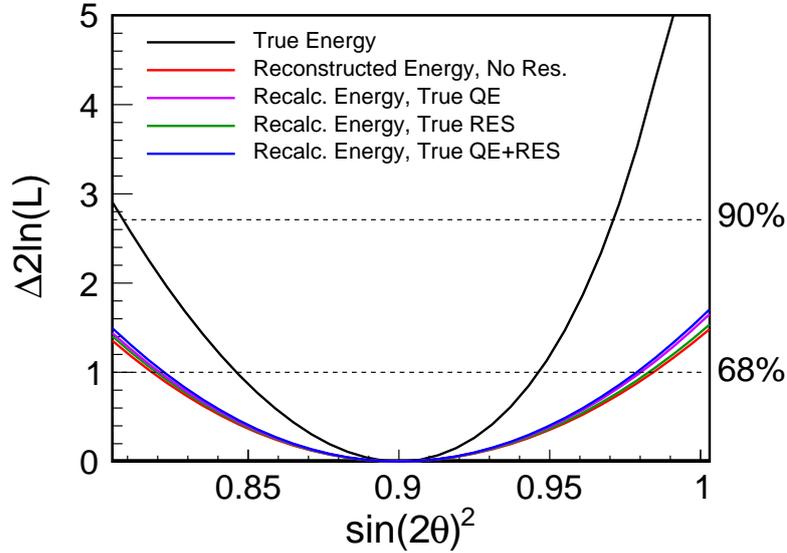


Figure 8.6: Projections onto $\sin^2(2\theta)$ over the statistical log-likelihood surfaces for oscillation fits performed with spectra binned in true energy (black) and reconstructed energy (red). The oscillation fits with true quasi-elastic events corrected (violet), true resonance events corrected (green), and both event types corrected (blue) are also shown. The energy correction was performed using Equation 8.1, for events with $E_{\text{true}} < 2$ GeV, and underlying oscillation parameters set to $\sin^2(2\theta) = 0.9$ and $\Delta m^2 = 2.32 \times 10^{-3} \text{ eV}^2$.

15% improvement in the level of exclusion of maximal mixing as defined above, and so is the maximum improvement that can be achieved using the kinematics of the interaction and a perfect event selection for the chosen oscillation parameters.

8.2.4 Correcting the k NN Shower Energy

The k NN shower energy reconstruction developed in [74] is summarised in Chapter 5, and has been seen to exhibit a hard cutoff below 0.25 GeV, where events cannot be assigned an energy greater than zero and less the 0.25 GeV; as shown in Figure 8.7. This impact of this feature is clearly visible, where the events with energies which would be reconstructed between 0 – 0.25 GeV have their energy artificially inflated or suppressed to zero, causing the double peak pathology in the reconstructed invariant mass W^2 shown in Figure 8.8.

To ensure the k NN shower energy was causing these features, the invariant mass

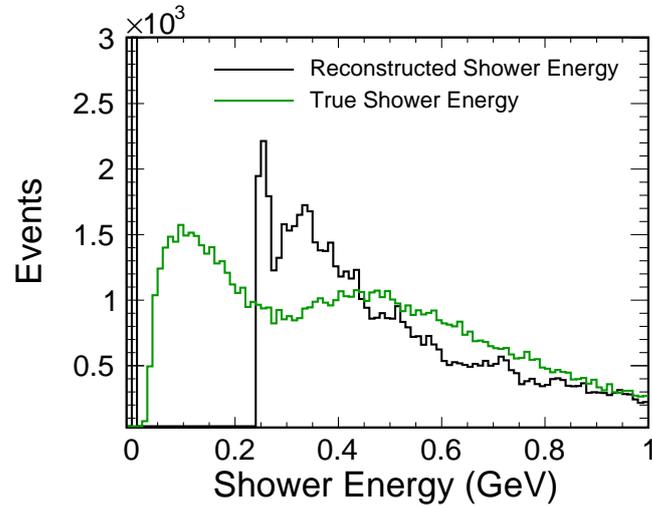


Figure 8.7: Reconstructed k NN shower energy distribution, for high statistics Far Detector Monte Carlo over all Runs.

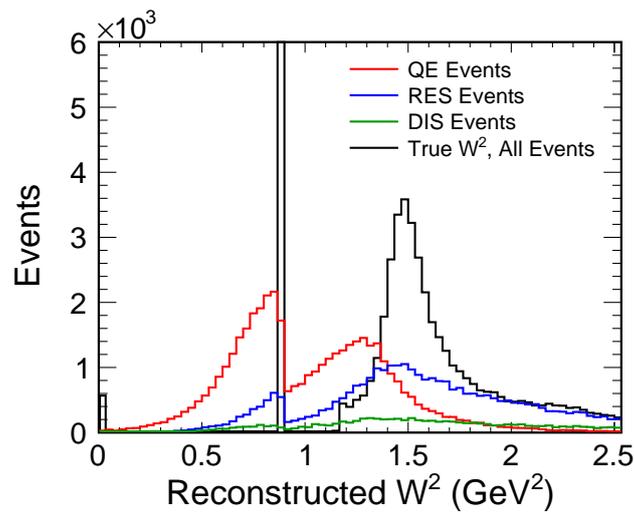


Figure 8.8: True and reconstructed invariant mass W^2 , split by event type for all selected events with $E_{\text{true}} < 2$ GeV.

distribution was calculated for Far Detector events with each variable in Equation 8.1 set to its true value; with one variable in turn set to its reconstructed value, to observe from which variable the double peak originates. The double peak does not appear when reconstructed muon energy or angle are substituted, but is present when the reconstructed neutrino energy is used. This confirms the reconstructed shower energy is responsible for the effects observed.

For the standard CC analysis the k NN shower energy artifact has little to no impact, as the first energy bin in the reconstructed spectra has a width of 500 MeV. However, for this analysis the double peaks need to be merged. As can be seen in Figure 8.8, the secondary peak in the QE reconstructed invariant mass overlaps with the position of the $\Delta(1232)$ resonance peak, and so if the energy remains uncorrected the ability to discriminate between QE and RES events when using the event kinematics will be diminished.

The procedure applied to each event to calculate a modified shower energy is summarised as follows:

- i. Select all events for which $E_{\text{shw}}^{\text{reco}} = 0$ GeV.
- ii. Parameterise the true shower energy distribution of these events with a combined Gaussian + Landau function, as shown in Figure 8.9.
- iii. Normalise the combined Landau and Gaussian function to one for use as a probability distribution.
- iv. For each event with $E_{\text{shw}}^{\text{reco}} = 0$, choose 100 values at random from the probability distribution. Each event becomes effectively 100 events, each with a new provisional reconstructed shower energy and a weight, $w = 1/100$.

The distributions of the reconstructed invariant mass squared W^2 for the three event types are shown in Figure 8.10, after the correction steps detailed above have been performed. The double peak feature has been removed. The following section describes how the kinematics of the interactions are utilised to calculate a final reconstructed shower energy.

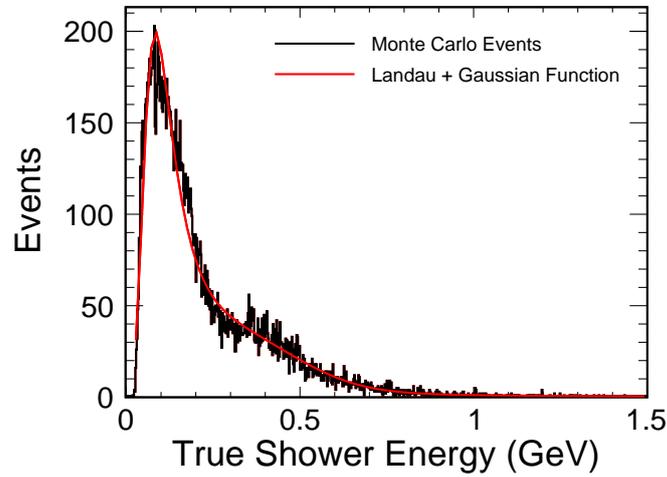


Figure 8.9: True shower energy distribution for events with $E_{\text{shw}}^{\text{reco}} = 0$.

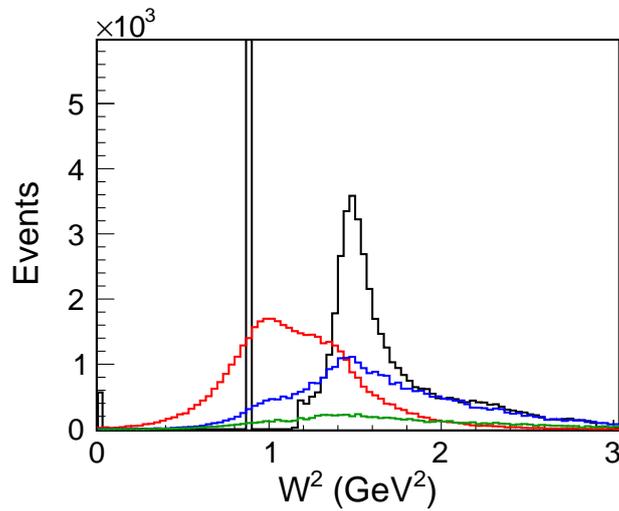


Figure 8.10: True and reconstructed invariant mass W^2 , split by event type for selected events in all Runs, with $E_{\text{true}} < 2$ GeV. Each event corresponds to a single shower energy selected from the probability distributions, weighted by a factor of $w = 1/100$. The double peak pathology introduced at low energy by the k NN shower energy has been removed.

8.2.5 Neutrino Energy Calculation

The approach taken to improve the neutrino energy estimator for events in the region of the oscillation dip is to define a χ^2 function, in terms of the reconstructed variables used in calculating the neutrino energy from Equation 8.1. The free variables are penalised when they deviate from the values calculated by the standard MINOS reconstruction. This function can then be minimised via a Lagrange multipliers method: $f = \chi^2 - \lambda \times g$, where Equation 8.1 acts as the constraint function g .

The value of the invariant mass W is taken in turn to be the proton mass and the mass of the $\Delta(1232)$, and two potential energies are calculated for each event, $E_{\text{reco}}^{\text{QE}}$ and $E_{\text{reco}}^{\text{RES}}$. A cut placed on the $\Delta\chi^2$ acts as the discriminant between QE-like and RES-like events. The sensitivity to non-maximal mixing can then be determined using the `GhostFitter` fitting package.

The χ^2 function to be minimised for each Monte Carlo event is defined below:

$$\chi^2 = \frac{(E_\mu - E_\mu^d)^2}{\sigma_{E_\mu}^2} + \frac{(E_{\text{shw}} - E_{\text{shw}}^d)^2}{\sigma_{E_{\text{shw}}}^2} + \frac{(\cos(\theta) - \cos(\theta)^d)^2}{\sigma_{\cos(\theta)}^2} \quad (8.3)$$

where E_μ , E_{shw} and $\cos(\theta)$ are the muon energy, shower energy and track angle calculated for the QE or RES hypotheses using Equation 8.1. E_μ^d , E_{shw}^d and $\cos(\theta)^d$ are the standard Monte Carlo reconstructed energies and angle. $\sigma_{E_{\text{shw}}}$ and σ_{E_μ} are taken as the event shower and track energy resolutions parameterised in Chapter 7. $\sigma_{\cos(\theta)}$ is derived in Section 8.2.6.

To determine $\Delta\chi^2$ for each event, the software minimises the function given below, using the GNU standard library's multi-dimensional root finder:

$$f = \frac{(E_\mu - E_\mu^d)^2}{\sigma_{E_\mu}^2} + \frac{(E_{\text{shw}} - E_{\text{shw}}^d)^2}{\sigma_{E_{\text{shw}}}^2} + \frac{(\cos(\theta) - \cos(\theta)^d)^2}{\sigma_{\cos(\theta)}^2} - \lambda * \left(E_\mu + E_{\text{shw}} - \frac{2m_n E_\mu - m_\mu^2 - m_n^2 + W^2}{2m_n - 2E_\mu + 2E_\mu \cos(\theta)} \right) \quad (8.4)$$

by solving the following system of equations with the additional constraint given by

setting Equation 8.1 equal to the sum of the track and shower energies (Equation 8.8).

$$\begin{aligned} \frac{\partial f}{\partial E_\mu} &= \frac{2(E_\mu - E_\mu^d)}{\sigma_{E_\nu}^2} \\ &\quad - \lambda \left[\frac{2m_n}{2m_n - 2E_\mu + 2E_\mu \cos(\theta)} \right. \\ &\quad \left. - \frac{(2m_n E_\mu - m_\mu^2 - m_n^2 + W^2)(2 \cos(\theta) - 2)}{(2m_n - 2E_\mu + 2E_\mu \cos(\theta))^2} - 1 \right] \end{aligned} \quad (8.5)$$

$$\frac{\partial f}{\partial E_{\text{shw}}} = \frac{2(E_{\text{shw}} - E_{\text{shw}}^d)}{\sigma_{E_{\text{shw}}}^2} + \lambda \quad (8.6)$$

$$\begin{aligned} \frac{\partial f}{\partial \cos(\theta)} &= \frac{2(\cos(\theta) - \cos(\theta)^d)}{\sigma_{\cos(\theta)}^2} \\ &\quad + \lambda \left[\frac{(2E_\mu)(2m_n E_\mu - m_\mu^2 - m_n^2 + W^2)}{(2m_n - 2E_\mu + 2E_\mu \cos(\theta))^2} \right] \end{aligned} \quad (8.7)$$

$$E_\mu + E_{\text{shw}} = \frac{2m_n E_\mu - m_\mu^2 - m_n^2 + W^2}{2m_n - 2E_\mu + 2E_\mu \cos(\theta)} \quad (8.8)$$

A cut is applied to the resulting $\Delta\chi^2$ value, to denote each event as ‘‘QE-like’’ or ‘‘RES-like’’. The event energy is recalculated using Equation 8.1 with the invariant mass value used based on whether the event is QE-like or RES-like.

8.2.6 Muon Angle Resolution

To determine the resolution of the angle of the muon track for use in Equation 8.3, three variables were considered to be likely to be correlated with the resolution of the angle of the muon track:

- i. Cosine of the muon angle, $\cos(\theta)$.
- ii. Number of track planes, N .
- iii. Momentum of the muon E_μ .

The resolution $\sigma_{\cos(\theta)}$ was defined as the Gaussian width of the $\cos(\theta)_{\text{reco}} - \cos(\theta)_{\text{true}}$ distribution in each bin of angle, number of planes or muon momentum, as was done for the shower and track energy resolutions in Section 7.1.

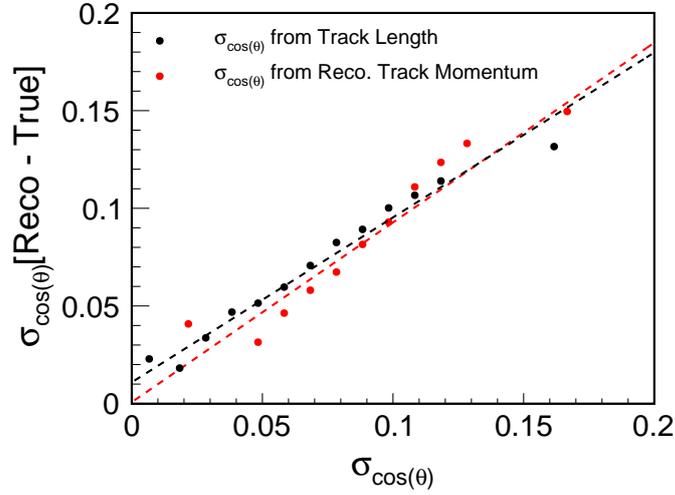


Figure 8.11: Distributions of Gaussian widths of $\cos(\theta)_{\text{reco}} - \cos(\theta)_{\text{true}}$ distributions in bins of $\sigma_{\cos(\theta)}$, with the resolution calculated using the track length in planes (black) and the reconstructed muon momentum (red). The dashed lines are linear fits, and give gradients of 0.84 and 0.92 respectively. A gradient of one would signify the parameterisation perfectly reconstructs the resolution of the angle.

To determine which choice of variable gives the best representation of the angular resolution, the best fit lines between the Gaussian widths of the $\Delta \cos \theta = (\cos \theta_{\text{reco}} - \cos \theta_{\text{true}})$ distributions in bins of $\sigma_{\cos \theta}$ were determined, and plotted in Figure 8.11. $\sigma_{\cos \theta}$ is calculated for each selected event with $E_{\text{true}} < 2 \text{ GeV}$ using two of the parameterisations of the resolution: in bins of track length in number of planes, N ; or reconstructed muon momentum, E_{μ} . The parameterisation exhibits a lack of structure in the reconstructed muon angle case and was not considered. The gradient of the best fit line equals $m = 0.92$ for the track momentum case, with $m = 0.84$ for the track length case. A gradient of one would signify the parameterisation perfectly reconstructs the resolution of the angle. Thus, the resolution parameterised in track momentum is utilised.

Figure 8.12 shows the resolution of the direction cosine of the muon angle as a function of the reconstructed muon momentum. The resulting resolution, with the function used to parameterise the distribution, is given below:

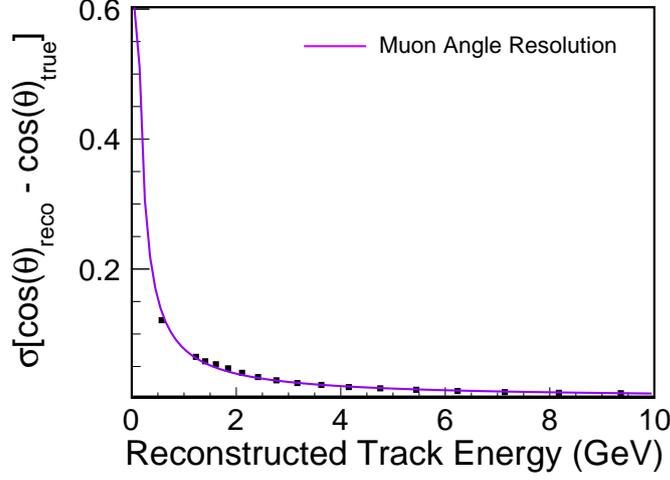


Figure 8.12: Muon angle resolution parameterisation as a function of reconstructed muon momentum, generated using selected Far Detector Monte Carlo events; for all neutrino event energies.

Reconstructed muon momentum, E_μ :

$$\sigma_{\cos(\theta)} = \frac{Ae^{-BN+C}}{N} \Rightarrow \sigma_{\cos(\theta)} = \frac{0.08e^{0.01E_\mu-0.01}}{E_\mu} \quad (8.9)$$

The decaying form was selected for the parameterisation as providing the best representation of the improvement in angle resolution with increasing muon momentum.

8.3 Results for Events Selected with $E_{\text{true}} < 2 \text{ GeV}$

Figure 8.13 shows the $\Delta\chi^2$ distribution for all Far Detector Monte Carlo events for all Runs, split into true QE and RES events. The k NN shower energy was corrected for events with $E_{\text{true}} < 2 \text{ GeV}$ as detailed in Section 8.2.4, and $\Delta\chi^2$ calculated by solving the system of equations given as Equations 8.5 to 8.8. Two distinct regions are visible, with predominantly QE like events above $\Delta\chi^2 = 0$, and RES events dominating below $\Delta\chi^2 = -0.5$. The selection cut applied to $\Delta\chi^2$ to distinguish QE-like and RES-like events is shown below:

$$\begin{aligned} \text{Resonance - like} &: \Delta\chi^2 < -0.2 \\ \text{Quasi - Elastic - like} &: \Delta\chi^2 \geq -0.2 \end{aligned} \quad (8.10)$$

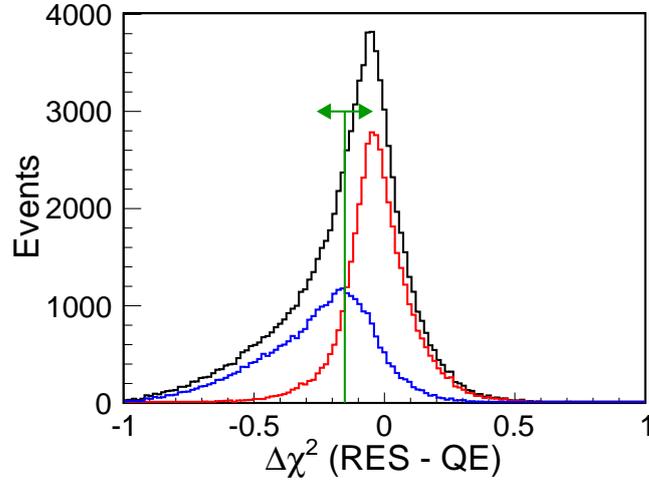


Figure 8.13: Average $\Delta\chi^2 = \chi^2_{RES} - \chi^2_{QE}$ distributions for events with corrected k NN shower energies for events with $E_{shw} = 0$. The average of the 100 $\Delta\chi^2$ values is calculated for each event. The green line denotes the position of the selection cut on the $\Delta\chi^2$ variable. All selected events with $\Delta\chi^2$ below the line have their energy calculated from Equation 8.1 assuming the $\Delta(1232)$ mass as the invariant mass W , and all above the line assuming the proton mass.

The projection onto $\sin^2(2\theta)$ over the statistical log-likelihood surface for the events with the modified energy reconstruction developed in this Chapter is shown in Figure 8.14, compared to an oscillation fit performed in true energy, and reconstructed energy without resolution binning applied. Correcting events using the $\Delta\chi^2$ variable improves the exclusion of maximal mixing, providing an improvement of 10% relative to the standard reconstruction; as defined by Equation 8.2. Correcting all true quasi-elastic and resonance events with neutrino energy below 2 GeV as shown in Section 8.2.3 resulted in an improvement of 15%. Alternative cut positions on $\Delta\chi^2$ were also attempted, yet did not provide a greater improvement than is seen using Equation 8.10.

The $(E_{reco} - E_{true})/E_{true}$ distributions for the corrected energy and the standard MINOS reconstructed energy are given in Figure 8.15, for all events with $E_{true} < 2$ GeV. The mean of the distribution of corrected true resonance events (plotted in blue in Figure 8.15(a)) is shifted low, with a mean of ~ -0.15 . This shift is not seen in Figure 8.15(b), where the distribution plotted is of resonance-like events, as identified by $\Delta\chi^2 < 0.2$. The offset of the true resonance distribution is due to the calculation of the energy of all resonance events with the reconstructed invariant mass: $W^2 = m_{\Delta(1232)}^2$, when multiple resonances with higher invariant masses than the $\Delta(1232)$ are present in

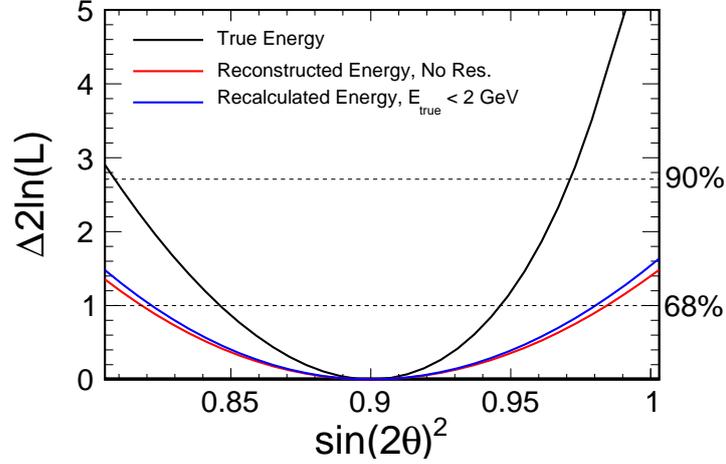


Figure 8.14: Projections onto $\sin^2(2\theta)$ over the statistical log-likelihood surfaces, for oscillation fits performed with spectra binned in true (black) and reconstructed energy (red), and energy recalculated using the new reconstruction method (blue) detailed in this chapter for events with $E_{\text{true}} < 2 \text{ GeV}$. The underlying oscillation parameters were set to be $\sin^2(2\theta) = 0.9$ and $\Delta m^2 = 2.32 \times 10^{-3} \text{ eV}^2$.

the Monte Carlo sample.

As a cross check, two alternative cases were considered: firstly, the $\Delta E/E$ distribution for events with their energy recalculated using the mean \overline{W}^2 for all true resonance events, where $\overline{W}^2 = 1.28 \text{ GeV}$; and secondly, determining the $\Delta E/E$ for true resonance events in the invariant mass peak corresponding to $\Delta(1232)$. In both cases, the mean of the $\Delta E/E$ for resonance events is centred on zero, and the overall distribution shifts to the right as expected. The oscillation sensitivity was also determined for the case where $\overline{W}^2 = 1.28 \text{ GeV}$ is used to correct the resonance-like events, but the sensitivity at maximal mixing worsened, and so $W^2 = m_{\Delta(1232)}^2$ was retained in the software framework.

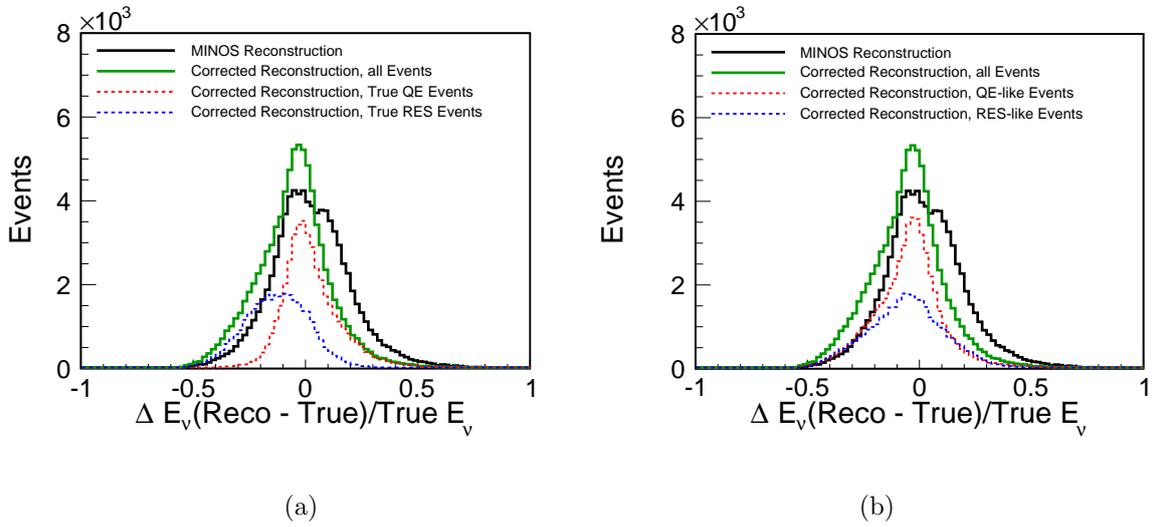


Figure 8.15: Ratio of $E_{\text{reco}} - E_{\text{true}}$ to E_{true} for the standard MINOS reconstruction (black) and the corrected energy reconstruction (green). The corrected energy reconstruction is also shown divided by (a) true quasi-elastic (red dashed) and true resonance (blue dashed), and (b) quasi-elastic-like (red dashed) and resonance-like (blue dashed), as defined by the $\Delta\chi^2$ of each event. The offset of the resonance distribution in (a) is due to the calculation of the energy of all resonance events with the reconstructed invariant mass: $W^2 = m_{\Delta(1232)}^2$, with higher mass resonances not taken into account.

8.4 Event Selection in Reconstructed Quantities

The event selection for the studies discussed previously in this Chapter was applied in true energy, for all events with $E_{\text{true}} < 2 \text{ GeV}$. Firstly, in order to determine the optimal cut on reconstructed energy if this variable is considered alone, the efficiency, ϵ , and purity, p , of the selected sample are considered, in terms of the number and proportion of events with true energies below 2 GeV:

$$p = \frac{\# \text{ Selected Events, } E_{\text{true}} < 2 \text{ GeV}}{\# \text{ Total Selected Events}}, \quad \epsilon = \frac{\# \text{ Selected Events, } E_{\text{true}} < 2 \text{ GeV}}{\# \text{ Total } E_{\text{true}} < 2 \text{ GeV Events}} \quad (8.11)$$

Figure 8.16 shows the distributions of efficiency, purity and $\epsilon \times p$ for a range of reconstructed energy cuts. Also shown is the distribution of the fractional improvement, I , in the sensitivity at maximal mixing as defined by Equation 8.2, for oscillation fits performed applying each of the energy cuts. The cut on $\Delta\chi^2$, as show in Figure 8.13 for the true energy selection, is re-evaluated for each energy cut at the point where the quasi-elastic-like and resonance-like distributions intersect. The best improvement in sensitivity is obtained with a cut in the region of $E_{\text{reco}} < 1.25 \text{ GeV}$. This corresponds to the approximate region of the peak of the sample purity, and indicates the purity of the selected sample is more significant than the efficiency when selecting events with $E_{\text{true}} < 2 \text{ GeV}$. Below 1 GeV the improvement falls to zero as size of the selected sample and efficiency tends to zero and no events are selected. The maximum percentage improvement in the sensitivity at maximal mixing obtained by a cut in reconstructed energy alone is thus 1.8%.

The sensitivity may also be improved by the inclusion of combined cuts on χ^2 and $\Delta\chi^2$. Events with higher true energy are more likely to be more accurately reconstructed (with higher energy tracks and showers). They may therefore exhibit larger values of χ^2 ; consisting of a higher proportion of Deep Inelastic Scattering events, and RES events with higher mass resonances. Figure 8.17 shows the distributions of χ_{QE}^2 for all events with $E_{\text{reco}} < 1.25 \text{ GeV}$, split by true energy below and above 2 GeV. Some discrimination between event types may be possible in the region with $\chi^2 > 0.3$, $\Delta\chi^2 < -0.3$, where the distribution of events with $E_{\text{true}} < 2 \text{ GeV}$ is peaked. The distribution in χ_{RES}^2 showed no discrimination between events above and below the true energy threshold, and so was not considered here.

The position of the selection cut was defined as the point where the values of χ_{QE}^2

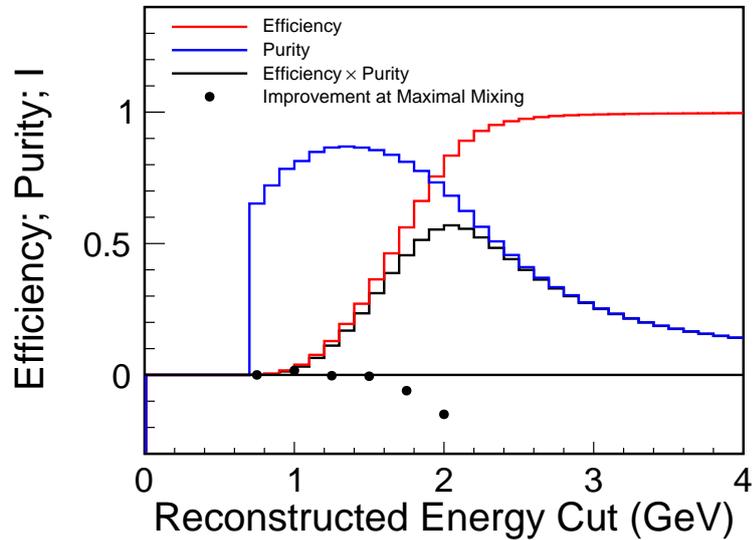
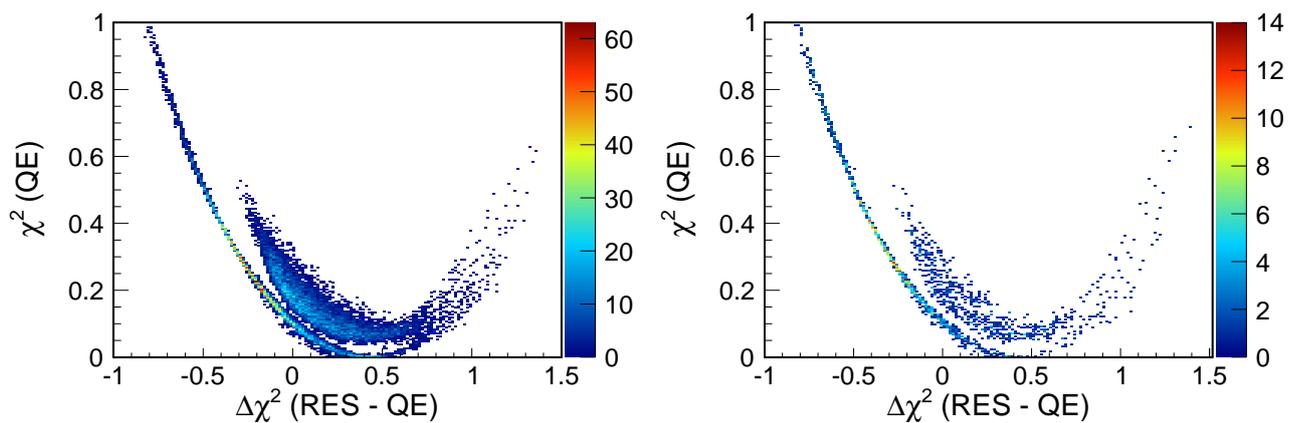


Figure 8.16: Efficiency (red), purity (blue) and efficiency \times purity (black) distributions for the range of reconstructed energy cuts (x -axis); for the selection of events with $E_{\text{true}} < 2$ GeV. The points show the fractional improvement in sensitivity at maximal mixing, I , for the oscillation fit, defined by Equation 8.2. The best sensitivity is obtained for a cut in the region of $E_{\text{reco}} < 1.25$ GeV, corresponding to the position of the peak of the sample purity. As the cut is increased to 2 GeV, the sensitivity worsens.



(a) Selected events with $E_{\text{true}} < 2$ GeV

(b) Selected events with $E_{\text{true}} > 2$ GeV

Figure 8.17: Distributions of χ^2_{QE} for all selected events with $E_{\text{reco}} < 1.25$ GeV, split by true energy (a) below and (b) above 2 GeV.

and $\Delta\chi^2$ maximised the purity of the selected sample, which as shown in Figure 8.16 correlates with an increased improvement in oscillation sensitivity at maximal mixing. A range of values for both variables was considered, and the maximal purity corresponds to a very loose cut on $\Delta\chi^2$, and a tighter cut on χ_{QE}^2 . The final selection in reconstructed variables is therefore defined as the combination of the following cuts:

Selection cut on energy:

- $E_{\text{reco}} < 1.25 \text{ GeV}$

Cuts to distinguish event type:

- Resonance-like: $\Delta\chi^2 < -0.134$
- Quasi-Elastic-like: $\Delta\chi^2 \geq -0.134$

Additional cuts:

- $\Delta\chi^2 > -0.994$
- $\chi_{QE}^2 < 0.31$

The projection onto $\sin^2(2\theta)$ of the log-likelihood surface for the final selected samples is given in Figure 8.18, compared to the oscillation fit performed in true energy and the standard MINOS reconstructed energy. The exclusion at maximal mixing increases from 1.8% when using the cut on energy only $E_{\text{reco}} < 1.25 \text{ GeV}$, to 2.0% when incorporating the cuts on χ^2 and $\Delta\chi^2$. A series of final cross checks were performed, to ensure a greater improvement on the exclusion at maximal mixing could not be made. The cut on χ_{QE}^2 was varied around the point of maximum purity and a series of oscillation fits were performed. All cases resulted in a worsened sensitivity compared to the cuts detailed above.

8.5 Summary

The improvement in the level of accuracy to which MINOS could exclude maximal mixing, for $\sin^2(2\theta) \neq 1.0$, is shown to be 2% when utilising event kinematics to correct low energy neutrino events. In comparison, the maximum improvement that could be

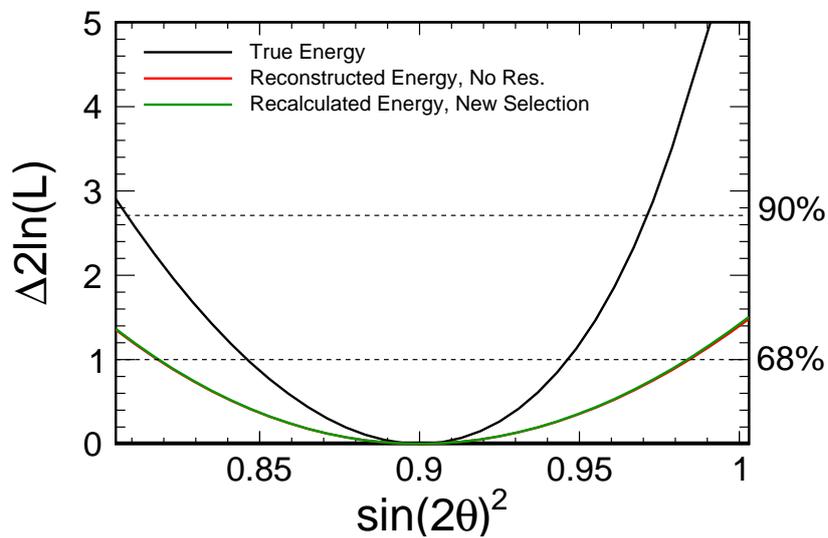


Figure 8.18: Projections onto $\sin^2(2\theta)$ over the statistical log-likelihood surfaces, for oscillation fits performed with spectra binned in true (black) and reconstructed energy (red), and energy recalculated using the new reconstruction method (green) detailed in this section. The underlying oscillation parameters were set to be $\sin^2(2\theta) = 0.9$ and $\Delta m^2 = 2.32 \times 10^{-3} \text{ eV}^2$. The $\Delta 2 \ln \mathcal{L}$ values corresponding to 68% confidence (1σ) and 90% confidence are shown by the dashed lines.

made using Monte Carlo truth information is 15%. The difficulty of identifying those events for which the reconstruction could make the most significant impact on sensitivity ($E_{\text{true}} < 2 \text{ GeV}$) illustrates the limitations of the MINOS detectors in accurately measuring events with very low true energies. Due to the combination of a minimal gain in sensitivity and a large increase in complexity, the alternative low energy reconstruction detailed here was not incorporated into the analysis detailed in this thesis.

The following Chapter summarises the results of the oscillation fits performed to the real data, including results published in [39]; incorporating the software detailed in Chapter 6 and the sensitivity improvement detailed in Chapter 7.

Chapter 9

Results

This Chapter presents the results of the analysis described in Chapters 5 to 7, published by the MINOS collaboration [39], with previous MINOS results introduced in Section 9.1. Work performed for this thesis is detailed in Sections 9.2 to 9.4, testing the performance of the event selection (Section 9.2) and cross checks of the `GhostFitter` software (Section 9.3), with the final results of the 2010 analysis as given by the `GhostFitter` software presented in Section 9.4.

9.1 Previous MINOS Oscillation Results

There have been two previous ν_μ disappearance analyses published by the MINOS experiment, in December 2006 after the culmination of Run I data taking [94], and in 2008 analysing the combined Run I + II dataset [69]. The exposure accumulated for the 2006 publication was 1.27×10^{20} PoT, and gave the results $\Delta m^2 = 2.74_{-0.26}^{+0.44} \times 10^{-3} \text{ eV}^2$, $\sin^2(2\theta) > 0.87$ at 68% confidence. The 2008 published analysis results were $\Delta m^2 = 2.43_{-0.13}^{+0.13} \times 10^{-3} \text{ eV}^2$, $\sin^2(2\theta) = 1.066$, with $\sin^2(2\theta) > 0.90$ at 90% confidence; for an exposure more than doubling the data published previously to 3.2×10^{20} PoT. The unphysical nature of the best fit in $\sin^2(2\theta)$ suppresses the $\sin^2(2\theta)$ contour in the physical region, producing limits smaller than might be expected. Fitting decay and decoherence models, each was excluded to 3.7σ and 5.7σ respectively (the exclusion of alternative models was discussed in Section 7.2.3).

The 2008 analysis utilised many of the same methods as the analysis described in this thesis; selecting events using the primary k NN PID selector (Section 5.2.3) and

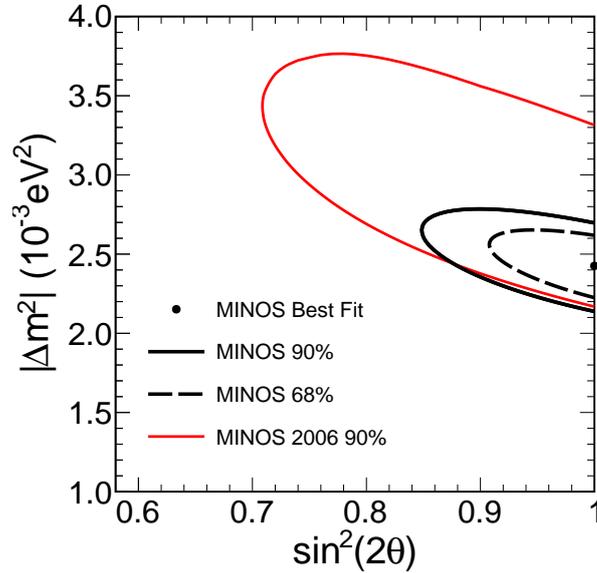


Figure 9.1: 90% and 68% confidence contours for the published 2006 and 2008 MINOS analyses, for exposures of 1.27×10^{20} PoT and 3.2×10^{20} PoT respectively.

extrapolating from Near to Far using the beam matrix method (Section 5.4). Three systematic uncertainties were fitted as nuisance parameters: the overall normalisation, the NC background and the absolute hadronic energy scale. The oscillation contours for the two previous analyses are shown in Figure 9.1.

A number of new approaches have been implemented for the analysis incorporating Run III. A new k NN method for calculating the shower energy, and an additional selector to aid alternative model discrimination are described in Chapter 5, and are the subject of theses in their own right [74, 78]. For this thesis, resolution binning has been implemented in the ν_μ fiducial sample (Chapter 7) and a fitting framework has been constructed to include a $\bar{\nu}_\mu$ sample and rock and anti-fiducial events, and fitting energy dependent systematic uncertainties extrapolated between the detectors (Chapter 6). Previous analyses considered systematic uncertainties in the Far Detector only.

9.2 Selected Events

The effects of the selection cuts detailed in Section 5.2 on the Far Detector data are shown in Table 9.1, which summarises the number of data events in each of the three samples as successive cuts are applied. The number of events selected in Runs I and II, already analysed and published in 2008, increases with the use of the new combined PID selector as low energy events are recovered. In total 62 events are added to the analysis sample, most with energies below 10 GeV.

Table 9.2 demonstrates how the selected CC ν_μ events are split by resolution bin, with a comparison made between the number of data events and the number of events expected assuming no oscillations in Table 9.3.

Selection Cut	Number of Events				
	Run I	Run II	Run I-pHE	Run III	Total
Raw Data	8886	8911	1616	13622	33035
Require a track	891	1401	331	2604	5227
<code>trkfitpass==successful</code>	886	1397	330	2595	5208
Fiducial Volume	433	694	178	1382	2687
Data Quality	428	682	172	1354	2636
Muon Direction	415	665	171	1306	2557
PID	318	511	129	1037	1995
Selected 2008 Events	282	448	118	-	848
Negative Curvature (ν_μ)	293	459	120	902	1774
Positive Curvature ($\bar{\nu}_\mu$)	25	52	9	135	221
Rock and Anti-Fiducial	357	555	128	977	2017

Table 9.1: Number of data events passing successive pre-selection and selection cuts, for the fiducial ν_μ and $\bar{\nu}_\mu$ samples. The Run I-pHE $\bar{\nu}_\mu$ sample containing 9 events is not included in the analysis or the total number of events to fit, being deemed insignificant before the data was inspected. The total number of events for each run in the RAF sample is also shown, as well as the total number of selected events by Run for the 2008 analysis [69].

Resolution Bin	Number of Events				
	Run I	Run II	Run I-pHE	Run III	Total
Res. Bin 0	63	106	25	168	362
Res. Bin 1	52	84	28	193	357
Res. Bin 2	59	87	25	150	321
Res. Bin 3	60	89	16	186	351
Res. Bin 4	59	93	26	205	383
Positive Curvature	25	52	-	135	212

Table 9.2: Number of selected data events in the ν_μ sample, split by resolution bin and compared to the total number of $\bar{\nu}_\mu$ events. The Run I-pHE $\bar{\nu}_\mu$ sample containing 9 events is not used in the analysis and so is not included in the total above.

Run Period	Observed (Far Detector)		Expected (No oscillations)		Expected (Oscillated)	
	Fiducial	RAF	Fiducial	RAF	Fiducial	RAF
Run I	318	357	426	375	335	329
Run II	511	555	639	565	503	497
Run I-pHE	120	128	134	136	126	130
Run III	1037	977	1252	1130	995	1001
Total	1986	2017	2451	2206	1959	1957

Table 9.3: Number of events observed and predicted in the absence of oscillations, for the fiducial and RAF classes. The final columns show the number of predicted events oscillated at the best fit point (fiducial + RAF fit, including all systematic uncertainties) of $\sin^2(2\theta) = 0.999$ and $\Delta m^2 = 2.317 \times 10^{-3} \text{ eV}^2$.

9.3 Cross Checks

Prior to the analysis being run on the full data set, the complete framework was finalised, checked and validated by the collaboration in a blessing package [95]. This ensures the analysis was performed in a “blind” fashion, unaltered after the dataset is examined to prevent unintentional biases. The final result was defined as a simultaneous fit to the ν_μ , $\bar{\nu}_\mu$ and RAF samples, with four systematic uncertainties fitted as nuisance parameters by the `GhostFitter` framework. The `NuSystFitter` used for the previous published analysis is unable to fit all systematic errors, but was used to validate the statistical fit and fits including the normalisation and neutral current systematic errors. The unblinding procedure was carried out in stages, the first being to re-analyse the 2008 published dataset of Runs I and II with and without analysis improvements.

9.3.1 Reproducing the 2008 Result

The first cross check performed with the `GhostFitter` software was to perform an oscillation fit utilising the calorimetric shower energy estimator, the selection and Monte Carlo used for the 2008 analysis, including the three systematic uncertainties that were considered as nuisance parameters: normalisation, neutral current background and relative shower energy. The best fit oscillation parameters and systematic shifts are given in Table 9.4 and the 90% sensitivity contour in Figure 9.2, compared to the fitting software used for the 2008 analysis, called the 2008 `Matrix Method`. The results obtained validated the `GhostFitter` software, and the next stage in the unblinding procedure, to introduce the k NN shower energy, the selection criteria and an updated version of the Monte Carlo was performed.

	$\Delta m^2 \times 10^{-3} \text{ eV}^2$	$\sin^2(2\theta)$	Norm.	NC	$E_{\text{shw}}^{\text{relative}}$
<code>GhostFitter</code>	2.43	1.00	+1.59%	-7.00%	-0.37σ
2008 <code>Matrix Method</code>	2.43	1.00	+1.61%	-8.16%	-0.36σ

Table 9.4: Best fits and nuisance parameters for oscillation fits to the 2008 data, fit using the `GhostFitter` software and the 2008 `Matrix Method` fitting software used for the published analysis. The fit includes the normalisation, neutral current and relative shower energy systematic errors, applied in the Far Detector and is constrained to the physical region in $\sin^2(2\theta)$.

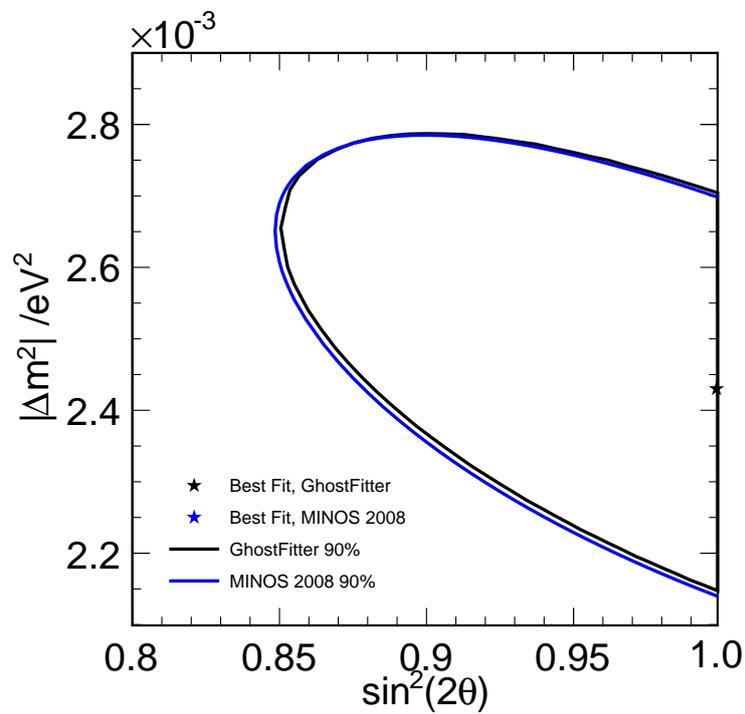


Figure 9.2: 90% confidence limits for the published 2008 analysis, compared to the best fit to the same data using the `GhostFitter` framework. The normalisation, neutral current and shower energy systematic errors were included as nuisance parameters. The 2008 default shower energy estimator, selection and Monte Carlo were used.

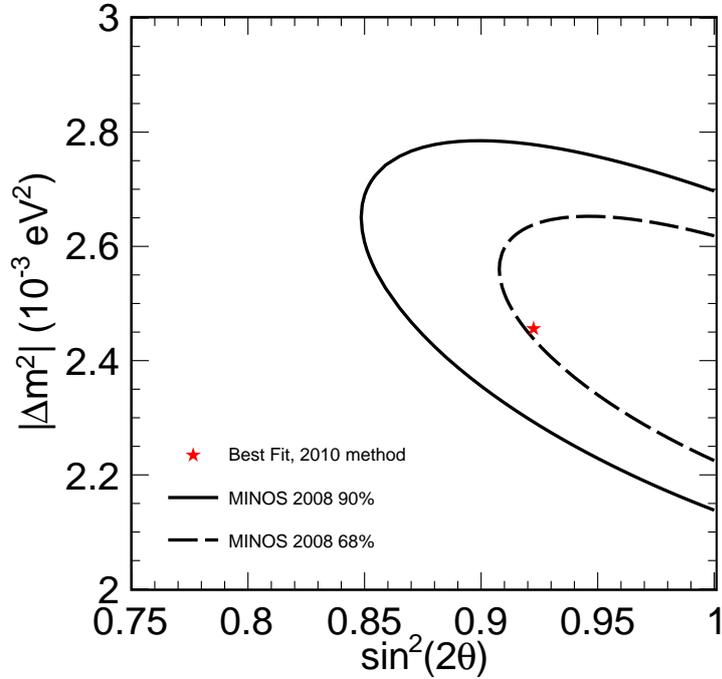


Figure 9.3: 90% and 68% confidence limits for the 2008 analysis, compared to the best fit to the same data with the new shower energy and event selection.

9.3.2 Fitting 2008 Data with Analysis Improvements

Including the new shower energy estimator and selection criteria are expected to change the results given in the previous section, as event energies change and more low energy NC events pass the selection. The $\bar{\nu}_\mu$ and RAF samples and resolution binning were not included to obtain a “2008-like” fit, which for validation was required by the working group to be required to lie within the 68% confidence limits of the previous result [69].

The result of the oscillation fit to the 2008 data set gave the oscillation parameters $\Delta m^2 = 2.456 \times 10^{-3} \text{ eV}^2$, $\sin^2(2\theta) = 0.923$, compared to the 2008 published contours in Figure 9.3. The best fit point moves away from maximal mixing, and lies within the 68% confidence limit as was required by the collaboration to proceed with the analysis.

9.3.3 Comparison of Fitting Software Results

As the third step in the unblinding procedure, a subset of the final results was validated by comparison to results from the `NuSystFitter`, for all samples and Runs, and the normalisation and neutral current systematic errors. Table 9.5 shows the deviations between the fits for various cases, in each case the difference is less than the statistical error. Exactly the same result cannot be obtained between the two methods due to differences in the way oscillations are implemented. The difference in sign between the systematic shifts is due to the fact the `NuSystFitter` applies systematic shifts to the data, where the `GhostFitter` applies them to the predicted Monte Carlo spectra.

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

Table 9.5: Comparing results from the GhostFitter and the NuSystFitter, for fiducial, RAF and combined fits for all Runs. No systematic uncertainties are included in the latter cases as the NuSystFitter cannot incorporate systematic uncertainties for the rock and anti-fiducial events.

9.4 Results

The results obtained to ν_μ only, fiducial ($\nu_\mu + \bar{\nu}_\mu$) and fiducial and RAF fits using the `GhostFitter` software are summarised in Table 9.6. For fits to fiducial events only, the binning schemes outlined in Section 6.1 result in 2300 energy bins, and when fitting for $(\sin^2(2\theta), \Delta m^2)$ 2298 degrees of freedom. Including RAF events increases the number of degrees of freedom in the fit to 2742, for four run periods and 111 energy bins.

The spectrum obtained when the fiducial sample is fully unblinded is shown in Figure 9.4, summed over all Runs and compared to the unoscillated prediction. A clear deficit in the number of expected neutrino events is observed in the region of the beam peak, shown in Table 9.3. The ratio of observed events to the unoscillated prediction is also shown in Figure 9.4, exhibiting a clear oscillation dip with the subsequent rise hinted at in the lowest two energy bins. The best fits to decay and decoherence models are also included, and discussed further in Section 9.4.2. Figure 9.5 shows the same spectrum split into resolution bins and separating out the $\bar{\nu}_\mu$ sample. The best fit oscillation parameters for the full fiducial sample including systematic uncertainties are $\Delta m^2 = 2.342 \times 10^{-3} \text{ eV}^2$ and $\sin^2(2\theta) = 1.001$, moving slightly into the unphysical region from maximal mixing.

Figure 9.6 shows the unblinded muon energy spectrum for the RAF sample, summed over all detector regions and compared to the unoscillated prediction. A deficit in the number of neutrino events is also seen here, though with much poorer resolution of the oscillation dip than is seen for fiducial events. The best fit oscillation parameters for the RAF samples, when the fit is constrained to the physical region ($\sin^2(2\theta) \leq 1.0$) in the fitting software, are $\Delta m^2 = 2.090 \times 10^{-3} \text{ eV}^2$ and $\sin^2(2\theta)$ at maximal mixing. The fit must be constrained when rock and anti-fiducial events are fitted alone, as the lack of sensitivity to $\sin^2(2\theta)$ seen in Monte Carlo studies permits highly unphysical values of $\sin^2(2\theta)$ when unconstrained. Figure 9.7 compares the 90% confidence contours for the fiducial and RAF fits run separately, showing the poor sensitivity to $\sin^2(2\theta)$ in the RAF event sample.

Performing the full fit to all samples and for all systematic uncertainties, with $\sin^2(2\theta)$ confined to the physical region, the resulting best fit point is:

$$|\Delta m^2| = 2.32_{-0.08}^{+0.12} \times 10^{-3} \text{ eV}^2 \quad (9.1)$$

$$\sin^2(2\theta) > 0.90 \text{ (90\% C.L.)} \quad (9.2)$$

with $\sin^2(2\theta) = 1.00_{-0.06}$. The quoted errors include statistical and systematic contributions, and are calculated from the one dimensional projections onto the $\sin^2(2\theta)$ and Δm^2 axes of the likelihood surface, where the error is defined by the values of the marginalised function at $\Delta \ln \mathcal{L} = 1.0$. Figure 9.8 shows the 90% and 68% confidence limits of the result compared to the Super-Kamiokande experiment, and Figure 9.9 compared to the previous MINOS published results described in Section 9.1.

Samples	Best Fit Values			Nuisance Parameters			
	$\Delta m^2 \times 10^{-3} \text{ eV}^2$ (10^{-3} eV^2)	$\sin^2(2\theta)$	$\Delta 2 \ln \mathcal{L} / \text{dof}$	Norm	NC	E_{trk}	E_{shw}
ν_μ	2.330	0.999	1828.1/1998	-	-	-	-
	2.346	0.999	1827.7/1998	+0.9%	-	-	-
	2.330	0.999	1828.1/1998	-	+0.4%	-	-
	2.322	0.999	1828.1/1998	-	-	+0.05 σ	-
	2.338	0.999	1826.3/1998	-	-	-	-0.28 σ
	2.350	0.999	1825.5/1998	+0.9%	-10%	+0.19 σ	-0.30 σ
Fiducial	2.330	0.995	2119.4/2298	-	-	-	-
	2.342	0.997	2119.1/2298	+0.8%	-	-	-
	2.330	0.997	2119.3/2298	-	+6%	-	-
	2.332	0.995	2119.4/2298	-	-	-0.01 σ	-
	2.344	0.995	2117.5/2298	-	-	-	-0.28 σ
	2.342	1.001	2117/2298	+0.8%	+5%	+0.13 σ	-0.29 σ
Fiducial + RAF (Unconstrained) (Constrained)	2.298	0.995	2633.3/2742	-	-	-	-
	2.310	0.995	2633.0/2742	+0.7%	-	-	-
	2.298	0.997	2633.0/2742	-	+6%	-	-
	2.302	0.993	2633.2/2742	-	-	-0.07 σ	-
	2.310	0.995	2630.9/2742	-	-	-	-0.29 σ
	2.314	1.001	2630.5/2742	+0.6%	+5%	+0.08 σ	-0.30 σ
2.317	0.999	2630.5/2742	+0.6%	+5%	+0.08 σ	-0.30 σ	

Table 9.6: Best fits and nuisance parameters for combinations of samples, fit with the `GhostFitter` software. Different combinations of systematic errors are included in the fitting software and summarised. The $\sin^2(2\theta)$ parameter is permitted to enter the unphysical region in all the fits given above.

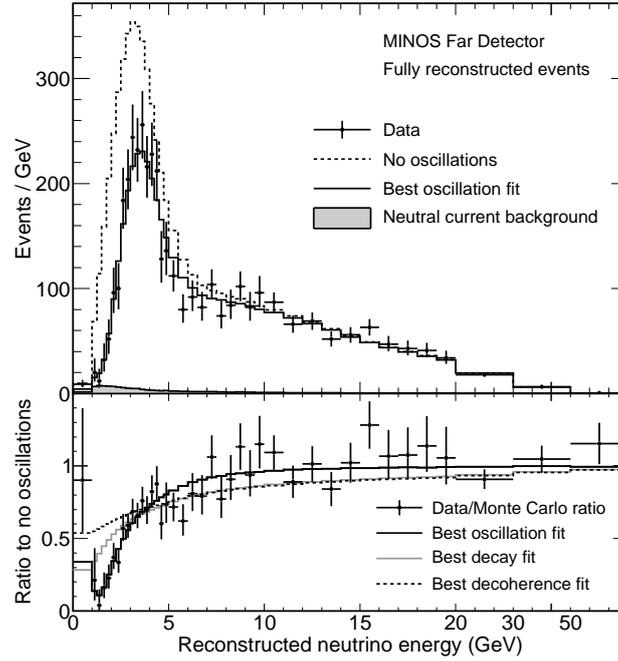


Figure 9.4: Spectra for the fiducial events, summed over all Runs. The points show the observed data, and the black line the prediction at the best fit parameters. The dashed line shows the prediction in the absence of oscillations, and the grey shaded area the neutral current background. The ratio of the background subtracted best fit to the unoscillated prediction is shown below, including the best fit to alternative models, discussed in Section 9.4.2.

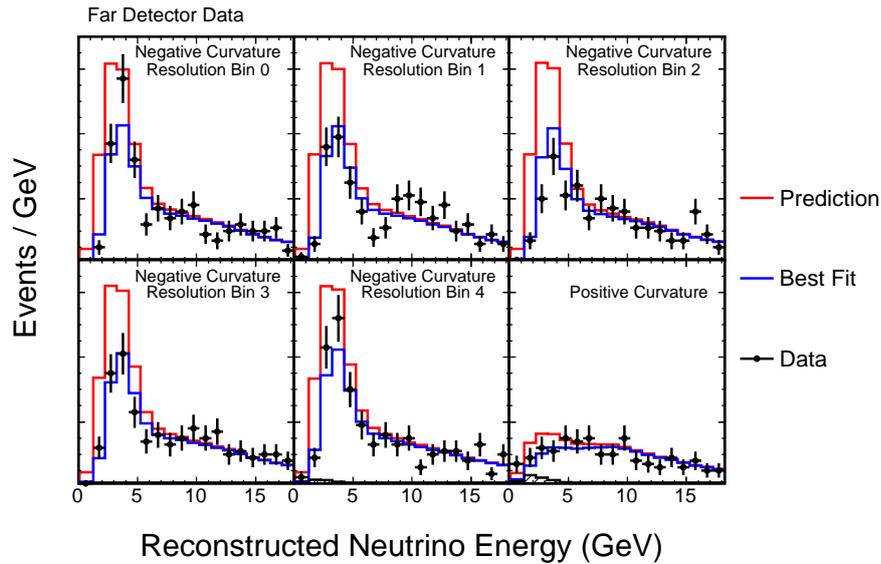


Figure 9.5: Spectra for the fiducial events, divided by resolution bin and the $\bar{\nu}_\mu$ sample. Black points show the data, compared to the unoscillated (red) and best fit (blue) energy spectra in each resolution bin.

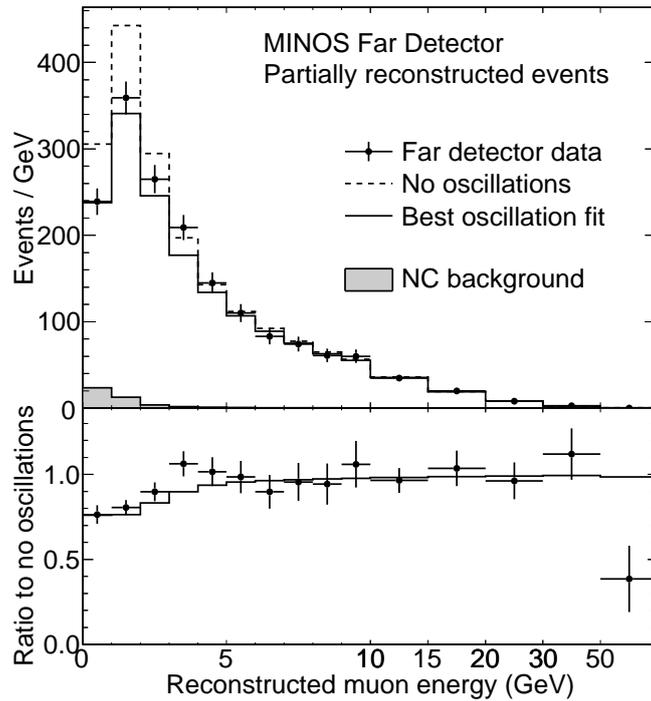


Figure 9.6: Spectra for the rock and anti-fiducial events, summed over all Runs and detector regions. The points show the observed data, and the black line the prediction at the best fit parameters for a full fiducial and RAF fit. The dashed line shows the prediction in the absence of oscillations, and the grey shaded area the neutral current background. The ratio of the background subtracted best fit to the unoscillated prediction is shown below.

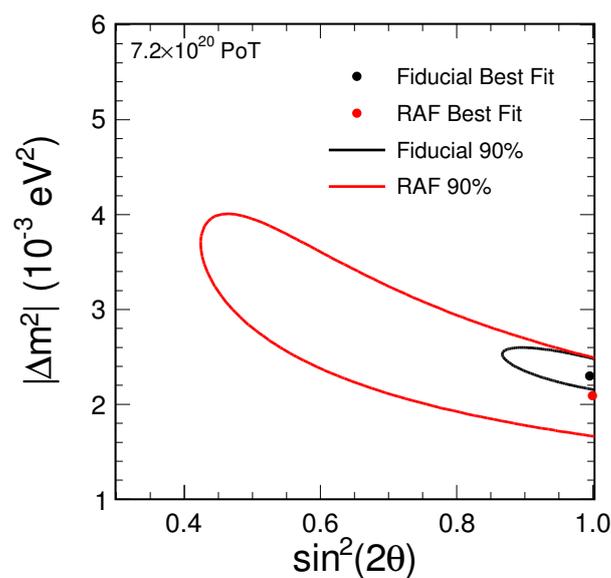


Figure 9.7: 90% confidence limits for statistical fits to the fiducial and RAF samples separately. The final result combines these samples into a single fit.

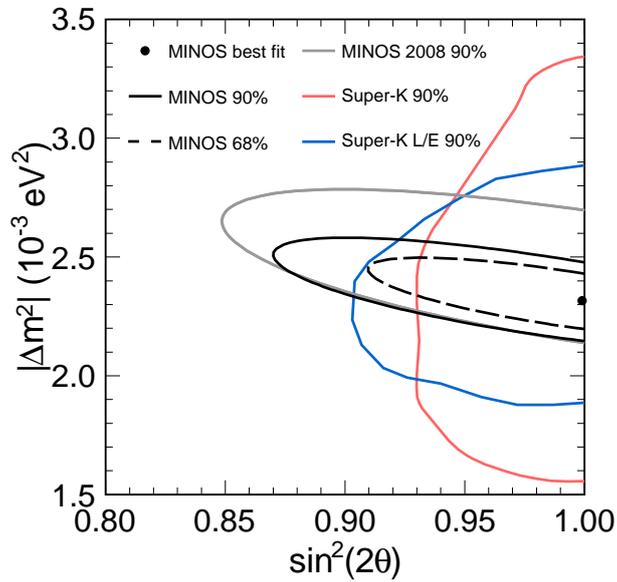


Figure 9.8: Best fit point and sensitivity contours for the ν_μ analysis (black), fit to all Runs and samples including systematic uncertainties. Also shown are results from the Super-Kamiokande 2004 L/E analysis (blue) [96] and 2005 zenith angle analysis (red) [37], in addition to the MINOS 2008 analysis [69].

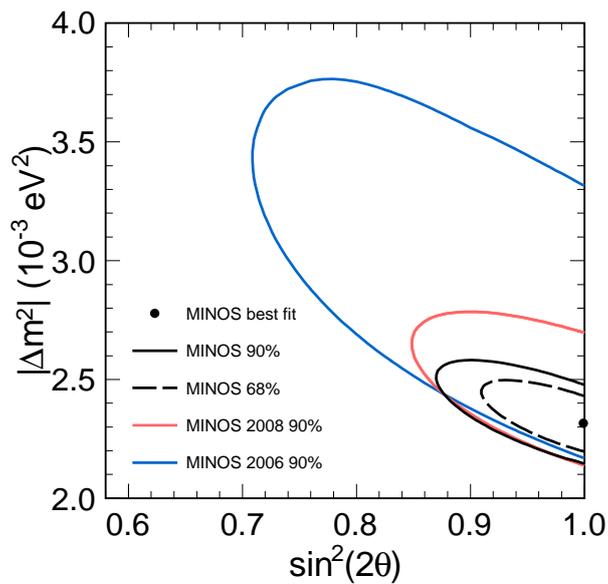


Figure 9.9: 90% confidence limits for the previous MINOS analyses published in 2006 [94] and 2008 [69], compared to the analysis presented here.

9.4.1 Cross Checks Performed Post-Unblinding

9.4.1.1 Fitting Systematic Uncertainties

The smoothness and absence of discontinuities in the values of the nuisance parameters across the oscillation parameter space serves as a useful cross check on the implementation of the systematic uncertainties in the fitting software. This cross check was performed with Far Detector Monte Carlo in Section 6.2.2, and is repeated here when fitting the unblinded data. The nuisance parameter surfaces in oscillation space are shown in Figure 9.10. No discontinuities are visible, and none of the systematic errors at the best fit point move more than 1σ from zero.

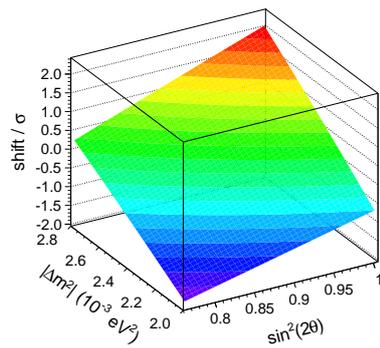
9.4.1.2 Goodness of Fit

In order to assess the statistical likelihood of the result presented here and to check for pathological problems, the $\ln \mathcal{L}$ value obtained from the fit to the data is compared to a large number of Poisson fluctuated fake datasets. This gives a measure of the degree to which the data collected has been subject to statistical fluctuations.

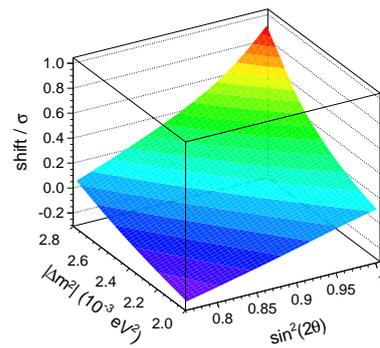
To perform this check, a Far Detector fake data set was selected, not including systematic uncertainties, scaled to the data exposure, and oscillated at the best fit parameters from real data. Each bin weight was then randomly Poisson fluctuated about the bin content, to produce a fake dataset with the same underlying true oscillation parameters smeared by statistical fluctuation, and a fit as performed with the `GhostFitter`. These steps were repeated 10,000 times, plotting the $\ln \mathcal{L}$ value for each fit.

Figure 9.11(a) shows the distribution of $\ln \mathcal{L}$ values for the fit to fiducial events. The $\ln \mathcal{L}$ for the fit to data is also marked, with $\ln \mathcal{L}/\text{dof} = 2119.4/2298$ better than 69% of the fluctuated fits. The fit to fiducial events therefore benefits from statistical fluctuations which reduce the $\ln \mathcal{L}$ value at the best fit point, which serves to reduce the size of the confidence contours compared to those that would be expected from the mean $\ln \mathcal{L}$ value.

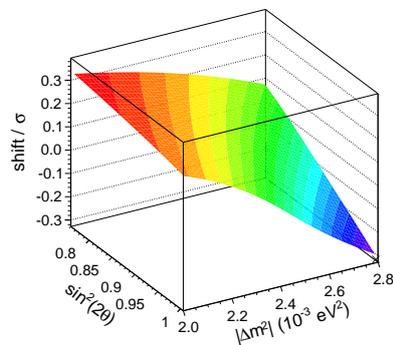
When the RAF sample is included in Figure 9.11(b), the $\ln \mathcal{L}/\text{dof} = 2633.3/2742$ is better than only 34% of fits, indicating the RAF sample incorporates more “unfavourable” statistical fluctuations, in the context of an increased confidence contour size compared to that which would result from the mean $\ln \mathcal{L}$ value.



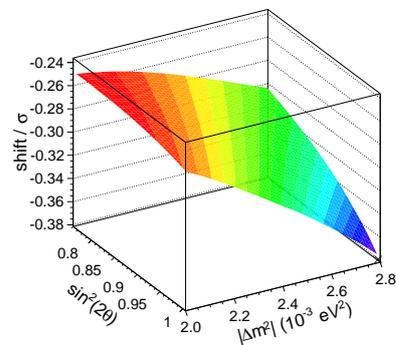
(a) Normalisation



(b) Neutral Current

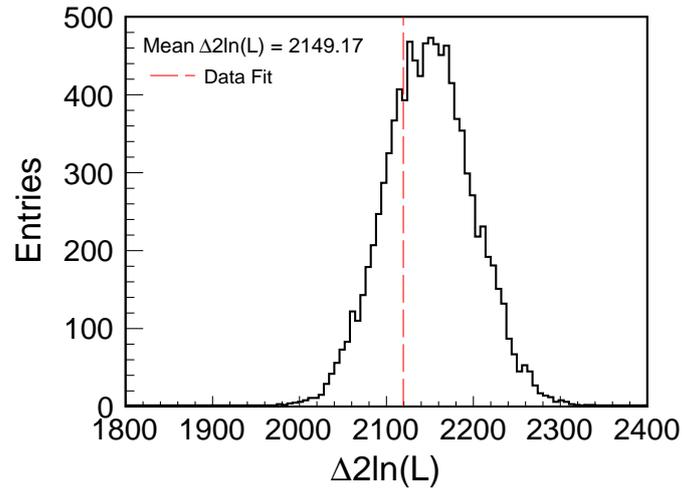


(c) Track Energy

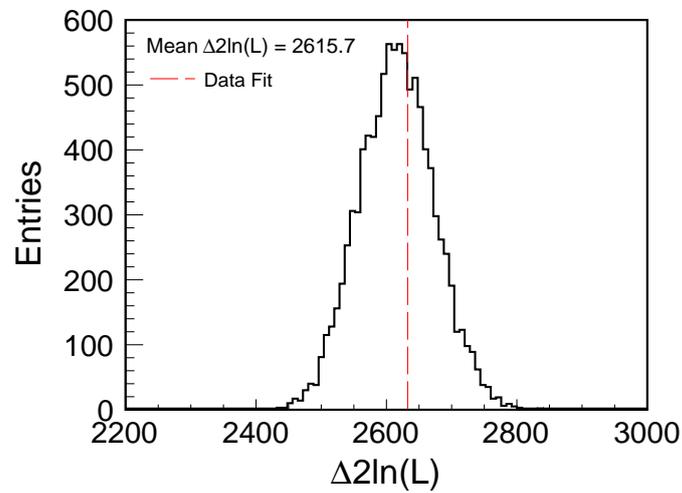


(d) Relative Shower Energy

Figure 9.10: Best fit nuisance parameter values in oscillation space, for the (a) normalisation, (b) neutral current, (c) track energy and (d) relative shower energy systematic errors. The fit includes all samples and Runs, with all systematic errors fit simultaneously.



(a)



(b)

Figure 9.11: Distributions of $\ln \mathcal{L}$ for 10,000 fake experiments for fits to the fiducial (a) and fiducial + RAF sample (b). The corresponding $\ln \mathcal{L}$ for the fits to data are shown by the red line in each case.

9.4.2 Alternative Models

In addition to the oscillation model, the best fit parameters for neutrino decay and decoherence were also obtained. The models are described in Chapter 2, and the mechanism of fitting them in Chapter 6. The models were fitted using the same samples and the four systematic uncertainties as the main analysis detailed in the previous sections, for parameters $(\sin^2(2\theta), \alpha)$ in the decay case and $(\sin^2(2\theta), \mu^2)$ for decoherence.

The best fits to the decay and decoherence model parameters are given in Table 9.7, with the 90% confidence limits for the statistical oscillation fit and fits including systematic errors are given in Figures 9.12 and 9.13. Ratios of the decay and decoherence best fit spectra to the unoscillated spectrum are also shown in Figure 9.4, showing a worse agreement with data at low energies than the oscillation hypothesis.

Using the definition of model discrimination given by Equation 7.10, the fiducial only fit to neutrino decay is disfavoured by 6.5σ from the oscillation case as defined by Equation 7.10, and the fiducial and RAF fit is disfavoured by 7.8σ . The fit to the decoherence hypothesis results in an exclusion of 8.8σ and 9.7σ for the fiducial and fiducial and RAF fits respectively.

	Oscillation		Decoherence		Decay	
	Fiducial	Fiducial+RAF	Fiducial	Fiducial+RAF	Fiducial	Fiducial+RAF
$\Delta m^2/10^{-3} \text{ eV}^2$	2.346	2.317	2.18	1.85	2.40	2.08
$\sin^2(2\theta)$	0.999	0.999	0.999	0.999	0.997	0.997
$\ln \mathcal{L}$	2116.95	2630.5	2194.94	2724.32	2159.46	2691.26
Exclusion	-	-	8.83σ	9.69σ	6.52σ	7.79σ

Table 9.7: Comparison of best fit parameters for alternative models of disappearance, for fits to the samples specified and all systematic uncertainties. The exclusion of each model is calculated as $\sqrt{\Delta \ln \mathcal{L}}$, with respect to the difference in $\ln \mathcal{L}$ relative to the oscillation hypothesis.

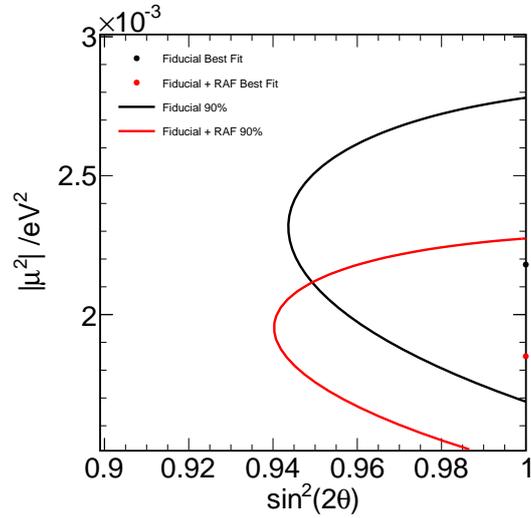


Figure 9.12: 90% confidence limits and best fit points for the fiducial only and fiducial + RAF fits to pure neutrino decoherence, including all four systematic uncertainties as nuisance parameters.

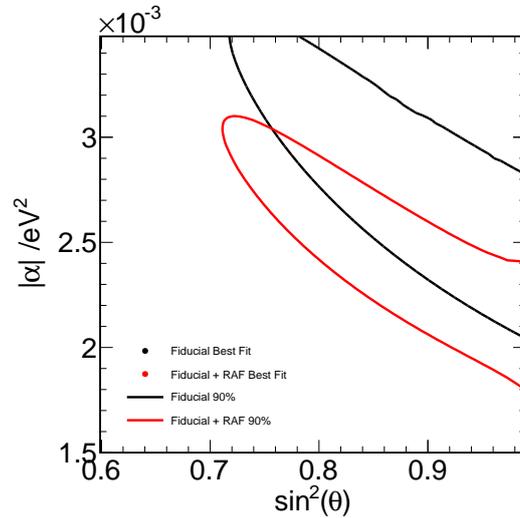


Figure 9.13: 90% confidence limits and best fit points for the fiducial only and fiducial + RAF fits to pure neutrino decay, including all four systematic uncertainties as nuisance parameters.

Chapter 10

Conclusion

An analysis of Charged Current ν_μ disappearance in the NuMI beam has been performed, using the Near and Far Detectors which comprise the MINOS Experiment. The experiment has been running reliably for a period of more than 6 years, during which time the energy spectra and rates of neutrino interactions have been recorded at distances 1 km and 735 km along the NuMI beamline. The Near Detector provides information on the initial beam composition and flux, and the Far Detector energy spectrum is compared to that expected in the absence of neutrino oscillations to extract a measurement of the atmospheric mixing angle and mass splitting.

This thesis focused on improvements made for the MINOS 2010 published result [39], an analysis which used data collected between May 2005 and June 2008. The beam, detectors and their monitoring were described in Chapters 3 and 4, and after data quality cuts provided a sample for analysis corresponding to 7.25×10^{20} Protons-on-Target (PoT).

Chapter 5 gave an overview of the analysis procedure, how samples of events are selected and how the Near Detector energy spectrum is used to constrain the Far Detector prediction. Two new analysis improvements were also summarised; a new energy estimator for the hadronic shower energy, and the inclusion of events from outside the detector fiducial region.

Chapter 6 and 7 described in detail two additional improvements for this analysis round, developed for this thesis. The former detailed a new fitting method, separate from the extrapolation procedure designed to incorporate multiple event samples, in addition to any number of energy based and scaling systematic across a subset or all of these samples.

The latter described a technique of estimating and dividing CC ν_μ events by energy resolution $\sigma_{E_\nu} = \sigma_{E_\mu} + \sigma_{\text{shw}}$. A significant improvement in sensitivity is obtained by the use of resolution information, representing a 6% improvement in sensitivity to $\sin^2(2\theta)$ and an 8% improvement in Δm^2 for five bins of resolution.

A possible improvement to the sensitivity to maximal mixing for a non-maximal value of $\sin^2(2\theta)$ was investigated in Chapter 8, developing a new method for reconstructing the neutrino energy of very low energy events utilising the event kinematics. The improvement in the exclusion of maximal mixing was found to be 2% when combined with a new low energy selection.

The results of the ν_μ disappearance analysis are given in Chapter 9. Extrapolating from the Near to the Far Detector, 4657 ± 68 events¹ were expected in the absence of oscillations, and 4003 were observed, thus the event rate alone strongly disfavours a no disappearance hypothesis. The best fit point and confidence contours from the dataset made up of data taken during Runs I-III are consistent with previous MINOS analyses, and provide the current world's most precise measurement of Δm_{atm}^2 .

The measured values of the atmospheric mixing parameters are given by $|\Delta m^2| = 2.32_{-0.08}^{+0.12} \times 10^{-3} \text{ eV}^2$ and $\sin^2(2\theta) > 0.90$ (90% C.L.), with the best fit $\sin^2(2\theta) = 1.00_{-0.06}$. The alternative disappearance models of neutrino decoherence and decay are disfavoured by 8.8σ and 6.5σ respectively for the fiducial only events, and increase to 9.7σ and 7.8σ when the Rock and Anti-Fiducial event sample is included in the oscillation fit.

The MINOS Experiment has also recently published an oscillation analysis of anti-neutrino data [97], taken between May and March 2010 as Run IV with an anti-neutrino beam. An unanticipated discrepancy between the neutrino and anti-neutrino oscillation parameters was presented, with best fit oscillation parameters $|\Delta \bar{m}^2| = 3.36_{-0.40}^{+0.46}(\text{stat.}) \pm 0.06(\text{syst.}) \times 10^{-3} \text{ eV}^2$ and $\sin^2(2\bar{\theta}) = 0.86_{-0.12}^{+0.11}(\text{stat.}) \pm 0.01(\text{syst.})$, and confidence limits shown in Figure 10.1. The probability the underlying ν_μ and $\bar{\nu}_\mu$ parameters were the same based on these results was 2%. A future publication based on Runs IV and VII, doubling the anti-neutrino dataset (see Table 10.1), will investigate this discrepancy further, and is expected to show a movement towards the neutrino oscillation parameters if the effect is due to a statistical fluctuation.

¹Statistical error on event count, not including systematic effects.

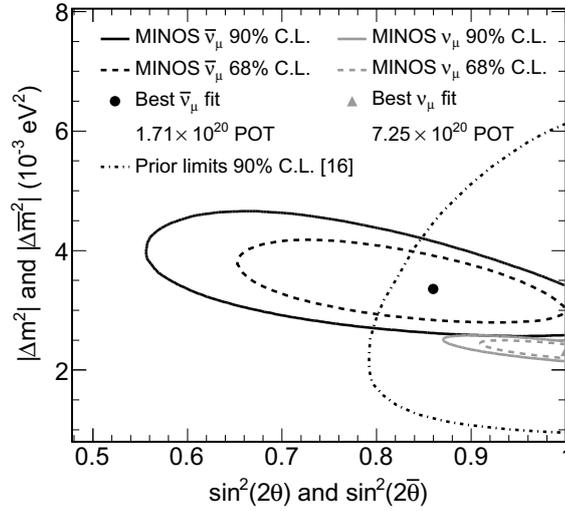


Figure 10.1: 90% (solid) and 68% (dashed) confidence contours for the published 2010 $\bar{\nu}_\mu$ MINOS analysis. The 2010 ν_μ analysis contours are also shown for comparison (grey). The dotted line shows the results of a global fit without MINOS data [98].

	Nature of beam	Start Date	End Date	Far Det. (Run/Subrun)		Far Det. PoT / 10^{20}
				Start	End	
Run IV	$\bar{\nu}_\mu$	29/09/2009	22/03/2010	44820/2	45480/12	1.71
Run V	ν_μ	22/03/2010	05/05/2010	45490/0	45702/9	0.46
Run VI	ν_μ	12/05/2010	12/07/2010	45783/0	46265/18	0.62
Run VII	$\bar{\nu}_\mu$	01/11/2010	18/03/2011	46887/1	47512/9	1.24
Run VIII	ν_μ	22/06/2011	11/07/2011	48005/19	48074/3	≈ 0.12
Run IX	$\bar{\nu}_\mu$	30/07/2011	15/09/2011	48174/0	48422/9	0.4

Table 10.1: Summary of Far Detector datasets taken after the thesis analysis, including future running.

10.1 Future Developments

Three parameters describing our current understanding of neutrino oscillations remain undetermined; the nature of the mass hierarchy, the size of the \mathcal{CP} violating phase δ and the mixing angle θ_{13} . Three reactor experiments currently under construction, Double-CHOOZ [47], Reno [48] and Daya Bay [49] aim to measure θ_{13} via $\bar{\nu}_e$ oscillations.

Two long baseline neutrino beam experiments, No ν A [41] and T2K [40] also have an

improved measurement of θ_{13} as their preliminary goal. The former is currently under construction, with a Far Detector sited in Northern Minnesota in an off-axis position to the NuMI beam. By positioning the detector away from the beamline, the flux is reduced but the neutrino spectrum is narrower and peaked at lower energies, providing an improved sensitivity to oscillations. The latter, an off-axis two detector experiment based in Japan has recently published preliminary limits on θ_{13} [45], observing $\nu_\mu \rightarrow \nu_e$ transitions to 2.5σ significance.

Both experiments also expect to make improved measurements of the atmospheric parameters, for both muon neutrinos and anti-neutrinos. With an extended period of $\bar{\nu}_\mu$ running in addition to ν_μ running, and an accurate measurement of $\bar{\nu}_\mu$ disappearance and $\bar{\nu}_e$ appearance, they may have some sensitivity to δ and the mass hierarchy. The proposed LBNE Experiment is intended to make precision measurements of interactions and investigate the value of δ , using a new highly intense neutrino beam from Fermilab to the Homestake mine, South Dakota and a 100 kt scale water-Čerenkov or Liquid-Argon detector.

The MINOS Experiment has also proposed to extend its running into the No ν A era, as the MINOS+ Experiment. The medium energy NuMI beam that will be delivered to No ν A will provide an intense neutrino flux to the MINOS Far Detector, peaking in the 4-10 GeV range. This will provide a unique sensitivity to non-standard models, the existence of sterile neutrinos and measurements of neutrino lifetime as well as improving the atmospheric parameter measurement.

The wide range of future beam and reactor experiments, coupled with a variety of double β -decay experiments attempting to improve sensitivity to the absolute mass scale and the tantalising brand new results from OPERA [99], the neutrino sector looks to be an exciting and rewarding field for decades to come.

Bibliography

- [1] W. Pauli, A letter to the Tübingen Conference on radioactivity, 1930.
- [2] C. L. Cowen, F. Reines, F. B. Harrison, H. W. Kruse, and A. D. McGuire, Detection of the free neutrino: A Confirmation, *Science* **124**, 103–104 (1956).
- [3] S. L. Glashow, Partial-symmetries of weak interaction, *Nucl. Phys.* **22**(4), 579–588 (1961).
- [4] S. Weinberg, A model of leptons, *Phys. Rev. Lett.* **19**(21), 1264–1266 (1967).
- [5] A. Salam and J. C. Ward, Electromagnetic and weak interactions, *Phys. Lett.* **13**, 168–171 (1964).
- [6] G. Danby *et al.*, Observation of high-energy neutrino reactions and the existence of two kinds of neutrinos, *Phys. Rev. Lett.* **9**, 36–44 (1962).
- [7] The DONUT Collaboration, First direct observation of the tau-neutrino, *Europhys. News* **32**, 56–57 (2001).
- [8] The LEP Collaborations, A combination of preliminary electroweak measurements and constraints on the standard model, (2005).
- [9] Z. Maki, M. Nakagawa, and S. Sakata, Remarks on the Unified Model of Elementary Particles, *Progress of Theoretical Physics* **28**(5), 870–880 (1962).
- [10] B. Pontecorvo, Neutrino experiments and the question of leptonic-charge conservation, *Sov. Phys. JETP* **26**, 984–988 (1968).
- [11] G. M. Tinti, *Sterile neutrino oscillations in MINOS and hadron production in pC collisions*, PhD thesis, University of Oxford, 2010.
- [12] J. J. Evans, *Measuring Antineutrino Oscillations with the MINOS Experiment*, PhD thesis, University of Oxford, 2009.

- [13] E. Lisi, A. Marrone, and D. Montanino, Probing Possible Decoherence Effects in Atmospheric Neutrino Oscillations, *Phys. Rev. Lett.* **85**(6), 1166–1169 (2000).
- [14] C. I Chang, W. Dai, X. Li, Y. Liu, F. Ma, and Z. Tao, Possible effects of quantum mechanics violation induced by certain quantum-gravity effects on neutrino oscillations, *Phys. Rev. D* **60**(3), 033006 (1999).
- [15] V. Barger, J. G. Learned, S. Pakvasa, and T. J. Weiler, Neutrino Decay as an Explanation of Atmospheric Neutrino Observations, *Phys. Rev. Lett.* **82**(13), 2640–2643 (1999).
- [16] J. N. Bahcall, A. M. Serenelli, and S. Basu, New Solar Opacities, Abundances, Helioseismology, and Neutrino Fluxes, *The Astrophysical Journal Letters* **621**(1), L85 (2005).
- [17] K. Nakamura and Particle Data Group, Review of Particle Physics, *Journal of Physics G: Nuclear and Particle Physics* **37**(7A), 075021 (2010).
- [18] B. T. Cleveland, T. Daily, R. Davis, Jr., J. R. Distel, K. Lande, C. K. Lee, P. S. Wildenhain, and J. Ullman, Measurement of the Solar Electron Neutrino Flux with the Homestake Chlorine Detector, *The Astrophysical Journal Letters* **496**, 505–+ (1998).
- [19] J. N. Bahcall, N. A. Bahcall, and G. Shaviv, Present Status of the Theoretical Predictions for the ^{37}Cl Solar-Neutrino Experiment, *Phys. Rev. Lett.* **20**(21), 1209–1212 (1968).
- [20] J. N. Abdurashitov *et al.*, Measurement of the solar neutrino capture rate by the Russian-American gallium solar neutrino experiment during one half of the 22-year cycle of solar activity, *J. Exp. Theor. Phys.* **95**, 181–193 (2002).
- [21] W. Hampel *et al.*, GALLEX solar neutrino observations: Results for GALLEX IV, *Phys. Lett.* **B447**, 127–133 (1999).
- [22] M. Altmann *et al.*, Complete results for five years of GNO solar neutrino observations, *Phys. Lett.* **B616**, 174–190 (2005).
- [23] J. N. Bahcall, S. Basu, and M. H. Pinsonneault, How uncertain are solar neutrino predictions?, *Physics Letters B* **433**(1-2), 1 – 8 (1998).
- [24] J. N. Bahcall, M. H. Pinsonneault, S. Basu, and J. Christensen-Dalsgaard, Are Standard Solar Models Reliable?, *Phys. Rev. Lett.* **78**(2), 171–174 (1997).

- [25] J. Boger *et al.*, The Sudbury Neutrino Observatory, *Nucl. Instrum. Meth.* **A449**, 172–207 (2000).
- [26] B. Aharmim *et al.*, Electron energy spectra, fluxes, and day-night asymmetries of ^8B solar neutrinos from measurements with NaCl dissolved in the heavy-water detector at the Sudbury Neutrino Observatory, *Phys. Rev. C* **72**(5), 055502 (2005).
- [27] L. Wolfenstein, Neutrino oscillations in matter, *Phys. Rev. D* **17**(9), 2369–2374 (1978).
- [28] S. P. Mikheev and A. Yu. Smirnov, Resonance enhancement of oscillations in matter and solar neutrino spectroscopy, *Sov. J. Nucl. Phys.* **42**, 913–917 (1985).
- [29] B. Aharmim *et al.*, An Independent Measurement of the Total Active ^8B Solar Neutrino Flux Using an Array of ^3He Proportional Counters at the Sudbury Neutrino Observatory, *Phys. Rev. Lett.* **101**, 111301 (2008).
- [30] S. Abe *et al.*, Precision Measurement of Neutrino Oscillation Parameters with KamLAND, *Phys. Rev. Lett.* **100**, 221803 (2008).
- [31] T. K. Gaisser and M. Honda, Flux of atmospheric neutrinos, *Ann. Rev. Nucl. Part. Sci.* **52**, 153–199 (2001).
- [32] J. Marshall, *A study of muon neutrino disappearance with the MINOS detectors and the NuMI neutrino beam*, PhD thesis, University of Cambridge, 2008.
- [33] R. Becker-Szendy *et al.*, A Search for muon-neutrino oscillations with the IMB detector, *Phys. Rev. Lett.* **69**, 1010–1013 (1992).
- [34] K. S. Hirata *et al.*, Observation of a small atmospheric ν_μ/ν_e ratio in Kamiokande, *Physics Letters B* **280**(1-2), 146 – 152 (1992).
- [35] W. Anthony Mann, Atmospheric neutrinos and the oscillations bonanza, in *The Proceedings of the 19th International Symposium on Lepton and Photon Interactions at High-Energies (LP 99)*, 1999.
- [36] Y. Fukuda *et al.*, The Super-Kamiokande detector, *Nucl. Instrum. Meth.* **A501**, 418–462 (2003).
- [37] Y. Ashie *et al.*, A Measurement of Atmospheric Neutrino Oscillation Parameters by Super-Kamiokande I, *Phys. Rev.* **D71**, 112005 (2005).

-
- [38] M. H. Ahn *et al.*, Measurement of Neutrino Oscillation by the K2K Experiment, *Phys. Rev.* **D74**, 072003 (2006).
- [39] P. Adamson *et al.*, Measurement of the neutrino mass splitting and flavor mixing by MINOS, *Phys. Rev. Lett.* **106**, 181801 (2011).
- [40] The T2K Collaboration, The T2K Experiment, *Nucl. Instrum. Meth.* **In Press**, **Uncorrected Proof** (2011).
- [41] The NO ν A Collaboration, NO ν A: Proposal to build a 30 kiloton off-axis detector to study $\nu_\mu \rightarrow \nu_e$ oscillations in the NuMI beamline, *hep-ex/0503053*, 112 (2004).
- [42] C. Giganti, New results from the T2K experiment, in *The Proceedings of the European Physical Society Conference on High Energy Physics (EPS HEP 2011) (In Preperation)*, 2011.
- [43] M. Apollonio *et al.*, Search for neutrino oscillations on a long base-line at the CHOOZ nuclear power station, *Eur. Phys. J.* **C27**, 331–374 (2003).
- [44] P. Adamson *et al.*, Improved search for muon-neutrino to electron-neutrino oscillations in MINOS, *Submitted to Phys. Rev. Lett.* (2011).
- [45] M. Apollonio *et al.*, Indication of Electron Neutrino Appearance from an Accelerator-Produced Off-Axis Muon Neutrino Beam, *Phys. Rev. Lett* **107**(4), 041801 (2011).
- [46] M. Apollonio *et al.*, Limits on neutrino oscillations from the CHOOZ experiment, *Physics Letters B* **466**(2-4), 415 – 430 (1999).
- [47] T. Lasserre and M. Goodman, Double Chooz, A Search for the Neutrino Mixing Angle θ_{13} , *hep-ex/0606025* (2006).
- [48] J. K. Ahn *et al.*, RENO: An Experiment for Neutrino Oscillation Parameter θ_{13} Using Reactor Neutrinos at Yonggwang, *arXiv:1003.1391v1* (2010).
- [49] The Daya Bay Collaboration, A Precision Measurement of the Neutrino Mixing Angle θ_{13} using Reactor Antineutrinos at Daya Bay, *hep-ex/0701029* (2007).
- [50] P. Adamson *et al.*, Study of muon neutrino disappearance using the Fermilab Main Injector neutrino beam, *Phys. Rev. D* **77**(7), 072002 (2008).
- [51] K. Anderson *et al.*, Technical Report, Fermi National Accelerator Laboratory, (2002).

-
- [52] V. Garkusha *et al.*, Design Study of the NuMI Low Energy Target for Higher Power beam, *MINOS-doc-2778* (2007).
- [53] A. Blake, *A study of atmospheric neutrino oscillations in the MINOS far detector*, PhD thesis, University of Cambridge, 2005.
- [54] A. Weber, Recent Results from the MINOS Experiment, *MINOS-doc-7876* (2010).
- [55] A. Ferrari, A. Sala, P.R. and Fasso, and J. Ranft, FLUKA: A multi-particle transport code, *CERN-2005-010* (2005).
- [56] S. Agostinelli *et al.*, Geant4 - a simulation toolkit, *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment* **506**(3), 250 – 303 (2003).
- [57] D. G. Michael *et al.*, The magnetised steel and scintillator calorimeters of the MINOS experiment, *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment* **596**(2), 190 – 228 (2008).
- [58] C. L. F. Howcroft, *Atmospheric Neutrinos in the MINOS Far Detector*, PhD thesis, University of Cambridge, 2004.
- [59] A. Cabrera, *Systematic Comparison of the MINOS Near and Far Detector Readout Systems*, PhD thesis, University of Oxford, 2005.
- [60] The MINOS Calibration Group, 2009 Position Paper on Calibration of Runs I-II-III, *MINOS-doc-6717* (2009).
- [61] M. Campanella, A. Ferrari, P. R. Sala, and S. Vanini, First Calorimeter Simulation with the FLUGG Prototype, *Technical Report ATL-SOFT-99-004 CERN Geneva* (1999).
- [62] A. Himmel, The NuMI Beam Simulation with Flugg, *MINOS-doc-6316* (2010).
- [63] H. Gallagher, The NEUGEN neutrino event generator, *Nucl. Phys. Proc. Suppl.* **112**, 188 (2002).
- [64] R. Brun *et al.*, GEANT Detector Description and Simulation Tool, *CERN Program Library Long Writeup W0513* (1993).

-
- [65] C. Zeitnitz and T.A. Gabriel, The GEANT-CALOR interface and benchmark calculations of ZEUS test calorimeters, *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment* **349**(1), 106 – 111 (1994).
- [66] M. Thomson, Alternative DeMuxer, *MINOS-doc-2535* (2006).
- [67] D. Groom, N. Mokhov, and S. Striganov, Muon stopping power and range tables 10 MeV - 100 TeV, *Atomic Data and Nuclear Data Tables* **78**(2), 183–356 (2001).
- [68] P. Adamson *et al.*, Improved search for muon-neutrino to electron-neutrino oscillations in MINOS, *Submitted to PRL* (2011).
- [69] P. Adamson *et al.*, Measurement of Neutrino Oscillations with the MINOS Detectors in the NuMI Beam, *Phys. Rev. Lett.* **101**(13), 131802 (2008).
- [70] P. Pawloski, Results from the MINOS Experiment, *MINOS-doc-7867* (2011).
- [71] T. Cundiff *et al.*, The MINOS Near Detector Front End Electronics, *IEEE Transactions on Nuclear Science* **53**, 1347–1355 (2006).
- [72] D. Petyt, ND Data Quality position paper for the Summer 2008 CC analysis, *MINOS-doc-4289* (2008).
- [73] A. Blake, Far Detector Data Quality Software, *MINOS-doc-3486* (2007).
- [74] C. Backhouse, *Measuring neutrino oscillation parameters using ν_μ disappearance in MINOS*, PhD thesis, University of Oxford, 2011.
- [75] M. Strait, *Measurement of Neutrino Oscillation Parameters Using Anti-fiducial Charged Current Events in MINOS*, PhD thesis, University of Minnesota, 2010.
- [76] S. Kopp and L. Loiacono, Beam Quality Cuts, *MINOS-doc-2747* (2007).
- [77] R. Ospanov, *A Measurement of muon neutrino disappearance with the MINOS detectors and NuMI Beam*, PhD thesis, University of Texas at Austin, 2008.
- [78] J. Ratchford, PhD thesis, University of Texas at Austin, 2011.
- [79] D. Petyt, PID sensitivity study status, *MINOS-doc-2743* (2007).
- [80] C. Backhouse, Estimating shower energies using a kNN based method, *MINOS-doc-6868* (2010).

-
- [81] C. Backhouse, Why we can't use NC events in the oscillation fit, *MINOS-doc-7138* (2010).
- [82] M. Strait, Results with Energy Only RAF Binning, *MINOS-doc-7484* (2010).
- [83] F. James and M. Roos, MINUIT - A system for function minimisation and analysis of the parameter error and correlations, *Computational Physics Communications* **10**, 353 (1975).
- [84] J. Evans, Systematic uncertainties for the 2010 CC analysis, *MINOS-doc-7145* (2010).
- [85] S. J. Coleman, *A measurement of neutrino oscillations with muon neutrinos in the MINOS experiment*, PhD thesis, College of William and Mary, 2011.
- [86] R. Hatcher *et al.*, Range Curvature Task Force Position Paper, *MINOS-doc-3134* (2007).
- [87] P.A. Rodrigues, *A sterile neutrino search for the MINOS experiment*, PhD thesis, University of Oxford, 2010.
- [88] J. Ratchford, Position Paper: MRCC motivated NC background, *MINOS-doc-7115* (2010).
- [89] M. Dorman, Beam Fit Position Paper, *MINOS-doc-7146* (2010).
- [90] J. Musser, Swimmer Energy Loss, *MINOS-doc-2081* (2006).
- [91] R. Wigmans, *Calorimetry, Energy Measurement in Particle Physics*, Oxford Science Publications, Oxford University Press, 2000.
- [92] S. Coleman, Position Paper on resolution binning and $\nu_\mu + \bar{\nu}_\mu$ extrapolation for the 7.2×10^{20} PoT CC analysis, *MINOS-doc-6874* (2010).
- [93] M. Dorman, *Cross section measurements for quasi-elastic neutrino-nuclear scattering with the MINOS Near Detector*, PhD thesis, University College London, 2008.
- [94] B.J. Rebel, First MINOS Results with the NuMI Beam, *Nuclear Physics B - Proceedings Supplements* **168**, 195 – 198 (2007).
- [95] J. Evans *et al.*, Charged Current Analysis Blessing Package, *MINOS-doc-7217* (2010).

-
- [96] Y. Ashie *et al.*, Evidence for an oscillatory signature in atmospheric neutrino oscillations, *Phys. Rev. Lett* **93**, 101801 (2004).
- [97] P. Adamson *et al.*, First direct observation of muon antineutrino disappearance, *Phys. Rev. Lett.* **107**(2), 021801 (2011).
- [98] M. C. Gonzalez-Garcia and M. Maltoni, Phenomenology with massive neutrinos, *Physics Reports* **460**(1-3), 1–129 (2008).
- [99] T. Adam *et al.*, Measurement of the neutrino velocity with the OPERA detector in the CGNS beam, (2011).

List of Figures

2.1	Neutral Current and Charged Current interactions	5
2.2	Permitted neutrino mass hierarchies	11
2.3	ν_μ survival probability as a function of neutrino energy	12
2.4	Solar neutrino fluxes and energy spectra as predicted by the SSM	17
2.5	SNO phase II flux results	20
2.6	Allowed regions for KamLAND and combined KamLAND and solar oscillation parameters	22
2.7	Ratio of KamLAND electron anti-neutrino spectrum to unoscillated prediction	22
2.8	Production of atmospheric neutrinos in the atmosphere	23
2.9	Summary of atmospheric neutrino double ratio measurements	24
2.10	Zenith angle distributions for selected classes of Super-Kamiokande events	26
2.11	Allowed oscillation parameters for $\nu_\mu \leftrightarrow \nu_\tau$ transitions from Super-K data.	27
2.12	K2K reconstructed neutrino energy spectrum	28
2.13	Allowed oscillation parameters for $\nu_\mu \leftrightarrow \nu_\tau$ oscillations from K2K data.	28
2.14	CHOOZ reconstructed neutrino energy spectrum	30
2.15	Allowed parameter region for θ_{13} and Δm_{23}^2 from CHOOZ data.	30
2.16	Confidence limits for $2 \sin^2(\theta_{23}) \sin^2(2\theta_{13})$ as a function of δ from MINOS	31
2.17	Confidence limits for $\sin^2(2\theta_{13})$ as a function of δ from T2K	32

3.1	The MINOS neutrino beamline	35
3.2	Schematic of the NuMI beam facility	36
3.3	Schematic of the NuMI target	37
3.4	Magnetic horn focusing system for the NuMI beam	37
3.5	Expected rate of ν_μ interactions in the Near Detector for different target/horn configurations	38
3.6	steel-scintillator plane orientations in the MINOS detectors	41
3.7	Cross section through a scintillator strip	41
3.8	Optical readout for the detectors	42
3.9	The Far Detector	43
3.10	The Far Detector coordinate system	44
3.11	The MINOS Calibration Procedure	46
3.12	Stopping power for muons in the Far Detector	50
3.13	Protons-on-Target as a function of time	52
4.1	The Near Detector	55
4.2	Near Detector plane configurations	56
4.3	Responses of the QIE electronics	58
4.4	Data flow through the Near Detector electronics	59
4.5	Distributions of low rate MENUs in the Near Detector by subrun	64
4.6	Distributions of low rate MINDERS in the Near Detector by subrun	64
4.7	Distribution of malfunctioning electronics Crates in the Near Detector	65
4.8	Distributions of CAPID errors in the Near Detector by subrun	66
4.9	Distributions of PARITY errors in the Near Detector by subrun	67
4.10	Distributions of MISCOUNT errors in the Near Detector by subrun	67
4.11	Distributions of high average ADC counts in the Near Detector by subrun	68

4.12	Beam spill timing distributions in the Near Detector	69
4.13	Distribution of PoT taken between March 2005 to February 2006	71
4.14	Distribution of PoT taken between June 2006 to July 2007	72
4.15	Distribution of PoT taken between November 2007 to June 2008	72
4.16	Distribution of PoT taken between July 2008 to June 2009	73
5.1	Fiducial volumes in the Near and Far Detectors	77
5.2	Distribution of the PID selection variable for data and MC in the Near Detector	79
5.3	Distribution of the low energy optimised PID selection variable for data and MC in the Near Detector	81
5.4	Efficiency and contamination of the CC selection in the Far Detector	81
5.5	Comparison of oscillation sensitivities with the standard shower energy reconstruction and improved energy resolution	83
5.6	Shower energy distributions for different methods of determination	84
5.7	Distributions of calorimetric and k NN shower energy resolutions in bins of 500 MeV.	85
5.8	Bias from the uncorrected k NN shower energy and polynomial correction	85
5.9	Neutrino parent decays in the NuMI beam pipe	87
5.10	Relationship between neutrino energy spectra at the Near and Far Detectors.	87
5.11	Beam Matrix for neutrinos from Run III Monte Carlo	88
5.12	Flowchart summarising the beam matrix extrapolation procedure	90
5.13	Distribution of the true vertices of events in the rock and anti-fiducial sample	90
5.14	Diagram of the anti-fiducial regions used for the RAF analysis	92
6.1	Examples of Monte Carlo template histograms for the ν_μ sample	95

6.2	Schematic of the C++ sample and vessels container within the fitting software	96
6.3	Sensitivity confidence intervals for different energy spectra binning in the oscillation fit	97
6.4	Ratio of oscillated to unoscillated MC events in true vs reconstructed energy	99
6.5	Shifts in the Monte Carlo best fit point induced by systematic errors . .	103
6.6	Sensitivity confidence intervals for individual and combined normalisation systematics	105
6.7	Interpolation weights for input histograms to generate systematically shifted templates	107
6.8	Effect of changing systematics simultaneously and summing separate residuals on predicted neutrino energy spectra	109
6.9	Best fit nuisance parameter surface and distribution of log-likelihood against track energy systematic error	110
6.10	Best fit nuisance parameter surface using non-zero initial values in MINUIT111	
6.11	Ratios of oscillated true CC ν_μ events in Run III Far Detector Monte Carlo to the unoscillated true energy spectrum for different fitting software	113
6.12	90% Sensitivity confidence interval for different methods of incorporating oscillations	113
6.13	Distribution of best fit points for 100 MDC fake experiments	115
6.14	Sensitivity confidence intervals for the MDC	115
6.15	Sensitivity confidence intervals for individual and combined Runs for fiducial events	116
6.16	Sensitivity confidence intervals for individual and combined Runs for rock and anti-fiducial events	117
6.17	Sensitivity confidence intervals combining the $\bar{\nu}_\mu$, and RAF events with the ν_μ sample	117

6.18	Sensitivity confidence intervals for the statistical fit and the fit including the four main systematics	118
6.19	Marginalisation of the log-likelihood surface with and without systematics	118
6.20	Sensitivity confidence intervals for decay and decoherence models of neutrino disappearance	120
7.1	Correlation between $E_{\text{reco}} - E_{\text{true}}$ and σ_{E_ν} for events with muon momentum measured from track curvature	123
7.2	Energy resolution parameterisations for shower energy and muon momentum from range	124
7.3	Energy resolution parameterisation for muon momentum measured from curvature	124
7.4	Shower and track energy fractional resolutions	126
7.5	The fractional resolution σ_{E_ν}/E_ν showing division into quantiles	127
7.6	Sensitivity confidence intervals comparing different parameterisations of shower energy resolution	129
7.7	Efficiency in quantiles of resolution	130
7.8	Reconstructed vs true neutrino energy in the best and worst quantiles of resolution	132
7.9	Sensitivity confidence intervals in bins of resolution	133
7.10	Projections of the log-likelihood surface for one, two and ten bins of resolution	134
7.11	Ratios of oscillated to unoscillated neutrino spectra for best and worst resolution bins	135
7.12	Shower energy distributions split by resolution bin	135
7.13	90% Confidence limits for fake data fits in each individual bin of resolution, with and without $E_{\text{shw}} = 0$ GeV events	136
7.14	Best fit points for events with a smeared reconstructed shower or track energy	138

7.15	Best fit points for events with a biased reconstructed shower or track energy	139
7.16	Best fit points for an increased proportion of low-y events	140
7.17	The shower fractional resolution $\sigma_{E_{\text{shw}}}/E_\nu$ and total fractional resolution σ_{E_ν}/E_ν for neutral current events	142
7.18	Best fit points for events with a scaled up NC background	143
7.19	MINOS sensitivity to oscillation parameters with successive analysis improvements	145
8.1	Projection of statistical log-likelihood surfaces onto $\sin^2(2\theta)$, for oscillation fits in true and reconstructed energy	147
8.2	Oscillation dip for oscillation fits in true and reconstructed energy	147
8.3	Reconstructed vs true neutrino energy distribution	148
8.4	Distribution of true event types in the oscillation dip region in reconstructed neutrino energy	149
8.5	$\Delta 2 \ln \mathcal{L}$ and projection onto $\sin^2(2\theta)$ over the likelihood surface for oscillation fits using true and reconstructed energy	151
8.6	Projections onto $\sin^2(2\theta)$ of statistical log-likelihood surfaces for oscillation fits with kinematically corrected energy reconstruction	152
8.7	Reconstructed shower energy distribution	153
8.8	Invariant mass distribution	153
8.9	True shower energy distribution for events with $E_{\text{shw}}^{\text{reco}} = 0$	155
8.10	True and reconstructed invariant mass W^2 after preliminary energy correction	155
8.11	Correlation between $\cos(\theta)_{\text{reco}} - \cos(\theta)_{\text{true}}$ and $\sigma_{\cos(\theta)}$	158
8.12	Muon angle resolution parameterisations	159
8.13	$\Delta\chi^2 = \chi_{RES}^2 - \chi_{QE}^2$ distributions	160
8.14	Projections onto $\sin^2(2\theta)$ of statistical log-likelihood surfaces for oscillation fits in true, reconstructed and recalculated energy	161

8.15	Ratio of $E_{\text{reco}} - E_{\text{true}}$ to E_{true} for the standard MINOS reconstruction and the corrected energy reconstructions	162
8.16	Distributions of efficiency, purity and the improvement in sensitivity at maximal mixing for reconstructed energy selections	164
8.17	Distributions of χ_{QE}^2 for all selected events with $E_{\text{reco}} < 1.25$ GeV, split by true energy below and above 2 GeV	164
8.18	Projections onto $\sin^2(2\theta)$ of statistical log-likelihood surfaces for oscillation fits in true, reconstructed and recalculated energy	166
9.1	Confidence contours for the published 2006 and 2008 MINOS analyses	169
9.2	Confidence limits for the 2008 analysis, compared to the best fit from the GhostFitter software	173
9.3	Confidence limits for the 2008 analysis, compared to the best fit with the new shower energy and event selection	174
9.4	Data, best fit and unoscillated spectra for the fiducial events	179
9.5	Data, best fit and unoscillated spectra for the fiducial events divided by resolution bin	179
9.6	Data, best fit and unoscillated spectra for the rock and anti-fiducial events	180
9.7	90% confidence limits for the 2010 analysis, for the fiducial and RAF samples	180
9.8	Confidence contours for the 2010 analysis, compared to other experiments	181
9.9	Confidence contours for the published 2006, 2008 and 2010 MINOS analyses	181
9.10	Nuisance parameter surfaces in oscillation space	183
9.11	Distribution of $\ln \mathcal{L}$ for 10,000 fake experiments for fits to the fiducial and fiducial + RAF samples	184
9.12	Confidence contours for fiducial and fiducial and RAF fits to the decoherence model	187
9.13	Confidence contours for fiducial and fiducial and RAF fits to the decay model	187

10.1 Confidence contours for the published 2010 $\bar{\nu}_\mu$ analysis 190

List of Tables

2.1	Solar neutrino producing reactions as predicted by the SSM	17
3.1	Summary of the Far Detector analysis dataset	53
5.1	Summary of pre-selection and selection cuts applied to Monte Carlo and data samples	82
6.1	Details of event samples and vessels	94
6.2	Shifts in the best fit point due to the dominant systematic errors	104
6.3	Average time per grid point to run the fitting software, by interpolation method between systematics templates	108
6.4	Mock Data Challenge best fit points	114
7.1	Shower energy resolution parameterisations	128
7.2	Event numbers by quantile for differently binned σ/E distributions	130
7.3	Neutral Current event counts in bins of resolution	141
7.4	Improvements in alternative model discrimination	144
9.1	Number of data events passing successive selection cuts	170
9.2	Number of selected data events, split by resolution bin	171
9.3	Number of events observed and predicted in the absence of oscillations	171
9.4	Best fits and nuisance parameters for oscillation fits to the 2008 data	172

9.5	Comparing results from the two fitting methods	176
9.6	Best fits and nuisance parameters for combinations of samples	178
9.7	Comparison of best fit parameters for alternative models of disappearance	186
10.1	Summary of Far Detector datasets taken after the thesis analysis, and future running	190