



A University of Sussex DPhil thesis

Available online via Sussex Research Online:

<http://sro.sussex.ac.uk/>

This thesis is protected by copyright which belongs to the author.

This thesis cannot be reproduced or quoted extensively from without first obtaining permission in writing from the Author

The content must not be changed in any way or sold commercially in any format or medium without the formal permission of the Author

When referring to this work, full bibliographic details including the author, title, awarding institution and date of the thesis must be given

Please visit Sussex Research Online for more information and further details

Muon Antineutrino Disappearance in the MINOS Experiment

NICHOLAS DEVENISH

Submitted for the degree of Doctor of Philosophy

May 2015

US
University of Sussex

Abstract

Antineutrinos from the set of 7.09×10^{20} Protons-on-target (POT) of ν_μ -dominant and 3.40×10^{20} POT of $\bar{\nu}_\mu$ -enhanced MINOS beam data is analysed to extract the antineutrino atmospheric-scale oscillation parameters $|\Delta\bar{m}^2|$ and $\sin^2(2\bar{\theta})$. Using the Feldman-Cousins statistical technique to account for systematic and statistical errors, $|\Delta\bar{m}^2| = 2.58^{+0.26}_{-0.17} \times 10^{-3} \text{ eV}^2$ and $\sin^2(2\bar{\theta}) = 0.96^{+0.04}_{-0.10}$ were measured. This measurement of $|\Delta\bar{m}^2|$ improves on the previous world limits from Super-K by a factor of five. Any observed difference between the neutrino and antineutrino oscillation parameters could be an indication of CPT violation, or otherwise physics beyond the standard model. This measurement is consistent with the global limits on $|\Delta m^2|$ and $\sin^2(2\theta)$ for neutrinos.

For Charlotte,
and for Sky,
Whom I both love,

And without whom,
I would not be here today.

Preface

The work in this thesis arose in participation with the antineutrino oscillation working group in the MINOS collaboration. As such, the thesis contains a mixture of original work, and also builds on work performed by, or in collaboration with, others. Much of the work developed for this thesis was used in the publications [1, 2].

Chapter 1 is the introduction, and contains a brief summary of motivations for the work carried out in the course of this thesis.

Chapter 2 is a summary of the theory and history of experimental neutrino physics, and so contains no original workings, save several diagrams constructed or reproduced to illustrate the discussion.

Chapter 3 contains summaries of the NuMI beamline and MINOS experiment, which are not original work. The chapter also contains a description of recently published results from the MINOS collaboration, which I was involved with. In particular, I performed much of the original work on which this thesis is based for the $\bar{\nu}_\mu$ oscillation result.

Chapter 4 describes the extrapolation method used to generate predictions, which was the work of others. However, I performed much optimisation of the procedure, needed for the computationally intensive Feldman-Cousins analysis, and performed the work on improving methods of calculating the oscillation probabilities, described in section 4.2.5.

Chapter 5 describes the selection used to identify charged-current $\bar{\nu}_\mu$ interactions in the detector. The development of the PID and selection variables was inherited, and not original work, but I was heavily involved in derivation and optimisation of the combination of the

variables used for the MINOS analysis selections, and the further work performed to optimise the selection used in this thesis is entirely my own. I performed the analysis of the effect of systematics on the results presented here, using the scales of the systematics quantified in studies largely performed by others.

Chapter 6 describes the Feldman-Cousins statistical method, which is not my own original concept, but the implementations, analysis, fitting approaches and method for integrating systematics into the method is entirely my own work.

Chapter 7 considers the results of the analysis, and so is entirely based on work performed for this thesis, save the data used in the comparisons with world experimental limits, which come from their respective experiments.

Chapter 8 is the conclusion, and so is entirely based on work undertaken in this thesis.

Acknowledgements

The list of people who have helped me to accomplish this thesis is long and varied - and should probably involve almost everybody I have interacted with during the process. Here is a small subset.

Foremost, I would like to thank my supervisor Jeff Hartnell for his endless patience, encouragement and support throughout the (somewhat more turbulent than expected) process of writing this thesis. His professionalism and persistent faith should be a model to all, and one which I can only hope to emulate in the future. Thanks also to the staff, postdocs, and other students at Sussex throughout - particularly Lisa Falk, Phil Harris, Marta Tavera and Gwenaelle Lefeuvre, who were there when I needed assistance - whether for friendship, physics, or just for help navigating the institutional bureaucracy.

Everyone on the MINOS experiment deserves thanks, but in particular the antineutrino working group heads Justin Evans and Donna Naples, and Jeff de Jong from the calibration working group. Additionally, Jess Mitchell, Phil Rodrigues, Chris Backhouse, Zeynep Isvan and Ruth Toner were all fellow students whom helped keep me sane, especially during long collaboration meetings, or when trying in vain to wrestle the ROOT framework into submission.

I would like to thank my parents for always supporting my love of science and computing, and the extraordinarily generous financial support throughout my entire university career.

I would like to thank Charlotte, who became my wife shortly after submitting, for supporting me in every way possible throughout the ups and downs of the entire process, and having faith in me when it felt like everything was grinding to a halt.

Finally, I would like to thank Angela Miller, Charlotte's mother, who passed away in 2014, but whose legacy allowed me to work to completion. This is a debt I can only ever hope to repay, by being a loving husband to Charlotte.

Contents

1	Introduction	1
2	History and Theory	4
2.1	Neutrino Mixing Theory	5
2.1.1	PMNS Mixing Matrix	9
2.1.2	Two-Flavour approximation	10
2.1.3	CPT Transformation and Conservation	11
2.2	The Solar Neutrino Problem	12
2.3	The Atmospheric Neutrino Anomaly	17
2.3.1	Super-Kamiokande	19
2.4	Recent Measurements	22
2.5	Unresolved Properties of the Neutrino	22
3	The MINOS Experiment	24
3.1	The NuMI Neutrino Beam	24
3.2	Detectors	28
3.2.1	Scintillator	28
3.2.2	Near Detector	29
3.2.3	Far Detector	32
3.2.4	Magnetisation	33
3.3	Calibration	35
3.3.1	The Light Injection System	35
3.3.2	Muon tracks	37
3.4	Monte Carlo Simulation	38
3.4.1	Flux Tuning	39
3.5	Event Reconstruction	40
3.6	Previous MINOS Results	42
4	Extrapolation	47
4.1	Predicting the Far Detector Energy Spectrum	47
4.2	Beam Matrix Method	49
4.2.1	Step-by-step	50
4.2.2	Beam Matrix	54
4.2.3	Converting Far Detector Flux to an Energy Spectrum	54
4.2.4	Purity Corrections: Background calculation	57
4.2.5	Oscillation	63

5	Selection	67
5.1	Preselection	68
5.2	Signal Selection	71
5.2.1	NC Discrimination	71
5.2.2	Charge Sign Selection	73
5.3	Choosing the selection variables	76
5.3.1	Antineutrino Beam Selection	78
5.4	Systematic Uncertainties	80
5.5	Actual effect of systematics	85
5.6	Data Validation	88
6	Application of the Feldman-Cousins Method	92
6.1	Statistical Corrections	95
6.1.1	Applying the correction	96
6.1.2	Fitting individual experiments	98
6.1.3	Aggregating Multiple Runs	101
6.2	Systematic Corrections	102
6.2.1	Experiment Generation	105
6.2.2	Far Detector Fake Experiment Generation	107
6.2.3	Systematic Shifts and Far Detector Generation	109
6.2.4	Systematic grids	111
7	Results	114
7.1	Near Detector	114
7.2	Far Detector	115
7.3	Feldman-Cousins Corrected Results	120
7.4	Measuring Individual Parameters	122
7.5	Comparisons	126
8	Conclusion	129

List of Figures

2.1	Two-flavour oscillation survival probability	11
2.2	Global fit from 2008 on antineutrino vs neutrino oscillation	13
2.3	Standard solar model prediction of neutrino energy from the sun	14
2.4	Schematic of Sudbury Neutrino Observatory, and the SNO detector	16
2.5	Measured flux of ν_e vs $\nu_\mu + \nu_\tau$ by the SNO detector	17
2.6	Neutrino production by cosmic rays. Based on a diagram in [9].	18
2.7	The Super-Kamiokande Detector	19
2.8	The Super-K antineutrino results	21
3.1	The NuMI facility	24
3.2	Stages in the target hall	25
3.3	Neutrino flux with different target hall configurations	26
3.4	The structure of a scintillator strip	30
3.5	A short scintillator strip illuminated by a blue LED	30
3.6	Photograph and diagram of the near detector	31
3.7	Diagram and Photograph of the far detector	32
3.8	The magnetic coils	34
3.9	A cross section of the light injection module	36
3.10	Detector average PMT 'gain' responses	36
3.11	Relative drift of detector response	38
3.12	Types of neutrino interaction and topologies in the detectors	41
3.13	MINOS ν_e appearance results	43
3.14	The reconstructed energy spectrum of selected NC events at the far detector	44
3.15	Confidence limits for the ν_μ disappearance oscillation fit.	44
3.16	Confidence limits for the MINOS $\bar{\nu}_\mu$ and ν_μ fit	46
4.1	Far/Near ratio for selected run I MC events	48
4.2	Converting measured ND data to an FD Prediction	49
4.3	Example of the purity correction applied in the near detector.	51
4.4	Converting reconstructed energy events to their original truth	52
4.5	Near detector selection efficiency	53
4.6	Far detector selection efficiency	55
4.7	Converting true energy events to their equivalent reconstructed energy	56
4.8	Calculating the FD spectrum background events	58
4.9	Efficiency of selecting wrong-sign μ tracks	59
4.10	FD Purity and NC selection Efficiency	61
4.11	Cross-section and Efficiency of $\bar{\nu}_\tau$ interactions	61

4.12	Comparison between methods of calculating oscillation probability	64
4.13	Integrated error of different oscillation methods.	65
5.1	Fiducial volumes in the near and far detectors	70
5.2	Distribution of the NC discrimination PID S_1	72
5.3	Distribution of the NC discrimination PID S_2	73
5.4	Calculation of the relative angle selection variable	75
5.5	The distribution of the relative angle variable	75
5.6	Distribution of the fitting certainty estimation for PQ events	76
5.7	Optimisation of the PID S_1 for run 1 Monte Carlo	78
5.8	Statistical sensitivity to oscillations with ν_μ -dominant beam	79
5.9	Optimisation of the PID S_2 for $\bar{\nu}_\mu$ -enhanced Monte Carlo	80
5.10	Downstream event detector effects	84
5.11	Band plot of systematic effects	86
5.12	Star plot of the systematic effects	87
5.13	Star plot of the systematic effects vs sensitivity	88
5.14	Data/MC comparison plots for ν_μ -dominant beam detector variables	90
5.15	Data/MC comparison plots for $\bar{\nu}_\mu$ -enhanced beam detector variables	91
6.1	Probability density of the $\Delta\chi^2$ distribution	93
6.2	Sample $\Delta\chi^2_{\text{fit}}$ distribution from the FC method for ν_μ -dominant beam	96
6.3	The $\Delta\chi^2$ surface for ν_μ -dominant beam and the associated FC grid	97
6.4	Distribution of best fit points for an FC grid point	99
6.5	The 68% FC-correction surfaces calculated by applying the two fitting scenarios.	99
6.6	The points considered in the coarse grid search	100
6.7	FC correction grids for separate run 3 and 1+2+3 MC	103
6.8	FC Correction surfaces for ν_μ and $\nu_\mu + \bar{\nu}_\mu$ modes	104
6.9	Systematic Shifts are chosen in a normal distribution	106
6.10	68% and 90% FC Correction Grids	112
6.11	Sensitivity comparison of Gaussian, Statistical FC and Full FC	113
7.1	ND positive charge data	115
7.2	FD Antineutrino results per run	117
7.3	FD Antineutrino results	118
7.4	Gaussian confidence limits of antineutrino fit	119
7.5	Confidence limits of the measurement as a function of $ \Delta\bar{m}^2 $ and $\sin^2(2\bar{\theta})$	121
7.6	FC 68% correction grid for marginalisation of $\sin^2(2\bar{\theta})$	123
7.7	1D Marginalisation corrections	123
7.8	FC 68% correction grid for marginalisation of $ \Delta\bar{m}^2 $	125
7.9	Comparison of confidence limits with other experiments	126

Chapter 1

Introduction

Neutrinos have been of great interest to physicists ever since they were first postulated in the thirties, not least because they have defied expectations at every turn. From being proclaimed “impossible to observe”, it was twenty years before they were first detected from nuclear reactors, and when used to probe the inner workings of the sun there appeared to be a large portion missing!

Eventually, it was discovered that the disappearing neutrinos were actually undergoing flavour change - confirming the oscillation hypothesis, and hence that neutrinos must have a mass. This is the strongest current experimental evidence of physics beyond the standard model.

Since then, many of the parameters describing the mixing between flavours have been measured, using both neutrinos and antineutrinos. This includes the neutrino mass squared differences, Δm_{21}^2 and $|\Delta m_{32}^2|$, and recently, a non-zero value for θ_{13} . For measurements of antineutrinos, reactors supply a steady and useful supply of $\bar{\nu}_e$, whereas the sun provides a source of ν_e . Atmospheric neutrinos provide a mixed flux of ν_μ and $\bar{\nu}_\mu$, as well as ν_e and $\bar{\nu}_e$. In contrast, accelerators provide beams enhanced in ν_μ or $\bar{\nu}_\mu$ and the associated experiments measure L/E extremely precisely. Long-baseline accelerator neutrino oscillation experiments have started operating with $\bar{\nu}_\mu$ -enhanced beams, and the data used in this thesis is from the first such operation of a long baseline experiment.

Prior to the analysis of this data the range of possible “atmospheric” $|\Delta \bar{m}^2|$

values was [3]

$$(1 < |\Delta\bar{m}^2| < 6) \times 10^{-3} \text{ eV}^2 \quad (90\% \text{ C.L.}).$$

In this thesis, data from exposure of the large steel scintillator-based MINOS detectors to 10.49×10^{20} protons-on-target from the NuMI (anti)neutrino beam in Fermilab, IL, USA is used to measure the antineutrino oscillation parameters $|\Delta\bar{m}^2|$ and $\sin^2(2\bar{\theta})$. The magnetised MINOS detectors are used to select a high purity $\bar{\nu}_\mu$ sample. The Feldman-Cousins method has been used to more accurately estimate the effects of systematic and statistical uncertainties on the final results.

Chapter 2 starts with an overview of the history of the discovery and measurements of the neutrino that led to the knowledge we have today, including a summary of the most important experiments. It then discusses the theory of neutrino oscillation, and looks at what measuring a difference between neutrinos and antineutrinos would imply.

Chapter 3 continues by describing the NuMI beam at Fermilab, the (anti)neutrino source for the MINOS Experiment, and then describes both the near and far detectors. The chapter concludes by summarizing recently published results by the experiment.

Chapter 4 describes the extrapolation method used to create predictions of the neutrino energy spectrum at the far detector, using data from the near detector. This procedure allows the significant reduction of systematic errors due to flux and cross-section uncertainties, by turning an absolute measurement into a relative one.

In chapter 5 the best way to select the purest possible set of antineutrino events amongst varying backgrounds is investigated. These backgrounds include wrongly-identified muons from ν_μ -CC interactions, and hadronic showers containing tracks. The largest sources of systematic uncertainty are then investigated to determine what is important to account for in the final result.

Chapter 6 describes the Feldman-Cousins method, a process by which intervals for both the sensitivity of the analysis and the final confidence levels can be calculated, under conditions where the errors in the oscillation parameters are not distributed according to Gaussian statistics. The Feldman-Cousins technique is also used to account for the effect of systematic uncertainties in these intervals.

In chapter 7, the techniques discussed in chapters 4-6 are applied to the actual data collected in the MINOS experiment, and a measurement of both $|\Delta\bar{m}^2|$ and $\sin^2(2\bar{\theta})$ is made. These results are then compared to those from Super-K, which prior to MINOS provided the world's best measurement of $|\Delta\bar{m}_{32}^2|$. Finally, chapter 8 concludes the thesis and looks towards the future measurement of these parameters.

Much of the work developed for this thesis was used in the MINOS publications [2] and [1], with a subset of the full beam data set that is analysed here.

Chapter 2

History and Theory

The history of the neutrino starts at the turn of the 20th century; with the discovery of penetrating rays from radioactive sources in 1896, and their identification as two separate products by Rutherford in 1899, α and β particles were found, swiftly joined by γ rays.

In 1914 it was discovered by Chadwick that the β decay spectrum was continuous [4], in contrast to the discrete energies of α and γ decay products. This was confirmed in 1927 when Ellis and Wooster [5] measured the heating from a radioactive sample in a thick-walled calorimeter to confirm that this continuity was due to the nature of the decay, ruling out other known energy loss mechanisms.

In order to resolve this problem, which was starting to cause questions on the universality of conservation of energy, W. Pauli wrote an open letter to a 1930 Tübingen physics conference. Starting “Dear Radioactive Ladies and Gentlemen” [6], he tentatively proposed a new, neutral fermion inside the nucleus, emitted in β -decay, which he suggested calling a *neutron*. He suggested that the particle would be about the same mass as the electron, and would have up to ten times more penetrating power than γ rays.

When the particle now known as the neutron was subsequently discovered by Chadwick [7] in 1932, Fermi renamed this particle the ‘Neutrino’, meaning ‘Little neutral one’ in Italian. This particle was subsequently included when Fermi for-

mulated a successful theory in 1934 [8] to explain β -decay, which is now known as Fermi theory. Although the scientific community was now convinced on the existence of the neutrino [9], none had been observed, and predictions on the cross-section for interaction by Bethe and Peierls [10] led them to suggest that it was “Absolutely impossible to observe processes of this kind with the neutrinos created in nuclear transformations”.

It was twenty years before Reines and Cowan first detected the neutrino in 1953 [11], and then confirmed in 1956 [12] - using a nearby nuclear reactor, they detected signal coincidence in liquid scintillator via inverse beta decay, now known to be the reaction

$$\bar{\nu}_e + p^+ \rightarrow N^0 + e^+,$$

for which Reines subsequently shared a part of the 1995 Nobel prize for physics.

After the discovery of the muon in 1937 [13] the existence of a second neutrino was discovered, ν_μ in 1962 [14] by searching for lepton flavour number violation via the process $\pi^+ \rightarrow \mu^+ + \nu$ and then $\nu + p \rightarrow n + e^+$. With the subsequent discovery of the τ lepton in 1975 [15] and the measurement of the invisible width of the Z boson [16] in 1989, the number of neutrinos is known to be three. The ν_τ was first observed in 2000 [17, 18] by the DONUT experiment.

2.1 Neutrino Mixing Theory

In the standard theory of neutrino oscillations, any neutrino created in a charged-current weak interaction is created in a weak flavour-eigenstate, which is a sum of mass eigenstates

$$|\nu_\alpha(t=0)\rangle = \sum_k U_{\alpha k}^* |\nu_k\rangle \quad (\alpha = e, \mu, \tau), \quad (2.1)$$

where U is a unitary rotational matrix (satisfying $U^\dagger U = \mathbb{1}$) that defines the scale of any mixing between the mass and flavour eigenstates. The actual form of this

matrix for neutrino mixing will be discussed in section 2.1.1, but the calculations here only depend on the general properties.

These mass eigenstates can be considered eigenstates of the Hamiltonian

$$\mathcal{H} |\nu_k\rangle = E_k |\nu_k\rangle,$$

and the Schrödinger equation

$$i \frac{d}{dt} |v_k(t)\rangle = \mathcal{H} |v_k(t)\rangle$$

can be solved in a way that the neutrino mass states can be expressed as an evolution over time as plane waves:

$$|v_k(t)\rangle = e^{-iE_k t} |v_k\rangle,$$

which, using equation (2.1), lets us describe how the flavour eigenstates evolve:

$$|\nu_\alpha(t)\rangle = \sum_k U_{\alpha k}^* e^{-iE_k t} |v_k\rangle. \quad (2.2)$$

Because of the unitarity of U , we can invert equation (2.1) to give the mass eigenstate in terms of a mixture of flavour eigenstates:

$$|\nu_k\rangle = \sum_\alpha U_{\alpha k} |\nu_\alpha\rangle.$$

Substituting this back into equation (2.2) then gives us:

$$|\nu_\alpha(t)\rangle = \sum_{\beta=e,\nu,\tau} \left(\sum_k U_{\alpha k}^* e^{-iE_k t} U_{\beta k} \right) |\nu_\beta\rangle.$$

Hence, if the mixing matrix U is not diagonal, the pure flavour $|\nu_\alpha(0)\rangle$ becomes at some time t a superposition of different flavour states. Using the projec-

tion of one of these flavour states,

$$\langle \nu_\beta | \nu_\alpha(t) \rangle = \sum_k U_{\alpha k}^* U_{\beta k} e^{-iE_k t},$$

we can describe the probability of transition from flavour $\alpha \rightarrow \beta$ at some time t as

$$P_{\nu_\alpha \rightarrow \nu_\beta}(t) = |\langle \nu_\beta | \nu_\alpha(t) \rangle|^2 = \sum_{k,j} U_{\alpha k}^* U_{\beta k} U_{\alpha j} U_{\beta j}^* e^{-i(E_k - E_j)t}. \quad (2.3)$$

Performing a binomial expansion on the energy, assuming that the neutrinos are ultra-relativistic and thus $m_k \ll |\vec{p}|$ (removing any terms of m_k^n higher than $n = 1$)

$$\begin{aligned} E_k &= \sqrt{\vec{p}^2 + m_k^2} \\ &\approx E + \frac{m_k^2}{2E}, \end{aligned}$$

where $E = |\vec{p}|$, we can thus express the energy difference $E_k - E_j$ as

$$E_k - E_j \approx \frac{\Delta m_{kj}^2}{2E}.$$

Together with the fact that since the neutrinos are travelling at approximately c , $t \simeq L$, (where L is the distance travelled between the source of the neutrino at $t = 0$ and the interaction point), this lets us approximate the transition probability in equation (2.3) as

$$P_{\nu_\alpha \rightarrow \nu_\beta}(L, E) = \sum_{k,j} U_{\alpha k}^* U_{\beta k} U_{\alpha j} U_{\beta j}^* \cdot \exp\left(-i \frac{\Delta m_{kj}^2 L}{2E}\right). \quad (2.4)$$

Because the squared-mass differences and members of the mixing matrix are all physical constants, only the properties L and E are available to be controlled in an experiment designed to observe oscillation.

It turns out to be useful to separate the real and imaginary parts of equation

(2.4). Using $e^{-ix} = \cos(x) + i \sin(x)$, and the fact that, from the unitarity of U ,

$$\sum_k |U_{\alpha k}|^2 |U_{\beta k}|^2 = \delta_{\alpha\beta} - 2 \sum_{k>j} \Re [U_{\alpha k}^* U_{\beta k} U_{\alpha j} U_{\beta j}^*],$$

we can write the oscillatory probability as

$$\begin{aligned} P_{\nu_\alpha \rightarrow \nu_\beta}(L, E) &= \delta_{\alpha\beta} - 2 \sum_{k>j} \Re [U_{\alpha k}^* U_{\beta k} U_{\alpha j} U_{\beta j}^*] \left[1 - \cos \left(\frac{\Delta m_{kj}^2 L}{2E} \right) \right] \\ &\quad + 2 \sum_{k>j} \Im [U_{\alpha k}^* U_{\beta k} U_{\alpha j} U_{\beta j}^*] \left[\sin \left(\frac{\Delta m_{kj}^2 L}{2E} \right) \right], \end{aligned}$$

and finally, using

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

we can conclude that the probability of oscillation can be expressed as:

$$\begin{aligned} P_{\nu_\alpha \rightarrow \nu_\beta}(L, E) &= \delta_{\alpha\beta} - 4 \sum_{k>j} \Re [U_{\alpha k}^* U_{\beta k} U_{\alpha j} U_{\beta j}^*] \left[\sin^2 \left(\frac{\Delta m_{kj}^2 L}{4E} \right) \right] \\ &\quad + 2 \sum_{k>j} \Im [U_{\alpha k}^* U_{\beta k} U_{\alpha j} U_{\beta j}^*] \left[\sin \left(\frac{\Delta m_{kj}^2 L}{2E} \right) \right]. \end{aligned}$$

For the kind of disappearance analysis we are performing in MINOS, and in this thesis, we are particularly interested in the case where $\alpha = \beta = \mu$, which is usually called the *survival probability*. With this condition, the product $U_{\alpha k}^* U_{\beta k} U_{\alpha j} U_{\beta j}^*$ becomes both real and equal to $|U_{\alpha k}|^2 |U_{\alpha j}|^2$. In addition, moving away from natural units where $c = \hbar = 1$ so that E is in units of GeV, L in km and Δm_{ij}^2 is in eV^2 , the relevant survival probability becomes

$$P_{\nu_\alpha \rightarrow \nu_\alpha}(L, E) = 1 - 4 \sum_{k>j} |U_{\alpha k}|^2 |U_{\alpha j}|^2 \sin^2 \left(\frac{1.27 \Delta m_{kj}^2 L}{E} \right). \quad (2.5)$$

2.1.1 PMNS Mixing Matrix

The modern formulism of neutrino oscillations was first first put forward by Maki, Nakagawa and Sakata [19] in order to explain the mixing proposed by Pontecorvo [20].

Relating the lepton flavour eigenstates which experience the weak interaction, $|\nu_e\rangle$, $|\nu_\mu\rangle$ and $|\nu_\tau\rangle$, to the mass eigenstates $|\nu_1\rangle$, $|\nu_2\rangle$ and $|\nu_3\rangle$ which propagate through the vacuum is the mixing matrix U , performing a similar role to the CKM [21, 22] matrix for quark-mixing. This matrix is known as the PMNS matrix.

From the maximum possible eighteen parameters, nine are removed by the unitarity condition, and a further five are relative phases between the lepton fields that they describe. This leaves four parameters free in U , which are normally parametrised as three mixing angles θ_{12} , θ_{23} and θ_{13} , and a complex phase δ :

$$U_{\text{PMNS}} = \begin{pmatrix} U_{e1} & U_{e2} & U_{e3} \\ U_{\mu1} & U_{\mu2} & U_{\mu3} \\ U_{\tau1} & U_{\tau2} & U_{\tau3} \end{pmatrix} = \begin{pmatrix} c_{12}c_{13} & s_{12}c_{13} & s_{13}e^{-i\delta} \\ -s_{12}c_{23} - c_{12}s_{23}s_{13}e^{i\delta} & c_{12}c_{23} - s_{12}s_{23}s_{13}e^{i\delta} & s_{23}c_{13} \\ s_{12}s_{23} - c_{12}c_{23}s_{13}e^{i\delta} & -c_{12}s_{23} - s_{12}c_{23}s_{13}e^{i\delta} & c_{23}c_{13} \end{pmatrix},$$

where $s_{ij} = \sin \theta_{ij}$ and $c_{ij} = \cos \theta_{ij}$. The complex phase δ can, if $\delta \neq 0, \pi$, give rise to CP violation in the lepton sector. For this reason, it is often referred to as δ_{CP} . In addition to this form, for oscillation phenomenology it can be useful to separate this into three separate matrices:

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

where it can be seen that the first and third are only dependent on θ_{23} and θ_{12} , respectively, and the second matrix is only dependent on θ_{13} and δ . Since θ_{13} is known to be small, from experimental data - $\sin^2 2\theta_{13} = 0.089 \pm 0.01$ [23], the second matrix can be approximated by the identity matrix. In addition, the scales of Δm_{23}^2 and Δm_{13}^2 have been shown to be sufficiently different that certain experiments involving θ_{23} and θ_{12} can be considered essentially decoupled. θ_{12} is usually associated with the study of solar neutrino oscillations, and θ_{23} with the study of atmospheric-scale oscillations.

2.1.2 Two-Flavour approximation

This naturally occurring simplification means that the oscillation probability formula can usually be reduced to a two-flavour case. When only considering oscillation between two flavours, the mixing matrix becomes:

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

It can be shown that the two-flavour survival probability can be expressed by substituting into (2.5) as:

$$P_{\nu_\alpha \rightarrow \nu_\alpha}(L, E) = 1 - \sin^2(2\theta) \sin^2\left(\frac{1.27 \Delta m^2 L}{E}\right), \quad (2.6)$$

where E is in units of GeV, L in km and Δm_{ij}^2 is in eV^2 . It can also be shown [9] that the general form reduces to this when using two-flavours.

Figure 2.1 shows the form of this probability, in terms of the combined L/E value, for the approximate value of $|\Delta m^2|$ relevant to the measurements being made in this thesis.

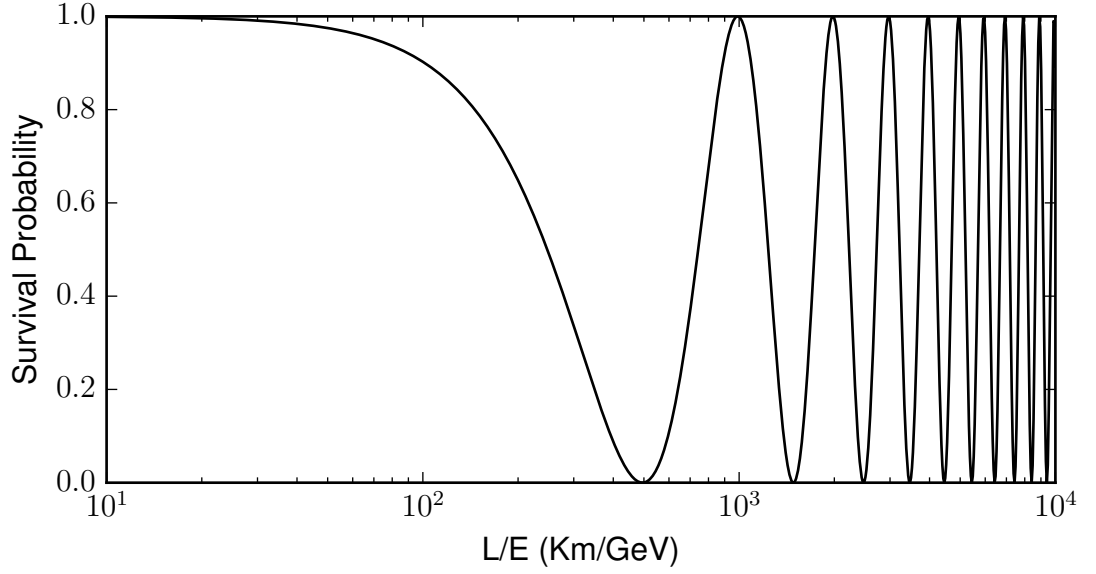


Figure 2.1: Two-flavour oscillation survival probability (equation (2.6)) for $\sin^2(2\theta) = 1.0$ and $|\Delta m^2| = 2.5 \times 10^{-3} \text{ eV}^2$, in terms of the ratio L/E .

2.1.3 CPT Transformation and Conservation

Under the action of charge conjugation (C) alone, left-handed neutrinos¹ are transformed into left-handed antineutrinos. However, only right-handed antineutrinos are observed in nature [26], along with left-handed neutrinos. Therefore, parity transformation (P) must also be applied to produce a state that interacts. This means that neutrinos and antineutrinos are related by a CP transformation

$$\nu_\alpha \xleftrightarrow{CP} \bar{\nu}_\alpha.$$

The addition of a time reversal operation (T) reverses the initial and final states. Therefore, under CP and T transformation a flavour change transforms as:

$$\nu_\alpha \rightarrow \nu_\beta \xleftrightarrow{CPT} \bar{\nu}_\beta \rightarrow \bar{\nu}_\alpha.$$

¹ The term *left-handed* refers to the *chirality* of the particle, which is invariant under Lorentz frame-boosting, and is identical to the *helicity* for massless particles. With massive neutrinos, a particle with specific chirality may have both left and right-handed helicity components, the measured value of which depends upon the frame of measurement [9, 24, 25].

Since CPT is a symmetry for any local quantum field theory [27], the oscillation probabilities under CPT must be equivalent, so

$$P_{\nu_\alpha \rightarrow \nu_\beta} = P_{\bar{\nu}_\beta \rightarrow \bar{\nu}_\alpha}. \quad (2.7)$$

For the special case of the survival probability, equation (2.7) becomes

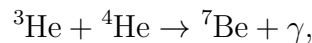
$$P_{\nu_\alpha \rightarrow \nu_\alpha} = P_{\bar{\nu}_\alpha \rightarrow \bar{\nu}_\alpha},$$

and so measuring a difference in oscillation probabilities between neutrinos and antineutrinos in a disappearance channel could reveal a violation of CPT symmetry, or some other new physics.

In 2008 Gonzalez-Garcia and Maltoni [3] analysed the world data from neutrino oscillation experiments in the context of CPT violation where ν and $\bar{\nu}$ are allowed to oscillate independently. Figure 2.2 shows their conclusions for sensitivity to neutrino vs antineutrino oscillation parameter differences in the CPT-violating scenario. The large discrepancy in the precision on Δm_{31}^2 was motivation for the analysis carried out in this thesis, and the associated publications [2, 1, 28].

2.2 The Solar Neutrino Problem

In 1964, motivated by theories developed to explain the thermonuclear reactions in the sun, the Homestake experiment was proposed [29]. In particular, there were new, higher predictions of the rate of the reaction



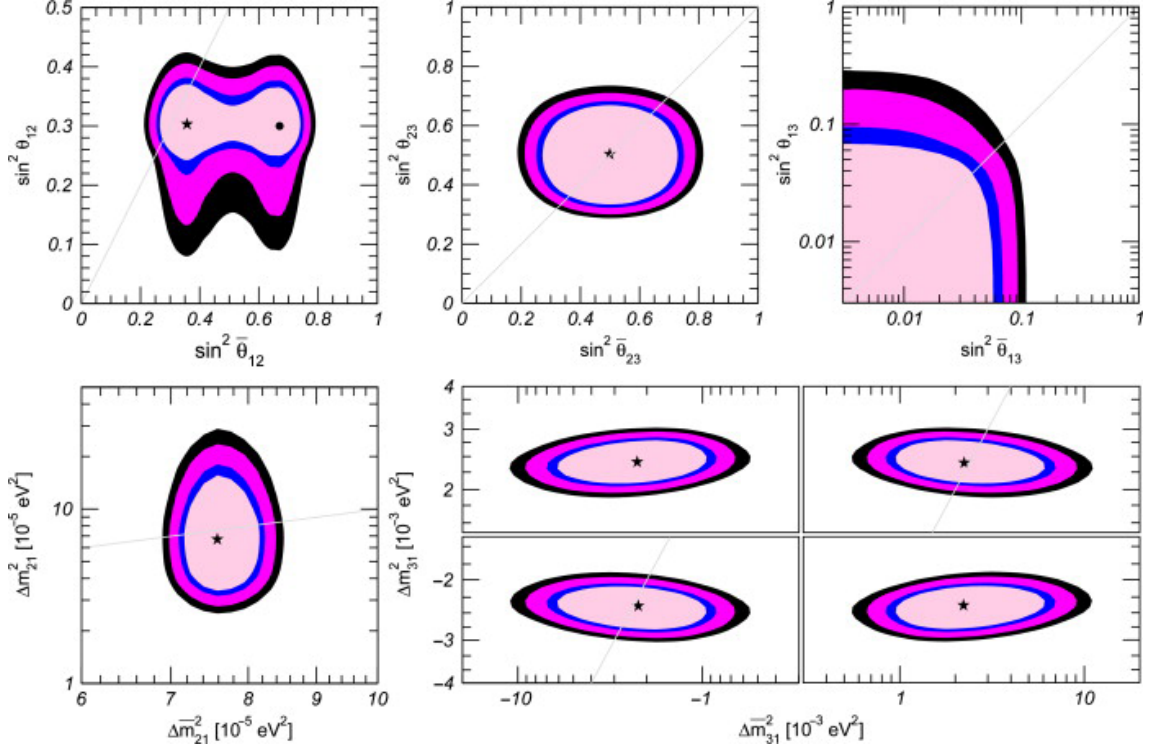


Figure 2.2: Global fit (from 2008) of measurements on ν vs $\bar{\nu}$ oscillation parameters that shows the opportunity for a greatly improved measurement on $|\Delta\bar{m}_{\text{atm}}^2|$. Figure from [3] and shows 90%, 95%, 99% and 3σ contours for neutrino and anti-neutrino mass splittings and mixing angles.

leading to the ${}^7\text{Be}$ and ${}^8\text{B}$ neutrino-producing reactions in the *Standard Solar Model* (SSM) (see figure 2.3), making it possible to observe the neutrinos via the reaction

$$\nu_e + {}^{37}\text{Cl} \rightarrow {}^{37}\text{Ar} + e^-, \quad (2.8)$$

which has a threshold neutrino energy of 0.814 MeV.

Located 1478 m below the surface in the Homestake gold mine, at Lead in South Dakota, and built over 1965-1967 [31], this radiochemical experiment consisted of a 6×10^5 litre tank of tetrachloroethylene (C_2Cl_4) providing a potential 2.16×10^{30} ${}^{37}\text{Cl}$ targets. Every two months, the argon was extracted and the radioactive ${}^{37}\text{Ar}$ was counted via its decay.

When first results came out in 1968 [32], they measured a capture rate of less than $3 \times 10^{-36} \text{ sec}^{-1} \text{ atom}^{-1}$ (with the unit of $1 \times 10^{-36} \text{ sec}^{-1} \text{ atom}^{-1}$ known as a SNU, or *Solar Neutrino Unit*), about a third of the eventually predicted rate [33] of

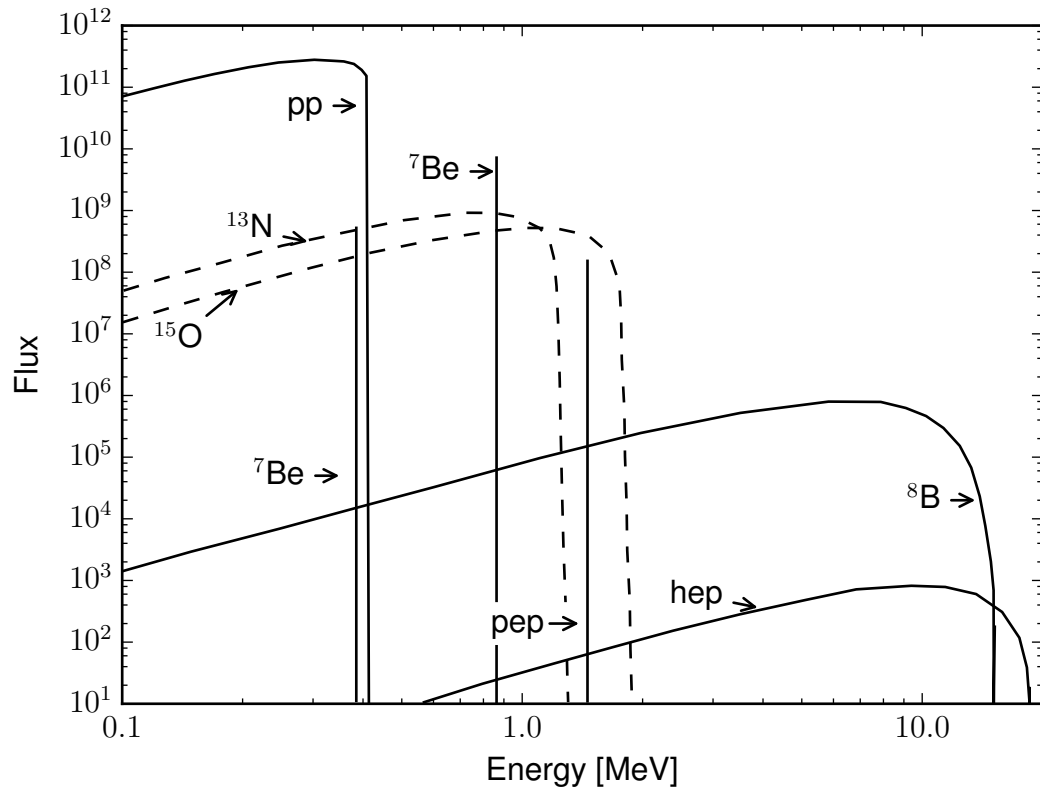
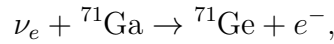


Figure 2.3: Standard solar model prediction of neutrino energy fluxes from from the pp (solid) and CNO (dashed) energy production chains. Data from [30]. The pp chain is the primary sequence of stellar thermonuclear reactions, starting with the $p+p \rightarrow {}^3\text{H}+e^++\nu_e$ (pp) and $p+e^-+p \rightarrow {}^2\text{H}+\nu_e$ (pep) reactions. The CNO cycle is a chain of reactions that converts between Carbon, ^{13}N and ^{15}O , also producing neutrinos.

9.3 ± 1.3 SNU. Further measurement and technique refinement over the next 26 years led to their final measurement of 2.56 ± 0.23 SNU, a discrepancy of over 3σ from the SSM prediction.

Eventually, with mounting evidence from Homestake other experiments started taking measurements of solar neutrinos. The water-Cherenkov detector Kamiokande was enlarged to Kamiokande-II in 1986, and the energy threshold lowered to the point where ^8B solar neutrinos could be detected [34]. After taking data through to 1995, they reported [35] a measured flux of $\Phi_{^8\text{B}}^{\text{Kam}} = (2.80 \pm 0.38) \times 10^6 \text{ cm}^{-2}\text{s}^{-1}$, about half of that predicted by the SSM, by more than 2σ .

In addition, a generation of Gallium-based detectors were constructed, in GALLEX/GNO [36] and SAGE [37]. These detected solar neutrinos via the reaction [38]



which has a lower threshold than (2.8), at $E_\nu^{\text{th}} = 0.233$ MeV - allowing neutrinos from all SSM sources (figure 2.3) to be detected. Crucially, the pp flux is both higher than the ^8B and ^7Be flux, and better known (to within 3% [39]) - as the primary pp reaction $p+p \rightarrow {}^2\text{H} + e^+ + \nu_e$ is well constrained by the total luminosity of the sun. Each of these experiments found [40, 41]:

$$R_{^{71}\text{Ga}}^{\text{GALLEX/GNO}} = 69.3 \pm 5.5 \text{ SNU},$$

$$R_{^{71}\text{Ga}}^{\text{SAGE}} = 70.8_{-6.1}^{+6.5} \text{ SNU},$$

which are both in remarkable agreement, both being about half of the flux predicted by the SSM, at over 5σ significance.

The experimental verification of the solution to these problems came from the *Sudbury Neutrino Observatory* (SNO). SNO was a water Cherenkov detector located in an active mine sited near Sudbury, (Ontario, Canada). A kiloton of heavy water (D_2O) is used as the detector medium, contained in a 12m diameter, spherical transparent-acrylic vessel, itself submerged in a large cavity of ultra-

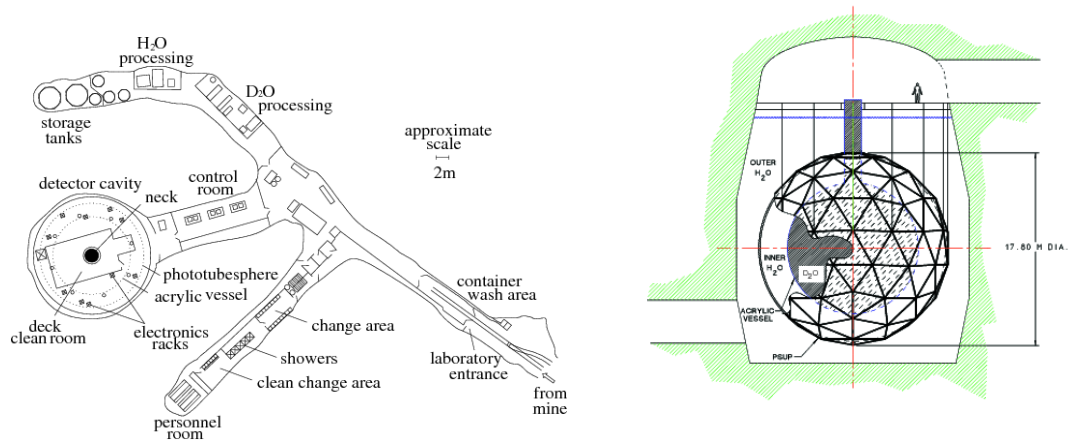


Figure 2.4: The layout of the Sudbury Neutrino Observatory laboratory (left) and schematic of the SNO detector (right). Figures from [42].

pure H₂O, all at a depth of 2092 m, equivalent to 6010 metres of water. The vessel is surrounded by 9456 *PhotoMultiplier Tubes* (PMTs), to observe the neutrino reactions. A schematic of the laboratory, and detector, can be seen in figure 2.4.

Using heavy water, SNO was sensitive to three channels of neutrino interaction:

$$\nu_x + e^- \rightarrow \nu_x + e^-, \quad (\text{ES})$$

$$\nu_e + d \rightarrow e^- + p + p, \quad (\text{CC})$$

$$\nu_x + d \rightarrow \nu_x + n + p, \quad (\text{NC})$$

and, because of energy thresholds, all three channels are only sensitive to ⁸B neutrinos, with the lowest threshold of $E_\nu^{\text{NC}} > 2.224 \text{ MeV}$. Crucially, the NC (and ES) processes are sensitive to all neutrino flavour², so SNO can not only measure the CC flux (which one can assume is almost entirely ν_e), but can also measure the total flux of active neutrinos, regardless of any oscillations that may have occurred. SNO also ran in several different phases, altering the composition of additives to the heavy water in order to enhance various signals.

Figure 2.5 shows the results from SNO, and the final combined measurement

²However, ES is more sensitive to ν_e because the ν_e ES cross-section is much higher than for ν_μ ES interactions

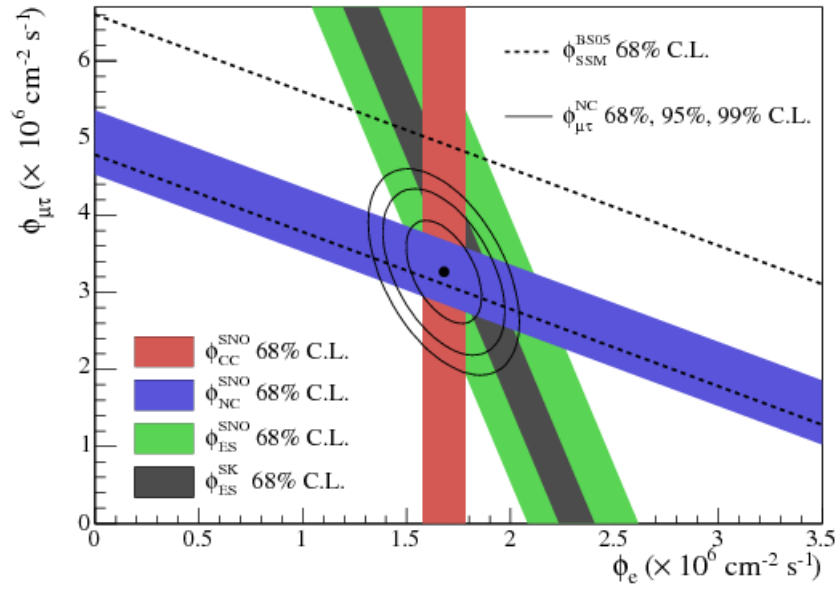


Figure 2.5: Measured flux of ν_e vs $\nu_\mu + \nu_\tau$ by the SNO detector [45]. The CC, NC and ES measurements are shown separately, as is the SSM [44] prediction (within the dashed lines). The Super-Kamiokande result from [46] is also shown.

from SNO gave a flux of [43]:

$$\Phi_{sB} = (5.25 \pm 0.21) \times 10^6 \text{ cm}^{-2} \text{ s}^{-1},$$

which is compatible with the various SSM predictions listed in [44], which range $(4.59 < \Phi_{sB} < 5.79) \times 10^6 \text{ cm}^{-2} \text{ s}^{-1}$. This provided comprehensive evidence that the solar models were accurate, and so neutrinos do, indeed, change flavour, which in turn means that they have a non-zero mass.

However, solar neutrinos were not the only class of neutrino to cause problematic discrepancies between prediction and measurement, as we shall see in section 2.3.

2.3 The Atmospheric Neutrino Anomaly

Neutrinos are produced by cosmic rays incident on the atmosphere. Cosmic rays consist mainly of protons, which cover a vast range of energies up to 10^{20} eV.

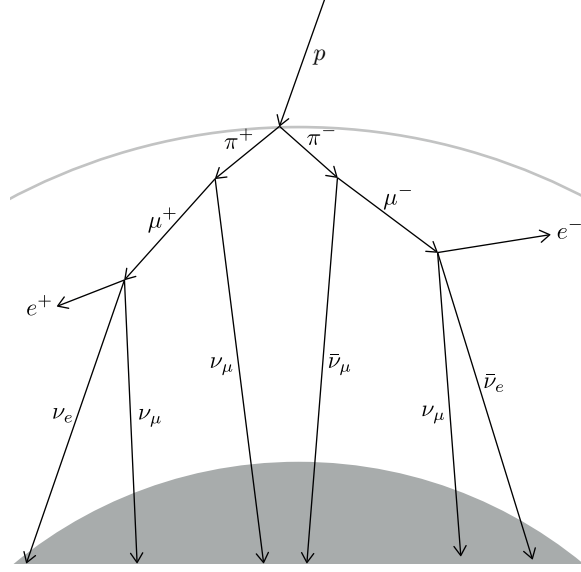


Figure 2.6: Neutrino production by cosmic rays. Based on a diagram in [9].

When interacting with the nuclei in the atmosphere, many secondary hadrons are produced, including many pions which decay to produce neutrinos via

$$\begin{aligned}
 \pi^+ &\rightarrow \mu^+ + \nu_\mu, & \pi^- &\rightarrow \mu^- + \bar{\nu}_\mu, \\
 \mu^+ &\rightarrow e^+ + \nu_e + \bar{\nu}_\mu, & \mu^- &\rightarrow e^- + \bar{\nu}_e + \nu_\mu,
 \end{aligned}$$

and figure 2.6 shows a diagram of these decays.

These *atmospheric* neutrinos can be detected by underground detectors (which would otherwise be overwhelmed by cosmic ray residuals that hit the ground before decaying). Atmospheric neutrinos were first detected in the 1960's by scintillator detectors at the Kolar Gold Field in South India [47] and at the East Rand Proprietary mine in South Africa [48, 49], which could only separate secondary cosmic ray muons from neutrino interactions in the rock surrounding the detectors by measuring the horizontal flux.

In the late 1980's, water Cherenkov detectors in the form of Kamiokande and IMB, designed to search for proton decay, started taking measurements of atmospheric neutrino fluxes, which were an important proton-decay background. These experiments could detect both neutrino interactions inside the detector,

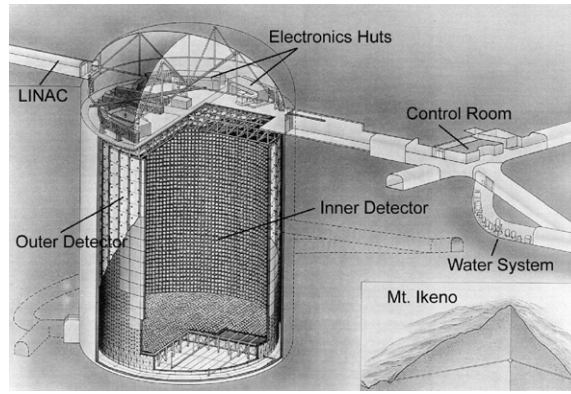


Figure 2.7: The Super-K detector. Image from [54].

and measure incoming muons from neutrino interactions with the rock surrounding the detector. Both experiments measured a significantly smaller total flux of muon neutrinos that was expected [50, 51], with IMB reporting $(26 \pm 3)\%$ out of an expected $(34 \pm 1)\%$ of events with identifiable muons, whilst Kamiokande reported detecting $(59 \pm 7)\%$ of the expected number of muon events. To complicate matters further, two fine-grained iron calorimeter experiments, NUSEX [52] and Frejus [53] found no evidence of this anomaly in their data.

This apparent deficit was known as the *atmospheric neutrino anomaly*, and was resolved by the Super-Kamiokande experiment.

2.3.1 Super-Kamiokande

Super-Kamiokande (Super-K) [54] is a 50 kton water Cherenkov detector, placed about 500 m from where the Kamiokande detector was operated. It consists of two separate detector volumes, a cylindrical inner detector with a height of 36.2 m and diameter of 33.8 m, contained within a cylindrical outer detector 42 m high and 39.3 m in diameter. The outer volume contains 1885 PMTs for veto of particles and to determine containment, and the inner volume contained 11146 PMTs, until an accident before phase II in 2002 caused a reduction to 5182 PMTs. In 2006, for phase III, the PMTs were replaced and Super-K has been running with 11129 inner PMTs since.

With the wide range of energies possessed by atmospheric neutrinos and the

effective baseline $L = 2 * R_{\oplus} * \sin(\theta_z/2)$ (where R_{\oplus} is the radius of the earth, and θ_z is the zenith angle from the horizon at the experiment to the source of the neutrinos) varying from ~ 15 km to 12,000 km, a very large area of L/E parameter space can be sampled by Super-K. Atmospheric ν_{μ} from directly above the detector will have travelled a very short L , and so will not have had a chance to disappear into ν_{τ} , regardless of energy³.

In contrast, the neutrinos from the other side of the earth will have travelled a very large L , and so the survival probability will be highly dependent on both the L and the neutrino energy, to the extent that without extremely precise energy and angle resolution, the average disappearance probability will be measured.

In 1998 the Super-K collaboration published the results of 535 days of exposure, [55], in terms of an asymmetry in up and down-going muons:

$$A_{\mu}^{\text{up-down}} = -0.296 \pm 0.048 \pm 0.01,$$

which was interpreted as a 6σ model-independent confirmation that part of the upward-going neutrino flux disappears, thus concluding that the neutrino has mass. Because Super-K could not easily measure the flux of ν_{τ} to establish that the ν_{μ} were transmuting into another active flavour, a measurement of flavour change (not just disappearance) would have to wait until the SNO results were published in 2002 [56].

In 2011 Super-K published [57] an analysis of their atmospheric neutrino data where neutrinos and antineutrinos were allowed oscillate separately in their fit. They found that, to 90% confidence, $|\Delta\bar{m}^2| = 2.0_{-0.7}^{+2.0} \times 10^{-3} \text{ eV}^2$ and $\sin^2(2\bar{\theta}) > 0.83$. The results can be seen in figure 2.8.

³For the large majority of energy ranges being measured by the detector, at least

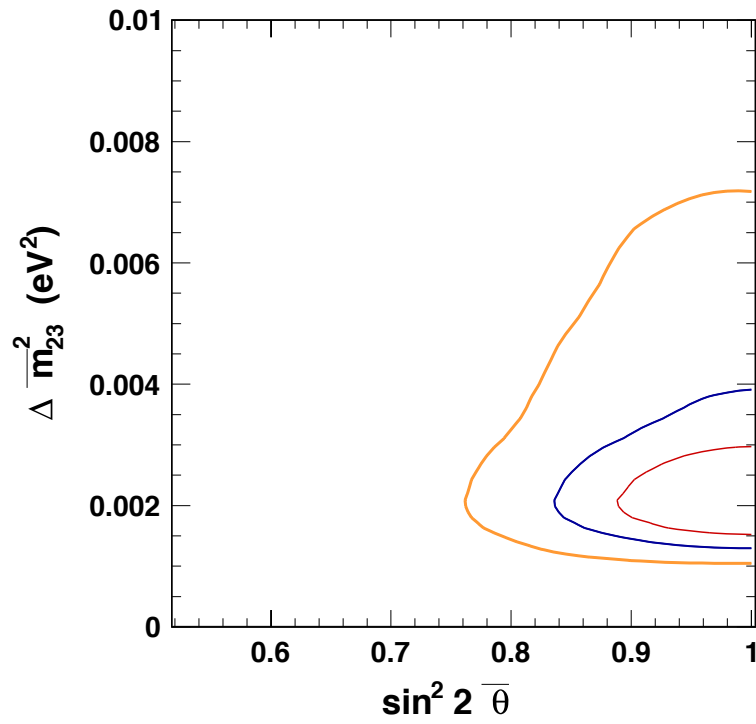


Figure 2.8: Results of Super-K analysis of antineutrino oscillation parameters. The lines are the 68%, 90% and 99% confidence allowed regions. Image from [57]. The areas excluded are those to the left of the contours.

2.4 Recent Measurements

One of the most exciting measurements of recent years in particle physics has been the discovery of a non-zero θ_{13} . Daya Bay consists [58] of six antineutrino detectors, combining central Gd-loaded liquid-scintillator volumes with water Cherenkov-based shields, placed at several different distances from six nuclear power plants. In 2012 Daya Bay released [23] an analysis of the first 55 days of data taking, and measured a non-zero θ_{13} to over 5σ significance - measuring $\sin^2(2\bar{\theta}_{13}) = 0.092 \pm 0.016 \text{ (stat)} \pm 0.005 \text{ (syst)}$.

Since this important measurement, this subdominant oscillation has been confirmed by RENO [59], and lately T2K - operating a muon neutrino beam 295 km from Super-K at the J-PARC accelerator facility, T2K recently measured [60]

$$\sin^2(2\theta_{13}) = 0.140^{+0.038}_{-0.032} \left(0.170^{+0.045}_{-0.037} \right),$$

for normal (inverse) neutrino hierarchy. This result gave a significance of 7.3σ over the $\sin^2 2\theta_{13} = 0$ hypothesis. The discovery of the third and last PMNS mixing angle has opened a door to measuring CP violation, the mass ordering and the octant of θ_{23} .

2.5 Unresolved Properties of the Neutrino

Despite great strides recently in narrowing down the values of parameters in the PMNS matrix, there is still much about the neutrino to discover.

There are two possible descriptions of the fundamental nature of the neutrino as a fermion, Dirac and Majorana. All of the other fermions in the standard model are of Dirac type - distinct particles and antiparticles. In the Majorana case, the neutrino would be its own antiparticle; that is - the only difference between a neutrino and what we currently call an 'antineutrino' would be helicity.

One way of testing the nature of the neutrino, would be searching for the ex-

istence of neutrinoless double-beta decay. Regular double beta decay transforms

$$2n \rightarrow 2p + 2e^- + 2\bar{\nu}_e,$$

where the two neutrinos escape. With Majorana neutrinos, $\nu = \bar{\nu}$, and so it becomes possible for the neutrino emitted by one proton to be absorbed by the other. If observed, this would prove the Majorana nature of the neutrino, and also provide a measurement of the absolute mass - since the rate of this interaction is related to the square of the absolute mass.

One other way to determine the absolute mass scale of the neutrino (independent of its nature) is from measuring the endpoint of the beta decay spectrum - since in the reaction $n \rightarrow p + e^- + \bar{\nu}_e$ all energy except that of the neutrino is directly measurable. Measuring the end of the beta decay spectrum should allow a measurement of the mass of the neutrino (or at least, a measurement of the mix of mass eigenstates created).

Much of the other information that we are missing - the value of the CP violating phase δ_{CP} , the quadrant of θ_{23} and the sign of Δm_{32}^2 should hopefully be within the reach of the next generation of accelerator experiments - NO ν A [61], DUNE [62], Hyper-K [63]. If not directly measurable, these experiments will have significantly enhanced capabilities to exclude large areas of the parameter space for δ_{CP} and θ_{23} .

Chapter 3

The MINOS Experiment

3.1 The NuMI Neutrino Beam

In MINOS, the neutrinos are provided by the *Neutrinos at the Main Injector* (NuMI) beam. The layout of the NuMI facility is detailed in figure 3.1 and can be described in three stages - the main injector and NuMI transfer line, target hall and decay pipe.

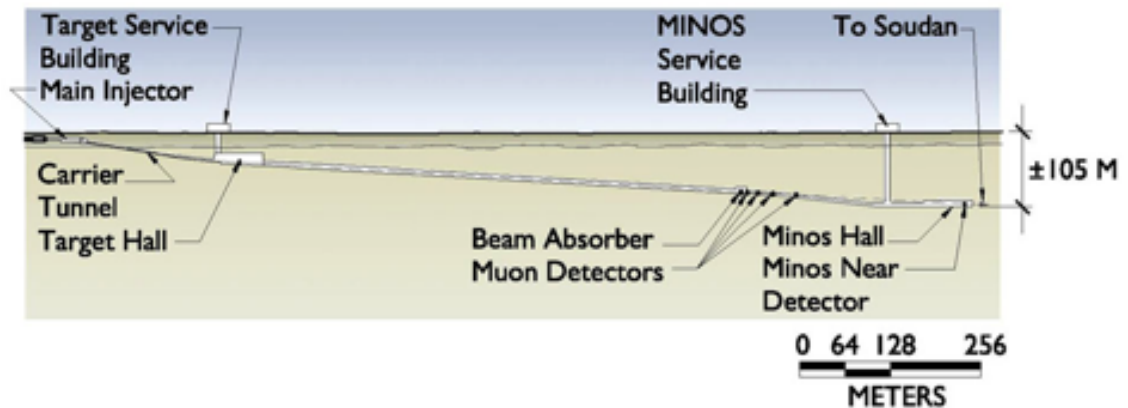


Figure 3.1: Layout of the NuMI facility. From [64].

The Main Injector

The *Main Injector* (MI) at Fermilab is part of the accelerator chain. It accepts 8 GeV protons and accelerates them up to 120 GeV ready for injection into either

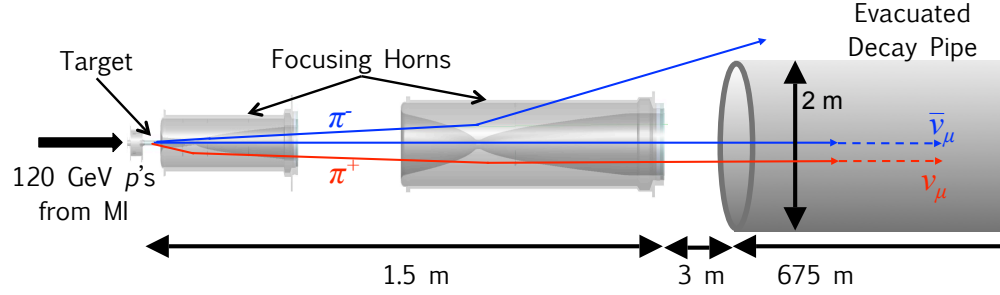


Figure 3.2: Stages the proton beam pass through in the target hall. This diagram shows the horns configured for antineutrino focusing. Image from [66].

the main Tevatron ring or onto the anti-proton production target until 2011, when the main Tevatron ring was shut down, or for use in NuMI. A comprehensive overview of the Main Injector and the accelerator chain is given in [65].

Protons are extracted from the MI by a system of kicker and Lambertson magnets, and then steered underground at a steep 156 mrad gradient, in order to avoid the local aquifer layer. Once below the aquifer the protons are aimed towards Soudan and the far detector, at an angle of 58 mrad, before travelling along a chain of magnets to the target hall.

The kicker operates in a pulsed mode, transferring approximately 3×10^{13} protons from the MI every 1.9 s. The duration of any single pulse, or *spill*, is between 8-10 μs , depending on how many batches of protons were designated for use in NuMI. The beam typically operated at a power of 300 kW.

Target Hall

The target hall is where the proton beam is used to produce the mesons, which then pass into the decay pipe where they decay into muons, neutrinos and other particles. A schematic showing the main components of the beamline in the target hall is given in figure 3.2.

The first device the protons pass through in the target hall is the *baffle*. This is positioned just upstream of the target, and is a 1.5 m long graphite ‘shield’ for the target and horn systems. Because of the power of the proton beam, a mis-

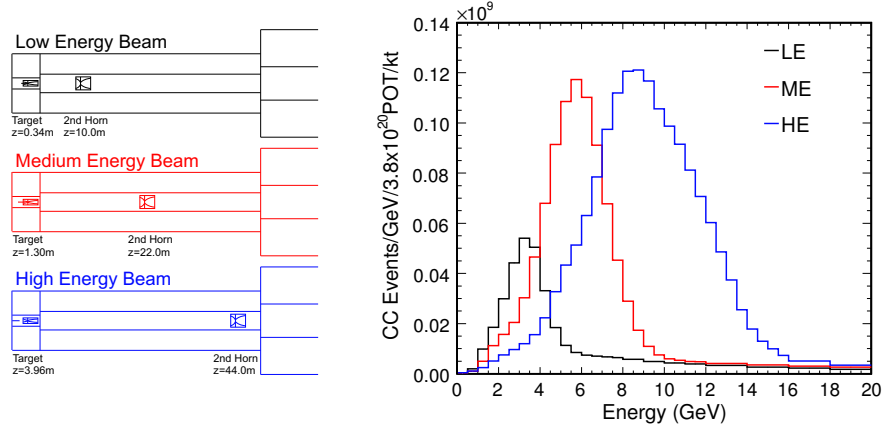


Figure 3.3: Three possible configurations of the target and second magnetic horn, and the resultant changes to the neutrino energy spectra. Diagram from [68].

steered spill could do considerable damage to the horn edges and target cooling systems. The baffle consists of a graphite core with an 11 mm aperture for the correct beam path, and is specifically designed to withstand enough full intensity pulses to give time to shut down the beam.

After the baffle, the protons strike the meson production target, a water-cooled collection of 47 rectangular graphite plates arranged in a series to give approximately 1.9 interaction lengths of material in total. Graphite is used because of its strength and ability to withstand the high, near-instantaneous heating of a spill interaction [67]. The primary product of the target interactions are pions, which decay into muons and neutrinos with 99.99% probability, but there are also other contributions to the neutrino beam composition, primarily from kaons.

The pions and kaons exiting the target proceed into two magnetic focussing horns. The use of such horns is a method pioneered by CERN [69] whereby intense, pulsed magnetic fields are used to focus and select the interaction products coming off the target. Only particles of one charge-sign are focussed, determined by the direction of the current through the horn. The NuMI horns are pulsed with up to 200 kA of current, giving a peak magnetic field of around 3 Tesla.

The longitudinal and transverse momenta of the mesons focussed by the horns is tunable by adjustment of the position of the target and second horn. Figure 3.3 shows several possible configurations, and the corresponding neutrino energy

spectra. Almost all running has been in a variant of the low energy (LE) configuration. To lower the risk of beam line component failure the target was positioned 10 cm further upstream from the lowest energy position (in the direction away from the first horn) and the horns were run at the slightly lower current of 185 kA. This variant is known as *LE-10*.

Decay Pipe

After production and focussing, the mesons enter a 677 m long decay pipe, where they can decay into neutrinos. The decay pipe is a 2 m diameter steel pipe, embedded in at least 1.4 m of concrete shielding.

For the first two years of data taking, the decay pipe was evacuated to minimise further interactions that would occur with air, however in the 2007 accelerator shutdown the decay pipe was filled with helium. The change was made over concerns at potential damage that might be sustained by the entry window [70], which may have caused the entrance to the decay pipe to implode. Because of the positioning and high radiation environment of the decay pipe, access is extremely limited, so this scenario would have been catastrophic for the continued running of NuMI. The helium in the decay pipe contributes an extra 0.17 interaction lengths of material.

At the end of the decay pipe, there is a hadron absorber made of aluminium, steel and concrete. This absorber is designed to stop any hadrons that have reached the end of the decay pipe. Particles absorbed include the 25% of protons that passed through the target, and secondary pions.

The neutrinos and muons pass through the absorber, and then face 300 m of rock before reaching the detector hall. The rock stops any remaining muons ¹, leaving only neutrinos to interact in the detector.

¹From the decay pipe - the neutrinos can interact in this rock to produce new muons, which may reach the detector.

3.2 Detectors

The MINOS experiment uses data from two steel-scintillator sampling calorimeter detectors. The *Near* detector is based at Fermilab and measures the energy spectrum and composition of the neutrino beam, while the *Far* detector sits 735 km downstream in the inactive iron mine at Soudan, MN. A third, calibration detector, known as *CalDet*, was constructed at CERN and exposed to test beams to evaluate the calorimetric response of the detector technology. Exhaustive details of all three detectors design and construction can be found in [71].

All detectors are constructed to be as functionally similar as possible, in order to reduce the systematic uncertainties associated with neutrino cross sections, detector response, and acceptance. Each detector is constructed of many planes of steel plates, of 2.54 cm average thickness², interleaved with 1 cm thick planes of scintillator and an air gap for mechanical tolerance, taking the plane pitch up to 5.95 cm.

The near and far detectors are toroidally magnetized, in order to provide a momentum measurement via curvature and aid in the containment of negatively, or positively charged muons. CalDet was not magnetized, as its primary aim was to measure the hadronic response of the technology, and the momentum and charge sign of the particles in the test beam was known *a priori*. The magnetics of the near and far detectors are discussed in section 3.2.4.

3.2.1 Scintillator

Each scintillator plane is constructed from a series of scintillator strips of identical transverse construction, varying only in length. The solid scintillator strips are constructed of polystyrene (doped with the organic scintillator³) extruded into strips up to 8 m long, with a cross section of 1.0 cm x 4.1 cm. To maximise light

²CalDet is a slight exception, as it built in Europe where 2.50 cm was the nearest obtainable standard

³The organic scintillator used is 1.0% by weight POP (2,5-diphenyloxazole), and 0.03% by weight POPOP (1,4-bis(5-phenyloxazol-2-yl) benzene).

collection, the scintillator strips are co-extruded with a reflective outer layer of TiO_2 . A cross-sectional diagram of a scintillator strip can be seen in figure 3.4, and a photograph in figure 3.5.

Because the absorption length of the scintillator is on the order of 20 cm, an extraction method is needed to get the light out. Each strip has a groove along one side down which a 1.2 mm diameter *Wavelength-shifting* (WLS) fibre-optic cable is glued. This fibre absorbs blue photons from the scintillator, and re-emits green photons isotropically (average wavelength $\lambda = 530$ nm). Any photons re-emitted along the direction of the fibre will be internally reflected, and carried away to the ends of the strip. The glue has a similar refractive index to the scintillator, and the WLS are double-cladded to minimize transmission of the green photons back into the scintillator.

Multiple strips of scintillator are packaged into a single aluminium encased *module*, that provides mechanical support and a light-tight enclosure. At the ends of each module there is a manifold where the WLS fibres are coupled to clear fibres, which are used to carry the light away from the detector to the PMTs. Multiple modules are then used to construct a single plane of scintillator.

3.2.2 Near Detector

The 980 metric ton near detector (seen in figure 3.6) is located in an underground cavern 100 m below the surface at Fermilab, offering an overburden equivalent to 225 metres of water. The detector sits at a distance of around 1 km downstream from the NuMI production target, and measures the characteristics of the intense neutrino beam. The detector is shaped as a squashed octagon, with a width of 4.8 m wide and a height of 3.8 m. Because the high neutrino flux at this position yields approximately 25 neutrino interactions per 3×10^{13} proton spill, the near detector is optimised for a relatively small target fiducial volume, while keeping it as similar to the far detector as possible.

The near detector is constructed of 282 planes of steel, but only 153 of these

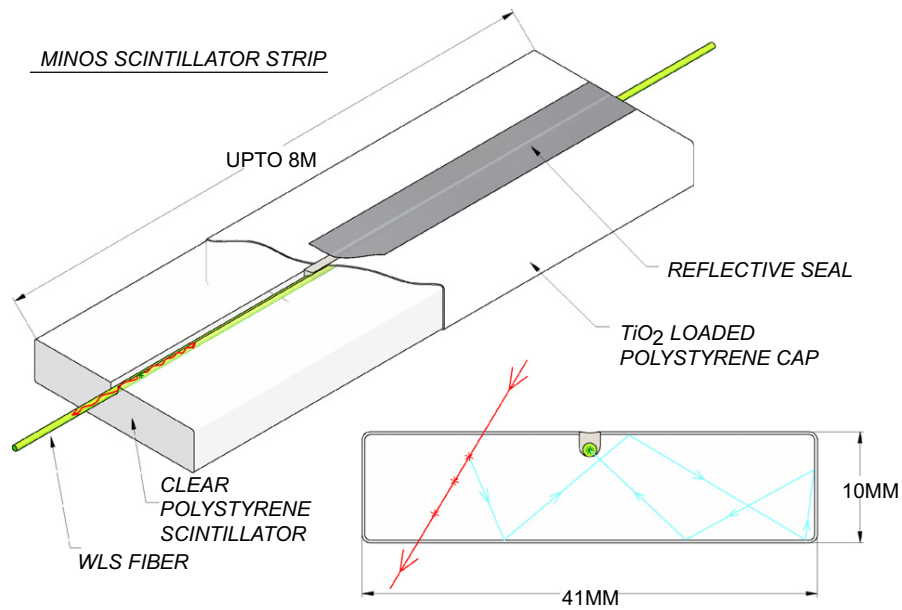


Figure 3.4: Drawing of a cut-away scintillator strip. Light produced by ionizing particles is reflected by the internal reflective coating, and may be absorbed by the wavelength-shifting fibre. It is then re-emitted isotropically. Any photons emitted along the direction of the fibre are trapped and routed out of the detector, to the photomultipliers. From [71].



Figure 3.5: A short scintillator strip being illuminated by a blue LED [72].

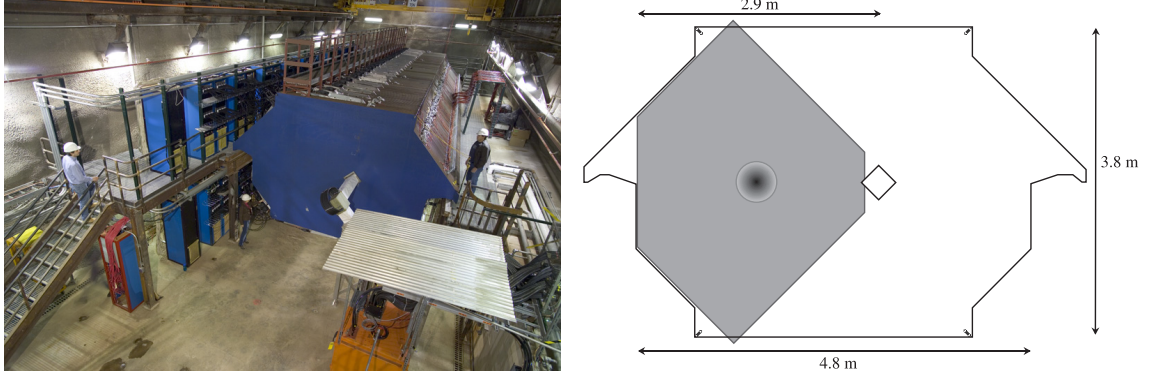


Figure 3.6: Photograph of the near detector looking away from the beam [71] (left) and a diagram of a partially instrumented plane [73] (right). In the diagram, the dark spot is the centre of the neutrino beam, the diamond the magnetic coil hole, and the shaded area shows the instrumented section.

planes are instrumented with scintillator. For data analysis, the first 120 planes are split up logically into three sections: *veto* (21 planes) to discriminate against upstream tracks entering the detector, *target* (40 planes) from where all the interactions of interest should be contained, and finally a *hadron calorimeter* (60 planes) to measure and contain the hadronic showers from any neutrino interactions. Because of cost considerations, only every fifth plane in these sections is fully covered with scintillator, and the other four planes have scintillator covering only the area around the beam spot. The scintillator placement for the partially instrumented planes can be seen in the diagram in figure 3.6. The beam spot is offset 1.48 m from the magnetic coil hole.

The last 161 planes compose the *spectrometer*, to track and ‘range out’ the muons from any neutrino interactions in the target area. This matches the behaviour of the far detector, allowing more of a direct comparison. Because it is not necessary to exactly measure the range of high energy muons, this rear section of the detector is only instrumented every five planes; there is no partial instrumentation.

Because of the smaller size of the near detector (relative to the far detector) and therefore of the scintillator strips, light is read out from only one end, with the other terminated with a reflective cover. In the near detector, Multi-anode Hamamatsu M64 PMTs with 64 pixels are used to read out the scintillator strips.

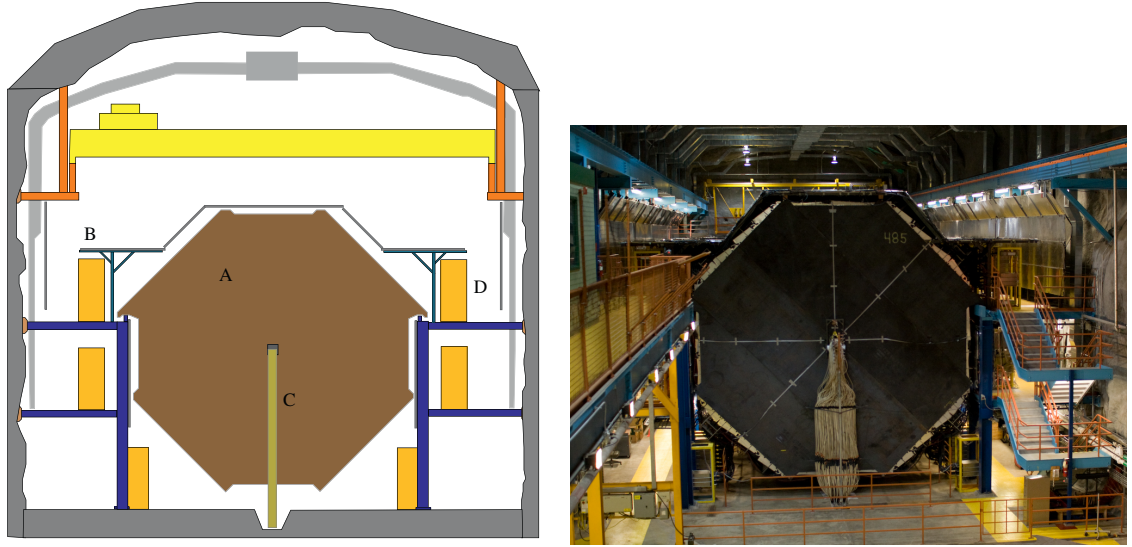


Figure 3.7: Diagram and Photograph of the far detector looking towards the beam. The labelled points on the diagram are A) The furthest downstream steel plane B) The cosmic ray veto shield c) The end of the magnetic coil and D) one of the electronics racks. From [71].

3.2.3 Far Detector

The 5400 metric ton far detector is 735.3 km downstream from the NuMI beam target, in an inactive iron mine at Soudan, Minnesota. The detector sits in a specially-excavated cavern 705 m underground, offering an overburden equivalent to 2070 metres of water. The detector's primary purpose is to look for disappearance of the muon neutrino flux from the beam, relative to that observed at the near detector. Because of the distance from NuMI, the detector only observes of order a few beam-neutrino interactions a day.

The detector is constructed from 486 octagonal steel planes, 8 m in width. Between each of the planes are layers of plastic scintillator. The detector is split longitudinally into two “supermodules”, each independently magnetised. The first (southernmost) supermodule consists of 249 planes, for a length of 14.78 m. The second supermodule is constructed from the remaining 237 planes, for 14.10 m length. The modules are separated by a gap of 1.15 m, taking the total length of the detector to 30.03 m.

Because of the size of the mineshaft, access to the cavern is limited to items of

maximum dimension $9 \times 2 \times 1 \text{ m}^3$ and weight less than 5.5 metric tons, and so the far detector was constructed in-situ from modular components. The scintillator strips were grouped into modules (described in section 3.2.1), and the steel planes were each built by plug-welding eight 2 m wide, 1.27 cm thick plates.

Each scintillator strip is read out at both ends, in contrast to the near detector. The far detector uses 1452 16-anode Hamamatsu M16 PMTs, with three PMTs housed in a single enclosure, called a “MUX” box. In order to save costs, the signals from eight scintillator strips are optically summed (inside the MUX boxes) and then fed to a single PMT pixel; different summing patterns on the opposite ends of the scintillator allow unambiguous reconstruction of scintillator hits.

Cosmic Veto Shield

The planar structure of the detector and its orientation provides challenges when measuring atmospheric neutrinos. Cosmic-ray muons can pass between scintillator planes and penetrate into the centre of the detector before interacting with the scintillator. To aid in the rejection of this background, a veto shield surrounds the upper part of the detector (visible in the diagram and photograph in figure 3.7).

The veto shield is constructed from the same types of scintillator modules as used in the main detector, and uses the same electronics, with an extra 64 PMTs dedicated to reading out the shield. The scintillator strips are still read out at both ends, but the same strips are optically summed at both ends, so demultiplexing is not possible.

3.2.4 Magnetisation

The toroidal magnetisation of the near and far detectors allows momentum measurement and charge-sign determination of the muons due to their curvature in the field, and aids containment of negative or positively charged muons, depending on the sign of the magnetic field.

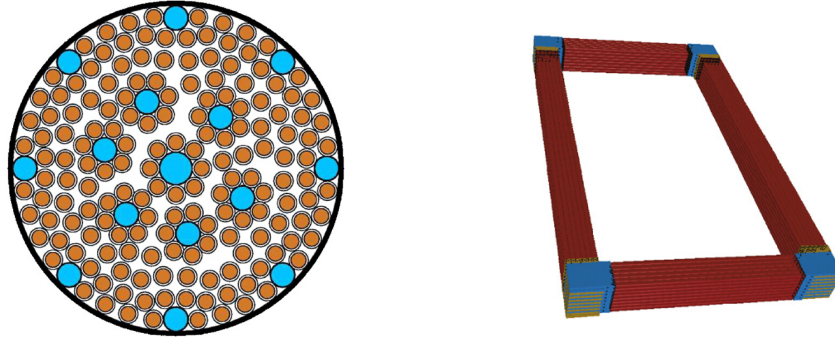


Figure 3.8: Cross-section of the far detectors magnetic coil (left) and a sketch of the four legs of the near detectors coil (right). The far detector coil consists of 190 turns of 8.252 mm stranded copper wire, with cooling pipes embedded. The near detector coil is constructed from six layers of individually wound coil assemblies. Both from [71].

Each detector is required to have a similar strength magnetic field, to help reduce systematic uncertainties. The field strength averaged over the fiducial volumes is 1.28 Tesla in the near detector, and 1.42 Tesla in the far detector. Differences in the topology and structure of the detectors required different designs for the magnetic field systems at each detector.

In the far detector, each supermodule has its own, independently controlled magnetic coil. The coil consists of 190 turns of 1/0 gauge (8.252 mm) diameter stranded copper wire, and consists of a bore leg that passes through the centre of the detectors, a return leg that runs down a trench below the detector, and two end legs, to connect the bore to the return. An 80 A power supply gives a 15.2 kA-turn current, providing the 1.28 T magnetic field. To deal with the 20 KW of power being dissipated by each coil, water cooling pipes are embedded within the copper turns in the bore and return legs (see figure 3.8 for a cross-sectional diagram of the coil). The heat is then removed from the underground laboratory by a secondary heat-exchange system.

In the near detector, a single coil is used. The coil consists of 48 aluminium conductors, each with a central channel for water cooling. Groups of six conductors are formed into a single ‘plank’, and carry current in parallel. Eight of these planks are used to form the coil. A 5 kA power supply provides a 40 kA-turn current, which gives a similar strength magnetic field to the far detector. A sketch of

the four legs of the assembled near detector coil can be seen in figure 3.8.

3.3 Calibration

One of the aims of the MINOS experiment is to measure the energy-dependent disappearance of neutrinos. As such, the level of systematic uncertainty and energy resolution when measuring the energy of particles in the detector is important. Since MINOS measures a relative disappearance it is necessary to calibrate both the relative and absolute energy scales of the detectors. An optical *Light Injection* (LI) system, described in section 3.3.1, measures the linearity and PMT response over time of the readout system in all of the detectors, including CalDet. Muons are used to measure the calorimetric response of the scintillator, calibrate the relative timing response of the detector readout channels, and provides a way to perform the relative calibration between detectors.

Absolute calibration allows the values obtained from the detector readout systems to be translated into an absolute energy value. This is achieved by using measurements from placing a calibration detector, CalDet, in a test beam at CERN, combined with simulations of the particle showers.

3.3.1 The Light Injection System

The *Light Injection* (LI) system [74] is used in the near, far and CalDet detectors to monitor the response of the PMTs and electronics over time, and the integrity of the readout systems.

The WLS fibres from the scintillator modules are illuminated by pulses of light from UV LEDs, replicating a potential signal from the scintillator. The LEDs are housed in groups of 20 or 40 in “pulser boxes”. From the pulser boxes, optical fibres carry the light to LI modules - part of the scintillator module manifolds, where light from a single fibre can illuminate multiple WLS readout fibres. A cross-section of this configuration can be seen in figure 3.9. The intensity of the

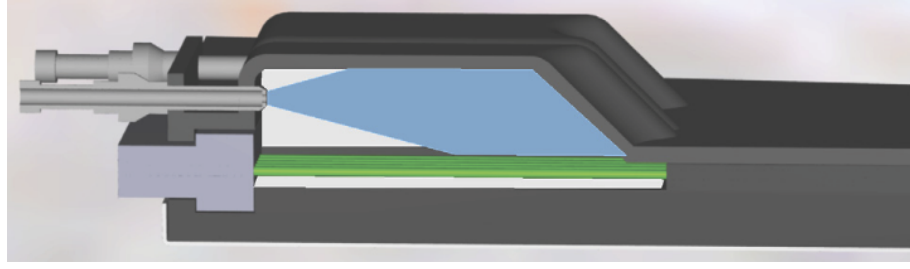


Figure 3.9: A cross section of the LI module, in the manifold at the top of each scintillator module. Up to ten WLS fibres are illuminated by the light from the LI system LEDs. From [74].

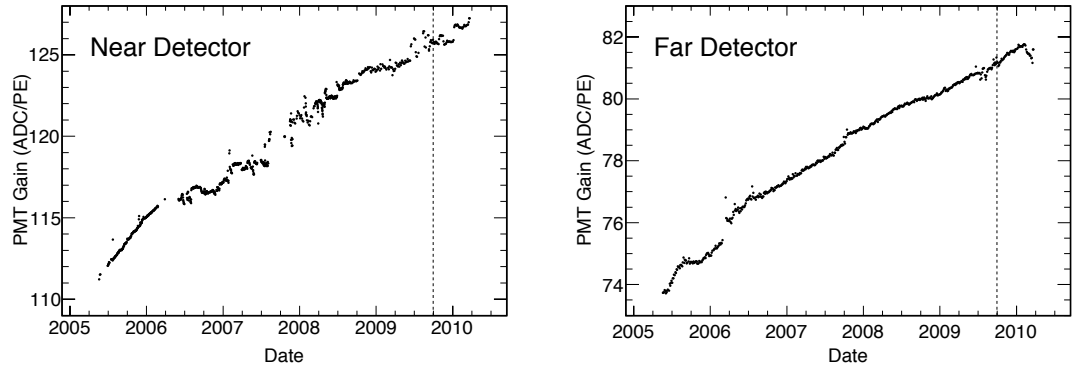


Figure 3.10: Detector average PMT ‘gain’ responses, over the data taking periods, for the near (left) and far (right) detectors. Originally generated for [75].

light from the LEDs is measured by PIN photodiodes⁴, to account for any change in the performance of the LED over time.

During normal operations, each strip end is pulsed around 300 times an hour in the far detector, and 1000 times an hour in the near detector. Both ends of the far detector strips are pulsed, allowing a double check of response using light that leaks to the opposite strip end. Every three days, data from these flashes are analysed and the average response per photoelectron of each PMT channel is calculated to give a set of PMT gains. These values are used by the Monte Carlo simulations to more accurately represent the detector response, and the reconstruction to aid in rejection of background noise. Figure 3.10 shows the detector average gains, over the data taking periods.

The PMT response becomes slightly non-linear at higher light levels (5-10%

⁴A PIN photodiode is like a regular PN diode, but with an undoped semiconducting region between the two doped PN regions. When reverse biased, photons arriving in the I region can cause an electron-hole pair that results in a small, measurable current.

at levels of approximately 100 photoelectrons [71]). Once a month, the linearity of the PMT response is measured by injecting light ranging from just a few to hundreds of photoelectrons. Combined with data from PIN diodes (themselves shown [74] to be linear to 1%), the response of each PMT is parametrized as a function of true illumination. This is then used before reconstruction to calculate the true light level, compensating for any non-linear effects.

3.3.2 Muon tracks

Whereas the LI system measures the time variation of the readout system, it cannot be used to measure the scintillator response over time, or the relative calibration between detectors. For this, we use muons as a standard candle. At the far detector, through-going cosmic muons have an average energy of approximately 200 GeV and a rate of around 0.5 Hz. At the near detector the mean energy is 55 GeV and 10 Hz, from calculations of known cosmic muon flux [76]. These average values stay relatively constant with time.

To calculate the daily change in the response, or *drift* of each detector, the pulse height per plane of each through-going cosmic muon is calculated, and then the daily median of these values is compared to the median of an arbitrary reference point, giving a single value representing the change in detector response. This measurement encompasses changes in the entire calorimeter and readout system. In contrast, the gain measurement made by the LI system is only sensitive to the readout systems. Figure 3.11 shows the measured drift values over the data taking periods.

In addition to the whole-detector response, the mean response of each individual channel is measured, and corrected to the detector average. Track information about each cosmic muon is used to calculate the position and angle of incidence, and corrections are applied to each hit so that the calibration constant is calculated based on a perpendicular hit in the centre of the strip.

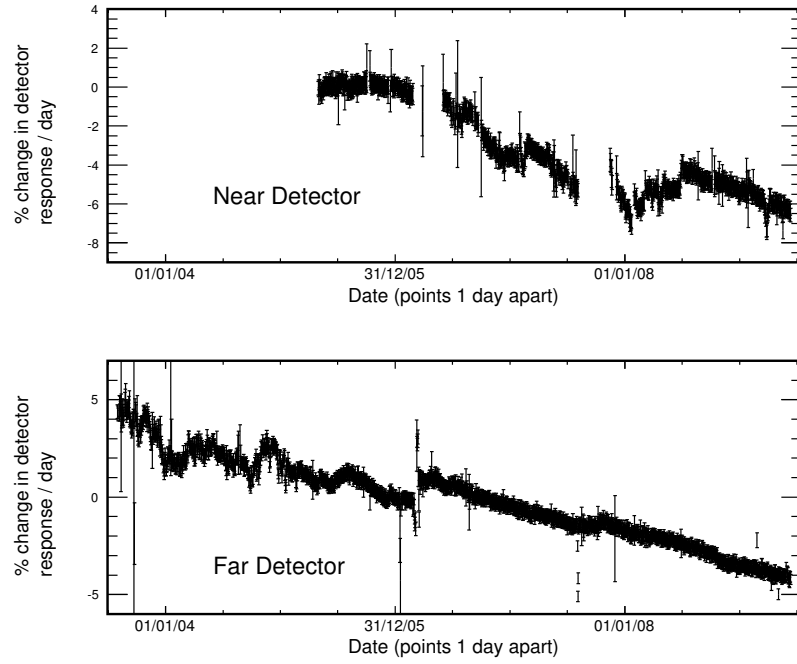


Figure 3.11: Relative drift of detector response, compared to an arbitrarily chosen reference point. From [77].

Relative Detector Calibration

The previously discussed LI and cosmic muon methods calibrate the response of a single detector. In order to use CalDet as an absolute calibration standard, the detectors must be calibrated relative to each other. This is achieved using cosmic ray muons that stop in the detector [72] - they are abundant enough in each, and the energy they deposit in each plane can be accurately determined from range measurements, and the Bethe-Bloch formula.

3.4 Monte Carlo Simulation

The Monte Carlo simulation in MINOS is split into two distinct units - the beam, and the detector simulation. For the beam simulations, a package called FLUGG [78] is used, which combines the FLUKA [79, 80] particle interaction and transport package with the Geant4 [81], for its geometry handling. The beam is simulated separately so that we can simulate a detailed model of the beamline, and study

in detail effects such as shifts in the target position, and the effect on the neutrino energy distribution of helium in the decay pipe. The output from the beam simulation is a set of ‘flux files’ that contain details of the hadron and muon decays for a given configuration. Each distinct data-taking run in the experiment has its own flux files, allowing changes between runs in the beam to be accounted for.

The output from the beam simulation is then fed into the detector simulation. The NEUGEN [82] package is used to simulate the neutrino interactions, both in the rock surrounding the detector caverns and the detectors themselves. The resulting particles are propagated through the detector geometry with the GEANT3 [83] library, accounting for any physical effects of propagation and recording the energy deposits. Showers are simulated using the GCALOR [84] package, which was chosen because it agreed well with CalDet data. It is at this stage that multiple interactions are combined for the near detector, because of the high instantaneous event rate in a beam spill.

The response of the scintillator to the energy deposition, the subsequent propagation of this light to the PMTs is handled by an internal MINOS C++ package called `PhotonTransport`. Another internal package called `DetSim` simulates the PMT response and the effects of the readout electronics. Inverse calibration constants are applied by these packages, such that the MC can be processed by the exact same processing chain as real data - when the reconstruction applies the calibration, the Monte Carlo ends up with neutral calibration. A random date from the running period that is being simulated is chosen for each overlayed set of events.

3.4.1 Flux Tuning

Comparisons between reconstructed ND data and MC simulation have shown that the beam simulation has significant uncertainties. In particular, comparison between different beam configurations show that the high energy tail of the neutrino energy spectrum is consistently mismodelled. This discrepancy stems from

uncertainty in the levels of hadron production from the target, expressed in p_t (transverse) and p_z (longitudinal) momentum parameter space.

To compensate for this, real data is used to carry out a tuning [85] of the hadron production parameters. CC- ν_μ selected data is used to constrain π^+ and K^+ production, whilst CC- $\bar{\nu}_\mu$ selected data is used together with the π^+/π^- production ratio measurements from the NA49 experiment, to constrain π^- and K^- production. The result of this tuning is a re-weighting value in p_t, p_z , which is applied to the Monte Carlo. As a result of this tuning, the agreement between data and Monte Carlo is greatly improved.

3.5 Event Reconstruction

The MINOS C++ based reconstruction software exists to estimate the energy of, and discriminate between, muon tracks from CC interactions, and showers from both CC and NC interactions (see figure 3.12 for typical interactions). Additionally, in the near detector it is the reconstruction that is responsible for splitting the multiple interactions seen in a single beam spill. This is done by applying a set of timing and topology cuts that create *slices* of hits that are grouped in space and time.

Tracks are found by looking for clusters of hits with correlated positions. A Kalman Filter [86] is then used to estimate the path of the track, accounting for any curvature in the detector's magnetic field. The track fitter outputs an estimate of the track momentum from this curvature, and an estimate from the track range. In the analysis stage, the momentum from range is used for muons that stop inside the detector.

The shower reconstruction looks for clustered groups of hits that appear to belong to the same shower. Single hits giving a pulse height of less than 2 photoelectrons are ignored by the reconstruction, as comparisons between data and Monte Carlo showed discrepancies for low level noise (such as PMT crosstalk)

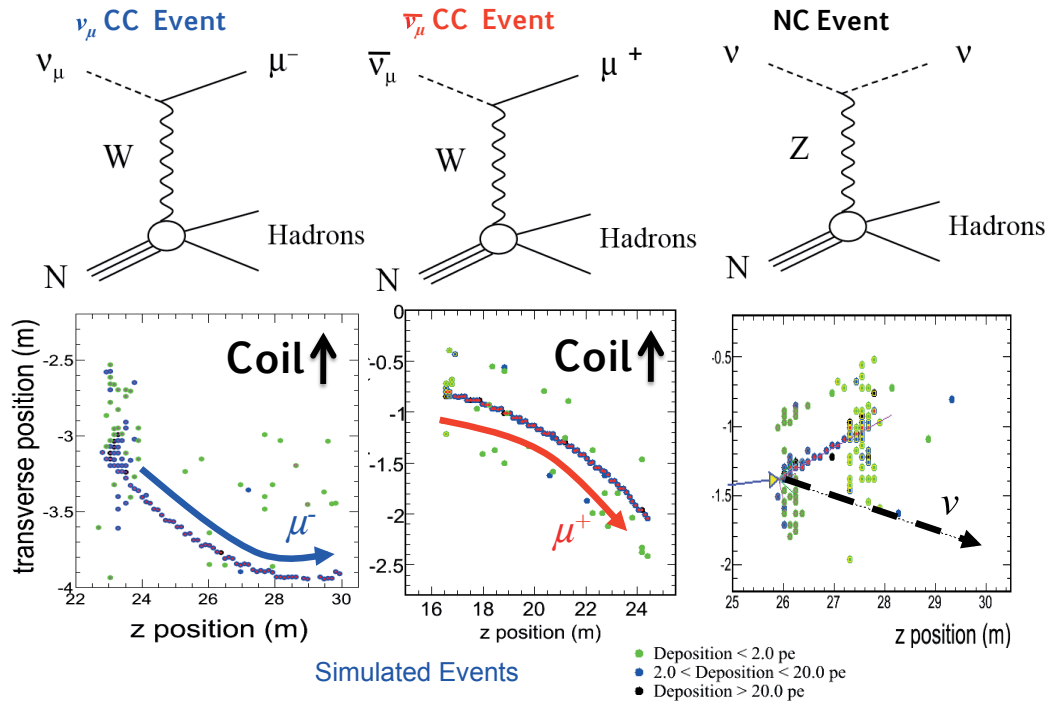


Figure 3.12: The three classes of event topology relevant to this analysis; CC- ν_μ (left), CC- $\bar{\nu}_\mu$ (center) and NC (right). The Feynman diagram for each interaction is shown in the top row, and a typical event from simulated data for the interaction type below. CC-events are recognised by the long muon track, curving in the magnetic field, and NC events typically have only a hadronic shower, from which short tracks can sometimes be found as a background. From [66].

that were not well understood. The energy of the shower is then estimated from the sum of calibrated pulse heights for all hits in the shower.

3.6 Previous MINOS Results

Being a mature experiment, MINOS has already performed multiple analyses exploring several areas of neutrino physics. A large part of this has been in the form of oscillation analyses.

The MINOS analyses encompass several classes of particle interaction. The ν_e appearance analysis looks for Charged Current (CC) events which contain both a hadronic shower, and an electromagnetic shower from the electron. Neutral Current (NC) interactions result in only a hadronic shower in the detector, and contain no flavour information. The ν_μ and $\bar{\nu}_\mu$ disappearance analyses identify CC interactions of the incident neutrinos, with a long track from the muon, and a possible hadronic shower at the event vertex.

ν_e Appearance

MINOS has searched for the appearance of electron neutrinos at the far detector through the $\nu_\mu \rightarrow \nu_e$ channel. This measurement is heavily background dominated, as hadronic showers (from NC and ν_μ -CC) and electromagnetic showers (from the ν_e) can look very similar in the coarse sampling of the far detector, which was designed to measure ν_μ -CC events. In order to discriminate against these background events, 11 variables characterising the shape of the energy deposition in the detector were used in an artificial neural network [87, 88]. Out of an expectation⁵ of 145.2 ± 12.0 (stat) ± 4.9 (syst) events with no ν_e appearance events, 172 events are observed. This allows limits to be placed on a measurement of the θ_{13} parameter. Using the scenario where the CP violating phase $\delta_{\text{CP}} = 0$,

⁵Calculated by adding the errors for the ν_e and $\bar{\nu}_e$ backgrounds in [89] in quadrature.

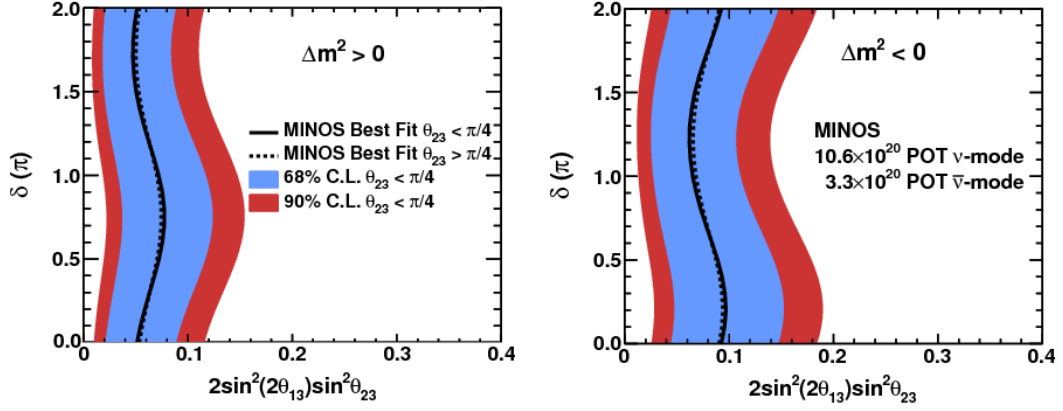


Figure 3.13: ν_e appearance results, showing the confidence limits for both normal (left) and inverted(right) hierarchy scenarios, with $\sin^2(2\theta_{23}) = 0.957$, $|\Delta m_{32}^2| = 2.39 \times 10^{-3} \text{eV}^2$ and as a function of the CP violating phase δ_{CP} . From [89].

$\sin^2(2\theta_{23}) = 0.957$ and $|\Delta m_{32}^2| = 2.39 \times 10^{-3} \text{eV}^2$, the limits are:

$$\begin{aligned} 0.01 &< 2 \sin^2(2\theta_{13}) \sin^2(\theta_{23}) < 0.12 && 90\% \text{ C.L. in the normal mass hierarchy,} \\ 0.03 &< 2 \sin^2(2\theta_{13}) \sin^2(\theta_{23}) < 0.18 && 90\% \text{ C.L. in the inverted mass hierarchy.} \end{aligned}$$

The confidence contours for this data can be seen in figure 3.13. Using recent values for $\sin^2(2\theta_{13}) = 0.098 \pm 0.013$ from reactor experiments, MINOS also demonstrated the capability of such long-baseline experiments to put constraints on the value of δ_{cp} , the octant of θ_{23} and the neutrino mass hierarchy. More details on this analysis can be found in reference [89].

Sterile Search

Because the NC interaction is independent of neutrino flavour, the rate should not change in the standard three-flavour oscillation model. Any measured difference could indicate oscillations to sterile neutrino flavour. MINOS has carried out an analysis searching for such a deficit [90]. Because any ν_e CC events would be included with the selected NC sample, the result depends on the possibility of ν_e appearance. Out of an expectation of $754 \pm 28(\text{stat}) \pm 37(\text{syst})$ events, 802 events were observed, shown in figure 3.14. This puts a limit on the fraction of neutrinos

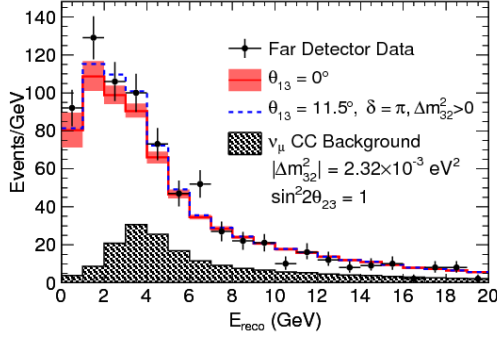


Figure 3.14: The reconstructed energy spectrum of selected NC events at the far detector. The dashed blue line shows the expectation with ν_e appearance - the solid red line without. [90].

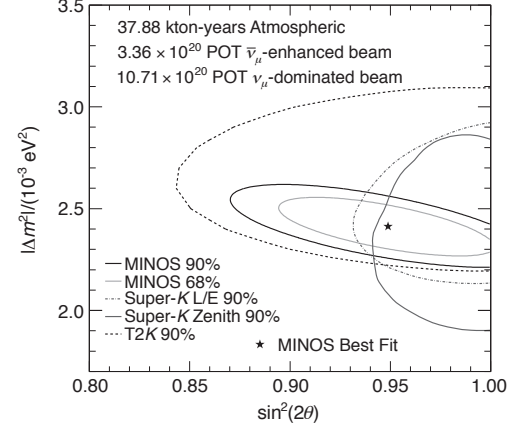


Figure 3.15: Confidence limits for the ν_μ disappearance oscillation fit [28]. Previous experimental limits are also shown.

oscillating to sterile flavours as:

$$f_s \equiv \frac{P_{\nu_\mu \rightarrow \nu_s}}{1 - P_{\nu_\mu \rightarrow \nu_\mu}} < \begin{cases} 0.22 & \text{no } \nu_e \text{ appearance} \\ 0.40 & \text{with } \nu_e \text{ appearance} \end{cases} \quad \text{at 90\% C.L.,}$$

where the fraction with ν_e appearance is assuming $\sin^2(2\theta_{13}) = 0.15$, slightly higher than the currently measured value.

ν_μ Disappearance

MINOS measures the atmospheric-scale oscillation parameters $|\Delta m^2|$ and $\sin^2(2\theta)$ through observing disappearance of ν_μ flux between the two detectors, over the baseline of the experiment. In a two-flavour approximation⁶, the survival probability of a ν_μ with an energy E [GeV] travelling a distance L [km] is given by:

$$P(\nu_\mu \rightarrow \nu_\mu) = 1 - \sin^2(2\theta_{23}) \sin^2\left(\frac{1.27 \Delta m_{32}^2 L}{E}\right),$$

where Δm_{32}^2 is the atmospheric neutrino mass splitting, and θ_{23} is the mixing angle.

⁶See section 2.1 for a more thorough description of neutrino mixing

MINOS has published ([28]) results for 10.71×10^{20} POT of ν_μ -dominant and 3.36×10^{20} POT of $\bar{\nu}_\mu$ -enhanced beam, combined with 37.88 kton years of exposure to atmospheric neutrinos. Many analysis improvements over the previous results ([91, 92]) were made, including updated simulation and reconstruction, selection improvements with improved efficiency, improved shower energy resolution, and reduced systematic errors. Data events are separated into bins of both energy and resolution, increasing the sensitivity of the measurement.

Purely in terms of ν_μ in the ν_μ -dominant beam, 2579 events were selected, out of a no-oscillation expectation of 3201 events. In addition to these, there were 905 atmospheric events of 1100 expected for the no-oscillation scenario, along with smaller $\bar{\nu}_\mu$ samples and large (but with relatively low contribution to the end result) antifiducial samples - events for which the primary interaction vertex was outside of the fiducial volume.

The best two-flavour neutrino oscillation fit to all data measures

$$|\Delta m^2| = (2.41_{-0.10}^{+0.09}) \times 10^{-3} \text{ eV}^2,$$

$$\sin^2(2\theta) = 0.950_{-0.036}^{+0.035}.$$

The confidence limit contours for the best fit to the data, compared to the previous results, are shown in figure 3.15. The MINOS results are highly compatible with the 90% Super-K results for mixing angles $\sin^2(2\theta) > 0.95$, including the best fit.

$\bar{\nu}_\mu$ Disappearance

In addition to the two-parameter fit to beam + atmospheric data sets described above, a four-parameter oscillation fit assuming separate $|\Delta m^2|$ and $|\Delta \bar{m}^2|$ (and $\sin^2(2\theta)$ and $\sin^2(2\bar{\theta})$) was done. This analysis builds on the work done for this thesis and earlier associated publications [2, 1].

$\bar{\nu}_\mu$ events are selected based on the reconstructed charge of the muon. Dis-

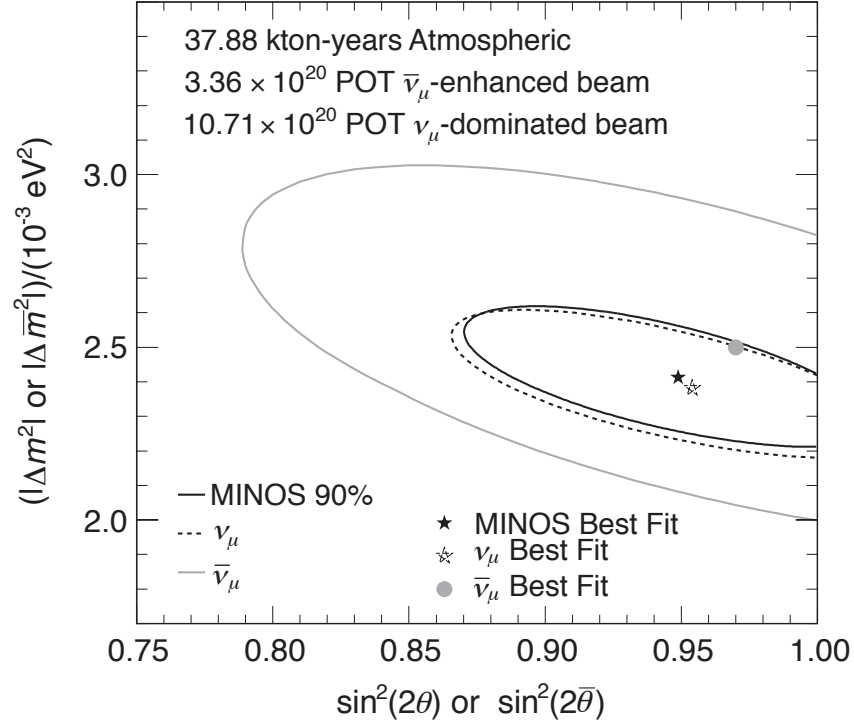


Figure 3.16: Confidence limits for the MINOS $\bar{\nu}_\mu$ and ν_μ oscillation fit, compared to the combined $\nu_\mu = \bar{\nu}_\mu$ fit displayed in figure 3.15. From [28].

crimination is achieved against background NC events in the $\bar{\nu}_\mu$ -enhanced beam by using a k-Nearest-Neighbour algorithm - data is compared to Monte- Carlo events in four parameters: track length, mean energy of track hits, energy fluctuations along the track, and transverse track profile. These discriminations allow MINOS to attain a high antineutrino purity in the final data set of $> 98\%$.

With an expectation of 364 (313) events in the ν_μ -dominant beam and $\bar{\nu}_\mu$ -enhanced beam respectively, 312 (226) events were measured. The four-parameter fit finds

$$|\Delta\bar{m}^2| = (2.50^{+0.23}_{-0.25}) \times 10^{-3} \text{ eV}^2,$$

$$\sin^2(2\bar{\theta}) = 0.97^{+0.03}_{-0.08}.$$

The confidence limits from this analysis, compared to those for the ν_μ disappearance analysis, can be seen in figure 3.16.

Chapter 4

Extrapolation

4.1 Predicting the Far Detector Energy Spectrum

Measurement of the ν_μ disappearance parameters in MINOS consists of looking for a deficit in the energy spectrum of ν_μ interactions in the FD, compared to that extrapolated from the measurement in the ND. Because of differences between the detector acceptances and in beam line geometry, the relative shape of the neutrino flux energy spectrum differs by up to 30%.

As explained in section 3.4, simulated data for both detectors are generated using a beam line simulation. In the simplest method of absolute prediction, this Monte Carlo could be used directly. This has several disadvantages, not least of which is a direct sensitivity to multiple sources of systematic uncertainty. The purpose of the extrapolation is to use the ND data to build a prediction of the event spectrum at the far detector that is relatively insensitive to systematics due to flux, cross-section and detector systematic uncertainties, by turning an absolute measurement into a relative one.

This chapter will start out by examining several ways of using the ND data to more accurately predict the spectrum of interactions in the FD, and will then explain in detail the method that has been used for past analyses. The changes to this method will then be examined, in order to give a complete picture of how the FD energy spectrum is predicted for the analysis in this thesis.

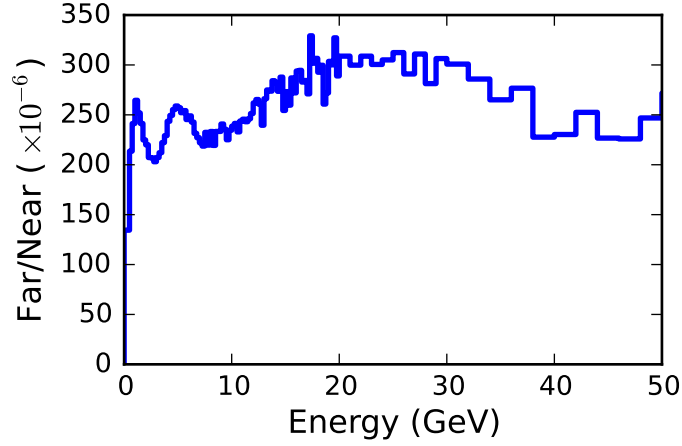


Figure 4.1: Far/Near ratio for selected run I MC events, in reconstructed energy. Each detector has been normalised to the same POT. The small nature of the absolute y scale is due to the approximate $1/r^2$ dependence of the neutrino flux over the 735 km baseline.

Far/Near

The fact that the neutrino spectra shapes are similar between the two detectors suggests that we can use the differences between ND data and MC spectra to correct the FD prediction, which naturally accounts for differences between the detectors, such as energy resolution and efficiency differences. We do this using the Monte Carlo for each detector, running the fully simulated events through the selection criteria to get a simulated energy spectrum F_{MC} for the far detector, and N_{MC} for the near detector. We can then use the ratio of these [93] to correct the ND data, creating a prediction for the far detector

$$F_{\text{predicted}} = \frac{F_{\text{MC}}}{N_{\text{MC}}} \times N_{\text{data}},$$

where N_{data} is the measured near detector data and $F_{\text{predicted}}$ is the resultant far detector prediction. The ratio F is shown in figure 4.1. To create an oscillated far detector prediction, the F_{MC} may be re-weighted with the desired oscillation parameters prior to calculating the ratio.

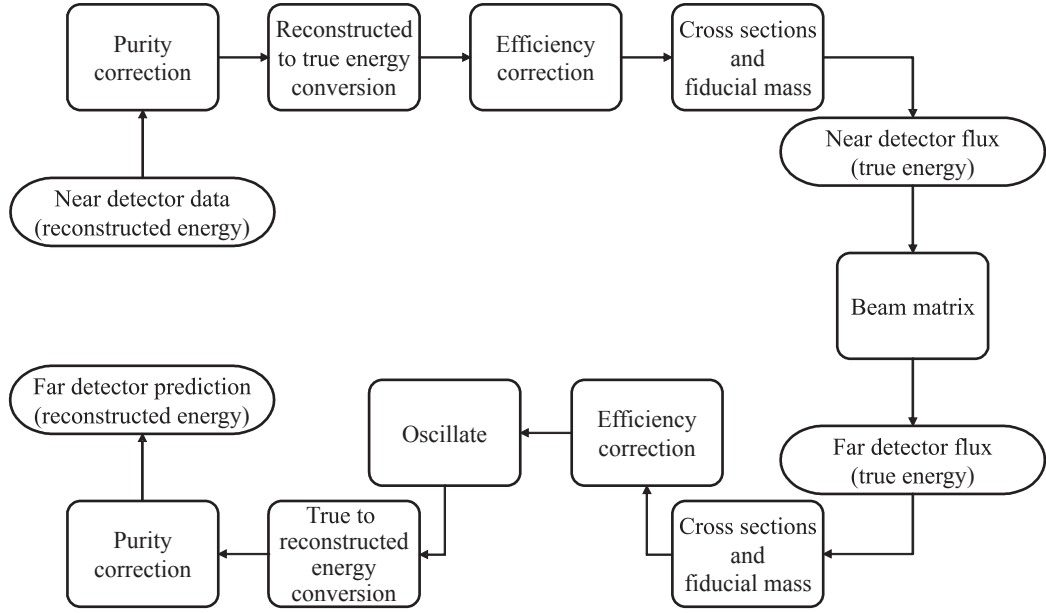


Figure 4.2: The steps to convert measured ND data to a FD prediction with the matrix method. [68, 94]

4.2 Beam Matrix Method

The beam matrix method is an the extrapolation method used for all the CC ν_μ and $\bar{\nu}_\mu$ disappearance studies [68, 94, 92], and is the extrapolation method used for the analysis in this thesis. A key element of the method is a 2D transfer matrix that translates a flux spectrum at the near detector to a flux spectrum at the far detector. The specifics of this mechanism are discussed in section 4.2.1. A flowchart showing the steps involved in the matrix method can be seen in figure 4.2. Because of the use of a transfer matrix, this method inherently works with binned data. By choosing bins of an appropriate size, the effect on the accuracy of the output of this method compared to a more precise method (e.g. more finely binned data, or unbinned extrapolation methods) can be shown to be marginal.

A general overview of the process is presented here, followed by an in-depth description in section 4.2.1.

Starting from the spectrum of selected near detector reconstructed energy events, a purity correction is applied to remove the effects of impurities, by subtracting the expected backgrounds. This pure ν_μ -CC reconstructed energy spectrum is then translated into an equivalent true energy spectrum, by using a de-

convolution matrix. The effect of efficiency from selection is then taken out, followed by interaction cross-section, fiducial mass and then POT exposure from the beam. This results in a neutrino flux distribution, describing the absolute number of neutrinos travelling through a unit section of the detector, per unit of beam.

After application of the beam matrix transformation, giving the corresponding neutrino flux for the far detector, these steps are applied in reverse order; the effects of POT, cross-section, far detector fiducial mass and the selection efficiency are multiplied back in. It is at this stage that the effects of neutrino oscillation are applied. After this, the spectrum undergoes another 2D matrix correction to account for energy smearing in the detector. Finally, a purity correction is made to account for the various background signals. The method of this correction is discussed in section 4.2.4.

Adaption for $\bar{\nu}_\mu$ Analysis

The matrix method was used for the $\bar{\nu}_\mu$ disappearance analysis, detailed in [68]. This required a few changes to the general prescription listed above. The main change for the analysis was the separation of events originating from ν_μ and $\bar{\nu}_\mu$ interactions. To allow the potential separate measurement of ν_μ and $\bar{\nu}_\mu$ oscillation parameters, the ν_μ and $\bar{\nu}_\mu$ spectra were individually extrapolated, with their own efficiencies, purities, energy resolution matrices and beam matrices. In addition, the way the far detector purity correction was made changed, due to the fact that mis-classified events from ν_μ interactions is a background for the $\bar{\nu}_\mu$ analysis.

4.2.1 Step-by-step

In this section, each step of the matrix method is examined in detail. The starting point is the near detector CC energy spectrum, the selection of which is examined in detail in chapter 5.

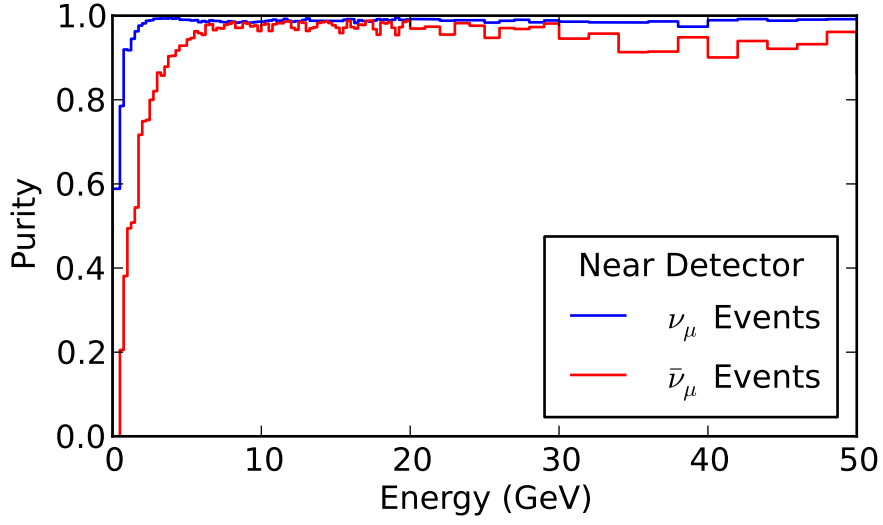


Figure 4.3: Example of the purity correction applied in the near detector. Each experimental run has its own purity correction histograms for ν_μ and $\bar{\nu}_\mu$ selections.

Near detector Purity

The first step of the matrix method is to remove the effect of impurities from the near detector data sample. Impurities are from event interactions that are misclassified, and that pass through the selection criteria. The definition of purity as applied to the selection is calculate using MC simulated data as:

$$P_i^N = \frac{(\text{Number of CC signal events selected})_i}{(\text{Total number of events selected})_i}, \quad (4.1)$$

where P_i^N indicates the purity in the i^{th} energy bin of the histogram for the near detector, signal events means ν_μ or $\bar{\nu}_\mu$ events depending on whether you are in the neutrino or antineutrino half of the extrapolation, and the total number of selected events is the literal number of events selected. A histogram as an example this correction, for neutrinos and antineutrinos, can be seen in figure 4.3.

The correction is applied as:

$$R_i^N \rightarrow R_i^{N,\text{pure}} = R_i^N P_i^N,$$

where R indicates the reconstructed energy spectrum, i is the bin in question, N

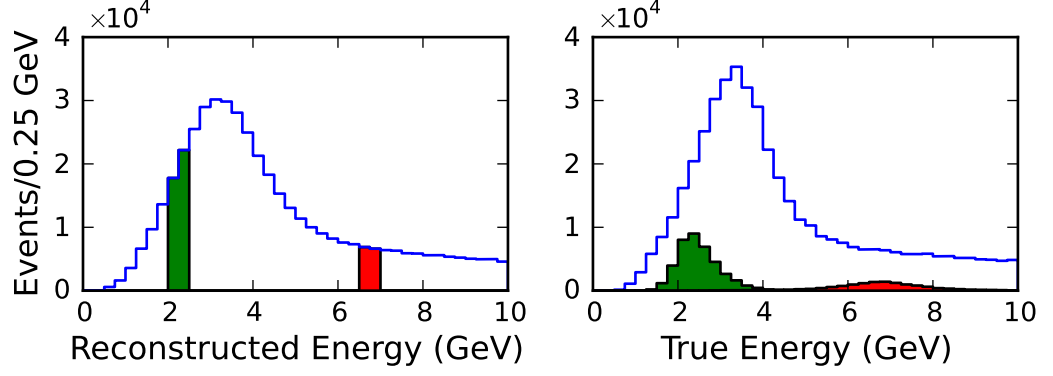


Figure 4.4: Reconstructed energy of near detector events (left) and the Monte Carlo truth energy of the interaction that caused these events (right). A single bin of reconstructed energy is contributed to by a range of true energies.

indicates we are operating on the near detector, and P is the purity correction defined in equation (4.1).

Converting Reconstructed to True Energy

At this point the energy spectrum is still in reconstructed energy, the energy determined by the reconstruction software. This is related to the true energy, but has the effects of the detector energy resolution folded in, causing a smearing of the energy spectrum, with many events at a given reconstructed energy coming from interactions over a range of true energy. The effects of this are demonstrated in figure 4.4. In order to deconvolve this spectrum, we use Monte Carlo truth information to build a 2D matrix \mathcal{M}_{ij}^N , where the i axis represents reconstructed energy, and the j axis represents true energy. This matrix is normalised:

$$\sum_i \mathcal{M}_{ij}^N = 1, \quad (4.2)$$

such that for every true energy bin event, only one reconstructed energy event is contributed. The transformation is then applied as a matrix multiplication:

$$R_i^{N,\text{pure}} \rightarrow T_j^N = \sum_i R_i^{N,\text{pure}} \mathcal{M}_{ij}^N.$$

Since this transformation is operating on a pure $\nu_\mu/\bar{\nu}_\mu$ -CC reconstructed en-

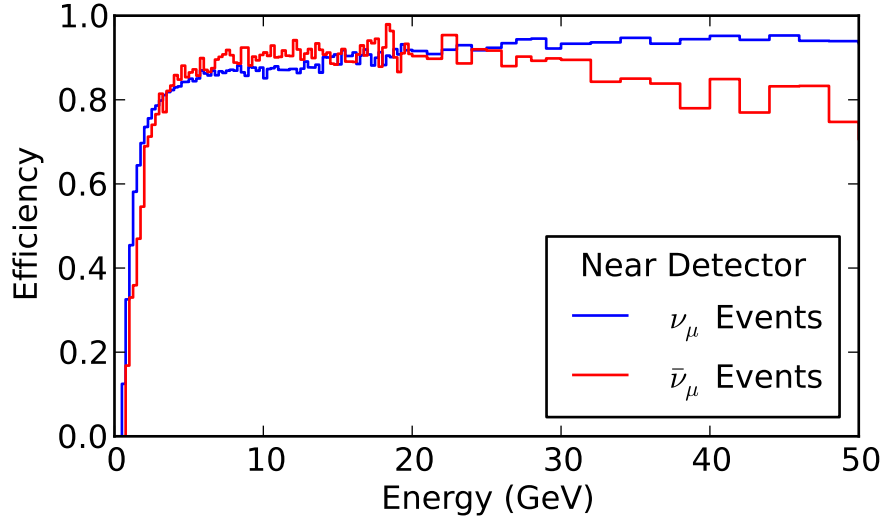


Figure 4.5: The efficiency of selection at the far near detector, for both ν_μ and $\bar{\nu}_\mu$ samples.

ergy spectrum, only these signal events are considered in the building of the matrix.

Efficiency correction

The true energy spectrum is then corrected for reconstruction and selection efficiency. The efficiency of the selection is defined using MC simulated data as:

$$E_j^N = \frac{(\text{Number of signal events selected})_j}{(\text{Total number of true signal events})_j}, \quad (4.3)$$

where E_j^N is the efficiency of the selection at the near detector, for a particular true energy bin j , and the total number of signal events is the total number of true $\nu_\mu/\bar{\nu}_\mu$ -CC events interacting in the fiducial volume. Figure 4.5 shows the typical near detector selection efficiency. This correction compensates for an imperfect selection, and is applied thusly, where $T_j^{N,E}$ indicates the efficiency-corrected true-energy near detector spectrum:

$$T_j^N \rightarrow T_j^{N,E} = \frac{T_j^N}{E_j^N}.$$

Cross-section, Fiducial mass, and POT

The final steps before we have a neutrino flux suitable for extrapolation with the beam matrix, are to correct for the effects of neutrino cross-section, fiducial mass, and POT exposure. Fiducial mass and POT are scalar constants which apply evenly to the entire spectrum, whereas cross-section is energy and particle-dependent (it differs between ν , $\bar{\nu}$ and flavours). Given the true-energy, efficiency corrected histogram $T_j^{N,e}$, mass of the fiducial volume at the near detector m_N , the exposure of protons-on-target at the near detector p_N and the neutrino species cross-section σ_j , the near detector neutrino flux is calculated as:

$$T_j^{N,e} \rightarrow \Phi_j^N = \frac{T_j^{N,e}}{m_N p_N \sigma_j}. \quad (4.4)$$

4.2.2 Beam Matrix

The beam matrix transforms the normalised near detector flux, to the far detector flux, in a similar way to the Reco-to-true matrices. The matrix itself is calculated using kinematic information about the distribution of interactions in the NuMI target and beamline, deriving from simulation and fundamental principles. The construction of this matrix is detailed in Justin Evans' thesis [68].

4.2.3 Converting Far Detector Flux to an Energy Spectrum

Once we have our far detector flux, we proceed with the reverse of the steps that we took to convert the near detector data spectrum into a near detector flux. One exception, is that the flux is also used in calculation of the backgrounds for the purity contribution; this process is discussed in section 4.2.4.

Firstly, the cross-section, fiducial mass and POT are multiplied back in. Although both detectors are subjected to the same beam, the time periods for which each detector is active and without errors causes either detector to have a different exposure. With m_F the fiducial mass of the far detector, p_F the exposure at

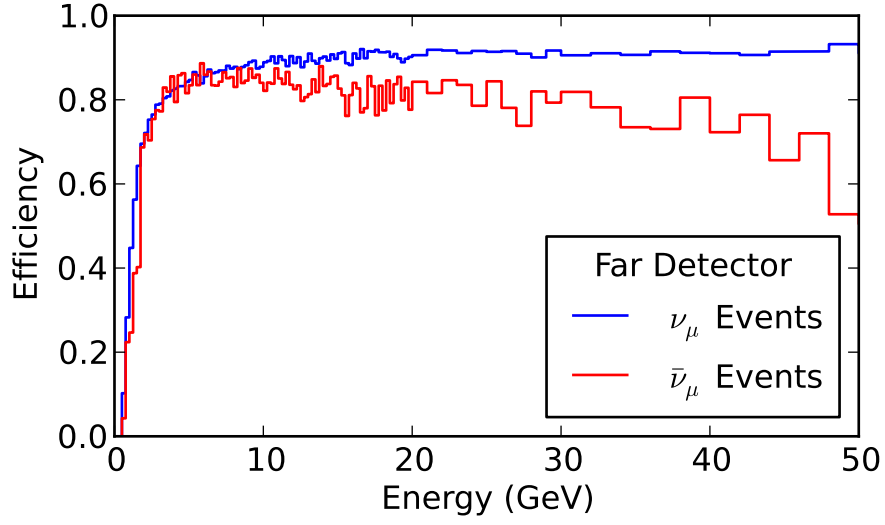


Figure 4.6: The efficiency of selection at the far detector, for both ν_μ and $\bar{\nu}_\mu$ samples.

the far detector and σ_j being the average cross-section of the energy bin under consideration, the far detector reverse of equation (4.4) is as follows:

$$\Phi_j^F \rightarrow T_j^{F,e} = \Phi_j^F m_F p_F \sigma_j.$$

The selector efficiency correction for the far detector is then made. Although the method of calculating the efficiency, E_j^F is the same as defined in equation (4.3) - using MC simulated data - here the intention is to reproduce the inefficiencies of the selection. The inefficiency is thus applied via:

$$T_j^{F,e} \rightarrow T_j^F = T_j^{F,e} E_j^F,$$

where $T_j^{F,e}$ is the complete far detector event sample, T_j^F represents the actual pure subsample selected, and E_j^F represents the efficiency that this data sample is selected at the far detector.

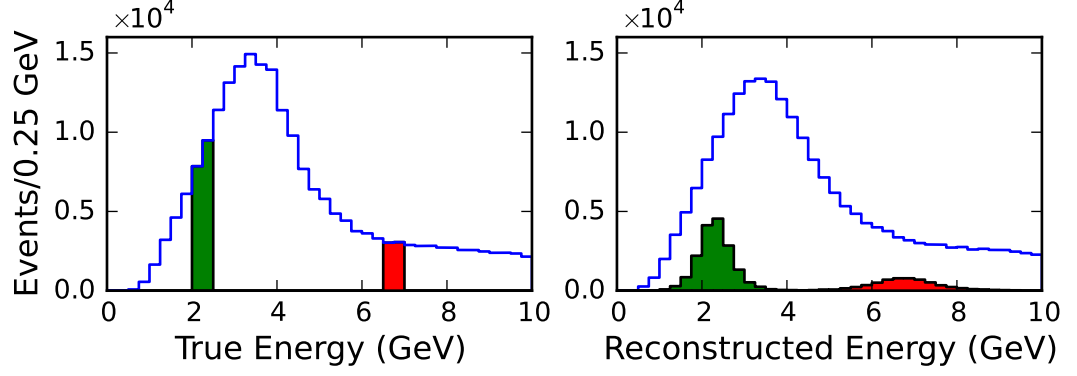


Figure 4.7: True energy of Monte Carlo far detector events (left) and the reconstructed energy of the interaction that caused these events (right).

Neutrino Oscillation

It is at this stage of the process that the effects of oscillation are introduced to the FD spectrum. The correction is made in the form:

$$T_j^F \rightarrow T_j^{F,\text{osc}} = T_j^F P_j,$$

where P_i is the oscillation probability of each energy bin. The oscillation probability is a function of energy and oscillation parameters. The oscillation parameters are input as required (for example, in the oscillation fit) to give the predicted FD spectrum with oscillation. Because this is a larger topic, detailed discussion of the calculation of this probability is deferred until section 4.2.5.

Converting true to reconstructed energy

The transformations that need to be applied to the true energy spectrum are now complete. The next step is to convert the true energy to reconstructed energy, in the reverse of the process carried out at the near detector. It is at this stage that the effects of energy resolution and detector imperfections are applied. For each set of events of a given true energy bin, the reconstruction of these events will span a range of reconstructed energy bins, as demonstrated in figure 4.7. This convolution is achieved by using FD Monte Carlo to build a matrix \mathcal{M}_{jk}^F where the j axis represents true energy, and the k axis represents reconstructed

energy. As in equation (4.2), this matrix is normalised, but in this instance along the reconstructed energy axis, so that each true energy event results in a single reconstructed energy event. The transformation is then applied as

$$T_j^{F,\text{osc}} \rightarrow R_k^{F,\text{pure}} = \sum_j T_j^{F,\text{osc}} \mathcal{M}_{jk}^F,$$

where $R_k^{F,\text{pure}}$ is the reconstructed energy spectrum, $T_j^{F,\text{osc}}$ is the oscillated true energy spectrum and the oscillation marker has been removed for clarity. At this stage the impurities have not been mixed in to the reconstructed spectrum, so the reconstructed energy spectrum is representative of a pure sample of the type of signal being extrapolated ($\nu_\mu, \bar{\nu}_\mu$). The addition of this impurity is the subject of the next section.

4.2.4 Purity Corrections: Background calculation

Unlike the near detector, where we simply reduce the measured event spectrum by the expected amounts due to background signals, in the far detector some of these signals are dependent on the oscillation parameters, and we add each background signal explicitly. The background signals to a $\bar{\nu}_\mu$ energy spectrum are: NC events wrongly reconstructed as CC, events from ν_τ -CC interactions (that appeared from $\nu_\mu \rightarrow \nu_\tau$ oscillations) and ν_μ -CC events that have been wrongly reconstructed as $\bar{\nu}_\mu$ -CC events.

The methods of calculation of each of these is shown in figure 4.8, and are described in detail in the following sections. To avoid confusion, only the process for the $\bar{\nu}_\mu$ backgrounds is explicitly discussed. The process for the calculation of the backgrounds to the ν_μ -CC sample in this method is identical, except the ν and $\bar{\nu}$ signals are swapped.

Each of the processes start from the predicted flux of neutrinos at the far detector, with the effects of POT exposure and detector fiducial mass accounted for.

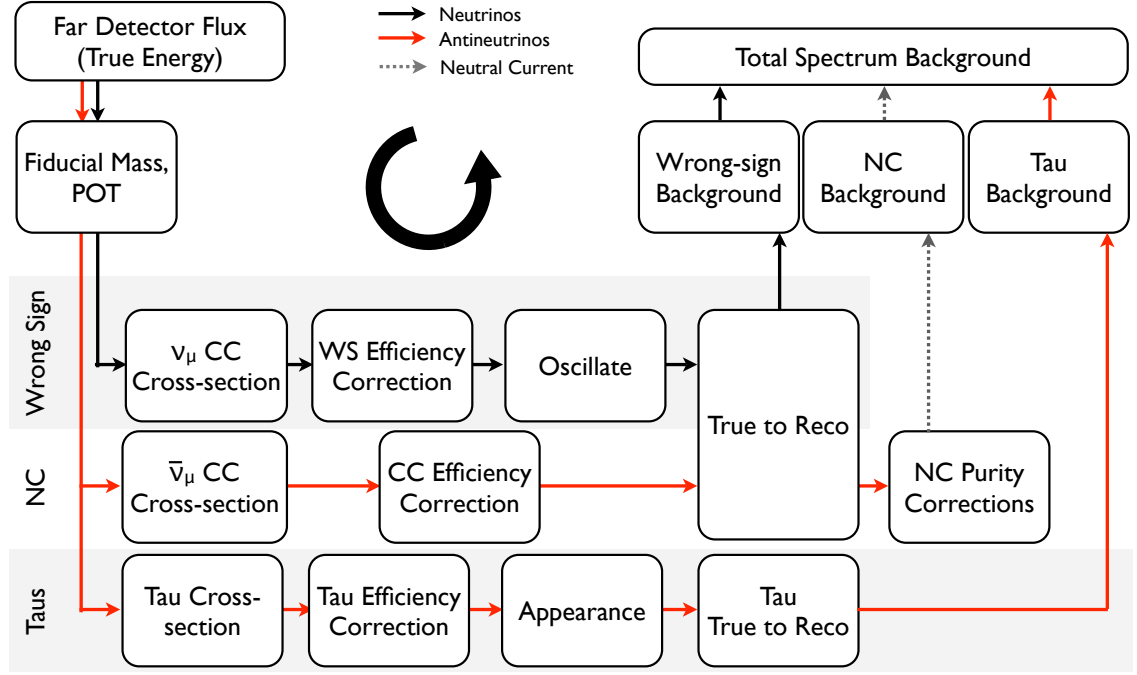


Figure 4.8: Calculation of the FD background events, for the $\bar{\nu}_\mu$ backgrounds. Both the $\bar{\nu}_\mu$ and ν_μ neutrino flux spectra are used in the calculation of the $\bar{\nu}_\mu$ backgrounds.

ν_μ -CC with Wrongly Identified μ Tracks

The first background to consider, is that of wrongly identified tracks. With the existence of a magnetic field in the far detector, most μ tracks have their charge-sign correctly identified through their curvature. A small fraction are assigned an incorrect charge. Low energy events are often too short to accurately determine curvature, and can be biased through effects such as multiple scattering. High energy tracks can be too straight for the reconstruction to reliably identify the type of muon.

Because the background to the $\bar{\nu}_\mu$ signal is due to ν_μ interactions, we must use the flux prediction for the ν_μ signal. The process followed is similar to the calculation of the ν_μ signal described in section 4.2.3. Firstly, the effects of ν_μ cross-section are applied. After this, an efficiency correction is made - but in this case the efficiency is defined as:

$$E_{ws} = \frac{\text{Number of wrong-signed signal events selected}}{\text{Total number of signal events}},$$

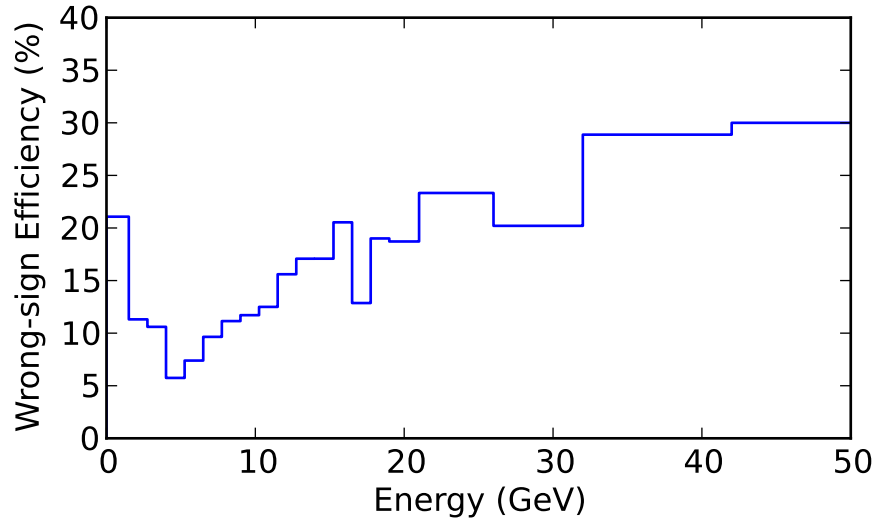


Figure 4.9: Efficiency of selecting a μ^- track as having the same curvature as a μ^+ track. Short low energy tracks, and straight high energy tracks, are the main contributors. Data from $\bar{\nu}_\mu$ selection of run I Monte Carlo events. MC statistics < 10 GeV are high enough that the errors are not visible on these scales.

where the number of wrong-signed signal events is calculated by inspecting the Monte Carlo truth for the charge of each track muon, and comparing it to the reconstructed charge. Events with a mismatch, and that pass the $\bar{\nu}_\mu$ selection, are counted. Figure 4.9 shows this efficiency for one of the data run periods, and is a typical example.

After this efficiency correction, the spectrum is oscillated with the ν_μ oscillation parameters, and converted to units of reconstructed energy through the ν_μ true-to-reconstructed matrix. The resulting energy spectrum is ready to be added to the $\bar{\nu}_\mu$ signal spectrum as a background.

Neutral Current interactions

Neutral current interactions can produce tracks that are selected as CC muon tracks by the reconstruction algorithm. Because these tracks are from shower remnants, this contamination occurs mostly at low energy.

As indicated by the flow of data in figure 4.8, this background is calculated by using the NC purity of the CC selection, calculated from the MC, after the

spectrum has been converted to reconstructed energy. The difference between the calculation of the $\bar{\nu}_\mu$ CC-signal and the NC contamination background, is that for the NC calculation the true energy spectrum is not oscillated.

Using Monte Carlo, the NC selection efficiency is calculated, as defined:

$$E_{\text{NC}} = \frac{\text{Number of NC signal events selected}}{\text{Total number of signal events}},$$

but we cannot apply this directly to the calculated $\bar{\nu}_\mu$ -CC spectrum, as it is pure by construction and does not contain any NC events. Thus, a purity correction comparable to that made in the near detector (equation (4.1)) is made, calculated using Monte Carlo as:

$$P^F = \frac{\text{Number of CC signal events selected}}{\text{Total number of events selected}}.$$

The typical content, using MC of run I data, of these purity and efficiency corrections is shown in figure 4.10. This purity correction is applied to the reconstructed energy spectrum as:

$$R_k^{F,\text{pure}} \rightarrow R_k^F = \frac{R_k^{F,\text{pure}}}{P_k^F},$$

and then the efficiency correction is made to extract the NC portion of this spectrum:

$$R_k^F \rightarrow R_k^{\text{NC}} = R_k^F E_{\text{NC},k}.$$

Tau neutrino background

The final background to include in the matrix method is events from $\bar{\nu}_\tau$ interactions. The majority of $\bar{\nu}_\mu$ that undergo neutrino oscillations are expected to reappear as $\bar{\nu}_\tau$, because $\sin^2(2\theta_{23})$ is measured to be maximal. These $\bar{\nu}_\tau$ particles can then interact in the detector. Because $\bar{\nu}_\tau$ can interact in a way that produces a μ^+ 18% of the time [95], and therefore a real, but unwanted positive track, $\bar{\nu}_\tau$ can

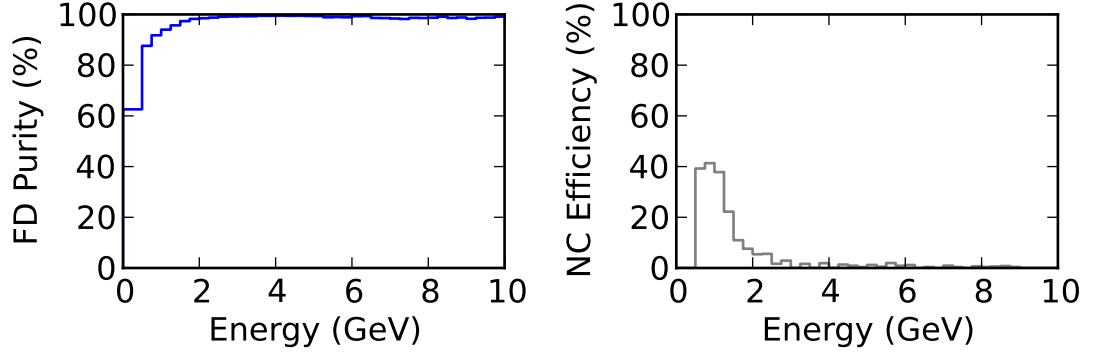


Figure 4.10: FD Purity (left) and Efficiency of falsely selecting an NC interaction shower track as a CC interaction muon (right). Because these tracks originate from shower fragments, low energy tracks are dominant. Data from $\bar{\nu}_\mu$ selection of run I Monte Carlo events.

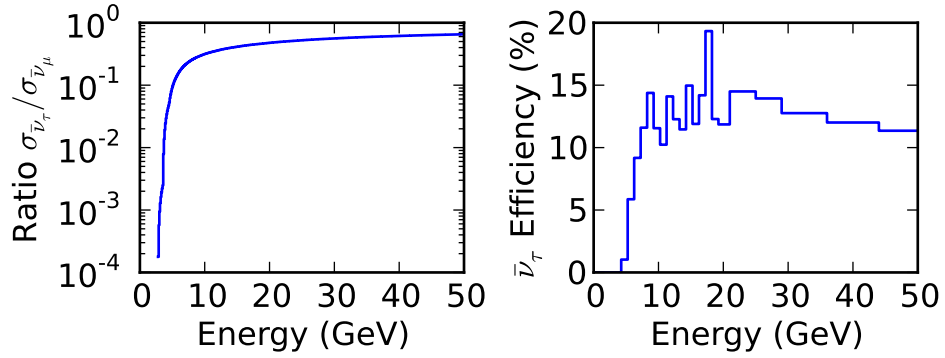


Figure 4.11: Ratio of total cross sections $\sigma_{\bar{\nu}_\tau}/\sigma_{\bar{\nu}_\mu}$ (left) and efficiency of selection of $\bar{\nu}_\tau$ events (right).

produce a hard to distinguish background. In practice, the low $\bar{\nu}_\tau$ cross-section (see figure 4.11) around the NuMI beam peak energy, and that this is a second-order effect only on events that have already undergone oscillation, mean that this background is minimal - with only an expectation of around 0.3 events at our analysis exposure level.

After application of the $\bar{\nu}_\tau$ cross-section to the FD flux, an efficiency correction is made in true energy. This efficiency is defined as:

$$E_{\bar{\tau}} = \frac{\text{Number of } \bar{\nu}_\tau \text{ events selected}}{\text{Total number of events selected}},$$

and determined using Monte Carlo. Figure 4.11 shows the $\bar{\nu}_\tau$ selection efficiency. These corrections are applied to the (exposure and fiducial mass corrected) FD

flux spectrum $\Phi_j^{F,p,m}$ as:

$$\Phi_j^{F,p,m} \rightarrow T_j^{F,\bar{\tau}} = \Phi_j^{F,p,m} \sigma_{j,\bar{\tau}} E_{\bar{\tau},j},$$

where $\sigma_{j,\bar{\tau}}$ is the $\bar{\nu}_\tau$ cross-section.

After the efficiency correction, the effects of oscillation are introduced to the $\bar{\nu}_\tau$ spectrum. Where $\bar{\nu}_\mu$ events are oscillated away according to the oscillation formula P , $\bar{\nu}_\tau$ events appear. To represent this, the $\bar{\nu}_\tau$ are ‘inverse’ oscillated, with a probability of $1 - P$:

$$T_j^{F,\bar{\tau}} \rightarrow T_j^{F,\bar{\tau},\text{osc}} = T_j^{F,\bar{\tau}} (1 - P_j),$$

where P_j is the oscillation probability of each bin, as a function of the neutrino oscillation parameters.

The $\bar{\nu}_\tau$ spectrum is then converted to reconstructed energy, using a $\bar{\nu}_\tau$ -only true-to-reconstructed matrix calculated using the Monte Carlo.

Combining the backgrounds

The final FD predicted energy spectrum is formed by summing the pure predicted CC signal $\bar{\nu}_\mu$ spectrum $R_k^{F,\text{pure}}$ with each of the backgrounds. Thus:

$$R_k^{F,\text{pure}} \rightarrow R_k^F = R_k^{F,\text{pure}} + R_k^{\text{WS}} + R_k^{\text{NC}} + R_k^{\bar{\nu}_\tau},$$

where R_k^{WS} is the wrong-sign background, R_k^{NC} the NC background, and $R_k^{\bar{\nu}_\tau}$ the $\bar{\nu}_\tau$ background.

4.2.5 Oscillation

As discussed in chapter 2, the two-flavour survival probability of a ν_μ beam can be expressed as:

$$P(\nu_\mu \rightarrow \nu_\mu) = 1 - \sin^2 2\theta_{23} \cdot \sin^2 \left(1.27 |\Delta m^2| \frac{L}{E} \right),$$

with L in km and E in GeV. When applying this probability to a set of data binned in energy, we have to apply a single probability to events covering a range of energies. In the simplest method, the energy at the centre of the bin can be used. This method is appropriate when we already know that the value of $|\Delta m^2|$ is low enough that the oscillation probability does not change significantly over a single bin, but for an unbiased search of unknown $\bar{\nu}_\mu$ oscillation parameter space, we need to be able to examine $|\Delta m^2|$ values where this is not the case, and multiple complete oscillation cycles can occur in a single bin. The error incurred through using this method with such oscillation values can be seen in figure 4.12.

During the process of attempting to solve this, a method of interpolation was developed [96] that sampled the contents of each bin in linear energy divisions, averaging the probability from each sample, and used this information to calculate a weighted average oscillation value in steps of 1×10^{-3} GeV.

In order to test the effectiveness of these methods, it is possible to calculate exactly what the correctly oscillated histogram should look like, by using individual Monte Carlo events and oscillating them according to their true energy. Figure 4.13 shows the error of the linear interpolation method, in comparison to the central-bin method, as an integral of the absolute differences between the two histograms, i.e. the error E is calculated:

$$E = \sum_j \left| \text{Approximation}_j - \text{Exact}_j \right|.$$

Unfortunately, this increased accuracy is paid for by increasing the running time of the algorithm at least 600 fold, as shown in table 4.1, and this became

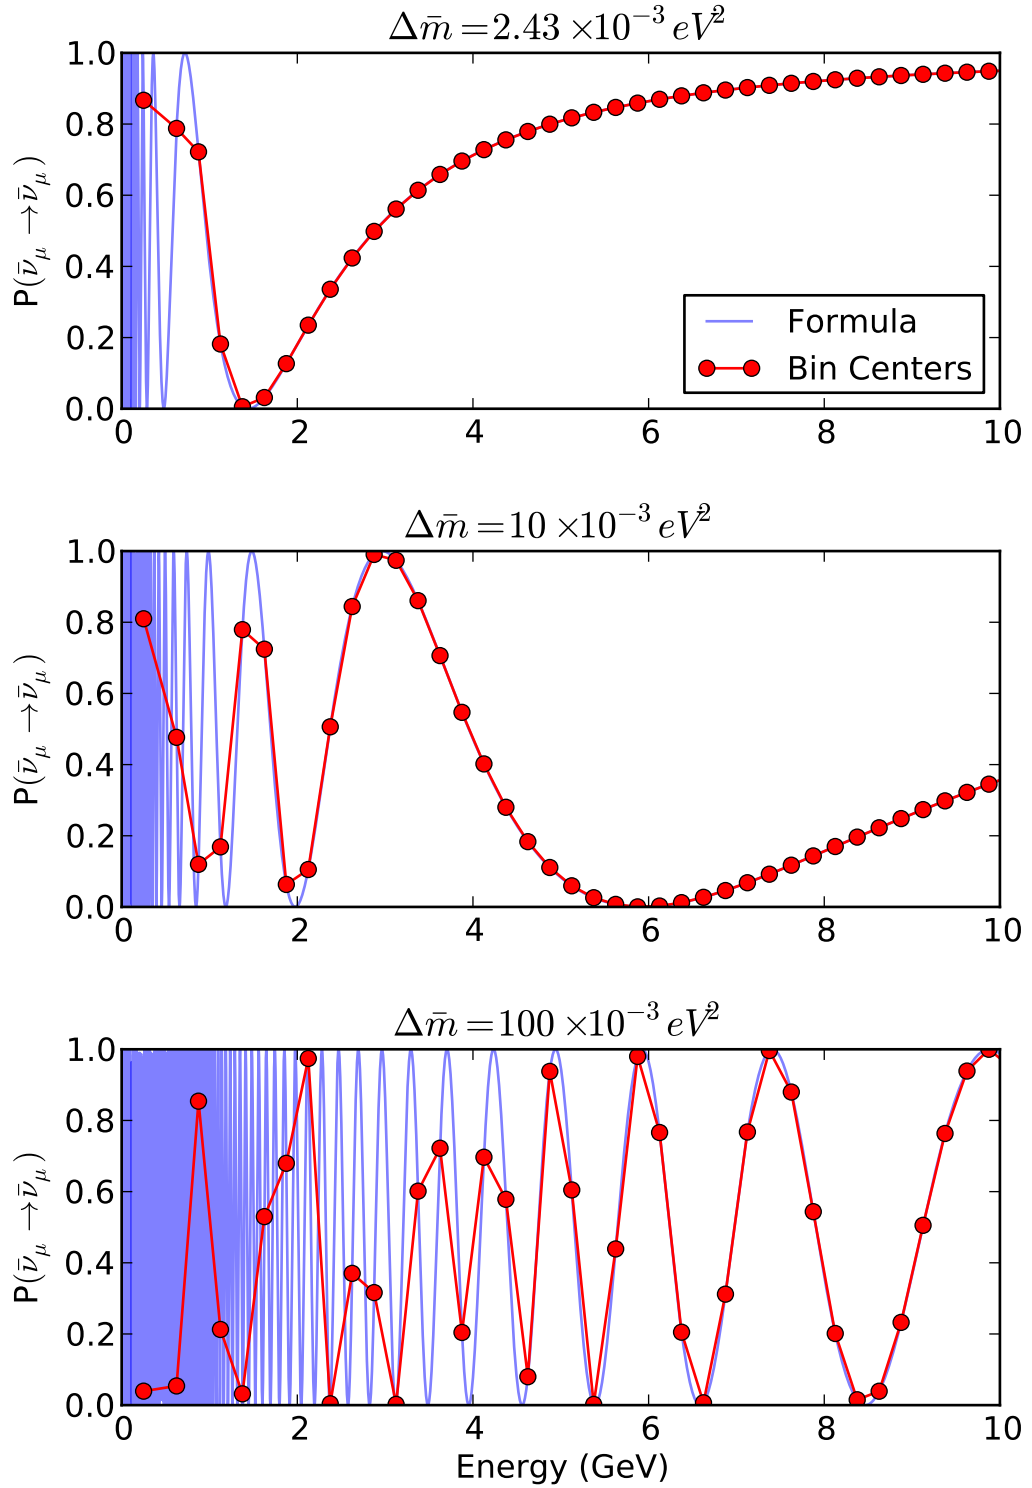


Figure 4.12: Comparison between exact evaluation of neutrino oscillation probability, and points at which bin oscillation probability is calculated using bin centres. Three values for the mass difference are shown, and $\bar{\theta} = 1.0$ for all plots.

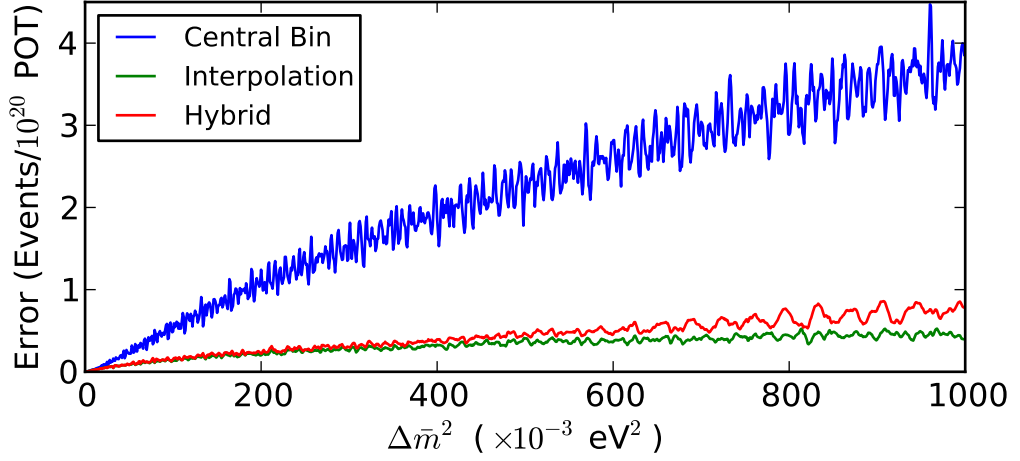


Figure 4.13: Integrated error of different oscillation methods.

$ \Delta\bar{m}^2 $	Centre	Linear	Hybrid
0	0.001	7.426	0.013
100	0.019	12.080	1.362
1000	0.021	12.122	4.332

Table 4.1: Time in seconds taken to oscillate 1000 spectra, with different methods and different values for $|\Delta\bar{m}^2|$ ($\times 10^{-3} \text{ eV}^2$)

a serious problem in the running time of the analysis. To compromise between these methods, a hybrid approach was developed for the analysis in this thesis. Using the fact that the oscillation minima and maxima are found where

$$\frac{\partial}{\partial E} P(\nu_\mu \rightarrow \nu_\mu) = 0,$$

and rearranging in terms of energy, it can be shown that the minima and maxima are found at:

$$E = \frac{2}{N\pi} 1.27 \Delta m_{32}^2 \cdot L \quad N = 1, 2, \dots, \infty.$$

Using this, the hybrid method steps sequentially over the maxima and minima, and identifies the energy at which the energy spacing drops below a threshold of 0.5 GeV, which corresponds to two bins in the energy range of interest. Bins with centres above this energy have their oscillation probabilities calculated on the bin centre alone. The lower energy bins have their oscillation probabilities calculated using the linear average method.

As shown in figure 4.13, this method matches very well with the accuracy of the linear interpolation method, and works at a much higher speed, as shown in table 4.1. One side effect of the hybridisation is that the algorithm has a $\Delta\bar{m}_{32}^2$ dependent running time, but outperforms the plain interpolation method at all oscillation parameter values.

Chapter 5

Selection

A crucial part of any neutrino analysis is selection. This is where the full set of data, including unwanted background events, is reduced to a smaller and purer set of signal events. For the analysis in this thesis, it is $\bar{\nu}_\mu$ -CC events that are of interest.

The selection of signal events can be separated into two distinct phases. The first phase of pre-selection, attempts to purge the data set of any events that do not originate from detector neutrino interactions. This is discussed in section 5.1. The second phase of selection aims to further reduce this to our desired signal interactions, and is discussed in section 5.2.

Throughout the selection process, two quantities are used to evaluate the success of the selection - the efficiency, and purity of selection. Each is estimated by applying the selection to the full set of Monte Carlo events, and comparing the selected set of events with the truth of all events. Efficiency E is defined as:

$$E = \frac{\text{Number of signal events selected}}{\text{Total number of true signal events}},$$

and is a measure of the proportion of events of interest that are included in the data set by the selection. Its complement, $1 - E$, is a measure of how many events of interest are mistakenly cut out of the data set by the selection. In contrast to

this, the purity P is defined as:

$$P = \frac{\text{Number of signal events selected}}{\text{Total number of events selected}},$$

and its complement, $1 - P$, is a measure of how many non-signal events are inadvertently let into the event sample by the selection.

Because of the much lower number of interactions at the far detector compared to the near detector, and the importance of these interactions to the measurement we are attempting to perform, the selections are optimised for selecting events at the far detector, and an identical selection is used at the near detector. The price for this simpler analysis is that the efficiency at the near detector may not be the highest possible, but this is more than compensated for by the fact that having the same selection cuts at both detectors minimises systematic errors.

5.1 Preselection

The preselection phase aims to cut out all events that are not detector neutrino interactions from the NuMI beam. We start by cutting out any data which is known to be bad - when there was a problem with the electronics, magnetic coil, or beam. This is done by comparing the events' timestamp to an automatic data quality database [97]. Additionally, any events marked as originating from the light injection calibration system are removed.

To look for muon candidates, we only consider events where the reconstruction has identified a track, so events without a track are removed. Additionally, even when a track is detected, a cut is applied on the fitting of the track, because it is possible for the reconstruction to identify a track but fail to fit it. Where multiple tracks are identified in a single event (such as tracks from particles in the hadronic shower), the longest track is used.

To eliminate events whose origin is outside of the detector, a fiducial volume cut is applied on the origin of the muon track, the *track vertex*. The most common

sources for these outside events are cosmic muons, and particles from beam neutrino interactions in the rock surrounding the detector. These rock-originating muons interact an unknown distance away from the detector and lose energy in their journey to the detector, so only a lower limit on the the energy of the neutrino can be set. In addition, particles from interactions near the edge of the detector may be lost, and so can also not be measured well.

In the far detector, the fiducial volume only includes events with vertices within a radius of 3.742 m of the centre of the detector, and additionally no closer than 0.42 m, to remove events most likely to lose energy in the uninstrumented magnetic coil. The fiducial volume runs from (and inclusive of) the 4th to the 239th plane of the first supermodule, and the 253rd to the 464th plane of the second supermodule. The gap in the fiducial volume, between the supermodules, is to reduce the possibility of contamination from external events that come in at an angle in the gap, and enter the detector without passing through the radial edge of the fiducial volume. A visual representation of the fiducial volume of both detectors can be seen in figure 5.1.

In the near detector, we require the track vertex to be within a cylinder of radius 0.8 m around the centre of the neutrino beam, and spanning the 14th to 68th planes, entirely within the fully instrumented region of the detector. The interaction rate in the near detector is high enough that the small fiducial volume has no significant impact on the result, due to the statistics.

Cosmic Muon Elimination

The largest source of non-beam tracks are from cosmic muons. We make two cuts at the far detector to keep this interference to a minimum, optimised by evaluation of backgrounds using Monte Carlo, and fake spill triggers when there is no real beam [98].

Two cuts are made to remove these events. The first is a cut on timing - any events outside of a window of time from $2\ \mu\text{s}$ before the spill to $12\ \mu\text{s}$ after the spill

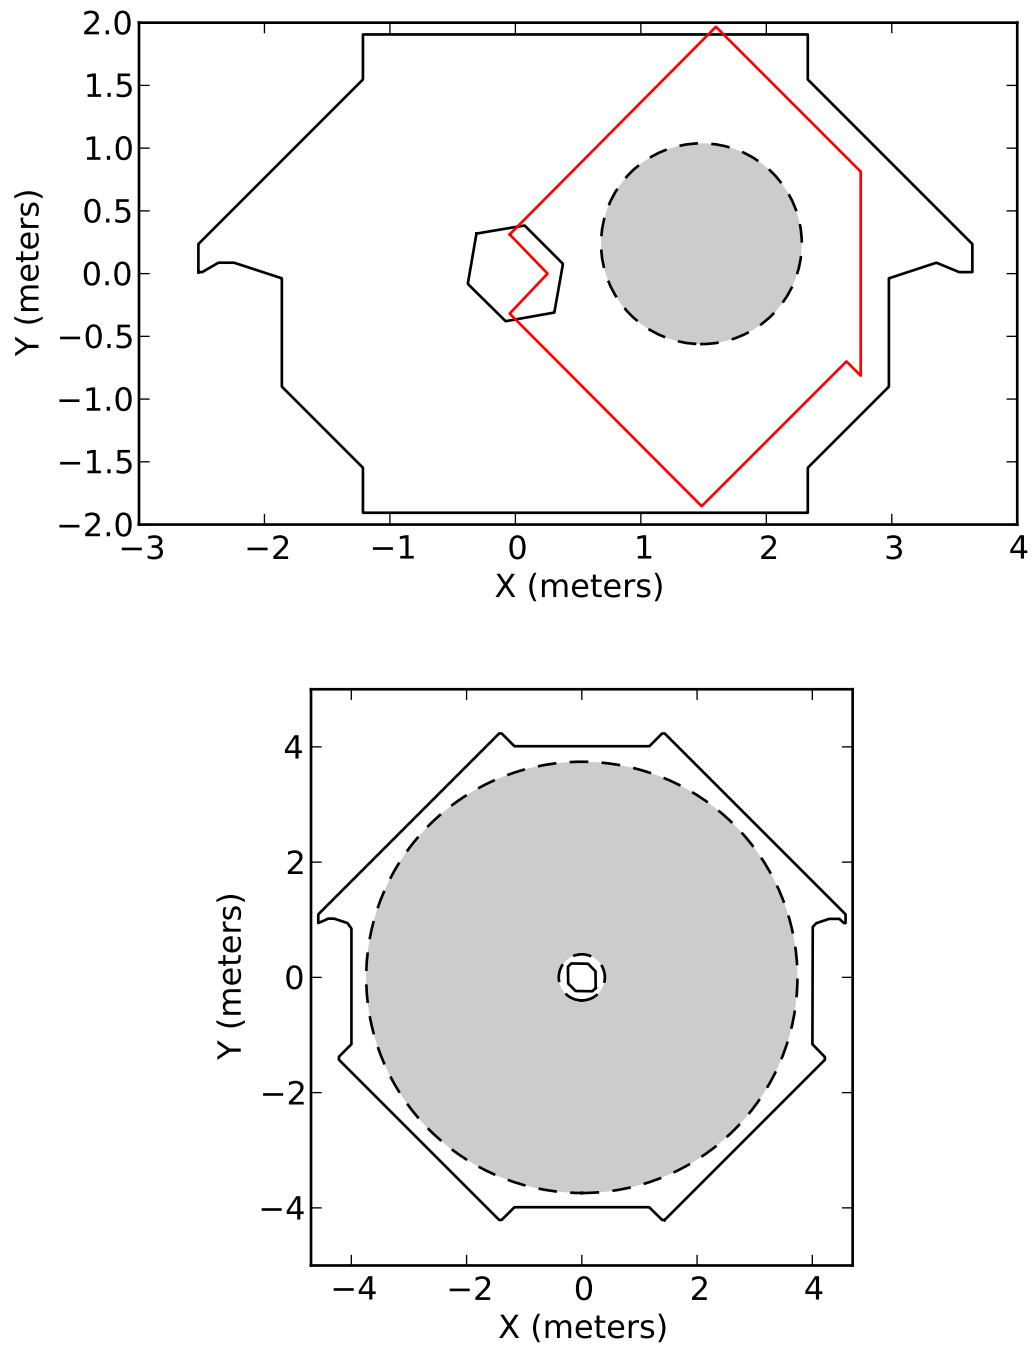


Figure 5.1: Fiducial Volumes of the ND (top) and FD (bottom), in X and Y . The detector outline and coil holes are indicated with black lines, the fiducial volumes are indicated by the grey areas, and the red line on the ND is the outline for the partially instrumented planes. Each of these outlines are for a U-plane view.

are discarded. The second is a cut on track direction - any events where the cosine of the angle between the beam and track direction is $\cos \theta > 0.6$ are ignored.

At the near detector these cuts are not necessary as spill events are directly tagged by the DAQ, the ratio of beam-originating to cosmic neutrinos is much larger, and there is more area around the fiducial volume for identification of external signals.

5.2 Signal Selection

While the job of preselection is to filter out non-beam events, a further level of discrimination is required to select the desired $\bar{\nu}_\mu$ -CC signal events from the background of NC and ν_μ -CC events. This section details the selection variables used to discriminate events, and section 5.3 details the method used to optimise the cut values used with these variables.

5.2.1 NC Discrimination

Two *Particle Identification* (PID) algorithms are used to distinguish NC events from CC events. The first is used for discrimination on the neutrino beam data, and is the PID used for [91, 94]. Three variables are used - the length of the event in planes, the fraction of the event energy in the measured track, and the average pulse height per plane in the hits attributed to the track. Each of these variables is used to populate a probability density function (PDF) for Monte Carlo truth CC and NC interactions, and a probability for each event being CC or NC in origin is derived for each data event according to:

$$P_{\text{CC,NC}} = \prod_{i=1}^3 f_i(x_i)_{\text{CC,NC}} ,$$

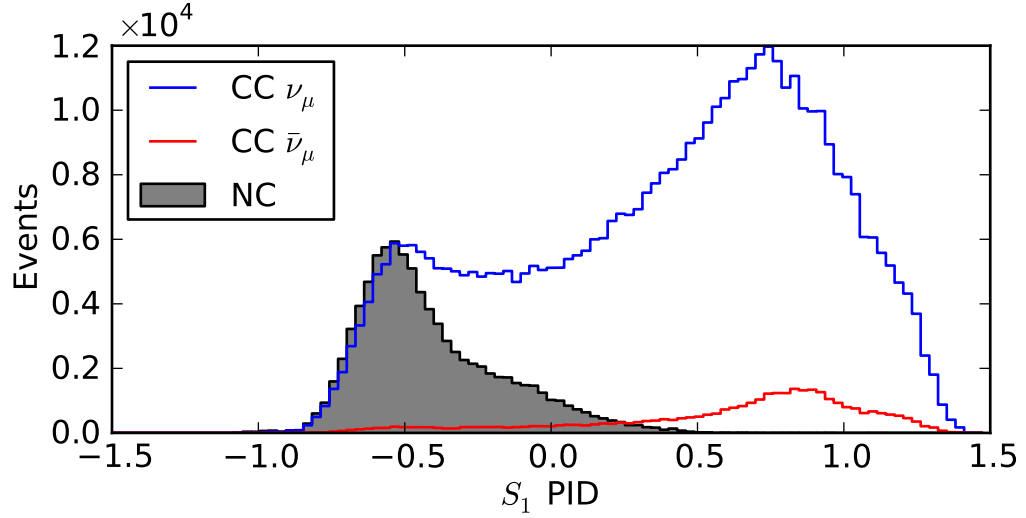


Figure 5.2: Distribution of the NC discrimination PID S_1 for Monte Carlo true NC, CC- ν_μ and CC- $\bar{\nu}_\mu$ events. An exposure of 1×10^{23} POT of run 1 far detector MC are shown.

where $f_i(x_i)$ are the individual PDF's for each variable. These two probabilities are combined into a single event selection parameter S_1 :

$$S_1 = - \left(\sqrt{-\ln(P_{CC})} - \sqrt{-\ln(P_{NC})} \right).$$

The distribution of this variable for NC and CC Monte Carlo truth events can be seen in figure 5.2 . There is a clear distinction between the majority of CC events and the NC events.

Antineutrino Beam NC discrimination

The second PID, S_2 is used for discrimination with the data set taken with a $\bar{\nu}_\mu$ -enhanced beam. The PID was first used in [92]. Four variables are used [99]: The length of the event, the average pulse height per plane of hits along the reconstructed track, the transverse energy deposition profile of the track, and a measure of the fluctuation of the energy deposited in scintillator strips along the track. A k-Nearest-Neighbour algorithm (described in depth in [100]) is used to combine an event's variables into a single multidimensional space, and find a subset k of the 80 nearest Monte Carlo events in this space. This sample can be

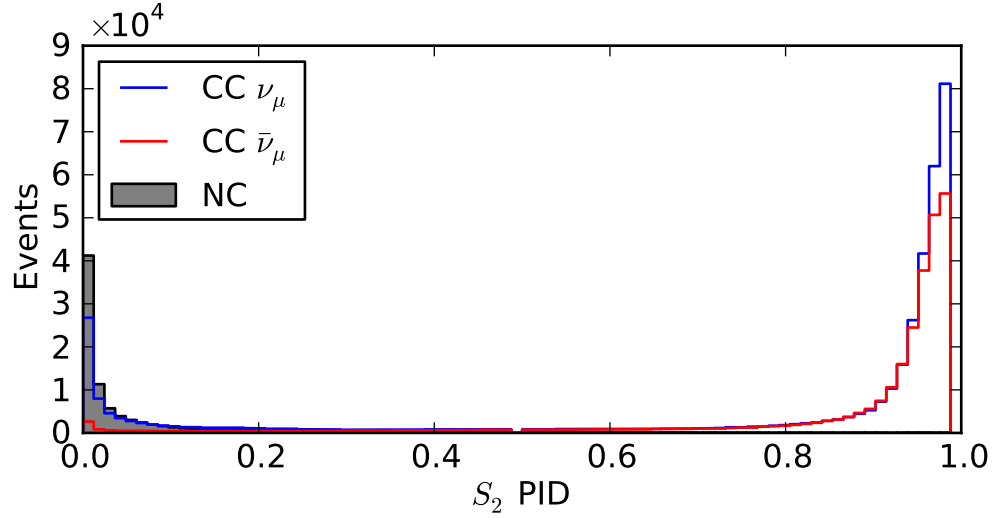


Figure 5.3: Distribution of the NC discrimination PID S_2 for Monte Carlo true NC, CC- ν_μ and CC- $\bar{\nu}_\mu$ events. Far detector MC for an exposure of 3×10^{23} POT with the $\bar{\nu}_\mu$ -enhanced beam is shown. The cut is made classifying events with $S_2 > 0.3$ as CC.

considered a sum of signal and background event sets:

$$k = k_s + k_b,$$

and so the probability of any randomly chosen event being a signal event can be expressed as:

$$S_2 = P_s = \frac{k_s}{k}.$$

This probability is used directly as the PID value S_2 , and represents the probability of an event being due to a CC interaction. A sample of the distribution of this variable for Monte Carlo can be seen in figure 5.3.

5.2.2 Charge Sign Selection

For the data taken with a ν_μ -dominant beam, this PID alone is not enough - the prevalence of ν_μ interactions provides a large background of neutrino events where the track fitter assigns an incorrect charge-sign. This situation can occur for cases where the muon Coulomb scatters off a nucleus and changes its direction, events with high inelasticity where energy from the collision results in an

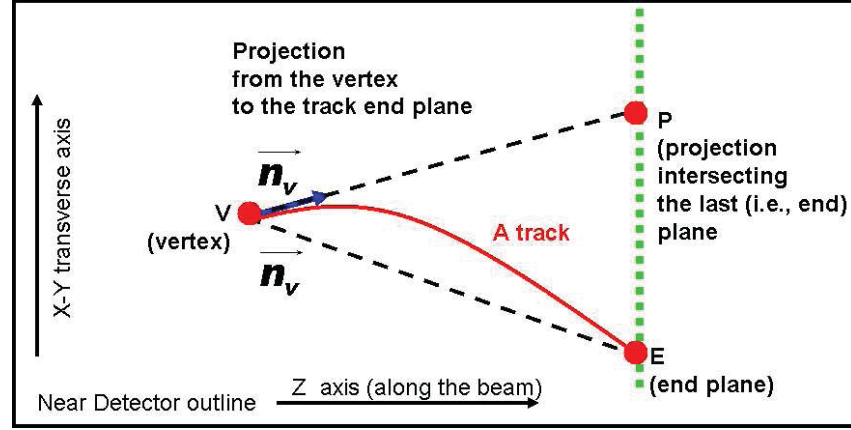
electromagnetic shower, the wrong track being selected, or plain failure of the reconstruction algorithm. To compensate for this, two additional variables are used to improve the estimation of the charge-sign of the interacting muon.

The first variable, developed in [99], is called *relative angle*. The initial direction of the reconstructed track is projected in a straight line to the plane at which the track ends (P), demonstrated in figure 5.4. The differences between this point P and the reconstructed end point of the track, E are used to construct a set of variables describing the relationship between these two points. The variable used by this analysis is the relative azimuthal angle ϕ , the polar coordinate angle of vector \mathbf{PE} in the coordinate system with P at the origin, with x aligned with the radial vector of the detector. The distribution of this variable for FD Monte Carlo events reconstructed as μ^+ tracks is shown in figure 5.5. Since we do not care which transverse direction the end point may have ended relative to the projected direction, we use the absolute offset $|\text{relativeAngle} - \pi|$. Because of the sign of the magnetic field with the ν_μ -dominant beam, μ^+ will tend to be bent outwards, away from the magnetic coil. Thus, μ^+ events will tend to have a value for the relative angle of π , as evident in figure 5.5.

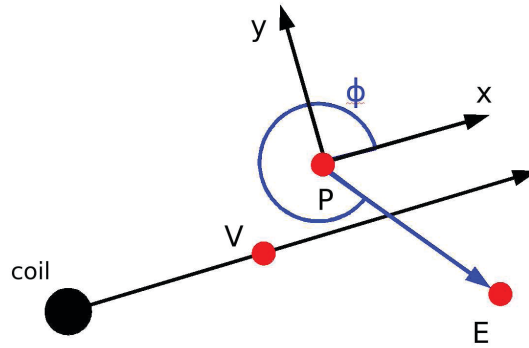
The second variable used for charge-sign selection is a measure of the certainty of the measured charge / momentum of the track,

$$\frac{q/p}{\sigma(q/p)},$$

where q is the reconstructed charge of the track, p the reconstructed momentum, and $\sigma(q/p)$ is the estimated uncertainty in this measurement. The distribution of this variable for Monte Carlo events reconstructed with μ^+ tracks can be seen in figure 5.6.



(a) The side-view geometry of how the relative angle positions are calculated



(b) End-on view of the sub-variable calculations

Figure 5.4: Geometry of the calculation of the relative angle selection variable. The calculation of the reference points is shown in (a), and the end-on view of the calculation variables is shown in (b). Diagrams from [99].

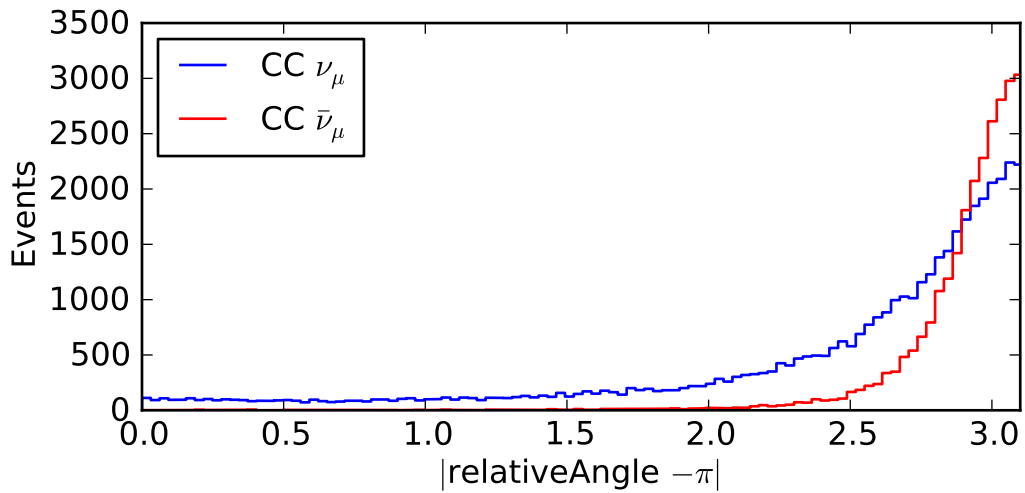


Figure 5.5: The distribution of the relative angle variable for events reconstructed as having a positive charge, used to discriminate against ν_μ events in the $\bar{\nu}_\mu$ sample. An exposure of 1×10^{23} POT of ν_μ -optimised far detector MC are shown.

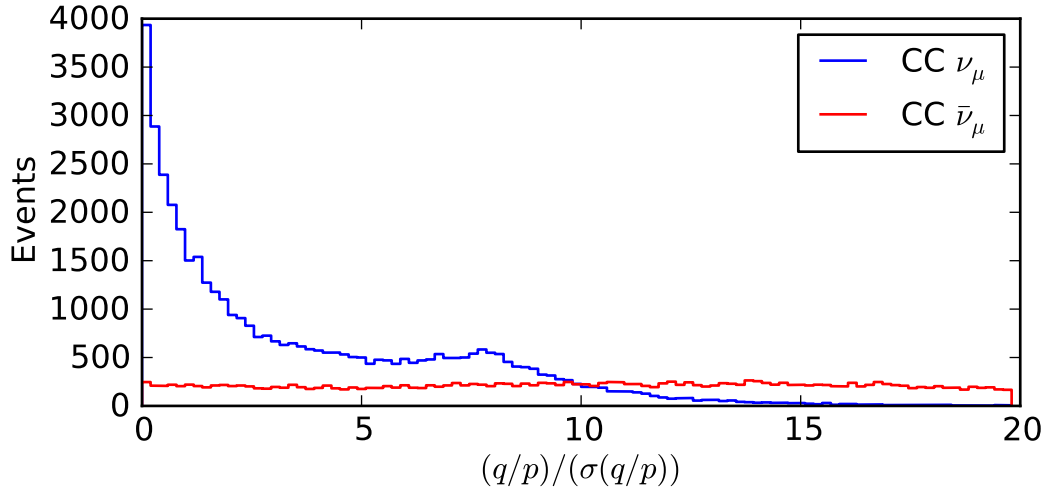


Figure 5.6: Distribution of the fitting certainty estimation for events reconstructed as having a positive charge after pre-selection cuts were applied. An exposure of 1×10^{23} POT of ν_μ -optimised far detector MC are shown.

5.3 Choosing the selection variables

To decide upon the values for our selection cuts for the first analysis of the $\bar{\nu}_\mu$ -CC events in the ν_μ -beam, potential selectors were evaluated based on several criteria [101]. First, a large number of potential selector combinations were generated manually by observing the variable distributions and liberally combining variables that appeared to be complimentary. These were then reduced down to a set of six potential selections based on simplicity and by maximizing the Efficiency \times Purity ($E \times P$) of far detector Monte Carlo event selection. The top five selections were then blindly evaluated by several teams based on efficiency, purity, and sensitivity to systematic errors [102]. The sensitivity to systematic errors was evaluated based on the shift of the best fit to fake-oscillated Monte Carlo that resulted from the $\pm 100\%$ application of each systematic. Each selector's sensitivity to all systematics was then evaluated visually.

The effectiveness of the $E \times P$ metric was also checked by cross-checking against the sensitivity of the analysis to $|\Delta \bar{m}^2|$ and $\sin^2(2\bar{\theta})$ using the different selections.

Based on the conclusions of the studies, it was decided that the following

variables and cut values would be used to select signal events in these two sets:

$$\begin{aligned}
 \text{CC/NC Discrimination: } S_1 &> 0.25, \\
 \text{Charge-sign Selection: } \frac{q}{p} / \sigma_{\frac{q}{p}} &> 3.5, \\
 |\text{relativeAngle} - \pi| &> 2.12,
 \end{aligned}$$

as the sensitivity to systematics was much lower than the other candidates, and the $E \times P$ was 0.67, for events with true energy less than 5 GeV, which did not vary very much from selector to selector.

There are two things that have changed since this analysis was performed. Firstly, the reconstruction software has advanced, changing some of the methods of calculation - this may cause small discrepancies in the reconstructed variables. Secondly, our knowledge of the ν_μ oscillation parameters has improved [103]. Because of these changes, the values for these cuts have been re-examined for this thesis.

Using the same PID and charge selection variables, a multidimensional maximisation was performed using the MINUIT2 package. Events from Monte Carlo were selected using the exact same preselection criteria as described in section 5.1, and then oscillated using the ν_μ parameters $\Delta m^2 = 2.32 \times 10^{-3} \text{ eV}^2$ and $\sin^2 2\theta = 1.0$. For each set of parameters, the $E \times P$ for events with true energy $< 5 \text{ GeV}$ was calculated. The results of this minimisation are displayed in table 5.1. A plot of the variation of E , P and $E \times P$ as the S_1 PID cut value is changed is shown in figure 5.7.

Variable		Run 1	Run 2	Run 3
S_1	$>$	0.39	0.37	0.36
$\frac{q}{p} / \sigma_{\frac{q}{p}}$	$>$	3.07	3.30	3.30
$ \text{relativeAngle} - \pi $	$>$	2.32	2.00	2.03
Efficiency		0.75	0.76	0.77
Purity		0.94	0.92	0.91
Efficiency \times Purity		0.71	0.69	0.70

Table 5.1: Optimised selection thresholds for the three ν_μ -beam data runs analysed in this thesis.

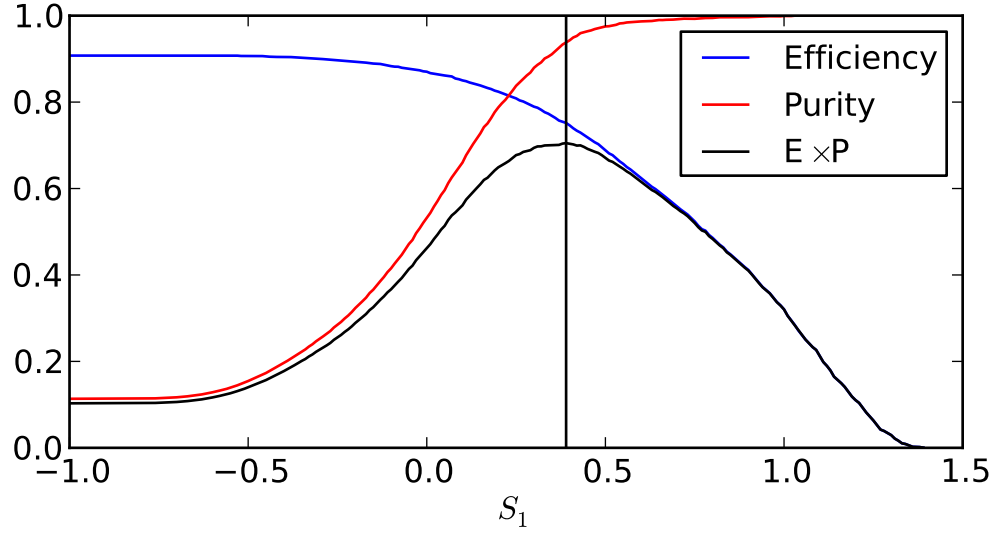


Figure 5.7: Optimisation of the PID S_1 for run 1 Monte Carlo, oscillated with $|\Delta\bar{m}^2| = 2.32 \times 10^{-3} \text{ eV}^2$. The charge selection cuts are at their optimal value.

The optimised cut values for runs 2 and 3 are very similar. Because of this similarity, and the fact that the exposure of run 3 is larger than runs 1 and 2 together, for this thesis the run 3 selection values are used for all three neutrino beam runs.

To confirm that this increase in $E \times P$ corresponds to an increase in statistical sensitivity, oscillated and scaled Monte Carlo was used to generate sensitivity contours, and the results can be seen in figure 5.8. A slight but definite improvement can be seen over all of the examined parameter space.

5.3.1 Antineutrino Beam Selection

With the data, taken with the $\bar{\nu}_\mu$ -enhanced beam, the number of antineutrino interactions in the region of interest is higher than those from neutrino interactions. This means that the additional charge sign selectors are not required to get a pure event sample. This data set has been analysed on its own [2], where for simplicity and compatibility with the neutrino analysis the only signal selection applied was:

$$S_2 > 0.3,$$

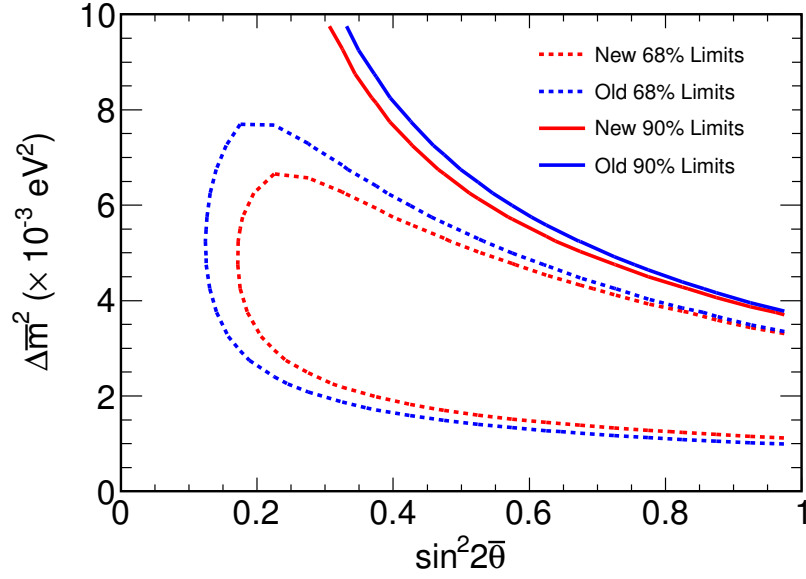


Figure 5.8: Statistical sensitivity to oscillations with $|\Delta\bar{m}^2| = 2.32 \times 10^{-3}$, $\sin^2(2\bar{\theta}) = 1.0$ for 7×10^{20} POT of run 3 MC. The old limits are using the previous values for the cuts, and the new limits are using the newly optimised values.

with S_2 being the same PID as described in 5.2.1. This choice of PID gives a base $E \times P = 0.917$, for events with true energy < 5 GeV. A similar optimisation of this selection was undertaken, both for the PID variable on its own, and in addition to the two charge-sign selection variables. The results of this optimisation can be seen in table 5.2. Neither of the additional CSS cuts gave improvement to the $E \times P$, with both additional variables cuts being optimised away.

Figure 5.9 shows the optimisation of the PID variable alone, with no additional cuts applied. The best results are found with a slightly harsher cut of

Variable		Previous	Optim.
S_2	$>$	0.30	0.40
$\frac{q}{p} / \sigma_{\frac{q}{p}}$	$>$	-	0
$ \text{relativeAngle} - \pi $	$<$	-	π
Efficiency		0.957	0.950
Purity		0.957	0.970
Efficiency \times Purity		0.917	0.921

Table 5.2: Optimised selection thresholds for the single $\bar{\nu}_\mu$ -beam data run analysed in this thesis. The previous analysis did not make any charge-sign selector cuts, and they were found to be optimal as null cuts in this optimisation.

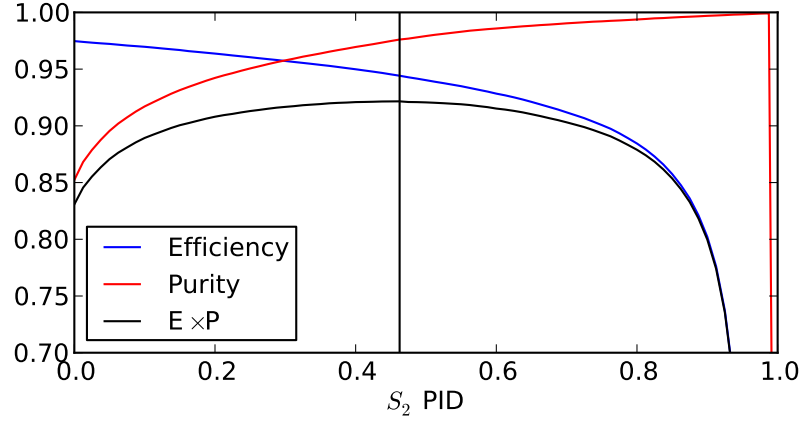


Figure 5.9: Optimisation of the PID S_2 for run 4 (antineutrino beam) Monte Carlo, oscillated with $|\Delta\bar{m}^2| = 2.32 \times 10^{-3} \text{ eV}^2$. There are no additional charge selection cuts applied.

$S_2 > 0.4$, but at a negligible improvement of $E \times P = 0.921$. Because this difference is so small, in this thesis the same value as used in the MINOS ν_μ -CC analyses of $S_2 > 0.3$ is used without alteration, for analysis of $\bar{\nu}_\mu$ -enhanced data runs.

5.4 Systematic Uncertainties

In this section, the systematic uncertainties that affect the analysis in this thesis are described, and the effects of the systematics on the best fit are evaluated. A summary of all the systematics errors considered in this analysis and the scale of their shifts can be seen in table 5.3.

A proper consideration of the effect of the systematics on the final confidence intervals will be discussed in chapter 6.

Track Energy Scale

Reconstructed tracks whose energy is determined by range have a systematic uncertainty of 2% on the measured energy, determined using CalDet. Tracks whose energy is measured using curvature have an energy uncertainty of 3% - determined by comparing the range and curvature measurements of tracks which stop

Systematic	Scale of Shift
Track Energy Scale from Range	2%
Track Energy Scale from Curvature	3%
ND Shower Energy Scale	2.0%
FD Shower Energy Scale	1.0%
Absolute Shower Energy Scale	(complex - see notes)
NC & ν_μ -CC background	50%
Near-to-Far Normalisation	1.54%
Downstream Events	+50%/-100%
Cross Sections	1σ
Flux Modelling	1σ

Table 5.3: Summary of systematic shifts considered in this analysis. Cross-section σ is relative to world data on neutrino cross sections. Flux modelling σ is relative to the error on the beam flux fitting.

within the detectors.

Relative Shower Energy Scale

The relative shower energy scale systematics are estimated by comparing data with MC simulations of cosmic ray muons [104]. By adding all the contributions from various calibration steps an uncertainty can be derived of 2.0% in the near detector and 1.0% in the far detector, with a total Near-to-Far Shower energy scale of 2.3%.

Absolute Shower Energy Scale

Two components contribute the the absolute energy scale uncertainty. The first is a 5.7% contribution from extrapolating CalDet measurements to any other detector. In addition, there is an energy-dependent contribution [105] encompassing uncertainties in hadron production and intra-nuclear effects. This systematic takes an overall energy-dependent form of:

$$\sigma_{\text{shw}} = 6.6\% + (3.5\%) \times e^{\frac{E_{\text{shw}}}{1.44\text{GeV}}}.$$

NC & ν_μ -CC background

The NC and ν_μ -CC background systematic comes from a mis-estimation of the quantities of NC and ν_μ -CC events that are wrongly classified as $\bar{\nu}_\mu$ -CC events. Two separate approaches were used to estimate the potential scales of any background systematic, for both beam configurations.

For ν_μ -dominant beam data, where the data is dominated by ν_μ events, data and MC with a large NC background, $0 < S_1 < 0.25$ (see figure 5.2) were compared. Because of the higher proportion of energy in the shower rather than the track in ν_μ samples when compared to $\bar{\nu}_\mu$ samples, due to the left-handed nature of the weak interaction, this portion of the NC background sample also contains approximately equal amounts of misidentified ν_μ background. The scale required to account for any discrepancies with data was 50%.

For $\bar{\nu}_\mu$ -enhanced beam data we are using the S_2 PID (figure 5.3) which does not give us this convenient dual measurement. Using samples of $S_2 < 0.3$ the systematic uncertainty was evaluated as 20% for NC background. Similarly, a sample with a high wrong-sign component, $\frac{q}{p} / \sigma\left(\frac{q}{p}\right) > 2.3$ was used to evaluate the ν_μ background uncertainty at 30%.

For simplicity, the analysis in this thesis uses the largest of these uncertainties - 50%, to cover all of the uncertainty in the NC and wrong-sign CC- ν_μ backgrounds.

Near-to-Far normalisation

The systematic uncertainties that contribute to a difference in the relative number of events expected at the two detectors [106] are summarised in table 5.4. The uncertainty is dominated by a 1.3% selection bias, derived by manually classifying by inspection, data and MC events in both detectors [107].

The contributions named “ND Fiducial Bias” are a reflection of the differences between simulation and data in the non-uniformity of the vertex distributions, due to the geometry of the near detector.

Systematic	Uncertainty
Steel Thickness	0.2%
Scintillator Thickness	0.2%
FD Live Time	0.32%
ND Fiducial Bias (z)	0.43%
ND Fiducial Bias (y)	0.14%
ND Fiducial Bias (x)	0.53%
N/F Selection Bias	1.3%
Total	1.54%

Table 5.4: Summary of the contributions to the relative normalisation of event numbers between the detectors. From [106].

Downstream Events

Primarily, the parents of neutrinos and antineutrinos are focussed by magnetic horns - giving ν_μ in $\bar{\nu}_\mu$ -enhanced mode, and $\bar{\nu}_\mu$ in ν_μ -enhanced mode. When in ν_μ -dominant mode, $\bar{\nu}_\mu$ parent hadrons are specifically not focussed. An important source of $\bar{\nu}_\mu$ arises from primary protons interacting in areas of the beamline other than the target - particularly, interactions with the walls of the decay pipe. These interactions can occur with protons that scatter in the target, but also secondary hadrons that come from target interactions.

Because these interactions happen closer to the near detector - over the entire length of the decay pipe, the ND/FD ratio is significantly larger than for hadrons born in the target, causing a difference in expected measurement. Figure 5.10 shows the scale of this difference between target-area and decay-pipe $\bar{\nu}_\mu$ parents.

Being approximately 7% of the beam for the ν_μ -dominant case, and concentrated near the oscillation maximum, analyses of the $\bar{\nu}_\mu$ are potentially sensitive to mismodelling of this effect, which can imitate the low energy spectrum difference between the detectors coming from oscillations.

To estimate the scale of the uncertainty [66], it was assumed that all ND MC-data discrepancy was due to this systematic, and the scale adjusted so that the number of low-energy events, both with and without other systematics, agreed. This gave a worst-case systematic of $^{+50\%}_{-100\%}$, which results in a similar effect on fitting as the other systematics, which is minor compared to the statistical error.

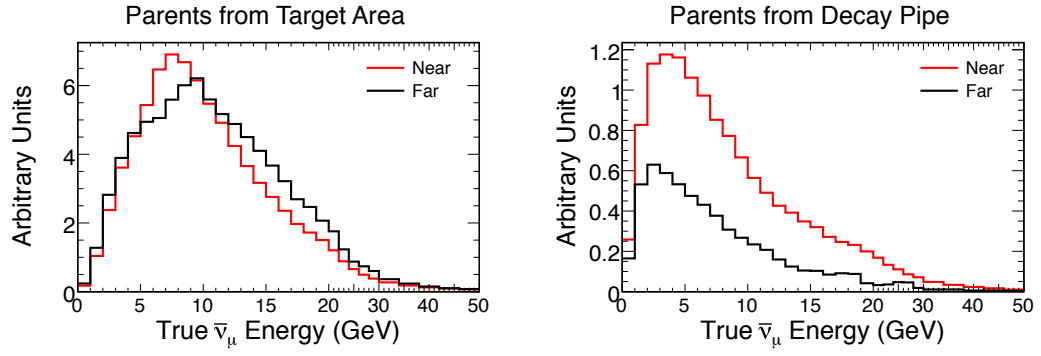


Figure 5.10: MC true $\bar{\nu}_\mu$ spectra seen at the Near (red) and Far (black) detectors, for parents produced in the target area (left) and decay pipe (right). The target area can include parents produced in the horns and other areas of the target assembly. The histograms are normalised for cross-comparison. From [66].

Cross-sections

The levels of uncertainty in the neutrino cross-sections is contributed to by the overall cross-section, but also by various NEUGEN parameters controlling the interaction models. Most of the difference that would be seen between the detectors are eliminated due to the near-identical construction of the detectors [94]. The remaining differences are due to the flux differences between the detectors causing a change in the composition of interaction types.

Flux Modelling

There are a number of contributions that can cause systematic differences in estimation of the neutrino flux. This includes uncertainties in hadron production, beam optics, the exact positioning of the target and the detectors. The fit parameters used for beam tuning were adjusted within their uncertainties to give an overall estimate on the error in the flux. This process is described in [94, p.19].

5.5 Actual effect of systematics

The estimated effect of these systematics on the end results can be assessed in a couple of different ways - separate from the Feldman-Cousins approach discussed in chapter 6. The first is by estimating the error introduced on the predicted data spectra, shown for both ν_μ -dominant and $\bar{\nu}_\mu$ -enhanced beam configurations in figure 5.11. For each systematic, sample ND and FD data was generated by applying the systematic positively, and negatively, to MC events - to simulate data taking under the $\pm 1\sigma$ scenarios. Oscillation was also applied to the FD fake data at a value of $|\Delta m^2| = |\Delta \bar{m}^2| = 2.43 \times 10^{-3} \text{eV}^2$, $\sin^2(2\theta) = \sin^2(2\bar{\theta}) = 1.0$.

The ND systematically shifted spectra were then extrapolated using this known oscillation value to generate a FD prediction, and then compared to the energy spectrum built from systematically shifting FD MC events directly. This gives a histogram describing the discrepancy, under this systematic, of the extrapolation procedure. Finally, any difference incurred under the nominal scenario (e.g. the extrapolation error without any systematics applied, $\sim 1\%$ max), is removed.

This gives the calculation:

$$\Sigma = \frac{\left(\frac{E_{\text{prediction}}}{E_{\text{summary}}} \right)}{\left(\frac{N_{\text{prediction}}}{N_{\text{summary}}} \right)},$$

where Σ is a spectrum of the fractional error introduced for a particular direction, of a specific systematic, E is the energy spectrum from applying that systematic shift, N is the energy spectrum from the nominal, no-systematics case, *prediction* indicates the spectrum is constructed from extrapolating ND MC, and *summary* indicates the spectrum is constructed by shifting FD MC events and then building a spectrum of the result.

Once these errors have been calculated, they are split into positive and negative effects, and summed in quadrature to give the final fractional error bands shown in figure 5.11.

The effects of systematics from ν_μ -dominant and $\bar{\nu}_\mu$ -enhanced are shown sep-

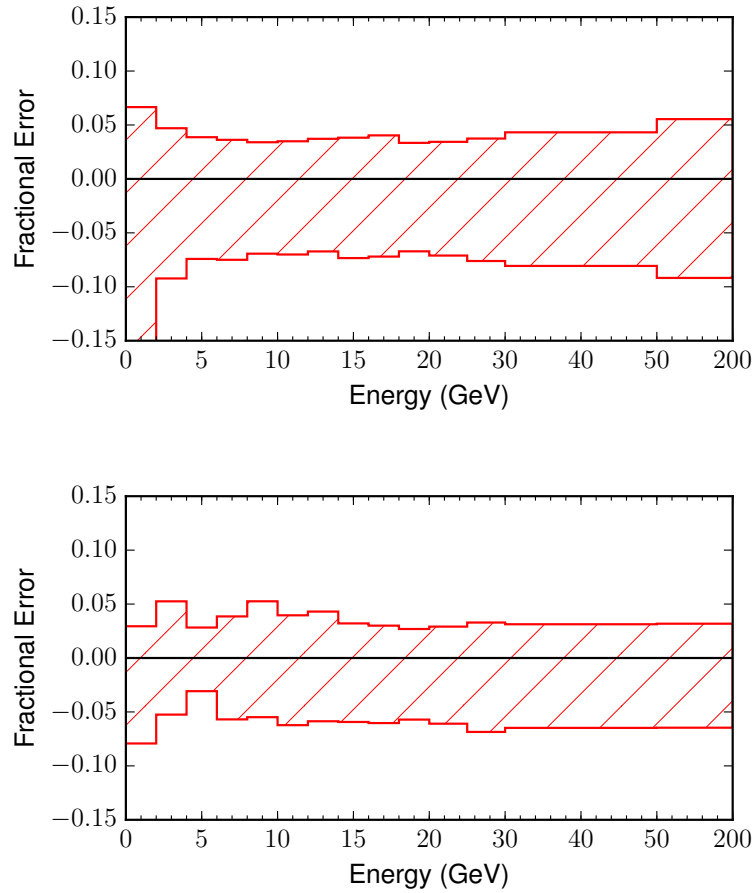


Figure 5.11: Systematic error bands on the predicted FD energy spectrum for ν_μ -dominant (top) and $\bar{\nu}_\mu$ -enhanced (bottom) beam configurations. Data generated by comparing the quadrature-added effect after extrapolation, of applying each systematic.

arately, as the beam spectra is qualitatively different between them and is how the actual data will be studied - separated by beam configuration. The dip at the lowest energy bin in ν_μ -dominant mode is due to the measured track energy from range systematic moving events into the lowest extrapolation bins, where the modelling is poor. The asymmetrical bias towards negative errors is due to the asymmetry of the systematic describing neutrino parents originating from within the decay pipe.

The second method employed to estimate the systematic errors is by observing the effect of applying systematics to the MC, on the best fit oscillation parameters to the energy spectra built with this shifted MC.

For each systematic under consideration, shifted ND and FD data were generated, as above, for $\bar{\nu}_\mu$ -enhanced and ν_μ -dominant configurations. These spectra were then scaled to the expected POT of the real data, and then fit simultaneously, in the same manner as the final data will be. The difference of the results of this fit, compared to the nominal case where there are no systematics applied, can be used as a visual guide to the scale of the effect of each systematic. This is because we have far more MC events than data, such that statistical effects become negligible. The results from these fits can be seen in figure 5.12.

When fitting, the best fit was allowed to move into unphysical territory, as it is not the absolute value of each fit that is of interest, but the difference from the nominal, unshifted scenario.

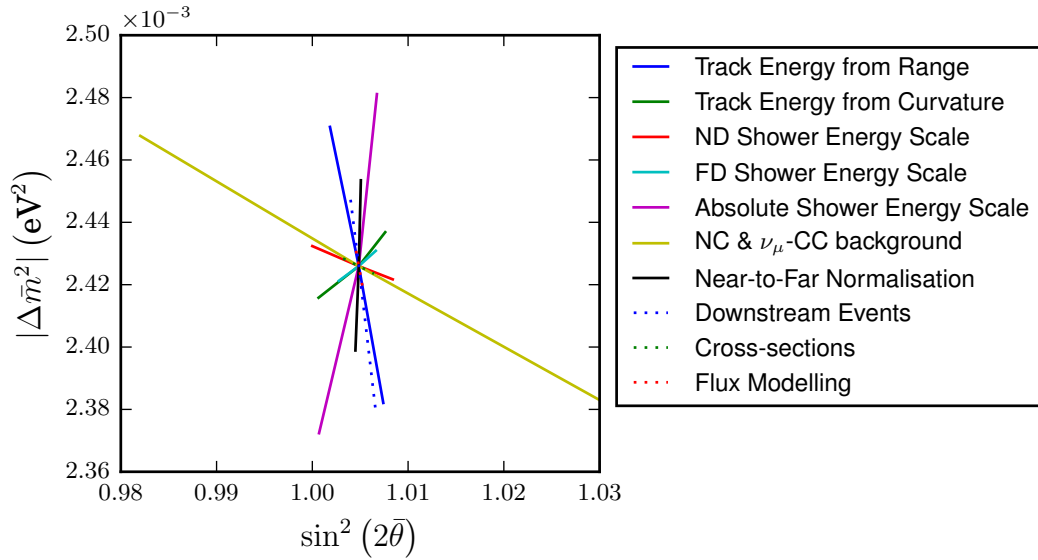


Figure 5.12: Effect of applying each systematic to the best fit results of scaled high-statistics MC, for the full, combined ν_μ -dominant + $\bar{\nu}_\mu$ -enhanced fit. For each \pm value of each systematic, the systematic is applied to MC events, and the resultant spectra are fitted as though they were data. A line is then drawn from each subsequent best fit to that obtained by fitting nominal, unshifted MC.

For better understanding of the scale of these systematics, when compared to the statistical uncertainty, the results in figure 5.12 have been superimposed on the Gaussian sensitivity contours for the combined results. This can be seen in figure 5.13. The nominal central point has been shifted to the true MC oscillation

value, for illustration. It is clear that for this analysis, statistics are dominant.

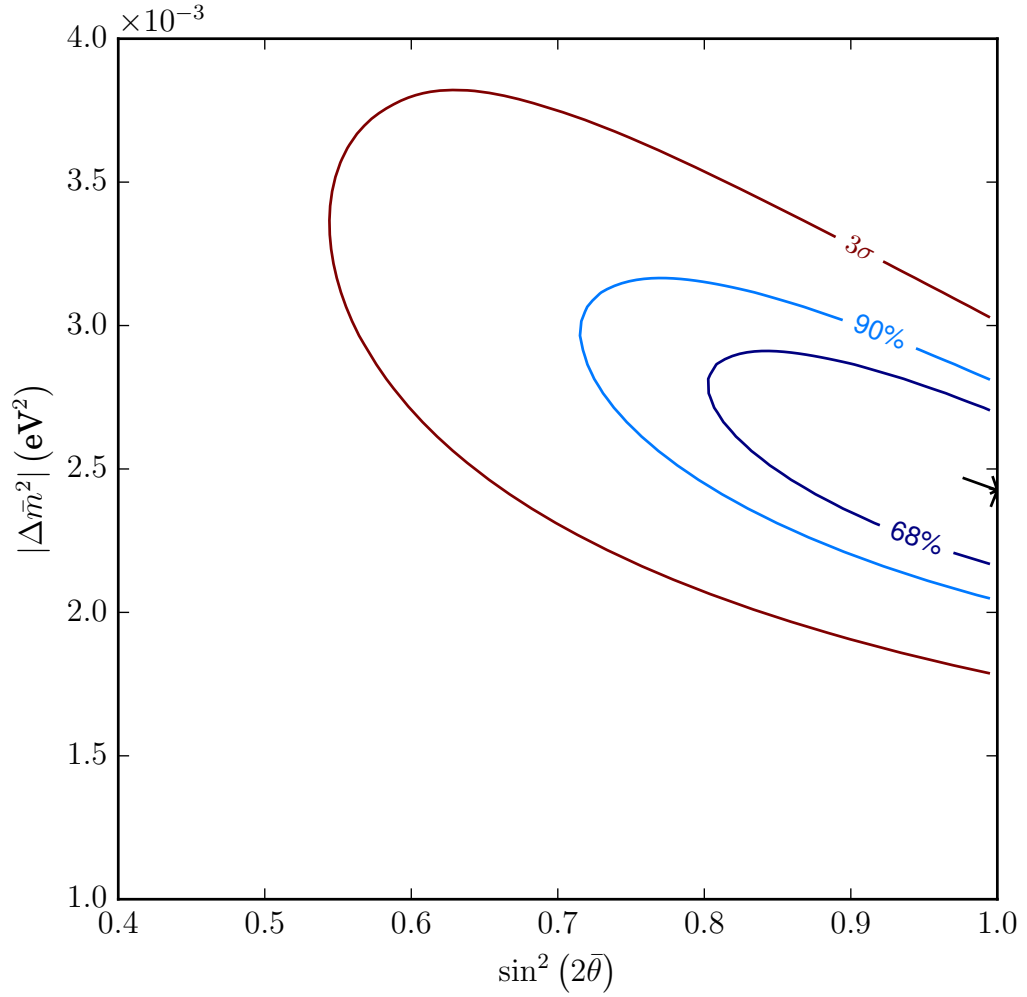


Figure 5.13: Effect of the systematic shifts on the best fit, compared with the Gaussian sensitivity contours.

5.6 Data Validation

Before fitting and final analysis, the data and MC needs to be checked for compatibility. This is done for the selection variables used in this thesis in figures 5.14 and 5.15 for ν_μ -dominant and $\bar{\nu}_\mu$ -enhanced beam modes respectively.

The near detector plots show the systematic error band on the MC for the near detector, from uncertainties on flux, neutrino cross-section and contributions from the decay pipe. The far detector plots show the (dominant) statistical error

for the selection variables, where the MC has been processed assuming oscillation parameters of $|\Delta\bar{m}^2| = 2.4 \times 10^{-3} \text{ eV}^2$ and $\sin^2(2\bar{\theta}) = 1.0$.

All the rest of the variables were validated as part of the MINOS results validation, see [106].

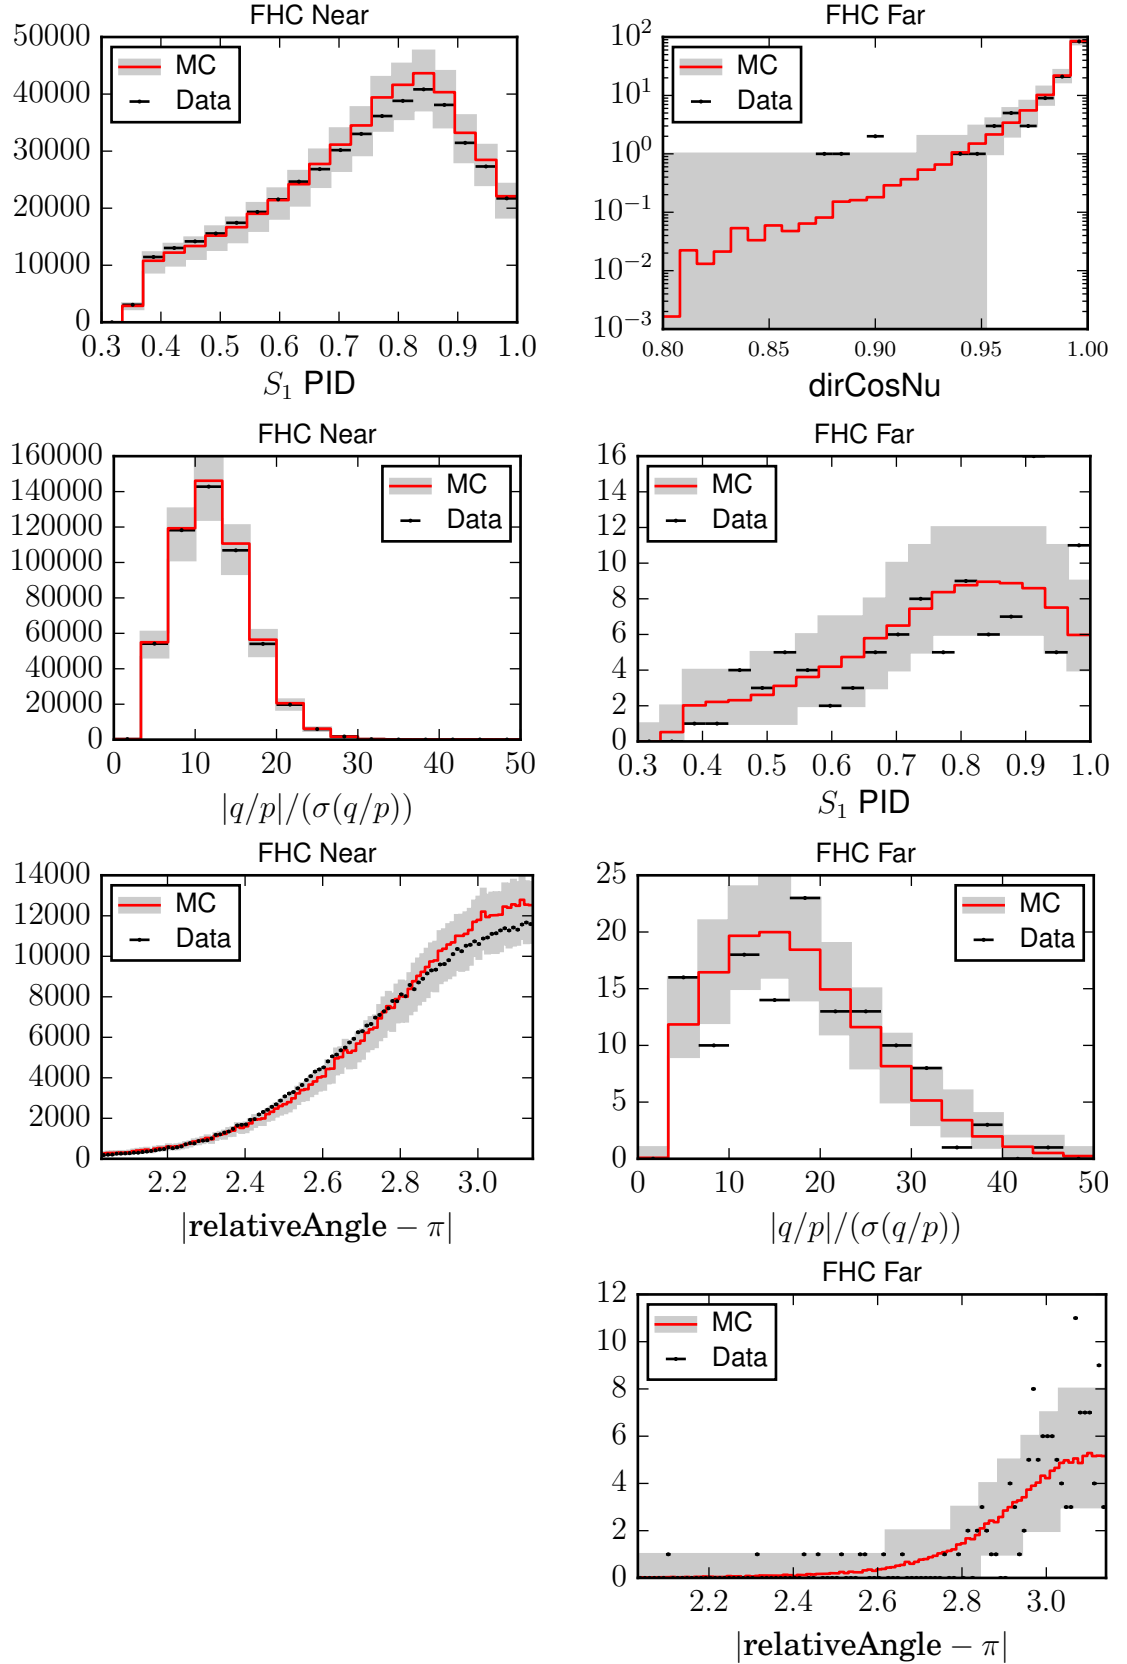


Figure 5.14: Data/MC comparison plots for ν_μ -dominant beam near detector variables (left) and far detector variables (right). The Systematic error bands are shown on the near detector plots, and Poisson statistical errors on the far detector plots, which have been generated with a $|\Delta\bar{m}^2| = 2.4 \times 10^{-3} \text{ eV}^2$ at maximal mixing.

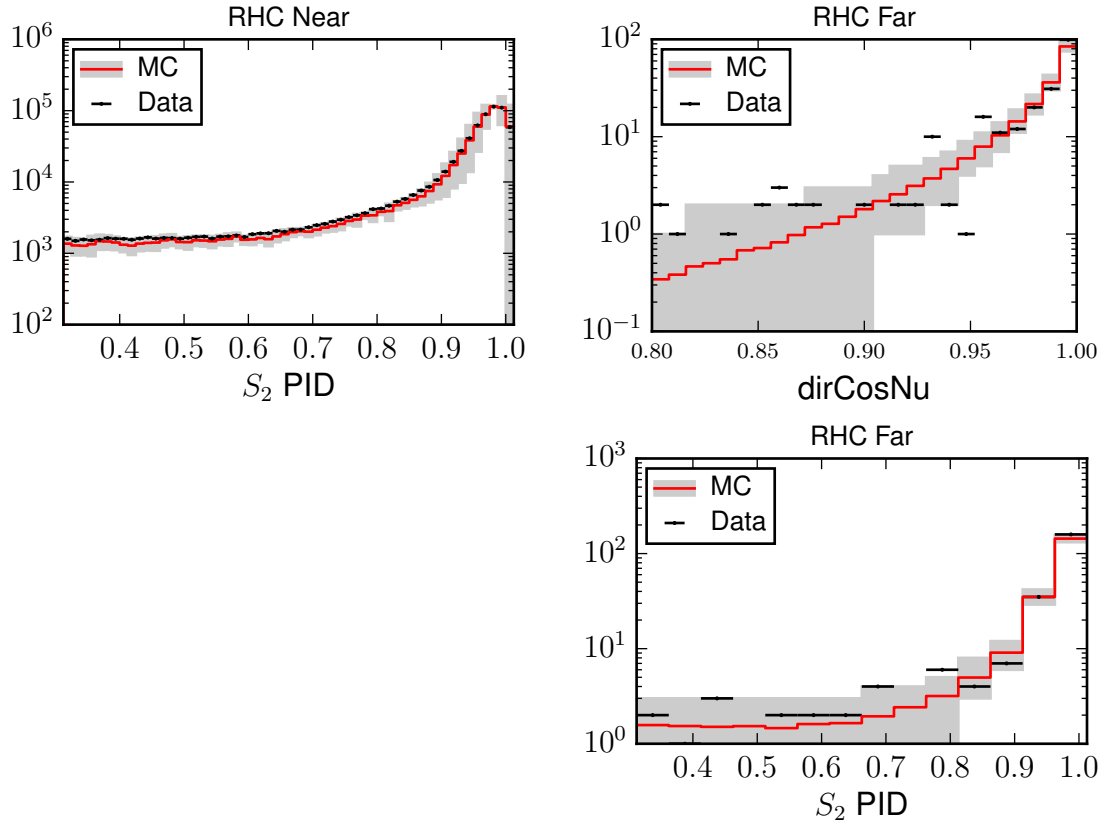


Figure 5.15: Data/MC comparison plots for $\bar{\nu}_\mu$ -enhanced beam near detector variables (left) and far detector variables (right). The Systematic error bands are shown on the near detector plots, and poisson statistical errors on the far detector plots, which have been generated with a $|\Delta\bar{m}^2| = 2.4 \times 10^{-3} \text{ eV}^2$ at maximal mixing.

Chapter 6

Application of the Feldman-Cousins Method

A crucial part of calculating the results of any experiment is the computation of the confidence intervals, which allows the reporting of errors on the results of experiments, attempting to take into account the known and unknown variations that affect every experiment.

One common approach, used in previous MINOS analyses, is directly based on the likelihood-ratio method. For a single set of measured events, a likelihood ratio can be calculated using the ratio of measured (N) to expected values (N_{expect}), of a single bin in the energy spectrum. Assuming that the distributions of errors for the model parameters are Gaussian, we can calculate a χ^2 value summed over all bins i :

$$\chi^2 = 2 \sum_i N_{\text{expect},i} - N_i + N_i \ln \left(\frac{N_i}{N_{\text{expect},i}} \right), \quad (6.1)$$

the form of which comes from [108]. By minimising this measure, we are maximizing the likelihood over parameter space, and so a best fit likelihood χ_{best}^2 can be found. Again using the assumption that the distribution of the errors in the measured parameters is Gaussian, the difference between this maximal likeli-

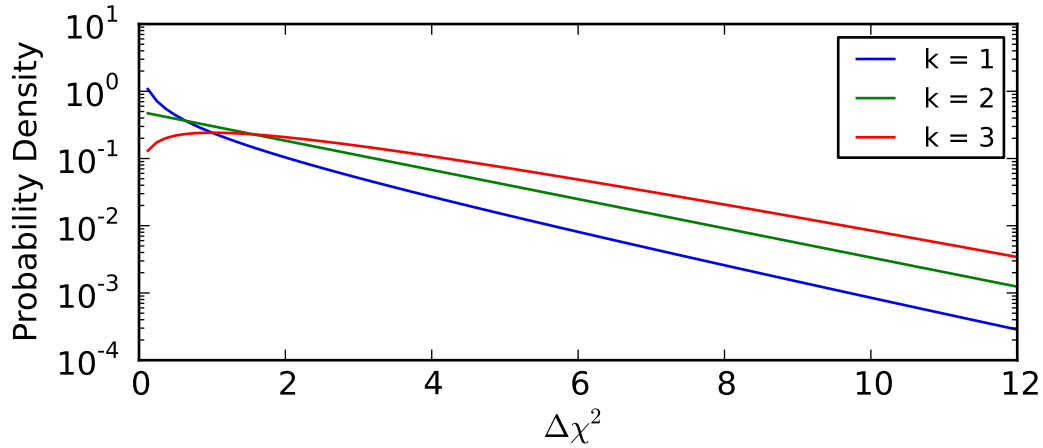


Figure 6.1: Probability density of the $\Delta\chi^2$ distribution, for $k = 1, 2, 3$ degrees of freedom.

hood and that from any other set of parameters:

$$\Delta\chi^2 = \chi^2 - \chi_{\text{best}}^2,$$

is distributed according to the k -parameter chi-squared distribution (shown in figure 6.1 for $k = 1, 2, 3$). The properties of this distribution are well understood, such that we can calculate the values of $\Delta\chi^2$ which correspond to certain integrated probability values (see table 6.1).

By constructing a k -dimensional surface of $\Delta\chi^2$ values over each of the fitted parameters, we can trace the contours at these pre-defined values in order to calculate the confidence intervals. For example, when measuring neutrino oscillation parameters, we are interested in fitting two parameters tracing the 2D surface - the mixing angle $\sin^2(2\theta)$ and the mass difference $|\Delta m^2|$. After finding the best fit, we can construct the $\Delta\chi^2$ surface and trace the contour at $\Delta\chi^2 = 11.83$ to find the 3σ confidence limits.

However, this simple process breaks down if the core assumption of Gaussian-distributed errors no longer holds. This can happen for several reasons, but in this analysis there are two primary causes.

The first, is that of low statistics - because of the low number of antineutrinos being produced (compared to neutrinos, for which the experiment was de-

Cum. Probability (%)	k = 1	k = 2
68.27	1.00	2.30
90.	2.71	4.61
99.	6.63	9.21
99.73	9.00	11.83

Table 6.1: Value of the $\Delta\chi^2$ distribution for which the integrated probability is certain, useful values. Values from [95]

signed) the number of events being selected for analysis is small, and the number of events in each bin is very often expected to be < 3 . This puts the quantities firmly in the realm of Poisson, not Gaussian statistics.

This effect is small compared to the second reason that the parameter errors are non-Gaussian; that of physical boundaries. Consider the case of true oscillation parameters near maximal mixing. Under the assumption of Gaussian errors, the confidence intervals would encompass unphysical areas of parameter space (that of $\sin^2(2\bar{\theta}) > 1$), a measurement of which would be meaningless. But enforcing this physical boundary changes the available parameter space, so the mixing angle would no longer be free to vary according to Gaussian statistics.

The Feldman-Cousins method [109] provides a way to reconcile this conflict. It proposes a method for calculating the precise confidence intervals through an ordering method, based on the likelihood ratios. Effectively, the method advocates generating a large number of randomly fluctuated experiments, and using the likelihood-ratio method as a measure whose distribution is sampled empirically.

For the analysis in this thesis, the Feldman-Cousins method has been implemented to accurately estimate the confidence interval results of the oscillation analysis. This work is described in section 6.1. The method has also been applied to the calculation of the effects of known systematic errors on the contour. This is described in section 6.2. In addition, the FC software and techniques I developed were used for MINOS publications [2] and [1].

6.1 Statistical Corrections

In this section, the actual process used to calculate the FC-corrections, and the method of applying these corrections will be summarised. One key part of the process not discussed here, is the actual method used to generate each random experiment. This is the part that required most adaption for the inclusion of systematic errors, and is detailed in section 6.2.

The starting point for the FC method is a set of randomly generated experiments with a single, known set of oscillation parameters $|\Delta\bar{m}^2|_{\text{truth}}, \sin^2(2\bar{\theta})_{\text{truth}}$. Each of these experiments consists of a near and far detector spectrum, though because of the level of statistics available at the near detector, the spectrum does not need to be treated in the same way as the far - . The far detector spectrum, containing a much more statistics-limited measurement, has to be sampled as if it were actually measured from real events. Practically, this means integer event counts, which can be achieved by Poisson-fluctuating each POT-scaled Monte Carlo bin of the extrapolated prediction. The near detector sample is large enough that having integer event-counts is not important.

Once each experiment has been generated, it can be fitted, to find the set of oscillation parameters $T_{\text{best}} = |\Delta\bar{m}^2|_{\text{fit}}, \sin^2(2\bar{\theta})_{\text{fit}}$ which gives best agreement between oscillated Monte Carlo and the individual fake experiments. This can then be used to calculate the ordering measure, for which we use the $\Delta\chi^2$:

$$\Delta\chi_{\text{fit}}^2 = \chi^2(|\Delta\bar{m}^2|_{\text{truth}}, \sin^2(2\bar{\theta})_{\text{truth}}) - \chi^2(|\Delta\bar{m}^2|_{\text{fit}}, \sin^2(2\bar{\theta})_{\text{fit}}).$$

After doing this for every generated experiment, we can use this set of $\Delta\chi^2$ values to calculate the value of $\Delta\chi_{\text{fit}}^2$ that encompasses each probability that we are interested in. This is best demonstrated in figure 6.2, where the distribution of these $\Delta\chi_{\text{fit}}^2$ for a sampled set of experiments with a single pair of true values is compared with the distribution that would be expected for the Gaussian case. In this plot, generated at maximal mixing, the the physical boundary causes fit

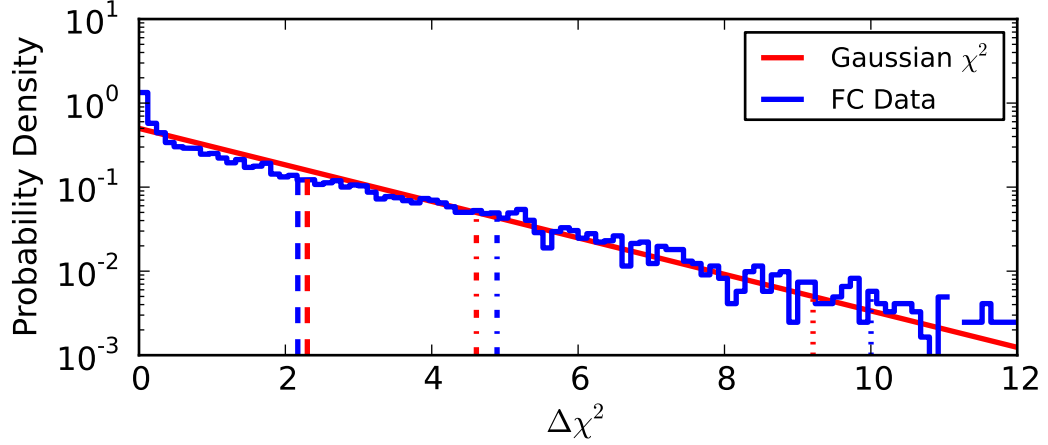


Figure 6.2: Sample $\Delta\chi_{\text{fit}}^2$ distribution from the FC method for 7×10^{20} POT of ν_μ -dominant beam Monte Carlo, oscillated at $|\Delta\bar{m}^2| = 2.32 \times 10^{-3} \text{ eV}^2$, $\sin^2(2\bar{\theta}) = 1.0$. The vertical markers are at cumulated probability of 68%, 90%, 99%.

results to ‘pile up’ against $\sin^2(2\bar{\theta}) = 1$, causing an excess of fits having a very low $\Delta\chi_{\text{fit}}^2$ (compared to the Gaussian case).

The number of experiments to run this process on is defined by the level of probability that you want to probe accurately; the number of experiments required to accurately find the $\Delta\chi_c^2$ sampling 68% of the distribution is far less than that required to find the 99.73% (3σ) sample. In this thesis, the process has been run until the statistical fluctuations between differing sets of true parameters is reduced to acceptable levels, such that the measured $\Delta\chi_c^2$ for a particular probability appear relatively smooth across multiple true values.

6.1.1 Applying the correction

Once the distribution of $\Delta\chi_{\text{fit}}^2$ values has been sampled, we can use this to calculate correct confidence intervals, without over- or under-coverage. After analysing the data we wish to calculate the contours for, including the finding of a best fit, we can construct the usual surface of

$$\Delta\chi^2(|\Delta\bar{m}^2|, \sin^2(2\bar{\theta})) = \chi^2(|\Delta\bar{m}^2|, \sin^2(2\bar{\theta})) - \chi_{\text{best}}^2$$

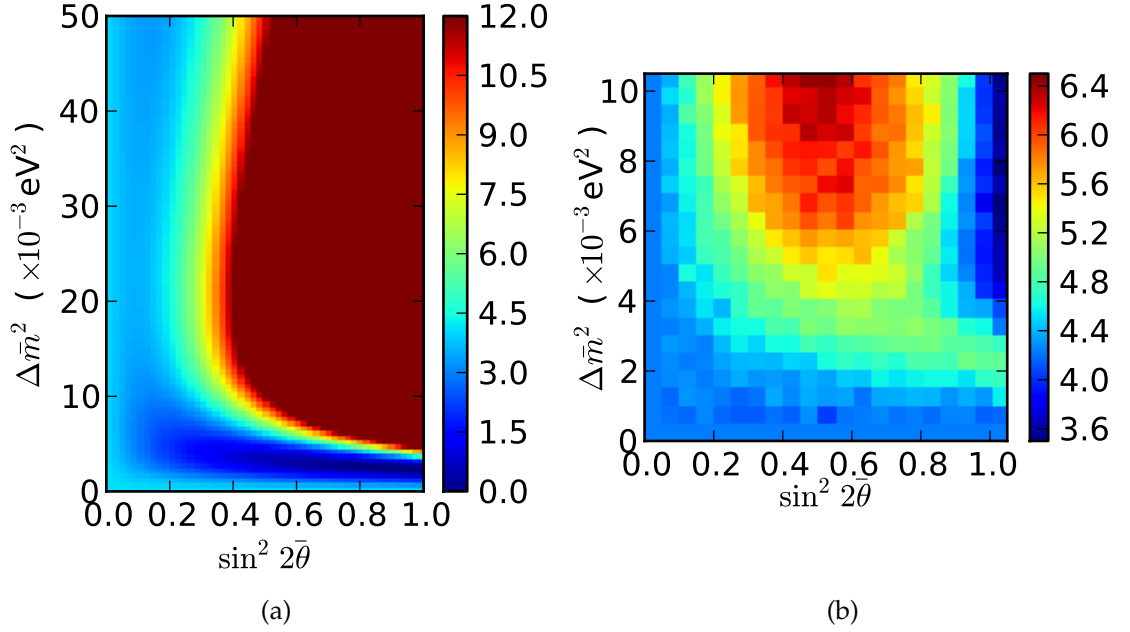


Figure 6.3: (a) The $\Delta\chi^2$ surface for 7×10^{20} POT of scaled ν_μ -dominant beam Monte Carlo, oscillated at $|\Delta\bar{m}^2| = 2.32 \times 10^{-3} \text{ eV}^2$, $\sin^2(2\bar{\theta}) = 1.0$ and (b) The corresponding FC 90% statistical correction grid for the same.

values at all parameters in the range of interest, relative to the best fit point χ_{best}^2 . Continuing the example from figure 6.2, figure 6.3a shows the $\Delta\chi^2$ surface for the same configuration, using equivalent data to 6.2 for many different oscillation parameters.

We then need to calculate the correction to this from our FC simulated data. After separately simulating experiments at all points in the range of parameters, we can calculate a value at each point on the surface $\Delta\chi_c^2(|\Delta\bar{m}^2|, \sin^2(2\bar{\theta}))$ such that α of the experiments have a value of $\Delta\chi_{\text{fit}}^2 < \Delta\chi_c^2$. This set of values can be formed into a grid, and this grid can be seen in figure 6.3b for $\alpha = 90\%$, for the ν_μ -dominant beam example we are considering.

The confidence intervals are then easy to calculate: For the $\Delta\chi^2$ surface, a point is in the acceptance region if:

$$\Delta\chi^2(|\Delta\bar{m}^2|, \sin^2(2\bar{\theta})) < \Delta\chi_c^2(|\Delta\bar{m}^2|, \sin^2(2\bar{\theta})),$$

and the contours defining the edge of this region can be found where $\Delta\chi^2 = \Delta\chi_c^2$.

Practically, this can be achieved by subtracting the correction grid from the $\Delta\chi^2$ likelihood surface and tracing the contours where $\Delta\chi^2 - \Delta\chi_c^2 = 0$.

6.1.2 Fitting individual experiments

Special care needs to be taken when fitting individual experiments that bias is not introduced via the introduction of artificial boundaries. The problem is illustrated in figure 6.4a, where a hard maximum of $10 \times 10^{-3} \text{ eV}^2$ has been imposed on the fitting. Even if the distribution of fits followed the perfectly Gaussian case, we would expect experiments to be distributed according to the likelihood surface e.g. that shown in figure 6.3a. When constrained, many experiments end up clustered around the artificial boundary instead of fitting higher in $|\Delta\bar{m}^2|$. As can be seen in the corresponding 68% FC correction histogram for this scenario (figure 6.5a), this boundary causes an additional lowering of the 68% measurement at the high- $|\Delta\bar{m}^2|$ border, and away from the borders the behaviour is similar to the Gaussian case (because the $\Delta\chi^2$ between the best fit and true oscillation point is being kept artificially low).

Even when not fitting with an artificial boundary, a similar effect can be observed when the fitting procedure finds a local minimum of $\Delta\chi^2$ instead of the global minimum. This problem can be identified easily through manual inspection, but is impractical for the tens of thousands of experiments being analysed in the FC method. For the fitting in this thesis and the published MINOS analyses, a fitting method was developed in order to attempt to find the global minimum automatically.

Firstly, a coarse grid search designed to probe as much of the likelihood surface as possible attempts to find an approximate minimum. The points considered in this search are displayed in figure 6.6, and the density of points was determined by manual inspection of a large number of random experiments and chosen to cover the size of most local minima. Once a minimum has been found this way, MINUIT is seeded with this minimum and used to find a more precise

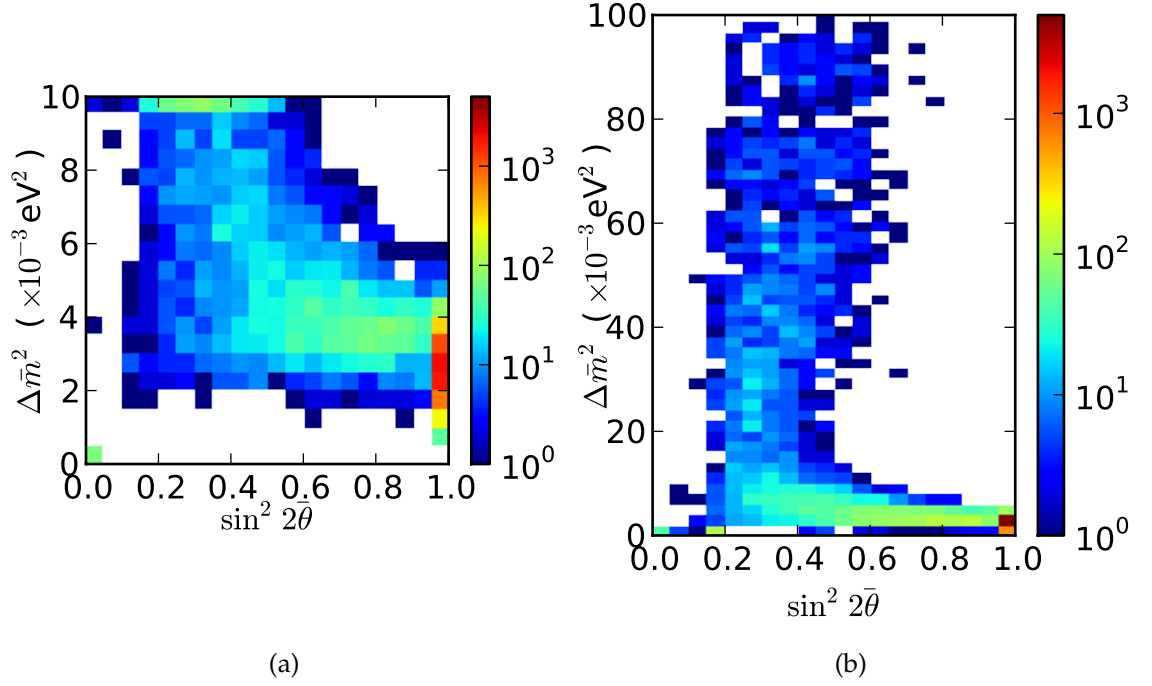


Figure 6.4: The distribution of best fit points for 10 k experiments where the fitting is (a) constrained to $10 \times 10^{-3} \text{ eV}^2$ and (b) fit globally. 7×10^{20} POT of ν_μ -dominant-beam MC was oscillated at $|\Delta \bar{m}^2| = 2.5 \times 10^{-3} \text{ eV}^2$, $\sin^2(2\bar{\theta}) = 1.0$

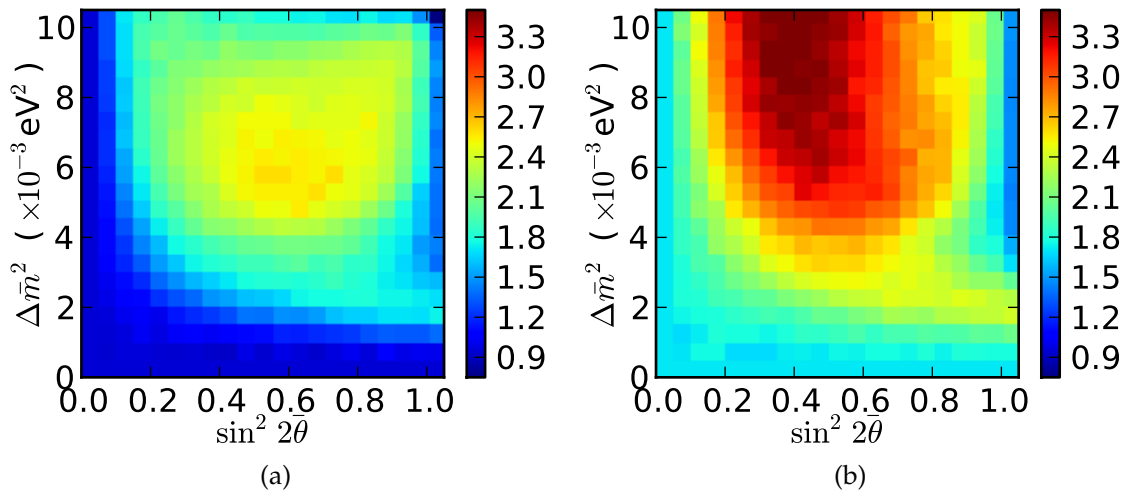


Figure 6.5: The 68% FC-correction surfaces calculated by applying the two fitting scenarios in figure 6.4 to a full range of parameters. (a) is with the fit constrained to $10 \times 10^{-3} \text{ eV}^2$, and (b) fitting globally.

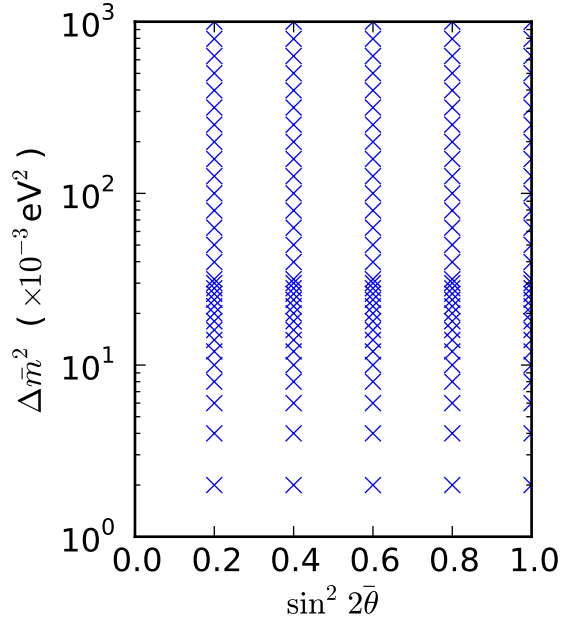


Figure 6.6: Points considered in the coarse grid search stage of the global fitting procedure. The likelihood is sampled linearly up to $|\Delta\bar{m}^2| = 30 \times 10^{-3} \text{ eV}^2$, and logarithmically thereafter. The zero-oscillation point is also considered.

location.

There are several ways in which this fitting method can fail, and several methods of recovery. In the event that MINUIT fails, or finds a less likely minimum than the seed value, a high resolution grid search is carried out around the best fit coarse grid search point. The best fit from MINUIT is also compared to two “known” points; the null-oscillation point, and the “known” true oscillation point.

The use of this latter, true value point, can be justified by considering what would happen with real data - a high resolution grid would be generated, followed by a manual inspection; if a better fit is found by manual inspection then the fit would be restarted. Therefore, in order to make sure the fitting procedures are as close to the way the final result would be fitted as possible, it is valid to use any information possible, in lieu of an exhaustive global search (which would *always* cover the point of true oscillation).

If either of these give a better fit than the found minimum, then MINUIT is

re-seeded from these points, with the same compensation for failure modes as outlined above. This method finds a global minimum, or one very close to it, extremely reliably, with all observed mis-fitting falling within the first 68% of experiments.

Being close to the global minimum, and not exactly on it, in difficult fitting cases, is not a problem, because it is the $\Delta\chi^2$ between the minimum and truth point that is important. As long as this difference falls within the first 68% of experiments, it will not affect the final result. From study of sets of sample experiments, the number of results with negative $\Delta\chi^2$ is $< 0.5\%$, and of those, all studied would have fallen within the first 68% of the $\Delta\chi^2$ spectrum, had they fitted the true global minimum.

The distribution of fits for the same scenario, but using the global fitting method is shown in figure 6.4b. The method is clearly finding fits over large areas of the parameter space. The resultant 68% FC-correction grid can be seen in figure 6.5b. There is no artificial boundary with this method, and the effects of non-Gaussian fit distributions is clearly visible, by the region in the centre of the displayed parameter space being $\gg 2.3\sigma$.

6.1.3 Aggregating Multiple Runs

Data from MINOS has been taken over many years and beam configurations. The biggest difference was in switching NuMI from ν_μ -dominant to $\bar{\nu}_\mu$ -enhanced mode, since this causes the shape of the $\bar{\nu}_\mu$ energy spectrum to change drastically.

Because of these differences, data from each run period is separately simulated, and this in turn leads to the requirement that each period of running must be considered separately when performing the FC method.

However, the process requires all data to be fitted simultaneously into a single $\Delta\chi^2$ value - so there is no way to combine the effects of multiple runs after the process - it must all be done simultaneously. To do this, fake data for each run is generated separately for every FC experiment, and then the fitting is done over

all runs simultaneously. The likelihood χ^2 , used for fitting, for the combined data set can be calculated as:

$$\chi^2 = \sum_j \chi_j^2,$$

where j indicates a sum over separate run periods, and χ_j^2 is calculated the same as equation (6.1). This expression follows from the basic principles of log-likelihood calculation.

The effects on the correction surface when combining multiple runs depends heavily on what is being combined. Figure 6.7 shows the differences between building the ν_μ -dominant- beam correction surface using run 3 MC only, and correctly combining runs 1-3 MC in the correct POT proportions. The difference in the end result is minimal, an expected result given the similarity between the spectra of the three ν_μ -dominant runs, and the statistical dominance of the run 3 POT. This means that we can safely use run 3 MC to represent all ν_μ -dominant running.

In contrast, figure 6.8 shows the effect from combining the ν_μ -dominant and $\bar{\nu}_\mu$ -enhanced MC. Figure 6.8a shows the FC correction grids for 3.4×10^{20} POT of ν_μ -dominant beam MC. Figure 6.8b shows the correction grids when combining this MC with 7.1×10^{20} POT of $\bar{\nu}_\mu$ -enhanced beam MC to represent the full 10.5×10^{20} POT of running. Clearly, adding the extra statistics has had the effect of constraining the extent to which the best fits to individual experiments can move away from the Gaussian expectation.

6.2 Systematic Corrections

An essential part of any analysis is the compensation for unknown but bounded experimental errors. For the analysis in this thesis, the FC method has been adapted to include the effects of systematics through the Bayesian procedure of

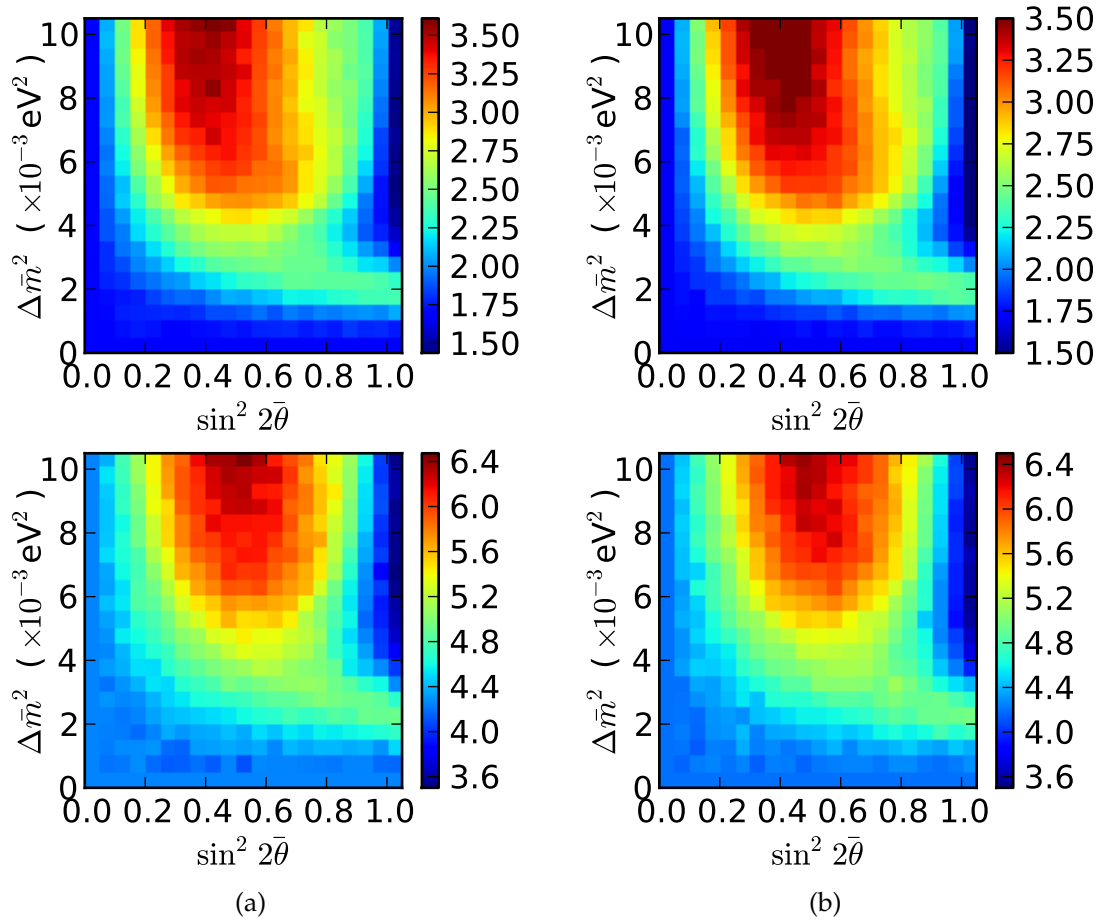


Figure 6.7: The 68% (top) and 90% (bottom) FC-correction surfaces calculated from 7×10^{20} POT of (a) run 3 MC and (b) separate run 1, 2 and 3 MC in the same proportions as actual data taking.

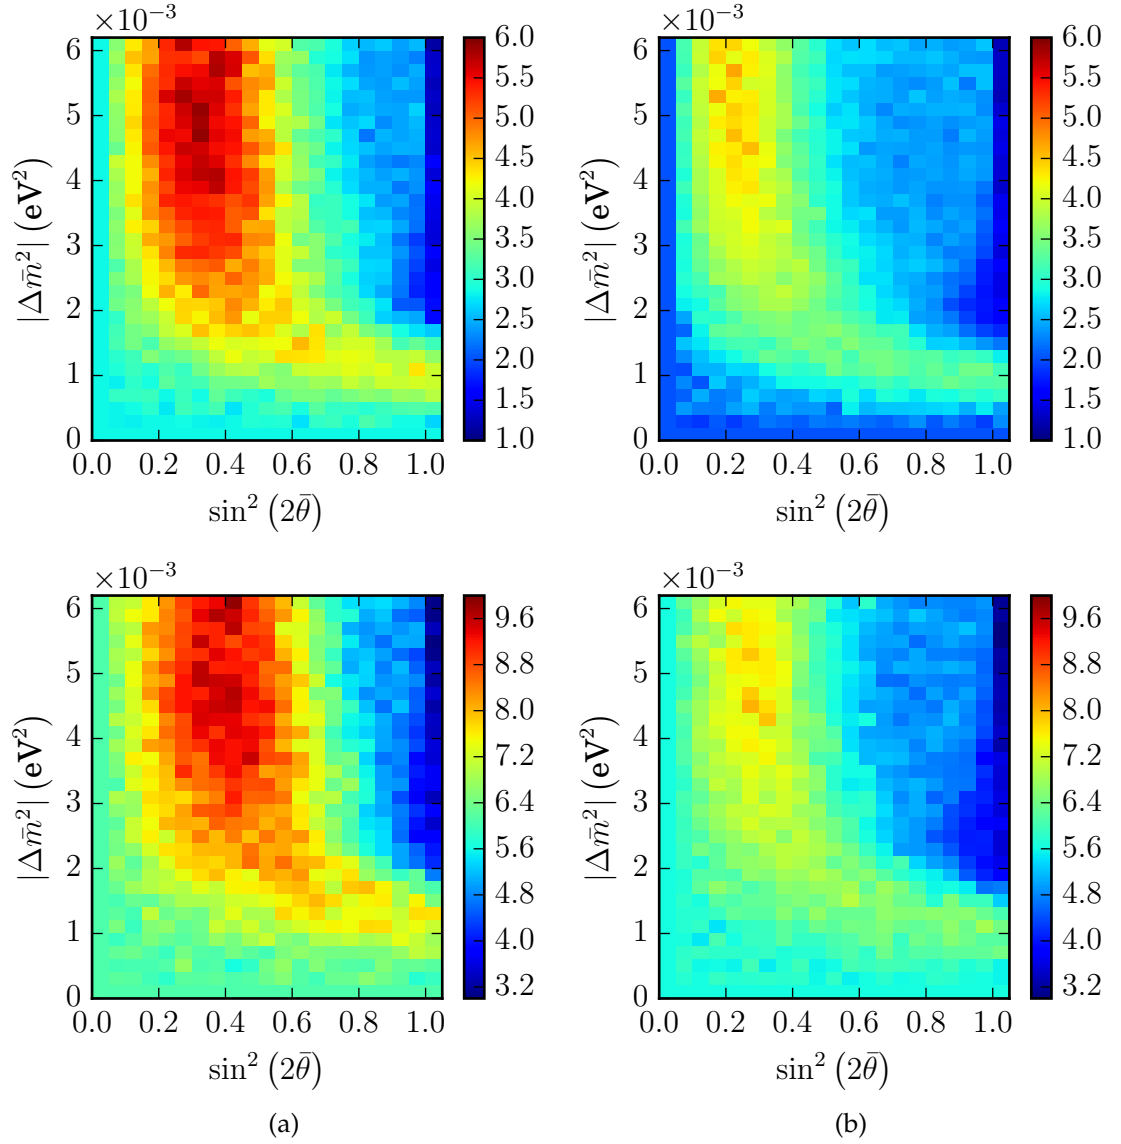


Figure 6.8: The 68% (top) and 90% (bottom) FC-correction surfaces for (a) 3.4×10^{20} POT of $\bar{\nu}_\mu$ -enhanced beam, and (b) aggregated correction for this $\bar{\nu}_\mu$ -sample with the 7.1×10^{20} POT of ν_μ -dominant beam (all runs combined).

marginalization [110]:

$$\text{prob}(X|I) = \int_{-\infty}^{+\infty} \text{prob}(X, Y|I) dY, \quad (6.2)$$

where X is the distribution that you are testing the probability of, I is the underlying assumption and Y is the parameter that you are uncertain about. This procedure allows us to eliminate nuisance parameters, such as systematic effects, by integrating out the effect they have on the final result.

Because the calculation of this integral directly for each experiment would require prohibitive quantities of computing power, especially when considering multiple systematic shifts, a different approach is made in this analysis.

Upon generation of each FC experiment, a single set of exact values for each of the systematic shifts is randomly generated, and used to build the data spectrum for that experiment. By accumulating many thousands of experiments generated this way, the effect is of performing the integral in equation (6.2) numerically.

6.2.1 Experiment Generation

The best way to accurately calculate the effect of multiple, possibly-correlated systematics, is to calculate the exact effects by accounting for the shift on every single Monte Carlo event. This is the approach used for this analysis.

In the statistical case, the ND spectrum does not need to be regenerated for every experiment, as it has enough data that the statistical fluctuations do not have a noticeable effect on the FD spectrum, and is otherwise identical for each experiment. However, when simulating systematic effects there are systematics that can affect the near and far detectors differently; thus affecting the difference between predicted and measured FD spectra. Therefore, we need to regenerate the ND spectrum for every experiment.

The systematic shifts are re-chosen for every experiment. Each shift has a defined variance, and a random number generator is used to generate a value on a

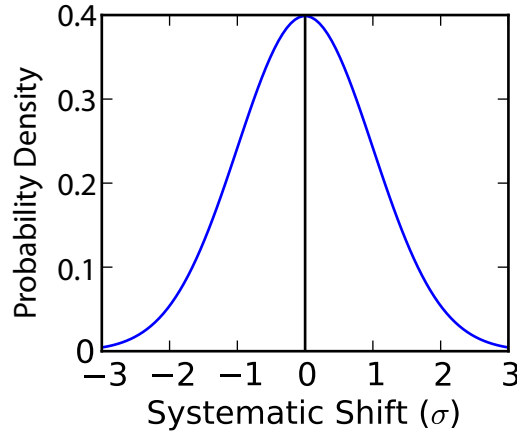


Figure 6.9: The magnitude of each systematic shift is chosen according to a normal distribution, based on predefined variance values.

normal distribution, as seen in figure 6.9. The value is chosen for each systematic i as:

$$\Delta_i = \sigma_i \times X_i,$$

where Δ_i is the magnitude of the systematic shift to be used in the experiment, σ_i is the magnitude that the shift should fall within 68% of the time, and X is a random variable distributed as $X_i \sim \mathcal{N}(0, 1)$. With this approach, the stated magnitudes for each of the systematic shifts become statements about the confidence we have, rather than a limit to the magnitude of the shift.

Once we have the set of systematic shifts, then we can generate the ND spectrum. Since we have enough near detector data as to be relatively insensitive to statistical fluctuations, the exact level of statistics at the near detector does not matter - we can use the entire set of ND MC events to build the reconstructed energy spectrum R^N . Each of the systematic shifts S_i are applied sequentially to every MC event:

$$R^N = S_1 \circ S_2 \circ \dots \circ S_n(\{\text{MC}_{\text{ND}}\}, \Delta_n).$$

This collection of events is then used to build the ND histograms used in the extrapolation.

6.2.2 Far Detector Fake Experiment Generation

Generation of the FD histograms is more complicated than generating the ND. The low event count makes it very important to correctly sample the correct level of statistics, from a high-statistics set of MC events. Additionally, the effects of neutrino oscillation can be observed at the far detector, causing an energy-dependant distortion of the MC spectrum. The nature of this oscillation function can make it computationally expensive to calculate the effects of oscillation on every MC event. Finally, the effects of systematic shifts can cause the distribution and normalisation of events to change, in ways that can be hard to approximate when combining multiple systematics.

This section describes how the FD histogram is built efficiently through usage of the extrapolation, how rejection sampling was used to compensate for high computation requirements, and how the FD histograms were eventually generated in this thesis.

The simple way

Before exploring the alternate methods that were developed for the analysis in [2], it is illustrative to examine how, in principle, systematics are included in the FC process, along with the reasons other methods were explored.

The simplest way of generating a fake-data systematically-shifted FD energy spectrum would be to iterate over every single event in the MC, applying the oscillation and chosen set of systematic shifts to every event, and then building a full-MC-statistics histogram with the results. This could then be scaled down to the desired level of POT, and each bin then fluctuated according to Poisson statistics, so that the final histogram has integer events.

This process is time consuming, because of the combination of number of MC events and the amount of data associated with each event, making it very computationally intensive to process even a single event. Worse, this process needs to be done thousands of times because every FC experiment needs to have a different

set of systematics applied to it, in order to correctly apply the marginalization. These originally combined to make this approach untenable - and so the rejection sampling method described in the next section was developed.

When constructing the final analysis for this thesis, however, this process was heavily optimised and through a combination of data reduction, and rewriting the systematic shifting and oscillation algorithms to minimize jumping around memory for data access (potentially reducing the effects described in [111]). Through these optimisations, approximately a 100x speed-up was observed in the rate of processing, which was cross checked by running the old and new methods side-by-side.

A faster way - Rejection Sampling

The speed problem can be solved by using rejection sampling [112]. Rejection sampling allows an arbitrary probability distribution function $f(x)$ to be sampled from another, usually simpler, distribution $g(x)$. For the purposes of this analysis, the distribution to sample is the systematically shifted, oscillated neutrino true-energy spectrum at the FD (T_j^F), and the actual distribution of events to sample from are the systematically-neutral FD events from simulation ($\{\text{MC}_{\text{FD}}\}$):

$$\begin{aligned} f(j) &\equiv T_j^F, \\ g(j) &\equiv F_j^{\{\}} = F(j, \{\text{MC}_{\text{FD}}\}), \end{aligned}$$

where j indicates that the probability density functions are binned, and $F_j^{\{\}}$ indicates the true energy far detector spectrum obtained by applying a binning function F to the far detector MC event data $\{\text{MC}_{\text{FD}}\}$. Any desired number of events can then be drawn from the simulation data by selecting a single event at random, and then deciding to throw it away and draw another by using the ratio of these functions to calculate the probability of acceptance such that for any bin j :

$$P_{\text{accept}}(j) = K \cdot \frac{g(j)}{f(j)}, \quad (6.3)$$

where K is an arbitrary normalization constant, chosen such that the maximum possible value of P_{accept} , P_{max} , satisfies

$$0.0 < P_{\text{max}} \leq 1.0.$$

The normalization of this ratio does not matter - the number of events is calculated independently. Usually, K would be chosen such that $P_{\text{max}} = 1.0$, to avoid unnecessary discarding of events. However, in this analysis a value of K such that $P_{\text{max}} = 0.87$ is used, in order to compensate for systematic shifts, discussed below.

Once the probability of acceptance has been calculated for each prospective event, a uniform random number can be generated. By comparing this random number with the probability for the bin that the events' energy falls into, a decision on acceptance can be made. However, first, systematic shifts need to be taken into account.

6.2.3 Systematic Shifts and Far Detector Generation

The procedure becomes slightly more complicated once the potential shifts from systematic errors are taken into account. Because of the computing time involved, we cannot reprocess the entire set of events for every set of systematic shifts, but because the effect of any particular shift could be highly dependent on each event, the shifts due to systematics needs to be calculated for each event.

Thus, rejection sampling needs to be adapted to account for the effect of the systematic shifts on the relative contribution of each event on the final histogram. The important property of consideration is each event's re-weight value, r_w . Each event is assigned a weight by the MC process that represents its proportional contribution to the experimental results. When systematic shifts are applied, they

can affect the relative probability of each event being observed, and so the weight associated with individual events can shift.

Because the systematic shifts can affect the normalisation of events, the FC process needs to account for the fact that a systematic shift can account for a change in the number of observed events (relative to the nominal case).

For each event n randomly picked from the MC, a uniform random variable $t_n \sim \mathcal{U}(0, 1)$ is drawn, and compared with the probability of acceptance calculated from the spectrum oscillation, such that the boolean acceptance $A_n \in \mathbb{B}$ of the event is:

$$A_n = t_n \geq r_n \cdot P_{\text{accept}}(j),$$

where the bin j is calculable from the event properties, and the weight of the event $r_n \equiv 1$, because any variation would be taken out by the constant in the acceptance probability, equation (6.3).

After application of the systematic shifts, the event weight changes such that

$$r_n \rightarrow r'_n = f(n, S),$$

and there is a new acceptance case $B \in \mathbb{B}$:

$$B_n = t_n \geq r'_n \cdot P_{\text{accept}}(j),$$

with three possible consequences; Either $A_n = B_n$, in which case the reweighting of the event has caused no change to its acceptance, $A_n = \emptyset, B_n = \mathbf{I}$ in which case the reweighting has caused an event that would have been previously rejected to be accepted, or $A_n = \mathbf{I}, B_n = \emptyset$, in which case an event that would have been previously accepted, is now rejected.

These last two cases, by virtue of letting events through that would be excluded by the rejection sampling process, affect the normalisation. Since we are using rejection sampling, we can deduce that while we are in the process of draw-

ing events, the number of events total that we are trying to draw, changes. For each event allowed through by the systematic reweighting that would have otherwise been rejected, we increase the total number of drawn events by one. For each event rejected that would have been allowed, we reduce the total number of events by one (with the corresponding handling of the fact that if we had drawn the last event, we would now be complete).

This gives us all of the knowledge we need in order to generate the entire set of systematically shifted FC experiments.

6.2.4 Systematic grids

The actual grids calculated from this process are shown in figure 6.10, for 68% and 90% confidence corrections. These were both generated with a minimum of 5000 experiments at every data point - and once the final result was approximately known, further statistics was gathered for the area covered by up to the 3σ Gaussian contours, giving the region covering the 3σ contour at least two million sample points each.

The result of applying the FC corrections to the Gaussian sensitivity is shown in figure 6.11, both for the statistical and statistical+systematic correction grids. It can be seen that the effect on the end result of the systematic errors is minimal, only slightly expanding the contours in $|\Delta\bar{m}^2|$ and $\sin^2(2\bar{\theta})$.

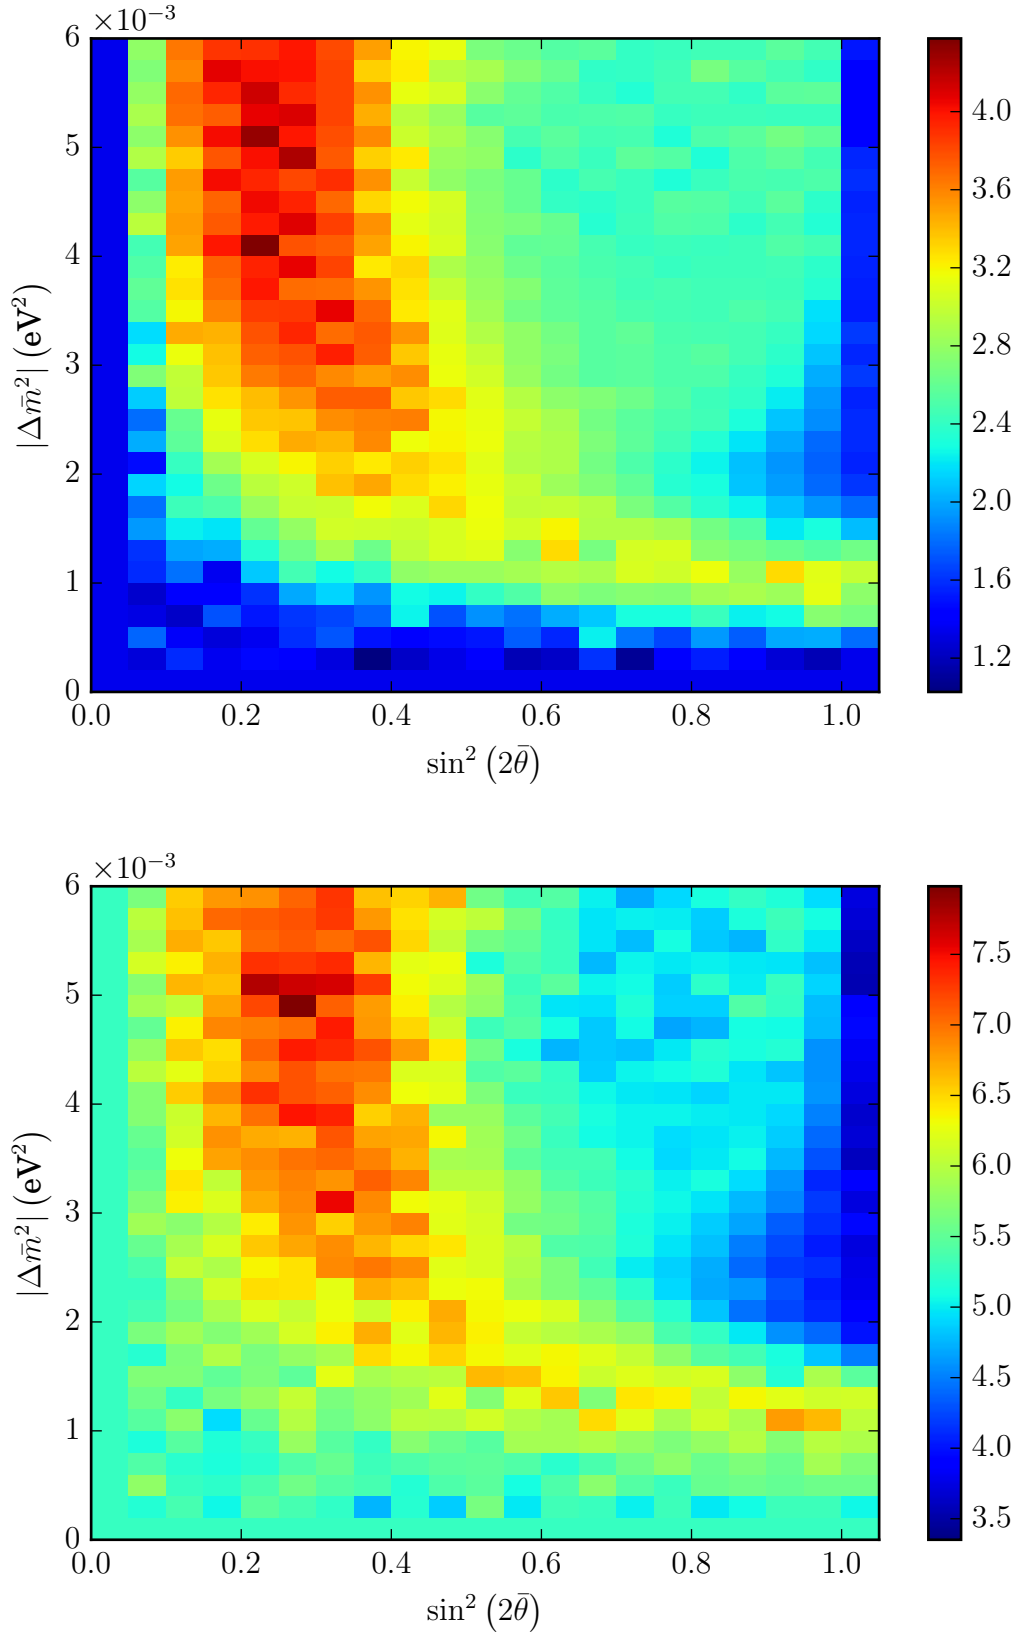


Figure 6.10: 68% (top) and 90% (bottom) Feldman-Cousins correction grids for full ν_μ -dominant $+\bar{\nu}_\mu$ -enhanced samples, with systematic shifts applied.

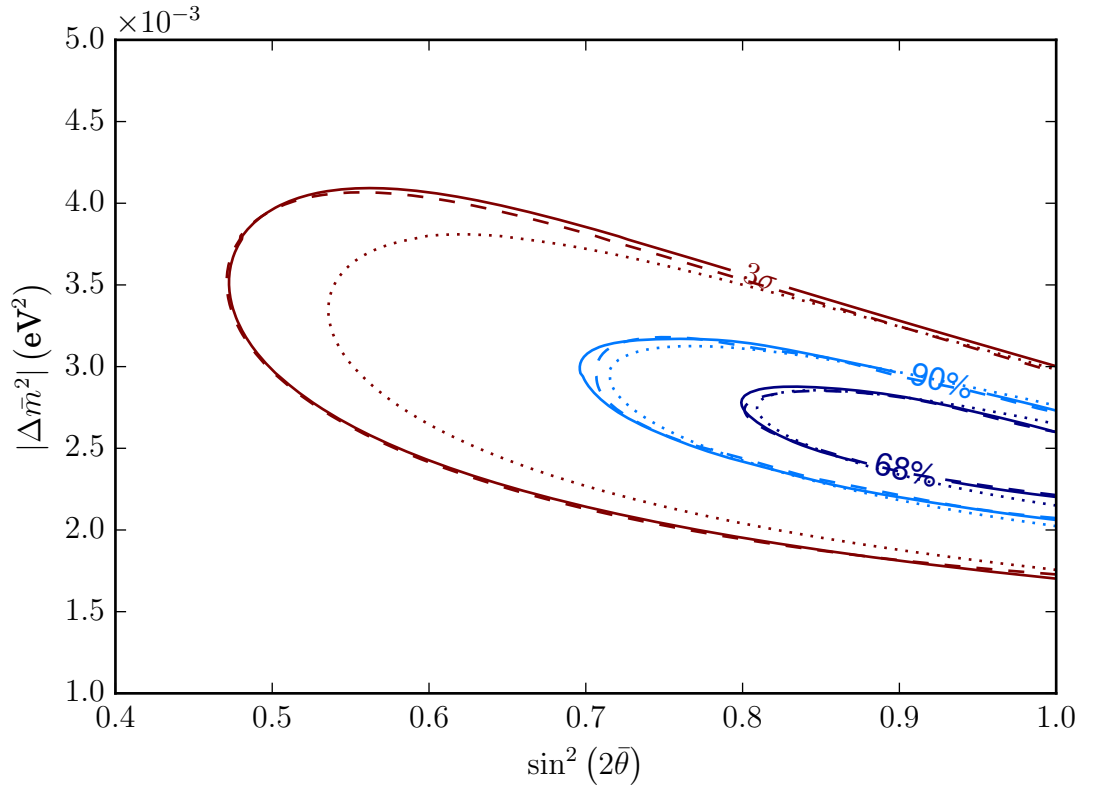


Figure 6.11: Antineutrino oscillation parameter sensitivity to oscillations at $|\Delta\bar{m}^2| = 2.4 \times 10^{-3} \text{ eV}^2$ and $\sin^2(2\bar{\theta}) = 1.0$ for Gaussian (dotted), Statistical-only FC (dashed) and Statistical+Systematics FC (solid), for the full POT sample. Contours are shown for 3σ , 90% and 68%.

Chapter 7

Results

Having established a method for accurately evaluating the systematic and statistical limits of any measured result with the Feldman-Cousins method, it becomes time to actually look at and compare the MC simulations and extrapolated predictions with actual data directly.

7.1 Near Detector

Over the entire running time of the experiment under consideration 8.48×10^{20} POT of data was collected at the near detector. This is less than at the far detector, due to a lower live-time, various special near-detector only runs, and data quality cuts, but is acceptable because the near detector is not measuring oscillations, but rather constraining our understanding of various systematics including beam flux, cross-sections and detector uncertainties (as discussed in chapter 5), and because of the distance to the beam and various analysis differences, the near detector sees over 10^3 more ν events than the far detector.

After applying the antineutrino selections detailed in chapter 5 to these data samples, the number of events selected as muon antineutrinos, along with the size of the data sample for that run, is summarized in table 7.1.

Figure 7.1 collects these data as an energy spectrum, split into separate spectra for ν_μ -dominant and $\bar{\nu}_\mu$ -enhanced mode running. The systematic error on the MC

Beam	Run	ND POT	$\bar{\nu}_\mu$ Events
ν_μ -dominant	1	1.19×10^{20}	109970
	2	1.62×10^{20}	150168
	3	2.48×10^{20}	245074
$\bar{\nu}_\mu$ -enhanced	4	1.66×10^{20}	350252
	7	1.08×10^{20}	227500
	9	0.44×10^{20}	88720
Total		8.48×10^{20}	1171684

Table 7.1: Collected POT and reconstructed events selected as $\bar{\nu}_\mu$ for each sample run at the near detector.

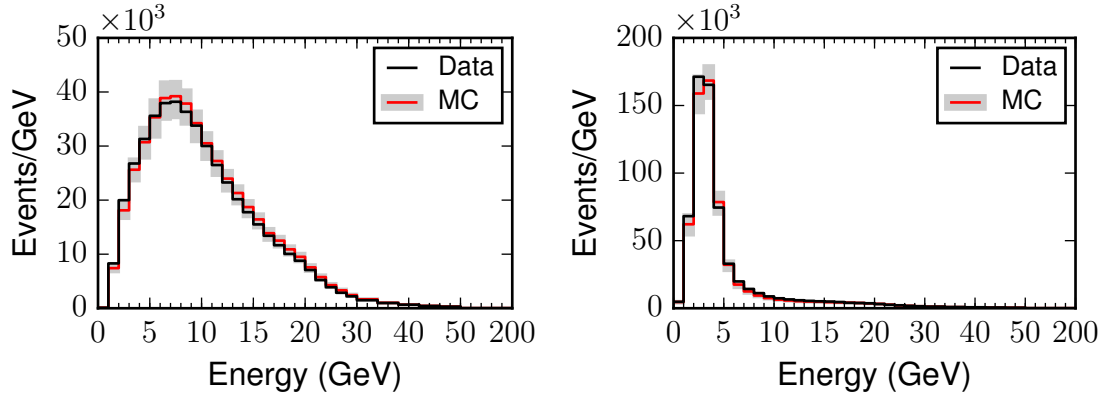


Figure 7.1: Near detector data events (black) selected as $\bar{\nu}_\mu$ for both ν_μ -dominant beam (left) and $\bar{\nu}_\mu$ -enhanced beam (right). The POT-scaled MC prediction is shown in red, with the systematic errors.

expectation is generated by fluctuating every systematic to its $\pm 1\sigma$ values, and adding the results in quadrature. The statistical error on each bin is negligible, and not visible on this scale.

The data appears to agree with the MC expectation within systematic errors, and any differences are fed in as a correction to the beam matrix extrapolation, so we can safely continue with inspecting the far detector data.

7.2 Far Detector

At the far detector, 10.49×10^{20} POT of data passed all the analysis cuts, corresponding to 7.09×10^{20} POT with the beam in ν_μ -dominant mode, and 3.40×10^{20} POT of data collected in $\bar{\nu}_\mu$ -enhanced mode. A summary of the per-run POT and num-

Beam	Run	FD POT	$\bar{\nu}_\mu$ Events	No Osc. MC
ν_μ -dominant	1	1.27×10^{20}	15	24.9
	2	1.94×10^{20}	27	38.0
	3	3.88×10^{20}	89	81.0
$\bar{\nu}_\mu$ -enhanced	4	1.71×10^{20}	99	160.0
	7	1.24×10^{20}	98	115.5
	9	0.45×10^{20}	29	40.7
Total		10.49×10^{20}	357	460.1

Table 7.2: Collected POT and reconstructed events selected as $\bar{\nu}_\mu$ for each sample run at the far detector. The expected number of events, from simulation, in the absence of neutrino disappearance is also shown, for comparison.

ber of reconstructed events selected as $\bar{\nu}_\mu$ interactions can be seen in table 7.2.

Also shown in this table, is the expected number of events in the absence of oscillations - comparing the scale of the numbers, it is clear that we are observing some sort of disappearance effect, the scale of which we shall now quantify.

Using the near detector data for the separate extrapolation of summed ν_μ -dominant and $\bar{\nu}_\mu$ -enhanced data, the minimisation procedure developed in 6.1.2 for fitting individual FC experiments was used to determine a best fit of oscillation parameters as $|\Delta\bar{m}^2| = 2.58 \times 10^{-3} \text{ eV}^2$ and $\sin^2(2\bar{\theta}) = 0.962$, for this observed data.

The energy spectra for each run individually is shown in figure 7.2, compared to the predicted energy spectra for no antineutrino oscillations, and the energy spectra of the extrapolated data at the overall best fit parameters (the individual best fit for each separate run is not calculated in the process of the final fit). Similarly, figure 7.3 shows the same data but summed over all ν_μ -dominant and $\bar{\nu}_\mu$ -enhanced samples, and compared to the CPT-conserving neutrino oscillation $|\Delta m^2| = 2.43 \times 10^{-3} \text{ eV}^2$. The blue band on all of these plots represents the 68% Poisson-statistics per-bin error on the best fit prediction, calculated based on the smallest-interval approach to Poisson coverage [113].

The contours for the Gaussian-distributed interpretation of the confidence limits can be seen in figure 7.4 for the 68%, 90% and 3σ contours, drawn at $\Delta\chi^2 = 2.30, 4.61, 11.83$ respectively, and compared to the point of best fit.

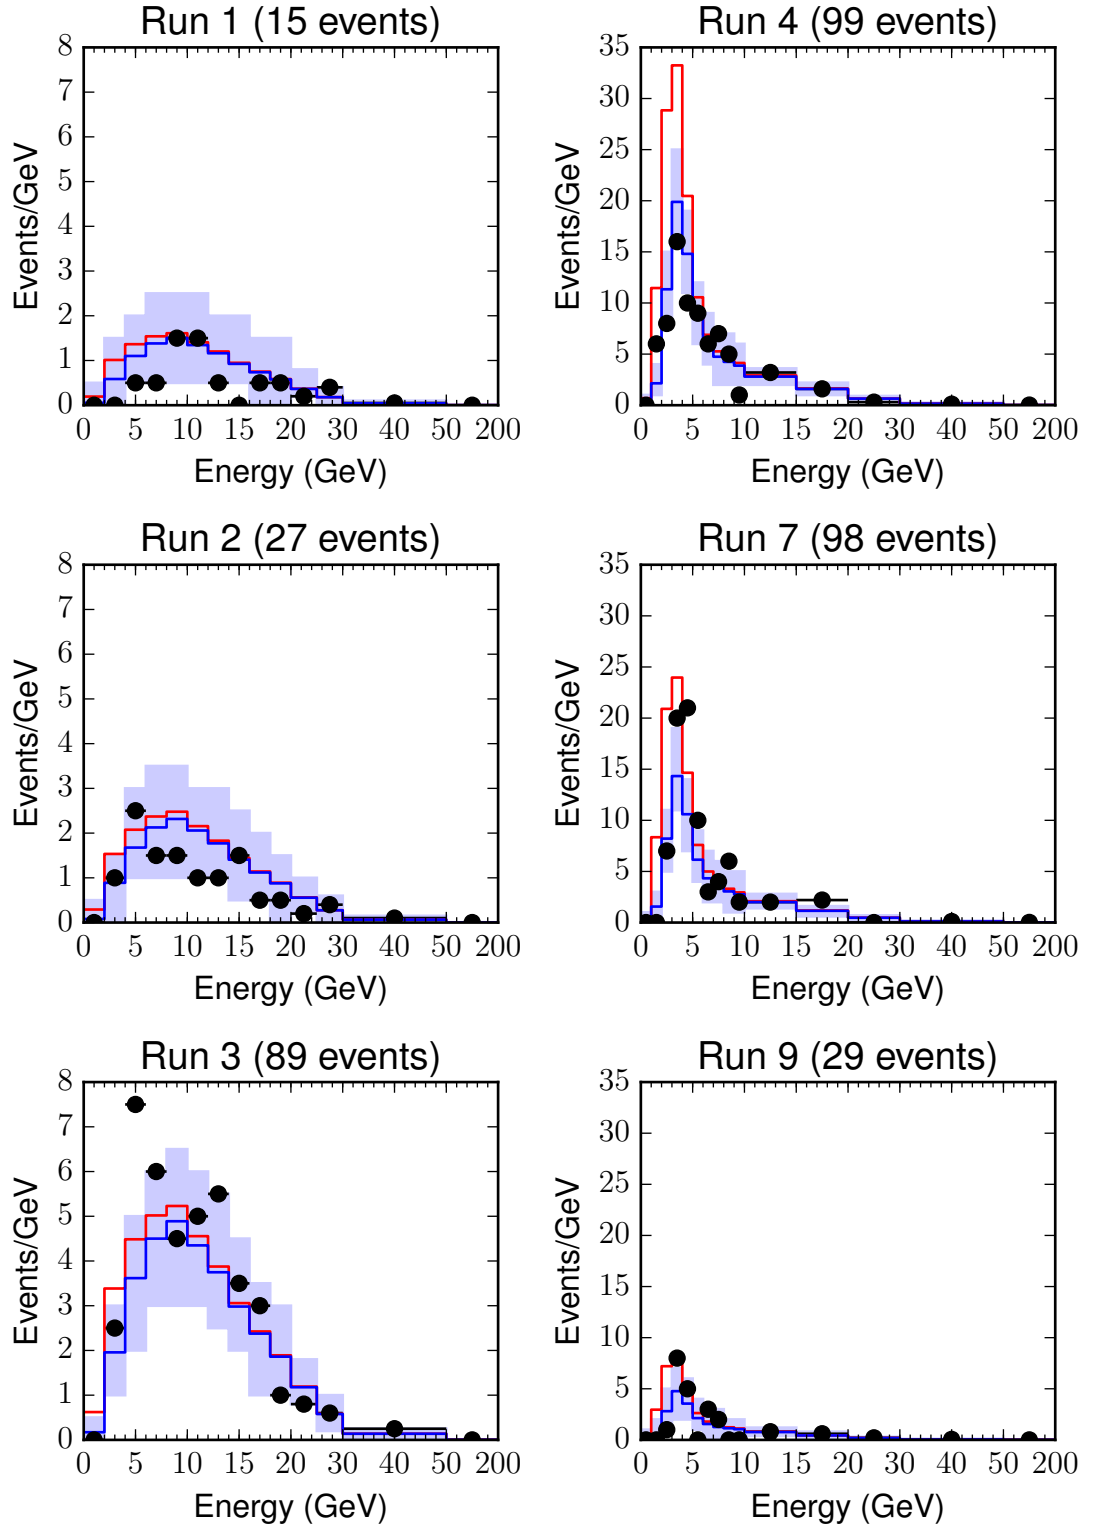


Figure 7.2: Far detector reconstructed energy data events for individual ν_μ -dominant runs (left) and $\bar{\nu}_\mu$ -enhanced runs (right). Each run shows data (black), the unoscillated prediction extrapolated from the ND data for the specific run (red) and the extrapolation of the same ND data at the global best fit (blue). The blue band shows the expected 1σ error on the contents of each bin, from Poisson statistics.

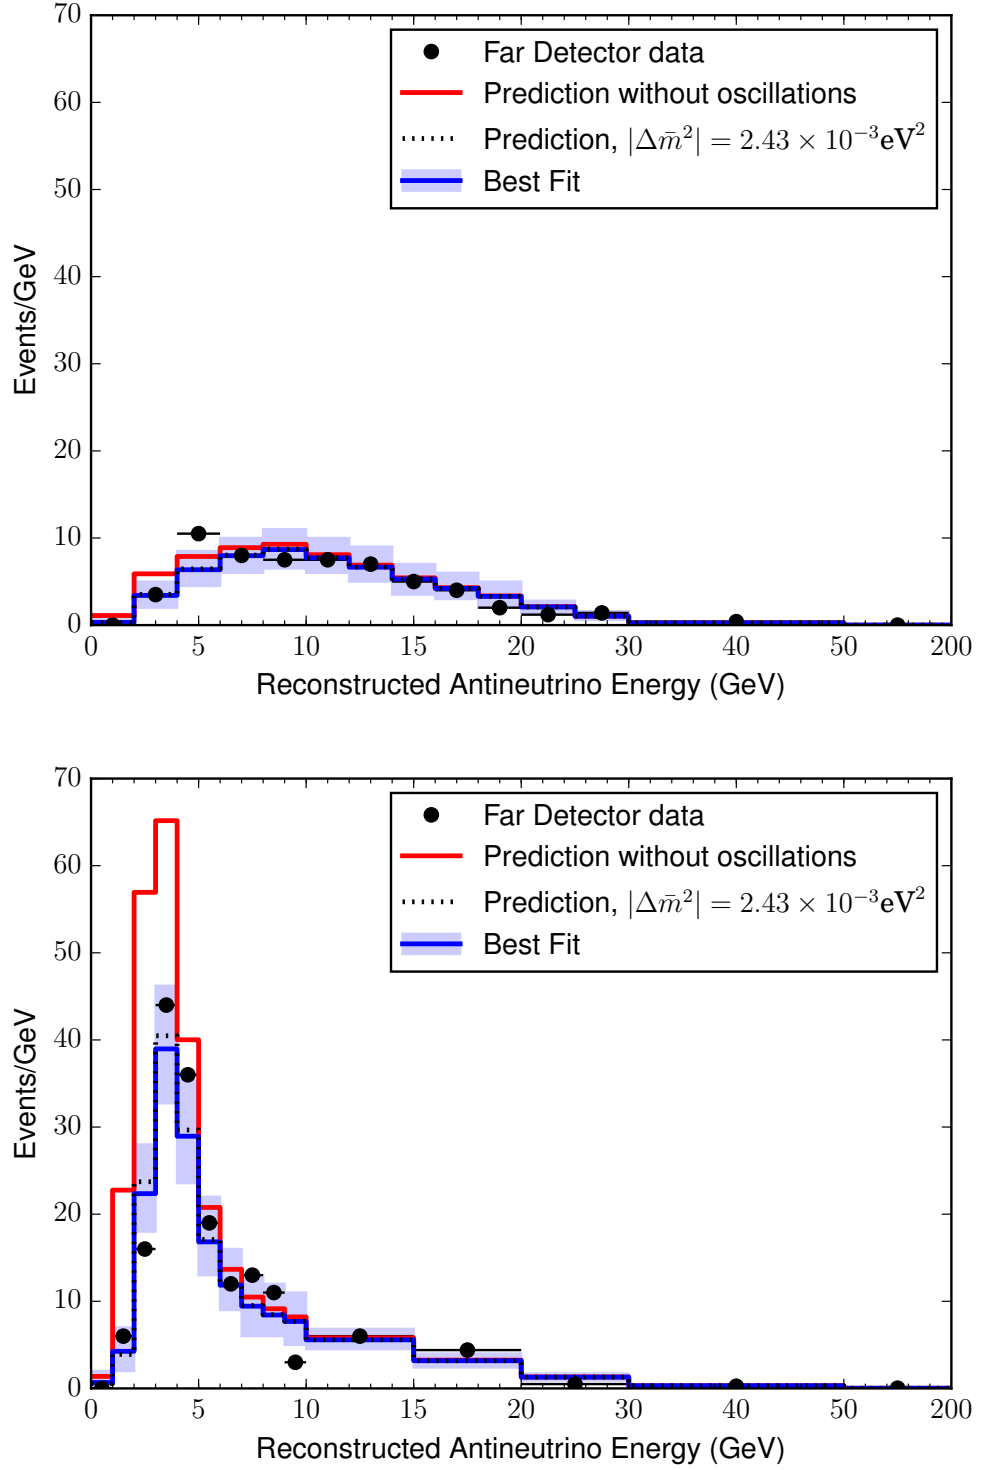


Figure 7.3: Far detector reconstructed energy data events for 7.09×10^{20} POT of ν_μ -dominant (top) and 3.40×10^{20} POT of $\bar{\nu}_\mu$ -enhanced (bottom) beam configurations. The dashed line shows the prediction at the CPT-conserving $|\Delta m^2| = 2.43 \times 10^{-3} \text{ eV}^2$, $\sin^2(2\theta) = 1.0$. The blue line shows the best fit, and the blue band shows the 1σ sensitivity for each bin's prediction, from Poisson statistics.

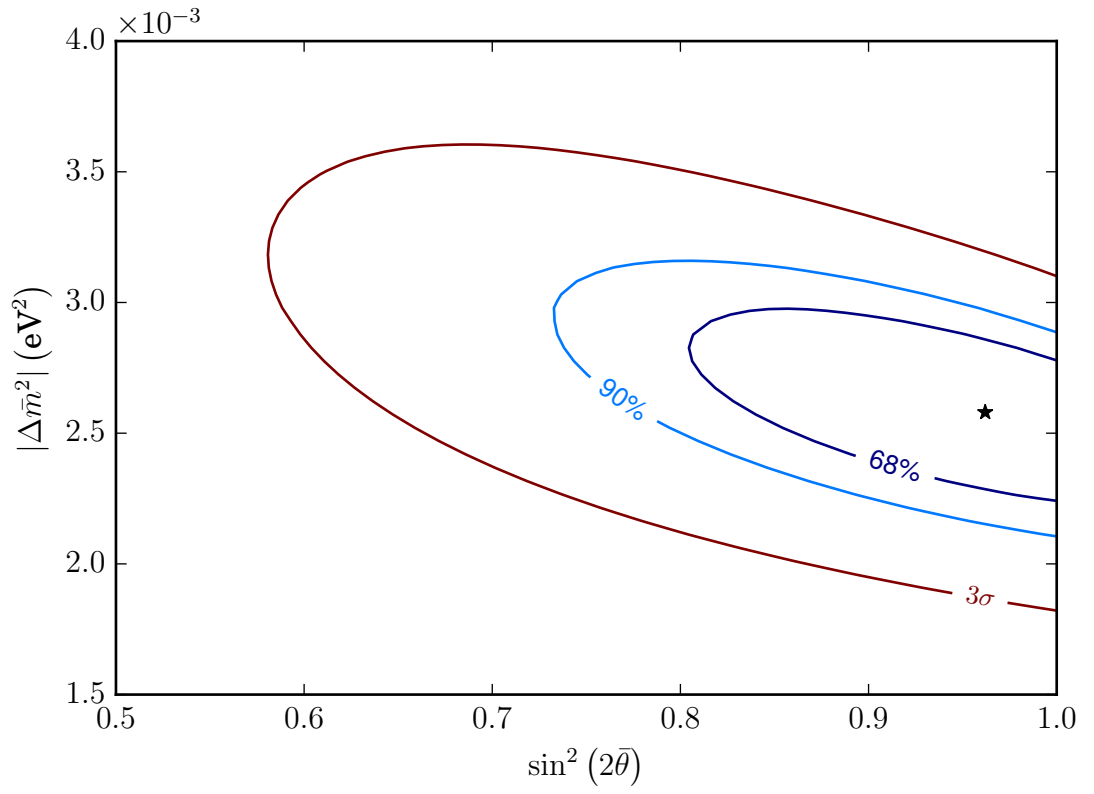


Figure 7.4: Gaussian confidence limits of fit to combined far detector reconstructed antineutrino data over all beam modes. The best fit point is indicated by the star marker.

7.3 Feldman-Cousins Corrected Results

Before using these results to draw conclusions, we need to apply the Feldman-Cousins method, as described in chapter 6. The results of this application is shown in figure 7.5, compared to the Gaussian sensitivity, the same as shown in figure 7.4. The effect of the procedure is clear - where the distribution of fits are more likely to encounter physical boundaries, at maximal mixing, the $\Delta\chi^2$ representing 68% coverage is reduced, as expected.

In contrast, at a lower mixing angle the contours are expanded in area - this represents the difficulty that the experiment has in distinguishing between different $|\Delta\bar{m}^2|$ at lower $\sin^2(2\bar{\theta})$ - a lower mixing angle means less disappearance, which means that it is harder to distinguish between the case where $|\Delta\bar{m}^2|$ would cause a single survival probability dip over the peak, and the case where multiple survival probability dips exist. This is because, in the case where $|\Delta\bar{m}^2|$ is high enough to cause multiple oscillation dips, the first dip would move into the high energy tail of the spectrum, where there is less resolving power.

This inability to distinguish between multiple regions of $|\Delta\bar{m}^2|$ means that small fluctuations can cause the best fit to randomly move between the regions, resulting in wider distribution of $\Delta\chi_{\text{exp}}^2$ than expected from a naively Gaussian perspective, thus causing a lower corresponding cumulative probability for the same $\Delta\chi_{\text{result}}^2$.

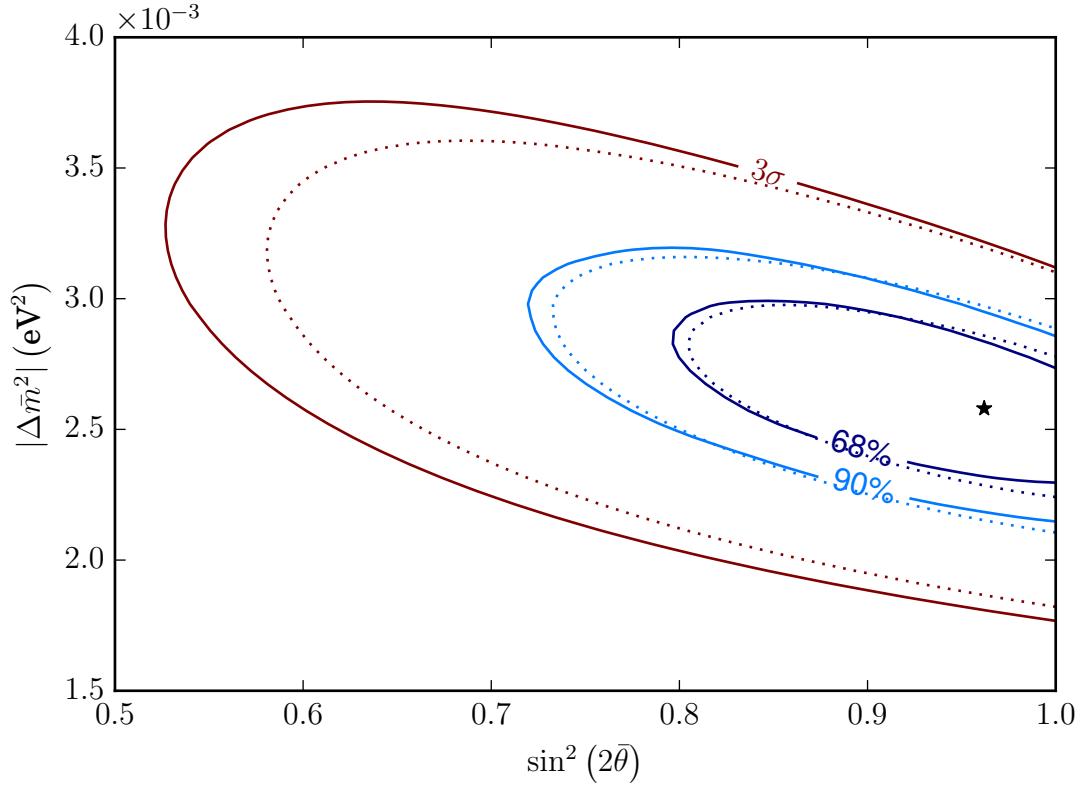


Figure 7.5: Confidence limits of the measurement as a function of $|\Delta\bar{m}^2|$ and $\sin^2(2\bar{\theta})$. The effect on the 68%, 90% and 3σ statistical-only Gaussian contours (dotted) is shown, of applying the Feldman-Cousins corrections to generate a more representative contour (solid) that includes the effects of systematic uncertainties. The best fit point is indicated by the ★ marker, and is by definition unchanged by the procedure.

7.4 Measuring Individual Parameters

It is useful to give the results of such oscillation experiments in terms of the individual parameters e.g. $|\Delta\bar{m}^2|$ (and $\sin^2(2\bar{\theta})$). In order to do this, we scan over fixed values of $|\Delta\bar{m}^2|$ and then fit for $\sin^2(2\bar{\theta})$. The resultant one-parameter $|\Delta\bar{m}^2|$ - $\Delta\chi^2$ confidence surface is then used to measure the interval of $|\Delta\bar{m}^2|$ which is within the desired $\Delta\chi^2$ confidence. In the simple case of our single minimum, and 68% confidence, this is equivalent to taking the range of $|\Delta\bar{m}^2|$ values for which $\Delta\chi^2 \leq 1$ on the contour plot.

We can correct the result from this using the Feldman-Cousins method, and get a similar result. However, if the fitting approach from the generation of the experiments that make up the grid is not identical to the way that the result is generated, then application of the grid is invalid. Therefore, we need to generate a special FC-correction grid where the fitting for every experiment is marginalised.

The results of this grid generation can be seen in figure 7.6. Whilst we only technically require a 1D grid, along the actual values of the marginalised results, generating a 2D grid allows us to generate the grid beforehand, and select the appropriate grid entries once the marginalisation values have been determined.

The result of applying this correction to the 68% 1D marginalised likelihood can be seen in figure 7.7, which compares the naive, Gaussian likelihood surface for marginalised $\sin^2(2\bar{\theta})$ with the FC-corrected surface for the same operation. Note that the corrected surface can only be interpreted in terms of above/below the 68% coverage limit - the specific shape can not be used to interpret other confidence intervals.

As expected, at low $|\Delta\bar{m}^2|$ where the marginalisation results in a maximal $\sin^2(2\bar{\theta})$, the confidence intervals have been shrunk to compensate for the lack of parameter space (because of the physical boundary), whilst the upper limit is left almost uncorrected - with this specific fit, there is no large $|\Delta\bar{m}^2|$ parameter space for the fits to escape into, because moving the $|\Delta\bar{m}^2|$ higher starts moving the oscillation probability minimum over the peak of the energy spectrum, where

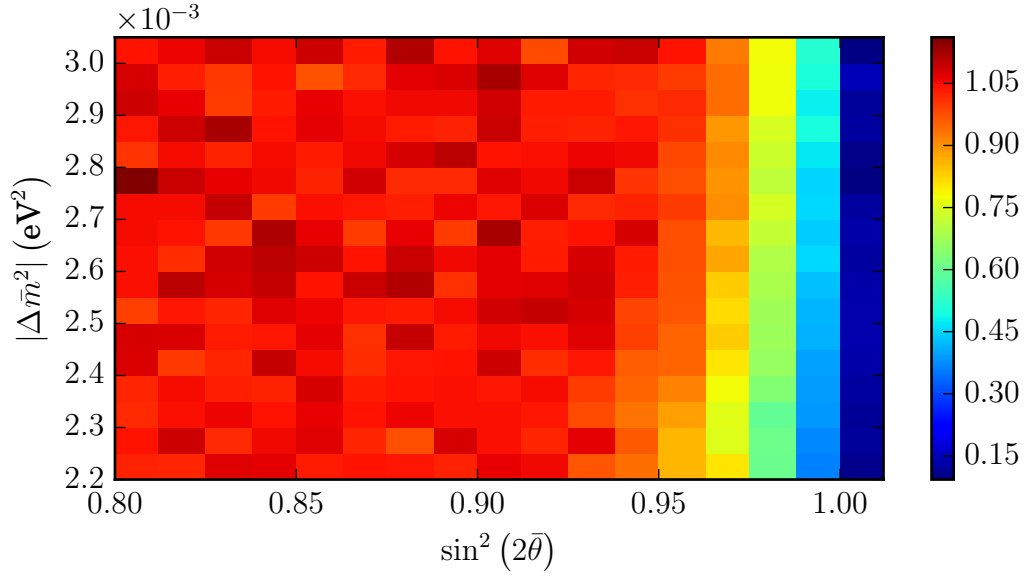


Figure 7.6: The FC 68% correction grid for 1-D marginalisation of $\sin^2(2\bar{\theta})$ in order to calculate a marginalised error bar for $|\Delta\bar{m}^2|$. Every experiment at each grid point was fitted one-dimensionally, only allowing the $\sin^2(2\bar{\theta})$ parameter to vary.

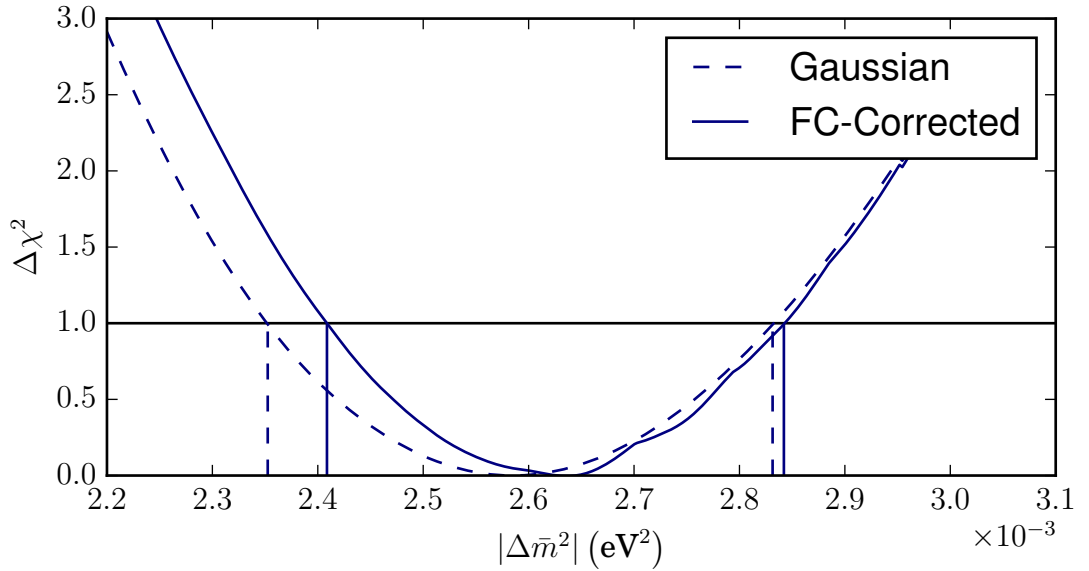


Figure 7.7: The result of applying the 1D FC correction (solid) to the $\sin^2(2\bar{\theta})$ -marginalised Gaussian likelihood (dashed). This FC-corrected likelihood is only valid for reading off the 68% confidence intervals, which for convenience of comparison with the Gaussian case has been set here to 1.0, and marked for both cases with vertical lines

there are lots of statistics and thus lots of statistical power.

Using the results from this, we can determine that within 1σ the confidence on the measurement of $|\Delta\bar{m}^2|$ is:

$$2.41 \times 10^{-3} \text{ eV}^2 \leq |\Delta\bar{m}^2| \leq 2.84 \times 10^{-3} \text{ eV}^2,$$

and therefore the final measurement from this thesis of $|\Delta\bar{m}^2|$ is:

$$|\Delta\bar{m}^2| = 2.58_{-0.17}^{+0.26} \times 10^{-3} \text{ eV}^2.$$

Similarly, we can generate a result for $\sin^2(2\bar{\theta})$ from marginalising $|\Delta\bar{m}^2|$. The FC grid for this process is shown in figure 7.8, and we can see that without the immediate physical constraint of $\sin^2(2\bar{\theta})$ the fit points are distributed in a reasonably Gaussian way - with the average being slightly higher than $\Delta\chi^2 = 1$. This is due to experiments 'Jumping' out of any minima around the true oscillation point, and thus ending up with a higher $\Delta\chi^2$ than cases where there are not multiple minima.

We can thus determine the 1σ confidence on the measurement of $\sin^2(2\bar{\theta})$ is:

$$0.857 \leq \sin^2(2\bar{\theta}) \leq 1.000.$$

Therefore, the final measurement of $\sin^2(2\bar{\theta})$ in this thesis is:

$$\sin^2(2\bar{\theta}) = 0.96_{-0.10}^{+0.04}.$$

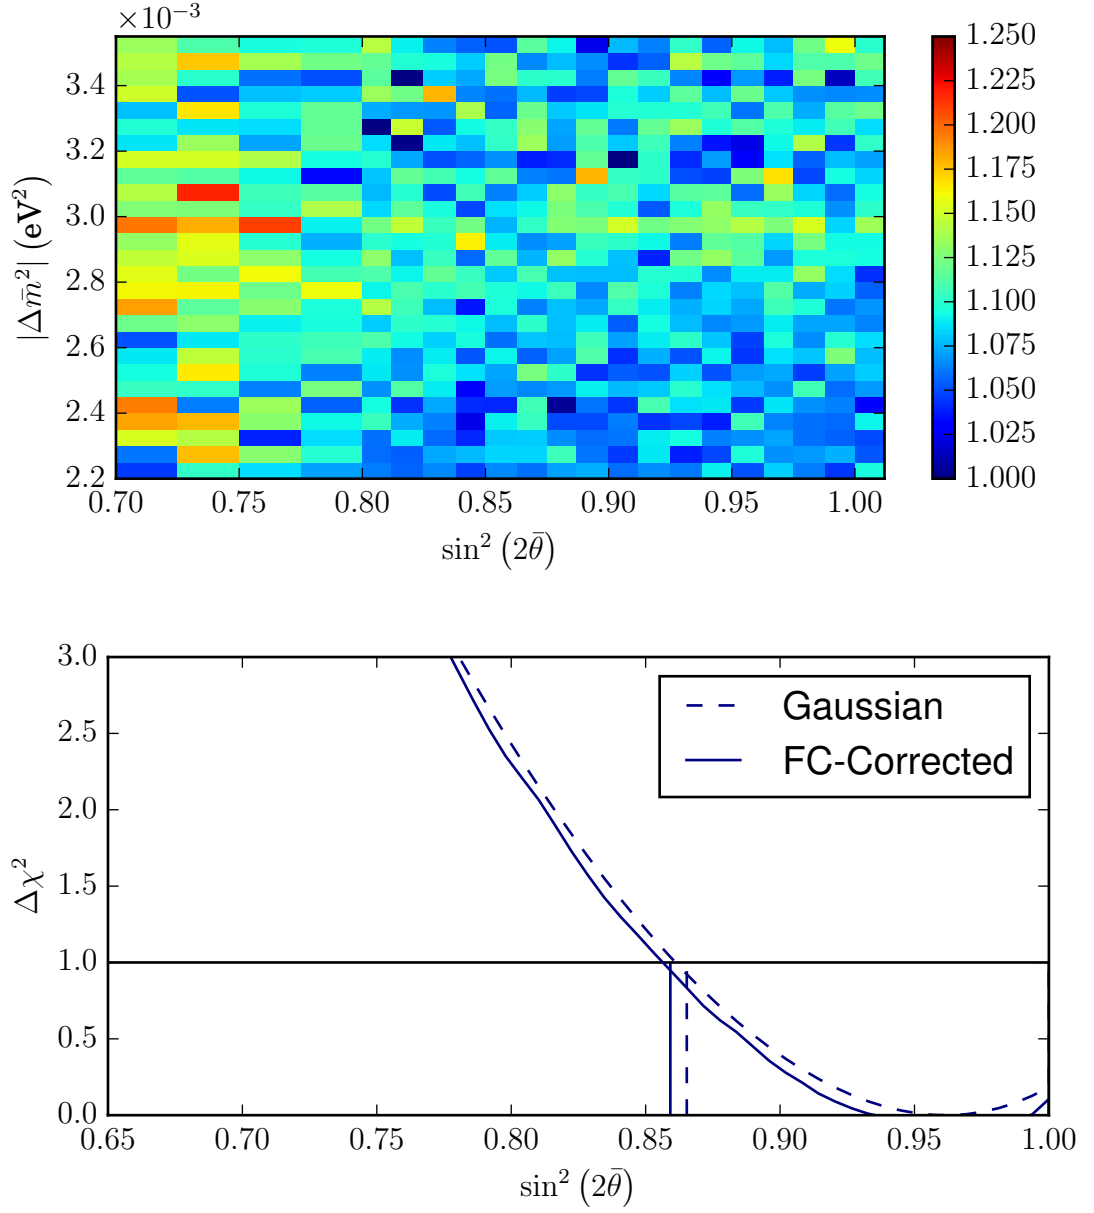


Figure 7.8: The FC 68% correction grid for 1-D marginalisation of $|\Delta\bar{m}^2|$ in order to calculate a marginalised error bar for $\sin^2(2\bar{\theta})$ (top), and the derivation of the confidence limits by applying this to the final data (bottom).

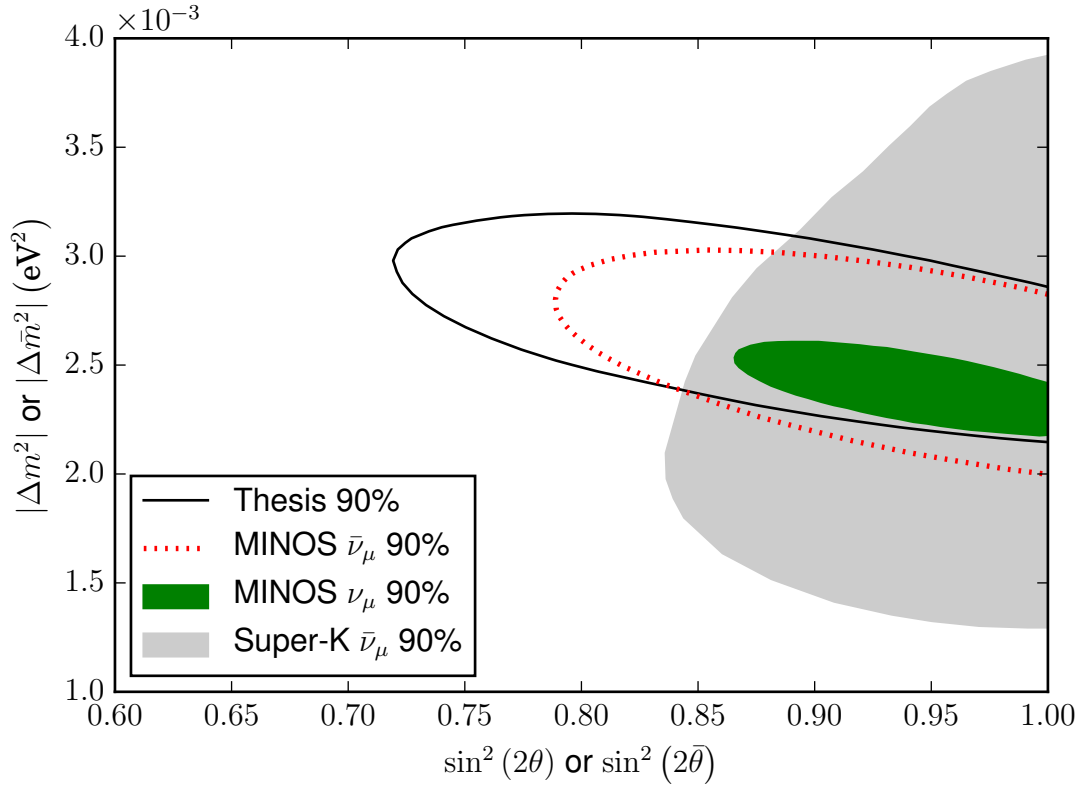


Figure 7.9: Confidence limits as a function of $|\Delta m^2|$, $|\Delta \bar{m}^2|$, $\sin^2(2\theta)$ and $\sin^2(2\bar{\theta})$ for the analysis in this thesis, MINOS 90% data from [28], and Super-K antineutrino data from [57].

7.5 Comparisons

Figure 7.9 shows these results in comparison with three other datasets; The MINOS $\bar{\nu}_\mu$ and ν_μ results, from analysing beam and atmospheric data [28], and the latest $\bar{\nu}_\mu$ -only measurement results from the Super-K experiment [57]. All of the results shown are to 90% confidence limits.

The result in this thesis is clearly both compatible with, and a great improvement on the measurement of $|\Delta \bar{m}^2|$ provided by the Super-K experiment, but has a worse measurement of $\sin^2(2\bar{\theta})$. This difference can be explained mainly by the difference in approach - since Super-K is measuring atmospheric neutrinos (see section 2.3.1) the experiment has a very large average sample of areas of the L/E parameter space where the resolution on L/E is not sufficient to see the oscillation pattern.

The high statistics measurement of this averaged out region gives a good resolving power of the $\sin^2(2\bar{\theta})$ parameter, but less information on $|\Delta\bar{m}^2|$. The resolving power that Super-K has for $|\Delta\bar{m}^2|$ will come from relatively smaller slices of zenith angle where the L/E oscillatory pattern can be resolved.

For calculation of any kind of global limits, Super-K is complementary to the results in this thesis, because of the entirely different source of neutrinos, different L/E , different baseline L and different energy range.

When comparing with the MINOS $\bar{\nu}_\mu$ result [28], the result in this thesis has a smaller limit on $|\Delta\bar{m}^2|$, but a slightly worse measurement on $\sin^2(2\bar{\theta})$. Part of this difference will be due to the more exhaustive calculation of limits using the Feldman-Cousins method - which has the end effect in this analysis of expanding the limits on $\sin^2(2\bar{\theta})$, and slightly reducing the limits on $|\Delta\bar{m}^2|$ at maximal mixing. In addition, whilst this thesis uses a slightly larger set of accelerator data, an extra 0.45×10^{20} POT, the MINOS analysis includes 37.88 kton yr of atmospheric neutrino data and antifiducial events, in the form of muons from beam neutrinos that interact in the rock upstream of the far detector.

The atmospheric sample gives the MINOS result extra power to constrain $\sin^2(2\bar{\theta})$, for the same reasons as Super-K. The antifiducial events will add a small global effect, but because of the very low energy resolution their power to constrain the end result is limited, especially in terms of $\sin^2(2\bar{\theta})$.

Clearly, the MINOS ν_μ measurement is much better in terms of equivalent limits on $|\Delta m^2|$ and $\sin^2(2\theta)$ - due to larger beam exposure in POT and ν_μ vs $\bar{\nu}_\mu$ cross-sections. But the result is clearly compatible with this $\bar{\nu}_\mu$ measurement - when comparing the measurements for compatibility in the case of $|\Delta m^2| \equiv |\Delta\bar{m}^2|$, there is a level of agreement with the mass differences of

$$\langle \Delta m_{32}^2 - \Delta\bar{m}^2 \rangle = 0.16 \pm 0.33 \times 10^{-3} \text{ eV}^2,$$

and, when comparing with $\sin^2(2\theta_{23}) \equiv \sin^2(2\bar{\theta}_{23})$,

$$\langle \sin^2(2\theta_{23}) - \sin^2(2\bar{\theta}_{23}) \rangle = 0.01 \pm 0.11.$$

Chapter 8

Conclusion

In this thesis, 10.49×10^{20} POT of ν_μ -dominant and $\bar{\nu}_\mu$ -enhanced NuMI beam data, collected over a period of 5 years using the magnetised MINOS detectors was analysed. Direct disappearance of $\bar{\nu}_\mu$ was observed, resulting in a measurement of the atmospheric-scale antineutrino oscillation parameters:

$$|\Delta \bar{m}^2| = 2.58_{-0.17}^{+0.26} \times 10^{-3} \text{ eV}^2,$$
$$\sin^2(2\bar{\theta}) = 0.96_{-0.10}^{+0.04}.$$

This measurement of $|\Delta \bar{m}^2|$ is the most precise ever made, surpassing the Super-K result by a factor of about five. The measurement of $\sin^2(2\bar{\theta})$ by Super-K remains the world's most precise, although accelerator experiments are likely to substantially improve on this measurement in the future.

This result allows the world's most precise comparison of $|\Delta m^2|$ and $|\Delta \bar{m}^2|$, with the worlds best measurement of $|\Delta m^2|$, from [28], as:

$$|\Delta m_{32}^2| = (2.38 \pm 0.085) \times 10^{-3} \text{ eV}^2,$$

giving a level of agreement with the mass differences of

$$\langle \Delta m_{32}^2 - \Delta \bar{m}_{32}^2 \rangle = 0.20 \pm 0.32 \times 10^{-3} \text{ eV}^2.$$

Thus, no evidence for a difference between neutrino and antineutrino atmospheric-scale oscillation parameters was found.

In the near future, this measurement is likely to be improved by the current long baseline neutrino oscillation experiments $\text{NO}\nu\text{A}$ and T2K. T2K is currently running in $\bar{\nu}_\mu$ -enhanced mode, and expected to be able to make a competitive measurement soon. $\text{NO}\nu\text{A}$ is likely to switch their beam to $\bar{\nu}_\mu$ within the next year. In the longer term, the DUNE and Hyper-K experiments will be able to make significant improvements to these measurements, and are expected to start in the 2020s.

Bibliography

- [1] P. Adamson et al. “Search for the disappearance of muon antineutrinos in the NuMI neutrino beam”. In: *Phys. Rev. D* 84 (7 Oct. 2011), p. 071103. DOI: 10.1103/PhysRevD.84.071103.
- [2] P. Adamson et al. “First direct observation of muon antineutrino disappearance”. In: *Phys.Rev.Lett.* 107 (2011), p. 021801. DOI: 10.1103/PhysRevLett.107.021801. arXiv: 1104.0344 [hep-ex].
- [3] M.C. Gonzalez-Garcia and Michele Maltoni. “Phenomenology with Massive Neutrinos”. In: *Phys.Rept.* 460 (2008), pp. 1–129. DOI: 10.1016/j.physrep.2007.12.004. arXiv: 0704.1800 [hep-ph].
- [4] J. Chadwick. “Intensitätsverteilung im magnetischen Spektren der β -Strahlen von Radium B + C”. German. In: *Verh.Phys.Gesell.* 16 (1914), pp. 383–391.
- [5] C. D. Ellis and W. A. Wooster. “The Continuous Spectrum of β -Rays”. In: *Nature* (1927-04-16).
- [6] Wolfgang Pauli. *Pauli letter collection: letter to Lise Meitner*. Typed copy. 1930. URL: <https://cds.cern.ch/record/83282>.
- [7] J. Chadwick. “Possible Existence of a Neutron”. In: *Nature* 129 (1932), p. 312. DOI: 10.1038/129312a0.
- [8] E. Fermi. “Versuch einer Theorie der β -Strahlen. I”. German. In: *Zeitschrift für Physik* 88.3-4 (1934), pp. 161–177. ISSN: 0044-3328. DOI: 10.1007/BF01351864.
- [9] Carlo Giunti and W. Chung Kim. *Fundamentals of neutrino physics and astrophysics*. Oxford: Oxford University Press, 2007. ISBN: 978-0198508717.
- [10] H. Bethe and R. Peierls. “The ‘neutrino’”. In: *Nature* 133 (1934), p. 532. DOI: 10.1038/133532a0.
- [11] F. Reines and C.L. Cowan. “Detection of the Free Neutrino”. In: *Phys. Rev.* 92 (1953), pp. 830–831. DOI: 10.1103/PhysRev.92.830.
- [12] C.L. Cowan et al. “Detection of the Free Neutrino: A Confirmation”. In: *Science* 124 (1956), pp. 103–104. DOI: 10.1126/science.124.3212.103.
- [13] J. C. Street and E. C. Stevenson. “New Evidence for the Existence of a Particle of Mass Intermediate Between the Proton and Electron”. In: *Phys. Rev.* 52 (9 Nov. 1937), pp. 1003–1004. DOI: 10.1103/PhysRev.52.1003.
- [14] G. Danby et al. “Observation of High-Energy Neutrino Reactions and the Existence of Two Kinds of Neutrinos”. In: *Phys. Rev. Lett.* 9 (1 July 1962), pp. 36–44. DOI: 10.1103/PhysRevLett.9.36.

- [15] M. L. Perl et al. "Evidence for Anomalous Lepton Production in $e^+ - e^-$ Annihilation". In: *Phys. Rev. Lett.* 35 (22 Dec. 1975), pp. 1489–1492. DOI: 10.1103/PhysRevLett.35.1489.
- [16] D. DeCamp et al. "Determination of the number of light neutrino species". In: *Physics Letters B* 231.4 (1989), pp. 519–529. ISSN: 0370-2693. DOI: [http://dx.doi.org/10.1016/0370-2693\(89\)90704-1](http://dx.doi.org/10.1016/0370-2693(89)90704-1).
- [17] Fermilab. *Physicists Find First Direct Evidence for Tau Neutrino at Fermilab*. Press Release. July 20, 2000.
- [18] K. Kodama et al. "Observation of tau neutrino interactions". In: *Phys.Lett. B* 504 (2001), pp. 218–224. DOI: 10.1016/S0370-2693(01)00307-0. arXiv: hep-ex/0012035 [hep-ex].
- [19] Ziro Maki, Masami Nakagawa, and Shoichi Sakata. "Remarks on the unified model of elementary particles". In: *Prog.Theor.Phys.* 28 (1962), pp. 870–880. DOI: 10.1143/PTP.28.870.
- [20] B. Pontecorvo. "Mesonium and anti-mesonium". In: *Sov.Phys.JETP* 6 (1957), p. 429.
- [21] Nicola Cabibbo. "Unitary Symmetry and Leptonic Decays". In: *Phys.Rev.Lett.* 10 (1963), pp. 531–533. DOI: 10.1103/PhysRevLett.10.531.
- [22] Makoto Kobayashi and Toshihide Maskawa. "CP Violation in the Renormalizable Theory of Weak Interaction". In: *Prog.Theor.Phys.* 49 (1973), pp. 652–657. DOI: 10.1143/PTP.49.652.
- [23] F.P. An et al. "Observation of electron-antineutrino disappearance at Daya Bay". In: *Phys.Rev.Lett.* 108 (2012), p. 171803. DOI: 10.1103/PhysRevLett.108.171803. arXiv: 1203.1669 [hep-ex].
- [24] *Spinors, Chirality, and Majorana Mass*. Jan. 4, 2008. URL: <https://fliptomato.wordpress.com/2008/01/04/spinors-chirality-and-majorana-mass/> (visited on 2015-11-12).
- [25] Flip Tanedo. *Helicity, Chirality, Mass, and the Higgs*. June 19, 2011. URL: <http://www.quantumdiaries.org/2011/06/19/helicity-chirality-mass-and-the-higgs/> (visited on 2015-11-12).
- [26] M. Goldhaber, L. Grodzins, and A. W. Sunyar. "Helicity of Neutrinos". In: *Phys. Rev.* 109 (3 Feb. 1958), pp. 1015–1017. DOI: 10.1103/PhysRev.109.1015.
- [27] R. F. Streater. *PCT, spin and statistics, and all that*. Princeton, N.J: Princeton University Press, 2000. ISBN: 978-0691070629.
- [28] P. Adamson et al. "Measurement of Neutrino and Antineutrino Oscillations Using Beam and Atmospheric Data in MINOS". In: *Phys. Rev. Lett.* 110 (25 June 2013), p. 251801. DOI: 10.1103/PhysRevLett.110.251801.
- [29] Raymond Davis. "Solar Neutrinos. II. Experimental". In: *Phys. Rev. Lett.* 12 (11 Mar. 1964), pp. 303–305. DOI: 10.1103/PhysRevLett.12.303.
- [30] V. Castellani et al. "Solar neutrinos: Beyond standard solar models". In: *Phys.Rept.* 281 (1997), pp. 309–398. DOI: 10.1016/S0370-1573(96)00032-4. arXiv: astro-ph/9606180 [astro-ph].

- [31] Bruce T. Cleveland et al. "Measurement of the Solar Electron Neutrino Flux with the Homestake Chlorine Detector". In: *The Astrophysical Journal* 496.1 (1998), p. 505.
- [32] Jr. Davis Raymond, Don S. Harmer, and Kenneth C. Hoffman. "Search for neutrinos from the sun". In: *Phys.Rev.Lett.* 20 (1968), pp. 1205–1209. DOI: 10.1103/PhysRevLett.20.1205.
- [33] John N. Bahcall, M. H. Pinsonneault, and G. J. Wasserburg. "Solar models with helium and heavy-element diffusion". In: *Rev. Mod. Phys.* 67 (4 Oct. 1995), pp. 781–808. DOI: 10.1103/RevModPhys.67.781.
- [34] K. S. Hirata et al. "Observation of ^8B solar neutrinos in the Kamiokande-II detector". In: *Phys. Rev. Lett.* 63 (1 July 1989), pp. 16–19. DOI: 10.1103/PhysRevLett.63.16.
- [35] Y. Fukuda et al. "Solar Neutrino Data Covering Solar Cycle 22". In: *Phys. Rev. Lett.* 77 (9 Aug. 1996), pp. 1683–1686. DOI: 10.1103/PhysRevLett.77.1683.
- [36] P. Anselmann et al. "Solar neutrinos observed by GALLEX at Gran Sasso." In: *Phys.Lett.* B285 (1992), pp. 376–389. DOI: 10.1016/0370-2693(92)91521-A.
- [37] J.N. Abdurashitov et al. "Solar neutrino flux measurements by the Soviet-American Gallium Experiment (SAGE) for half the 22 year solar cycle". In: *J.Exp.Theor.Phys.* 95 (2002), pp. 181–193. DOI: 10.1134/1.1506424. arXiv: astro-ph/0204245 [astro-ph].
- [38] V.A. Kuzmin. "Detection of solar neutrinos by means of the $^{71}\text{Ga}(\nu, e)^{71}\text{Ge}$ reaction". Trans. by J. G. Adashko. In: *Sov.Phys.JETP* 22.5 (1966), pp. 1051–1056.
- [39] John N. Bahcall and M. H. Pinsonneault. "What Do We (Not) Know Theoretically about Solar Neutrino Fluxes?" In: *Phys. Rev. Lett.* 92 (12 Mar. 2004), p. 121301. DOI: 10.1103/PhysRevLett.92.121301.
- [40] M. Altmann et al. "Complete results for five years of GNO solar neutrino observations". In: *Phys.Lett.* B616 (2005), pp. 174–190. DOI: 10.1016/j.physletb.2005.04.068. arXiv: hep-ex/0504037 [hep-ex].
- [41] J.N. Abdurashitov et al. "Measurement of the response of the Russian-American gallium experiment to neutrinos from a Cr-51 source". In: *Phys.Rev.* C59 (1999), pp. 2246–2263. DOI: 10.1103/PhysRevC.59.2246. arXiv: hep-ph/9803418 [hep-ph].
- [42] J. Boger et al. "The Sudbury neutrino observatory". In: *Nucl.Instrum.Meth.* A449 (2000), pp. 172–207. DOI: 10.1016/S0168-9002(99)01469-2. arXiv: nucl-ex/9910016 [nucl-ex].
- [43] N. Barros. "Final results from SNO". In: *Nucl.Phys.Proc.Suppl.* 237-238 (2013), pp. 107–110. DOI: 10.1016/j.nuclphysbps.2013.04.069.
- [44] John N. Bahcall, Aldo M. Serenelli, and Sarbani Basu. "New solar opacities, abundances, helioseismology, and neutrino fluxes". In: *Astrophys.J.* 621 (2005), pp. L85–L88. DOI: 10.1086/428929. arXiv: astro-ph/0412440 [astro-ph].

- [45] B. Aharmim et al. "Electron energy spectra, fluxes, and day-night asymmetries of B-8 solar neutrinos from measurements with NaCl dissolved in the heavy-water detector at the Sudbury Neutrino Observatory". In: *Phys.Rev.* C72 (2005), p. 055502. DOI: 10.1103/PhysRevC.72.055502. arXiv: nucl-ex/0502021 [nucl-ex].
- [46] S. Fukuda et al. "Determination of solar neutrino oscillation parameters using 1496 days of Super-Kamiokande I data". In: *Phys.Lett.* B539 (2002), pp. 179–187. DOI: 10.1016/S0370-2693(02)02090-7. arXiv: hep-ex/0205075 [hep-ex].
- [47] C. V. Achar et al. "Detection of muons produced by cosmic ray neutrinos deep underground". In: *Phys.Lett.* 18 (1965), pp. 196–199. DOI: 10.1016/0031-9163(65)90712-2.
- [48] F. Reines et al. "Evidence for high-energy cosmic ray neutrino interactions". In: *Phys.Rev.Lett.* 15 (1965), pp. 429–433. DOI: 10.1103/PhysRevLett.15.429.
- [49] F. Reines et al. "Muons produced by atmospheric neutrinos: Experiment". In: *Phys.Rev.* D4 (1971), pp. 80–98. DOI: 10.1103/PhysRevD.4.80.
- [50] T. J. Haines et al. "Calculation of Atmospheric Neutrino-Induced Backgrounds in a Nucleon-Decay Search". In: *Phys. Rev. Lett.* 57 (16 Oct. 1986), pp. 1986–1989. DOI: 10.1103/PhysRevLett.57.1986.
- [51] K.S. Hirata et al. "Experimental Study of the Atmospheric Neutrino Flux". In: *Phys.Lett.* B205 (1988), p. 416. DOI: 10.1016/0370-2693(88)91690-5.
- [52] M. Aglietta et al. "Experimental study of atmospheric neutrino flux in the NUSEX experiment". In: *Europhys.Lett.* 8 (1989), pp. 611–614. DOI: 10.1209/0295-5075/8/7/005.
- [53] Christoph Berger et al. "Study of Atmospheric Neutrino Interactions with the Frejus Detector". In: *Phys.Lett.* B227 (1989), p. 489. DOI: 10.1016/0370-2693(89)90968-4.
- [54] S. Fukuda et al. "The Super-Kamiokande detector". In: *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment* 501.2–3 (2003), pp. 418–462. ISSN: 0168-9002. DOI: [http://dx.doi.org/10.1016/S0168-9002\(03\)00425-X](http://dx.doi.org/10.1016/S0168-9002(03)00425-X).
- [55] Y. Fukuda et al. "Evidence for oscillation of atmospheric neutrinos". In: *Phys.Rev.Lett.* 81 (1998), pp. 1562–1567. DOI: 10.1103/PhysRevLett.81.1562. arXiv: hep-ex/9807003 [hep-ex].
- [56] Q.R. Ahmad et al. "Measurement of the rate of $\nu_e + d \rightarrow p + p + e^-$ interactions produced by ^8B solar neutrinos at the Sudbury Neutrino Observatory". In: *Phys.Rev.Lett.* 87 (2001), p. 071301. DOI: 10.1103/PhysRevLett.87.071301. arXiv: nucl-ex/0106015 [nucl-ex].
- [57] K. Abe et al. "Search for Differences in Oscillation Parameters for Atmospheric Neutrinos and Antineutrinos at Super-Kamiokande". In: *Phys. Rev. Lett.* 107 (24 Dec. 2011), p. 241801. DOI: 10.1103/PhysRevLett.107.241801.

- [58] F.P. An et al. "A side-by-side comparison of Daya Bay antineutrino detectors". In: *Nucl.Instrum.Meth.* A685 (2012), pp. 78–97. DOI: 10.1016/j.nima.2012.05.030. arXiv:1202.6181 [physics.ins-det].
- [59] J.K. Ahn et al. "Observation of Reactor Electron Antineutrino Disappearance in the RENO Experiment". In: *Phys.Rev.Lett.* 108 (2012), p. 191802. DOI:10.1103/PhysRevLett.108.191802. arXiv:1204.0626 [hep-ex].
- [60] K. Abe, J. Adam, et al. "Observation of Electron Neutrino Appearance in a Muon Neutrino Beam". In: *Phys. Rev. Lett.* 112 (6 Feb. 2014), p. 061802. DOI: 10.1103/PhysRevLett.112.061802.
- [61] D.S. Ayres et al. "NOvA: Proposal to build a 30 kiloton off-axis detector to study $\nu(\mu) \rightarrow \nu(e)$ oscillations in the NuMI beamline". In: (2004). arXiv: hep-ex/0503053 [hep-ex].
- [62] K. Scholberg. "The Long-Baseline Neutrino Experiment". In: *Nucl.Phys.Proc.Suppl.* 237-238 (2013), pp. 184–186. DOI: 10.1016/j.nuclphysbps.2013.04.085.
- [63] K. Abe et al. "Physics Potential of a Long Baseline Neutrino Oscillation Experiment Using J-PARC Neutrino Beam and Hyper-Kamiokande". In: *PTEP* (2015). arXiv: 1502.05199 [hep-ex].
- [64] J Hylen et al. *NuMI Technical Design Handbook*. Tech. rep. Internal NuMI report, 2002.
- [65] Bob Zwaska. "Accelerator Systems and Instrumentation for the NuMI Neutrino Beam". PhD thesis. The University of Texas at Austin, 2005.
- [66] A. Himmel. "Antineutrino Oscillations in the Atmospheric Sector". PhD thesis. California Institute of Technology, 2011.
- [67] A Abramov et al. *Dynamic Stress Calculations for ME and LE targets and Results of Prototyping for the LE Target*. Tech. rep. MINOS Internal Note 675, 2000.
- [68] J Evans. "Measuring Antineutrino Oscillations with the MINOS Experiment". PhD thesis. Jesus College, Oxford, 2008.
- [69] S. Van der Meer and B. de Raad. *Proposal for an Enhanced Neutrino Beam*. Tech. rep. CERN-61-07. CERN, 1961.
- [70] J Hylen. "NuMI Target Hall Update". Internal MINOS DocDB-3565. Sept. 2007.
- [71] D.G. Michael et al. "The Magnetized steel and scintillator calorimeters of the MINOS experiment". In: *Nucl.Instrum.Meth.* A596 (2008), pp. 190–228. DOI:10.1016/j.nima.2008.08.003. arXiv:0805.3170 [physics.ins-det].
- [72] J Hartnell. "Measurement of the Calorimetric Energy Scale in MINOS". PhD thesis. St John's College, Oxford, 2005.
- [73] MINOS Collaboration. *The MINOS Detectors Technical Design Report*. Tech. rep. NuMI-L-337. 1998.
- [74] P. Adamson et al. "The MINOS light injection calibration system". In: *Nucl. Instrum. Meth.* A492 (2002), pp. 325–343. DOI: 10.1016/S0168-9002(02)01293-7. arXiv: hep-ex/0204021 [hep-ex].

- [75] N Devenish. "Gain calibration- Run IV and Continuity". Internal MINOS DocDB-8263. Mar. 2010.
- [76] K.A. Olive et al. "Review of Particle Physics". In: *Chin.Phys.* C38 (2014), p. 090001. DOI: 10.1088/1674-1137/38/9/090001.
- [77] MINOS Collaboration. "2009 Position Paper on Calibration of Runs I-II-III". Internal MINOS DocDB-6717. Jan. 2010.
- [78] *FLUGG: FLUKA + Geant4 Geometry for Simulation in HEP*. 2003. URL: <http://www.fluka.org/content/tools/flugg/index.html>.
- [79] A. Fasso et al. *FLUKA: A Multi-particle Transport Code*. Tech. rep. CERN-2005-10 INFN/TC_05/11, SLAC-R-773. 2005.
- [80] G. Battistoni et al. "The FLUKA code: Description and benchmarking". In: *AIP Conference Proceeding* 896 (2007), pp. 37–49.
- [81] S. Agostinelli et al. "Geant4: a simulation toolkit". In: *Nucl. Instrumen. Meth.* A506.3 (2003), pp. 250–303. ISSN: 0168-9002. DOI: DOI : 10.1016/S0168-9002(03)01368-8.
- [82] H. Gallagher. "The NEUGEN neutrino event generator". In: *Nucl. Phys. B Proc.* 112 (2002), pp. 188–194. ISSN: 0920-5632. DOI: DOI : 10.1016/S0920-5632(02)01775-9.
- [83] Rene Brun, Federico Carminati, and Simone Giani. "GEANT Detector Description and Simulation Tool". In: (1994).
- [84] C. Zeitnitz and T.A. Gabriel. "The GEANT-CALOR interface and benchmark calculations of ZEUS test calorimeters". In: *Nucl. Instrumen. Meth.* A349.1 (1994), pp. 106–111. ISSN: 0168-9002. DOI: DOI : 10.1016/0168-9002(94)90613-0.
- [85] Mark Dorman. "Beam Fit Position Paper". Internal MINOS DocDB-7146. Apr. 2010.
- [86] R. Frühwirth. "Application of Kalman filtering to track and vertex fitting". In: *Nucl. Instrumen. Meth.* A262 (1987), pp. 444–450. ISSN: 0168-9002. DOI: 10.1016/0168-9002(87)90887-4.
- [87] R Toner. "Measuring Theta13 via Muon Neutrino to Electron Neutrino Oscillations in the MINOS Experiment". PhD thesis. University of Cambridge, 2011.
- [88] J P Ochoa. "A Search for Muon Neutrino to Electron Neutrino Oscillations in the MINOS Experiment". PhD thesis. California Institute of Technology, 2009.
- [89] P. Adamson et al. "Electron neutrino and antineutrino appearance in the full MINOS data sample". In: *Phys.Rev.Lett.* 110.17 (2013), p. 171801. DOI: 10.1103/PhysRevLett.110.171801. arXiv: 1301.4581 [hep-ex].
- [90] P. Adamson et al. "Active to sterile neutrino mixing limits from neutral-current interactions in MINOS". In: *Phys.Rev.Lett.* 107 (2011), p. 011802. DOI: 10.1103/PhysRevLett.107.011802. arXiv: 1104.3922 [hep-ex].
- [91] D. G. Michael et al. "Observation of Muon Neutrino Disappearance with the MINOS Detectors in the NuMI Neutrino Beam". In: *Phys. Rev. Lett.* 97.19 (Nov. 2006), p. 191801. DOI: 10.1103/PhysRevLett.97.191801.

- [92] P. Adamson et al. "Measurement of Neutrino Oscillations with the MINOS Detectors in the NuMI Beam". In: *Phys. Rev. Lett.* 101.13 (Sept. 2008), p. 131802. DOI: 10.1103/PhysRevLett.101.131802.
- [93] Philip A. Rodrigues. "A Sterile-Neutrino Search with the MINOS Experiment". PhD thesis. Lincoln College, Oxford, 2010.
- [94] P. Adamson et al. "Study of muon neutrino disappearance using the Fermilab Main Injector neutrino beam". In: *Phys. Rev. D* 77.7 (Apr. 2008), p. 072002. DOI: 10.1103/PhysRevD.77.072002.
- [95] K Nakamura and Particle Data Group. "Review of Particle Physics". In: *Journal of Physics G* 37.7A (2010), p. 075021.
- [96] Istvan Danko. Internal MINOS DocDB-4838. July 2008.
- [97] Andy Blake, Gwen Lefevre, and Jess Mitchell. "Data validation for the Near and Far Detector". Internal MINOS DocDB-6324. Sept. 2009.
- [98] A. Blake. "Cosmic Muon and Rock Muon Backgrounds in the CC Event Selection". Internal MINOS DocDB-3386. July 2007.
- [99] Rustem Ospanov. "A measurement of muon neutrino disappearance with the MINOS detectors and NuMI beam". PhD thesis. University of Texas at Austin, 2008.
- [100] R. Tibshirani T. Hastie and J. Friedman. *The Elements of Statistical Learning: Data Mining, Inference, and Prediction*. Springer, 2001. ISBN: 978-0387952840.
- [101] Jeff Hartnell. "NuMI Muon Antineutrinos in MINOS". In: *Fermilab Joint Experimental-Theoretical Seminar*. 2009.
- [102] N Devenish et al. "NuMuBar Selector Evaluation". Internal MINOS DocDB-4829. Dec. 2008.
- [103] P. Adamson et al. "Measurement of the neutrino mass splitting and flavor mixing by MINOS". In: *Phys.Rev.Lett.* 106 (2011), p. 181801. DOI: 10.1103/PhysRevLett.106.181801. arXiv: 1103.0340 [hep-ex].
- [104] The MINOS Calibration Group. "Run I-VII Calibration Position Paper". Internal MINOS DocDB-8263. July 2011.
- [105] Mike Kordosky, Hugh Gallagher, and Steve Dytman. "Shower Energy Scale Uncertainty for the Run I+II CC Analysis". Internal MINOS DocDB-4287. Mar. 2008.
- [106] Justin Evans et al. "Position paper on the analysis of the 2.95×10^{20} PoT $\bar{\nu}_\mu$ -beam dataset". Internal MINOS DocDB-8195. Feb. 2012.
- [107] J. M. Paley. "CC Handscan Results". Internal MINOS DocDB-5613. 2009.
- [108] Steve Baker and Robert D. Cousins. "Clarification of the use of CHI-square and likelihood functions in fits to histograms". In: *Nuclear Instruments and Methods in Physics Research* 221.2 (1984), pp. 437–442. ISSN: 0167-5087. DOI: DOI:10.1016/0167-5087(84)90016-4.
- [109] Gary J. Feldman and Robert D. Cousins. "Unified approach to the classical statistical analysis of small signals". In: *Phys. Rev. D* 57.7 (Apr. 1998), pp. 3873–3889. DOI: 10.1103/PhysRevD.57.3873.

- [110] D. S. Sivia. *Data Analysis: A Bayesian Tutorial*. Second. Oxford Science Publications, 2006. ISBN: 978-0198568322.
- [111] Igor Ostrovsky. *Gallery of Processor Cache Effects*. 2010-01-19. URL: <http://igoro.com/archive/gallery-of-processor-cache-effects/>.
- [112] J. von Neumann. "Various Techniques Used in Connection with Random Digits". In: *National Bureau of Standards Applied Mathematics Series 12* (1951), pp. 36–38.
- [113] R. Aggarwal and A. Caldwell. "Error Bars for Distributions of Numbers of Events". In: *Eur.Phys.J.Plus* 127 (2012), p. 24. DOI: 10.1140/epjp/i2012-12024-0. arXiv: 1112.2593 [physics.data-an].