

A Sterile-Neutrino Search with the MINOS Experiment

Philip A. Rodrigues

Lincoln College, Oxford



Thesis submitted in partial fulfilment of the requirements for the
degree of Doctor of Philosophy at the University of Oxford

Trinity Term, 2010

A Sterile-Neutrino Search with the MINOS Experiment

Philip A. Rodrigues
Lincoln College, Oxford

Thesis submitted in partial fulfilment of the requirements for the degree of
Doctor of Philosophy at the University of Oxford

Trinity Term, 2010

Abstract

The MINOS experiment is a long-baseline neutrino oscillation experiment in the NuMI beamline at Fermilab, USA. Using a near detector at 1 km distance from the neutrino production target, and a far detector at 735 km from the target, it is designed primarily to measure the disappearance of muon neutrinos.

This thesis presents an analysis using MINOS data of the possibility of oscillation of the neutrinos in the NuMI beam to a hypothetical sterile flavour, which would have no Standard Model couplings. Such oscillations would result in a deficit in the neutral current interaction rate in the MINOS far detector relative to the expectation derived from the near detector data.

The method used to identify neutral current and charged current events in the MINOS detectors is described and a new method of predicting and fitting the far detector spectrum presented, along with the effects of systematic uncertainties on the sterile neutrino oscillation analysis.

Using this analysis, the fraction f_s of the disappearing neutrinos that go to steriles is constrained to be below 0.15 at the 90% confidence level in the absence of electron neutrino appearance in the NuMI beam. With electron appearance at the CHOOZ limit, $f_s < 0.41$ at 90% C.L.

Acknowledgements

Like any doctoral thesis, this one owes its completion to the help and support of many people. First and foremost, my supervisor Alfons Weber. He has allowed me space to develop my ideas while always being available with an insightful comment or suggestion. Giles Barr also deserves thanks for much support and many helpful discussions.

I owe my induction into the practice of high-energy physics research to my MINOS colleagues at Oxford, who were generous with both their time and their expertise, and from whom I have learned a great deal: Chris Backhouse, Justin Evans, Jeff de Jong, Phill Litchfield, Robert Pittam, Tobi Raufer, Alex Sousa and Gemma Tinti – it was a pleasure working with all of you.

Amid the various frustrations I would never have stayed sane without good friends, so thanks Max, Lauren, Sarah, Amit, Alex, Iain, Vikki, Tony and everyone who threw a disc with me on Team Shark and OW! You guys made the last four years fun.

And lastly, thanks to my family, who I can always rely on for their love and support.

Contents

1	History and Theory of Neutrino Physics	3
1.1	The Theory of the Weak Interaction	6
1.2	The Types of Neutrino	9
1.3	Neutrino Mass and Oscillations	10
1.3.1	Two-flavour oscillations	13
1.3.2	Neutrino Oscillations in Matter	15
1.4	The Solar Neutrino Sector	19
1.5	The Atmospheric Neutrino Sector	22
1.6	Three-Flavour Oscillations	27
1.6.1	θ_{13} , \mathcal{CP} violation and the hierarchy	30
1.7	Sterile Neutrinos	32
1.7.1	The LSND Anomaly	33
1.7.2	Sterile Neutrinos in the Atmospheric Sector	34
1.7.3	Parametrizing Four-Flavour Oscillations	36
1.8	Non-oscillation Neutrino Physics	37
1.8.1	Absolute Neutrino Mass	37
1.8.2	Majorana Neutrinos	38
2	The MINOS Experiment	39
2.1	Overview	39
2.2	The NuMI beamline	40

2.3	The MINOS Detectors	41
2.3.1	Scintillators and Light Detection	45
2.3.2	Electronics	46
2.3.3	Triggering	48
2.3.4	Calibrating the Detectors	49
2.4	MINOS Event Simulation	52
2.5	MINOS Event Reconstruction	53
2.5.1	Flux tuning	54
2.6	MINOS Physics	56
3	Identifying NC and CC Events in the MINOS Detectors	59
3.1	Neutrino event types in MINOS	59
3.1.1	ν_μ CC events in the MINOS detectors	59
3.1.2	NC events in the MINOS detectors	60
3.1.3	ν_e CC events in the MINOS detectors	61
3.2	Event preselections	63
3.2.1	Obtaining a clean sample in the near detector	64
3.2.2	Obtaining a clean sample in the far detector	68
3.3	Separating NC and CC events using an ANN	72
3.3.1	Artificial neural networks for event classification	73
3.3.2	ANN input variables	75
3.3.3	ANN performance	78
4	Predicting and Fitting the Far Detector Spectrum	82
4.1	Blind analysis	82
4.2	Fitting with separate NC and CC spectra	83
4.3	An energy <i>vs</i> PID spectrum	84
4.3.1	Motivation	84
4.4	Predicting the far detector spectrum	86

4.4.1	Producing an oscillated prediction	92
4.5	Sensitivity to neutrino oscillation parameters	95
4.5.1	Statistical sensitivity	95
4.5.2	Reconstructing neutrino energy	97
4.5.3	Choice of binning in PID	98
4.6	Mock data studies	100
4.6.1	Input data sets	100
4.6.2	Results	103
4.7	Effect of ν_e appearance	104
5	Systematic Uncertainties	109
5.1	Sources of uncertainty	109
5.1.1	Relative near/far normalization	109
5.1.2	Shower energy scales	113
5.1.3	Muon Energy Scale	115
5.1.4	Misidentified components	115
5.1.5	Uncertainties due to preselections	119
5.1.6	Cross section uncertainties	120
5.1.7	Flux uncertainties	122
5.1.8	ANN shape uncertainty	123
5.2	Constraining the reconstruction uncertainty	125
5.2.1	Muon Removal	126
5.2.2	Calculating efficiencies	127
5.2.3	Results	132
5.3	Systematic shifts	133
5.4	Systematic uncertainties in the fit	136
5.4.1	Producing systematically shifted predictions	137
5.4.2	Effect on oscillation parameter limits	138

6	Results and Discussion	141
6.1	Fit results	141
6.1.1	Fitting for f_e	146
6.2	Future Directions	149
6.3	Summary	150
A	Mock Data Plots	151
	Bibliography	155

List of Figures

1.1	Neutrino CC and NC interactions in the Standard Model	6
1.2	Results from SNO on the solar flux of ν_e and $\nu_\mu + \nu_\tau$	22
1.3	KamLAND oscillation probability	23
1.4	Δm_\odot^2 and $\tan^2 \theta_\odot$ bounds from KamLAND and solar expts	24
1.5	L/E dependence of the atmospheric neutrino flux from SuperK	25
1.6	MINOS charged current results	28
1.7	Flavour content of neutrino mass states	30
1.8	LSND and MiniBooNE mixing parameters	34
2.1	Components of the NuMI beamline	41
2.2	NuMI beam configurations	42
2.3	MINOS detector schematic	43
2.4	Near detector planes	44
2.5	Far detector planes	44
2.6	MINOS scintillator strip arrangement	46
2.7	Change in MINOS detector response with time	50
2.8	Effect of flux tuning procedure	55
2.9	ν_μ survival probability at the far detector	56
2.10	MINOS ν_e appearance result	57
3.1	Typical ν_μ CC events in the MINOS far detector	60
3.2	Typical NC events in the MINOS far detector	62

3.3	A typical beam ν_e CC event in the MINOS far detector	63
3.4	A typical oscillated ν_e CC event in the MINOS far detector	64
3.5	Near detector reconstructed over true shower energy by position . . .	65
3.6	Split event in the near detector	66
3.7	Leakage event in the near detector	67
3.8	Slice pulse height fraction in the near detector	68
3.9	Maximum consecutive planes in the near detector	69
3.10	Effect of cleaning on NC-like spectrum	70
3.11	Effect of far detector preselection	73
3.12	Example of a feedforward ANN	74
3.13	Variables used in the ANN preclassification step	76
3.14	Input variables for the ANN	79
3.15	Output of the ANN in the near detector	81
3.16	ANN performance <i>vs</i> cut position	81
4.1	Visible energy spectrum of truly NC events with sterile mixing . . .	84
4.2	NC and CC fractions as a function of ANN output value	85
4.3	True energy distributions for different PID ranges	86
4.4	Geometrical effects of finite decay pipe and detector size	87
4.5	Data/MC ratios for selected events in the near detector	90
4.6	Near and far detector energy projections in various bins of PID . . .	91
4.7	Accuracy of the oscillation calculation approximation	94
4.8	f_s limit as a function of number of PID bins	99
4.9	f_s <i>vs</i> Δm_{atm}^2 contour compared to published analysis	99
4.10	Energy projections of spectrum used in sensitivity calculation . . .	101
4.11	Best-fit points for mock data set 1	105
4.12	Best-fit distributions for mock data set 1	106
4.13	Best-fit points for mock data set 2	107

4.14	Best-fit distributions for mock data set 2	108
4.15	PID distribution in the far detector with ν_e appearance	108
5.1	Shower energy scale uncertainty from hadronic shower modelling . .	114
5.2	Spectra used to estimate CC background uncertainty	118
5.3	Corrections to NC and CC components	118
5.4	Uncertainty on NC-like events from near detector preselection . . .	120
5.5	Uncertainty on NC-like events from far detector preselection	121
5.6	Flux uncertainty as a function of energy	123
5.7	Toy demonstration of ANN shape systematic	124
5.8	Normalized cumulative ANN output distributions in data and MC .	125
5.9	ANN shift value in ANN shape systematic	126
5.10	Event before and after muon removal	127
5.11	pre- and post-MRCC spectra	129
5.12	MRCC efficiencies in near and far detectors	130
5.13	Systematic shifts in fit parameters	134
5.14	Weights used for systematic interpolation	138
5.15	Effect of systematics on $\Delta\chi^2$ projections	140
6.1	Energy projections of FD data spectrum	143
6.2	FD data spectra integrated over PID and energy	144
6.3	$\Delta\chi^2$ projections from data fit	145
6.4	$\Delta\chi^2$ projection of f_e when fit	147
6.5	$\Delta\chi^2$ projection of f_s when f_e is fit	148
A.1	Sample of $(f_s, \Delta m^2)$ contours from mock data set 1	152
A.2	Sample of $(f_s, \Delta m^2)$ contours from mock data set 2	152
A.3	Example spectrum from mock data set 1	153
A.4	Example spectrum from mock data set 2	154

List of Tables

3.1	Fractions of events in each preclassification category in near detector MC.	81
4.1	$\Delta\chi^2$ values for various confidence levels	96
4.2	Mock data oscillation parameters	100
5.1	Contributions to normalization uncertainty pre-2010	110
5.2	Updated ND fiducial bias numbers with statistical uncertainty. . . .	112
5.3	NC spectrum-weighted efficiencies from MRCC	132
5.4	Input oscillation parameters for systematics studies	133
5.5	Systematic shifts in fit parameters	135
6.1	Best-fit values of oscillation and systematic parameters	142
6.2	Best-fit values of oscillation and systematic parameters with f_e fit .	148

Introduction

From both an intuitive and a scientific point of view, neutrinos are intriguing particles; neutral and weakly interacting, they pass largely unhindered through any material we are able to put in their way. We know that they have nonzero masses, but we are unable (so far) to measure them directly, and we do not know why the masses are so small in comparison with other fundamental particles.

The MINOS experiment is designed to elucidate some of the properties of neutrinos. Using the NuMI neutrino beam at Fermilab, and two detectors separated by 734 km, MINOS is able to study (among other things) the disappearance of muon neutrinos, the possible appearance of electron neutrinos and the possibility of oscillation to exotic sterile neutrinos, with no coupling to the Z .

It is the last of these that this thesis is concerned with. A mixing between the three known flavours of neutrino, which couple to the Z boson, and a fourth *sterile* flavour with no coupling to the Z would lead to a deficit of neutral current interactions in the MINOS far detector relative to the prediction produced from measurements at the near detector, where no oscillations are assumed to have occurred. This thesis presents an analysis that takes advantage of the continuous output of an artificial neural network which discriminates between neutral current and charged current events to produce a slightly improved sensitivity to sterile neutrino mixing.

This thesis is organised as follows: Chapter 1 gives a brief overview of the history of neutrino physics, with the experimental and theoretical results which

have led to the current picture of neutrino masses and oscillations. The theoretical motivation for sterile neutrinos and an empirical parametrization for oscillations between active and sterile neutrinos is also presented.

A description of the MINOS experiment is given in Chapter 2, starting with the NuMI beamline which produces the neutrino beam measured at the near and far detectors, which are also described.

Chapter 3 describes the methods used to distinguish charged current (CC) and neutral current (NC) interactions in the MINOS detectors from one another and from non-neutrino-induced backgrounds, and Chapter 4 presents a method of predicting and fitting the far detector spectrum taking advantage of the NC/CC separation available from the artificial neural network described in Chapter 3.

With this method of predicting the far detector spectrum, Chapter 5 presents the effects of a number of systematic uncertainties on the sensitivity to sterile neutrino oscillation parameters, and a new, reduced uncertainty on the near/far normalization is presented. Finally, in Chapter 6, the analysis is run on MINOS near and far detector data to determine limits on the parameters governing oscillations to sterile neutrinos.

Chapter 1

History and Theory of Neutrino Physics

Most histories of the neutrino begin with the observation that most histories of the neutrino begin with Wolfgang Pauli's famous 1930 letter to Lise Meitner and the participants of a conference at Tübingen, addressed to "Dear Radioactive Ladies and Gentlemen" [1]. In it, he proposed a solution to the problem of β -decay, in which observation of the continuous spectrum of the electron made it appear that energy was not conserved. While Niels Bohr had suggested that the result be taken at face value, and the law of energy conservation be weakened, Pauli postulated a different "dangerous remedy:" the existence of a new particle, existing in the nucleus, with spin $\frac{1}{2}$ and zero electric charge. The emission of these particles in addition to the electron in β -decay would allow energy conservation.

Pauli's *Neutronen*, as he called them, were not quite the neutrino we know today; rather, they played roles of both the neutrino and the neutron. What we now call the neutron was discovered in 1932 by James Chadwick [2], leading Enrico Fermi to rechristen the particle responsible for rescuing energy conservation the *neutrino*.

The theory of the interaction governing β -decay moved forward almost im-

mediately, with Fermi's theory being published in 1934 [3]. This theory put the neutrino on good theoretical ground, treating the β -decay as analogous to electromagnetic radiation, and allowing other, as yet unobserved, processes such as electron capture to be predicted [4]. Fermi's theory also provided background for Hans Bethe and Rudolf Peierls to calculate the cross section for the inverse β -decay reaction in which neutrinos could be directly detected [4]. They concluded that the cross section was far too small to ever allow the reaction to be observed, and indeed no strong direct experimental observation of the neutrino came for 22 years. The first direct observation of neutrinos in 1953 by Fred Reines and Clyde Cowan [5] was tentative, and the experiment was later moved from Hanford to Savannah River, where the desired confirmation was obtained in 1956 [6].

In both cases, the experiment searched for the inverse β -decay reaction of reactor antineutrinos, $\bar{\nu} + p \rightarrow n + e^+$, in cadmium chloride-loaded water targets. The final state positron annihilated immediately with an electron in the water, producing a pair of gamma rays, while the neutron was moderated by the water until being captured by the cadmium around $10\mu\text{s}$ later, producing an unstable nucleus which decayed giving off a gamma ray. The water tanks were sandwiched by liquid scintillator to detect the gamma rays, with the signature of the inverse β -decay reaction being a pair of signals in the scintillator (from the electron-positron annihilation), followed some microseconds later by a “delayed coincidence” signal from the neutron capture. Reines and Cowan observed a rate of neutrino interactions consistent with that expected from the reactor.

After the long gap between the neutrino's theoretical postulation and its experimental verification, the field advanced more quickly: in 1955 Ray Davis performed a search for the reaction $\bar{\nu} + {}^{37}\text{Cl} \rightarrow e^- + {}^{37}\text{Ar}$, which was not observed [7]. Davis pointed out that such a reaction would be possible if the neutrinos given off in β -decay and inverse β -decay were indistinguishable in their reactions. The null result of the experiment was able to set an upper bound on the $\bar{\nu}$ capture cross-

section, but the constraint was not strong enough to conclusively show that ν and $\bar{\nu}$ differed in their reactions.

Further elucidation of the nature of the neutrino was quick in coming: only a year after Reines and Cowan's Savannah River experiment, Maurice Goldhaber *et al.* demonstrated extremely elegantly that the neutrino has negative helicity [8]. They observed decays of ^{152}Eu , with spin 0, to $^{152}\text{Sm}^*$, with spin 1, via electron capture. Since the neutrino has spin $\frac{1}{2}$, conservation of angular momentum implies that its spin must be antiparallel to that of the $^{152}\text{Sm}^*$ nucleus. Since the nucleus and the neutrino have equal and opposite momenta, antiparallel spins lead to the nucleus having the same helicity as the neutrino. When the $^{152}\text{Sm}^*$ nuclei decay to ^{152}Sm by photon emission, those photons which are emitted parallel to the direction of motion of the nucleus then have the same helicity as the neutrino. These photons also have the useful property that they are slightly boosted by the recoil momentum of the nucleus, and therefore have exactly the right energy to be absorbed and rescattered in the process $^{152}\text{Sm} + \gamma \rightarrow ^{152}\text{Sm}^* \rightarrow ^{152}\text{Sm} + \gamma$. These *resonant scatters* therefore select photons with the same helicity as the neutrino. The final step was to measure the helicity of these resonantly scattered photons, which was done by passing them through magnetized iron. Absorbing the photon on an electron in the iron is possible if the photon and electron spins are antiparallel, but not if the spins are parallel. Therefore the amount of absorption of photons in the iron indicates the neutrino helicity. Goldhaber *et al.* found that the neutrino has negative helicity, i.e., spin antiparallel to its momentum.

On another experimental front, Lederman, Schwartz, Steinberger *et al.* demonstrated in 1962 that the neutrino produced in pion decay in association with a muon is different from the neutrino produced in association with an electron in β -decay [9]. They used 15 GeV protons from the AGS accelerator at Brookhaven, incident on a beryllium target. The pions produced in this interaction decayed in flight to $\mu\nu_\mu$, with the decay products passing through an absorber to produce




Figure 1.1: Charged current (left) and neutral current (right) interactions of the neutrino in the Standard Model.

a beam consisting only of neutrinos. Neutrino-induced events in a downstream spark chamber produced muons and not electrons, showing that ν_μ and ν_e are distinct particles.

1.1 The Theory of the Weak Interaction

Over the decades following the publication of Fermi’s theory of β -decay, the theory of the weak interaction developed along several fronts: the problem of unitarity violation at high energies was identified and solved by addition of a mediating boson, as suggested by Yukawa, and later extended by Glashow, Salam and Weinberg; the violations of parity and charge-parity were discovered and incorporated into the theory; and the existence of a weak neutral current interaction was posited and confirmed.

The first of these developments began with the observation by Heisenberg in 1936 that the Fermi theory violates unitarity at high energies [10]. This problem could be ameliorated using Hideki Yukawa’s concept of an exchange boson [11], and this concept was extended in a series of papers by Sheldon Glashow, Abdus Salam and Steven Weinberg in the 1960s [12, 13, 14]. In these papers, the weak interaction was unified with the electromagnetic interaction, and a new, neutral boson was predicted, having couplings to all fermions.

There were therefore two sets of predictions of the Glashow-Salam-Weinberg

model: firstly, the existence of massive gauge bosons, whose masses were predicted in Weinberg's paper [14]; and secondly, the existence of *neutral current* (NC) interactions mediated by the new neutral boson, as shown in Figure 1.1. In 1973, the first experimental verification of neutral current interactions was made in the Gargamelle bubble chamber at CERN, where the NC electron-scattering process $\nu_\mu + e^- \rightarrow \nu_\mu + e^-$ and the NC hadronic interaction $\nu_\mu + N \rightarrow \nu_\mu + X$ were observed [15, 16]. Direct observation of the massive charged and neutral bosons (named the W and Z respectively) would have to wait until the next decade, when the UA1 and UA2 experiments on the CERN Sp \bar{p} S saw both particles and made measurements of their masses [17, 18, 19, 20].

A second thread of understanding of the nature of the weak interaction developed alongside this, namely its properties with respect to the discrete symmetries of parity, \mathcal{P} , and charge conjugation, \mathcal{C} . For some time, both were considered inviolate, until in 1956 Tsung-Dao Lee and Chen-Ning Yang pointed out that parity conservation in weak decays was “an extrapolated hypothesis unsupported by experimental evidence” [21]. They suggested a simple experiment to unambiguously search for parity violation in the weak interaction: a measurement of the angular distribution between the nuclear spin $\boldsymbol{\sigma}$ of a β -decaying nucleus and the momentum \mathbf{p} of the electron given off in the decay. If a component of the distribution proportional to the pseudoscalar $\boldsymbol{\sigma} \cdot \mathbf{p}$ is found, then parity is not conserved, since \mathbf{p} changes sign under parity, while $\boldsymbol{\sigma}$ does not.

The experiment was realized the same year and was published in 1957. Chen-Shiung Wu used a sample of ^{60}Co with their spins aligned by a magnetic field at low temperature, and observed the angle of the β -decay electrons emitted relative to the nuclear spin, θ . An observed asymmetry in the distribution between θ and $180^\circ - \theta$ indicated the presence of a pseudoscalar term, and therefore parity violation.

Incorporating parity violation into the theoretical description of the weak in-

teraction was done in 1957 by a number of authors [22, 23]. They added a term corresponding to an axial component to the existing vector term in the description of the weak current, leading to the so-called $V - A$ theory. This predicts maximal parity violation, and the correct helicity for the neutrino.

While the left-handedness of neutrinos demonstrates parity violation, in accordance with the $V - A$ theory, the combined operation of charge conjugation and parity, \mathcal{CP} , leaves the system unchanged, since antineutrinos are right-handed. \mathcal{CP} is known to be violated in other weak interactions, however: neutral kaons can decay via the weak interaction into either two or three pions in states which are eigenstates of \mathcal{CP} . The 2π decay mode occurs much more quickly due to the larger phase space, so a beam of neutral kaons allowed to decay will eventually produce only 3π decays, assuming \mathcal{CP} is conserved. However, in 1964, James Cronin *et al.* performed exactly this experiment and observed a peak in the invariant mass of two of the pions from the decay: a characteristic of a two-body decay [24]. Therefore \mathcal{CP} is violated in the weak decays of neutral kaons.

While \mathcal{CP} is conserved in β -decay, the question of whether \mathcal{CP} is conserved throughout the whole neutrino sector is an open one that will be dealt with in Section 1.3.

One further discrete symmetry is of interest: that of time reversal \mathcal{T} . It is a result of quantum field theory that any local, unitary, Lorentz invariant field theory must be invariant under \mathcal{CPT} [25]. Therefore any violation of \mathcal{CP} must be accompanied by a corresponding violation of \mathcal{T} to ensure that \mathcal{CPT} is conserved. In the neutrino sector, tests of \mathcal{CPT} can be made by comparing oscillation probabilities which are \mathcal{CPT} conjugates of one another, such as $P(\nu_\alpha \rightarrow \nu_\beta)$ and $P(\bar{\nu}_\beta \rightarrow \bar{\nu}_\alpha)$, or $P(\nu_\alpha \rightarrow \nu_\alpha)$ and $P(\bar{\nu}_\alpha \rightarrow \bar{\nu}_\alpha)$, and a number of experiments have made such tests [26].

1.2 The Types of Neutrino

In 1953, Konopinski and Mahmoud introduced the concept of conserved lepton number, L , in which the negative leptons and their associated neutrinos are assigned $L = +1$, while their antiparticles are assigned $L = -1$ [27]. This conservation law served to explain the null result in Davis's search for $\bar{\nu} + {}^{37}\text{Cl} \rightarrow e^- + {}^{37}\text{Ar}$, since the hypothesized process violates lepton number.

The Brookhaven experiment of Lederman *et al.* [9] showed that there were at least two types of neutrino, each associated with a particular lepton. Along with the non-observation of the process $\mu \rightarrow e\gamma$, which is allowed by L , this provided evidence for lepton *family number*, which is conserved separately for each family (lepton and associated neutrino).

The discovery of the τ lepton in e^+e^- collisions in 1975 [28] started the search for a corresponding neutrino ν_τ for symmetry with the already-known leptons. The difficulties of producing the tau neutrino and unambiguously identifying its interactions meant that the first strong evidence came indirectly from the LEP collider at CERN [29]. By comparing the total width of the Z resonance to the sum of the partial widths to visible final states, the partial width to 'invisible' states (i.e., neutrinos) can be inferred. Comparing this invisible width to the theoretical prediction for the Z partial width for one neutrino yields the number of neutrinos that have mass less than $M_Z/2$ and couple to the Z , which is found to be 2.9840 ± 0.0082 , in excellent agreement with the number of known generations of charged fermions.

Direct observation of the tau neutrino has since been made by the DONUT collaboration at Fermilab [30]. A ν_τ -enriched beam was created and impinged on a detector consisting of emulsion plates sandwiched between steel sheets, along with various active detectors for triggering and identifying which emulsions to remove and develop. Selecting charged current (CC) ν_τ interactions requires identifying

the presence of a τ lepton. At DONUT energies, the τ decays with a path length ~ 2 mm, necessitating a detector with extremely high position resolution as well as the large mass always required for neutrino interaction experiments. This difficulty made emulsion the only practical choice for the detector technology used. In their final results, the DONUT collaboration reported 9 ν_τ CC candidates above a background of 1.5 events [31].

1.3 Neutrino Mass and Oscillations

One of the largest and most productive avenues of investigation in neutrino physics is the phenomenon of neutrino flavour change, almost certainly due to quantum mechanical oscillations. (Regardless of the mechanism, flavour change implies a nonzero neutrino mass, since in order to change, the neutrino must have a rest frame.)

In the oscillation scheme, neutrinos are massive (although the masses may still be extremely small), and the mass eigenstates are not equivalent to the flavour states. Since neutrinos are produced and detected in weak interaction processes that pick out flavour eigenstates, whereas propagation is controlled by the mass eigenstates, a neutrino born as one flavour may, after propagation, become a neutrino of another flavour.

The standard treatment of neutrino oscillations in vacuum is presented below, following [32]. A more correct treatment is given by the use of wave packets, or even quantum field theory, but this simpler approach reproduces the same results [33].

We begin with a neutrino flavour state $|\nu_\alpha\rangle$ that is a superposition of mass states $|\nu_i\rangle$ related by a matrix U :

$$|\nu_\alpha\rangle = \sum_i U_{\alpha i}^* |\nu_i\rangle. \quad (1.1)$$

This matrix is known as the PMNS matrix, for Maki, Nakagawa and Sakata, who were instrumental in formulating the theory of neutrino oscillations, and Pontecorvo, who first posited the change in flavour of massive neutrinos [34, 35]. For N neutrino states it is an $N \times N$ matrix satisfying the unitary condition $U^\dagger U = \mathbb{1}$. The relation between flavour and mass states for antineutrinos is given by Equation 1.1 with $U_{\alpha i}^* \rightarrow U_{\alpha i}$.

From this equation we can find the amplitude $\mathcal{A}(\nu_\alpha \rightarrow \nu_\beta)$ for a neutrino created in flavour state α to be detected later in a flavour state β

$$\mathcal{A}(\nu_\alpha \rightarrow \nu_\beta) \equiv \langle \nu_\beta | \nu_\alpha \rangle \quad (1.2)$$

$$= \sum_i \sum_j \langle \nu_j | U_{\beta j} \mathcal{P} U_{\alpha i}^* | \nu_i \rangle \quad (1.3)$$

$$= \sum_i U_{\beta i} \mathcal{P} U_{\alpha i}^* \quad \text{since } |\nu_i\rangle\text{s orthogonal} \quad (1.4)$$

where \mathcal{P} is the propagator for the appropriate mass eigenstate, $\mathcal{P} = \exp(-ip \cdot x)$ where p and x are the momentum and spacetime four-vectors of the mass eigenstate, respectively. In the lab frame, $p \cdot x = Et - |\mathbf{p}|L$, where t is the time of measurement, \mathbf{p} is the momentum three-vector, and L is the distance between source and measurement. Since neutrino masses are very small, $v \approx c$ and so $t \approx L$ in natural units, giving $p \cdot x = (E - |\mathbf{p}|)L$. Using the relation $E^2 - |\mathbf{p}|^2 \equiv (E - |\mathbf{p}|)(E + |\mathbf{p}|) = m^2$, we obtain

$$p \cdot x \approx \frac{m^2 L}{E + |\mathbf{p}|} \quad (1.5)$$

$$\approx \frac{m^2 L}{2E} \quad (1.6)$$

where in the second step, $E \approx |\mathbf{p}|$ for ultrarelativistic neutrinos has been used.

We can put this expression for the propagator back into the amplitude, and take its modulus-squared to obtain the probability to observe a neutrino originally

of flavour ν_α in flavour state ν_β after propagating a distance L :

$$P(\nu_\alpha \rightarrow \nu_\beta) = \sum_i \sum_j U_{\beta j}^* U_{\alpha j} e^{im_j^2 L/2E} U_{\beta i} U_{\alpha i}^* e^{im_i^2 L/2E} \quad (1.7)$$

$$= \sum_i \sum_j U_{\beta j}^* U_{\alpha j} U_{\beta i} U_{\alpha i}^* e^{i\Delta m_{ij}^2 L/2E} \quad (1.8)$$

where $\Delta m_{ij}^2 \equiv m_j^2 - m_i^2$. With a little manipulation this can be rearranged into a form that is easier to deal with. We add and subtract a carefully chosen term from the right-hand side and regroup the resulting terms:

$$\begin{aligned} P(\nu_\alpha \rightarrow \nu_\beta) &= \sum_i \sum_j U_{\beta j}^* U_{\alpha j} U_{\beta i} U_{\alpha i}^* e^{i\Delta m_{ij}^2 L/2E} \\ &\quad + \sum_i \sum_j U_{\beta j}^* U_{\alpha j} U_{\beta i} U_{\alpha i}^* - \sum_i \sum_j U_{\beta j}^* U_{\alpha j} U_{\beta i} U_{\alpha i}^* \end{aligned} \quad (1.9)$$

$$\begin{aligned} &= \sum_i \sum_j U_{\beta j}^* U_{\alpha j} U_{\beta i} U_{\alpha i}^* \left(e^{i\Delta m_{ij}^2 L/2E} - 1 \right) \\ &\quad + \sum_i \sum_j U_{\beta j}^* U_{\alpha j} U_{\beta i} U_{\alpha i}^* \end{aligned} \quad (1.10)$$

$$= \sum_i \sum_j U_{\beta j}^* U_{\alpha j} U_{\beta i} U_{\alpha i}^* \left(e^{i\Delta m_{ij}^2 L/2E} - 1 \right) + \delta_{\alpha\beta} \quad (1.11)$$

where the last term uses the unitarity of U , i.e., $\sum_i U_{\alpha i}^* U_{\beta i} = \delta_{\alpha\beta}$. Call the first term on the right hand side B_{ij} . Then terms B_{ii} are identically zero, since $\Delta m_{ii}^2 = 0$, and terms with $i \neq j$ satisfy $B_{ij} = B_{ji}^*$ and therefore $B_{ij} + B_{ji} = 2\Re(B_{ij})$.

Putting all this together and using Euler's formula gives

$$P(\nu_\alpha \rightarrow \nu_\beta) = \delta_{\alpha\beta} + 2 \sum_{i>j} \Re [U_{\beta j}^* U_{\alpha j} U_{\beta i} U_{\alpha i}^*] \times (\cos \Delta m_{ij}^2 L/2E + i \sin \Delta m_{ij}^2 L/2E - 1) \quad (1.12)$$

$$= \delta_{\alpha\beta} + 2 \sum_{i>j} \Re [U_{\beta j}^* U_{\alpha j} U_{\beta i} U_{\alpha i}^*] (\cos \Delta m_{ij}^2 L/2E - 1) - 2 \sum_{i>j} \Im [U_{\beta j}^* U_{\alpha j} U_{\beta i} U_{\alpha i}^*] \sin \Delta m_{ij}^2 L/2E \quad (1.13)$$

$$= \delta_{\alpha\beta} + 4 \sum_{i>j} \Re [U_{\beta j}^* U_{\alpha j} U_{\beta i} U_{\alpha i}^*] \sin^2 \Delta m_{ij}^2 L/2E - 2 \sum_{i>j} \Im [U_{\beta j}^* U_{\alpha j} U_{\beta i} U_{\alpha i}^*] \sin \Delta m_{ij}^2 L/2E \quad (1.14)$$

1.3.1 Two-flavour oscillations

The simplest nontrivial case is for two neutrinos. This is particularly interesting as it is a limiting case for other schemes under certain conditions. It also illustrates many salient features relevant to neutrino oscillation experiments.

The mixing matrix in two dimensions can be parametrized by a single *mixing angle* θ , thus:

$$U = \begin{pmatrix} \cos \theta & \sin \theta \\ -\sin \theta & \cos \theta \end{pmatrix}. \quad (1.15)$$

The oscillation probabilities become

$$P(\nu_\alpha \rightarrow \nu_\beta) = \sin^2 2\theta \sin^2(1.27 \Delta m^2 L/E) \quad (1.16)$$

for $\nu_\alpha \neq \nu_\beta$ (the *appearance probability*) and

$$P(\nu_\alpha \rightarrow \nu_\alpha) = 1 - \sin^2 2\theta \sin^2(1.27 \Delta m^2 L/E) \quad (1.17)$$

for the *survival probability*. In these equations, Δm^2 is given in eV^2/c^4 , L is given

in km and E in GeV, which accounts for the factor of 1.27.

Some features of neutrino oscillations follow immediately from these equations:

- No oscillation will be observed when the mixing angle θ is zero. In this case the mass and flavour eigenstates coincide. The largest effect will be seen when $\sin^2 2\theta = 1$, a case referred to as *maximal mixing*.
- The largest flavour change effect occurs when $1.27\Delta m^2 L/E = \pi/2$. The experimental parameter L/E therefore determines the mass-squared splitting to which an experiment is sensitive.
- In any case where $\Delta m^2 L/E$ is large, the oscillation probability varies rapidly with energy, and cannot be resolved by an experiment at a fixed baseline L . When this is the case, the measured appearance probability averages to $\frac{1}{2} \sin^2 2\theta$ and the survival probability to $1 - \frac{1}{2} \sin^2 2\theta$.
- Since $\sin^2(1.27\Delta m^2 L/E)$ is an even function of Δm^2 , the oscillation probabilities for Δm^2 and $-\Delta m^2$ are equivalent, and so two-neutrino oscillations contain no information about the ordering of the two mass states.

There are two different cases in which the two-flavour oscillation formula is regained from the general n flavour formula. The first occurs when the effect of one of the squared-mass splittings dominates over the others. In this case, it can be shown [32] that the oscillation probability reduces to the two-flavour formula with the replacements

$$\sin^2 2\theta \equiv 4 \left| \sum_{i \in \text{Up}} U_{\alpha i}^* U_{\beta i} \right|^2 \quad (1.18)$$

for appearance experiments, and

$$\sin^2 2\theta \equiv 4 \sum_{i \in \text{Up}} |U_{\alpha i}|^2 \left(1 - \sum_{i \in \text{Up}} |U_{\alpha i}|^2 \right) \quad (1.19)$$

for disappearance. In these formulae, the “ i Up” denotes summation over either the set of states lying above the large mass splitting, or the set of states lying below it.

The two-flavour oscillation formula is also regained if the flavour states being considered are coupled to only two of the mass eigenstates.

The last four decades of neutrino oscillation experiments have, with one exception¹, been consistent with the existence of three neutrino states, as measured by LEP, separating into two quasi-independent two-flavour sectors. There is one additional feature of neutrino oscillations required to interpret the results so far, namely the propagation of neutrinos in matter.

1.3.2 Neutrino Oscillations in Matter

Neutrinos propagating in matter experience an alteration in the oscillation probabilities due to the different interactions of electron neutrinos and other flavours with the electrons and neutrons in the matter. All flavours experience coherent scattering via the neutral current interaction, while electron neutrinos experience an additional coherent scattering from the charged current interaction with the electrons. This alters the phase of electron neutrinos relative to the other flavours, and therefore changes the oscillation probabilities, a feature known as the MSW effect, for Mikheyev, Smirnov and Wolfenstein, who were the first to consider it [36, 37].

A sketch of the matter effects can be given in a two-flavour scenario, with a neutrino state

$$\begin{pmatrix} \nu_e \\ \nu_\alpha \end{pmatrix}. \quad (1.20)$$

The Hamiltonian in matter consists of two parts: the vacuum Hamiltonian with

¹The LSND experiment, Section 1.7.1

an additional potential caused by the matter

$$\mathcal{H}_M = \mathcal{H}_0 + V \quad (1.21)$$

where the vacuum Hamiltonian for the flavour state is

$$\mathcal{H}_0 = U^\dagger \begin{pmatrix} \frac{m_1^2}{2E} & 0 \\ 0 & \frac{m_2^2}{2E} \end{pmatrix} U \sim U^\dagger \begin{pmatrix} 0 & 0 \\ 0 & \frac{\Delta m^2}{2E} \end{pmatrix} U \quad (1.22)$$

where U is the two-flavour mixing matrix from Equation 1.15, and on the right of the \sim we have subtracted a multiple of the unit matrix, which has no effect on the physics.

The matter-induced potential V consists of the NC and CC parts V_{NC} and V_{CC} . The NC scattering depends on the number density of neutrons in the matter, having strength $\mp G_F N_n / \sqrt{2}$, where G_F is the Fermi constant, while the CC scattering has strength $\pm \sqrt{2} G_F N_e$ where N_e is the number density of electrons. In each case, the upper sign is taken for neutrinos and the lower sign for antineutrinos. Taking into account the types of neutrino on which these act, we find the matter potential

$$V = V_{\text{NC}} + V_{\text{CC}} \quad (1.23)$$

$$= \begin{pmatrix} \mp G_F N_n / \sqrt{2} & 0 \\ 0 & \mp G_F N_n / \sqrt{2} \end{pmatrix} + \begin{pmatrix} \pm \sqrt{2} G_F N_e & 0 \\ 0 & 0 \end{pmatrix}. \quad (1.24)$$

The term V_{NC} is a multiple of the unit matrix, and will therefore be dropped,

giving the matter Hamiltonian

$$\mathcal{H}_M = U^\dagger \begin{pmatrix} 0 & 0 \\ 0 & \frac{\Delta m^2}{2E} \end{pmatrix} U + \begin{pmatrix} \pm\sqrt{2}G_F N_e & 0 \\ 0 & 0 \end{pmatrix} \quad (1.25)$$

$$= \frac{\Delta m^2}{2E} \begin{pmatrix} \sin^2 \theta + \epsilon & -\sin \theta \cos \theta \\ -\sin \theta \cos \theta & \cos^2 \theta \end{pmatrix} \quad (1.26)$$

where $\epsilon = \pm\sqrt{2}G_F N_e/(\Delta m^2/2E)$.

If we can put this Hamiltonian into the form

$$\mathcal{H}_M = U_M^\dagger \begin{pmatrix} 0 & 0 \\ 0 & \frac{m_M^2}{2E} \end{pmatrix} U_M \quad (1.27)$$

$$= \frac{\Delta m_M^2}{2E} \begin{pmatrix} \sin^2 \theta_M & -\sin \theta_M \cos \theta_M \\ -\sin \theta_M \cos \theta_M & \cos^2 \theta_M \end{pmatrix} \quad (1.28)$$

where U_M is U with θ replaced by θ_M , we will have the same form as vacuum oscillations, but with altered mass splitting and mixing angle. Comparing the off-diagonal terms in Equations 1.26 and 1.28 we see that

$$\frac{\Delta m_M^2}{2E} \sin \theta_M \cos \theta_M = \frac{\Delta m^2}{2E} \sin \theta \cos \theta. \quad (1.29)$$

To obtain the effective mass splitting in matter, we find the difference in the eigenvalues of the matter Hamiltonian by diagonalizing the matrix in Equation 1.26. By comparison with Equation 1.22 we see that this difference will be equal to $\Delta m_M^2/2E$. The eigenvalues λ are given by

$$\lambda = \frac{\Delta m^2}{2E} \cdot \frac{1}{2} \left(\epsilon + 1 \pm \sqrt{(1 + \epsilon)^2 - 4\epsilon \cos^2 \theta} \right) \quad (1.30)$$

so that

$$\Delta m_M^2 = \Delta m^2 \sqrt{(1 + \epsilon)^2 - 4\epsilon \cos^2 \theta} \quad (1.31)$$

$$= \Delta m^2 \sqrt{(\epsilon - \cos 2\theta)^2 + \sin^2 2\theta} \quad (1.32)$$

and using Equation 1.29

$$\sin \theta_M \cos \theta_M \sqrt{(\epsilon - \cos 2\theta)^2 + \sin^2 2\theta} = \sin \theta \cos \theta \quad (1.33)$$

which can be rearranged to

$$\sin^2 2\theta_M = \frac{\sin^2 2\theta}{(\epsilon - \cos 2\theta)^2 + \sin^2 2\theta}. \quad (1.34)$$

There are several interesting experimentally observable consequences for neutrinos which propagate in matter:

- The mixing parameters depend on the energy of the neutrino via the parameter $\epsilon = \pm \sqrt{2} G_F N_e / (\Delta m^2 / 2E)$. Thus, the behaviour of low and high energy neutrinos may differ.
- Since ϵ changes sign for antineutrinos, the behaviour of neutrinos and antineutrinos will be different.
- ϵ also depends on the sign of Δm^2 , and not just its absolute value (as is the case in vacuum oscillations). This means that the ordering of the mass states may be determined.
- If $\epsilon = \cos 2\theta$, then $\sin^2 2\theta_M$ will be maximal regardless of the value of $\sin^2 2\theta$. That is, a small vacuum mixing can result in a large matter mixing: this is known as the *MSW resonance*.

In the Sun, a further special case applies. ν_e are born in nuclear processes in

the core, where the electron density is extremely high. From Equation 1.34, this means that $\sin^2 2\theta_M \rightarrow 0$ and so the mass and flavour states are coincident. On their journey to the surface of the Sun, the neutrinos encounter an adiabatically changing density, and so remain in the mass eigenstate in which they are born. Once they have escaped the Sun, these neutrinos remain in the (now vacuum) mass eigenstate, which is no longer pure ν_e , and so the ν_e flux on Earth is lower than would be expected in the absence of neutrino mass.

1.4 The Solar Neutrino Sector

The mechanisms that could drive thermonuclear energy production in stars were well established theoretically by the 1960s, but an unambiguous demonstration that the processes were present in the otherwise inaccessible stellar core would require the observation of the neutrinos produced in the weak interaction processes. As John Bahcall pointed out in 1964, “only neutrinos, with their extremely small interaction cross-sections, can enable us to see into the interior of a star.” [38]. Although initially motivated by considerations of nuclear and astrophysics, the experimental study of these stellar neutrinos opened the door on the field of neutrino mass and flavour change.

The energy generation in the Sun mainly proceeds via processes known as the *pp* chain and the CNO cycle, in which heavier nuclei are built out of hydrogen. A number of stages in this chain produce electron neutrinos with either a continuous spectrum (in the case of three-body decays) or a unique energy (for two-body decays). These neutrinos have mean energy 0.6 MeV [32], up to a maximum of around 10 MeV. Knowledge from nuclear physics and observations of the surface of the Sun allowed theorists such as Bahcall [38] and others to calculate the neutrino fluxes on Earth resulting from some of these nuclear processes.

The first experiment aimed at verifying these predictions was set up at the

Homestake mine in South Dakota by Ray Davis, using the same technique with which he had attempted to detect $\bar{\nu}_e$ from reactors in 1955. He observed the process $\nu_e + {}^{37}\text{Cl} \rightarrow e^- + {}^{37}\text{Ar}$ (with threshold 814 keV) in a tank of 390,000 litres of C_2Cl_4 . The presence of ${}^{37}\text{Ar}$ was inferred by flushing the tank with helium to remove the argon, and waiting to observe the decay of the ${}^{37}\text{Ar}$ to its ground state [39]. The first results set an upper limit on the solar neutrino flux at Homestake around one-seventh of the flux predicted by the models, while later results [40] obtained a measurement around one-third the predicted flux. This was the first hint that either the models were wrong, or that neutrinos were changing flavour between production and their detection on Earth.

For a long time, Davis was the only experimenter studying solar neutrinos, until 1989 when the Kamiokande experiment in Japan published results on the flux of neutrinos from the decay of ${}^8\text{B}$ in the pp chain; the same flux to which Davis's experiment was sensitive [41]. The Kamiokande detector consisted of a large water tank watched by photomultiplier tubes which measured the Čerenkov light emitted by electrons from the process $\nu_e e^- \rightarrow \nu_e e^-$. They measured a value of the ${}^8\text{B}$ neutrino flux around half of that expected from the Standard Solar Model (SSM), and consistent with the result from Davis's chlorine experiment.

The next wave of experiments in the early 1990s used ${}^{71}\text{Ga}$ as the target, as its threshold for the reaction ${}^{71}\text{Ga} + \nu_e \rightarrow {}^{71}\text{Ge} + e^-$ (233 keV) is low enough to allow detection of neutrinos from the initial $p + p \rightarrow {}^2\text{H} + e^+ + \nu_e$ stage of the pp chain, whose flux is much more precisely determined in the SSM than the ${}^8\text{B}$ flux. SAGE in 1991 [42] and GALLEX in 1992 [43] published results using this reaction, both showing fluxes below that predicted by the SSM.

While Kamiokande and the radiochemical experiments continued to take data and were able to improve the constraints on the solar neutrino fluxes, the most convincing evidence for neutrino flavour change came from the Sudbury Neutrino Observatory (SNO) experiment, which was able to measure the neutrino flux for

all active flavours and thereby confirm the solar neutrino flux predictions of the SSM.

The SNO experiment consisted of a tank containing 1,000 tons of heavy water (D_2O) watched by photomultiplier tubes. In addition to the electron scattering channel $\nu_x e^- \rightarrow \nu_x e^-$, (which is primarily sensitive to electron neutrinos) SNO observed the charged current reaction $\nu_e + d \rightarrow p + p + e^-$ and the neutral current reaction $\nu_x + d \rightarrow p + n + \nu_x$. It is this last reaction that gave SNO sensitivity to the total flux of active neutrinos from the Sun, since the NC reaction does not distinguish between flavours. In addition to the deuterium target, SNO also added two tonnes of NaCl to the tank to improve the NC measurement during one running period. This improvement came about for three reasons: firstly the cross section for neutron capture on Cl is large, increasing the fraction of neutrons that are captured in the detector; secondly the energy of the γ rays emitted by Cl after neutron capture is higher than the corresponding γ for capture on the deuteron; and lastly, the isotropy of the multiple γ emission from Cl allows better separation between NC events and electron scattering events. With this setup SNO was able to measure a total neutrino flux in agreement with the SSM, along with the now-expected reduced flux of ν_e [44]. These results are shown in Figure 1.2.

With the results from radiochemical experiments and SNO, the phenomenon of flavour change in solar neutrinos was confirmed, but in interpreting the results in terms of neutrino mass states and mixing angles, there remained some ambiguity. The Kamioka Liquid Antineutrino Detector (KamLAND) aimed to resolve this, using a liquid scintillator detector to observe the antineutrinos from 55 reactors in Japan. The reactor neutrinos had an energy spectrum peaking at around 3 MeV and extending up to 8 MeV. The KamLAND detector's location made it sensitive to the oscillation parameters which were favoured at the time, and the lack of matter effects made the KamLAND results complementary to the solar experiments [45].




Figure 1.2: Results from SNO on the solar flux of ν_e (abscissa) and $\nu_\mu + \nu_\tau$ (ordinate). The red band shows the constraint from CC events, which determine the ν_e flux only; blue is from NC events which constrain the total flux of all three flavours, and green shows the constraint from electron scattering events. The best-fit point is also shown, falling between the dashed lines, where the fluxes sum to the SSM values. Figure taken from [44].

Figure 1.3 shows the energy-dependent spectral distortion characteristic of neutrino oscillations seen by KamLAND, a feature not directly seen in the solar experiments. Taking the KamLAND results together with the solar experiments gives a best-fit mass splitting and mixing for the solar neutrino sector of $\Delta m^2 = 7.59^{+0.21}_{-0.21} \times 10^{-5} \text{ eV}^2/\text{c}^4$ and $\tan^2 \theta = 0.47^{+0.06}_{-0.05}$. The constraints from these experiments on the mixing parameters are shown in Figure 1.4.

1.5 The Atmospheric Neutrino Sector

A second thread of evidence for neutrino flavour change developed alongside the results from solar neutrino experiments. In the 1980s, a large experimental effort was underway to search for proton decay (a possible signal of physics beyond the Standard Model) in large underground detectors. An important background for these studies was the interaction of neutrinos produced in the atmosphere by the




Figure 1.3: The survival probability of $\bar{\nu}_e$ at a solar baseline as a function of $\langle L \rangle / E$, as measured by KamLAND. $\langle L \rangle$ is the flux-weighted average baseline. The data are shown in black and the best-fit oscillation hypothesis is shown in blue. The sinusoidal $\langle L \rangle / E$ dependence of the survival probability can be clearly seen. Figure from [45].

decay products of cosmic ray interactions, with typical energies around 1 GeV and higher. Kamiokande, in 1988 [46], and IMB in 1992 [47], both measured the fluxes of ν_μ and ν_e , finding a significant deficit in the ν_μ flux but no corresponding deficit in the ν_e flux, a phenomenon dubbed the *atmospheric neutrino anomaly*.

The picture was complicated slightly by a pair of iron tracking detectors, Fréjus and the bawdily-named NUSEX, which observed no deficit in the muon neutrino flux [48, 49].

Firm evidence that the atmospheric neutrino anomaly was due to neutrino oscillations came from the SuperKamiokande (“SuperK”) experiment, a next-generation version of Kamiokande containing 50 kton of water with 11,000 PMTs used to observe the Čerenkov light. In 1998 SuperK published results showing an asymmetry in the upward and downward fluxes of muon neutrinos [50]. Since upward-going atmospheric neutrinos originate at the far side of the Earth, they have a much longer path length ($\sim 12,000$ km) to the detector than downward-




Figure 1.4: Constraints on the solar mixing parameters Δm_{\odot}^2 and $\tan^2 \theta_{\odot}$ from KamLAND and solar neutrino experiments. The KamLAND contours, shown as the coloured regions, constrain the mass splitting very precisely, while the solar experiments' contours, shown as lines, better constrain the mixing angle. From [45].




Figure 1.5: Ratio of the atmospheric neutrino flux to that expected in the absence of oscillations as a function of the distance travelled divided by the incident neutrino energy, as measured by SuperKamiokande. The black points are the data and the red line is the best oscillation fit. Figure adapted from [51].

going neutrinos (~ 10 km), which are produced in the atmosphere above the detector. Thus, the angle of the neutrino (as estimated from the angle of the lepton it produces in a CC interaction) gives information about the distance the neutrino has travelled, L . Two-flavour neutrino oscillations as described by Equation 1.17 therefore predict a deficit of up-going ν_μ but no deficit for down-going ν_μ (for certain values of the mass splitting Δm^2). With their first results, SuperK was able to constrain the mass splitting and mixing angle. It also favoured the oscillation $\nu_\mu \rightarrow \nu_\tau$ over $\nu_\mu \rightarrow \nu_e$ by observing no asymmetry in the ν_e flux.

Later results from SuperK [51] tested the oscillation hypothesis more directly by measuring a value of L/E for each neutrino. Since this value occurs directly in the oscillation probability, the shape of the observed energy spectrum can be compared to the shape predicted by oscillations. The result is shown in Figure 1.5, in which the data agree well with the prediction from oscillations. These results also allowed SuperK to disfavour other explanations for neutrino flavour change such

as neutrino decay (by 3.4σ) and neutrino decoherence (by 3.8σ). The constraints on the oscillation parameters found were $1.9 \times 10^{-3} < \Delta m^2 < 3.0 \times 10^{-3} \text{ eV}^2$ and $\sin^2 2\theta > 0.90$ at 90% C.L.

At around the same time, the Soudan 2 collaboration published results [52] giving evidence from an iron detector for atmospheric neutrino oscillations, which had not been found in NUSEX or Fréjus. The observed deficit of ν_μ -like events, when interpreted as neutrino oscillations, was consistent with the mixing parameters allowed by SuperK.

Confirmation that oscillations at the atmospheric L/E were not due to $\nu_\mu \rightarrow \nu_e$ came from the CHOOZ experiment in France. A liquid scintillator detector located around 1 km from a nuclear reactor, CHOOZ aimed to observe disappearance of the reactor $\bar{\nu}_e$. The baseline was chosen to give sensitivity to the atmospheric oscillation parameters (in contrast to KamLAND, whose baseline gave sensitivity to the solar oscillation parameters). Final results [53] found no evidence of a deficit, favouring $\nu_\mu \rightarrow \nu_\tau$ as the oscillation mode in atmospheric neutrino oscillations and limiting the mixing angle θ_{13} (see Section 1.6) to $\sin^2 2\theta_{13} < 0.19$.

While SuperK continues to take atmospheric neutrino data, a more recent focus of studies in the atmospheric neutrino sector has been on manmade beams of ν_μ produced in accelerators. The first such experiment was KEK to Kamioka (K2K), which produced a beam of ν_μ with 12 GeV protons incident on an aluminium target. The beam was directed towards the SuperK detector 250 km away, with a suite of smaller detectors located along the beamline around 300 m from the target. These detectors, which included a 1kton water Čerenkov detector (essentially a smaller version of SuperK), served to measure the properties of the neutrino beam at its production. K2K observed 112 beam neutrino events in the SuperK detector, against a no-oscillation expectation of 158 events, and used a smaller sample of muon-like events to provide evidence of spectral distortion. With these data, they were able to make some constraint on the oscillation parameters, consistent with

the SuperK L/E result [54].

The OPERA experiment also uses a manmade neutrino beam originating at CERN, with the aim of directly observing the ν_τ produced from $\nu_\mu \rightarrow \nu_\tau$ oscillations at the atmospheric baseline. Like the DONUT experiment, OPERA uses layers of nuclear emulsion interspersed with lead plates to provide the high spatial resolution necessary to identify τ leptons. In 2010, OPERA published evidence for a CC ν_τ interaction candidate [55] and the experiment continues to take data.

The MINOS experiment which is the topic of this thesis is also a long-baseline accelerator neutrino experiment. It differs from K2K chiefly in the energy of the incident protons (120 GeV instead of 12 GeV) and the corresponding baseline (735 km), as well as in the detector technology. MINOS uses magnetized steel-scintillator tracking calorimeters optimized for the identification and measurement of muon tracks. With this setup, the MINOS experiment has provided the world-best measurement of the atmospheric mass splitting along with a bound on the mixing angle [56] as shown in Figure 1.6.

1.6 Three-Flavour Oscillations

All of the results discussed so far, when taken individually, can be interpreted in terms of oscillations between two neutrino flavour eigenstates coupled to two mass eigenstates. In this section a description of the general oscillation probability (Equation 1.14) when specialized to three flavours is given and the conditions, realized in Nature, which make the two-flavour approximation appropriate for both solar and atmospheric neutrinos, are discussed.

A common parametrization of the mixing matrix for three flavours separates

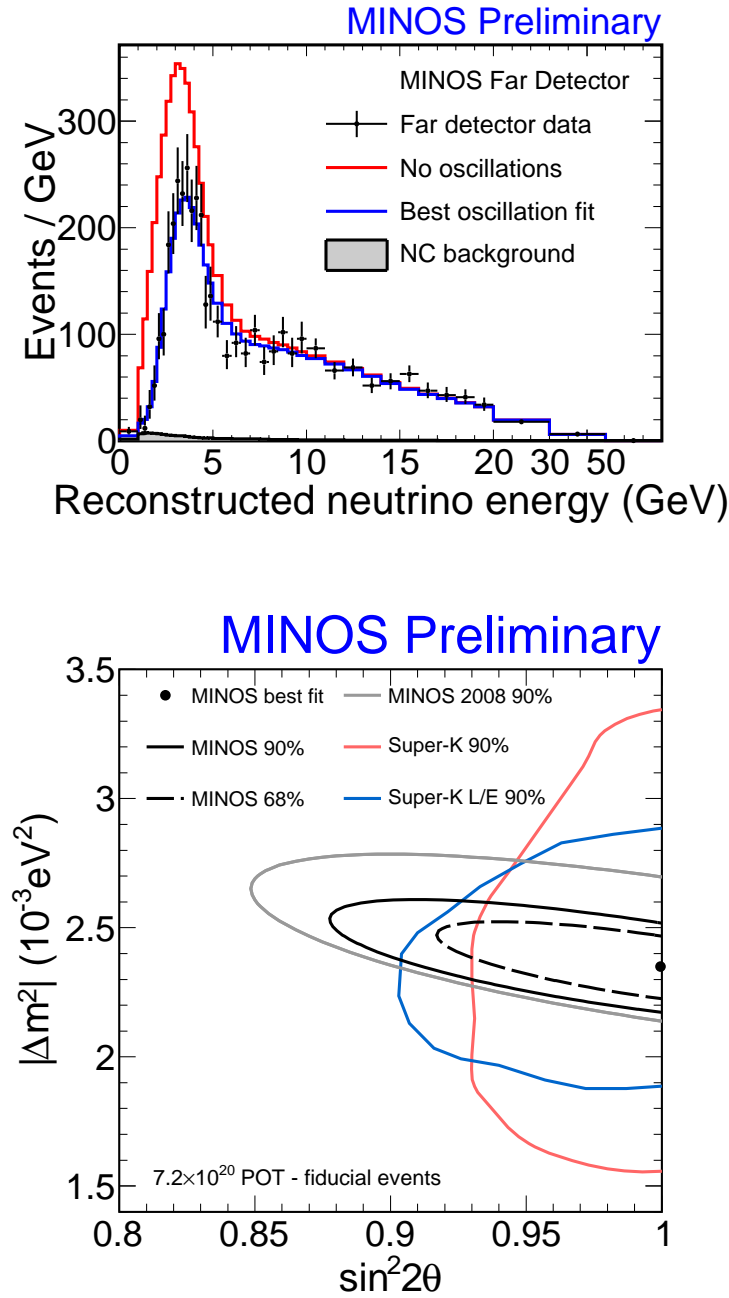


Figure 1.6: MINOS charged current energy spectrum (top) and confidence limits in the $(\Delta m^2_{\text{atm}}, \sin^2 2\theta)$ plane. In the top plot the black points are the data and the solid red line is the prediction in the absence of oscillations, while the solid black line shows the best fit to oscillations. The lower plot contains confidence limits from MINOS and other experiments, showing the world-best constraint on Δm^2_{atm} . From [57].

out the rotations in three planes:

$$U = \begin{pmatrix} 1 & 0 & 0 \\ 0 & c_{23} & s_{23} \\ 0 & -s_{23} & c_{23} \end{pmatrix} \begin{pmatrix} c_{13} & 0 & s_{13}e^{-i\delta} \\ 0 & 1 & 0 \\ -s_{13}e^{i\delta} & 0 & c_{13} \end{pmatrix} \begin{pmatrix} c_{12} & s_{12} & 0 \\ -s_{12} & c_{12} & 0 \\ 0 & 0 & 1 \end{pmatrix} \quad (1.35)$$

$$= \begin{pmatrix} c_{12}c_{13} & s_{12}c_{13} & s_{13}e^{-i\delta} \\ -s_{12}c_{23} - c_{12}s_{23}s_{13}e^{i\delta} & c_{12}c_{23} - s_{12}s_{23}s_{13}e^{i\delta} & s_{23}c_{13} \\ s_{12}s_{23} - c_{12}c_{23}s_{13}e^{i\delta} & -c_{12}s_{23} - s_{12}c_{23}s_{13}e^{i\delta} & c_{23}c_{13} \end{pmatrix} \quad (1.36)$$

where $c_{ij} = \cos \theta_{ij}$ and $s_{ij} = \sin \theta_{ij}$. The parameter δ is the only complex phase in the matrix, and is known as the \mathcal{CP} violating phase. It should be noted that if neutrinos are Majorana particles (that is, if $\nu = \bar{\nu}$), then additional complex phases enter, but these have no effect on oscillations, so they are ignored here.

The parametrization of U in Equation 1.35 has the desirable feature of separating out the parameters relevant to the two sectors discussed so far. The leftmost matrix, depending on θ_{23} , drives the atmospheric oscillation. From SuperK and MINOS, it is known that this mixing is close to maximal [56] (i.e., $\sin^2 2\theta_{23}$ is close to 1). The rightmost matrix contains the parameter relevant to solar neutrino oscillations, θ_{12} , whose value is large but not maximal [45]. The central matrix contains the parameters relevant to $\nu_\mu \rightarrow \nu_e$ oscillations at the atmospheric baseline, but the CHOOZ result shows that this angle must be small; specifically $\sin^2 \theta_{13} < 0.10$ [53].

With these mixing angles and the mass splittings for the solar and atmospheric regimes we can represent the flavour composition of the mass eigenstates graphically, as in Figure 1.7. The sign of the atmospheric mass splitting has not been determined, and both possibilities are shown in the figure. The scheme with the closely-spaced doublet at lower mass is called *normal hierarchy*, for its similarity to the quark mass spectrum, while the other scheme is known as *inverted hierarchy*.

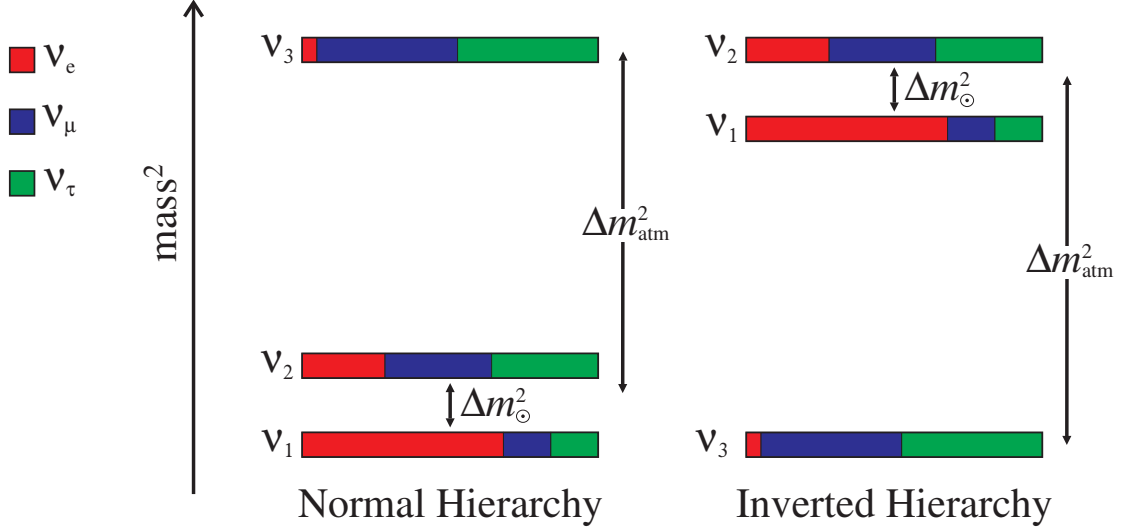


Figure 1.7: The flavour content of the neutrino mass states, approximately as found from current data, with U_{e3} set to $0.05e^{i\pi/2}$. Normal and inverted hierarchies are shown. Figure taken from [58]

From this cartoon we can see the decoupling of the solar and atmospheric regimes: neutrinos produced as ν_e (solar ν , reactor ν) are a superposition of the mass eigenstates ν_1 and ν_2 , with no contribution from ν_3 , and so the two-flavour approximation applies. It is exact when $\theta_{13} = 0$. In the atmospheric sector, all three mass eigenstates participate, but as one mass splitting is much larger than the other, the “one mass scale dominance” scheme of Equations 1.18 and 1.19 is regained. In this case, the effective mass splitting $\Delta m_{\text{atm}}^2 \approx \Delta m_{13}^2 \approx \Delta m_{23}^2$, and the effective mixing angle can be found from Equation 1.19 to be

$$\sin^2 2\theta_{\text{atm}} = |U_{\mu 3}|^2(1 - |U_{\mu 3}|^2). \quad (1.37)$$

1.6.1 θ_{13} , \mathcal{CP} violation and the hierarchy

While the two-flavour approximation has so far been successful in describing the results of experiments (except for LSND), the features unique to the full three-flavour treatment are still of great interest. These are the size of the mixing angle θ_{13} , the possible presence of \mathcal{CP} violation in the neutrino sector (controlled by

the parameter δ), and whether the atmospheric hierarchy is normal or inverted. Each of these may be investigated using *subdominant* oscillations between ν_μ and ν_e at the atmospheric baseline. Two oscillation probabilities are experimentally relevant: the appearance probability $\bar{\nu}_\mu \rightarrow \bar{\nu}_e$, which can be investigated in long-baseline (LBL) accelerator experiments; and the disappearance probability $\bar{\nu}_e \rightarrow \bar{\nu}_e$ which is investigated in shorter baseline reactor neutrino experiments. The observation of any ν_e appearance in LBL experiments, or any $\bar{\nu}_e$ disappearance in reactor experiments, if occurring at the atmospheric L/E , would signal a nonzero value of θ_{13} . As well as being of interest as the only unmeasured angle in the PMNS matrix, θ_{13} has further-reaching consequences. Since the \mathcal{CP} -violating phase δ only appears in the PMNS matrix multiplied by $\sin \theta_{13}$, a nonzero value of the mixing angle opens up the possibility of \mathcal{CP} violation in the neutrino sector.

The last piece of the three-flavour picture is the hierarchy. Determining this parameter requires taking advantage of matter effects which differ for ν_e and $\bar{\nu}_e$, causing the oscillation probabilities $\nu_\mu \rightarrow \nu_e$ and $\bar{\nu}_\mu \rightarrow \bar{\nu}_e$ to differ measurably.

A thriving experimental programme is underway to investigate this sector. The T2K experiment [59], which started running in 2009, uses a beam of ν_μ produced at JPARC and directed towards the SuperK detector 295 km away, looking for ν_e appearance. The beam is not in fact aimed directly towards SuperK, but at a small angle away from it. Because of the secondary pion decay kinematics, this results in a more sharply-peaked energy spectrum at SuperK and therefore reduces the background from high-energy NC interactions which produce $\pi^0 \rightarrow \gamma\gamma$, faking ν_e . The NO ν A experiment [60] at Fermilab will use a similar *off-axis* principle with a longer baseline of 810 km, switching between a beam of mainly ν_μ and a beam of mainly $\bar{\nu}_\mu$, with the aim of observing \mathcal{CP} -violating effects.

Reactor neutrino experiments aiming to measure θ_{13} are also under construction worldwide, from Double Chooz in France [61] to Daya Bay in China [62] and RENO in Korea [63]. These experiments will improve on the previous θ_{13} limits by

the use of near detectors which will measure the flux of $\bar{\nu}_e$ close to the production point, thereby driving down systematic errors.

1.7 Sterile Neutrinos

While the LEP result provides a strong constraint on the number of neutrinos directly coupling to the Z , the possibility of other types of neutrinos exists. Such particles would be without any Standard Model couplings at all and are thus known as *sterile* neutrinos. A sterile neutrino arises naturally when the Standard Model (SM) is extended to account for neutrino mass, as required by the phenomenon of neutrino flavour change: the SM contains only a left-handed neutrino field, but a mass term requires a right-handed field to couple this to. This right-handed field would have no SM interactions, and so would fit the definition of a sterile neutrino².

The implications of sterile neutrinos have been widely investigated in astrophysical contexts, where they could make up part of the ‘dark matter’ believed to account for around 25% of the matter density of the Universe [64]. The ultra-high density environment present in supernova explosions and neutron stars can be opaque to active neutrinos, while sterile neutrinos are able to pass unhindered. For this reason, the effects of sterile neutrinos on these environments is a topic of active research [65, 66]. In the context of high-energy physics, however, the strongest hint of the existence of sterile neutrinos came from the LSND experiment.

²A popular but slightly more complicated scheme for including neutrino masses in the SM is the so-called *see-saw* model, which also contains a right-handed neutrino field.

1.7.1 The LSND Anomaly

The three flavour picture presented so far accounts for the results of all recent experiments apart from one. In 1996 the Liquid Scintillator Neutrino Detector (LSND) at Los Alamos observed a signal which did not fit into the three flavour scheme, and which has led to a number of further experiments aiming to confirm or refute it.

The LSND experiment [67] consisted of a tank of 167 tons of mineral oil doped with scintillator, which produced both scintillation light and Čerenkov light, detected by an array of 1220 PMTs. This detector was exposed to a beam of $\bar{\nu}_\mu$ produced in proton collisions on a water target. The 800 MeV protons produced π^+ which came to rest inside the target, decaying to produce $\mu^+\nu_\mu$ and in turn $e^+\nu_e\bar{\nu}_\mu$. The energies of the $\bar{\nu}_\mu$ ranged up to 52.8 MeV and the detector was located about 30 m from the neutrino source. The 1996 LSND results [68] found a significant excess of $\bar{\nu}_e$ events which they interpreted as a signal of $\bar{\nu}_\mu \rightarrow \bar{\nu}_e$ oscillations. Given the energy and baseline of the LSND experiment, such oscillations would require a new mass splitting of the order of $1 \text{ eV}^2/c^4$. Such a mass splitting would imply the presence of a sterile neutrino in order to be compatible with the LEP Z width result. LSND published further results in 1998 on $\nu_\mu \rightarrow \nu_e$ oscillations, finding an excess consistent with the previous results [69].

Other experiments with sensitivity to $\bar{\nu}_\mu \rightarrow \bar{\nu}_e$ oscillations observed no signal, excluding a large portion of the LSND allowed parameter space, as shown in Figure 1.8. Nonetheless, the LSND result was not entirely ruled out, prompting the construction of the MiniBooNE experiment at Fermilab, aimed at making a definitive statement about the LSND anomaly. MiniBooNE consists of an 800 ton tank of mineral oil watched by PMTs. The detector is 541 m from the source of ν_μ , whose energy peaks at 700 MeV, giving a similar L/E to LSND. The first MiniBooNE results, published in 2008, found no excess of ν_e in their signal region,




Figure 1.8: Regions of parameter space allowed by LSND and excluded by MiniBooNE. The MiniBooNE 90% C.L. is shown as a solid black line, ruling out the LSND 90% allowed region (dark blue) and most of the 99% region (light blue).

ruling out the LSND 90% confidence region at 90% C.L., under the assumption of two-neutrino oscillations [70], as shown in Figure 1.8. More complex schemes in which there are 2 new sterile states in addition to the three active neutrinos allow compatibility between LSND and MiniBooNE via \mathcal{CP} violation, since the principal LSND results observed $\bar{\nu}_\mu \rightarrow \bar{\nu}_e$, whereas MiniBooNE’s results pertain to $\nu_\mu \rightarrow \nu_e$ [71]. However, even in this scheme there is significant tension between the results from appearance experiments (MiniBooNE and LSND) and disappearance experiments (Bugey, Chooz and Palo Verde). To investigate this sector further, MiniBooNE has run in antineutrino mode, publishing results that are so far inconclusive in ruling out oscillations at the LSND mass scale [72].

1.7.2 Sterile Neutrinos in the Atmospheric Sector

Oscillations of atmospheric neutrinos demonstrate the disappearance of ν_μ , and the results of CHOOZ [53] and SuperK [50] have shown that the dominant channel

cannot be $\nu_\mu \rightarrow \nu_e$. The most likely candidate is therefore $\nu_\mu \rightarrow \nu_\tau$, but the threshold for τ production in $\nu_\tau n \rightarrow \tau^- p$ is around 3 GeV, making ν_τ effectively invisible at the energies of most of the atmospheric flux. For this reason, no direct observation of ν_τ resulting from oscillations has been made, leaving open the possibility of oscillation to another invisible channel, namely a sterile neutrino $\nu_\mu \rightarrow \nu_s$.

Although directly detecting the oscillation products is difficult for ν_τ and impossible for ν_s , several indirect approaches have been taken. SuperK have compared the predictions for pure $\nu_\mu \rightarrow \nu_\tau$ and pure $\nu_\mu \rightarrow \nu_s$ with their data [73]. Two features allow this discrimination: firstly, oscillations to ν_τ leave the rate of neutral current (NC) interactions unchanged, because of lepton universality, whereas oscillations to ν_s reduce the rate of NC events. Secondly, active neutrinos passing through the Earth experience matter effects from the NC interaction which ν_s do not. In an entirely analogous way to the difference in CC interactions between ν_e and other flavours, this alters the effective oscillation parameters as in Equation 1.34, but with the number density of electrons replaced by the number density of neutrons (and a numerical factor). The result is a suppression of the oscillation probability for neutrinos of energy above ~ 15 GeV. By combining an NC-enriched data set with a data set sensitive to the high energy tail of the atmospheric flux, SuperK was able to disfavour pure $\nu_\mu \rightarrow \nu_s$ against $\nu_\mu \rightarrow \nu_\tau$ at the 99% confidence level.

While this analysis excluded $\nu_\mu \rightarrow \nu_s$ as the only oscillation channel, the possibility of $\nu_\mu \rightarrow \nu_s$ as a subdominant channel remains open. SuperK has also investigated this case, considering oscillations of $\nu_\mu \rightarrow (\cos \xi \nu_\tau + \sin \xi \nu_s)$. With this parametrization, they found $\sin^2 \xi < 25\%$ at the 90% C.L. [74].

1.7.3 Parametrizing Four-Flavour Oscillations

In trying to incorporate the LSND anomaly into the picture of neutrino oscillations, a fourth mass eigenstate is needed to provide the necessary large Δm^2 . There are two categories of schemes which provide this. In the first category, the pair m_1 and m_2 provide the solar mass splitting, while the gap between m_3 and m_4 provides the atmospheric mass splitting. The LSND mass splitting is provided by the gap between these two (relatively) closely-spaced pairs, leading to the designation “2 + 2” for such models. Alternatively, the LSND mass splitting may be far above or below the three other states, a scheme known as “3 + 1”. Ignoring the LSND result, one may also choose to place a fourth mass eigenstate degenerate with either the m_1, m_2 doublet or the m_3 state.

In each of these cases, the full oscillation probabilities at the atmospheric L/E have a rather complicated dependence on the parameters of the 4×4 mixing matrix, which now contains five mixing angles, three mass-squared differences and three \mathcal{CP} -violating phases. The MINOS NC analysis is not sensitive to this number of parameters, but under the assumption that only one mass-squared difference has a non-negligible value of $\Delta m^2 L/E$, the probabilities take on the simple forms given in Equations 1.18 and 1.19. A more empirical way to write the oscillation probabilities in this case is [75]

$$P(\nu_\mu \rightarrow \nu_\mu) = 1 - f_\mu \sin^2 \Delta \quad (1.38)$$

$$P(\nu_\mu \rightarrow \nu_e) = f_e f_\mu \sin^2 \Delta \quad (1.39)$$

$$P(\nu_\mu \rightarrow \nu_\tau) = (1 - f_s - f_e) f_\mu \sin^2 \Delta \quad (1.40)$$

$$P(\nu_\mu \rightarrow \nu_s) = f_s f_\mu \sin^2 \Delta. \quad (1.41)$$

where $\Delta = 1.27 \Delta m^2 L/E$. The parameter f_s may be interpreted as the fraction of the disappearing ν_μ (as measured in the MINOS CC analysis) that are sterile,

and is limited to the range $0 \leq f_s \leq 1 - f_e$.

For a given model with one mass scale dominance, the empirical parameters may be related to the mixing matrix elements by comparing Equations 1.38–1.41 to Equations 1.18 and 1.19.

1.8 Non-oscillation Neutrino Physics

While studies of neutrino oscillations have provided many insights into the nature of the neutrino, there are important questions that cannot be resolved by oscillation experiments. The mass-squared splittings of the neutrinos are now well-known, but they give no information on the *absolute* mass scale, and the question of whether neutrinos are Majorana or Dirac particles is unmeasurable in oscillation experiments.

1.8.1 Absolute Neutrino Mass

The current most stringent limits on the absolute neutrino mass come from cosmology, where nonzero neutrino masses have effects on cosmic microwave background anisotropies and the formation of cosmological structures [76]. Although these limits can be very tight, they are also strongly dependent on the other astrophysical assumptions that go into the calculation. To reflect this uncertainty, the PDG [32] gives

$$\sum_i m_i < (0.17 - 2.0) \text{ eV}. \quad (1.42)$$

Studying the spectrum of the β electron in β -decay allows a more direct neutrino mass measurement. The maximum possible energy of the electron is the Q^2 value of the decay minus the neutrino mass, and so by examining the endpoint of the electron energy spectrum, the neutrino mass may be constrained. There is a further complication in that the neutrino produced in β -decay is in a flavour

eigenstate, and therefore a superposition of mass eigenstates, and so the quantity actually constrained in these measurements is [77]

$$m_\beta^2 = \sum_i |U_{ei}|^2 m_{\nu_i}^2. \quad (1.43)$$

where U is the PMNS matrix. This method was in fact pointed out by Hans Bethe and Rudolf Peierls in 1934 [4], but no direct observation of nonzero neutrino mass has yet been made in this way. Nonetheless, stringent limits have been set from the β -decay of ^3H , with a world-best value of $m_\beta^2 < 2\text{eV}^2/\text{c}^4$ at 95% C.L. [32].

1.8.2 Majorana Neutrinos

In the Standard Model, neutrinos are Dirac particles, meaning that ν and $\bar{\nu}$ are distinct. However, it is possible that neutrinos are Majorana fermions, with particle and antiparticle equivalent. In this case, what we commonly call $\bar{\nu}$ is just a Majorana neutrino with right-handed chirality, while ν is the left-handed version. The difference between Majorana and Dirac neutrinos is not experimentally observable in most cases, in particular in neutrino oscillations, but a Majorana mass term in the Lagrangian would allow lepton-number violating processes such as neutrinoless double β -decay, in which, for example, a nucleus decreases its mass number by two with the emission of two electrons but no neutrinos. The experimental signature of such a decay is a coincidence of electrons with total energy equal to the Q^2 value of the nuclear transition.

One experiment has produced a claim of observation of neutrinoless double β -decay [78], although it has not been independently confirmed. Investigations are continuing with experiments such as SNO+ [79], EXO [80], SuperNEMO, COBRA [81] and CUORE [82] aiming to provide improved sensitivity.

Chapter 2

The MINOS Experiment

This chapter gives an overview of the MINOS experiment. The NuMI beamline which produces the mainly ν_μ beam studied in MINOS is described in Section 2.2, and the two detectors located in this beam are described in Section 2.3. A brief description of the software used to simulate and analyze the data from the detectors is given in Sections 2.4 and 2.5, and finally an overview of the physics results from the MINOS experiment is given in Section 2.6. Comprehensive descriptions of the NuMI beamline can be found in [83, 84], while the MINOS detectors are described fully in [85].

2.1 Overview

MINOS is a long-baseline neutrino oscillation experiment studying ν_μ oscillations at the atmospheric Δm^2 . It consists of two detectors located in the NuMI beamline, the *near detector* at Fermilab, around 1 km from the neutrino production point, and the *far detector* located 735 km away in the Soudan mine in northern Minnesota, at a depth of 705 m (2070 metres water equivalent). The primary aim of the experiment is to make precision measurements of the atmospheric neutrino mixing parameters Δm_{atm}^2 and $\sin^2 2\theta_{23}$. In order to achieve this goal, the mainly

ν_μ neutrino beam is measured at the near detector, where no oscillations are assumed to have occurred. This measurement is used to predict the far detector spectrum, where an energy-dependent deficit of ν_μ is the signal for oscillations. This deficit and its spectral shape provide the information necessary to constrain Δm_{atm}^2 and $\sin^2 2\theta_{23}$.

The use of two detectors is a fundamental feature of the experiment, aimed at reducing systematic uncertainties. The prediction of the far detector spectrum from a combination of Monte Carlo and near detector data allows for the cancellation of uncertainties such as the neutrino interaction cross section and beam flux. To this end, the detectors are constructed to be as close to functionally identical as possible: both use steel planes as target material, attached to planes of scintillator strips for particle detection.

2.2 The NuMI beamline

MINOS uses a flux of neutrinos produced in the Neutrinos at the Main Injector (NuMI) beamline at Fermilab. In the NuMI beamline, 120 GeV protons are extracted from the Main Injector and directed onto a 940 mm-long graphite target. Each extraction of the beam, or *spill*, lasts 10 μs and typically contains $\sim 10^{13}$ protons. Beam spills occur every 2.2 s.

The primary and secondary interactions of protons in the target produce mainly pions and kaons. The π^+ and K^+ are focused by a pair of magnetic horns separated by 10 m before decaying into neutrinos inside a 675 m-long decay pipe. For the initial data-taking the decay pipe was evacuated, but the risk of implosion due to radiation damage and corrosion of an inspection window led to the pipe being filled with helium in the most recent data-taking period. Any remaining hadrons and muons which have not decayed are absorbed by an aluminium and steel beam dump followed by a further 300 m of rock, leaving a flux entirely of




Figure 2.1: The components of the NuMI beamline.

neutrinos. A schematic of the NuMI beamline is shown in Figure 2.1.

The energy spectrum of this neutrino flux can be modified by varying the position of the target relative to the first horn and by altering the current in the horns. The energy spectra for three different configurations of the beamline are shown in Figure 2.2. The majority of the data has been taken in the low-energy configuration with the target fully inserted into the first horn and the horn current at 185 kA, as this maximizes the neutrino interaction rate at the oscillation minimum at the far detector baseline. All data in this thesis uses this beam configuration. Other beam configurations have been used to provide constraints on the beam flux as described in Section 2.5.1 and to improve discrimination between oscillations and other hypotheses for neutrino disappearance [56].

In the low-energy configuration, simulations indicate that the beam consists of 92.9% ν_μ , 5.8% $\bar{\nu}_\mu$ and 1.3% $\nu_e + \bar{\nu}_e$. The NuMI beam has also been run with the horn current reversed, so that π^- and K^- are focused, resulting in a beam consisting of mostly $\bar{\nu}_\mu$. Only ‘forward’ horn current running (i.e., focussing positively-charged particles) is used in this thesis.

2.3 The MINOS Detectors

The MINOS detectors are nearly functionally identical steel-scintillator tracking calorimeters optimized for the detection and measurement of neutrino in-

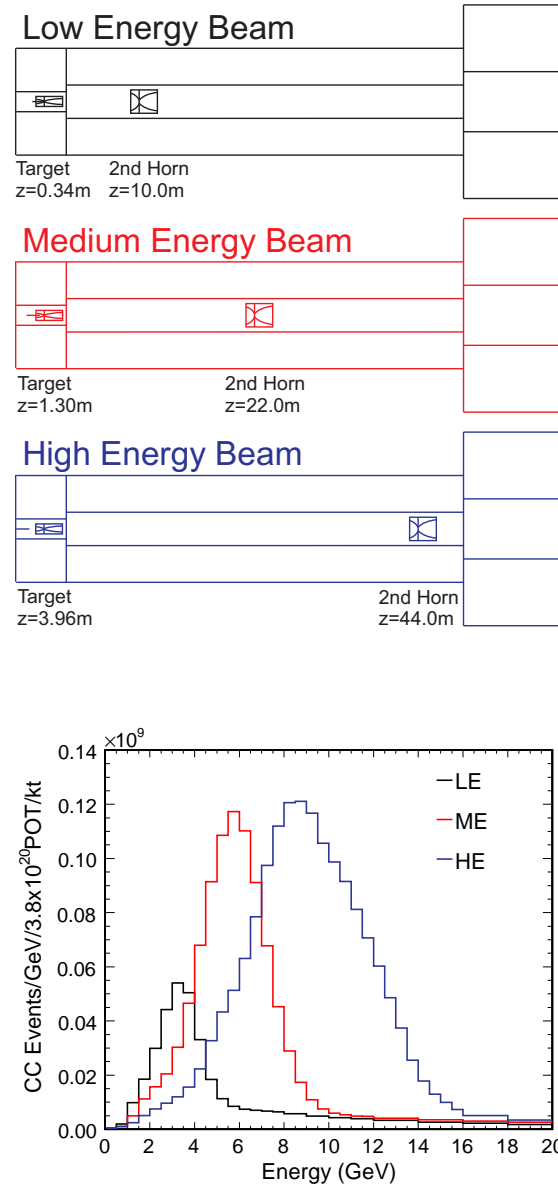


Figure 2.2: Three configurations of the NuMI beamline, and the resulting spectrum of charged current ν_μ interactions at the MINOS near detector. The long tail in the low energy beam arises from hadrons which pass down the beam axis and are therefore not affected by the focussing. Data taken in a configuration similar to the “Low Energy Beam” is used in this thesis.




Figure 2.3: Schematic of the MINOS detector arrangement, showing the scintillator (white) attached to the steel planes (grey). Not to scale.

teractions in the few GeV range. Each detector consists of 2.54 cm-thick steel planes. Mounted on these steel planes are planes of plastic scintillator, divided into $4.1 \text{ cm} \times 1 \text{ cm}$ strips. As shown schematically in Figure 2.3, the strips in alternating planes are oriented at 90° to one another, to allow for three-dimensional event reconstruction. Each steel plane is separated from the next by an air gap to give a total plane separation of 5.94 cm.

The near detector, with a total mass of 980 ton, consists of 282 steel planes, each having the shape of a “squashed octagon”. The front 120 planes form the *calorimeter* section, intended for measuring the energy of hadronic showers. In this section, each plane is instrumented with scintillator, maximizing the sampling of hadronic showers and therefore the shower energy resolution. Of these planes, four out of five are *partially instrumented*, having 64 strips covering a 6.0 m^2 area, while the remaining one in five planes are *fully instrumented*, having 96 strips covering the full 13.2 m^2 of the steel plane. The downstream *spectrometer* section is intended for measuring the momentum of muons, and so is fully instrumented on every fifth plane, with no partially instrumented planes.

The far detector, with a larger total mass of 5.4 kton reflecting the much lower neutrino flux at Soudan, consists of 486 steel planes in the shape of a regular octagon, divided into two almost equal *supermodules*. Every plane except the ends of the supermodules is instrumented with scintillator across the full face of




Figure 2.4: The near detector planes. The diamond in the centre is the coil hole, and the circle offset from it is the beam centre. Partially instrumented planes have scintillator covering the shaded region only.




Figure 2.5: Far detector even-numbered (left) and odd-numbered (right) planes. The width of each plane is 8 m. The scintillator strips are manufactured in modules of six different types labelled A–F, each containing 20 or 28 strips.

the plane. The near and far detector shapes are shown in Figures 2.4 and 2.5.

A current-carrying coil passing through the centre of each plane of each detector produces toroidal magnetic field with mean strength 1.4 T. The curvature of muons in the field provides discrimination between μ^+ and μ^- , and allows a determination of the momentum of muons which do not range out in the detector. For muons which do stop in the detector, a more precise estimate of the momentum is given by the range of the muon, which is not as sensitive to multiple scattering as the momentum determined from curvature. Therefore the momentum estimate from

range is used for all stopping muons, with the curvature only used when the muon escapes.

For most of the data-taking period, the detectors were run in forward-field mode, in which the B -field focuses μ^- , keeping them within the detector, and defocuses μ^+ . Some data has been taken with the detector magnetic field reversed (including all of the reverse horn current running), but is not used in this thesis.

2.3.1 Scintillators and Light Detection

The scintillator strips are made of extruded polystyrene doped 1% by weight with the organic scintillators PPO and POPOP. Each strip has a reflective coating of titanium-dioxide, which retains the scintillation light within the strip until it is captured by the wavelength-shifting (WLS) fibre that runs the length of the strip. The WLS fibre transports the light to the ends of the strip, where it is transferred to the pixel of a multi-anode photomultiplier tube (PMT) by a clear fibre. In the near detector, the fibres are read out at one end only, with the far end mirrored to maximize light output. In the far detector, the attenuation along the longer strip length is larger, and so the fibres are read out at both ends. The structure of the scintillator strips is shown in Figure 2.6.

The PMTs used for light detection differ between the near and far detectors but have some similarities. Both are multianode PMTs made by the Hamamatsu Corporation, and have multiple pixels sharing a common photocathode. The operating voltage is around 800 V, which provides a gain of order 10^6 . In the near detector, 64-anode PMTs are used [86], with each pixel reading out one fibre, while in the far detector, 16-anode PMTs [87] read out eight fibres per pixel. This optical summing scheme reduces the number of PMTs and associated readout electronics required, thereby lowering the cost of the detector. The ambiguity in which strip was hit is resolved at reconstruction time using the fact that the




Figure 2.6: A MINOS scintillator strip, showing scintillator, reflective coating and inlaid WLS fibre.

multiplexing pattern is different at each end of the strip.

With this setup, approximately 4 photoelectrons per minimum-ionizing particle are produced at each strip end in the far detector.

2.3.2 Electronics

The design of the MINOS detector electronics is driven by the rate of neutrino interactions at each detector. In the near detector, there are several events per $10\mu\text{s}$ spill, whereas in the far detector, the neutrino rate is much lower, and the data rate consists mostly of detector noise from the WLS fibres and PMTs.

The near detector front end electronics [88] digitize the input signal continuously at a frequency of 53 MHz. To provide an increased dynamic range, the input current is first split onto eight capacitors in fractions $\frac{1}{2}$, $\frac{1}{4}$, $\frac{1}{8}$ etc. The capacitors integrate the charge over the 18.8 ns “bucket”, with the result that only one of the capacitors has a voltage within the allowed limits of the analogue-to-digital

converter (ADC) residing on the same PCB. A chip known as a charge integration encoder (QIE) selects the appropriate capacitor for the ADC, also outputting a 3-bit number which indicates the range which was used. To allow deadtimeless readout of the electronics, each QIE contains four copies of the current splitter, which integrate the charge in adjacent clock cycles.

Each QIE chip resides on a *MENU* board along with the ADC and a large enough buffer to hold the output from a whole spill. The other electronics are then grouped hierarchically: 16 *MENU*s live on a *MINDER*, and up to eight *minders* live on a *MASTER*. The *MASTER* linearizes the readouts from the electronics using a lookup table, sending the results to a VME computer known as a *ROP*, whence they are sent to the data acquisition (DAQ) system.

In the far detector [89], the lower rate of interactions means that a larger fraction of deadtime can be tolerated. The arrangement of the electronics is again hierarchical: each M16 PMT is read out by a *VA* chip, three of which are mounted on a *VFB*. Twelve *VFB*s (36 *VA* chips) are read out by a *VARC*. The *VFB* houses the *ASDLite* chip, which reads out the dynodes of the PMTs. Each PMT has a single dynode which is common to all of the channels. The *ASDLite* produces a trigger signal if the dynode signal from a PMT exceeds a threshold set to about $\frac{1}{3}$ of the signal expected from a single photoelectron (pe) at the photocathode. The digitization is controlled by the *VARC*, which receives the dynode threshold signals from the *ASDLite* chips. On receiving such a trigger, the *VARC* signals the appropriate *VA* chip to hold its charge, which is digitized by an ADC on the *VARC*. The dynode trigger is also used by the *VARC* to produce a timestamp for the digits. Due to the high rate of noise from the WLS fibres, digitization is only initiated when two dynode triggers on a single *VARC* are received within 400 ns of one another. This so-called “2/36” trigger reduces the noise rate without affecting the recording of physics events.

2.3.3 Triggering

The DAQ system has the task of triggering the readout of the frontend electronics and transferring the desired data to a farm of PCs which send the data for archival at Fermilab. Alongside this, the DAQ provides control of calibration runs and online monitoring of the data.

In reading out the electronics, the DAQ applies a number of triggers to select data for further processing. Spill triggers in both detectors signal readout of the whole detector for a period of time around the spill. In the near detector, the spill trigger is produced by the Fermilab accelerator network, and lasts up to $18.8\mu\text{s}$, while in the far detector, the timestamps of spills are received via the internet: data is buffered for long enough to allow time for the spill trigger to arrive. Timestamping is provided by GPS receivers at each detector, which have a worst-case resolution of 200 ns. There is an additional 64 ns uncertainty in calculating the time offset between detectors because of uncertainty in hardware time delays. In the far detector, there are also *fake* spill triggers, which occur when no spill is present, but signal the same minimum bias readout of the detector. Information from these fake spill triggers is used to constrain the rate of backgrounds such as cosmic ray muons and detector noise in true spills.

There are also activity-based triggers which allow for the recording of events when no spill trigger is present. These capture cosmic ray muons, atmospheric neutrinos, and (in the far detector) spill events that occur when the remote spill trigger is not available (such as during an internet outage between Fermilab and Soudan). The “4/5” trigger requires four planes in a contiguous set of five to contain at least one hit, while an activity trigger requires activity in any 20 planes of the detector. Additionally, the far detector is triggered if four contiguous planes of the detector have a summed pulse height greater than 1500 ADC counts and a total of at least six hits.

2.3.4 Calibrating the Detectors

MINOS analyses rely on accurate determination of the energy in hadronic and electromagnetic showers, chiefly for the determination of Δm_{atm}^2 . There are two separate concerns: firstly that the calibrated response per unit of energy deposited in the scintillator is constant with time, position in the detector, and size of signal, and is the same for both detectors. Secondly, for showers, the calibrated response must be correctly matched to the energy in the shower. The first of these, the relative energy scale, is known to 2.1%, based on detector calibrations with cosmic ray muons and the detector light injection (LI) system, while the second, the absolute energy scale, is energy dependent and is calculated from the single-particle responses from the MINOS calibration detector along with shower modelling uncertainties from the NEUGEN neutrino event generator, as described in more detail in Section 5.1.2.

The calibrated detector response Q_{corr} in detector d at time t on position x in strip s is calculated from the raw response Q_{raw} as

$$Q_{\text{corr}} = Q_{\text{raw}} \times D(d, t) \times L(d, s, Q_{\text{raw}}) \times S(d, s, t) \times A(d, s, x) \times M(d). \quad (2.1)$$

D is the drift correction which corrects for changes in the overall detector response, L corrects for the nonlinearity in PMT and electronics response, S is a strip-to-strip calibration that removes differences between strips in the detector, and A corrects for light attenuation between production point and the strip end where it is detected. Additionally, M converts the near and far pulse heights into the same units.

The drift, strip-to-strip and attenuation calibrations use cosmic ray muons as a “standard candle” to provide a known input to which variations in the detector response can be compared. The rate of cosmic muons is approximately 10 Hz at the near detector and 0.5 Hz at the far detector, giving a high-statistics sample.




Figure 2.7: Change in detector response with time, relative to a reference date, for near detector (top) and far detector (bottom). The general downward trend is due to scintillator aging, and the discontinuities are due to retuning of the high voltage to the PMTs.

Also, the energy deposition of muons is well understood.

Correcting the overall detector response with time is done by the drift correction, which uses throughgoing cosmic muons. The spectrum of energy deposition from these muons does not change with time, and so any change in the median pulse height recorded from these throughgoing muons is attributed to detector effects due to scintillator aging, temperature changes and overall gain variations in the PMTs (such as when the voltage is changed over the whole detector). The median detector response with time, used in the drift calibration, is shown in Figure 2.7.

The strip-to-strip calibration takes the response from each strip and applies

a correction factor to correct the response of the strip to the mean response of the detector. In each strip, the response is measured at several positions. These measurements are fit to a double-exponential, and the values of the fit function are used to provide the correction at any position on the strip.

To correct for nonlinearity in the PMTs and electronics, a light injection system is used, which produces light pulses of various size that are fed to the PMTs via the WLS fibres. In the near detector, the light produced by the LI system is measured by PIN diodes, which have a linearity in response of around 1%, to enable comparison with the readout from the PMT. In the far detector, the LI signal is read out by the PMT at the far end of the detector, where the attenuated signal is within the range where nonlinearity has a smaller effect. There is still some nonlinearity on the end that is read out, so the process is iterated with the injected and read out ends reversed.

With these calibrations applied, the response in each detector is constant in time, space and signal size. Nonetheless, the response may still differ between detectors. To account for this, the inter-detector calibration $M(d)$ is used. This uses a sample of cosmic muons that stop in the detector to allow identification of hits with known energy deposition [90]. From the Bethe-Bloch formula, the energy deposition as a function of distance from the end of a muon track can be calculated. Using these hits therefore provides a sample with known energy deposition. Furthermore, the energy deposition of muons near the end of a muon track varies sharply with energy, but away from the end of the track it is approximately constant, and so hits near the track end are excluded from the calculation, meaning that a moderate uncertainty on the true track end position results in only a small uncertainty on the energy deposition.

The absolute energy scale was found with the calibration detector, a small-scale version of the MINOS detectors, using the same steel, scintillator and readout electronics [91]. In one of the configurations used, one end of each strip was read

out by the near detector electronics, the other end by the far detector electronics. The calibration detector was exposed to test beams at CERN containing hadrons, electrons and muons at energies up to 10 GeV. This allowed the detector response to individual leptons and hadrons to be determined. With this data, the energy resolution for hadron showers was determined to be $56\%/\sqrt{E/\text{GeV}} \oplus 2\%$, and the difference in near and far detector electronics was shown not to induce a bias on the shower energy scale [92].

2.4 MINOS Event Simulation

The Monte Carlo event simulation in MINOS divides into two sections: firstly, the NuMI beamline is simulated using the FLUGG package, which combines the FLUKA [93] code for particle interactions and propagation with the GEANT4 geometrical description package [94]. This section results in ‘flux files’ containing details of hadrons and muons that decay in the beamline producing neutrinos. In the second part, these neutrino fluxes are taken as input for the detector Monte Carlo, which simulates neutrino interactions in the detectors and the rock surrounding them, and simulates the transport of the resulting particles.

The use of the FLUGG package for the beam simulation allows for a detailed description of the beamline, including the 1cm target shift for part of the data-taking period, and the inclusion of the effects of helium in the decay pipe for running after November 2007. To take account of these effects accurately, and permit careful comparisons between data and Monte Carlo, three sets of neutrino fluxes are generated; one for each running period.

To simulate events in the detectors, the NEUGEN3 package [95] is used to simulate neutrino interactions based on the neutrino flux, and the resulting particles are propagated through the detectors using the GEANT3 [96] package. The effects of multiple neutrino interactions per spill and the addition of muons from the rock

surrounding the detector are taken account of at this stage. Hadronic showers are modelled using the GCALOR code [97], chosen for its good agreement with data from the calibration detector. The GEANT3 simulation outputs the energy depositions in the detector, which are converted to light which is propagated to the PMTs by the MINOS package PhotonTransport. A further MINOS package, DetSim, simulates the PMTs and readout electronics.

The conversion from energy to light multiplies by the reciprocals of the MINOS calibration constants so that when the calibration procedure is carried out on the simulated detector readout (by multiplying by the calibration constants themselves), the Monte Carlo is correctly calibrated. This procedure is necessary to take account of effects such as readout thresholds and the Poisson statistics of photoelectrons at the PMT face. Additionally, the calibration constants for each Monte Carlo file are taken from a randomly-chosen date in the data-taking period. In this way, the distribution of calibration constants in the data and the Monte Carlo are the same.

2.5 MINOS Event Reconstruction

The primary aims of the MINOS event reconstruction are to allow discrimination between charged and neutral current events and to accurately reconstruct the energy of hadronic showers and muons. The code can also reconstruct atmospheric neutrinos and cosmic ray muons, and provides information useful for the identification of electronic showers (used in the ν_e appearance analysis).

In the near detector where there are typically several neutrino interactions per spill, the reconstruction software must separate the hits due to different interactions. This is first done with a series of timing and topology cuts which group hits that are close in time and space into *slices*, with the intention that the hits in one slice are those due to an individual neutrino interaction. The hits in these slices

are passed to the later stages of the reconstruction.

Tracks in the detector are found and fitted in separate stages. Clusters of hits with correlated positions are first identified, and these hits are given to a fitting routine which uses a Kalman Filter to estimate the actual path of the track, taking into account the curvature of the particle in the detector magnetic field. The output from the fitting stage is used to identify any other hits that lie along the track and should be included. The fitter outputs an estimate of the track momentum obtained from curvature as well as an estimate based on the range of the track in the detector material. At the analysis stage, the range estimator is used for tracks which stop in the detector, while the curvature is used for exiting tracks.

The shower reconstruction forms clusters of hits and groups together clusters considered to originate from the same shower. Only hits with a pulse height above $2p_e$ are considered for inclusion, as the simulation of low pulse-height hits (due mainly to detector effects such as fibre noise and crosstalk between PMT pixels) shows some differences from the data that are not well understood. The energy of a shower is estimated from the total calibrated pulse height of the hits it includes.

A final event building stage groups together tracks and showers considered to come from the same neutrino interaction, also producing an estimate of the interaction vertex position.

2.5.1 Flux tuning

The energy and angular distribution of hadrons produced from the NuMI target is not extremely well known, and so near detector data are used to constrain the hadron production spectrum, expressed in terms of the longitudinal and transverse momenta of the hadrons, p_z and p_T . The method takes advantage of the multiple possible configurations of the NuMI beamline, which sample different regions of




Figure 2.8: Effect of flux tuning procedure on the Monte Carlo in low-energy beam configuration. The blue histogram is the MC before the beam fit, and red is after. The black points are the near detector data. The flux tuning gives a much improved agreement between data and Monte Carlo.

the (p_z, p_T) space. The method is described in detail in [98], and proceeds as follows: the (p_z, p_T) distributions for hadrons giving rise to neutrino interactions in the near detector are parametrized analytically with a number of free parameters allowing the distribution to be warped, and a simultaneous fit to near detector charged current selected spectra from several beam configurations is performed. A number of cross section- and detector-related parameters (such as energy scales) are included in the fit to ensure that mismodellings in these quantities do not affect the fit values for the flux parameters.

The result of the flux tuning procedure is a weight for each Monte Carlo event based on its parent hadron and the parent's p_z and p_T . In all spectra and fits in this thesis, these weights from the flux tuning procedure are applied. Figure 2.8 shows the effect of this tuning on the near detector charged current spectrum in the low-energy beam configuration which is used in this thesis. The peak at 3–4 GeV is reduced slightly, and the high energy tail in the 6–10 GeV region is increased relative to the nominal Monte Carlo, providing a much improved agreement with the data.




Figure 2.9: The ν_μ survival probability at the MINOS far detector, as a function of true neutrino energy. The position of the first oscillation minimum is determined by the parameter Δm_{atm}^2 and its depth by $\sin^2 2\theta_{23}$. The rapid oscillations at low energies are observed averaged over energy. The resulting energy spectrum is shown in Figure 1.6.

2.6 MINOS Physics

A wide range of physics has been undertaken with the MINOS detectors. The main result is the measurement of the atmospheric mixing parameters to high precision using the spectrum of charged current ν_μ events. At the far detector, there is an energy-dependent deficit of ν_μ , according to the survival probability as shown in Figure 2.9. A fit of the far detector charged current ν_μ spectrum to this shape constrains the oscillation parameters to $\Delta m_{\text{atm}}^2 = 2.35^{+0.11}_{-0.08} \times 10^{-3} \text{ eV}^2/c^4$ at 68% C.L. and $\sin^2 2\theta > 0.91$ at 90% C.L [57].

MINOS has also constrained the probability for $\nu_\mu \rightarrow \nu_e$ oscillations, thereby placing limits on the value of the last unmeasured mixing angle θ_{13} . As the MINOS detector is not optimized for the identification of electronic showers, the analysis is complex, requiring detailed particle identification (PID) methods to distinguish electronic from hadronic showers. The application of these PID methods to data and Monte Carlo in the near detector, where only backgrounds to the ν_e appearance search are present, results in differences of up to 20% in the region of interest, which are dealt with by two methods of decomposing the spectrum into its neu-




Figure 2.10: 90% C.L. contours in $2 \sin^2 \theta_{23} \sin^2 2\theta_{13}$ as a function of the \mathcal{CP} -violating phase $\delta_{\mathcal{CP}}$, from the MINOS ν_e appearance analysis. The contour under the assumption of normal mass hierarchy is in blue, while the red contour is for inverted hierarchy. The limit from the CHOOZ experiment is shown in green.

tral current, ν_μ charged current and beam ν_e components. These new background estimates are propagated to the far detector, where an excess in the data relative to the Monte Carlo would be an indication of ν_e appearance and therefore nonzero θ_{13} [99].

In an exposure of 7×10^{20} protons-on-target, 54 candidate ν_e events were observed, with a background of 49.1 ± 7.0 (stat.) ± 2.7 (syst.), corresponding to a 0.7σ excess. Therefore no strong evidence for ν_e appearance was found, and a constraint at 90% C.L. on the combination of mixing angles was given as:

$$2 \sin^2 \theta_{23} \sin^2 2\theta_{13} < 0.12 \text{ (0.20)} \quad (2.2)$$

for the normal (inverted) hierarchy, at $\delta_{\mathcal{CP}} = 0$. The limit is shown as a function of $\delta_{\mathcal{CP}}$ in Figure 2.10.

The NuMI beam has also been run with reversed horn current, providing a beam of mostly $\bar{\nu}_\mu$. With the charge separation capability of the MINOS detectors, this allows a direct measurement to be made of the parameters governing

antineutrino oscillations, and thus provides a test of \mathcal{CPT} conservation [57].

In the far detector, atmospheric neutrino interactions have been studied, providing a bound on oscillation parameters [100], and a number of results have been produced using cosmic ray muons. The μ^+/μ^- ratio measured constrains the production processes in the upper atmosphere [101], while correlations between stratospheric temperature and the underground muon rate have also been observed [102].

As well as neutrino oscillation physics, the MINOS near detector has been used for measurements of neutrino and antineutrino cross sections down to energies < 10 GeV [103].

Chapter 3

Identifying NC and CC Events in the MINOS Detectors

This chapter details the methods used to obtain the sample of neutrino-induced, well-reconstructed events to be used in the analysis, and to identify these events as due to either charged or neutral current interactions.

3.1 Neutrino event types in MINOS

3.1.1 ν_μ CC events in the MINOS detectors

The MINOS detectors are optimized for the identification and measurement of ν_μ charged current interactions. In these events, the incoming neutrino exchanges a W boson with a nucleus (usually iron) in the detector, producing a muon and accompanying hadronic shower. An example of such an event from the MINOS Monte Carlo is shown in Figure 3.1. The muon produces a long track which curves in the detector magnetic field and is easily identified. The energy of the incoming neutrino may be reconstructed by summing the calorimetric energy of the shower and the energy of the muon track, with the latter determined by range for tracks which stop in the detector, and by curvature for tracks which exit. The momentum

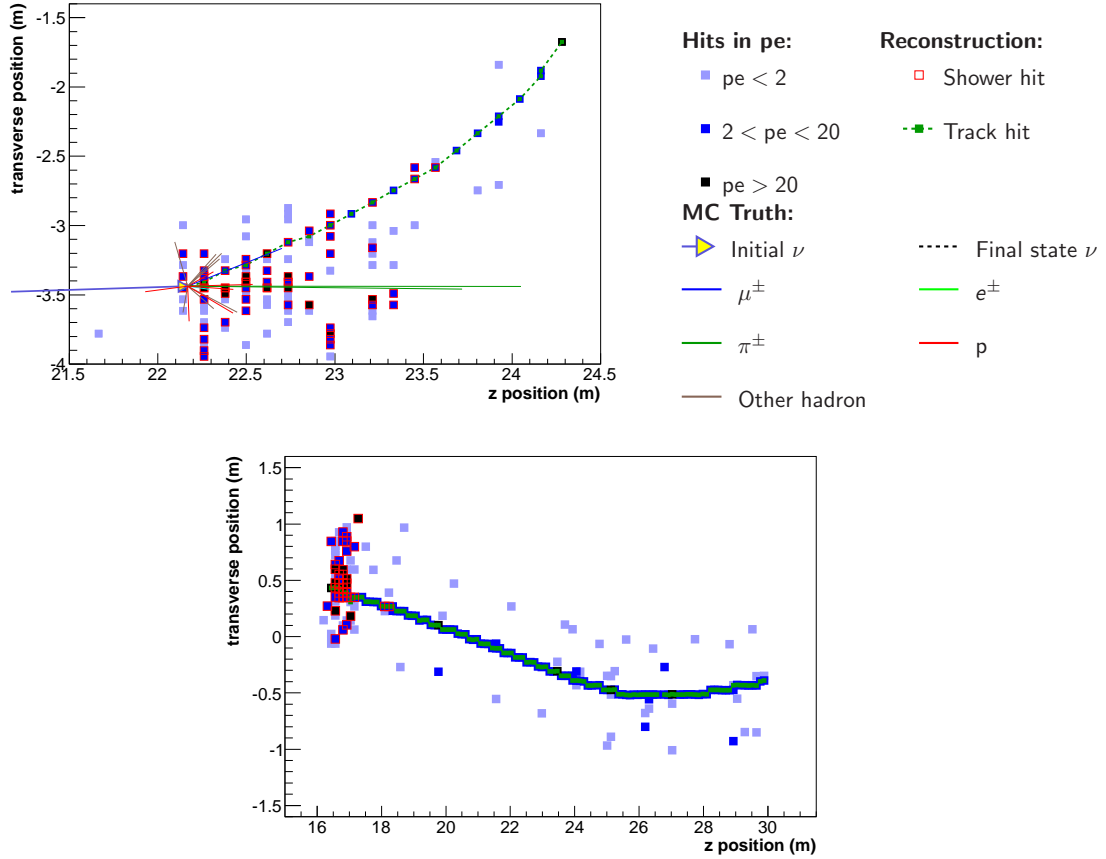


Figure 3.1: Top: A 16 GeV ν_μ CC event in the MINOS far detector Monte Carlo. The muon track is accompanied by hadronic activity around the interaction vertex. The legend is shown on the right. Bottom: A ν_μ CC candidate from the far detector data with reconstructed track plus shower energy of 21 GeV.

from range is used if possible as it has better resolution because it is not sensitive to multiple scattering of the muon.

3.1.2 NC events in the MINOS detectors

An example of a neutral current event is shown in Figure 3.2. In these events, a Z boson is exchanged with the nucleus, producing a hadronic shower. The initial neutrino remains in the final state, carrying away some fraction of the energy, with the result that the energy of the incoming neutrino cannot be reconstructed. A neutral current event depositing 1 GeV of energy in the detector (the signal region for this analysis) typically produces around 10 hits.

A major background to the identification of NC events is CC events with a high value of the inelasticity y , defined as

$$y = \frac{E_\nu - E_\mu}{E_\nu} \quad (3.1)$$

where E_ν is the energy of the incoming neutrino and E_μ is the energy of the final-state muon. Neglecting energy lost in intranuclear rescattering, this is approximately

$$y \approx E_{\text{shw}}/E_\nu \quad (3.2)$$

where E_{shw} is the energy of the hadronic shower. These high- y events have a short muon track which is entirely within the hadronic shower, and cannot be identified by the reconstruction.

3.1.3 ν_e CC events in the MINOS detectors

A second background for neutral current events (and a signal for the θ_{13} measurement) is ν_e charged current events. These have a dense electromagnetic shower at the vertex, which is typically surrounded by the sparser hadronic shower. Separating ν_e CC events from neutral current events in the MINOS detectors requires sophisticated techniques [99], and is not attempted in this thesis.

ν_e CC events in the far detector can arise from both the intrinsic ν_e contamination of the beam and from $\nu_\mu \rightarrow \nu_e$ oscillations driven by a nonzero value of θ_{13} . (Oscillations driven by the solar mixing angle θ_{12} have an effect smaller by a factor of order $\alpha = \Delta m_{\text{atm}}^2/\Delta m_{\odot}^2$: the Monte Carlo indicates approximately 0.4 events from this source for the exposure used in this thesis.) The beam ν_e s are mostly from K and μ decays and have higher energies than ν_μ in the beam. An example of such an event is given in Figure 3.3. ν_e from oscillations occur mostly at the oscillation peak energy of $E \approx 1.4 \text{ GeV}$, and so have much lower average energy




Figure 3.2: Top: A typical NC event in the MINOS far detector Monte Carlo consisting of a sparse hadronic shower. The neutrino of energy 8 GeV has deposited 4 GeV of energy in the detector. Bottom: An NC candidate from the far detector data with reconstructed shower energy of 4 GeV. For legend, see Figure 3.1.




Figure 3.3: A 29 GeV ν_e CC event in the MINOS far detector Monte Carlo arising from the intrinsic ν_e content of the beam. The shower profile is dense, particularly at the core. The reconstruction has “drawn on” a track over the shower. For legend, see Figure 3.1.

than the beam ν_e component. Because of the coarse granularity of the MINOS detectors, these lower energy ν_e events deposit only a small number of hits in the detector and are very similar to low-energy neutral current events. Figure 3.4 shows one such event.

3.2 Event preselection

In both detectors, events must have their reconstructed vertex within the interior of the detector—the *fiducial volume*—in order to be used in the analysis. This requirement helps to reject non-neutrino backgrounds such as cosmic ray muons and ensures that only events with accurately reconstructed energy are included. Events near the edges of the detector will tend to deposit a larger fraction of their energy outside the detector, and are also in the region of the detector which is least well calibrated. A fiducial cut removes these events.




Figure 3.4: A 3 GeV ν_e CC event in the MINOS far detector Monte Carlo arising from $\nu_\mu \rightarrow \nu_e$ oscillations. A small EM shower is present, but difficult to distinguish from the hadronic shower surrounding it. For legend, see Figure 3.1.

3.2.1 Obtaining a clean sample in the near detector

The high rate of neutrino interactions in the near detector drives the choice of preselection cuts. As the total number of events is very high, a tight fiducial volume cut may be made without adversely affecting the statistics used. There are also a larger number of reconstruction failures in which the energy of a shower is poorly reconstructed. The near detector event preselection aims to reduce the number of these poorly reconstructed events significantly.

Near detector fiducial volume

The near detector fiducial cut defines a region based on the distance from the nearest edge of a partially instrumented plane transversely, and on the distance from the ends of the calorimeter section longitudinally. Events must have their vertex more than 0.5 m from the outline defined by the scintillator in the partially instrumented plane, and have a vertex position in z (the longitudinal direction) of $1.7 \text{ m} < z < 4.737 \text{ m}$. The lower limit of this cut vetoes muons from neutrino interactions in the rock in front of the detector, while the upper limit of this




Figure 3.5: Average ratio of reconstructed shower energy to true shower energy as a function of position in the near detector. On the left is the ratio as a function of distance to the closest partial plane edge, and on the right as a function of position in z . The fiducial cut positions are shown by the red lines. In each case the cut in the other variable has been applied in making the plot.

cut corresponds to a distance of around 2.4 m from the back of the calorimeter, ensuring good containment of hadronic showers. Figure 3.5 shows the mean of the ratio of reconstructed to true hadronic shower energy as a function of distance to the edge of the plane and as a function of the z position. The shower energy is calibrated so that the $E_{\text{reco}}/E_{\text{true}}$ distribution peaks at 1, but the full distribution is asymmetric, with a larger fraction of events at $E_{\text{reco}}/E_{\text{true}} > 1$ than below 1. This causes the mean of the distribution to be greater than 1 for well reconstructed events, as can be seen in Figure 3.5.

Near detector event preselection

Because of the high rate of neutrino interactions in the near detector, some classes of reconstruction failures occur more often than in the far detector. Such poorly reconstructed events tend to be reconstructed as low energy showers, becoming a background for the true hadronic showers used in a neutral current analysis. Differences between the near and far detectors lead to increased systematic uncertainties in the analysis, and so reducing the level of these poorly reconstructed events improves the reliability of the analysis. Additionally, the reconstruction failures are concentrated in areas with poor agreement between data and Monte




Figure 3.6: A split event in the near detector, removed by the preselection. The hits highlighted in red are split off from the high-energy CC event. For legend, see Figure 3.1. The grey bars show the extent of the scintillator planes.

Carlo, which is improved on removing these events. Further details on the near detector preselection are given in [104].

Examples of poorly reconstructed events are shown in Figures 3.6 and 3.7. Figure 3.6 is an example of the “split event” class from the MINOS Monte Carlo, in which the hits due to a single neutrino interaction are reconstructed as two separate events. In this case, the original event is a high energy charged current interaction (as evidenced by the long muon track). A high energy pion has escaped the main hadronic shower, depositing energy downstream. The hits from this pion are reconstructed as a shower in a separate event, faking a low energy neutral current shower.

Figure 3.7 shows a “leakage” event, in which a neutrino interaction outside the fiducial volume is reconstructed inside it. This class of events accounts for a larger fraction of the events used in the near detector than the far, since the latter has a much smaller fiducial volume. Therefore reducing the number of such events reduces near/far differences in the analysis.

The near detector preselection removes events such as these using cuts on two variables for each event. The first is the fraction of the pulse height in the




Figure 3.7: A “leakage” event in the near detector, in which the event vertex is truly outside the fiducial volume, but is reconstructed inside. The event is removed by the preselection cuts. For legend, see Figure 3.1. The grey bars show the extent of the scintillator planes.

reconstructed slice which is included in the event. The reconstructed slice is defined as a region in time and space of detector activity that is separated in time and space from other regions of detector activity. The slice pulse height fraction variable has a low value for split events such as those in Figure 3.6, in which most of the pulse height in the reconstructed slice is reconstructed into a different event. Figure 3.8 shows the distribution of this variable in data and Monte Carlo, and also for poorly reconstructed events, defined as those which have

$$\frac{E_{\text{reco}}}{E_{\text{true}}} < 0.3 \quad (3.3)$$

where E_{reco} is the reconstructed shower energy and E_{true} is the true shower energy. Events with a slice pulse height fraction of less than 0.5 are rejected by the cut.

The second variable used is the maximum number of contiguous planes in the event. Since real showers tend to deposit energy in all of the scintillator planes they pass through, well-reconstructed events have a higher value of this variable than poorly reconstructed events. The distribution of this variable is shown in Figure 3.9. Events with fewer than 3 contiguous planes are rejected by

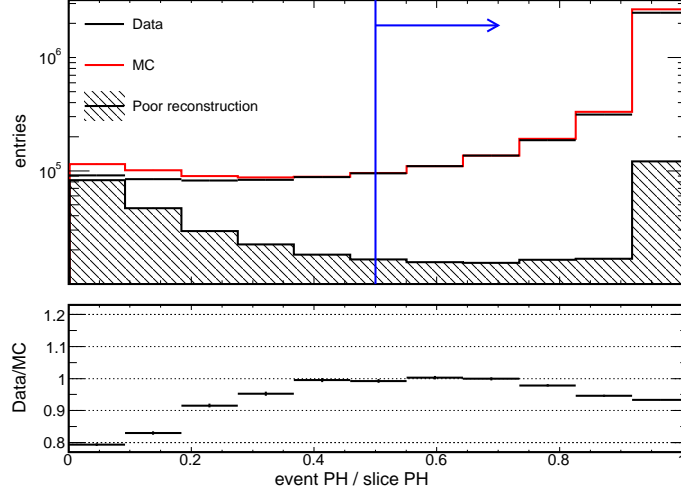


Figure 3.8: Distribution of the slice pulse height fraction in the near detector. The Monte Carlo is in red, with the poorly reconstructed component shown as a hatched area. The black points are the near detector data, and the arrow shows the events retained by the cut. Figure taken from [104].

the preselection.

The results of applying this preselection on the near detector NC-like shower energy spectrum are shown in Figure 3.10. The fraction of poorly reconstructed events is reduced significantly in the region below 1 GeV, from 37.5% to 11.5% [104] at the cost of a factor of 2 in efficiency below 2 GeV in reconstructed shower energy.

3.2.2 Obtaining a clean sample in the far detector

In the far detector, there are a different set of backgrounds to be removed. As the rate of neutrino-induced events is low (around 1 in every 10^5 spills), rejection of background events must be very stringent. As in the near detector, a fiducial volume cut is applied to all events, and a further preselection removes events due to non-neutrino backgrounds and poorly reconstructed neutrino-induced events.

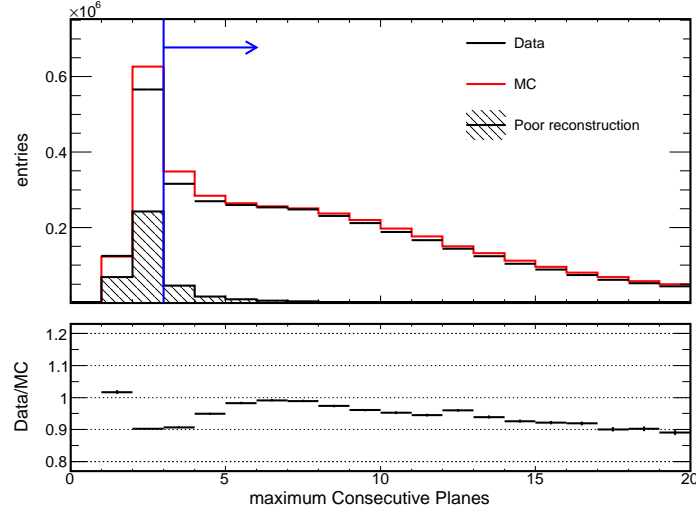


Figure 3.9: Distribution of the maximum consecutive planes variable in the near detector. The Monte Carlo is in red, with the poorly reconstructed component shown as a hatched area. The black points are the near detector data, and the arrow shows the events retained by the cut. Figure taken from [104].

Far detector fiducial volume

As in the near detector, the fiducial volume cut in the far detector ensures good hadronic shower containment by requiring the reconstructed vertex of an event to be in the interior of the detector. Additionally, despite the “2/36” trigger described in Section 2.3.2, detector noise may be read out by the electronics and reconstructed into an NC-like event. The multiplexing of fibres onto PMT pixels creates an ambiguity in the position of such events, which is resolved by the reconstruction. The algorithm used biases these noise events towards the edges of planes, where they are removed by the fiducial cut [58].

The far detector fiducial volume requires events to have their reconstructed vertex more than 0.6m from the coil hole and more than 0.4m from the outer edge of the detector. Longitudinally, the event vertex must be more than 4 planes from the front of supermodule 1, and 18 planes from the back of supermodule 2. In the second supermodule, the front 4 planes are excluded along with the back 18 planes. The cut on the front of the supermodule removes events which

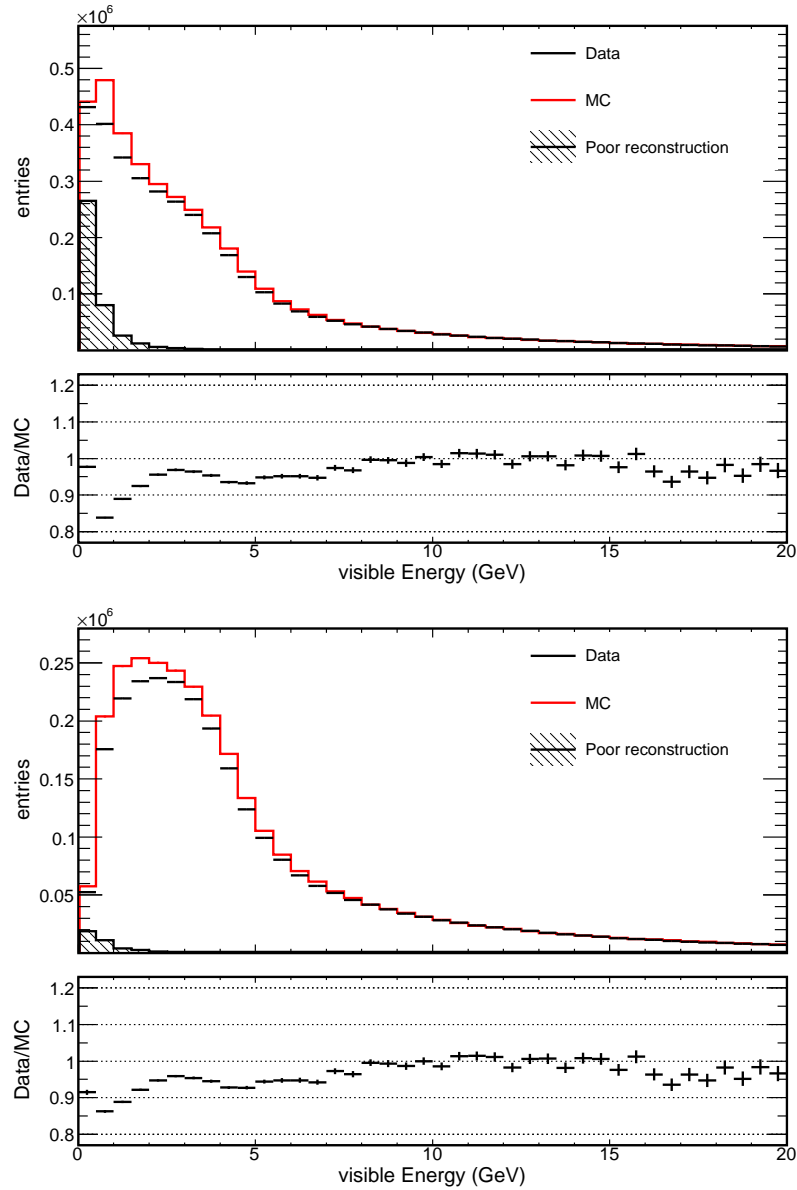


Figure 3.10: Near detector NC-like spectrum before (top) and after (bottom) preselection. The Monte Carlo is in red, with the poorly reconstructed component, which is reduced by the preselection, shown as a hatched area. The black points are the near detector data. Figure taken from [104]. NC-like events are selected using the selection method in [57].

interact in the first supermodule, but have their reconstructed vertex in the second supermodule.

Far detector event preselection

The far detector event preselection removes events due to detector noise, cosmic rays and flashes of the light injection system. The performance of the cuts is quantified using the fake spill triggers, which occur out of time with beam spills, and therefore contain only non-neutrino backgrounds. The far detector preselection is described in detail in [58] and [105].

The main non-neutrino induced background for the neutral current analysis is clusters of fibre noise reconstructed as a small event. These are removed with a cut on the number of strips in the event and its pulse height. All events with calibrated pulse height less than 2000 ADC (25 photoelectrons) are removed, along with events with fewer than 9 strips and calibrated pulse height less than 3750 ADC (47 photoelectrons).

Cosmic ray muons are generally reconstructed as tracks close to vertical (while tracks from beam neutrino events are close to horizontal). The z direction cosine, $|p_z|/E$, of the muon is therefore required to be greater than 0.4. Very steep cosmic ray muons may be reconstructed as showers having a relatively large number of strips in a small number of planes. Cuts on a *transverse RMS* variable and the ratio of strips in the event to number of event planes squared identify and remove these steep cosmics.

To reject spills containing light injection events, the primary method is a dedicated piece of hardware known as a trigger PMT (tPMT). The tPMTs are attached to the light injection system, providing a signal when the LI is flashing, and allowing those spills to be vetoed during later processing. A second method, known as the LISieve, is also employed in the neutral current analysis. This uses some features of events coming from light injection to distinguish them from beam-induced

events: essentially, the LI flashes strips on neighbouring electronics channels, while physics events produce hits that are close in space. The LISieve therefore looks for events in which a large number of strips per plane are hit; which have a high asymmetry in the pulse height read out at the two sides of the detector; and align along boundaries of the LI pulser box system.

Spills failing either the tPMT or the LISieve cuts are rejected. Since the methods are independent and each have an inefficiency of a few $\times 10^{-4}$, the combined cuts will miss only 1 LI spill in 10^7 .

The rate of beam-induced neutrino interactions in the far detector is around 1 in every 10^5 spills. Therefore the probability of obtaining two beam-induced events in one spill is negligible. Two reconstructed events may occur in the same spill due to either the addition of a non-beam background event alongside a beam event, or due to splitting of a beam event by the reconstruction. Both of these cases are dealt with by the “main event” cut, which removes events which contain less than 75% of the pulse height in the spill.

The spill trigger in the far detector causes readout of $100\ \mu\text{s}$ of detector activity. Since the actual beam spill lasts around $10\ \mu\text{s}$, a further cut requires events to be between $-2\ \mu\text{s}$ and $12\ \mu\text{s}$ of the beginning of the beam spill.

The result of the far detector preselection is shown in Figure 3.11. Non-beam induced backgrounds occur uniformly throughout the spill trigger window, while the events passed by the preselection are restricted to the region around the beam spill itself. In the full data sample the number of background events passing the preselection is expected to be approximately 3 [105].

3.3 Separating NC and CC events using an ANN

The analysis presented in this thesis uses the event separation method described in [106] to provide separation between NC and CC events. The method uses a set

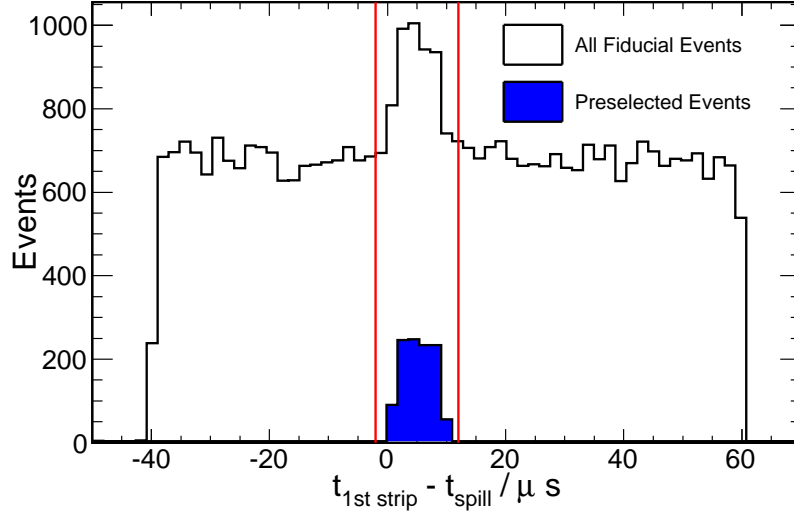


Figure 3.11: Effect of far detector preselection on fiducial events in data. The abscissa is the time relative to the beam spill. The black line shows the times of all fiducial events, while the blue-filled histogram contains the events passing the preselection. The position of the timing cut is shown by the red lines.

of cuts to preclassify a subset of events as NC or CC, and passes the remaining events to an artificial neural network (ANN) trained on MINOS Monte Carlo, which gives an output value which tends towards 1 for events that are CC-like, and towards 0 for events which are NC-like.

3.3.1 Artificial neural networks for event classification

Artificial neural networks are a general method of machine learning in which a real-valued function of a series of input parameters is output based on training data. This section follows the overview given in [107]. In this thesis, as in most high-energy physics contexts, an ANN is used for event classification.

The basic component of an ANN is the *neuron*. A neuron takes a number of inputs, x_i (in this case, variables describing an event), each with an associated weight w_i . The inputs and weights are combined to form the *activation* a , defined




Figure 3.12: Example of a feedforward ANN of the type used for NC/CC separation. The blue points are neurons, and the lines their inputs/outputs. Each input variable is fed into each neuron in the hidden layer, and the output of the single neuron in the output layer is the separation parameter.

as

$$a = \sum_i x_i w_i. \quad (3.4)$$

The neuron has an output y which is a function of the activation, and is commonly chosen to be the sigmoid function

$$y = \frac{1}{1 + e^{-a}}. \quad (3.5)$$

An artificial neural network uses a number of neurons connected together, with the outputs of some neurons being the inputs of others. The weights are then chosen so as to produce the desired function of the inputs, in this case to provide separation between classes of events. In the ANN used here, the neurons are connected as shown in Figure 3.12, with each input fed to the so-called “hidden layer” of neurons, and each neuron in the hidden layer passing its output to a single neuron whose output value is taken as the output of the network.

The weights for the neuron inputs are chosen by training the network on Monte Carlo events. For each input event an error function $G(\mathbf{w})$ can be defined based

on the information content of the network output:

$$G(\mathbf{x}, \mathbf{w}) = -[t \ln y(\mathbf{x}, \mathbf{w}) + (1 - t) \ln(1 - y(\mathbf{x}, \mathbf{w}))] \quad (3.6)$$

where t is the target value for this event: 1 for CC events and 0 for NC events. This function attains its lower bound when the sample is correctly classified, i.e., $y(\mathbf{x}, \mathbf{w}) = t$. Since y is a function of the weights w_i , the derivatives $\partial G / \partial w_i$ can be found:

$$\frac{\partial G}{\partial w_i} = -(t - y)x_i. \quad (3.7)$$

The weights are then adjusted to minimize G , by changing the weights by

$$\Delta w_i = \eta(t - y)x_i \quad (3.8)$$

where η is a constant chosen to determine the step size. The process is repeated for each next event. One loop over all events is known as a training *epoch*. The NC/CC separation ANN used 1000 epochs of training.

In training an ANN, as with other multivariate classification techniques, care must be taken not to *overtrain* the classifier such that it responds to statistical fluctuations in the training sample. To this end, a second independent sample of Monte Carlo is used as a test sample, alongside the training sample. The error summed over all events is calculated for the training and test samples separately. Overtraining will appear as an increase in the test sample error, as the network learns features in the training sample not present in the test sample.

3.3.2 ANN input variables

The NC/CC separation ANN uses 7 input variables. Prior to the ANN training step, events which are extremely likely to be NC or CC are preclassified and not input to the ANN itself. These consist of two classes: events having more than




Figure 3.13: Variables used in the preclassification step which determines whether events are input to the ANN. On the left is the number of planes in the event, and on the right is number of tracks in the event. The red histogram is the total number of events in near detector Monte Carlo, blue is the truly CC component and the black points are near detector data.

40 planes are almost all CC, while events with no track are 70% NC in the near detector. The distributions of these variables in the near detector are shown in Figure 3.13.

The input variables used for the ANN are chosen for their discriminating power between NC and CC events. Key to this is distinguishing muon tracks, the signal of a CC event, and hadronic tracks which are present in NC events. The variables are:

Event pulse height per plane Since muon track planes in general only have one hit with a pulse height equivalent to one minimum-ionizing particle (1 mip), CC events tend to have lower pulse height per plane.

Event pulse height per strip Similarly, hits from hadronic activity are higher pulse height than those from muons, so NC events in general have higher values of pulse height per strip.

Track pulse height fraction This is defined as the ratio $P_t/(P_t + P_a)$, where P_t is the pulse height in strips in the track and P_a is the pulse height in the strips in a window of ± 4 strips each side of the track. As muon tracks tend to be clean and deposit energy in only one strip in each plane, they have a

value of this variable close to 1. Tracks reconstructed in NC events are likely to be buried within a shower and therefore be surrounded by pulse height from the shower.

Number of track strips This is the number of active strips in the detector reconstructed into a track. Since muons are highly penetrating, large values of this variable indicate CC events.

Track vertex z cosine This variable is the cosine of the angle between the track and the detector z axis. The separating power of this variable can be understood in terms of the inelasticity y , which is approximately

$$y = E_{\text{had}}/E_{\nu} \quad (3.9)$$

where E_{had} is the energy of the recoiling hadronic system and E_{ν} is the energy of the incident neutrino. Although the CC y distribution is approximately flat, the reconstruction efficiency is higher for low- y events (with a clear track and little hadronic activity) than for high- y events. The muons from low- y events are very forward-going, since they take the full momentum of the incident neutrino. Therefore in CC events, the track tends to be along the z direction.

Additionally, tracks in NC events may be “drawn-on” rather than truly following particles, and these drawn-on tracks may have large angles to the z direction.

Track extension This variable is defined as the number of planes in the track minus the number of planes in the shower. CC events tend to receive a higher value of this variable.

Track χ^2 per degree of freedom This is a measure of the quality of the fit as

returned by the track fitting algorithm. Hadronic tracks and fake tracks in NC events tend to receive larger values of this variable than clean muon tracks.

The distributions of these variables in near detector data and MC are shown in Figure 3.14. The agreement is in general moderately good, although almost all of the variables show an excess in Monte Carlo relative to data in the region where more NC events lie.

3.3.3 ANN performance

The output of the ANN is shown in Figure 3.15. At high output values, the sample is almost pure CC, with all of the NC events clustered at low values. Nonetheless, the low-output region, considered as a sample of NC events, has much higher CC contamination than the high-output region has NC contamination. The data/MC disagreement in the NC-like region which appeared in the input variables remains, with data lower than MC. Chapter 4 discusses how the near detector data is taken into account in predicting the far detector spectrum in order to reduce the effect of this mismodelling in the Monte Carlo, and Chapter 5 details how the remaining uncertainty is incorporated into the limits quoted on the oscillation parameters.

By placing a cut on the ANN output value, an NC-enriched sample can be obtained. Two competing measures of quality of selection are efficiency, ε , and purity, p . The efficiency for the selection is defined as

$$\varepsilon = N_{\text{sel}}/N_{\text{tot}} \quad (3.10)$$

where N_{sel} is the number of truly NC events selected, and N_{tot} is the total number of truly NC events available to the selection (i.e., after preselection cuts have been

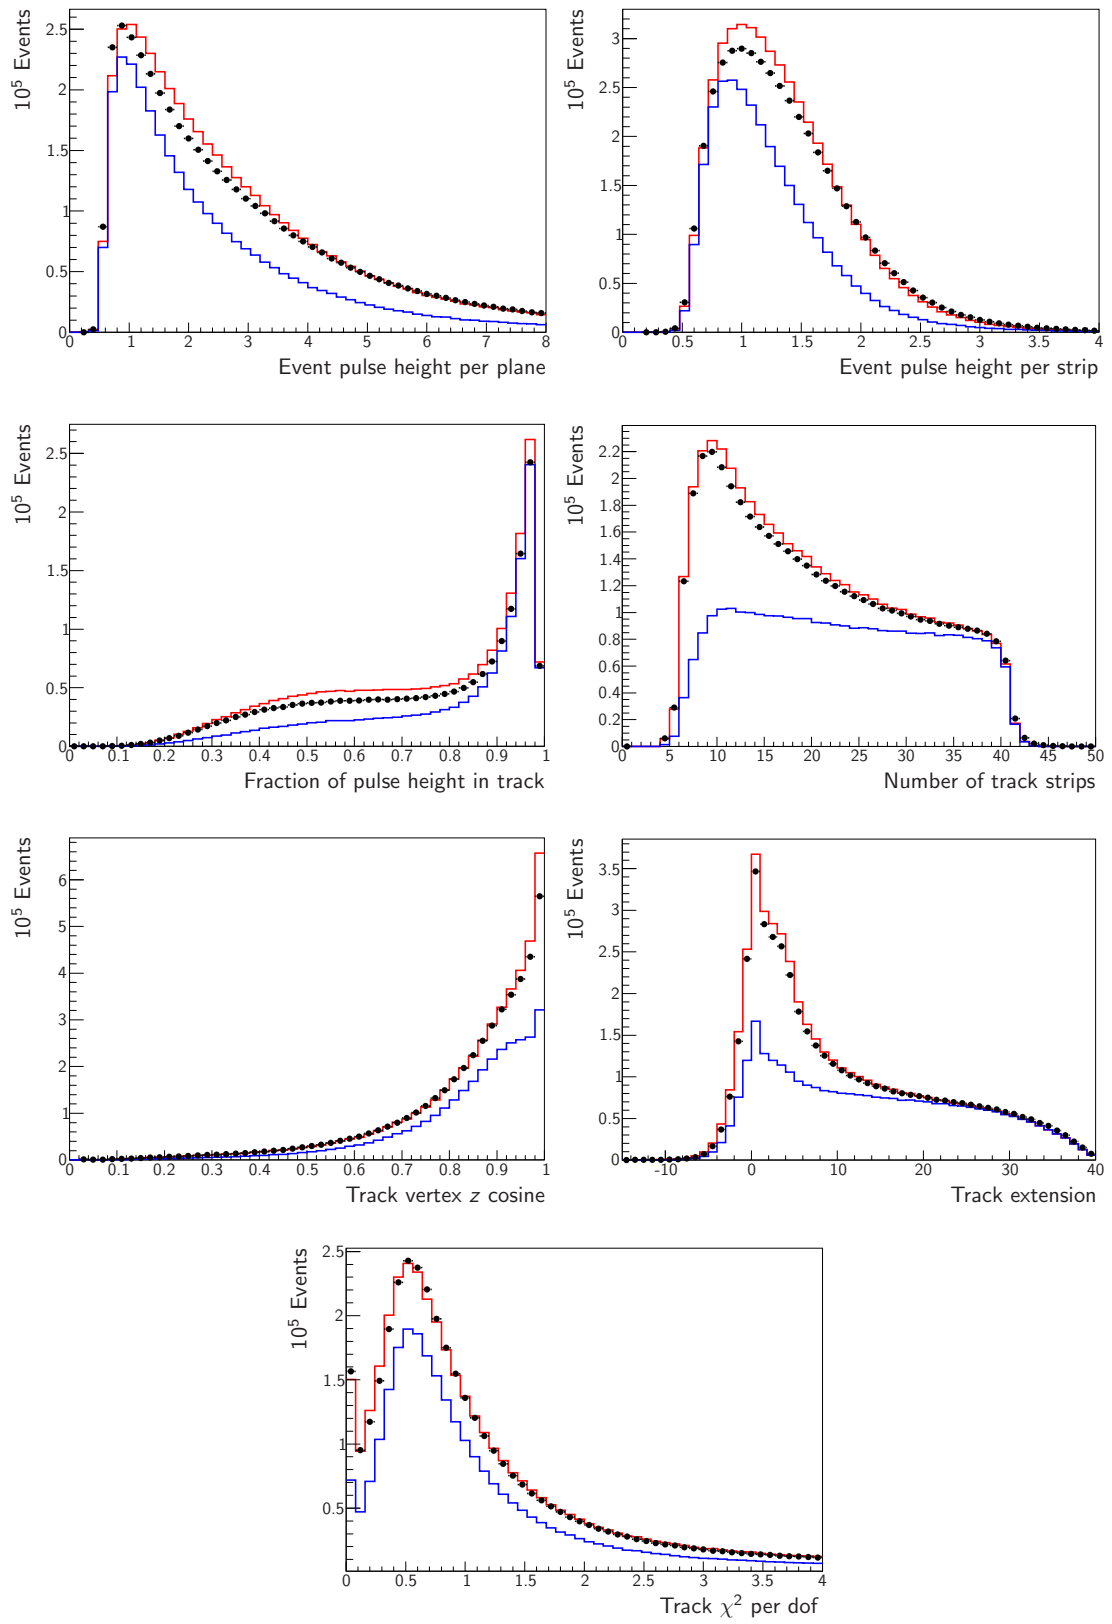


Figure 3.14: Input variables for the ANN. The red line is the near detector Monte Carlo and blue is the truly CC component. The black points are the near detector data.

applied). The purity is defined as

$$p = N_{\text{sel}} / (N_{\text{sel}} + C_{\text{sel}}) \quad (3.11)$$

where, as before, N_{sel} is the number of truly NC events selected, and C_{sel} is the number of truly CC events selected as NC.

Figure 3.16 shows the variation of efficiency and purity with cut position in unoscillated far detector Monte Carlo. A cut at 0 accepts only events without a track, and events with more than 40 planes will not be selected unless the cut is raised above 1.5. The figure shows that even cutting at very low values of the separation parameter produces a sample with around 30% CC contamination, highlighting the difficulty of separating high- y CC events (with a short muon track) from NC events in the MINOS detectors.

Although an NC/CC separation cut is not used *per se* in this thesis, it is sometimes necessary to define subsets of events that are NC or CC enriched. To this end, the cut value of 0.675 from [106] is used. This value was obtained by maximizing the figure of merit $\varepsilon p / (2 - p)$.

In the analysis presented in this thesis, the results of the preclassification step and ANN output are combined to form a “PID” variable thus:

$$\text{PID} = \begin{cases} 0 & \text{if } N_{\text{trk}} = 0 \\ 1.5 & \text{if } N_{\text{plane}} > 40 \\ \text{ANN output} & \text{otherwise.} \end{cases} \quad (3.12)$$

The fractions of events in each category are shown in Table 3.1, along with the NC/CC fractions in each.

Category	Fraction	Of which:	
		True NC	True CC
$N_{\text{trk}} = 0$	0.07	0.70	0.30
$N_{\text{plane}} > 40$	0.53	0.01	0.99
ANN output	0.40	0.35	0.65

Table 3.1: Fractions of events in each preclassification category in near detector MC.




Figure 3.15: Output of the ANN for near detector Monte Carlo and data. The red line is the Monte Carlo and blue is the truly CC component. The black points are the data.




Figure 3.16: Efficiency and purity of an NC selection using the ANN method as a function of cut position. Events below the cut value are selected as NC. The red line is efficiency and the blue line is purity. In each case, the preselection has been applied to both the numerator and the denominator.

Chapter 4

Predicting and Fitting the Far Detector Spectrum

This chapter presents a method of predicting and fitting the MINOS far detector spectrum in a two-dimensional scheme of energy *vs* the output of the ANN separation method described in Section 3.3.

4.1 Blind analysis

To avoid introducing subjective bias into the results of this analysis, it has been carried out in a ‘blinded’ fashion, in the same manner as published MINOS analyses [108, 109]: the details of the analysis method have been determined before looking at the far detector data¹. All near detector data was available for data/Monte Carlo comparisons and studies into systematic effects.

¹With one exception: after examining the far detector data, a technical problem relating to available Monte Carlo statistics was found, and the binning in energy was changed. See Section 4.5.2.

4.2 Fitting with separate NC and CC spectra

Previous analyses of the MINOS neutral current data [109, 106] have used two spectra, one selected to be mostly neutral current, and one mostly charged current, to fit neutrino oscillation parameters. The NC selected spectrum provides sensitivity to sterile neutrino oscillation parameters, while the CC spectrum constrains Δm^2 and f_μ . Although a measurement of these two parameters is not the principal aim of this thesis, constraining them precisely is vital for measurements of sterile neutrino oscillation parameters for two reasons: firstly, the charged current background in the NC-like spectrum oscillates according to Δm^2 and f_μ , so a more precise determination of these parameters leads to a more precise determination of the true neutral current component in the NC-like spectrum. Secondly, the deficit in the neutral current spectrum in the presence of active–sterile mixing depends on Δm^2 , as seen from Equation 1.41:

$$P(\nu_\mu \rightarrow \nu_s) = f_s f_\mu \sin^2(1.27 \Delta m^2 L / E). \quad (1.41)$$

Oscillations between active neutrinos leave the total flux of neutrinos coupling to the Z unchanged, so the signature of sterile neutrino mixing in a two-spectrum analysis is an energy-dependent deficit in the neutral current selected spectrum relative to the far detector prediction derived from near detector data. Equation 1.41 shows that this deficit has the same dependence on true neutrino energy as the deficit in the CC spectrum, but the energy carried away by the final state neutrino in an NC interaction results in a large smearing of the E_{reco} vs E_ν distribution, with $E_{\text{reco}} < E_\nu$ nearly always, and so the observed deficit with sterile neutrino mixing occurs at low energies as shown in Figure 4.1.




Figure 4.1: Reconstructed energy spectrum of truly NC events with sterile neutrino mixing for an exposure of 7.1×10^{20} protons-on-target (the data exposure used in this thesis). The black line is the reconstructed shower energy of truly NC events in the far detector in the absence of mixing to sterile neutrinos, i.e., with $f_s = 0$. The red line shows the corresponding spectrum with a large sterile mixing of $f_s = 0.5$, showing the energy-dependent deficit appearing at low energies.

4.3 An energy *vs* PID spectrum

4.3.1 Motivation

The analysis presented in this thesis aims to improve the sensitivity to sterile neutrino mixing parameters by taking advantage of information available from the PID value of events. In previous analyses, the PID is used only to make the binary decision of whether an event is placed in the NC or the CC spectrum. The continuous nature of the ANN output means that more information is available by dividing events into multiple energy spectra by PID value (i.e., by forming a two-dimensional histogram in event energy and PID value). Events in each of the PID bins can be expected to have different NC/CC probabilities, and different true energy distributions; both pieces of information which aid the oscillation fit. This can be seen by considering the oscillation fit as a measurement of the number of NC and CC events in the far detector spectrum: producing purer samples allows




Figure 4.2: Fraction of events that are truly NC as a function of ANN output value in unoscillated far detector MC. The dip near an ANN output of zero is an artefact of low statistics in that bin.

these numbers to be constrained more effectively. The effect of the differing true energy distributions is discussed below.

Figure 4.2 shows the first of these. At low ANN values, a fairly pure NC sample is obtained, while essentially all events with an ANN value near 1 are CC. At intermediate values, the NC/CC fraction varies between these two extremes. These pure NC and CC samples help to separately constrain f_μ , which controls the deficit in CC events, and f_s , which controls the deficit in NC events.

The true energy distributions for a few ranges of the PID value are shown in Figure 4.3. These show that events with a small PID value have, on average, a lower true energy than events with a larger PID value. Since the neutrino oscillation probability depends on the true energy, which is only weakly correlated with the reconstructed energy of the event for neutral current interactions, binning events in PID can allow the fit to better constrain the oscillation probability².

As described in Section 2.4, separate Monte Carlo samples have been generated for each of the MINOS data-taking periods, reflecting the differences in neutrino flux coming from a small target shift and from the inclusion of helium in the

²As an extreme case, imagine that the true energy distribution for a certain bin in PID was a delta function at energy E . Then a measurement of the number of events in that bin in the far detector would be a direct measurement of the oscillation probability at energy E .




Figure 4.3: True energy distributions for different PID ranges in the far detector. Events with a PID value of less than 0.33 are shown in black (this includes events with no track). PID values from 0.33 to 0.66 are shown in red, and PID values greater than 0.66 in blue. Each histogram is normalized to the same area.

decay pipe in the most recent data-taking period. In this analysis, each of the three running periods is included as a separate histogram in a joint fit, in order to maintain the maximum amount of information available for the fit. The total amount of Monte Carlo used is equivalent to approximately 6×10^{23} protons-on-target in the far detector and 2×10^{20} protons-on-target in the near detector.

4.4 Predicting the far detector spectrum

A key feature of the MINOS experiment is the use of two detectors to minimize systematic uncertainties. The large number of neutrino interactions in the near detector provides a high-statistics sample that can be used to predict the far detector spectrum with reduced dependence on the Monte Carlo simulation.

The near detector information can be incorporated into an analysis in more than one way, and several methods have been investigated for MINOS analyses [110]. These fall into two broad categories:

- ‘Direct’ methods, in which the near detector energy spectrum in data is directly used as the far detector prediction, with some extra correction taken




Figure 4.4: Schematic showing the geometrical effects giving rise to differences in near and far detector spectra. Neutrinos emitted at angles as large as θ_{near} may interact in the near detector, while only those with angle θ_{far} may interact in the far detector. Figure taken from [110].

from Monte Carlo.

- ‘Indirect’ methods, in which the near detector data is used to fit parameters in the Monte Carlo simulation (for example, cross sections or energy scales). The best-fit parameters from the near detector are then used in the far detector prediction.

Direct methods have the benefit of being able to correct for both known and unknown effects, while indirect methods, by their nature, can only correct for effects arising from the parameters used in the fit. On the other hand, if the information given to an indirect method is sufficient, it may be able to identify which physical parameters are being constrained by the near detector data, and so produce a far detector prediction that modifies only those parameters. The method used in this thesis is a direct method, chosen for its simplicity and ease of explanation.

In addition to the lower event rate at the far detector due to the divergence of the beam, there is a geometrical effect from the finite size of the detectors and the decay pipe in which parent mesons decay into neutrinos, shown schematically in Figure 4.4. The near detector subtends a much larger solid angle at the decay pipe, and so neutrinos produced at large angles to the beam direction may interact in the near detector, while only neutrinos in a much smaller range of angle will

intersect the far detector. Because of the kinematics of pion decay (the source of the majority of neutrinos in the NuMI beam), the energy of the neutrino E_ν depends on the angle to the pion direction as [110]

$$E_\nu \approx \frac{0.43E_\pi}{1 + \gamma^2 \tan^2 \theta} \quad (4.1)$$

where E_π is the parent pion energy, γ is the pion's Lorentz boost and θ is the angle in the lab between neutrino and pion directions. So the difference in range of angles interacting in the near and far detectors leads to a difference in energy spectra at the two.

This analysis corrects for these geometrical effects using the “Far/Near” method [109], which takes advantage of MINOS Monte Carlo, simulated with different neutrino fluxes at near and far detectors. In each (energy, PID) bin, the far detector prediction F is formed by multiplying the near detector data N by the ratio of far/near Monte Carlo:

$$F = \frac{F_{\text{MC}}}{N_{\text{MC}}} N. \quad (4.2)$$

A far detector prediction with oscillations is produced by reweighting the far Monte Carlo F_{MC} appropriately, as described in Section 4.4.1.

By rewriting Equation 4.2 in the following way, it can be seen that the Far/Near method can equally be thought of as correcting the far detector Monte Carlo prediction by the ratio of data/Monte Carlo in the near detector:

$$F = \frac{N}{N_{\text{MC}}} F_{\text{MC}}. \quad (4.3)$$

In this form, it is obvious that the method is able to correct for effects such as incorrect cross sections in the Monte Carlo. The Far/Near method also takes account of differences in reconstruction and selection efficiency between the detectors to the extent that they are modelled by the Monte Carlo, since the far and near

MC spectra used have the relevant selections applied (in this case, the near and far preselections described in Sections 3.2.1 and 3.2.2).

Figures 4.5 and 4.6 demonstrate the effect of the Far/Near method considered as Equations 4.3 and 4.2 respectively. In Figure 4.5 the near detector data/MC ratio is shown, demonstrating the level to which the Far/Near method “corrects” the far detector MC prediction. Data and MC agree well at energies above 10 GeV, while the data is lower than the Monte Carlo at lower energies. This effect is particularly noticeable in the region of low PID, which contains most of the NC events. A possible cause of this effect is a mismodelling of the neutral current cross section as a function of the inelasticity ($y \approx E_{\text{had}}/E_\nu$), with an excess of low- y events in the Monte Carlo, or simply a mismodelling of the cross section as a function of energy. The data/MC agreement for high PID values, which are dominated by charged current events, is good across all energies. This is because the beam fit described in Section 2.5.1 is done on a charged current selected spectrum.

Figure 4.6 shows energy projections in the near and far detectors, for the same bins in PID as in Figure 4.5. This demonstrates the level to which the Far/Near method does more than just scale the near detector data by $1/L^2$; as the plots are normalized to the same area, such a scaling would result in two lines lying on top of one another. The level of difference between the near and far spectra is small in the CC-like region, with larger differences occurring in the low energy region of the low PID projection, and in the NC/CC overlap region. The former is due to the near preselection, which has its largest effect in these bins. Additionally, any differences in reconstruction efficiency would be expected to occur in low energy NC-like events, which leave little pulse height in the detector over a small number of strips, and are therefore most susceptible to failing the reconstruction stage. Specific differences between the detectors which could lead to reconstruction efficiency differences for low energy NC-like events are light levels in the scintillator

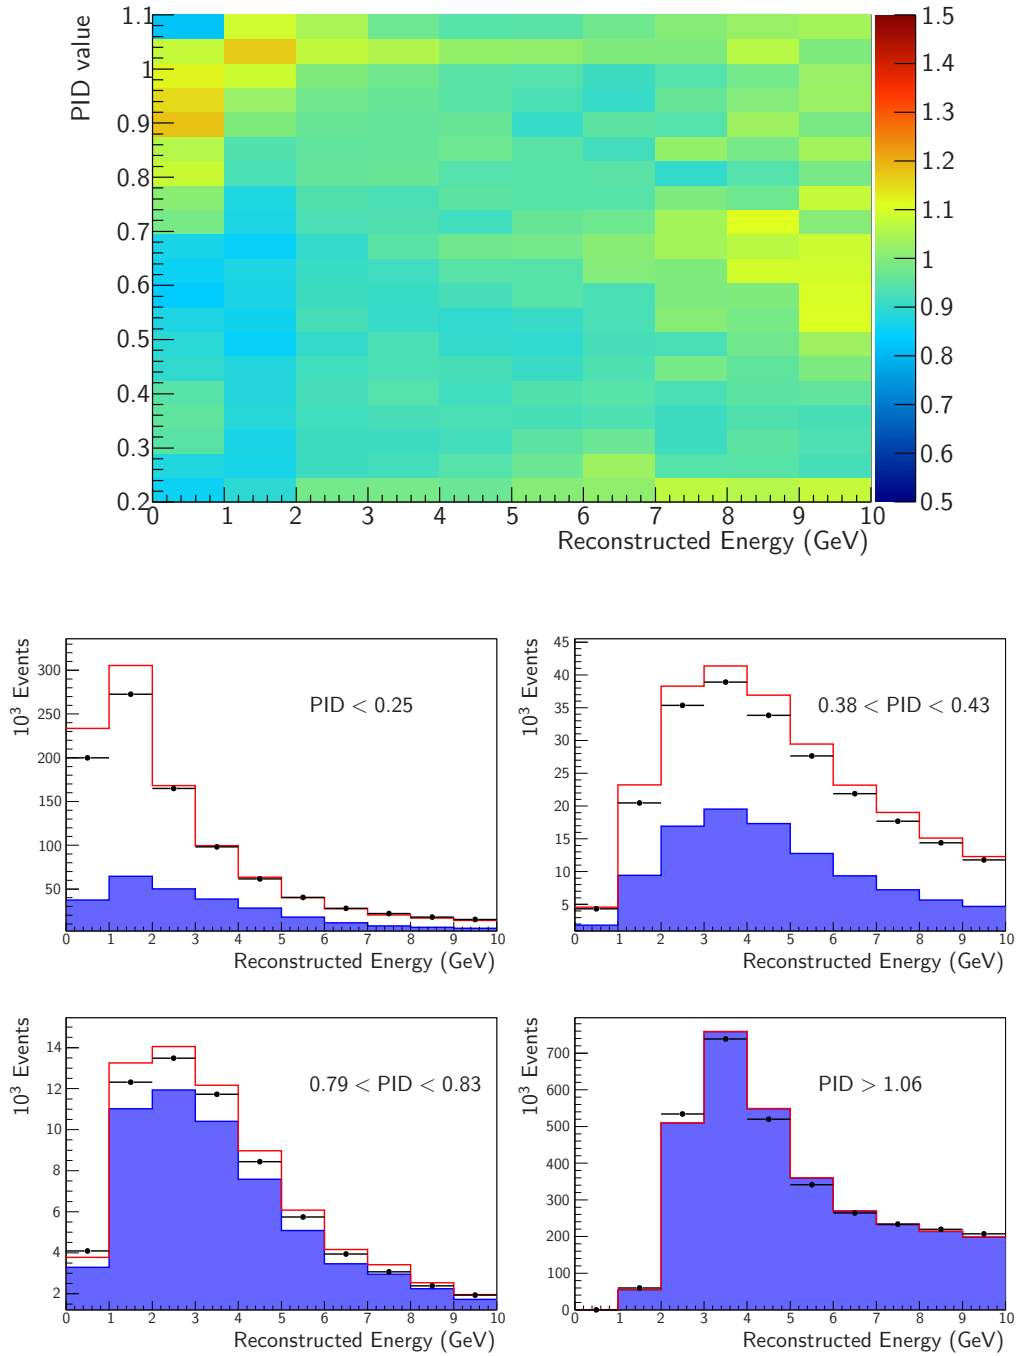


Figure 4.5: Data/MC ratios for selected events in the near detector, indicating the amount to which the Far/Near method “corrects” the far detector MC prediction. At the top is the ratio for all bins up to 10 GeV used in the analysis, while the four plots below it show the projection on the reconstructed energy axis for four bins in PID. The colour scale on the top plot shows the ratio, while on the lower plots, red is the total MC, blue is the truly CC component, and the black points are the data.

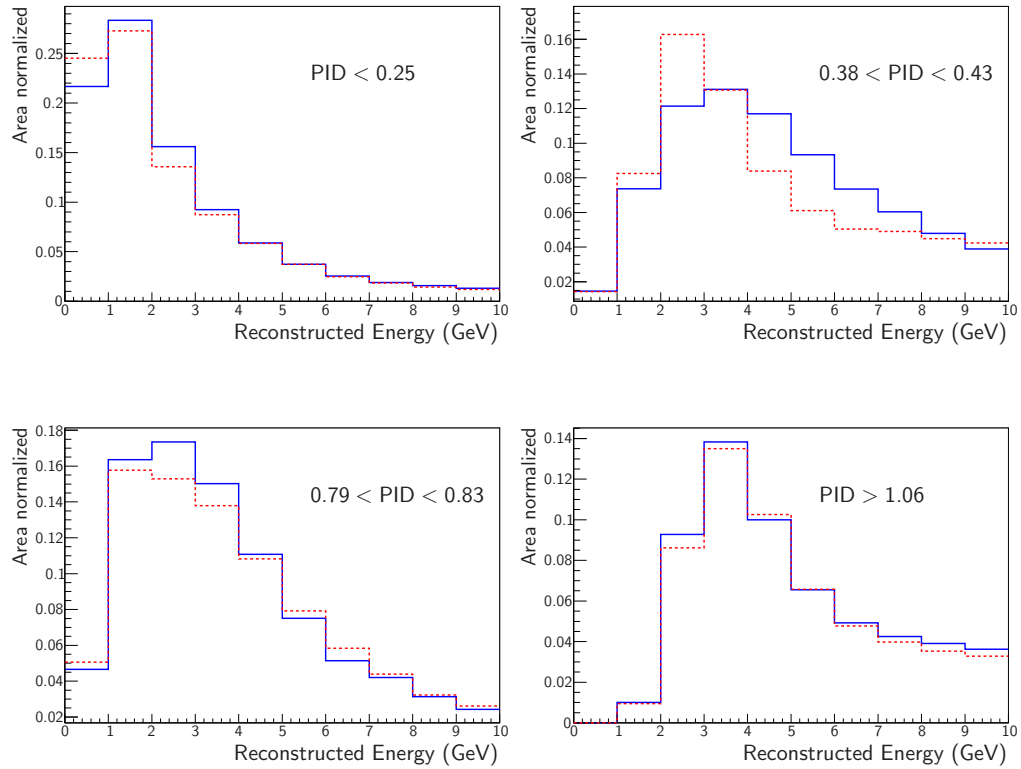


Figure 4.6: Near and far detector energy projections in various bins of PID, indicating the level of correction for flux differences and selection efficiencies provided by the Far/Near method. The near detector projection is shown in blue and the far detector MC projection in the absence of oscillations as the dashed red line. The plots are normalized to the same area.

and PMTs, and the demultiplexing performed by the reconstruction software in the far detector (which does not occur in the near detector). The light level has an effect because, although the overall energy scales are calibrated between the detectors, different light levels may lead to the same event passing a readout threshold in one detector and not the other.

4.4.1 Producing an oscillated prediction

The quantity F_{MC} in Equation 4.2 is calculated at certain values of the oscillation parameters as follows. The number of selected neutrino events of a certain flavour with reconstructed energy E_{reco} is

$$N(E_{\text{reco}}) = k \int S(E_{\text{reco}}, E) \Phi_0(E) \sigma(E) \epsilon(E) P(E, \boldsymbol{\nu}) dE \quad (4.4)$$

where E is the true neutrino energy, $\Phi_0(E)$ is the flux in the absence of oscillations, $\sigma(E)$ is the cross section for a neutrino interaction, $\epsilon(E)$ is the efficiency to select an interaction in the sample (including reconstruction efficiency) and $P(E)$ is the oscillation probability at oscillation parameters $\boldsymbol{\nu}$. $S(E_{\text{reco}}, E)$ is the probability that a neutrino with true energy E has reconstructed energy E_{reco} , and k is a scale factor for the appropriate beam exposure. The events available from the MINOS Monte Carlo simulation are reconstructed, and generated from the unoscillated beam flux generated with the best-known cross sections, so the number of MC events folds together several of the quantities into the number of events expected in the absence of oscillations as a function of true energy, $T(E)$, defined by

$$T(E) \equiv k \Phi_0(E) \sigma(E) \epsilon(E). \quad (4.5)$$

The most accurate way to calculate the integral in Equation 4.4 is via a Monte Carlo calculation. In terms of the MINOS simulation, this means taking all events

in the simulation and reweighting them by the appropriate oscillation probability $P(E)$. This would require a loop over all Monte Carlo events for each new set of oscillation parameters, so to reduce computational time used by the fit, the integral is approximated by a sum:

$$N_i(\boldsymbol{\nu}) = \sum_j S_{ij} T_j P_j(\boldsymbol{\nu}) \quad (4.6)$$

where indices j indicate true energy bins, and indices i indicate reconstructed energy bins. Only $P_j(\boldsymbol{\nu})$ must be recalculated at each new value of the oscillation parameters. It is this value $N_i(\boldsymbol{\nu})$ that is used as F_{MC} in Equation 4.2 for the far detector prediction. To provide a good approximation, the true energy bins must be narrow: the binning used in true energy is 10 times finer than the reconstructed energy binning. The oscillation probabilities in P_j are calculated at the energy corresponding to the bin centre. The two-dimensional histogram S_{ij} is normalized such that

$$\sum_j S_{ij} = 1 \quad \forall i. \quad (4.7)$$

Equation 4.6 is sufficient for only one flavour of neutrino. In the MINOS far detector, all three flavours of neutrino may interact: ν_μ and ν_e are present in the beam at production, while ν_τ are produced by oscillations of the ν_μ beam. Additionally, for nonzero θ_{13} (nonzero f_e in the four-flavour model used in this thesis), there are ν_e due to oscillations present. Each of these components is calculated separately with the appropriate oscillation probability, and the total bin content is the sum of the components. The default MINOS Monte Carlo uses the beam composition at production, and therefore provides the ν_μ and beam ν_e components in the far detector. Oscillated ν_e and ν_τ events are provided by special simulations in which the beam flux is converted entirely to ν_e or to ν_τ .

An indication of the accuracy of the approximation is shown in Figure 4.7. The




Figure 4.7: Indication of the accuracy of the oscillation calculation. The energy spectrum integrated over all PID values is shown. The dashed red line is the Monte Carlo spectrum without oscillations, the solid red line is the prediction at the true oscillation parameters, and the black histogram is the fake data with Gaussian error bars. The line on the lower plot is the ratio of prediction and fake data.

prediction at the true oscillation parameters is shown along with the fake data, and their ratio. Since the fake data is calculated by weighting events individually, it represents the most accurate oscillated spectrum available from the Monte Carlo, so if the calculated prediction is a good approximation, it will agree with the fake data everywhere. The ratio is 1 to much better than 1% across all energies. The slight deviation at high energies is because of the wide 10–120 GeV energy bin, which, like the other bins in energy, is split into 10 true energy bins. This results in wider true energy bins, and therefore a less precise approximation in Equation 4.6. Using this approximation, the total time to run the fits described in Chapter 6 is a matter of a few hours.

4.5 Sensitivity to neutrino oscillation parameters

4.5.1 Statistical sensitivity

This analysis is optimized by maximizing the statistical sensitivity to oscillation parameters. The sensitivity is calculated by fitting a ‘fake’ data set using the Monte Carlo. This fake data set is generated for the far detector by taking the Monte Carlo and reweighting each event by the appropriate oscillation probability (Equations 1.38–1.41) based on its true energy, and a factor to account for the exposure of the experiment (the exposure used for this thesis is 7.1×10^{20} protons-on-target). The number of fake data events in a bin may therefore be non-integer, and contains no statistical fluctuations. In the near detector, the fake data is simply the Monte Carlo scaled to the appropriate exposure. This means that the sensitivities derived from fake data studies take no account of the effects of systematic uncertainties, which are dealt with in Chapter 5.

For these sensitivity studies, a value of $f_s = 0$ is used, along with $f_e = 0$, the former chosen because of the strong SuperK bound on sterile neutrino mixing in the atmospheric sector [74], and the latter chosen for simplicity. The MINOS best-fit values for the active neutrino mixing parameters are also used [57]:

$$\Delta m^2 = 2.35 \times 10^{-3} \text{ eV}^2/\text{c}^4 \quad \sin^2 2\theta_{23} = 1 \quad (4.8)$$

The best-fit point is found by minimizing a χ^2 statistic for the far detector spectrum, depending on the oscillation parameters $\boldsymbol{\nu}$:

$$\chi_P^2(\boldsymbol{\nu}) = 2 \sum_i \left[e_i(\boldsymbol{\nu}) - o_i + o_i \ln \frac{o_i}{e_i(\boldsymbol{\nu})} \right] \quad (4.9)$$

where $e_i(\boldsymbol{\nu})$ is the number of events expected in bin i for the given set of oscillation

$(1 - \alpha)$ (%)	$m = 1$	$m = 2$
68	1.00	2.30
90	2.71	4.61

Table 4.1: $\Delta\chi^2$ values for confidence level $1 - \alpha$ with m parameters. Adapted from [32].

parameters, calculated as described in Section 4.4 above, o_i is the number of events observed in that bin, and the sum runs over all bins. Minimizing the χ_P^2 statistic is equivalent to maximizing the ratio of Poisson likelihoods [111]

$$\mathcal{L}(\boldsymbol{\nu}) = \prod_i \frac{\text{Po}(o_i; e_i(\boldsymbol{\nu}))}{\text{Po}(o_i; o_i)}, \quad \text{where } \text{Po}(x; \lambda) = \frac{e^{-\lambda} \lambda^x}{x!}. \quad (4.10)$$

For fake data as described here, the best-fit point will be the input point, while for real data it may differ from the true point. This property of the estimator—in the limit of infinite statistics, the estimator converges on the true point—is known as *consistency* [112].

With the best-fit point $\boldsymbol{\nu}^*$ found, limits on a parameter, or set of parameters, are determined by finding the points $\boldsymbol{\nu}$ where

$$\Delta\chi_P^2 \equiv \chi_P^2(\boldsymbol{\nu}) - \chi_P^2(\boldsymbol{\nu}^*) = C \quad (4.11)$$

where C is a constant that depends on the desired confidence level and the number of parameters being constrained. Some common values of C are shown in Table 4.1. When there are more parameters in the fit than are being constrained, the remaining parameters are *marginalized* over by minimizing the χ_P^2 in the remaining parameters while holding the desired parameters fixed. This allows the uncertainty in those parameters to be taken into account when quoting limits on the desired parameter(s).

4.5.2 Reconstructing neutrino energy

For CC ν_μ events in MINOS, the best estimator of the incident neutrino energy is the sum of track and shower energies, on the assumption that the track energy accurately estimates the muon energy, while the shower energy estimates the energy of the recoiling hadronic system. For NC events, the incident neutrino energy cannot be reconstructed due to the neutrino in the final state. The best energy estimator to use in an analysis is therefore ambiguous: the published MINOS NC analysis [109] uses the reconstructed shower energy for events classified as NC, for example.

For this analysis, no separate NC-like and CC-like spectra are formed, so the same energy estimator is used for all events. Precise neutrino energy reconstruction is vital in order to use the CC events to constrain the parameters Δm^2 and f_μ , and so the sum of track and shower energies is used. The interpretation of this quantity is not as transparent for neutral current events: tracks in NC events either follow a non-muon particle or are ‘drawn-on’ by the reconstruction, whereas the calculation of track energy assumes a muon track. Nonetheless, the same estimator is used in data and Monte Carlo.

The binning scheme in energy is as follows: 10 1-GeV bins from 0–10 GeV, and one overflow bin from 10–120 GeV. This binning scheme provides sensitivity to the energy dependence of the oscillations. The large overflow bin ensures that, even with a large number of bins in PID (see Section 4.5.3), the statistics in each (PID, energy) bin are not too small. This binning scheme was adopted after fitting the far detector data; previously the 1 GeV bins extended to 20 GeV, and the overflow bin covered the region 20–120 GeV. This resulted in low Monte Carlo statistics in some bins, and problems in the fitting of systematic parameters, so the current binning was adopted. All fits and spectra in this thesis use the new binning scheme for consistency.

4.5.3 Choice of binning in PID

From Figure 3.15, it can be seen that there are few events below an ANN value of 0.2, or above a value of 1.1. Therefore the lowest PID bin has its low edge at 0.2, and the highest has its high edge at 1.1. Events with PID values outside this range are placed in the highest or lowest bin as appropriate. This includes the preclassified events. In this way, all events are used in the analysis. Bin edges are equally spaced between the lower and upper limits.

Clearly, from the arguments in Section 4.3.1 above, increasing the number of bins in PID should increase the sensitivity. Nonetheless, beyond a certain number of bins, little extra information will be available to the fit, and the sensitivity increase will be small. Also, a large number of bins increases the computing time taken to run the fit, and may increase the dependence on Monte Carlo statistics. This is because each (energy, PID) bin has an associated true energy distribution which is used to calculate the oscillated spectrum as described in Section 4.4.1. The full Monte Carlo sample is therefore divided up among a large number of bins (of order 10^3), with the result that the number of events in each (reco energy, PID, true energy) bin may be small if the number of PID bins is large.

Figure 4.8 shows the dependence of the 90% sensitivity upper limit on f_s as a function of the number of bins in PID used. The limit becomes slightly tighter (by 2.6%) up to about 20 bins, but beyond this, there is no significant improvement. Therefore, for this thesis, 20 bins in PID are used.

Figure 4.9 shows the resulting sensitivity with 20 PID bins as a contour in f_s *vs* Δm^2 . Also shown is the same contour for the analysis method using one NC and one CC spectrum, as published in [109]. A moderate improvement is seen for the analysis presented in this thesis.

Energy projections of the spectrum used in the fake data fit are shown for each PID bin in Figure 4.10. These show the total Monte Carlo for no oscillations and




Figure 4.8: 90% confidence limit sensitivity on f_s as a function of number of PID bins. The points with labels show the values at which the limit was evaluated and the line is a linear interpolation between them. A value of 20 bins is used in the analysis.




Figure 4.9: 90% C.L. sensitivity contours on f_s vs Δm^2 for this analysis and the published MINOS NC analysis, with a simulated exposure of 7.1×10^{20} protons-on-target. The star shows the input value used.

	Δm^2 ($10^{-3} \text{ eV}^2/c^4$)	f_μ	f_s
Set 1	2.35	1	0
Set 2	2.35	0.84	0.2

Table 4.2: True oscillation parameters used in the generation of mock data sets 1 and 2.

for oscillations at the best fit point (which is equal to the true point for fake data), along with the fake data itself.

4.6 Mock data studies

The sensitivity calculation of the previous section provides an accurate guide to the average result expected from an ensemble of hypothetical experiments carried out when the true values of the parameters are as given. However, since the expected numbers of events in each bin in the far detector are small, the fractional statistical fluctuations in these numbers are large, and so the result of any individual experiment may differ by a large amount from the true value.

Producing an ensemble of such experiments and fitting them individually is instructive as a guide to what results from an actual experiment might look like. It also provides a method of checking that the analysis code is able to successfully fit something that looks much more like realistic data than the fake data of the previous section, without violating the blinding scheme by fitting real far detector data.

4.6.1 Input data sets

Two such ensembles of ‘mock data’ were produced, with true oscillation parameters as shown in Table 4.2. Set 1 uses the best-fit parameters from the MINOS CC analysis, with no sterile mixing, while set 2 includes sterile mixing at a small level, along with non-maximal mixing in f_μ .

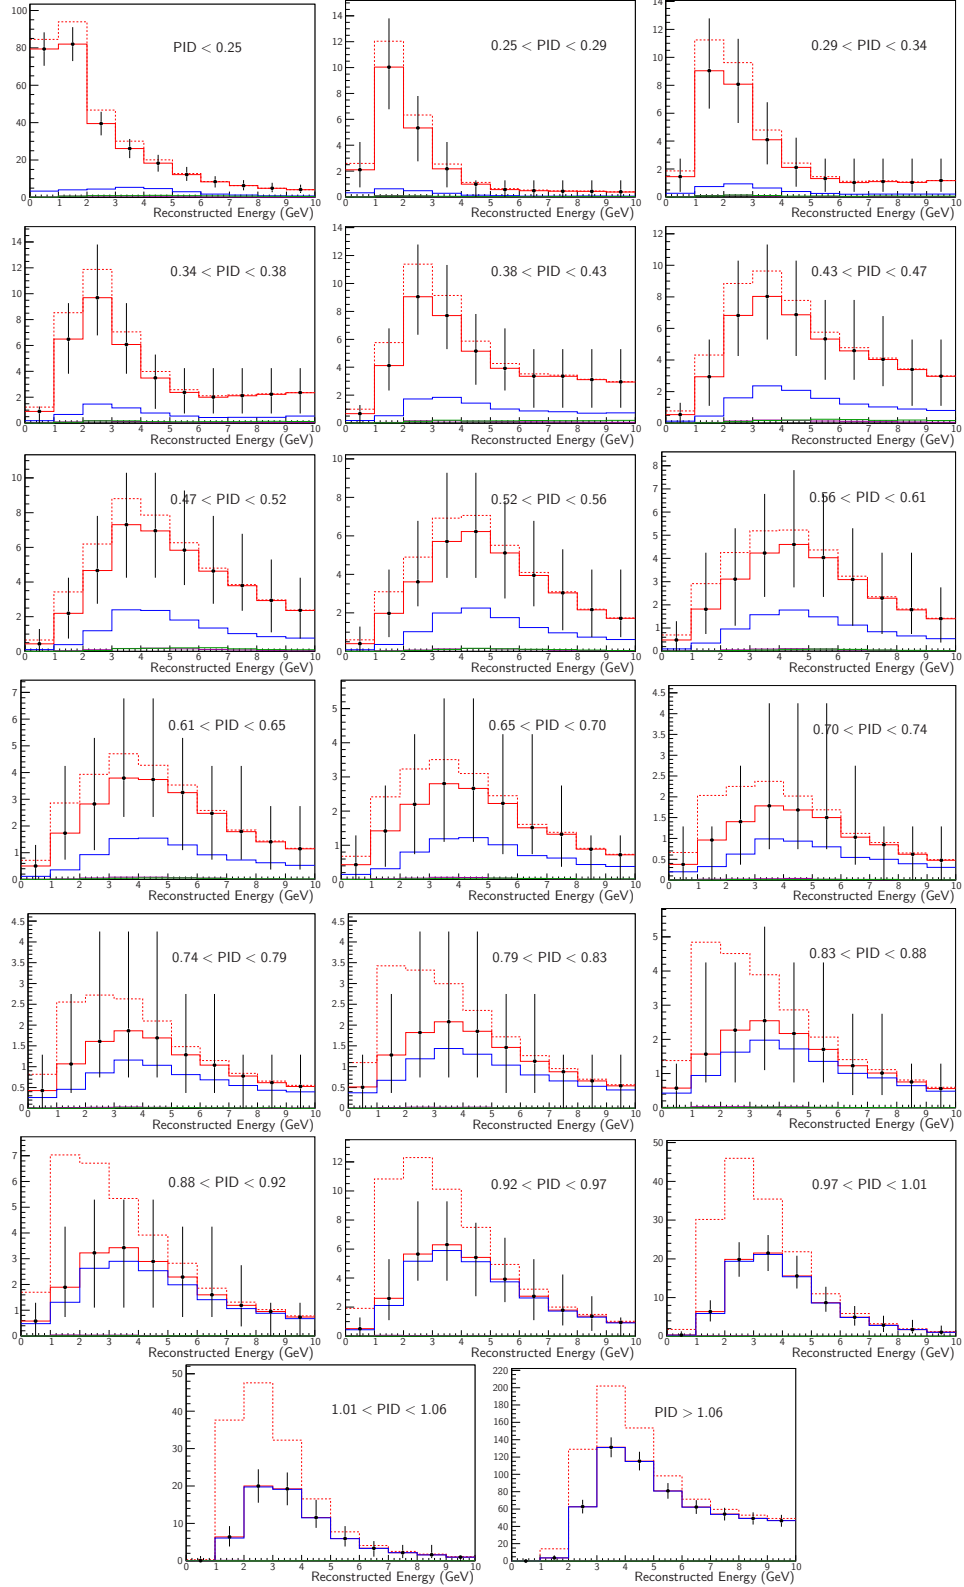


Figure 4.10: Energy projections of spectrum used in sensitivity calculation for each bin in PID. PID increases from left-to-right, top-to-bottom. The dashed red line is the Monte Carlo prediction for no oscillations, and the solid red line is the total Monte Carlo oscillated at the best fit point. Also shown (oscillated) are CC ν_μ events in blue, CC ν_τ events in purple and CC ν_e events coming from the beam in green. The latter two components are small. The fake data is shown in black with error bars appropriate for a Poisson variable.

Each set contains 100 statistically-independent subruns in the far detector with a number of events appropriate for the exposure used for the data (7.1×10^{20} protons-on-target, split among running periods the same way as in the data). That is, the total number of events varies between subruns, but the number is a sample from a Poisson distribution with the appropriate mean. For the near detector, the full Monte Carlo was used in the role of both Monte Carlo and data, with the result that this study is statistical only, and takes no account of systematics.

The far detector mock data files were produced by choosing events from the MINOS Monte Carlo at random, with a probability P defined by:

$$P = P_{\text{osc}} \times F \times W \quad (4.12)$$

where P_{osc} is the oscillation probability of the event, F is the event weight given by the MINOS flux tuning procedure, and W is the desired exposure in the subrun divided by the total exposure available in the Monte Carlo file. This procedure produces statistically independent subruns if $P \ll 1$, so that the overlap in events between subruns is small. For the Monte Carlo used, $P \sim 10^{-3}$, which is sufficiently small. Since the events are drawn from the same Monte Carlo used to fit the resulting spectra, the mock data sets are not strictly independent from the Monte Carlo, but the condition $P \ll 1$ ensures that any effect coming from this is small.

A simpler method of producing statistically-fluctuated mock data would be to start with the mean numbers of events in each histogram bin as determined from fake data, and then produce a mock data histogram with bin values that are picked from a Poisson distribution with this mean. This method would have the advantage that the samples would be guaranteed to be statistically independent, but has the disadvantage of working on histograms, rather than individual events, making the mock data less similar (in terms of analysis code) to the real data.

Therefore the method used in this thesis provides a more realistic test of the analysis code.

4.6.2 Results

The distributions of best-fit points are shown in Figures 4.11 and 4.12 for mock data set 1, and in Figures 4.13 and 4.14 for mock data set 2. Each shows the fit points clustered around the true parameter values, and shows no pathologies in fitting, indicating that the analysis code is sound.

Further plots can be found in Appendix A, showing some individual subruns and demonstrating the level of statistical fluctuation expected in 7.1×10^{20} protons-on-target of data.

It is worth noting that these distributions show that the estimators used are not *unbiased*, in the statistical sense of the word. An unbiased estimator $\hat{\theta}$ of a parameter with true value θ_0 is one for which the expectation of $\hat{\theta}$ is equal to the true parameter [112], i.e.:

$$E[\hat{\theta}] = \theta_0. \quad (4.13)$$

For the parameters f_μ and f_s , the best-fit point is clearly a biased estimator for the input values used; since the inputs are on a physical boundary, and the best-fit point is constrained to the physical region, the best-fit point can only be unbiased if it has no variance, which is clearly not the case. The best-fit point for Δm^2 is always far from the physical boundary, but still shows some bias; its mean value in mock data set 1 is $2.43 \times 10^{-3} \text{ eV}^2/\text{c}^4$. Although unbiasedness is a desirable property in an estimator, it is less important than consistency, the property that, in the infinite statistics limit, the estimator converges to the true point. Fits on high-statistics fake data produce a best-fit point equal to the true point, showing that the estimators used here are consistent.

The χ_P^2/dof variable peaks at around 1, with a mean value of 0.97, as expected

if χ_P^2 follows a χ^2 distribution with the appropriate number of degrees of freedom.

4.7 Effect of ν_e appearance

In the sensitivities and spectra presented so far, it has been assumed that the $\nu_\mu \rightarrow \nu_e$ probability, parametrized in Equation 1.39 by f_e , is zero. However, ν_e produced by $\nu_\mu \rightarrow \nu_e$ oscillations are a background to a neutral current analysis, since the ν_e populate the NC-like part of the PID spectrum. The current best upper limits for the parameter θ_{13} , which drives $\nu_\mu \rightarrow \nu_e$ oscillations in the three-flavour model, come from the CHOOZ reactor experiment [53] and the MINOS ν_e appearance analysis [99]. The CHOOZ 90% C.L. upper limit is $\sin^2 2\theta_{13} < 0.19$ [32]. This can be converted to a value of f_e by comparing Equation 1.39

$$P(\nu_\mu \rightarrow \nu_e) = f_e f_\mu \sin^2 \Delta \quad (1.39)$$

with the first-order approximation to the probability in the three-flavour model [113]:

$$P(\nu_\mu \rightarrow \nu_e) = \sin^2 \theta_{23} \sin^2 2\theta_{13} \sin^2 \Delta. \quad (4.14)$$

Assuming maximal mixing between m_2 and m_3 , $\sin^2 \theta_{23} = f_\mu = 1$, and so

$$f_e = \sin^2 2\theta_{13}. \quad (4.15)$$

Figure 4.15 shows the distribution of PID value in the far detector Monte Carlo with electron neutrino appearance at the CHOOZ limit. The oscillated ν_e component occurs almost entirely in the $\text{PID} < 0.6$ region which is dominated by NC events. To deal with this, all fits will be run at fixed values $f_e = 0$ and $f_e = 0.19$, and both results quoted.




Figure 4.11: Best-fit points for mock data set 1. The upper plot shows the points in f_s and Δm^2 , while the lower plot shows the points in f_μ and Δm^2 . The blue markers are the best fits from each of the 100 subruns, and the red star is the true input point. The black line is the statistical sensitivity at the true parameters for an exposure of 7.1×10^{20} protons-on-target.

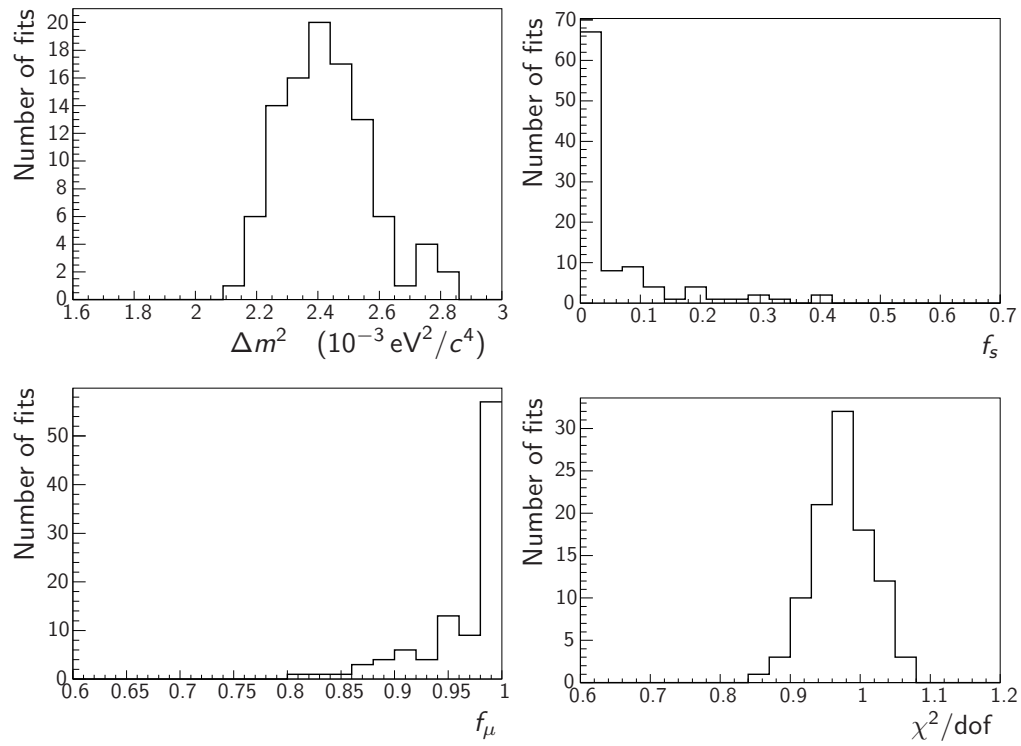


Figure 4.12: Distribution of parameter values at the best fit point for mock data set 1. From left-to-right, top-to-bottom, the parameters shown are Δm^2 , f_s , f_μ and the χ^2 per degree of freedom.




Figure 4.13: Best-fit points for mock data set 2. The upper plot shows the points in f_s and Δm^2 , while the lower plot shows the points in f_μ and Δm^2 . The blue markers are the best fits from each of the 100 subruns, and the red star is the true input point. The black line is the statistical sensitivity at the true parameters for an exposure of 7.1×10^{20} protons-on-target.




Figure 4.14: Distribution of parameter values at the best fit point for mock data set 2. From left-to-right, top-to-bottom, the parameters shown are Δm^2 , f_s , f_μ and the χ^2 per degree of freedom.




Figure 4.15: Oscillated far detector Monte Carlo PID distribution with electron neutrino appearance at the CHOOZ limit. The components are shown stacked, with colours as shown in the legend. The first and last bins have been reduced by a factor of 10 for display purposes.

Chapter 5

Systematic Uncertainties

The two-detector design of the MINOS experiment is intended to reduce the dependence of analyses on accurate Monte Carlo modelling of physics processes and detector response. Nonetheless, the far/near ratio method is not able to cancel out all of these uncertainties fully. This chapter describes the sources of uncertainty considered for this analysis, and presents an improved constraint on the near/far normalization uncertainty compared to previous analyses. The effect of the systematic uncertainties on the best-fit oscillation parameters is compared, and those with the largest effect are included in the fit.

5.1 Sources of uncertainty

5.1.1 Relative near/far normalization

The relative near/far normalization takes account of all sources of uncertainty that result in an overall scale in the near/far ratio. For MINOS analyses prior to 2010, a value of 4% was used for this uncertainty, made up of the components shown in Table 5.1.

The uncertainties on steel and scintillator thickness affect the fiducial mass and were determined from measurements taken during construction [115].

Systematic	Contribution to relative uncertainty (%)
Steel thickness	0.2
Scintillator thickness	0.2
FD live time	1.0
ND fiducial bias (z)	1.9
ND fiducial bias (y)	0.7
ND fiducial bias (x)	0.7
N/F reconstruction efficiency	3.0

Table 5.1: Contributions to the near/far normalization uncertainty used in pre-2010 MINOS analyses. Taken from [114].

The FD live time is around 99%, and the live time uncertainty was taken as 100% of the inefficiency, i.e., 1%. This is clearly a very conservative estimate; the upper bound it implies would correspond to the detector being live on every occasion it was recorded as down. This component also takes account of uncertainty in the total number of protons-on-target (POT) used in the analysis, which may arise from various technical issues such as files containing data failing to reconstruct, but having their associated exposure still included in the count of protons-on-target.

The ND fiducial bias was originally estimated in [116] as a way to take account of the deviations from uniformity in the mass distribution across the detector. The MINOS Monte Carlo contains no such nonuniformities—every plane has the same thickness and density—so any nonuniformity in the real detector induces a systematic error. The fiducial bias was determined in [116] as follows: the ND fiducial volume was divided in half in x , y or z , and the data/MC ratio in each was calculated. The asymmetries

$$R_{x,y,z} = \frac{|A - B|}{2(A + B)} \quad (5.1)$$

was calculated, where A was the data/MC ratio in the first half of the fiducial volume, and B was the data/MC ratio in the second half, in the appropriate co-

ordinate. The value used in the systematic was the central value plus its one-sigma statistical error, a conservative estimate.

The largest component of the near/far normalization uncertainty is the reconstruction efficiency. This component refers to a possible difference in reconstruction efficiency between the near and far detectors unaccounted for by Monte Carlo. The far detector prediction is essentially a correction on the far detector Monte Carlo based on the near detector data/Monte Carlo ratio, so there is no overall systematic effect if the double ratio of efficiencies,

$$R \equiv \frac{\epsilon_{\text{data}}^F / \epsilon_{\text{MC}}^F}{\epsilon_{\text{data}}^N / \epsilon_{\text{MC}}^N} = 1 \quad (5.2)$$

where N indicates near detector and F indicates far detector. Therefore any estimate of this contribution to the uncertainty must take account of all four combinations of near/far and data/MC. The value of 3% shown in Table 5.1 is derived from a hand scan of events described in [117].

Updated contributions

The values given in Table 5.1 are generally rather conservative estimates. For early MINOS analyses with low statistics, such conservative estimates were tolerable. With the 7.1×10^{20} protons-on-target exposure used in this thesis, the statistical errors are much smaller and so systematics have a larger relative effect. Therefore, the individual components of the normalization uncertainty are rederived here, mostly by combining work already done by others. The new values have also been used in recent MINOS analyses [99, 57].

Reducing the steel and scintillator thickness uncertainties would require new measurements of the detector itself, and so these contributions to the uncertainty remain unchanged.

The contribution from the FD live time uncertainty has been reduced signifi-

cantly, based on two strands of work. Firstly, an improved understanding of data quality in the Far Detector is shown in [118], which results in essentially negligible uncertainty in the actual livetime of the detector.

Secondly, an extremely thorough check of the counting of protons-on-target used in the MINOS neutral current analysis [109] was carried out [119]. It compared two methods of counting the protons-on-target: one from a count in the reconstructed files; and one from a count in a database table. These agreed well, except for a month with a known problem, for which the database numbers are more reliable. The ratio of POT from files to POT from the database (excluding the month with a known problem) is 0.9967, giving a contribution to the normalization uncertainty of 0.32%.

The toroids in the NuMI beamline which measure the number of protons hitting the target have a systematic uncertainty of 1% [120]. However, this cancels in the near/far ratio, and so does not contribute to the normalization uncertainty.

Ordinate	Asymmetry (%)
x	0.53 ± 0.15
y	0.14 ± 0.15
z	0.43 ± 0.15
Sum in quadrature	0.7

Table 5.2: Updated ND fiducial bias numbers with statistical uncertainty.

Table 5.2 shows updated numbers for the ND fiducial bias in each co-ordinate. These numbers were calculated by repeating the study in [116] with more recent, higher-statistics Monte Carlo and data. Events selected as CC by the method described in [109] were used. The Monte Carlo and reconstruction versions used in the study differed from those used in the rest of this thesis, but the conclusions should not be expected to change: the aim is to constrain effects resulting from different mass distributions in the data and Monte Carlo, which are independent of the software version used.

An updated contribution to the reconstruction efficiency uncertainty is pre-

sented in Section 5.2.

5.1.2 Shower energy scales

Two types of overall shower energy scale are considered: an *absolute* shower energy scale, common to both detectors; and a *relative* energy scale uncertainty between the two detectors.

Absolute shower energy scale

An error in the absolute shower energy scale means that a sample of showers of a given energy E will receive a different average reconstructed energy in data and Monte Carlo. A difference between true and reconstructed shower energy would not induce a systematic error provided that the difference is correctly modelled by the Monte Carlo.

There are two contributors to this uncertainty. The first is the uncertainty on the overall detector response to single hadrons derived from the MINOS calibration detector, which contributes a 5.7% uncertainty [121]. There is also an uncertainty that arises from hadronic shower modelling in the Monte Carlo: some of the energy transferred to the hadronic system is lost to intranuclear rescattering, and so does not appear in the detector. Although this effect is included in the simulation, the uncertainties are large and energy-dependent. An estimate of the uncertainty is given in [122], in which samples of Monte Carlo events were produced with the parameters controlling intranuclear rescattering altered within their uncertainties, and the effect on reconstructed shower energy quantified. The uncertainty parametrized as a function of reconstructed energy is used in this analysis, as shown in Figure 5.1. This is added in quadrature with the 5.7% single hadron response uncertainty.




Figure 5.1: Energy-dependent shower energy scale uncertainty from hadronic shower modelling, as parametrized for this analysis. The 1σ error band is shown. This uncertainty is added in quadrature with the constant 5.7% calibration uncertainty to give the overall absolute shower energy scale uncertainty.

Relative shower energy scale

An error in the relative shower energy scale occurs if the same shower occurring in the near and far detectors is given different reconstructed energies. The MINOS near and far detectors are calibrated to give the same response, and so relative energy scale uncertainties come from residual differences remaining after the calibration procedure. The total near-to-far relative shower energy systematic has been calculated to be 2.1% [121]. This value is the sum in quadrature of the systematic uncertainties on each stage in the calibration chain, the largest of which is the discrepancy between spill and cosmic muons, which contributes 1.5% to the total. Although the calibration uses cosmic ray muons, the method should produce the same results when run on muons from neutrino events from the NuMI beam, however, differences between the cosmic and beam samples such as the angular distribution of tracks and the μ^+/μ^- ratio lead to the 1.5% discrepancy between detectors.

The second largest contributor to the relative shower energy scale is the spatial

variation of the calibration, which was assessed by dividing the detector volume into separate regions longitudinally and transversely, and calculating the difference in mean calibrated response between the regions. This contributes a further 0.76% uncertainty in the far detector and 0.73% uncertainty in the near detector.

5.1.3 Muon Energy Scale

The reconstructed neutrino energy for each event is the sum of the reconstructed shower energy in the event and the reconstructed track energy, as described in Section 4.5.2. The systematic error on the track energy for muons is 2% [123], applied to both energy from range and energy from curvature. For the momentum from range, this uncertainty derives from uncertainty in the detector density and the Monte Carlo geometry, along with uncertainty in the path length of detector material traversed as calculated by the track fitting algorithm. The curvature uncertainty uses the sample of muons which stop in the detector, and therefore have both range and curvature momentum measurements available, which were compared to find the residual uncertainty [124].

5.1.4 Misidentified components

The published MINOS sterile neutrino search [109] uses two spectra: one selected to be mostly CC, and one selected to be mostly NC. Each spectrum contains some background, and there are systematic uncertainties on the level of these backgrounds. Since these consist of misidentified events, they might be expected to be at the tails of distributions, where the Monte Carlo modelling may be less accurate than in the bulk of the distribution. In any case, an incorrect estimate of the background will lead to an error in the calculated number of signal events, and thus in the fit parameters.

For the analysis presented here, there is not one NC-like and one CC-like

spectrum, but the same issues occur if the minority component in any of the energy spectra is mismodelled. Therefore, the uncertainties on the background components from the MINOS neutral current analysis are used. The systematic error on the NC component in the CC spectrum is applied to spectra with $\text{PID} < 0.675$, and the error on the CC component in the NC spectrum is applied to spectra with $\text{PID} > 0.675$, as described in Section 3.3.3.

NC Background

The error on the NC component in the CC spectrum (the “NC background”) is estimated using charged current events with the muon removed, producing a sample of NC-like events [125], as in Section 5.2 of this thesis. To estimate the NC background uncertainty, the charged current selection is applied to the events with their muon removed in both data and MC and the ratio of those samples is taken as the uncertainty, which is found to be 25%.

CC Background

The CC background uncertainty was found by considering data from alternate configurations of the NuMI beam, produced by varying the target position and/or horn current [126]. In these alternate configurations, the NC/CC ratio can be very different to the ratio in the normal “low energy” (LE) configuration, allowing the components to be calculated separately as follows: the NC and CC components in the LE and alternate beam in each bin of reconstructed shower energy are expressed as

$$N^{\text{LE}} = N_{\text{CC}}^{\text{LE}} + N_{\text{NC}}^{\text{LE}} \quad (5.3)$$

$$N^{\text{alt}} = N_{\text{CC}}^{\text{alt}} + N_{\text{NC}}^{\text{alt}} \quad (5.4)$$

where N^{LE} and N^{alt} are the numbers of NC-selected data events in the LE and alternate configurations, and N_{CC}^x and N_{NC}^x are the (unknown) numbers of truly CC and NC events respectively in beam configuration x . The number of unknowns can be reduced from four to two by re-expressing Equation 5.4 as

$$N^{\text{alt}} = r_{\text{CC}} N_{\text{CC}}^{\text{LE}} + r_{\text{NC}} N_{\text{NC}}^{\text{LE}} \quad (5.5)$$

where $r_{\text{CC}} = N_{\text{CC}}^{\text{alt}}/N_{\text{CC}}^{\text{LE}}$ and similarly for NC. The values of the ratios r are taken from the Monte Carlo. The two equations in two unknowns can now be solved, giving a meaningful answer under the assumption that the r values are better modelled than the NC and CC components themselves. Cross section uncertainties can be seen to cancel in the r values, as well as selection efficiencies. Additionally, contributions to the beam flux which are common to both beam configurations will cancel: this includes hadron production at the target, but does not include uncertainties in the beam focussing and positioning of the beamline components.

The solution of the simultaneous Equations 5.3 and 5.5 gives values for the truly NC and CC components $N_{\text{NC}}^{\text{LE}}$ and $N_{\text{CC}}^{\text{LE}}$ which can be used to correct the Monte Carlo prediction and the correction propagated to the far detector. This procedure is used in the MINOS ν_e appearance analysis [99], but for this analysis, the CC component from the Monte Carlo is used in the far detector, with the ratio of corrected and uncorrected CC component taken as an uncertainty.

Figures 5.2 and 5.3 show the result of running this method on the near detector data. For the alternate beam, the “horn-off” configuration, in which the focussing horns are turned off entirely, is used. This results in a much lower fraction of truly CC events in the NC-selected spectrum than in the “low-energy” spectrum. The horn-off spectrum also shows a better data/MC agreement than the low-energy spectrum, suggesting that the disagreement in the low-energy spectrum is due to the CC component.




Figure 5.2: NC-selected spectra used in the estimation of the CC background uncertainty. The spectrum on the left is from the normal low-energy running mode of the NuMI beam, and that on the right is from the configuration with the focussing horns turned off, giving a much smaller CC component. In both plots, red is the total MC, blue is the truly CC component and black is the data.




Figure 5.3: Corrections to the NC (left) and CC (right) components of the low-energy beam configuration derived from the method of alternate beams described in the text. The red line is a fit to a constant. The regions with large error bars either have low statistics or have $r_{CC} \approx r_{NC}$, which leads to a degeneracy in the equations.

Figure 5.3 shows that this method gives a 15% reduction in the true CC component in the NC-like region, which is taken as the systematic uncertainty. This is similar to the level of data/MC disagreement seen in the selection variables for the ANN (Figure 3.14) and the output of the ANN (Figure 3.15) and so is consistent with the disagreement stemming almost entirely from mismodelling of CC events which appear NC-like.

5.1.5 Uncertainties due to preselections

The preselection steps used to remove non-neutrino backgrounds and poorly reconstructed events differ between the detectors, and so any problem with the modelling of variables near the positions of cuts used by the preselections will not cancel in the near/far ratio.

In the near detector, the preselection uncertainty has been estimated by investigating the change in the spectrum under changes in the cut positions [104]. The cuts were moved until the fraction of events cut in data and Monte Carlo agreed, and the ratio of the NC-selected spectrum with the modified cuts to the spectrum at nominal cut positions was taken as the systematic.

The resulting uncertainty is large at low energies, reducing to zero above 2.5 GeV, and is shown in Figure 5.4. This uncertainty applies only to NC-like events as defined by $\text{PID} < 0.675$.

The uncertainty on the far detector preselection is divided into two components: one relating to uncertainties on the cuts intended to remove noise; and one for the cuts aimed at removing cosmics [105]. The uncertainties on the NC-like spectrum are determined by shifting the cut positions up and down by a value based on the RMS of the distribution of MC events in each cut variable. The ratio of reconstructed energy spectra with the shifted cut and with the nominal cut is taken as the systematic uncertainty.




Figure 5.4: 1σ uncertainty on NC-like events from near detector preselection.

As in the near detector, the uncertainties are energy-dependent and largest at low energies, where neutral current events deposit only a few hits in the detector and are more difficult to distinguish from non-neutrino backgrounds. The noise and cosmic systematic uncertainties are shown in Figure 5.5.

5.1.6 Cross section uncertainties

Cross section uncertainties are dealt with via tunable parameters available in the NEUGEN event generator used in the MINOS Monte Carlo. The NEUGEN code includes (among others) the three main classes of processes contributing to event rates in MINOS: quasielastic scattering, resonance production, and deep inelastic scattering.

In quasielastic scattering, a neutrino interacts with a single nucleon in the nucleus and produces a lepton. In MINOS, this means $\nu_\mu + n \rightarrow \mu^- + p$. Quasielastic scattering dominates at low energies ($E_\nu < 2$ GeV), and accounts for about one-third of the event rate in the region $2 < E_\nu < 4$ GeV.

At a few GeV, resonant production of excited states forms a major component of the cross section. The main contributor is the $\Delta(1232)$ which decays to a




Figure 5.5: 1σ uncertainty on NC-like events from far detector preselection. The top plot is the uncertainty due to the noise-removing cuts, while the lower plot is the uncertainty due to the cuts removing cosmic ray muons.

pion-nucleon final state.

At higher energies, deep inelastic scattering (DIS) takes over, in which the neutrino interacts with an individual parton, breaking up the nucleus and producing high multiplicity final states.

The NEUGEN code provides a large number of parameters to alter the cross sections of these processes, but three are investigated as likely to affect this analysis:

- The CC axial mass M_A is an effective parameter in the dipole form factor of the nucleon that contributes to one of the terms in the quasielastic cross section. It has an uncertainty of 15% [127]. A similar effective parameter M_A^{res} occurs in the resonance cross section, and is varied along with the quasielastic M_A within its 15% uncertainty.
- In the transition region between resonance and DIS, NEUGEN provides parameters to tune the cross section based on the interaction type (CC or NC), the struck nucleon (neutron or proton) and the final state multiplicity. The parameters σ_n^{CC} and σ_p^{CC} vary the cross section for CC interactions with neutrons and protons respectively which produce two hadrons in the final state. These two parameters are varied together within their 33% uncertainties.
- The corresponding parameters for neutral current events, σ_n^{NC} and σ_p^{NC} , are also varied together within their 33% errors.

5.1.7 Flux uncertainties

The MINOS flux tuning procedure fits the hadron production at the target to the near detector energy spectrum, using data from a variety of beam configurations [128]. The result is a better prediction of the NuMI neutrino flux (in the assumption that Monte Carlo cross sections are correct), which can be applied to the far detector prediction.




Figure 5.6: Uncertainty on the NuMI neutrino flux as a function of true neutrino energy.

The fit has some residual uncertainty coming from the hadron production spectrum itself as well as the functioning and composition of the components of the NuMI beamline [98]. There is also a component from POT counting, which will cancel in the Far/Near prediction. The size of the total flux uncertainty as a function of energy is shown in Figure 5.6.

5.1.8 ANN shape uncertainty

In contrast to an analysis which divides events into two samples according to their ANN output value, this analysis relies more strongly on an accurate modelling of the shape of the ANN output distribution. To investigate the dependence on this shape, a shifted fake data set is produced using events with their ANN output value altered. The size of the alteration is chosen such that normalized near detector data and Monte Carlo agree exactly. At face value, this involves some double counting, since other systematic uncertainties that have been considered certainly account for some of the disagreement. Nonetheless, if the effect of the shape uncertainty calculated in this way is small, the double-counting is benign.




Figure 5.7: Toy demonstration of ANN cumulative distributions, showing how the shift value $S(x)$ is calculated. Red is toy MC and black is toy data. For the case shown, $x = 1$ and $x' = 0.7$.

Calculating ANN shape shift

To determine the size of the fake data shift to be used to estimate the effect of the ANN shape uncertainty, the difference between near detector data and Monte Carlo is used. This technique produces a shift in ANN output value whose size reflects the mismodellings in the Monte Carlo. Technically, this is implemented by making the normalized cumulative distributions of the data and Monte Carlo match at each ANN output value. At each value of the ANN output x , the value of the cumulative distribution $C_{MC}(x)$ in MC is found. The value x' at which the *data* cumulative distribution $C_{data}(x') = C_{MC}(x)$ is found to give the distance $S(x)$ to shift the Monte Carlo event:

$$S(x) = x' - x. \quad (5.6)$$

Figure 5.7 shows a toy version of this, where the difference between data and Monte Carlo has been increased for display purposes.

In deriving the real shifts, events with a PID value between 0.2 and 1.1 are




Figure 5.8: Normalized cumulative ANN output distributions in data and Monte Carlo in the near detector. Data is in black and Monte Carlo is in red.

used: these are the events for which the PID value corresponds with the ANN output value, and which are thus relevant for the ANN shape systematic. The beam flux weights from the hadron production tuning are applied, as is done in the oscillation fit. The normalized cumulative distributions in near detector data and Monte Carlo are shown in Figure 5.8 and the ANN shift values in Figure 5.9. In calculating the shift in best-fit oscillation parameters, both near and far detector fake data events have their ANN output value shifted by the value shown in Figure 5.9, since both near and far detectors use the same ANN.

5.2 Constraining the near/far reconstruction uncertainty with MRCC events

The muon-removed charged current (MRCC) analysis technique begins with the observation that ν_μ charged current events with the muon removed contain only a hadronic shower, and are thus very similar to neutral current events. Since such events are available in both data and Monte Carlo, the analysis strategy is to




Figure 5.9: Shift in ANN output value as a function of original ANN value. The function is constrained to zero at ANN values of 0.2 and 1.1 by construction.

produce a sample of NC-like events which are independent from the actual NC spectrum and can be used to provide uncertainties on, or corrections to, the NC spectrum. This technique was previously used in the MINOS ν_e analysis as a method of correcting the NC component of the ν_e -selected spectrum, which forms a significant part of the background [129]. It was also used to constrain the NC background in the CC spectrum, as described in Section 5.1.4.

5.2.1 Muon Removal

The muon removal is implemented as a second reconstruction pass. Fully reconstructed events are taken, and those with a track have the track hits removed. Hits shared between a track and a shower have a pulse height equivalent to one minimum-ionizing particle (1 mip) subtracted, but are not removed. This procedure is performed for all events in the beam spill, and the remaining hits in the spill are passed through the full reconstruction chain. This second pass may result in the hits being reconstructed into a track or a shower, or not being reconstructed at all. The reconstruction keeps track of the event from which the track




Figure 5.10: Far detector Monte Carlo event before (left) and after (right) muon removal. In the right-hand plot, the event has been passed through the full reconstruction chain a second time after removal of the hits.

was removed, allowing the re-reconstructed events to be matched with their original event at analysis time. An example event before and after the muon removal is shown in Figure 5.10

As pointed out in [129], the distribution of CC hadronic showers with the muon removed is not expected to be exactly equivalent to the distribution of NC hadronic showers. The distribution of hadronic y is different, resulting in a different shower energy distribution for CC and NC showers in the same beam, and the differences in the charge of the mediating boson in the NC and CC interactions means that the net shower charge differs. These issues must be considered for the ν_e appearance analysis, where the aim is to correct the Monte Carlo NC component based on the MRCC analysis, but they have a smaller influence in the analysis presented here, where the aim is only to quantify an overall level of data/Monte Carlo disagreement. Therefore, provided that the same cuts and selections are applied to both data and Monte Carlo, the similarity between true NC showers and muon-removed CC events need not be exact.

5.2.2 Calculating efficiencies

The muon-removal technique provides a way to evaluate the efficiencies ϵ in Equation 5.2 by counting the fraction of CC events which are still reconstructed after

the track is removed, i.e.,

$$\epsilon = \frac{M}{N} \quad (5.7)$$

where N is the number of CC events used, and M is the number of those reconstructed after track removal. Both the non-muon-removed sample and the corresponding muon removed sample are used. The following cuts are applied to the non-muon-removed sample:

1. The event must have a vertex within the fiducial volume.
2. The event must have a value of the CC k NN-based PID [124] of k NN > 0.3 .
(A k NN is a multivariate classification technique in which each event is compared to a large sample of MC events to find the k nearest neighbours. The PID value for the event is the fraction of nearest neighbours which are of the desired type). This produces a sample which is 98.9% pure CC in both detectors.

Each muon-removed event is associated with its original event in the non-muon-removed sample, based on the number of shared hits [129]. In the muon-removed sample, the original event is required to pass the previous cuts, and additionally the muon-removed event must

1. have a vertex within the fiducial volume,
2. pass the preselection cuts in the relevant detector.

Applying the fiducial cut to the muon-removed sample means that events may migrate out of the sample after muon removal (because the vertex position is recalculated in the second reconstruction pass) even when they are successfully reconstructed. There is no corresponding inward migration, because the same fiducial requirement is applied to the original event associated with the muon-removed event. It would be possible to remove this latter requirement by including

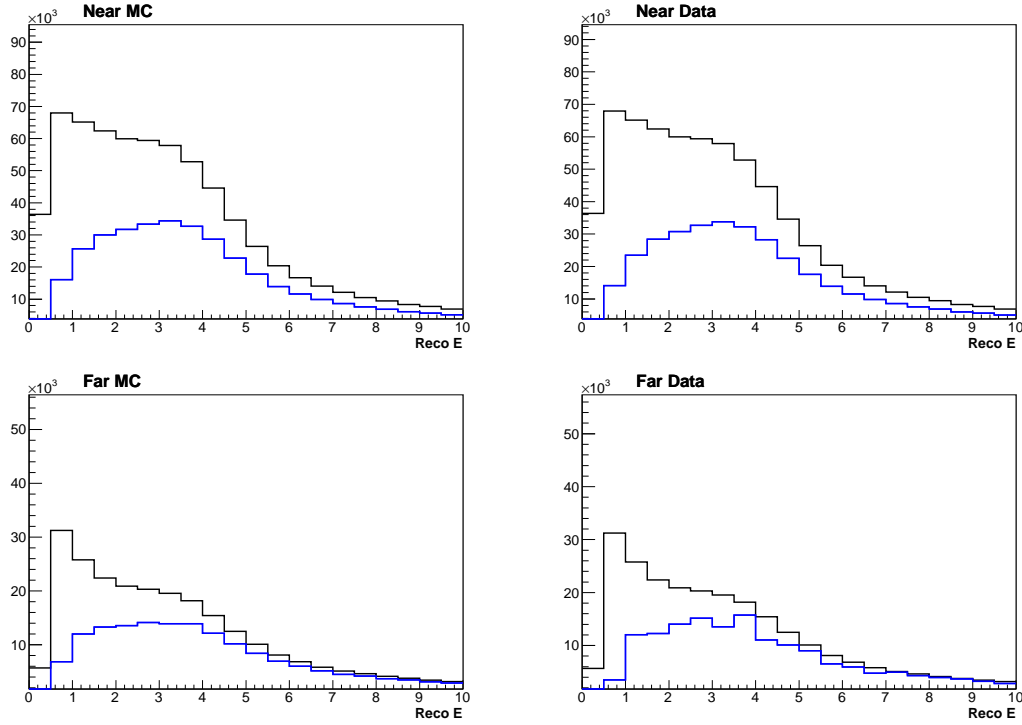


Figure 5.11: Shower energy spectra before and after muon removal for each combination of near/far, data/MC. The black histogram is events in the non-muon-removed sample, and the blue histogram is the events in the muon-removed sample. In each case the cuts described in the text are applied. The energy used for the muon-removed events is the shower energy of the *original* event (i.e., before the muon was removed).

events in the muon-removed sample whose original event was outside the fiducial volume. However, if the inward migration is greater than the outward migration, this could result in a value of ϵ which is greater than 1, and is therefore not so easily interpreted as a reconstruction efficiency.

The fiducial cut results in a rather low value of ϵ in the near detector data and MC, because the fiducial volume is small, and so a large fraction of events are close enough to the fiducial volume boundary to be moved outside it after the muon removal procedure. In the far detector, the fiducial volume is much larger, and so the fraction of events close to the fiducial volume boundary is correspondingly smaller.




Figure 5.12: Efficiencies ϵ_i as a function of shower energy of the original event. The black line is data and the red line is Monte Carlo. An additional energy bin from 10–120 GeV is used but not shown. The efficiency does not reach 1 at high energies because of the fiducial requirement discussed in the text.

The shower energy spectra before and after muon removal are shown in Figure 5.11, and the efficiencies as a function of energy are shown in Figure 5.12. Reasonable agreement is seen between data and Monte Carlo, with efficiencies rising from essentially zero at shower energies below 500 MeV to 0.8–0.9 at high energies. Low energy showers consist of a small number of hits, and so are easily lost when reconstructed a second time, while high energy showers are only lost when re-vertexed to outside the fiducial volume.

With these efficiencies calculated, the method can be applied to both near and far detector data and Monte Carlo, giving the four efficiencies required in Equation 5.2. However, the spectrum of MRCC events is not exactly the same as the actual NC-selected spectrum, and so the efficiency given in Equation 5.7, calculated over all energies, will not be the same as the corresponding efficiency for the NC-selected events. An improved estimate, E , can be found by using the MRCC spectrum to calculate the efficiencies per energy bin, ϵ_i , as above, and taking an average weighted by the fraction of NC-selected events in that energy bin, i.e.,

$$E = \frac{\sum_i \epsilon_i n_i}{\sum_i n_i} \quad (5.8)$$

where n_i is the number of NC-selected events in the energy bin. In the case of far detector Monte Carlo, the n_i are oscillation weighted at the MINOS best-fit CC values, assuming no NC disappearance. There is a choice of energy to use for the event (to decide which bin it appears in for the calculation of ϵ_i): on the second reconstruction pass, the shower energy will not in general be equal to the shower energy of the event before the muon removal. The efficiency in an energy bin is $\epsilon_i = M_i/N_i$, where M_i and N_i are the numbers of events after and before muon removal respectively, so changes in shower energy after the re-reconstruction may lead to movement between energy bins and therefore to values of $\epsilon_i > 1$. To avoid this, the shower energy of the *original* event (i.e., before muon removal) is used in

	Data	MC	Data/MC
Near	0.541	0.554	0.976
Far	0.694	0.697	0.996

Table 5.3: NC spectrum-weighted efficiencies, E , for near and far detector data and Monte Carlo.

calculating both the numerator and the denominator of ϵ_i .

A value of E is calculated for all four combinations near/far, data/MC using the corresponding ϵ_i , but in each case the n_i are taken from Monte Carlo for the relevant detector in order to reduce the total statistical error on the final estimate.

The best estimate of the reconstruction uncertainty is therefore given by the double ratio

$$R' = \frac{E_{\text{data}}^F/E_{\text{MC}}^F}{E_{\text{data}}^N/E_{\text{MC}}^N}. \quad (5.9)$$

5.2.3 Results

This method has been applied to the full dataset used in this analysis, giving the efficiencies shown in Table 5.3. With these numbers, the value of the double ratio is $R' = 1.02 \pm 0.017$, and the central value is taken as a normalization error of 2%. The error on R' is approximate, and derived from the value of ϵ as defined in Equation 5.7 for the far detector, assuming the errors on all other components are zero, and using the method for errors on a binomial parameter defined in [130].

This uncertainty (as part of the near/far normalization uncertainty) is applied across the spectrum in energy and PID, despite deriving from a method that applies only to NC events. This can be justified by noting that a hand-scan analysis of CC events carried out for the MINOS CC analysis resulted in a reconstruction-related uncertainty of 1.3% [131], so at worst the use of the 2% uncertainty across the spectrum is a conservative estimate.

Parameter	Input value
Δm^2	$2.35 \times 10^{-3} \text{ eV}^2/c^4$
f_μ	0.84
f_s	0.2
f_e	0

Table 5.4: Input oscillation parameters for systematics studies

5.3 Systematic shifts

To assess the effect of each systematic uncertainty on the final analysis result, systematically shifted fake data samples are produced in both detectors, for each systematic at $\pm 1\sigma$. The far detector prediction is calculated using the shifted near detector fake data and the Far/Near ratio from unshifted Monte Carlo, and this far detector prediction is fit to the shifted far detector fake data. This fit only includes oscillation parameters and not systematic parameters.

Since high-statistics fake data is being used, a systematic shift that is handled perfectly by the Far/Near prediction method will result in a best fit point that is equivalent to the input point. The larger the effect of the systematic shift, the larger the deviation from the input point. Systematics with a large effect will be included in the fit as constrained nuisance parameters, as described in Section 5.4 below.

The input parameters for these systematic studies are given in Table 5.4. f_μ and f_s are moved away from the physical boundary to allow the best-fit point to vary away from the true point in any direction. f_e is fixed to zero in the fit.

Figure 5.13 shows the size of the shifts in each fit parameter for each systematic, and Table 5.5 shows the same quantities, sorted by the size of the total shift in f_s . In general, systematics which do not cancel between the detectors have the largest effect, with flux and cross section uncertainties (the $\sigma_{n/p}^{\text{NC/CC}}$ parameters) cancelling almost entirely.

The near detector preselection dominates the f_s shift, since the effect of the




Figure 5.13: Shifts in the fit parameters for each of the systematics considered. The lower end of the error bar represents the shift in fit parameter for a -1σ shift in the given systematic, and the upper end is the shift for $+1\sigma$. From top to bottom, the plots show f_s , Δm^2 and f_μ .

Systematic	Total shift		
	$f_s \times 100$	Δm^2 ($10^{-5} \text{ eV}^2/c^4$)	$f_\mu \times 1000$
Near presel.	35.7 (86.2%)	2.0 (4.5%)	2.1 (1.1%)
Far presel. (cosmic)	18.0 (43.3%)	1.5 (3.5%)	6.1 (3.2%)
N/F normalization	15.8 (38.2%)	13.7 (31.6%)	5.1 (2.7%)
CC background	15.1 (36.3%)	1.3 (3.0%)	3.7 (2.0%)
Far presel. (noise)	7.1 (17.0%)	0.3 (0.7%)	1.0 (0.5%)
Shower E scale (abs)	6.4 (15.3%)	7.3 (16.8%)	8.5 (4.5%)
Shower E scale (rel)	5.6 (13.5%)	0.7 (1.5%)	23.5 (12.6%)
M_A QEL, res	4.8 (11.6%)	2.5 (5.9%)	1.9 (1.0%)
$\sigma_{n/p}^{\text{CC}}$	3.1 (7.6%)	0.2 (0.5%)	2.1 (1.1%)
ANN shape	3.1 (7.5%)	3.2 (7.4%)	13.5 (7.2%)
Flux uncertainty	2.6 (6.2%)	5.1 (11.7%)	6.5 (3.5%)
$\sigma_{n/p}^{\text{NC}}$	2.1 (5.0%)	0.6 (1.5%)	3.1 (1.7%)
Muon energy scale	1.4 (3.5%)	6.2 (14.3%)	7.3 (3.9%)
NC background	0.5 (1.2%)	1.8 (4.1%)	14.7 (7.9%)

Table 5.5: The total systematic shift in each fit parameter for each systematic, ordered by the size of the shift in f_s . The total shift is the sum of the absolute values of the $+1\sigma$ and -1σ shifts, and the quantities tabulated are multiplied by a factor to give numbers of order 1 for ease of comparison. The number in brackets is the size of the total systematic shift as a percentage of the 1σ statistical uncertainty on the parameter.

systematic is to reweight the low energy bins in the NC-like part of the spectrum (as shown in Figure 5.4), which is the same as the effect of sterile neutrino mixing. The far detector cosmic preselection systematic and Near/Far normalization systematic have similar shapes, extending to high energies at about the same level (see Figure 5.5), and so produce similar effects on the shift in f_s . The far cosmic preselection uncertainty has little effect on f_μ and Δm^2 because it only produces a change in the NC-like region, while the CC-like region determines f_μ and Δm^2 .

The absolute energy scales (on shower energy and muon energy) have large effects on the fit value of Δm^2 , which is expected since Δm^2 controls the position of the oscillation dip in energy.

5.4 Systematic uncertainties in the fit

The systematics with the largest effect are included in the fit as extra parameters, and constrained by a penalty term, by modifying the χ^2 function (Equation 4.9) to depend on the systematic shifts \mathbf{s} as follows:

$$\chi_P^2(\boldsymbol{\nu}, \mathbf{s}) = 2 \sum_i \left[e_i(\boldsymbol{\nu}, \mathbf{s}) - o_i + o_i \ln \frac{o_i}{e_i(\boldsymbol{\nu}, \mathbf{s})} \right] + \sum_j \frac{(s_j - s_j^0)^2}{\sigma_j^2} \quad (5.10)$$

where s_j is the value of the systematic parameter, s_j^0 is its default value, and σ_j is its 1σ error. The index j runs over all the systematics in the fit. The expected values $e_i(\boldsymbol{\nu}, \mathbf{s})$ are calculated at the given values \mathbf{s} of the systematics. The penalty term represents the fact that the true value of the systematic is unknown, but constrained to be around the nominal value s_j^0 . In producing contours and one-parameter projections, the systematics are always marginalized over, which has the effect of increasing the size of the contour/projection in a way that reflects the size of the systematic uncertainty.

The four systematics with the largest effect on f_s are chosen to be included

in the fit: near detector preselection, far detector cosmic preselection, near/far normalization, and CC background. In addition, the absolute shower energy scale is included, as the second-largest shift in Δm^2 . As shown in Table 5.5, the next largest shift in f_s (FD noise preselection) is half the size of the shift due to CC background. The five systematics chosen alter the spectrum in rather different ways, and so cover the space of possible changes to the spectrum well. They are also generally quite conservative estimates, and so their effect in inflating the quoted error should be broad enough to take account of other systematics which are not included directly.

5.4.1 Producing systematically shifted predictions

The production of systematically shifted predictions uses the same method as in the MINOS neutral current analysis [57], which uses interpolation between systematically shifted templates to produce a prediction at given values of the oscillation parameters $\boldsymbol{\nu}$ and systematic shifts \mathbf{s} . The method is described in more detail in [132].

Briefly, for each systematic in the fit, the systematics interpolation uses three spectra: one at the nominal value of the systematic, and one each at $\pm 1\sigma$. Each spectrum has an associated true energy distribution, and so can be used to produce a prediction at a given value of the oscillation parameters $\boldsymbol{\nu}$, as in Equation 4.6:

$$N_i(\boldsymbol{\nu}) = \sum_j S_{ij} T_j P_j(\boldsymbol{\nu}). \quad (4.6)$$

The systematically-shifted prediction at a value s of the systematic is a weighted sum of the three oscillated predictions:

$$N_i(\boldsymbol{\nu}, s) = w_{-\sigma}(s) N_i^{-\sigma}(\boldsymbol{\nu}) + w_0(s) N_i^0(\boldsymbol{\nu}) + w_{+\sigma}(s) N_i^{+\sigma}(\boldsymbol{\nu}) \quad (5.11)$$




Figure 5.14: Weights $w(s)$ used for systematics interpolation. The red line is the weight $w_{+\sigma}(s)$ given to the $+1\sigma$ spectrum, blue is the weight $w_{-\sigma}(s)$ given to the -1σ spectrum, and black is the weight $w_0(s)$ given to the nominal spectrum.

where the functions $w(s)$ are piecewise linear weighting functions which interpolate between the spectra, shown in Figure 5.14. At values $s = 0, \pm 1\sigma$, the corresponding weighting function has a value of 1 while the others are zero. The generalization to any number of systematics involves storing $\pm 1\sigma$ spectra for each systematic and making the weighting functions depend on all of the systematic shifts, i.e., $w(\mathbf{s})$. This scheme is necessarily an approximation, but has been shown to reproduce fake data systematic shifts without bias, and so the approximation is good enough.

5.4.2 Effect on oscillation parameter limits

With the chosen systematics included in the fit, the sensitivity can be re-evaluated in the same way as in Section 4.5.1, using high-statistics fake data. The fake data has no systematic shifts applied, so the true input point will be returned as the best fit point, but the inclusion of the systematic parameters will increase the size of the contour or projection relative to the fit with only oscillation parameters.

The effect of the five fit systematics on the projections for each oscillation

parameter is shown in Figure 5.15. By construction, the effect of the systematics is largest on f_s , with the near detector preselection uncertainty dominating. Since Δm^2 and f_μ are mostly constrained by the CC-like part of the spectrum, the systematics which affect only the NC-like part of the spectrum have no effect on the contours in those parameters, but the Near/Far normalization uncertainty and absolute shower energy scale make some difference.

The 90% confidence upper limit on f_s without systematics is 0.28, while with the systematics included, the limit is $f_s < 0.415$. As can be seen from Figure 5.15, the effect of the systematics on the Δm^2 and f_μ limits is relatively small.




Figure 5.15: Effect of systematics on $\Delta\chi^2$ projections in the fit parameters. The black line is the statistical-only projection, and each coloured contour shows the effect of adding one more systematic in addition to those used in the previous contour. From top to bottom, the contours in f_s , Δm^2 , and f_μ are shown. The $\Delta\chi^2$ values corresponding to 68% and 90% confidence levels are also shown.

Chapter 6

Results and Discussion

6.1 Fit results

This chapter contains the results of applying the analysis method described in this thesis to the far detector data. In the absence of neutrino oscillations, 2840.75 events are expected in the far detector; 2376 are observed, giving a clear signal of ν_μ disappearance. Figure 6.1 shows the fitted energy spectra in the far detector and Figure 6.2 shows the energy spectrum integrated over PID and the PID spectrum integrated over energy.

The best-fit oscillation and systematic parameters are shown in Table 6.1, along with the 68% C.L. stat.+syst. limits on the oscillation parameters. The fit is performed for f_e fixed at the values of 0 and 0.19, as described in Section 4.7. The χ^2/dof is 657.55/657 for $f_e = 0$ and 657.54/657 for $f_e = 0.19$, showing that spectra both with and without ν_e appearance are equally compatible with the data. In both cases, the best-fit value of f_s is zero, i.e., no sterile neutrino mixing, consistent with the standard three-flavour picture of neutrino oscillations. The upper limit on f_s depends on the presence or absence of ν_e appearance, since the far detector prediction with ν_e appearance is higher in the NC-like region than the

	$f_e = 0$	$f_e = 0.19$
f_s	$0.00^{+0.06}$	$0.00^{+0.23}$
$\Delta m^2 \quad (10^{-3} \text{ eV}^2/c^4)$	$2.37^{+0.13}_{-0.09}$	$2.33^{+0.11}_{-0.09}$
f_μ	$1.00_{-0.05}$	$1.00_{-0.05}$
Near presel.	-0.76	-0.61
Far presel. (cosmic)	0.61	-0.37
N/F normalization	0.88	0.00
CC background	-1.05	0.06
Shower E scale (abs)	-0.49	-0.57

Table 6.1: Best-fit values of the oscillation and systematic parameters with $f_e = 0$ and $f_e = 0.19$. The 68% C.L. limits (stat.+syst.) on the oscillation parameters are given, and the systematic best-fits are given as a fraction of the 1σ uncertainty on the parameter.

prediction without ν_e appearance¹. This means that the fit can tolerate a larger reduction in the NC-like region due to a larger value of f_s when $f_e > 0$.

At 90% C.L., the upper limit on f_s is 0.15 (0.41) with $f_e = 0$ (0.19).

Figure 6.3 shows the $\Delta\chi^2$ projections from the data fit for each of the oscillation parameters after marginalizing over the other oscillation and systematic parameters. The projection in Δm^2 shows that although the effect of ν_e appearance on the best-fit value is noticeable, there is relatively little effect on the shape or width of the contour. The lower left plot in Figure 6.2 shows that electron neutrinos appear in the spectrum at the same energies as the ν_μ that are disappearing². Since it is the position in energy of the disappearing ν_μ that constrains Δm^2 , it is therefore unsurprising that the appearance of ν_e can alter the best-fit Δm^2 . (This feature is irrelevant for the MINOS CC analysis, since it only uses a ν_μ CC selected spectrum which does not contain any ν_e CC events).

The parameters Δm^2 and f_μ are constrained by the CC-like region of the spec-

¹The total neutrino flux at the far detector is the same in each case, but with $f_e = 0$ there are more ν_τ , which are below threshold and do not produce CC interactions.

²The appearance ν_e come from oscillations of the ν_μ , so it is obvious that they must have the same *true* energy, but reconstruction efficiencies and differences in energy reconstruction mean that they need not be most significant in the *reconstructed* spectrum at the peak ν_μ disappearance energy.

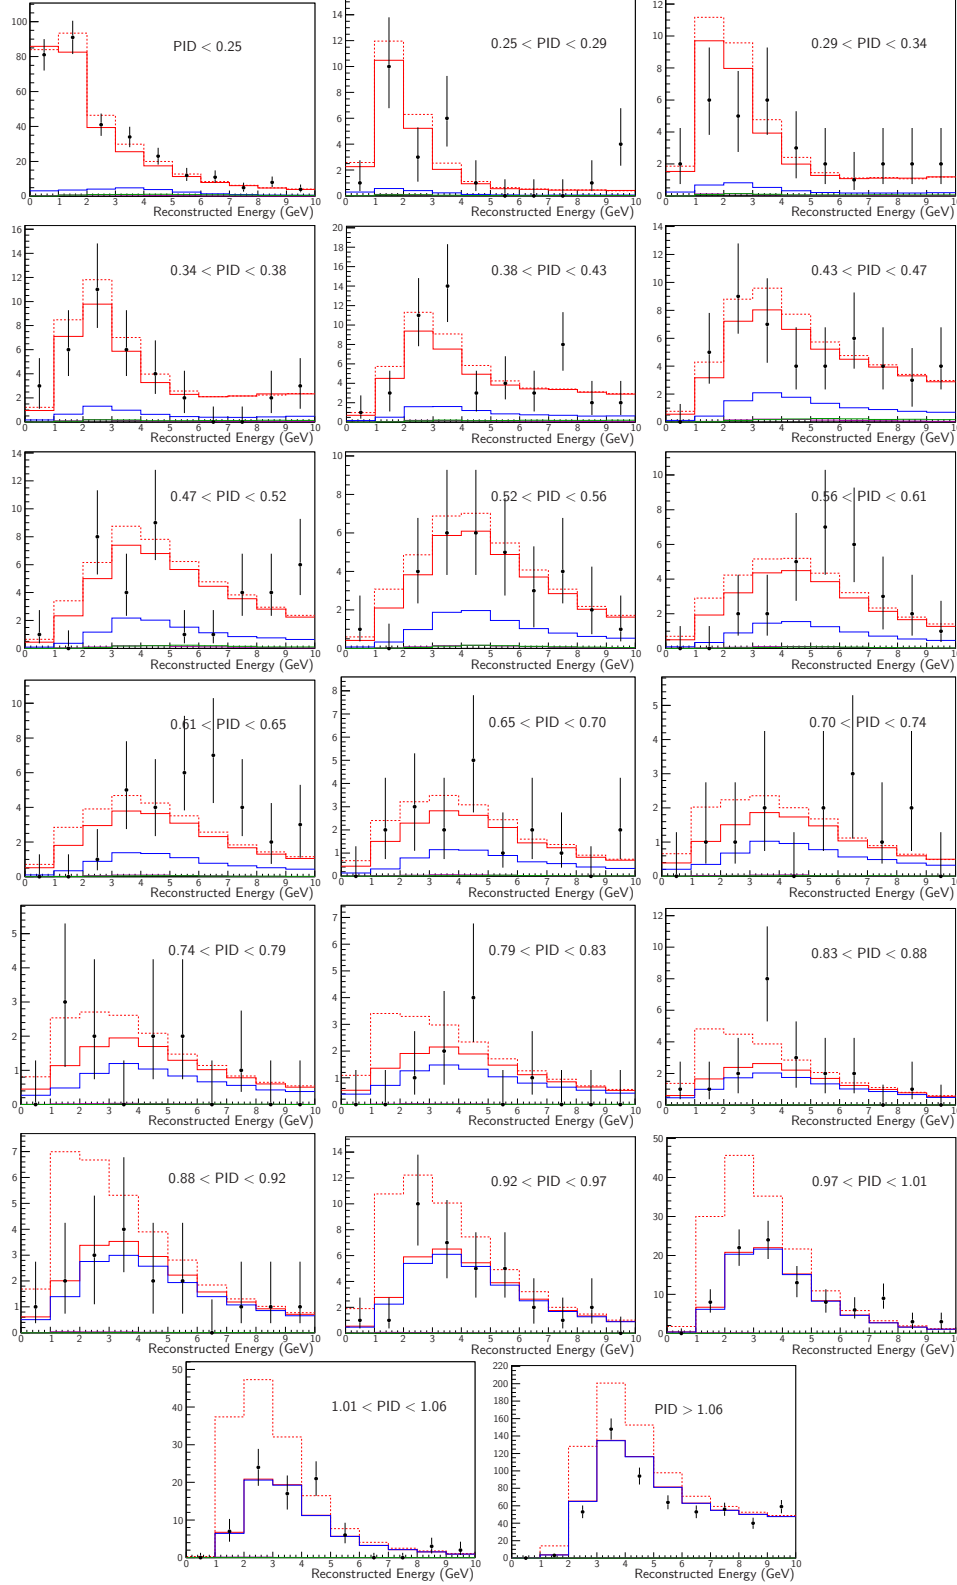


Figure 6.1: Energy projections of far detector data spectrum for each bin in PID. The black points are the data, the red dashed histogram is the prediction in the absence of oscillations, and the other histograms are the prediction at the best-fit point (including systematics) for $f_e = 0$: red is the total prediction, blue is ν_μ CC, green is ν_e CC and the small purple component is ν_τ CC.

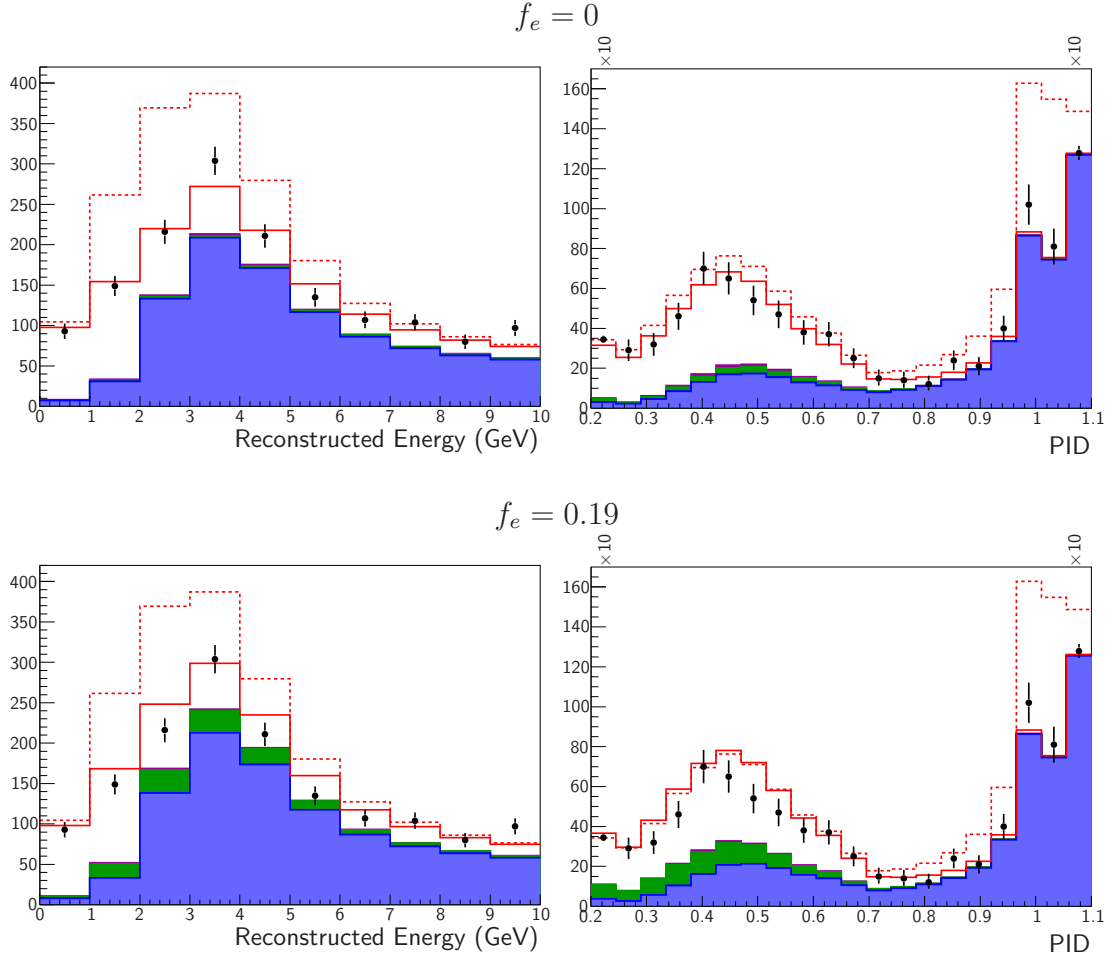


Figure 6.2: Far detector data spectra projected onto energy axis (left) and PID axis (right). The upper plots are for $f_e = 0$ and the lower plots are for $f_e = 0.19$. In each case the spectrum is integrated over the variable which is not shown. The black points are the data, the red dashed histogram is the prediction in the absence of oscillations, and the other histograms are the prediction at the best-fit point (including systematics): red is the total prediction, blue is ν_μ CC, green is ν_e CC and the small purple component is ν_τ CC. In the PID spectra, the end bins are reduced by a factor of 10 for display purposes.




Figure 6.3: $\Delta\chi^2$ projections in the oscillation parameters from the fit to the far detector data. In each case, the black line is from the fit with $f_e = 0$ and the blue line with $f_e = 0.19$. From top to bottom, the contours in f_s , Δm^2 , and f_μ are shown. The $\Delta\chi^2$ values corresponding to 68% and 90% confidence levels are also shown.

trum and so can be compared to the results obtained in the two-flavour framework of the MINOS charged current analysis, which are $\Delta m^2 = 2.35_{-0.08}^{+0.11} \times 10^{-3} \text{ eV}^2/\text{c}^4$ at 68% C.L. and $\sin^2 2\theta > 0.91$ at 90% C.L [57]. The best-fit values of Δm^2 obtained in this thesis, both with and without ν_e appearance, are consistent with the MINOS CC value, although the upper limit is slightly larger for $f_e = 0$. This is unsurprising, since the MINOS CC analysis was optimized for measuring this parameter. By comparing Equations 1.17 and 1.38, we see that $f_\mu = \sin^2 2\theta$, and so the results of this thesis and the MINOS CC analysis agree very well in this parameter.

As shown in Table 6.1, several of the systematic pulls are rather large: -1.05σ for the CC background and $+0.88\sigma$ for near/far normalization ($f_e = 0$), for example. Additionally, the pulls differ for the $f_e = 0$ and $f_e = 0.19$ cases. It is therefore worth checking the robustness of the result against the presence of systematic parameters. A fit to the data with only the oscillation parameters included in the fit gives a total $\chi^2/\text{dof} = 665/657$, and best-fit values $f_s = 0$, $\Delta m^2 = 2.29 \times 10^{-3} \text{ eV}^2/\text{c}^4$ and $f_\mu = 1$. Although Δm^2 is pulled by a large amount (nearly 1σ) relative to the stat.+syst. fit, the values of f_s and f_μ remain unchanged. Thus the conclusions on the absence of sterile neutrino mixing are robust, even though this analysis may not provide a truly reliable value for Δm^2 .

6.1.1 Fitting for f_e

Another feature of the $\Delta\chi^2$ projections shown in Figure 6.3 is that the shape of the projection in f_s for the $f_e = 0$ case is close to a straight line, rather than the parabolic shape which is usual, whereas the projection for $f_e = 0.19$ is much more parabolic with derivative zero (i.e., a true minimum) at $f_s = 0$. This suggests that the fit with $f_e = 0$ would prefer a negative value of f_s and therefore that the data do not come from a true distribution with $f_e = 0$. To be more quantitative about




Figure 6.4: $\Delta\chi^2$ projection of f_e when it is included in the fit to the far detector data. The $\Delta\chi^2$ values corresponding to 68% and 90% confidence levels are also shown.

this, a fit was performed with f_e as a free parameter and the $\Delta\chi^2$ projection in f_e produced, which is shown in Figure 6.4. The best-fit value and 68% C.L. limits are $f_e = 0.087^{+0.049}_{-0.045}$, with the other parameters as shown in Table 6.2. At $f_e = 0$, $\Delta\chi^2 = 4$, so there is a suggestion at the 2σ level for ν_e appearance.

Although they are suggestive, any numbers relating to ν_e appearance from this analysis should be taken circumspectly: the analysis has been tuned to consider neutral current events and not electron neutrino charged current events, so any systematics affecting CC ν_e and not NC have not been included (for example, ν_e reconstruction efficiency). Secondly, the fit for f_e has only been performed after examining the results from the far detector data, and so cannot be considered unbiased. The $\Delta\chi^2$ projection for f_s when f_e is fit is shown in Figure 6.5; its shape is still not parabolic. Nonetheless, the MINOS ν_e appearance analysis observes a small excess (0.7σ) of ν_e -like events over the expected background, and the analysis in this thesis uses the same data sample, so the presence of events which appear ν_e -like in one analysis implies the presence of such events in the other.

Parameter	Best-fit value
f_e	$0.09^{+0.05}_{-0.04}$
f_s	$0.00^{+0.11}$
Δm^2 ($10^{-3} \text{ eV}^2/c^4$)	$2.37^{+0.12}_{-0.09}$
f_μ	$1.00_{-0.05}$
Near presel.	-0.69
Far presel. (cosmic)	0.12
N/F normalization	0.56
CC background	-0.36
Shower E scale (abs)	-0.56

Table 6.2: Best-fit values of the oscillation and systematic parameters when f_e is included in the fit. 68% confidence level errors on the oscillation parameters are also shown. The systematic best-fits are given as a fraction of the 1σ uncertainty on the parameter.




Figure 6.5: $\Delta\chi^2$ projection of f_s when f_e is included in the fit to the far detector data. The $\Delta\chi^2$ values corresponding to 68% and 90% confidence levels are also shown.

6.2 Future Directions

The MINOS experiment is nearing the end of its life, and so the future increase in statistics is likely to be small. Therefore improvements in the limits on sterile neutrino mixing will have to come from analysis improvements. The statistical sensitivity to sterile mixing had only a moderate improvement from the PID-binning method used in this thesis (Figure 4.9), suggesting that there is little more information to be gleaned from individual events to improve the statistical sensitivity further. On the other hand, the effect of systematics on the size of the f_s limit is significant (Figure 5.15), so reductions in the systematic uncertainties would lead to improved limits on f_s .

The near detector preselection produces the largest increase in the size of the f_s limit, and is therefore the most obvious target for reduction in a future analysis. Since the preselection aims mostly to remove poorly reconstructed events, improvements in the reconstruction to reduce the number of such events produced in the first place (rather than just removing the events in the analysis) would be a possibility. Alternatively (or additionally), a more complete understanding of the near/far differences induced by the presence of poorly reconstructed events in the near detector could lead to a reduced assessment of this systematic. To this end, it has been suggested that the use of the same preselection cuts in near and far detectors may reduce the systematic uncertainties [104].

Another large contribution to the systematic uncertainty on f_s is the CC component in the NC-like part of the spectrum. The alternate beam configuration method used to assess the uncertainty on this component (Section 5.1.4) can be used to provide a *correction* to it, rather than just an uncertainty, which would lead both to a more accurate best fit value and to a smaller uncertainty on f_s . The same correction could also be found using the muon removed charged current (MRCC) events described in Section 5.2, as follows: the total number of NC-

like events in data is D , consisting of a truly NC component N and a truly CC component C , so that

$$D = N + C. \quad (6.1)$$

N can be found from MRCC events as $N = rN_{\text{MC}}$ where r is the data/MC ratio for the pseudo-NC sample provided by MRCC events. With N calculated, the truly CC component C can be inferred from the constraint in Equation 6.1. This method would leave a residual uncertainty due to the differences between real NC events and the pseudo-NC events produced by the MRCC method, which would require a more complete study than the one used in this thesis, where the MRCC method was used only to constrain an uncertainty, not provide a correction.

6.3 Summary

This thesis has presented a novel method of predicting and fitting the far detector energy spectrum in MINOS by binning in both energy and a parameter distinguishing NC and CC events. The method was shown to give slightly improved statistical sensitivity to sterile neutrino oscillations. The effect of systematic uncertainties on the method was examined, and the largest contributions to the uncertainty on f_s were included in the fit. The near/far normalization uncertainty was reduced by almost a factor of 2 by reassessing each of its components, with muon removed charged current events used to constrain the reconstruction efficiency component.

Fitting the far detector data using this analysis gave a result consistent with the standard three-flavour picture of neutrino oscillations, with no evidence for mixing to a sterile species, and the fraction of ν_μ oscillating to ν_s , f_s , limited at 90% C.L. to 0.15 (0.41) in the absence (presence) of $\nu_\mu \rightarrow \nu_e$ oscillations.

Appendix A

Mock Data Plots

This appendix presents additional plots from the mock data studies described in Section 4.6.

Figures A.1 and A.2 show a few 90% confidence level contours in f_s and Δm^2 from individual subruns in set 1 and 2 respectively, demonstrating how much the limits from different experiments with the same true value may change.

Finally, Figures A.3 and A.4 show example spectra integrated over energy and PID separately, demonstrating the level of statistical fluctuation in the data at 7.1×10^{20} protons-on-target.




Figure A.1: 90% confidence level contours in $(f_s, \Delta m^2)$ for three randomly chosen subruns in mock data set 1. The 90% C.L. sensitivity is shown in black, with the star showing the true input point. The red, green and blue contours and points are the individual subruns and their best-fit points.




Figure A.2: 90% confidence level contours in $(f_s, \Delta m^2)$ for three randomly chosen subruns in mock data set 2. The 90% C.L. sensitivity is shown in black, with the star showing the true input point. The red, green and blue contours and points are the individual subruns and their best-fit points.

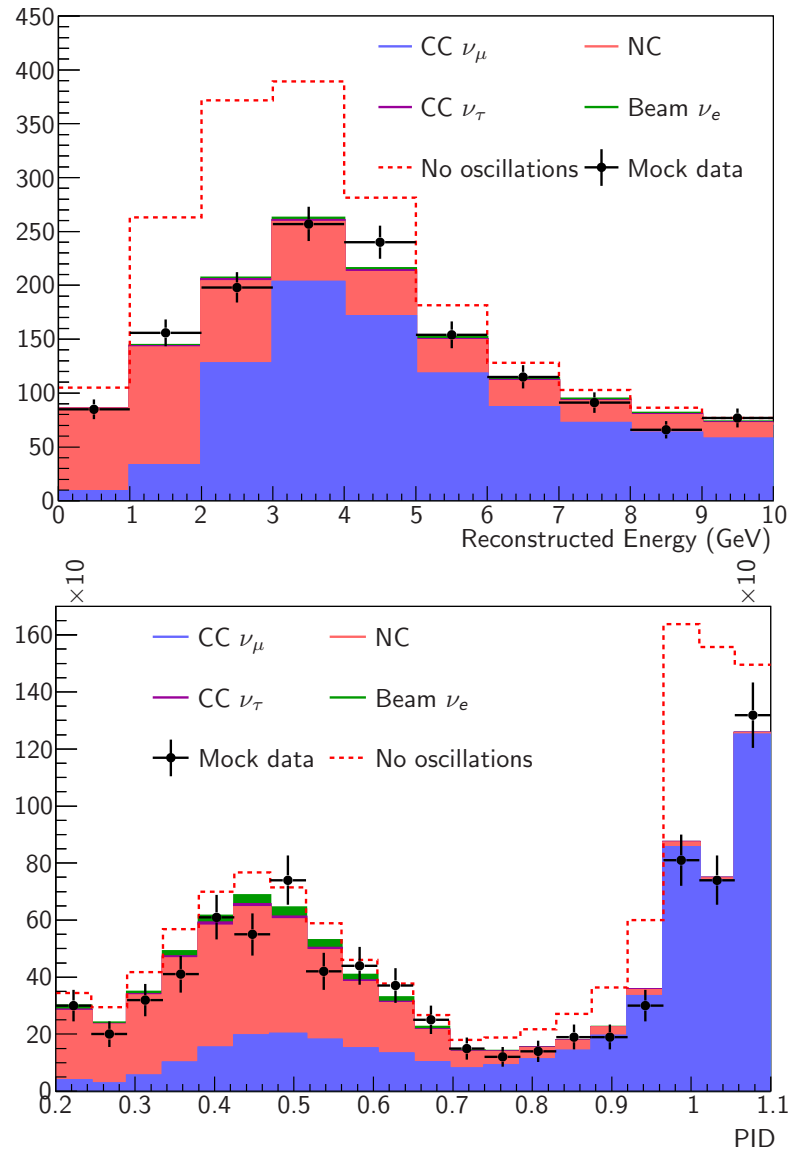


Figure A.3: Example spectra from one subrun in mock data set 1. This is the subrun whose $(f_s, \Delta m^2)$ contour is shown in red in Figure A.1. The upper plot is the energy spectrum from 0–10 GeV integrated over PID, and the lower plot is the PID spectrum integrated over energy. The components oscillated at the best fit values are shown stacked and filled, with the prediction for the total in the absence of oscillations as a dashed red line. The black points with errors are the mock data. In the lower plot, the first and last bins have been reduced by a factor of 10 for display purposes.

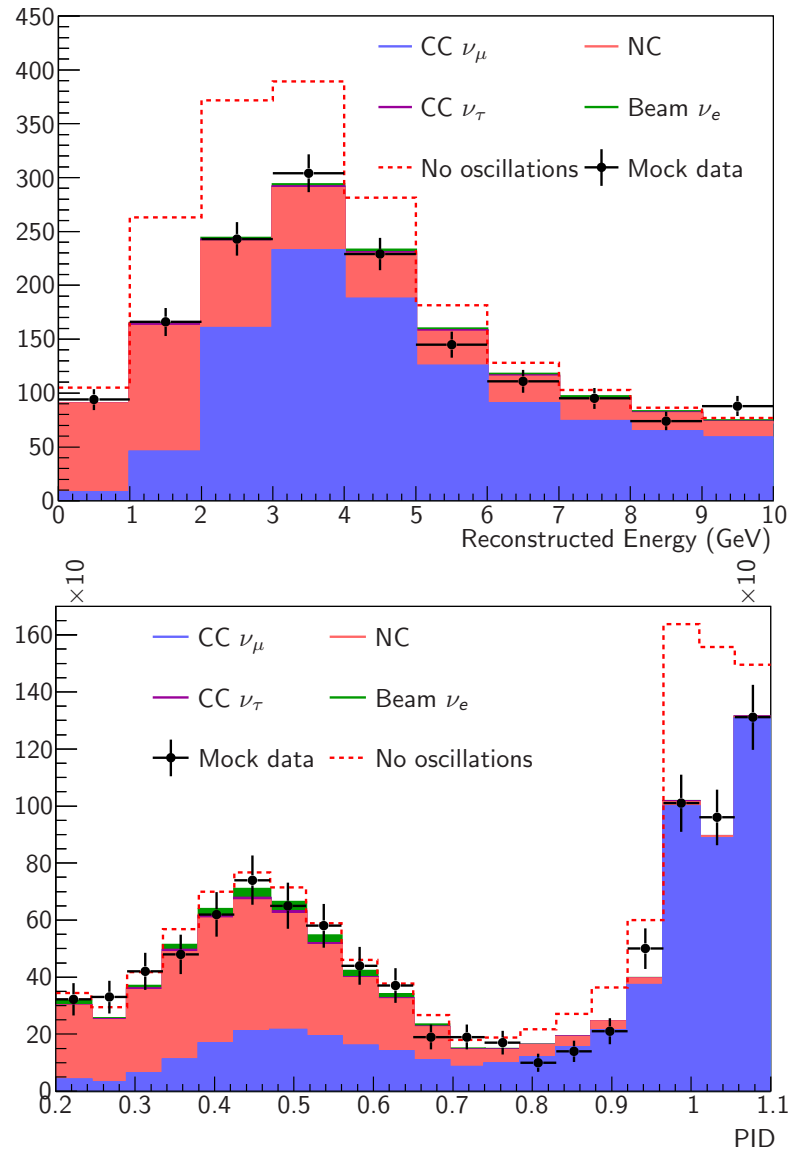


Figure A.4: Example spectra from one subrun in mock data set 2. This is the subrun whose $(f_s, \Delta m^2)$ contour is shown in red in Figure A.2. The upper plot is the energy spectrum from 0–10 GeV integrated over PID, and the lower plot is the PID spectrum integrated over energy. The components oscillated at the best fit values are shown stacked and filled, with the prediction for the total in the absence of oscillations as a dashed red line. The black points with errors are the mock data. In the lower plot, the first and last bins have been reduced by a factor of 10 for display purposes.

Bibliography

- [1] W. Pauli, Letter, addressed to participants of the Tübingen conference on radioactivity 1930, Available from the CERN Document server: <http://cdsweb.cern.ch/record/83282P>
- [2] J. Chadwick, The Existence of a Neutron, Proceedings of the Royal Society of London **A136**, 692 (1932)
- [3] E. Fermi, Versuch einer Theorie der β -Strahlen. I., Zeitschrift für Physik **88**, 161 (1934)
- [4] H. Bethe and R. Peierls, The Neutrino, Nature **133**, 689 (1934)
- [5] F. Reines and C. Cowan, Detection of the Free Neutrino, Phys. Rev. **92**, 830 (1953)
- [6] C. L. Cowan, F. Reines, F. B. Harrison, H. W. Kruse, and A. D. McGuire, Detection of the Free Neutrino: A Confirmation, Science **124**, 103 (1956)
- [7] R. Davis, Attempt to Detect the Antineutrinos from a Nuclear Reactor by the $\text{Cl}^{37}(\bar{\nu}, e^-)\text{A}^{37}$ Reaction, Phys. Rev. **97**, 766 (1955)
- [8] M. Goldhaber, L. Grodzins, and A. W. Sunyar, Helicity of Neutrinos, Phys. Rev. **109**, 1015 (1958)
- [9] G. Danby et al., Observation of High-Energy Neutrino Reactions and the Existence of Two Kinds of Neutrinos, Phys. Rev. Lett. **9**, 36 (1962)

- [10] W. Heisenberg, Zur Theorie der “Schauer” in der Höhenstrahlung, Zeitschrift für Physik **101**, 533 (1936)
- [11] H. Yukawa, On the Interaction of Elementary Particles, Proc. Phys. Math. Soc. Japan **17**, 48 (1935)
- [12] S. L. Glashow, Partial Symmetries of Weak Interactions, Nucl. Phys. **22**, 579 (1961)
- [13] A. Salam and J. C. Ward, Electromagnetic and weak interactions, Phys. Lett. **13**, 168 (1964)
- [14] S. Weinberg, A Model of Leptons, Phys. Rev. Lett. **19**, 1264 (1967)
- [15] Gargamelle Neutrino Collaboration: F. J. Hasert et al., Search for elastic muon neutrino electron scattering, Phys. Lett. **B46**, 121 (1973)
- [16] Gargamelle Neutrino Collaboration: F. J. Hasert et al., Observation of neutrino-like interactions without muon or electron in the Gargamelle neutrino experiment, Phys. Lett. **B46**, 138 (1973)
- [17] UA1 Collaboration: G. Arnison et al., Experimental Observation of Isolated Large Transverse Energy Electrons With Associated Missing Energy at $\sqrt{s} = 540$ GeV, Phys. Lett. **B122**, 103 (1983)
- [18] UA2 Collaboration: R. Banner et al., Observation of Single Isolated Electrons of High Transverse Momentum in Events with Missing Transverse Energy at the CERN $\bar{p}p$ Collider, Phys. Lett. **B122**, 476 (1983)
- [19] UA1 Collaboration: G. Arnison et al., Experimental observation of lepton pairs of invariant mass around $95 \text{ GeV}/c^2$ at the CERN SPS collider, Phys. Lett. **B126**, 398 (1983)

- [20] UA2 Collaboration: P. Bagnaia et al., Evidence for $Z^0 \rightarrow e^+e^-$ at the CERN anti-p p collider, Phys. Lett. **B129**, 130 (1983)
- [21] T. D. Lee and C. N. Yang, Question of Parity Conservation in Weak Interactions, Phys. Rev. **104**, 254 (1956)
- [22] E. C. G. Sudarshan and R. E. Marshak, Chirality Invariance and the Universal Fermi Interaction, Phys. Rev. **109**, 1860 (1958)
- [23] R. P. Feynman and M. Gell-Mann, Theory of the Fermi Interaction, Phys. Rev. **109**, 193 (1958)
- [24] J. Christensen, J. Cronin, V. Fitch, and R. Turlay, Evidence for the 2π Decay of the K_2^0 Meson, Phys. Rev. Lett. **13**, 138 (1964)
- [25] G. Lüders, Proof of the TCP Theorem, Annals of Physics **2**, 1 (1957)
- [26] S. Antusch and E. Fernandez-Martinez, Signals of \mathcal{CPT} Violation and Non-Locality in Future Neutrino Oscillation Experiments, Phys. Lett. **B665**, 190 (2008)
- [27] E. J. Konopinski and H. M. Mahmoud, The Universal Fermi Interaction, Phys. Rev. **92**, 1045 (1953)
- [28] M. L. Perl et al., Evidence for Anomalous Lepton Production in $e^+ - e^-$ Annihilation, Phys. Rev. Lett. **35**, 1489 (1975)
- [29] Precision electroweak measurements on the Z resonance, Physics Reports **427**, 257 (2006)
- [30] DONUT Collaboration: K. Kodama et al., Observation of Tau Neutrino Interactions, Phys. Lett. **B504**, 218 (2001)
- [31] K. Kodama et al., Final tau-neutrino results from the DONuT experiment, Phys. Rev. D **78**, 052002 (2008)

- [32] C. Amsler et al., Review of Particle Physics, Phys. Lett. **B667**, 1 (2008)
- [33] E. Akhmedov and J. Kopp, Neutrino oscillations: quantum mechanics vs. quantum field theory, Journal of High Energy Physics **2010**, 1 (2010)
- [34] B. Pontecorvo, Inverse beta processes and nonconservation of lepton charge, Sov. Phys. JETP **7**, 172 (1958)
- [35] Z. Maki, M. Nakagawa, and S. Sakata, Remarks on the unified model of elementary particles, Prog. Theor. Phys. **28**, 870 (1962)
- [36] L. Wolfenstein, Neutrino oscillations in matter, Phys. Rev. **D17**, 2369 (1978)
- [37] S. P. Mikheev and A. Y. Smirnov, Resonance enhancement of oscillations in matter and solar neutrino spectroscopy, Sov. J. Nucl. Phys. **42**, 913 (1985)
- [38] J. N. Bahcall, Solar Neutrinos. I. Theoretical, Phys. Rev. Lett. **12**, 300 (1964)
- [39] R. Davis, D. S. Harmer, and K. C. Hoffman, Search for neutrinos from the sun, Phys. Rev. Lett. **20**, 1205 (1968)
- [40] B. T. Cleveland et al., Measurement of the Solar Electron Neutrino Flux with the Homestake Chlorine Detector, The Astrophysical Journal **496**, 505 (1998)
- [41] Kamiokande-II Collaboration: K. S. Hirata et al., Observation of ^8B solar neutrinos in the Kamiokande-II detector, Phys. Rev. Lett. **63**, 16 (1989)
- [42] A. I. Abazov et al., Search for neutrinos from the Sun using the reaction $^{71}\text{Ga}(\nu_e, e^-)^{71}\text{Ge}$, Phys. Rev. Lett. **67**, 3332 (1991)
- [43] GALLEX Collaboration: P. Anselmann et al., Solar neutrinos observed by GALLEX at Gran Sasso., Phys. Lett. **B285**, 376 (1992)

- [44] SNO Collaboration: B. Aharmim et al., Electron Energy Spectra, Fluxes, and Day-Night Asymmetries of ^8B Solar Neutrinos from the 391-Day Salt Phase SNO Data Set, Phys. Rev. **C72**, 055502 (2005)
- [45] KamLAND Collaboration: S. Abe et al., Precision Measurement of Neutrino Oscillation Parameters with KamLAND, Phys. Rev. Lett. **100**, 221803 (2008)
- [46] Kamiokande Collaboration: K. S. Hirata et al., Experimental study of the atmospheric neutrino flux, Phys. Lett. **B205**, 416 (1988)
- [47] IMB Collaboration: R. Becker-Szendy et al., Electron- and muon-neutrino content of the atmospheric flux, Phys. Rev. D **46**, 3720 (1992)
- [48] NUSEX Collaboration: M. Aglietta et al., Experimental Study of Atmospheric Neutrino Flux in the NUSEX Experiment, Europhys. Lett. **8**, 611 (1989)
- [49] Fréjus Collaboration: K. Daum et al., Determination of the atmospheric neutrino spectra with the Fréjus detector, Z. Phys. **C66**, 417 (1995)
- [50] Super-Kamiokande Collaboration: Y. Fukuda et al., Evidence for oscillation of atmospheric neutrinos, Phys. Rev. Lett. **81**, 1562 (1998)
- [51] Super-Kamiokande Collaboration: Y. Ashie et al., Evidence for an oscillatory signature in atmospheric neutrino oscillations, Phys. Rev. Lett. **93**, 101801 (2004)
- [52] M. Sanchez et al., Measurement of the L/E distributions of atmospheric ν in Soudan 2 and their interpretation as neutrino oscillations, Phys. Rev. D **68**, 113004 (2003)

- [53] CHOOZ Collaboration: M. Apollonio et al., Search for neutrino oscillations on a long base-line at the CHOOZ nuclear power station, *Eur. Phys. J.* **C27**, 331 (2003)
- [54] K2K Collaboration: M. H. Ahn et al., Measurement of neutrino oscillation by the K2K experiment, *Phys. Rev.* **D74**, 072003 (2006)
- [55] N. Agafonova et al., Observation of a first ν_τ candidate event in the OPERA experiment in the CNGS beam, *Physics Letters B* **691**, 138 (2010)
- [56] MINOS Collaboration: P. Adamson et al., Measurement of Neutrino Oscillations with the MINOS Detectors in the NuMI Beam, *Phys. Rev. Lett.* **101**, 131802 (2008)
- [57] P. Vahle, New Results from MINOS, (2010), Proceedings of The XXIV International Conference on Neutrino Physics and Astrophysics (in preparation)
- [58] R. P. Litchfield, *Neutrino induced events in the MINOS detectors*, DPhil thesis, University of Oxford, 2008
- [59] Y. Itow et al., The JHF-Kamioka neutrino project, (2001), hep-ex/0106019
- [60] NOvA Collaboration: D. S. Ayres et al., NOvA proposal to build a 30-kiloton off-axis detector to study neutrino oscillations in the Fermilab NuMI beamline, (2004)
- [61] Double Chooz Collaboration: F. Ardellier et al., Double Chooz: A search for the neutrino mixing angle θ_{13} , hep-ex/0606025 (2006)
- [62] Daya Bay Collaboration: X. Guo et al., A precision measurement of the neutrino mixing angle θ_{13} using reactor antineutrinos at Daya Bay, hep-ex/0701029 (2007)

- [63] RENO Collaboration: J. Ahn et al., RENO: An Experiment for Neutrino Oscillation Parameter θ_{13} Using Reactor Neutrinos at Yonggwang, hep-ex/1003.1391 (2010)
- [64] S. Dodelson and L. M. Widrow, Sterile neutrinos as dark matter, Phys. Rev. Lett. **72**, 17 (1994)
- [65] J. Hidaka and G. M. Fuller, Sterile neutrino-enhanced supernova explosions, Phys. Rev. D **76**, 083516 (2007)
- [66] A. Kusenko, Sterile neutrinos, New Journal of Physics **11**, 105007 (8pp) (2009)
- [67] C. Athanassopoulos et al., The liquid scintillator neutrino detector and LAMPF neutrino source, Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment **388**, 149 (1997)
- [68] C. Athanassopoulos et al., Evidence for $\bar{\nu}_\mu \rightarrow \bar{\nu}_e$ Oscillations from the LSND Experiment at the Los Alamos Meson Physics Facility, Phys. Rev. Lett. **77**, 3082 (1996)
- [69] C. Athanassopoulos et al., Results on $\nu_\mu \rightarrow \nu_e$ Neutrino Oscillations from the LSND Experiment, Phys. Rev. Lett. **81**, 1774 (1998)
- [70] The MiniBooNE Collaboration: A. A. Aguilar-Arevalo et al., A Search for electron neutrino appearance at the $\Delta m^2 \sim 1 \text{ eV}^2$ scale, Phys. Rev. Lett. **98**, 231801 (2007)
- [71] M. Maltoni and T. Schwetz, Sterile neutrino oscillations after first MiniBooNE results, Phys. Rev. D **76**, 093005 (2007)
- [72] A. A. Aguilar-Arevalo et al., Search for Electron Antineutrino Appearance at the $\Delta m^2 \sim 1 \text{ eV}^2$ Scale, Phys. Rev. Lett. **103**, 111801 (2009)

- [73] Super-Kamiokande Collaboration: S. Fukuda et al., Tau neutrinos favored over sterile neutrinos in atmospheric muon neutrino oscillations, *Phys. Rev. Lett.* **85**, 3999 (2000)
- [74] Super-Kamiokande Collaboration: A. Habig, Discriminating between $\nu_\mu \leftrightarrow \nu_\tau$ and $\nu_\mu \leftrightarrow \nu_{\text{sterile}}$ in atmospheric ν_μ oscillations with the Super-Kamiokande detector, (2001), Contributed to 27th International Cosmic Ray Conferences (ICRC 2001), Hamburg, Germany, 7–15 Aug 2001.
- [75] P. Litchfield, On the generality of the f_s parametrisation, MINOS-doc-4618 (2008)
- [76] J. Lesgourgues and S. Pastor, Massive neutrinos and cosmology, *Physics Reports* **429**, 307 (2006)
- [77] C. Giunti and C. W. Kim, *Fundamentals of Neutrino Physics and Astrophysics* (Oxford University Press, 2007), First edition
- [78] H. Klapdor-Kleingrothaus and I. Krivosheina, The Evidence for the Observation of $0\nu\beta\beta$ Decay: The Identification of $0\nu\beta\beta$ Events from the Full Spectra, *Modern Physics Letters A* **21**, 1547 (2006)
- [79] C. Kraus, SNO with liquid scintillator: SNO+, *Progress in Particle and Nuclear Physics* **57**, 150 (2006), International Workshop of Nuclear Physics 27th course - Neutrinos in Cosmology, in *Astro, Particle and Nuclear Physics*
- [80] EXO Collaboration: N. Ackerman, Status of EXO-200, (2009)
- [81] K. Zuber, COBRA - Double beta decay searches using CdTe detectors, (2001)
- [82] C. Arnaboldi et al., CUORE: a cryogenic underground observatory for rare events, *Nuclear Instruments and Methods in Physics Research Section A*:

- Accelerators, Spectrometers, Detectors and Associated Equipment **518**, 775 (2004)
- [83] J. Hylen et al., NuMI Technical Design Handbook, Internal NuMI report (2003)
- [84] S. E. Kopp, The NuMI neutrino beam at Fermilab, (2005), Proceedings of the Particle Accelerator Conference (PAC 05), Knoxville, Tennessee.
- [85] MINOS Collaboration: D. G. Michael et al., The Magnetized Steel and Scintillator Calorimeters of the MINOS Experiment, Nucl. Instrum. Meth. **A569**, 190 (2008)
- [86] N. Tagg et al., Performance of Hamamatsu 64-anode photomultipliers for use with wavelength-shifting optical fibres, Nucl. Instrum. Meth. **A539**, 668 (2005)
- [87] K. Lang et al., Characterization of 1600 Hamamatsu 16-anode photomultipliers for the MINOS Far detector, Nucl. Instrum. Meth. **A545**, 852 (2005)
- [88] T. Cundiff et al., The MINOS near detector front end electronics, IEEE Trans. Nucl. Sci. **53**, 1347 (2006)
- [89] N. Felt et al., MINOS Far Detector Electronics User's Manual, Internal MINOS document (2001)
- [90] J. Hartnell, *Measurement of the Calorimetric Energy Scale in MINOS*, DPhil thesis, University of Oxford, 2005
- [91] A. Cabrera et al., Comparisons of the MINOS near and far detector read-out systems at a test beam, Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment **609**, 106 (2009)

- [92] A. Cabrera, *Systematic Comparison of the MINOS Near and Far Detector Readout Systems*, DPhil thesis, University of Oxford, 2005
- [93] A. Fasso, A. Ferrari, P. R. Sala, and J. Ranft, FLUKA: Status and prospects for hadronic applications, Contributed to the International Conference on Advanced Monte Carlo for Radiation Physics, Particle Transport Simulation and Applications (MC 2000), Lisbon, Portugal, 23–26 October 2000
- [94] S. Agostinelli et al., GEANT4—a simulation toolkit, Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment **506**, 250 (2003)
- [95] H. Gallagher, The NEUGEN neutrino event generator, Nucl. Phys. Proc. Suppl. **112**, 188 (2002)
- [96] GEANT—Detector description and simulation tool, CERN Program Library, Long Writeup, W5013, <http://wwwasd.web.cern.ch/wwwasd/geant/index.html>
- [97] C. Zeitnitz and T. A. Gabriel, The GEANT - CALOR interface and benchmark calculations of ZEUS test calorimeters, Nucl. Instrum. Meth. **A349**, 106 (1994)
- [98] M. Dorman, Beam fit position paper, MINOS-doc-7146 (2010)
- [99] MINOS Collaboration: P. Adamson et al., New constraints on muon-neutrino to electron-neutrino transitions in MINOS, (2010)
- [100] MINOS Collaboration: P. Adamson et al., First observations of separated atmospheric ν_μ and $\bar{\nu}_\mu$ events in the MINOS detector, Phys. Rev. **D73**, 072002 (2006)

- [101] MINOS Collaboration: P. Adamson et al., Measurement of the atmospheric muon charge ratio at TeV energies with MINOS, *Phys. Rev.* **D76**, 052003 (2007)
- [102] NCAS and MINOS Collaboration: S. Osprey et al., Sudden Stratospheric Warmings seen in MINOS Deep Underground Muon Data, (2009), Accepted for publication by *Geophys. Res. Lett.*
- [103] P. Adamson et al., Neutrino and antineutrino inclusive charged-current cross section measurements with the MINOS near detector, *Phys. Rev. D* **81**, 072002 (2010)
- [104] G. Tinti, *Sterile neutrino oscillations in MINOS and constraints on the neutrino flux from hadron production data in NA49*, DPhil thesis, University of Oxford, 2010
- [105] J. de Jong, Position Paper for the Runs I-III FD NC Cleaning, MINOS-doc-7165 (2010)
- [106] R. Pittam, *A search for sterile neutrinos at the MINOS experiment*, DPhil thesis, University of Oxford, 2010
- [107] D. Mackay, *Information Theory, Inference, and Learning Algorithms* (Cambridge University Press, 2003), First edition
- [108] MINOS Collaboration: D. G. Michael et al., Observation of muon neutrino disappearance with the MINOS detectors and the NuMI neutrino beam, *Phys. Rev. Lett.* **97**, 191801 (2006)
- [109] P. Adamson et al., Search for sterile neutrino mixing in the MINOS long-baseline experiment, *Phys. Rev. D* **81**, 052004 (2010)

- [110] MINOS Collaboration: P. Adamson et al., Study of Muon Neutrino Disappearance using the Fermilab Main Injector Neutrino Beam, *Phys. Rev. D* **77**, 072002 (2008)
- [111] S. Baker and R. Cousins, Clarification on the use of chi-square and likelihood functions in fits to histograms, *Nucl. Instrum. Meth.* **221**, 437 (1984)
- [112] F. James, *Statistical Methods in Experimental Physics* (World Scientific, 2006), Second edition
- [113] J. Boehm, *Measurement of Electron Neutrino Appearance with the MINOS Experiment*, PhD thesis Harvard University 2009
- [114] J. Nelson, Notes on the normalization systematic, MINOS-doc-2106 (2006)
- [115] D. Boehnlein et al., The thickness of the MINOS detector steel, MINOS-doc-1529 (2006)
- [116] M. Kordosky and M. Dorman, Flux measurements with the near detector, MINOS-doc-1547 (2006)
- [117] P. Litchfield, Scan comparison of data and MC, near and far detectors, MINOS-doc-1549 (2006)
- [118] A. Blake, Data Validation and PoT Counting, MINOS-doc-6322 (2009)
- [119] R. Pittam, POT counting for neutral current top-up analysis, MINOS-doc-5283 (2008)
- [120] P. Adamson et al., Measuring the Number of Protons-on-Target (POT) in the NuMI Beamline, MINOS-doc-1491 (2006)
- [121] T. M. C. Group, 2009 Position paper on calibration of runs I-II-III, MINOS-doc-6717 (2010)

- [122] H. Gallagher, S. Dytman, and M. Kordosky, Shower Energy Scale Uncertainty, MINOS-doc-4309 (2008)
- [123] A. Sousa et al., Neutral current analysis of 7.1×10^{20} POT of MINOS data: executive summary, MINOS-doc-7191 (2010)
- [124] R. Ospanov, *A Measurement of Muon Neutrino Disappearance with the MINOS Detectors and NuMI Beam*, PhD thesis University of Texas at Austin 2008
- [125] J. Ratchford, A data driven estimate to the NC background, MINOS-doc-7115 (2010)
- [126] L. Hsu and P. Rodrigues, Charged current background in the neutral current sample, MINOS-doc-3878 (2007)
- [127] C. Andreopoulos et al., Updated cross section model uncertainties for the charged current analysis, MINOS-doc-2989 (2007)
- [128] v. Z. Pavlović, *Observation of disappearance of muon neutrinos in the NuMI beam*, PhD thesis University of Texas at Austin 2008
- [129] A. Holin, Muon Removal Background Estimation Method, MINOS-doc-4141 (2008)
- [130] A. Agresti and B. A. Coull, Approximate Is Better than "Exact" for Interval Estimation of Binomial Proportions, *The American Statistician* **52**, 119 (1998)
- [131] J. M. Paley, Hand-scan estimation of the reconstruction-related systematic uncertainty in the near/far normalization of the MINOS CC analysis, MINOS-doc-5613 (2009)

- [132] C. Backhouse, Systematics interpolation for 7×10^{20} pot analysis, MINOS-doc-6048 (2009).