Measuring neutrino oscillation parameters using ν_μ disappearance in MINOS

Christopher James Backhouse

Oriel College, Oxford

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

Hilary Term, 2011

Measuring neutrino oscillation parameters using ν_μ disappearance in MINOS

Christopher James Backhouse
Oriel College, Oxford

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

Hilary Term, 2011

Abstract

MINOS is a long-baseline neutrino oscillation experiment. It consists of two large steel-scintillator tracking calorimeters. The near detector is situated at Fermilab, close to the production point of the NuMI muon-neutrino beam. The far detector is 735 km away, 716 m underground in the Soudan mine, Northern Minnesota.

The primary purpose of the MINOS experiment is to make precise measurements of the "atmospheric" neutrino oscillation parameters ($\Delta m^2_{\mathrm{atm}}$ and $\sin^2 2\theta_{\mathrm{atm}}$). The oscillation signal consists of an energy-dependent deficit of ν_μ interactions in the far detector. The near detector is used to characterize the properties of the beam before oscillations develop. The two-detector design allows many potential sources of systematic error in the far detector to be mitigated by the near detector observations.

This thesis describes the details of the ν_μ-disappearance analysis, and presents a new technique to estimate the hadronic energy of neutrino interactions. This estimator achieves a significant improvement in the energy resolution of the neutrino spectrum, and in the sensitivity of the neutrino oscillation fit. The systematic uncertainty on the hadronic energy scale was re-evaluated and found to be comparable to that of the energy estimator previously in use.

The best-fit oscillation parameters of the ν_μ-disappearance analysis, incorporating this new estimator were: $\Delta m^2 = 2.32^{+0.12}_{-0.08} \times 10^{-3}\,\mathrm{eV}^2$, $\sin^2 2\theta > 0.90$ (90% C.L.). A similar analysis, using data from a period of running where the NuMI beam was operated in a configuration producing a predominantly $\bar{\nu}_\mu$ beam, yielded somewhat different best-fit parameters $\Delta\bar{m}^2 = \left(3.36^{+0.46}_{-0.40}(\mathrm{stat.}) \pm 0.06(\mathrm{syst.})\right) \times 10^{-3}\mathrm{eV}^2$, $\sin^2 2\bar{\theta} = 0.86^{+0.11}_{-0.12}(\mathrm{stat.}) \pm 0.01(\mathrm{syst.})$. The tension between these results is intriguing, and additional antineutrino data is currently being taken in order to further investigate this apparent discrepancy.

Declaration

I declare that this thesis and the work presented in it are my own and were produced by me as the result of my own original research. The work was done while a candidate for a degree at the University of Oxford, and has not been submitted for any other qualification. Results and figures from published works have been clearly attributed.

The energy estimation technique presented in Chapter 5 of this thesis was my primary original contribution to the MINOS charged-current analysis. The analysis as a whole was performed by a group within the collaboration, and while I made significant contributions, the innovations described in Chapter 4 are the work of others (in many cases the subject of a thesis in their own right).

Acknowledgments

Over the course of $3^1/_2$ years I've collected quite a number of people without whom this thesis wouldn't have been possible (or at least severely delayed). So here goes, in roughly chronological order:

Thanks to my supervisor, Giles Barr, who never ceases to be optimistic, no matter how broken the analysis, or close the deadline. I also appreciated his offers (which I never had to take up) to throw his weight around and leave me to get on with writing, and his prompt proofreading (by usual advisor standards) that let me do so at a reasonable rate. Thanks also to my secondary supervisor, Alfons Weber, who supervised the fourth year project that got me hooked on neutrinos in the first place.

The generation of Oxford MINOS students above me: Jeff, Justin, Phill and Tobi, set a good example when I first arrived: working hard laying the foundations of all the analyses to come, and taking copious coffee breaks. They all then set another good example by graduating shortly after.

The bulk of my time, the group consisted of Alex, Bob, Gemma, and Phil. I really appreciated having so much expertise so near at hand. Long coffee breaks remained a productive way to hammer out physics problems, Friday pub lunch less so.

Thanks to the CC group for putting up with the disproportionate number of bugs I was responsible for. Congratulations on surviving the great box opening of 2010. Also, to members of the $\bar{\nu}_\mu$ and miscellaneous other groups for answering all my questions, and for making collaboration meetings fun.

With the "second generation" group all graduated, and me in the middle of writing up, the task of keeping my sane fell to: two great societies, the OUS and OS, who gave me something other than work to do, almost every evening; Abbey who looked sympathetic at a lot of tedious tribulations, and who I'm confident will bring some of the spirit of the old MINOS office to the T2K heathens; and Zeynep, who frequently managed to take my mind off writing via the medium of perplexing statistics questions.

Last, but not least, thank you to my family for putting up with so many years of physics (with more to come), and not seeming too incredulous at my claims that I find this stuff interesting.

Contents

List of Figures

List of Tables

Chapter 1

Introduction

With the discovery of the top quark at the Tevatron in 1995 [1, 2] the particle content of the Standard Model was complete (save for the ever-elusive Higgs boson). In the fifteen years of careful experimentation since then, no convincing evidence has yet been found for any behaviour outside of the Standard Model, except in the neutrino sector.

Evidence for deficits in the neutrino rate from the sun, and from the atmosphere, the "solar neutrino problem" and "atmospheric neutrino anomaly", was longstanding. However, the confirmation of the interpretation of these in terms of the phenomenon of neutrino oscillations has only come relatively recently. The evidence in favour of this interpretation is now very strong, and the existence of non-zero neutrino masses and of neutrino mixing is now well established. Chapter 2 describes the historical progress of the field of neutrino physics, and the theoretical basis and derivation of the standard neutrino oscillation formulae.

Much has been learnt about the properties of neutrinos in the last few years. We now have good measurements of two of the three neutrino mixing angles, and of the differences in their (squared) masses. Still unknown are: the magnitude of the final mixing angle; the question of whether neutrino oscillations violate $\mathcal{CP}$ symmetry, and if so the value of the corresponding angle; the absolute, rather than relative, values of the neutrino masses; and the correct resolution of the current remaining ambiguity in the ordering of the mass states. There of course remains the possibility of surprises, not captured in this orderly programme of measurements.

The MINOS experiment, consisting of two functionally identical detectors, 1 km and 735 km from the source of the NuMI neutrino beam, was constructed to make precise measurements of the "atmospheric" neutrino oscillation parameters: the mass splitting and mixing angle that govern $\nu_\mu \leftrightarrow \nu_\tau$ oscillations in the energy and distance regime $L/E \sim 500\text{km/GeV}$. The MINOS detectors, beam, Monte Carlo simulation and reconstruction software are described in Chapter 3.

Chapter 4 details the components of the primary MINOS analysis. The atmospheric mass splitting Δm^2_{atm} and mixing $\sin^2 2\theta_{\text{atm}}$ are determined from the deficit of charged-current ν_μ interactions observed in the far detector – specifically, the depth and position of the energy-dependent oscillation dip. This chapter describes the procedures used to select charged-current events from the neutral-current background, to use information gained from the near detector to produce a more accurate far detector prediction, and to find the oscillation parameters that best fit the observed data spectrum. Two analysis refinements: dividing events according to their estimated energy resolution, and including a new sample of events from outside the fiducial volume, are also described.

A significant analysis improvement, developed for this thesis, is described in detail in Chapter 5. The energy resolution of the hadronic (or "shower") component of neutrino interactions is substantially improved by using a multivariate estimator based on three variables characterizing the shower activity in an event. Improved energy resolution leads to a more precise determination of the details of the oscillation dip, and improved sensitivity to the neutrino oscillation parameters.

Chapter 6 presents the results of the oscillation fit performed using the techniques of Chapters 4 and 5. This analysis provides the world's most precise measurement of the mass splitting Δm^2. Also presented in this chapter, are the results of the similar $\bar{\nu}_\mu$-disappearance analysis, which provides a measurement of the same oscillation parameters from an antineutrino beam.

Finally, the conclusion in Chapter 7 summarizes MINOS oscillation results, and describes the plans for the future running of MINOS. It also gives an overview of the expected course of the next several years of neutrino oscillation physics, the reactor and accelerator neutrino experiments that aim to extend our knowledge of the neutrino sector.

Chapter 2

History and theory of neutrino physics

This chapter provides an overview of the history of neutrino physics, beginning with the proposal and discovery of the neutrino and the weak interaction in Section 2.1 and then describing, in Section 2.3, the evidence of anomalies in the rates of solar and atmospheric neutrinos, now explained by the phenomenon of neutrino oscillations. Section 2.4 derives the standard neutrino oscillation formula, and goes on to discuss various useful approximations, and the effect on the oscillatory behaviour caused by passage through matter. The following Sections 2.5 to 2.9 go on to describe significant oscillation experiments, and the best evidence about the properties of neutrinos that they have provided us. The chapter concludes by describing other significant constraints on the properties of neutrinos (particularly their absolute masses) obtained by non-oscillation experiments (Section 2.10), and from astrophysical evidence (Section 2.11).

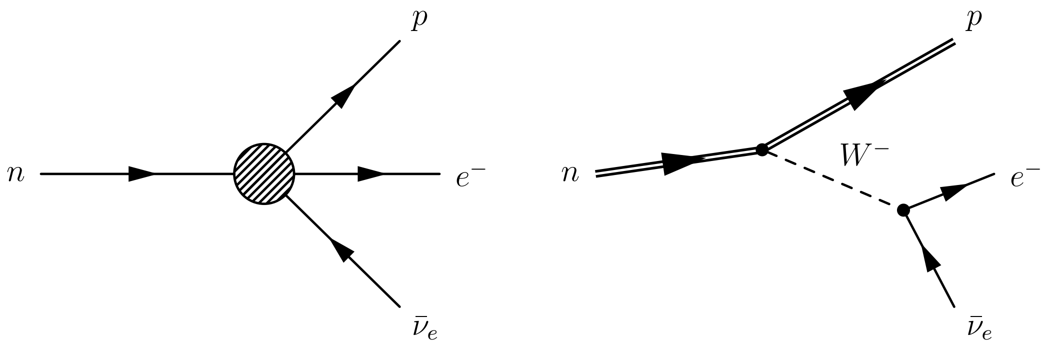

Figure 2.1: Beta decay according to the four-point interaction of Fermi theory (left) and via exchange of an intermediate W boson as in the Standard Model (right).

2.1 Early history of the neutrino and the weak interaction

The neutrino began life named the *neutron*, and was considered a “desperate remedy”, the proposal of a (virtually) undetectable particle being deemed the lesser of two evils compared to the prospect of energy non-conservation. The problem that caused Pauli to resort to such extreme measures in his famous 1930 letter to the Tübingen conference [3] was the observed continuous energy spectrum of electrons emitted in beta decay. He proposed a new uncharged, spin-$\frac{1}{2}$ particle, existing in the nucleus, and emitted undetected in conjunction with the electron. The discovery of the neutron proper by Chadwick in 1932 [4] prompted Fermi to rename Pauli’s particle the *neutrino* or “little neutral one”.

The theory of beta decay and the weak interaction was soon given some theoretical grounding by Fermi [5]. In Fermi theory, beta decay is a four-point interaction, as illustrated in the left half of Figure 2.1. Rearrangement of the incoming and outgoing particles allows the prediction of additional processes, such as electron capture, and the inverse beta-decay process by which the neutrino was first detected. Whilst this theory was not without its problems – it was quickly recognized that it violates unitarity at high energies [6] – it was sufficient for the inverse beta-decay cross-section to be calculated [7], and dismissed as too small to ever

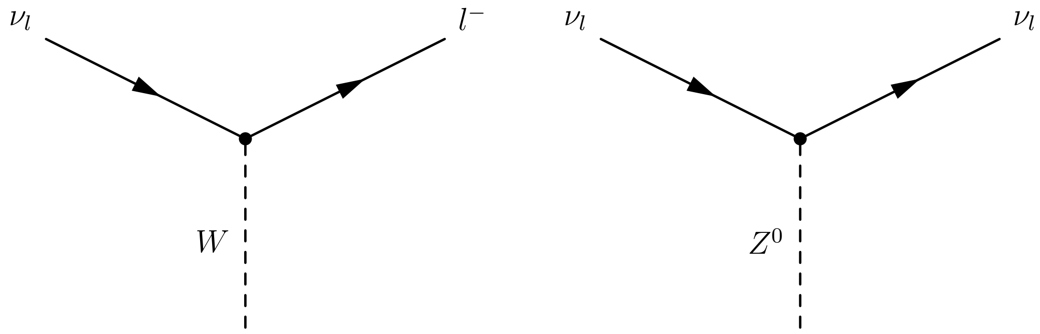

Figure 2.2: Feynman diagrams for the two possible neutrino interactions: a charged-current interaction via exchange of a $W^{\pm}$ (left), or a neutral-current interaction via the Z^0 boson (right).

observe.

It was indeed 22 years before the neutrino was experimentally detected at Hanford [8], and then more definitively at Savannah River [9], via the inverse beta reaction of neutrinos from a nuclear reactor in a scintillator detector. The measured rate was consistent with theoretical expectations for the cross-section.

In 1956, an assumption that had been made up to that point, that weak interactions conserve parity, was revealed to be just that: an assumption with no experimental justification [10]. A year later, parity violation was observed for the first time. The angular distribution of the electrons emitted in ^{60}Co decay was observed to be strongly correlated with the direction of the nuclear spin[1] [11]. The neutrino helicity was determined in 1958. The observed polarization of the photon emitted in the decay of ^{152}Eu requires the neutrino to be a left-handed particle [12].

The theory of weak interactions was finalized into the form it now takes in the Standard Model during the 1960s. Based on Yukawa's concept of an exchange boson [13], which solves the unitarity violation problem in Fermi theory, the $W^{\pm}$ and Z^0 bosons were introduced, their masses predicted, and the weak force unified

---

[1] The dot product of spin and momentum $\sigma \cdot p$ is a pseudoscalar quantity, and changes sign under parity transformations. If parity were conserved there should be no dependence of the rate on this quantity.

with electromagnetism [14, 15, 16]. Figure 2.2 shows the two interaction vertices for neutrinos in this theory. Beta decay is then mediated by the exchange of a W (right side of Figure 2.1). The $V - A$ form of the interaction predicts maximal parity violation and left-handed neutrinos. While $\mathcal{CP}$-violation has been observed in weak interactions, it has not yet been observed in the neutrino sector.

A key prediction of the Glashow-Salam-Weinberg model was the existence of “neutral-current” (NC) neutrino interactions, mediated by the Z-boson (right side of Figure 2.2). The Gargamelle bubble-chamber at CERN observed both NC electron-scattering [17] and hadronic [18] neutrino interactions in 1973. Direct discovery of the exchange bosons themselves, including confirmation of their masses, came in 1983 from the UA1 and UA2 experiments at CERN’s Sp$\bar{\text{p}}$S [19, 20, 21, 22].

2.2 The number of neutrino generations

The concept of a conserved lepton number was introduced in 1953 [23]. A 1962 experiment at Brookhaven [24] demonstrated that the neutrinos produced in pion decay, when interacting in a detector, produced muons and not electrons. This indicated the existence of two different types of neutrino, one coupling to the electron, the other to the muon. Together with the non-observation of the decay $\mu \rightarrow e + \gamma$, this led to the introduction of lepton generation number, conserved separately for electron- and muon-like leptons.

In 1975, events of the form $e^+ + e^- \rightarrow e^\pm + \mu^\mp +$ [missing energy] were observed at SLAC [25]. Interpreted as the production of a pair of particles which decay either to muons or to electrons, plus neutrinos, this was the first evidence for another generation of leptons. The mystery particle is what we now know as the τ. The probability of the existence of a corresponding tau neutrino was immediately raised but, due to the technical difficulties of producing and detecting them, the first evidence for their existence came via an indirect method, as follows.

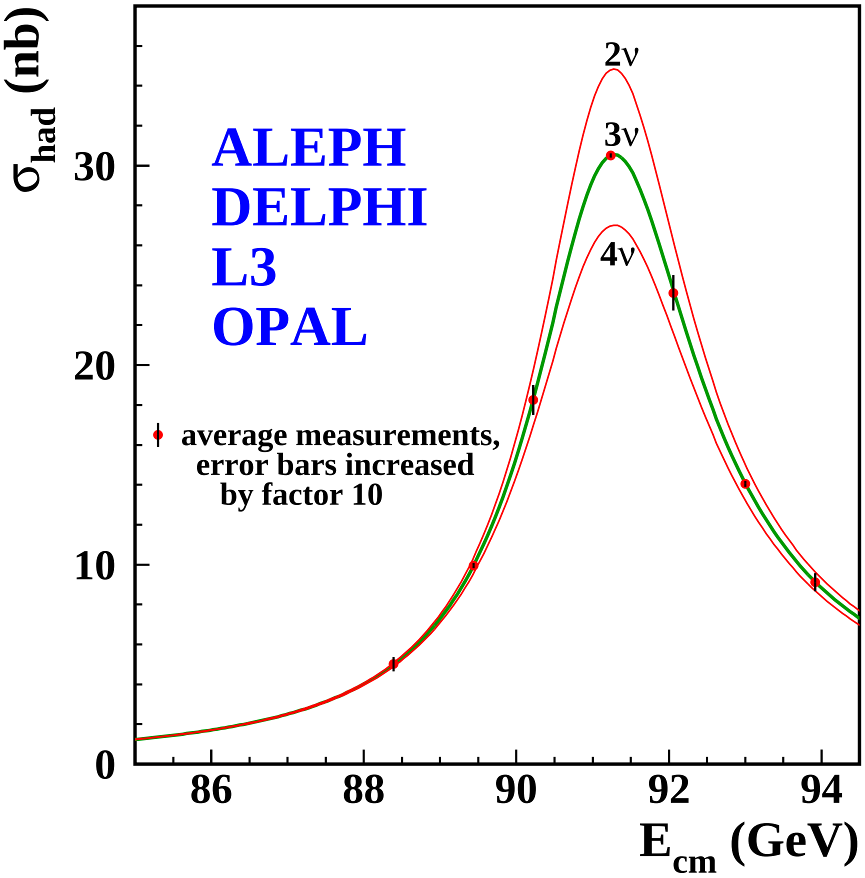

Figure 2.3: Cross section for $e^+ + e^- \rightarrow$ hadrons at centre-of-mass energies around the Z-mass from LEP. The prediction for three neutrino generations (green) fits the data much better than two or four (red). Taken from [26].

The Breit-Wigner formula for the cross-section σ_f of a process around a resonance E_0 as a function of energy is

$$\sigma_f(E) = \sigma_{\rm tot}(E_0)\frac{\Gamma_f^2/4}{(E-E_0)^2+\Gamma_{\rm tot}^2/4} \ , \qquad (2.1)$$

where Γ_f is the width to the final state in question, $\Gamma_{\rm tot}$ is the width to all final states, and E_0 is the mass of the virtual particle. The $\Gamma_{\rm tot}$ term in the denominator means that additional decay modes increase the width of the resonant peak for all other modes. Extremely accurate measurements of the properties of the Z-peak were made at LEP [26]. Figure 2.3 shows the cross-section obtained at different energies around the Z mass, compared to the prediction in the case of

2, 3, or 4 generations of neutrinos, which contribute to the Z width via the decay $Z^0 \rightarrow \nu + \bar{\nu}$. Leaving the number of neutrino generations free in the fit, a value of $N_\nu = 2.9840 \pm 0.0082$ was obtained, confirming the existence of the third, tau, neutrino and ruling out further generations beyond the three already known. The only way in which additional generations of neutrinos could exist is if either they exceed half the Z-mass (*i.e.* 45 GeV) or if they are of a novel "sterile" type which does not couple to the Z, and would thus have no Standard Model interactions at all, interacting with other matter only via gravity.

Direct confirmation of the existence of the ν_τ was not obtained until 25 years after the discovery of its counterpart the τ. The DONUT experiment at Fermilab used a neutrino beam produced by the decay of charmed mesons created by a 800 GeV proton beam from the Tevatron, incident on a tungsten target. The ν_τ component of the beam came primarily from the decay of D_S mesons, and the detector consisted of interleaved layers of steel and photographic emulsion. At these energies, the tau lepton formed in a charged-current ν_τ interaction only travels about 2 mm before it decays, so the signal of a ν_τ-event is a track with a sharp kink near the vertex, indicating the decay of the initial τ. In 2000 DONUT reported the first evidence for ν_τ events [27]. In their final dataset they observe 9 candidate tau events over an expected background of 1.5 events [28], obtaining conclusive evidence for ν_τ interactions and completing the fermion content of the Standard Model.

2.3 Evidence for neutrino oscillations

2.3.1 The solar neutrino problem

By the 1960s the mechanisms of energy production in the sun were well understood theoretically, but since the zone of energy production is hidden deep within the

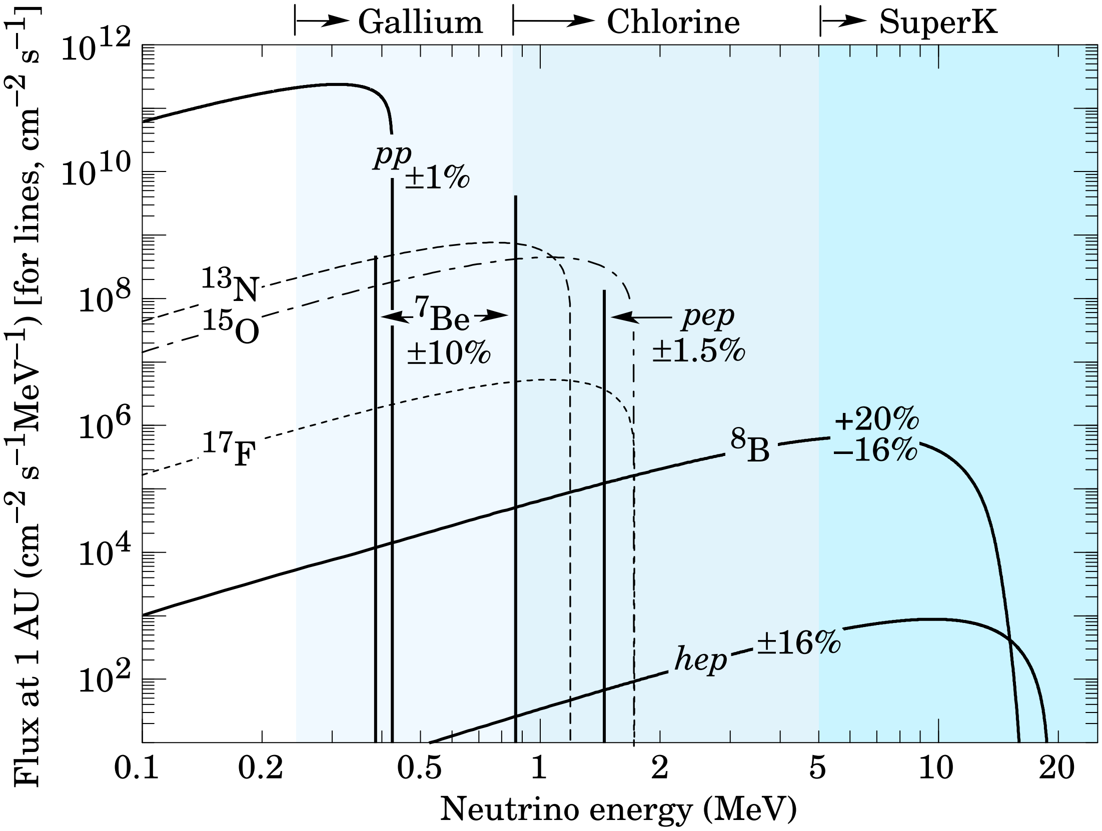

Figure 2.4: Prediction of the solar neutrino energy spectrum from the Standard Solar Model. The shaded regions indicate the energy ranges over which various classes of experiment are sensitive. Taken from [29].

sun at its core, direct confirmation had not been obtained. Neutrinos, with their ability to reach the earth directly from the solar core, held the potential to provide this direct evidence of the nuclear reactions occurring there. Attempts to use these solar neutrinos for astrophysics provided the first evidence for the phenomenon of neutrino oscillations.

According to the Standard Solar Model (SSM), energy generation in the sun occurs mainly via two processes: the *pp* chain, in which Hydrogen nuclei (protons) fuse together directly to ultimately become ^{4}He; and the CNO cycle, in which carbon, nitrogen and oxygen nuclei act as catalysts for the transformation $4p \to {}^4\text{He}$. Both processes contain steps that involve the emission of an electron neutrino. In our sun the *pp* chain dominates, and is responsible for the majority of the energy production. Figure 2.4 shows a modern prediction of the neutrino flux expected from the sun. The continuous spectra are due to neutrinos emitted in three-body decays, the lines from two-body decays.

The first experiment to search for, and successfully detect, these solar neutrinos was carried out by Ray Davis in the Homestake mine, South Dakota, starting in 1968 [30]. Using a 390,000 litre tank of tetrachloroethylene (C_2Cl_4), neutrino interactions occur via a charged-current interaction to form an unstable isotope of argon, with a neutrino energy threshold of 814 keV:

$$\nu_e + {}^{37}\mathrm{Cl} \rightarrow e^- + {}^{37}\mathrm{Ar}\ . \tag{2.2}$$

The argon atoms were flushed from the tank using helium gas, and detected by their decays back to ^{37}Cl. The data from the Homestake experiment consistently showed an unexpectedly low rate of neutrino interactions, approximately a third of the SSM prediction [31]. At the time, this anomaly was put down to either incorrect flux predictions from the solar models or to a problem with the experiment, which is technically challenging, searching for radioactive decays that happen at a rate on the order of one atom per week.

Further evidence for a deficit of electron neutrinos from the sun had to wait until 1989 when the Kamiokande-II water-Čerenkov experiment (originally built to search for proton decay) reported their observations of solar neutrinos [32]. The signal observed was the neutral-current interaction of a solar neutrino with an electron in the water, and the resulting Čerenkov light from the outgoing electron:

$$\nu_e + e^- \rightarrow \nu_e + e^-\ . \tag{2.3}$$

This reaction is sensitive to the ^{8}B decay step of the *pp*-III chain in the sun, one of the flux components observed by the Homestake experiment:

$$^{8}\mathrm{B} \to {}^{8}\mathrm{Be} + e^{+} + \nu_e + \gamma \ . \tag{2.4}$$

Compared to radiochemical experiments, Kamiokande had several advantages, including real-time read-out and the ability to determine the direction of the recoil electron, and hence obtain evidence that the detected neutrinos were coming from the direction of the sun. The results confirmed a deficit in the rate of solar neutrino interactions, finding a rate about half that expected. The deviations of the Homestake and Kamiokande observations from the SSM predictions became known as the "solar neutrino problem".

Independent confirmation of a deficit in solar neutrino rates came from two more experiments, using a third technique: SAGE [33] in 1991 and GALLEX [34] in 1992. Both experiments took advantage of the reaction

$$\nu_e + {}^{71}\mathrm{Ga} \to {}^{71}\mathrm{Ge} + e^{-} \tag{2.5}$$

using liquid gallium, extracting the germanium atoms, and observing their decay, analogously to the Homestake experiment. The advantage of using this reaction is that the threshold is sufficiently low (233 keV) that it is sensitive to neutrinos from the first step of the *pp* chain:

$$p + p \to {}^{2}\mathrm{H} + e^{+} + \nu_e \ . \tag{2.6}$$

The rate of this reaction is directly correlated with the overall rate of energy production in the sun, and the uncertainties on this flux in the SSM are therefore low, making the deficits observed by SAGE and GALLEX convincing evidence that the solar neutrino problem is due in some way to disappearance of electron neutrinos from the sun before they reach Earth.

2.3.2 The atmospheric neutrino anomaly

Another, parallel, strand of evidence for disappearance of neutrinos came from the study of atmospheric neutrinos. These neutrinos are produced by cosmic rays striking the atmosphere, and the decay of pions and kaons thus produced. The dominant process is:

$$\begin{aligned} \pi^+ &\to \mu^+ + \nu_\mu \\ &\qquad \hookrightarrow e^+ + \nu_e + \bar{\nu}_\mu \ , \end{aligned} \tag{2.7}$$

plus a similar contribution from π^- decays, and a smaller one from kaons. The neutrinos produced range from around the GeV scale up to the TeV scale.

These atmospheric neutrinos were a significant background to experiments in the 1980s searching for proton decay, and large samples of atmospheric neutrino interactions were taken as by-products. In 1988 the Kamiokande experiment, mentioned previously in relation to the solar neutrino problem, reported that, whilst their observed rate of atmospheric ν_e interactions was in agreement with Monte Carlo expectations, they saw a significant deficit of ν_μ events [35].

This deficit of atmospheric ν_μ was confirmed by the IMB 8 kton water-Čerenkov detector in 1992 [36]. Like Kamiokande, they saw the expected rate of ν_e interactions but only about $\frac{2}{3}$ of the expected ν_μ flux.

Whilst there are large uncertainties in the overall normalization of the expected atmospheric neutrino flux, the $\nu_\mu : \nu_e$ ratio is much better constrained. From (2.7) one would expect a $2:1$ ratio[2]. This is somewhat modified by kaon decays, and uncertainty in the K/π ratio leads to a few percent uncertainty in the $\nu_\mu : \nu_e$ ratio, insufficient to cover the observed discrepancies. This problem became know as the "atmospheric neutrino anomaly".

The situation was somewhat confused by results from the NUSEX and Fréjus

---

[2] When integrated over all energies. Coincidentally, the pion and muon masses are such that this relation also approximately holds at each particular neutrino energy.

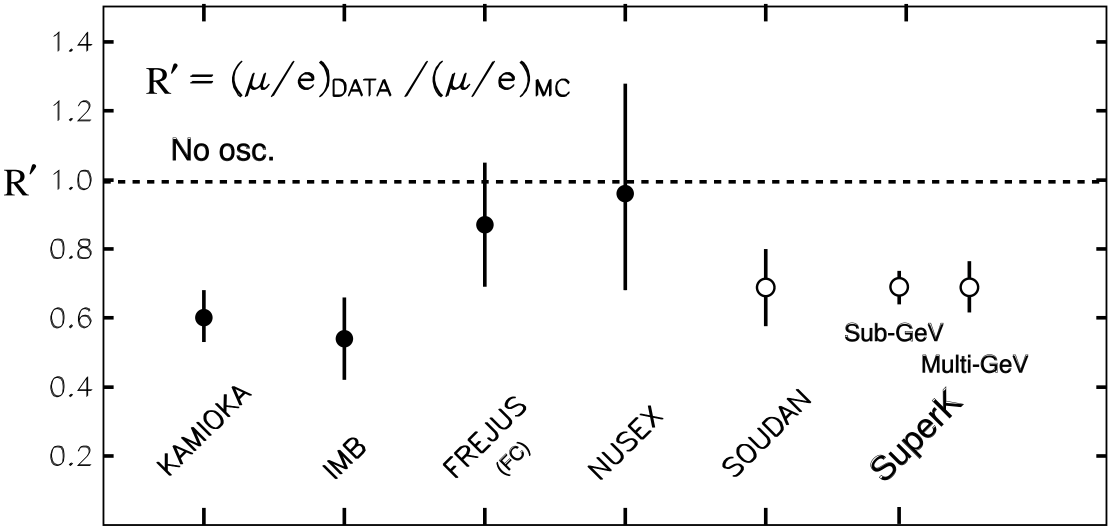

Figure 2.5: Summary of the deficit in atmospheric muon neutrino rates compared to electron neutrinos observed in various experiments. Although the NUSEX and Fréjus results are consistent with no deficit they are also consistent with the consensus of the other experiments $R' \approx \frac{2}{3}$. Figure taken from [37].

experiments, both iron calorimeter tracking detectors searching for proton decay. In contrast to Kamiokande and IMB, they reported no significant deviation of the $\nu_\mu : \nu_e$ ratio from expectations [38, 39]. At the time, doubts were cast on the water-Čerenkov technique of Kamiokande and IMB, but the observation of a comparable deficit by the iron calorimeter Soudan 2 experiment [40] eliminated the possibility that the Čerenkov experiments were simply suffering from some unmodelled systematic effect specific to their design. In the light of more recent measurements, and the large statistical error inherent in all the experiments, all the results are seen to be consistent with one another, and the claim of a ν_μ deficit in atmospheric neutrinos is vindicated (see Figure 2.5).

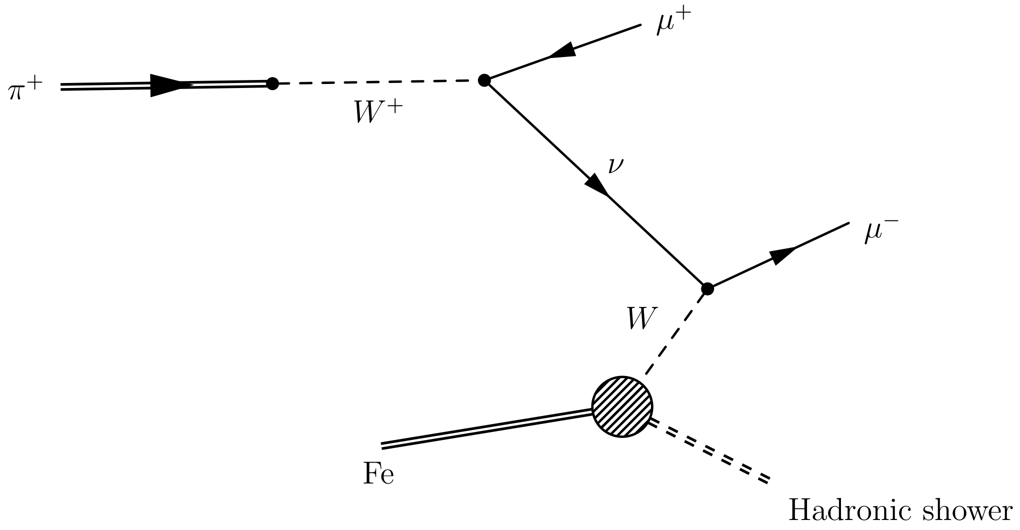

Figure 2.6: Feynman diagram showing the full history of a neutrino in a representative oscillation experiment (in this case MINOS). Here the neutrino is produced, from the decay of a W boson, in association with an anti-muon. After travelling some distance, the neutrino interacts via a charged-current interaction, producing a muon. In this case, the associated leptons at production and absorption are of the same flavour and we may loosely say that no oscillation has occurred. Note that the neutrino forms an internal line, and may be treated as a virtual particle.

2.4 Neutrino oscillation theory

2.4.1 Derivation of the oscillation formula

Of the various hypotheses put forward to explain the deficits in neutrino rates described in Sections 2.3.1 and 2.3.2, only one has stood the test of experiment.

The possibility that neutrinos have non-zero mass and may therefore change state in flight was first suggested by Pontecorvo in 1958 [41] although the proposal differs somewhat from our current understanding, advocating transitions between neutrinos and antineutrinos, rather than between neutrino flavours. The suggestion that electron and muon flavour neutrinos are formed from a mixing of two independent mass states was made by Maki, Nakagawa, and Sakata in 1962 [42].

The specific problem that we are interested in calculating probabilities for, is that of a neutrino being created, travelling some distance and later interacting and being destroyed again. Figure 2.6 shows a concrete example of such a process. In the case of massless neutrinos, lepton-generation number is conserved and the charged leptons associated with the neutrino's production and absorption must have the same flavour. The intermediate neutrino would then also be referred to as being a neutrino of that flavour. In the presence of nonzero neutrino masses, lepton flavour number conservation may be violated, and it is possible for the two associated leptons to have differing flavours. The probability for this to occur will be dependent on the neutrino energy and on the distance travelled.

With the introduction of non-zero neutrino masses, the relative coupling strength of each mass state to each of the charged leptons via the W must be specified. The matrix that defines these relative strengths is called the PMNS matrix, after Pontecorvo, Maki, Nakagawa and Sakata. The matrix element $U^{\star}_{\alpha i}$ gives the relative amplitude for a charged lepton of flavour α to couple to the ith neutrino mass state. For reasons of self-consistency the matrix U must be unitary, *i.e.* $UU^{\dagger} = I$.

The combination of mass states created in association with a lepton of flavour α

$$|\nu_\alpha\rangle = \sum_i U^{\star}_{\alpha i} |\nu_i\rangle \tag{2.8}$$

may be loosely referred to as a neutrino of flavour α.

The overall picture of the oscillation mechanism is straightforward. A lepton of definite flavour is produced in conjunction with a superposition of different neutrino mass states. Due to their differing masses, these propagate at varying speeds and so accumulate a phase difference between them, which increases with the distance travelled. After travelling some distance the neutrino interacts again, each mass state being responsible for creating a superposition of differing associated leptons. Due to the accumulated phase differences this final state is not identical to the initial lepton, and upon detection it may be found to be of a different flavour.

Several subtleties arise in carrying out the calculation itself however, especially in the simplest plane-wave version (as reproduced here). Problems of energy-momentum conservation due to the differing neutrino masses are resolved in principle by noting that the neutrino in all cases is only observed via its interaction products and thus forms a virtual particle where energy and momentum conservation are not required to strictly hold. In practice, a wave-packet or QFT approach handles this automatically. The essential issue with plane-wave calculations is that the plane waves by construction fill all space for all time. There is thus no concept of the neutrino's travel from production to detection, which has to be put in by hand, appealing to some reasonable assumption. This step may fall down, especially in the usual ultra-relativistic limit, due to failure of Lorentz invariance or use of unjustified approximations. These and other issues are covered in [43] for plane waves and [44] in the wave-packet description.

We can write the amplitude $A_{\alpha\beta}$ for an initial flavour α to produce a final flavour β as a sum over all possible intermediate mass states, with the appropriate couplings and phase factors. The different intermediate mass states are indistinguishable[3], so their contributions are summed coherently in the amplitude. Natural units $\hbar = c = 1$ are used throughout.

$$A_{\alpha\beta} = \sum_i U^\star_{\alpha i} e^{-i\tilde{p}_i \cdot \tilde{x}_i} U_{\beta i} \tag{2.9}$$

where $\tilde{p}_i = (E_i, \vec{p}_i)$ is the 4-momentum of mass state i and $\tilde{x}_i$ its position (relative to production at the origin) and we will use a $+---$ metric. Taking the modulus squared to find the transition probability:

$$P_{\alpha\beta} = \sum_{ij} U^\star_{\alpha i} U_{\beta i} U_{\alpha j} U^\star_{\beta j} e^{i(\tilde{p}_j \cdot \tilde{x}_j - \tilde{p}_i \cdot \tilde{x}_i)} \ . \tag{2.10}$$

---

[3] Attempts to distinguish them via their differing times of flight or by carefully observing the details of their production are frustrated by the uncertainty principle.

At this point we need to relate the energies, momenta and positions of the different mass components to each other in order to determine the relative phase accumulated between the components at the point of observation. This phase is the argument of the exponent:

$$\Delta\phi_{ij} = \tilde{p}_j \cdot \tilde{x}_j - \tilde{p}_i \cdot \tilde{x}_i \ . \tag{2.11}$$

Now we make the approximation that all the mass states are sufficiently close to having a common energy E. The same result can also be obtained by assuming equal momenta or equal velocities. Then, calculating the expression at a common time and distance from production L, we obtain an expression in terms of the magnitude of the three-momenta $p = |\vec{p}|$:

$$\Delta\phi_{ij} = p_i L - p_j L = \frac{p_i^2 - p_j^2}{p_i + p_j} L = \frac{m_j^2 - m_i^2}{2p_{\mathrm{avg}}} L \ , \tag{2.12}$$

where $p_{\mathrm{avg}} = \frac{1}{2}(p_i + p_j)$.

Finally we approximate $p_{\mathrm{avg}} \approx E$, justified by the high energies involved in experiments compared to the neutrino masses, and substitute $\Delta\phi_{ij}$ into (2.10).

$$P_{\alpha\beta} = \sum_{ij} U_{\alpha i}^{\star} U_{\beta i} U_{\alpha j} U_{\beta j}^{\star} \exp\left(i \frac{\Delta m_{ji}^2 L}{2E} \right) \tag{2.13}$$

with the notation $\Delta m_{ji}^2 = m_j^2 - m_i^2$. Expanding out the exponential using Euler's formula, and using the identity $\cos 2\theta = 1 - 2\sin^2\theta$:

$$\begin{aligned} P_{\alpha\beta} \ &= \ \sum_{ij} \Re(U_{\alpha i}^{\star} U_{\beta i} U_{\alpha j} U_{\beta j}^{\star}) \left(1 - 2\sin^2\left(\frac{\Delta m_{ji}^2 L}{4E} \right) \right) \\ &\quad - \sum_{ij} \Im(U_{\alpha i}^{\star} U_{\beta i} U_{\alpha j} U_{\beta j}^{\star}) \sin\left(\frac{\Delta m_{ji}^2 L}{2E} \right) \ . \end{aligned} \tag{2.14}$$

Using the property of unitary matrices that the rows form an orthonormal basis $\sum_i U_{\alpha i} U^{\star}_{\beta i} = \delta_{\alpha\beta}$ we arrive at the final transition probability:

$$\begin{aligned} P_{\alpha\beta} &= \delta_{\alpha\beta} - 4 \sum_{i>j} \Re(U^{\star}_{\alpha i} U_{\beta i} U_{\alpha j} U^{\star}_{\beta j}) \sin^2\left(\frac{\Delta m^2_{ij} L}{4E}\right) \\ &\quad + 2 \sum_{i>j} \Im(U^{\star}_{\alpha i} U_{\beta i} U_{\alpha j} U^{\star}_{\beta j}) \sin\left(\frac{\Delta m^2_{ij} L}{2E}\right) . \end{aligned} \tag{2.15}$$

If $\mathcal{CP}$ symmetry is obeyed in neutrino oscillations then $P_{\bar{\alpha}\bar{\beta}} \equiv P_{\alpha\beta}$, where the bars on the indices represent oscillations between antineutrinos. If $\mathcal{CP}$ symmetry is broken then $\mathcal{CPT}$ symmetry still implies $P_{\bar{\alpha}\bar{\beta}} \equiv P_{\beta\alpha}$ and thus $P_{\bar{\alpha}\bar{\alpha}} \equiv P_{\alpha\alpha}$. That is, survival probabilities should be identical between neutrinos and antineutrinos.

2.4.2 The two flavour approximation

In a world with only two neutrino generations the oscillation formulae simplify considerably. The mixing matrix can be expressed in terms of a single angle[4]:

$$U = \begin{bmatrix} \cos\theta & \sin\theta \\ -\sin\theta & \cos\theta \end{bmatrix} . \tag{2.16}$$

Substituting into equation (2.15), and making the conversion to practical units yields a much-simplified expression for the survival probability.

$$P_{\alpha\alpha} = 1 - \sin^2 2\theta \sin^2\left(1.267 \frac{\Delta m^2\,[\mathrm{eV}^2] L\,[\mathrm{km}]}{E\,[\mathrm{GeV}]}\right) . \tag{2.17}$$

Alternatively L may be measured in metres and E in MeV. The oscillatory term is an even function of Δm^2, so investigation of two-flavour oscillations of this form can only give information about the magnitude of the mass splitting, with no information about which of the two mass states is the heavier.

The assumption of two neutrino flavours is obviously unrealistic, but the sim-

---

[4] Technically there is also a $\mathcal{CP}$-violating Majorana phase, which we will gloss over here.

plification to only two oscillation parameters (Δm^2 and $\sin^2 2\theta$) and the survival probability (2.17) does have relevance to the real world. In the case of a neutrino mass spectrum where one Δm^2 is much larger than the others, and an experiment whose baseline is chosen such that, for all the small splittings, $\frac{\Delta m^2 L}{E} \ll 1$ the general oscillation formula (2.15) simplifies [45] to give a survival probability:

$$P_{\alpha\alpha} = 1 - 4\left(\sum_{i \text{ bunch}} |U_{\alpha i}|^2\right)\left(1 - \sum_{i \text{ bunch}} |U_{\alpha i}|^2\right)\sin^2\left(\frac{\Delta m^2 L}{4E}\right) \quad , \tag{2.18}$$

where the sums are over all the states below the large mass splitting, or alternatively all those above. In this case the survival probability can be parameterized in terms of two quantities: the largest mass splitting, and an effective mixing angle. The neutrino mass spectrum found in nature does indeed exhibit one dominant mass splitting (see Section 2.8), and the MINOS experiment is tuned to the relevant L/E, so the simple survival probability (2.17) may be adopted for the analysis of the observed muon-neutrino deficit.

2.4.3 Three flavour oscillations

Motivated by the simplified oscillation formula obtained in the case of two neutrino generations, where the mixing is parametrized by a single angle, the full PMNS matrix is often expressed in terms of angles[5]:

$$
\begin{aligned}
U &= \begin{bmatrix} 1 & 0 & 0 \\ 0 & c_{23} & s_{23} \\ 0 & -s_{23} & c_{23} \end{bmatrix} \begin{bmatrix} c_{13} & 0 & s_{13}e^{-i\delta} \\ 0 & 1 & 0 \\ -s_{13}e^{i\delta} & 0 & c_{13} \end{bmatrix} \begin{bmatrix} c_{12} & s_{12} & 0 \\ -s_{12} & c_{12} & 0 \\ 0 & 0 & 1 \end{bmatrix} \qquad (2.19) \\
&= \begin{bmatrix} c_{12}c_{13} & s_{12}s_{13} & s_{13}e^{-i\delta} \\ -s_{12}c_{23} - c_{12}s_{23}s_{13}e^{i\delta} & c_{12}c_{23} - s_{12}s_{23}s_{13}e^{i\delta} & s_{23}c_{13} \\ s_{12}s_{23} - c_{12}c_{23}s_{13}e^{i\delta} & -c_{12}s_{23} - s_{12}c_{23}s_{13}e^{i\delta} & c_{23}c_{13} \end{bmatrix}, \qquad (2.20)
\end{aligned}
$$

where $s_{ij} \equiv \sin\theta_{ij}$ and $c_{ij} \equiv \cos\theta_{ij}$ and the convention for the decomposition is that adopted by the Review of Particle Physics [46]. For historical reasons the angle θ_{23} may be referred to as the atmospheric mixing angle, and the corresponding mass-splitting Δm^2_{32} as $\Delta m^2_{\mathrm{atm}}$. Similarly, θ_{12} is known as the solar mixing angle and the mass splitting Δm^2_{12} can be written $\Delta m^2_{\mathrm{sol}}$ or $\Delta m^2_{\odot}$. θ_{13} is the only angle not to have been experimentally verified to be non-zero and is the focus of several current and future experiments. δ is the $\mathcal{CP}$-violating phase.

The most experimentally relevant channel which exhibits all the features of three-flavour oscillations is the appearance of electron-neutrinos in a muon-neutrino beam. The transition probability $P_{\mu e}$ can be written:

$$
P_{\mu e} = \left| U^{\star}_{\mu 1} e^{-im_1^2 L/2E} U_{e1} + U^{\star}_{\mu 2} e^{-im_2^2 L/2E} U_{e2} + U^{\star}_{\mu 3} e^{-im_3^2 L/2E} U_{e3} \right|^2 . \qquad (2.21)
$$

---

[5] The two Majorana phases α_1 and α_2 are not observable in oscillation experiments and have been omitted in this expression. If they are included an additional matrix $\mathrm{diag}(\exp(i\alpha_1/2), \exp(i\alpha_2/2), 1)$ is included in the product.

Using the orthogonality of the rows of U to eliminate the first term:

$$P_{\mu e} = \left| U^{\star}_{\mu 2} U_{e2} \left(e^{-i\frac{m_2 L}{2E}} - e^{-i\frac{m_1 L}{2E}} \right) + U^{\star}_{\mu 3} U_{e3} \left(e^{-i\frac{m_3 L}{2E}} - e^{-i\frac{m_1 L}{2E}} \right) \right|^2 \tag{2.22}$$

$$= \left| 2iU^{\star}_{\mu 2} U_{e2} \sin\left(\frac{\Delta m^2_{21} L}{4E}\right) e^{-i\frac{(m_1+m_2)L}{2E}} \right. \\ \left. + \, 2iU^{\star}_{\mu 3} U_{e3} \sin\left(\frac{\Delta m^2_{31} L}{4E}\right) e^{-i\frac{(m_1+m_3)L}{2E}} \right|^2 \tag{2.23}$$

$$= \left| 2U^{\star}_{\mu 3} U_{e3} \sin\left(\frac{\Delta m^2_{31} L}{4E}\right) e^{-i\frac{\Delta m^2_{32} L}{4E}} + 2U^{\star}_{\mu 2} U_{e2} \sin\left(\frac{\Delta m^2_{21} L}{4E}\right) \right|^2 . \tag{2.24}$$

Inserting the definition of U from (2.19), and approximating $\theta_{13} \ll 1$ in the second term, the probability becomes

$$P_{\mu e} \approx \left| \sqrt{P_{\mathrm{atm}}} e^{-i(\Delta m^2_{32} L/4E \pm \delta)} + \sqrt{P_{\mathrm{sol}}} \right|^2 \tag{2.25}$$

$$= P_{\mathrm{atm}} + 2\sqrt{P_{\mathrm{atm}} P_{\mathrm{sol}}} \cos\left(\frac{\Delta m^2_{32} L}{4E} \pm \delta\right) + P_{\mathrm{sol}} \; , \tag{2.26}$$

where

$$\sqrt{P_{\mathrm{atm}}} = \sin\theta_{23} \sin 2\theta_{13} \sin\left(\frac{\Delta m^2_{31} L}{4E}\right) \tag{2.27}$$

$$\sqrt{P_{\mathrm{sol}}} = \cos\theta_{23} \sin 2\theta_{12} \sin\left(\frac{\Delta m^2_{21} L}{4E}\right) \; . \tag{2.28}$$

That is, the oscillation occurs over two different characteristic length scales, atmospheric and solar. This general behaviour is illustrated in Figure 2.7. The effect of $\mathcal{CP}$-violation is seen in the cross term. Between them P_{atm} and P_{sol} contain sines of all the mixing angles, thus the value of δ is only observable if all these angles are non-zero. $\theta_{13} = 0$ remains an experimental possibility and so discovery of a non-zero θ_{13} is a prerequisite to studying $\mathcal{CP}$-violation in the lepton sector. This is the reason why δ is grouped with θ_{13} in (2.19).

Note that if any of the angles are zero then either $P_{\mathrm{atm}} = 0$ or $P_{\mathrm{sol}} = 0$ and the oscillation reduces to the two flavour form (2.17) (with a suitable definition of

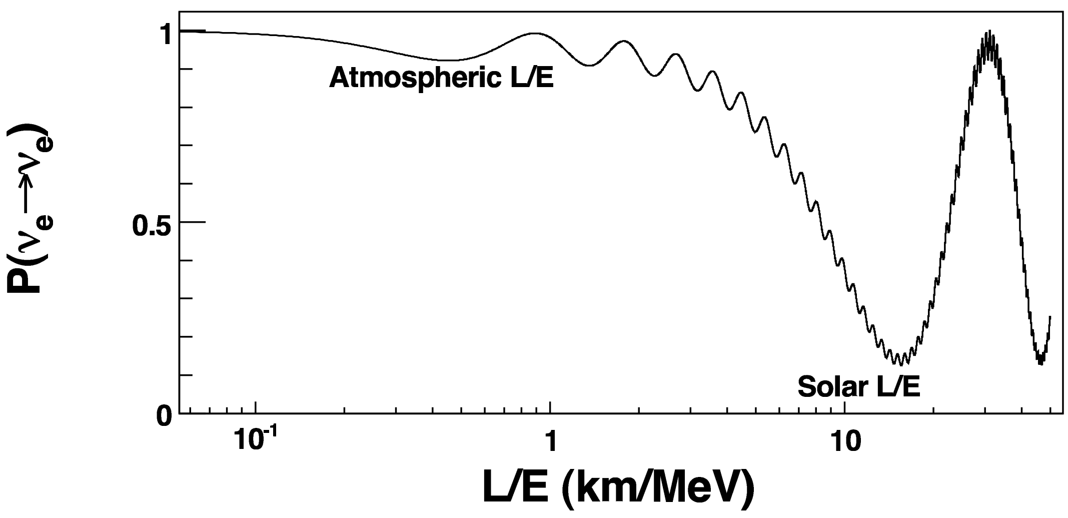

Figure 2.7: Survival probability for an electron neutrino as a function of L/E. The oscillation occurs over two distinct baselines. The first minimum at the atmospheric baseline is that probed by reactor neutrino experiments such as CHOOZ, whereas the first solar minimum has been measured by KAMLAND. Figure taken from [47].

the angle θ in terms of the two remaining angles). This is a general property of three-neutrino oscillations. If any angle is small then all oscillation probabilities can be approximated by a two-flavour form. In practice θ_{13} plays the role of small angle, simplifying the analysis of data from many experiments.

2.4.4 Matter effects

Neutrinos traveling through matter undergo a modification of their oscillatory behaviour as a result of their interactions with it. This phenomenon was first noted by Wolfenstein [48] and is known as the Mikheev-Smirnov-Wolfenstein (MSW) effect. It is of practical interest in the analysis of solar neutrino oscillations and for the future long-baseline experiments NOνA and LBNE.

First we must recast the general treatment of neutrino oscillations given in Section 2.4.1 into a more suitable form. The general expression for the transition amplitude $A_{\alpha\beta}$ (Equation (2.15)) can also be expressed in terms of a Schrödinger

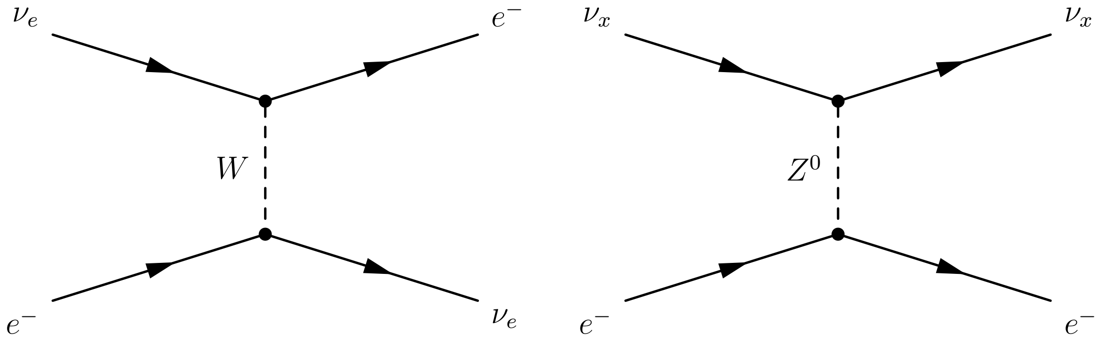

Figure 2.8: Coherent forward scattering of neutrinos from matter. Left: an electron neutrino scattering from an electron via a charged-current interaction. Right: a neutrino of any flavour scattering from an electron via a neutral-current interaction.

equation:

$$i\frac{\mathrm{d}}{\mathrm{d}L}A_{\alpha\beta}(L) = \sum_{i\gamma} U_{\beta i}\frac{m_i^2}{2E}U^\star_{\gamma i}A_{\alpha\gamma}(L)\ . \tag{2.29}$$

For the two-flavour case, choosing the flavours to be e and μ, and beginning in a ν_e state, neglecting an overall phase, this becomes

$$i\frac{\mathrm{d}}{\mathrm{d}L}\begin{bmatrix} A_{ee} \\ A_{e\mu} \end{bmatrix} = \frac{1}{4E}\begin{bmatrix} -\Delta m^2\cos 2\theta & \Delta m^2 \sin 2\theta \\ \Delta m^2 \sin 2\theta & \Delta m^2\cos 2\theta \end{bmatrix}\begin{bmatrix} A_{ee} \\ A_{e\mu} \end{bmatrix}\ . \tag{2.30}$$

For neutrinos traveling through matter there is a contribution to the Hamiltonian from coherent forward scattering from protons, electrons, and neutrons. The two possible diagrams for scattering from an electron are shown in Figure 2.8. Only electron neutrinos can interact via the charged-current in this way, while all flavours of neutrino can undergo neutral-current interactions. The contribution from neutral-current scattering is the same for all neutrino flavours, contributes nothing to the phase difference that drives oscillations, and is thus ignored. The contribution from charged-current interactions is

$$\mathcal{H}_{\mathrm{CC}} = \pm\sqrt{2}G_{\mathrm{F}}N_e\ , \tag{2.31}$$

where N_e is the number density of electrons in the medium and the positive sign is for neutrinos, the negative for antineutrinos.

The Schrödinger equation (2.30) then becomes

$$i\frac{\mathrm{d}}{\mathrm{d}L}\begin{bmatrix} A_{ee} \\ A_{e\mu} \end{bmatrix} = \frac{1}{4E}\begin{bmatrix} -\Delta m^2\cos 2\theta \pm 4\sqrt{2}EG_{\mathrm{F}}N_e & \Delta m^2 \sin 2\theta \\ \Delta m^2\sin 2\theta & \Delta m^2\cos 2\theta \end{bmatrix}\begin{bmatrix} A_{ee} \\ A_{e\mu} \end{bmatrix}. \tag{2.32}$$

Diagonalizing the Hamiltonian matrix we obtain

$$\mathcal{H} = \frac{1}{4E}\begin{bmatrix} \cos\theta_M & -\sin\theta_M \\ \sin\theta_M & \cos\theta_M \end{bmatrix}\begin{bmatrix} -\Delta m_M^2 & 0 \\ 0 & \Delta m_M^2 \end{bmatrix}\begin{bmatrix} \cos\theta_M & \sin\theta_M \\ -\sin\theta_M & \cos\theta_M \end{bmatrix}, \tag{2.33}$$

with

$$\Delta m_M^2 = \sqrt{(\Delta m^2\cos 2\theta \mp 2\sqrt{2}EG_{\mathrm{F}}N_e)^2 + (\Delta m^2\sin 2\theta)^2} \tag{2.34}$$

$$\tan 2\theta_M = \frac{\tan 2\theta}{1 \mp \frac{2\sqrt{2}EG_{\mathrm{F}}N_e}{\Delta m^2\cos 2\theta}}\,. \tag{2.35}$$

That is, the effect of the matter potential causes the neutrino oscillations to occur with effective oscillation parameters Δm_M^2 and θ_M which are formed by modifying the underlying parameters according to the neutrino energy and the matter density.

These equations contain an interesting resonance effect, first noted by Mikheev and Smirnov [49]. Around a certain electron density

$$N_e^{\mathrm{res}} = \frac{\Delta m^2\cos 2\theta}{2\sqrt{2}EG_{\mathrm{F}}} \tag{2.36}$$

the mixing becomes maximal at $\theta_M = 45°$, no matter how small the initial θ, allowing a much greater conversion between neutrino flavours (in this case from ν_e to ν_μ) than would otherwise be possible.

2.4.5 Neutrino decay and decoherence

In addition to neutrino oscillations driven by mass splittings, the phenomena of neutrino decay and decoherence have been proposed to explain observed neutrino deficits, either in whole or in part [50, 51].

In neutrino decay models, one of the neutrino states is taken to have a finite lifetime τ_0. This has the effect of including an additional term $\exp(-\tau/2\tau_0)$ in the phase factor in (2.9). In a two flavour model the survival probability then becomes

$$\begin{aligned} P_{\mu\mu} &= \sin^4\theta + \cos^4\theta\exp(-\alpha L/E) \\ &\qquad + 2\sin^2\theta\cos^2\theta\exp(-\alpha L/2E)\cos\left(\frac{\Delta m^2 L}{2E}\right) , \end{aligned} \tag{2.37}$$

where $\alpha = m/\tau_0$ for the decaying state. In the limit of stable neutrinos ($\tau_0 \to \infty$) this reduces to the usual two-flavour oscillation formula (2.17), and in the limit of small or zero mass splitting ($\Delta m^2 \to 0$) to a simpler form describing disappearance driven purely by neutrino decay:

$$P_{\mu\mu} = \left(\sin^2\theta + \cos^2\theta\exp\left(-\alpha L/2E\right)\right)^2 . \tag{2.38}$$

In decoherence models, the disappearance of one flavour is taken to be due to a "foamy" or "fuzzy" spacetime background, which makes path lengths indeterminate and so gradually destroys the phase relationship between the mass states. General arguments lead to a survival probability of the form

$$P_{\mu\mu} = 1 - \frac{1}{2}\sin^2 2\theta\left(1 - \exp(-\mu^2 L/2E)\cos\left(\frac{\Delta m^2 L}{2E}\right)\right) . \tag{2.39}$$

Again, taking the decoherence parameter μ to zero recovers the standard two-flavour oscillation form, and $\Delta m^2 \to 0$ yields pure decoherence:

$$P_{\mu\mu} = 1 - \frac{1}{2}\sin^2 2\theta \left(1 - \exp(-\mu^2 L/2E)\right) \ . \tag{2.40}$$

Initially the pure forms (2.38) and (2.40) seemed to provide good fits to the atmospheric data, but increased statistics has led to them being ruled out at high significance. Some sufficiently small admixture of either effect with the usual mass-driven oscillations will always remain a possibility, although upper limits can be set on the parameters α and μ^2.

2.5 Atmospheric neutrino oscillations

2.5.1 Super-Kamiokande

The Super-Kamiokande experiment is the successor to Kamiokande. It consists of 50 kton of pure water observed by 11,000 photomultiplier tubes. In 1998 the collaboration reported strong evidence for neutrino oscillations in the form of an up-down asymmetry in their observed ν_μ deficit [52, 53]. A significant deficit is observed in up-going muon events, where the neutrino has had to cross the entire diameter of the Earth, whereas no such deficit is observed for down-going muon events, where the neutrino has travelled only a few kilometres from its production point. Interpreted in terms of neutrino oscillations, these observations allow limits to be placed on the parameters governing atmospheric neutrino oscillations, which are found to occur over a baseline of approximately 500 km/GeV, and with near-maximal mixing. The absence of any deviation from expectations for electron-like events indicates that the oscillations must be $\nu_\mu \leftrightarrow \nu_\tau$. Updated results with a larger period of data taking are reported in [54] and the zenith angle distributions for different event classes shown in Figure 2.9.

An alternate analysis [55] of the Super-Kamiokande data bins the observed

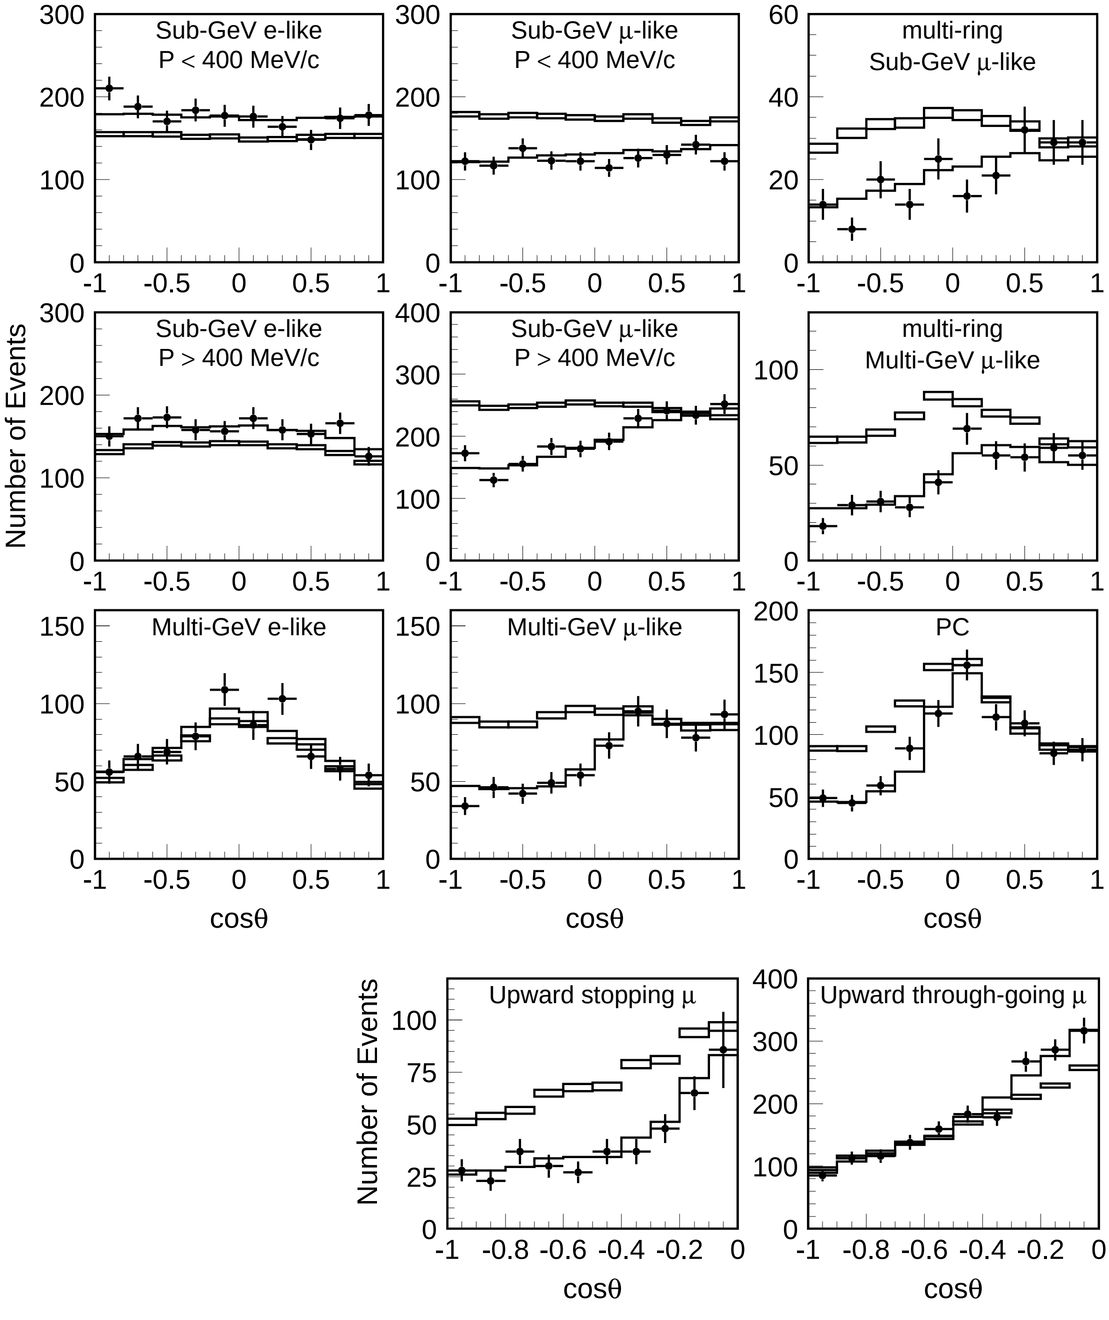

Figure 2.9: Comparison of zenith angle distributions of different event classes in Super-Kamiokande between data and Monte Carlo. The open boxes are the expectation with systematic error, the points are data and the line is the best oscillation fit. A deficit is visible in upward-going muon samples but not in electrons. Taken from [54].

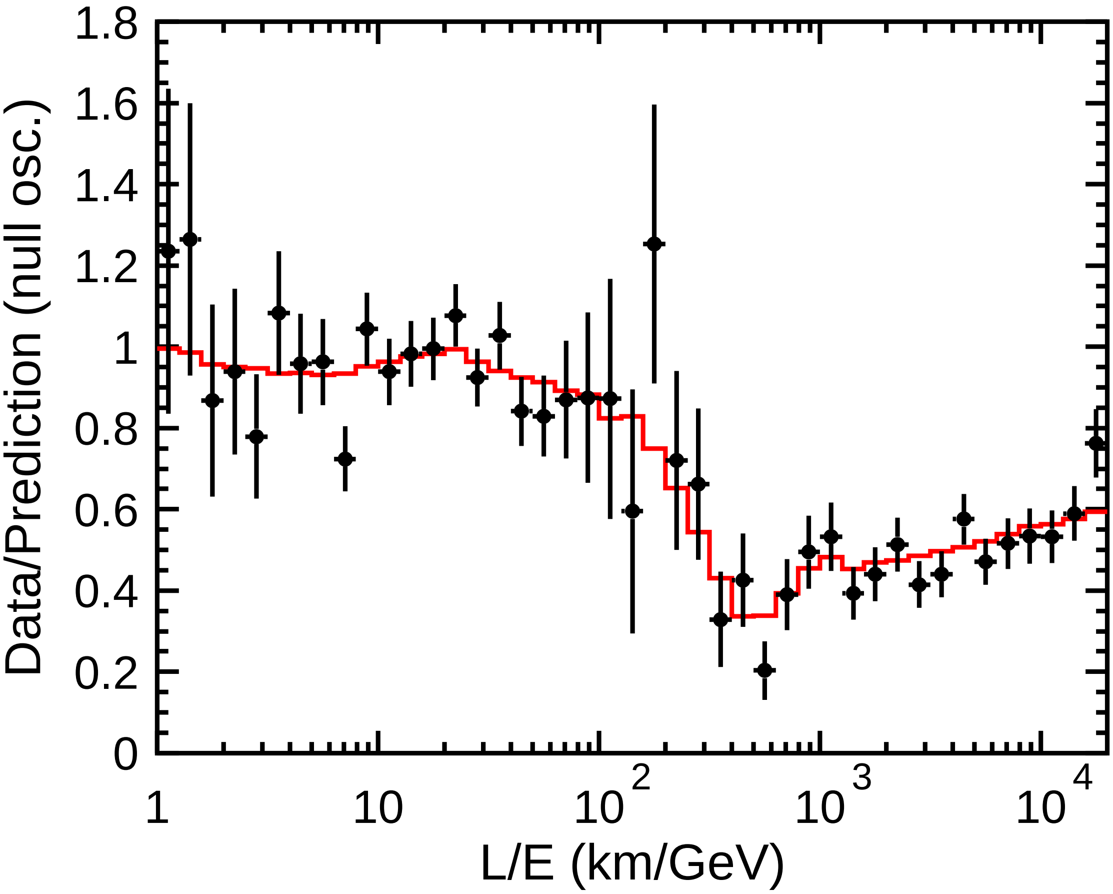

Figure 2.10: Ratio of muon-like events to expectation in Super-Kamiokande as a function of L/E. The red line indicates the best fit to two-flavour $\nu_\mu \leftrightarrow \nu_\tau$ oscillations. Taken from [55].

events by their estimated L/E value, having first selected only those events whose L/E is expected to be well estimated. This method allows the characteristic oscillation dip to be made out (Figure 2.10). The confidence limits obtained on the oscillation parameters are complementary to the standard zenith angle analysis, obtaining somewhat tighter limits in $\Delta m^2_{\mathrm{atm}}$ but a weaker limit on $\sin^2 2\theta_{\mathrm{atm}}$, as shown in Figure 2.11. The best fit parameters obtained are (at 90% confidence level):

$$1.9 \times 10^{-3}\,\mathrm{eV}^2 < |\Delta m^2_{\mathrm{atm}}| < 3.0 \times 10^{-3}\,\mathrm{eV}^2 \tag{2.41}$$

$$\sin^2 2\theta_{\mathrm{atm}} > 0.90\ . \tag{2.42}$$

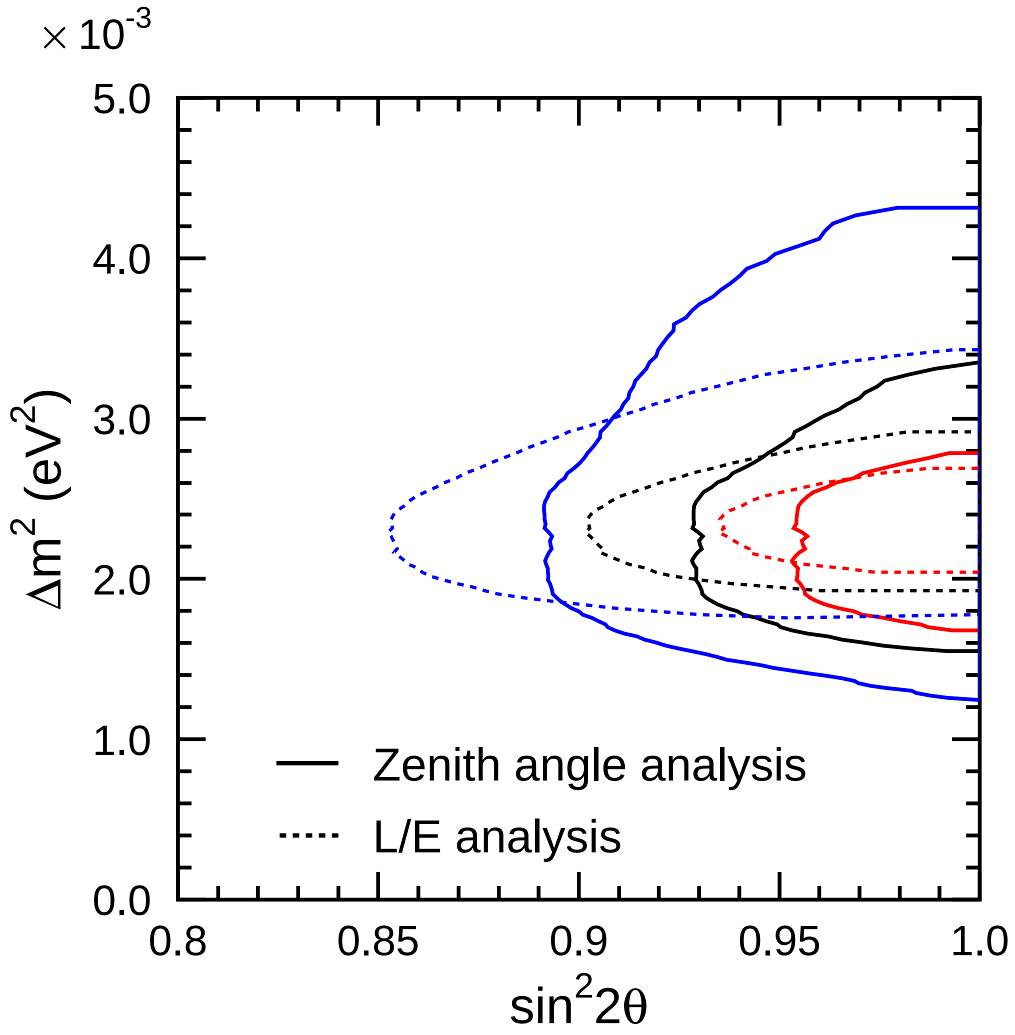

Figure 2.11: 68%, 90% and 99% (red, black and blue respectively) confidence contours for the atmospheric neutrino parameters from Super-Kamiokande. Solid lines indicate the zenith-angle analysis, dotted lines the L/E analysis. Taken from [54].

The rate of atmospheric neutrinos as a function of zenith angle has also been analyzed by Soudan 2 [56], MACRO [57], and MINOS [58, 59]. All favour $\nu_\mu \leftrightarrow \nu_\tau$ oscillations with parameters consistent with those found by Super-Kamiokande, though with significantly less sensitivity.

2.5.2 Accelerator neutrino experiments

Confirmation of neutrino oscillations with the atmospheric parameters, independent of detailed knowledge of the atmospheric neutrino flux, was provided by the

K2K experiment [60]. A beam of muon neutrinos was produced at KEK from a 12 GeV proton beam striking an aluminium target. This beam was directed towards the Super-Kamiokande detector 250 km away. Additional detectors, including a 1 kton water-Čerenkov detector, were placed 300 m downstream of the target in order to measure the beam properties before oscillations had occurred. Accelerator experiments have the advantage of knowing the baseline over which oscillations occur precisely, whilst atmospheric experiments must estimate this for each event based on the reconstructed neutrino direction. Results of the oscillation analysis were once again consistent with Super-Kamiokande's atmospheric fits, but with lower sensitivity.

The MINOS experiment uses the same technique of a man-made neutrino beam from an accelerator, and widely separated "near" and "far" detectors to reduce systematic uncertainties, to study $\nu_\mu \leftrightarrow \nu_\tau$ oscillations at the atmospheric baseline. The extraction of the neutrino oscillation parameters from MINOS data is the focus of later chapters in this thesis.

One other long-baseline accelerator neutrino experiment has thus far reported results. The OPERA detector was built specifically to test the hypothesis that the disappearing muon neutrinos oscillate to ν_τ. While the lack of a ν_e excess in atmospheric data rules out $\nu_\mu \leftrightarrow \nu_e$, in principle the possibility of decay or oscillation to an invisible particle is left open. A muon neutrino beam is sent from CNGS at CERN 730 km to the Gran Sasso laboratory where OPERA is located. The detector consists of 1300 tons of bricks, each containing a sandwich of lead plates and photographic emulsion. The extremely high spatial resolution allows detection of the "kink" in a track as a tau lepton decays to a muon, confirming that the incident particle was indeed a ν_τ. In 2010 OPERA reported the first tau-candidate event [61].

2.6 Solar neutrino oscillations

2.6.1 SNO

Confirmation that the flux prediction of the Standard Solar Model is correct, and that the solar neutrino problem is indeed also caused by oscillations between neutrino flavours, eventually came from the Sudbury Neutrino Observatory (SNO). The experiment consists of 1000 tons of heavy water (D_2O) in a spherical vessel, observed by photomultiplier tubes. The feature of SNO that allows it to make a definitive statement about the origin of the solar neutrino deficit is its sensitivity to solar neutrinos via three separate processes:

- Charged-current: $\nu_e + d \rightarrow p + p + e^-$. This reaction is sensitive only to electron neutrinos.
- Elastic scattering: $\nu_x + e^- \rightarrow \nu_x + e^-$. This reaction is sensitive to all flavours, but predominantly ν_e due to the additional s-channel diagram available to them.
- Neutral-current: $\nu_x + d \rightarrow p + n + \nu_x$. This reaction is sensitive to all neutrino flavours equally.

Sensitivity to all neutrino flavours via neutral-current interactions means that SNO is able to make a measurement of the total solar neutrino flux independent of any oscillations that may have occurred.

In the "salt phase" from 2001 to 2003, 2 tons of NaCl were added to the heavy water. This provides three advantages in the detection of the critical neutral-current interactions. The neutron capture cross-section is larger on chlorine, increasing the event rate; the photons emitted in neutron capture on chlorine are more energetic than those from deuterium, raising them well above the low energy backgrounds; and the isotropy of the several photons emitted contrasts well with

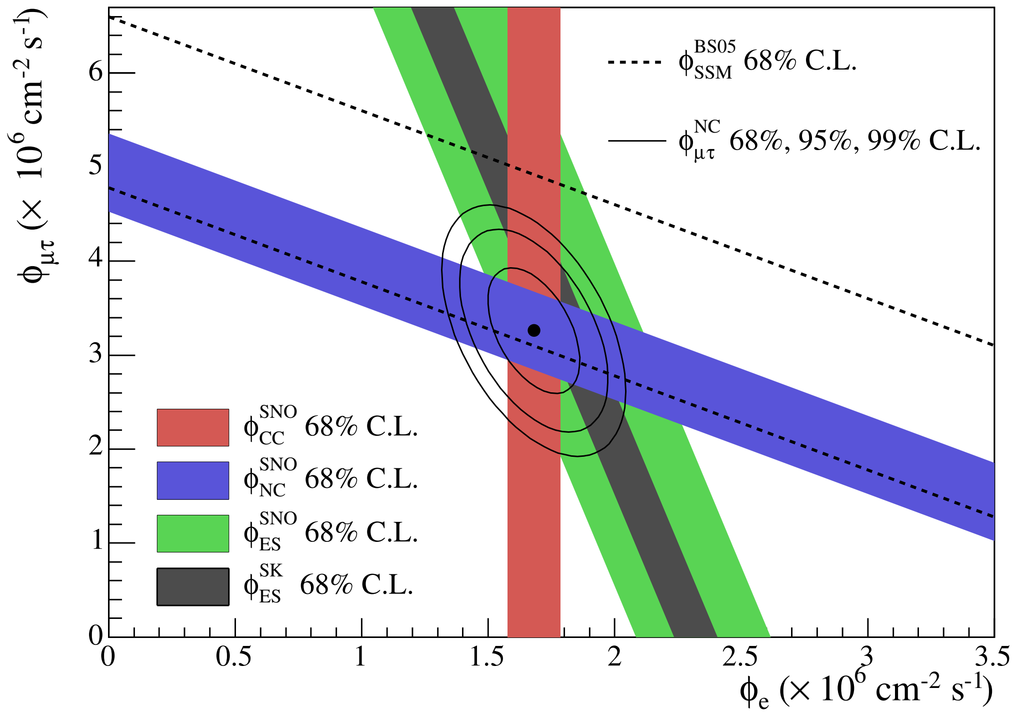

Figure 2.12: Measurement of the solar neutrino flux by SNO [62]. The x-axis represents the flux of electron flavour neutrinos and the y-axis the combined flux of μ and τ flavour neutrinos. The coloured bands indicate the constraints placed by the different event categories. The grey band is a comparable measurement of elastic scattering by Super-Kamiokande. The best fit flux is marked along with 68%, 95% and 99% confidence levels. The best fit is consistent with the SSM expectation (dotted lines) and requires a non-zero flux of non-electron neutrinos.

the single electron of the other two reactions, allowing for better separation of the NC signal.

Results from the SNO salt phase showed the same deficit of ν_e events as observed in previous experiments, but crucially the rate of NC interactions was consistent with the overall neutrino flux predicted by the SSM [62]. A combined fit to the rates of the three processes yields a non-zero flux of non-electron-flavour neutrinos and an overall rate consistent with the SSM prediction (Figure 2.12). The only conclusion is that the overall ν_e yield predicted by solar models is correct, but that approximately $\frac{2}{3}$ of these convert into ν_μ, ν_τ, or both, before reaching Earth.

2.6.2 Interpretation in terms of oscillations

Analysis of solar neutrino results is complicated by the possibility of matter effects in the sun, but the data now favour the "LMA" (large mixing angle) solution, where the oscillations are driven by an adiabatic passage through the MSW resonance within the sun. Oscillations in vacuum between the sun and the Earth are disfavoured due to the absence of a seasonal variation in the solar neutrino flux (as would be expected from the varying baseline caused by the eccentricity of the Earth's orbit) beyond simple $\frac{1}{r^2}$ scaling. The absence of any observable day-night variations (caused by matter effects as neutrinos pass through the Earth at night) rules out a region around $\Delta m^2_\odot \cos 2\theta_\odot \sim 10^{-6}\,\mathrm{eV}^2$.

In the LMA solution the density at the centre of the sun is far above the critical density ($N_e \gg N_e^{\mathrm{res}}$) and so from Equation (2.35) θ_M is nearly 90°. This means that at their production point the solar ν_e are almost entirely ν_2. The density in the sun varies smoothly over a long scale so the resonance is passed adiabatically and the neutrinos remain in the ν_2 state throughout. Exiting the sun into vacuum the state is a mixture of electron and other flavours $|\nu_2\rangle = \sin\theta_\odot\,|\nu_e\rangle + \cos\theta_\odot\,|\nu_x\rangle$, and so the fraction of the expected ν_e detected on earth is $\sin^2\theta_\odot$ (*i.e.* SNO measures $\sin^2\theta_\odot \sim \frac{1}{3}$ or $\tan^2\theta_\odot \sim \frac{1}{2}$). The constraints placed on $\Delta m^2_\odot$ are not so tight, requiring only that the resonance effect happens gradually somewhere between the core and surface of the sun. Note however that this solution only works if $\mathrm{sgn}(\Delta m^2_\odot) = +1$ since we take the negative sign in (2.35) as both the sun and the neutrinos are matter, and $\cos 2\theta_\odot \sim \frac{1}{3}$. By convention the heavier mass state is labelled ν_2, and the lighter ν_1.

2.6.3 KamLAND

The solar neutrino oscillation parameters have also been probed by one terrestrial experiment. KamLAND was a liquid scintillator detector observing electron-

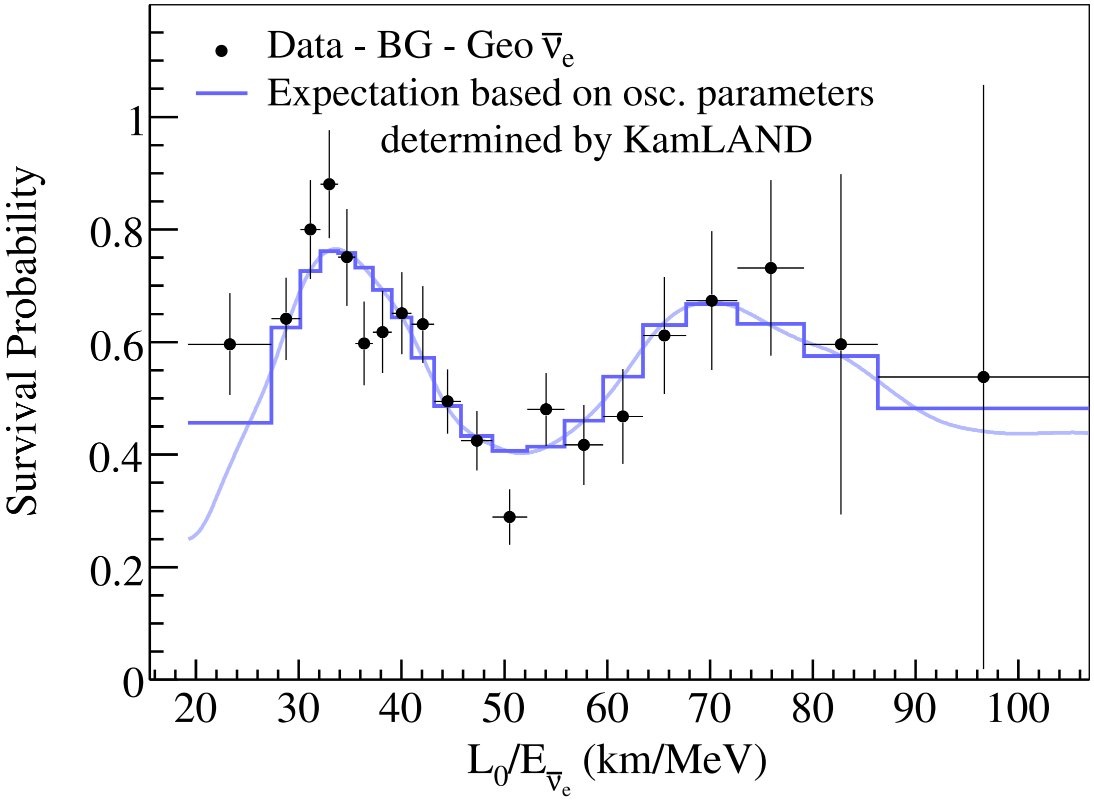

Figure 2.13: KamLAND ratio of observed $\bar{\nu}_e$ events to expectation as a function of L/E (L_0 is simply the flux-weighted average over all the reactors). The oscillation dip is clearly visible, and a fit to the two-flavour oscillation formula (2.17) is shown. Taken from [63].

antineutrinos produced by Japan's 53 nearby nuclear power stations. The average baseline is 180 km and KamLAND's energy threshold is 1.8 MeV making it sensitive to oscillations driven by the LMA mass splitting favoured by the solar experiments. In terms of the three-flavour oscillation probabilities (Figure 2.7) KamLAND sits at the bottom of the "solar L/E" dip and the smaller magnitude and rapidly oscillating behaviour of the atmospheric L/E term make it neglectable in the analysis. Plotting the ratio of observed $\bar{\nu}_e$ events to that expected as a function of their energy (Figure 2.13), a clear oscillatory pattern is observed [63]. A fit to the position and depth of this characteristic oscillation dip provides a measurement of the solar oscillation parameters.

Due to its observation of the detailed shape of the oscillation dip as a function

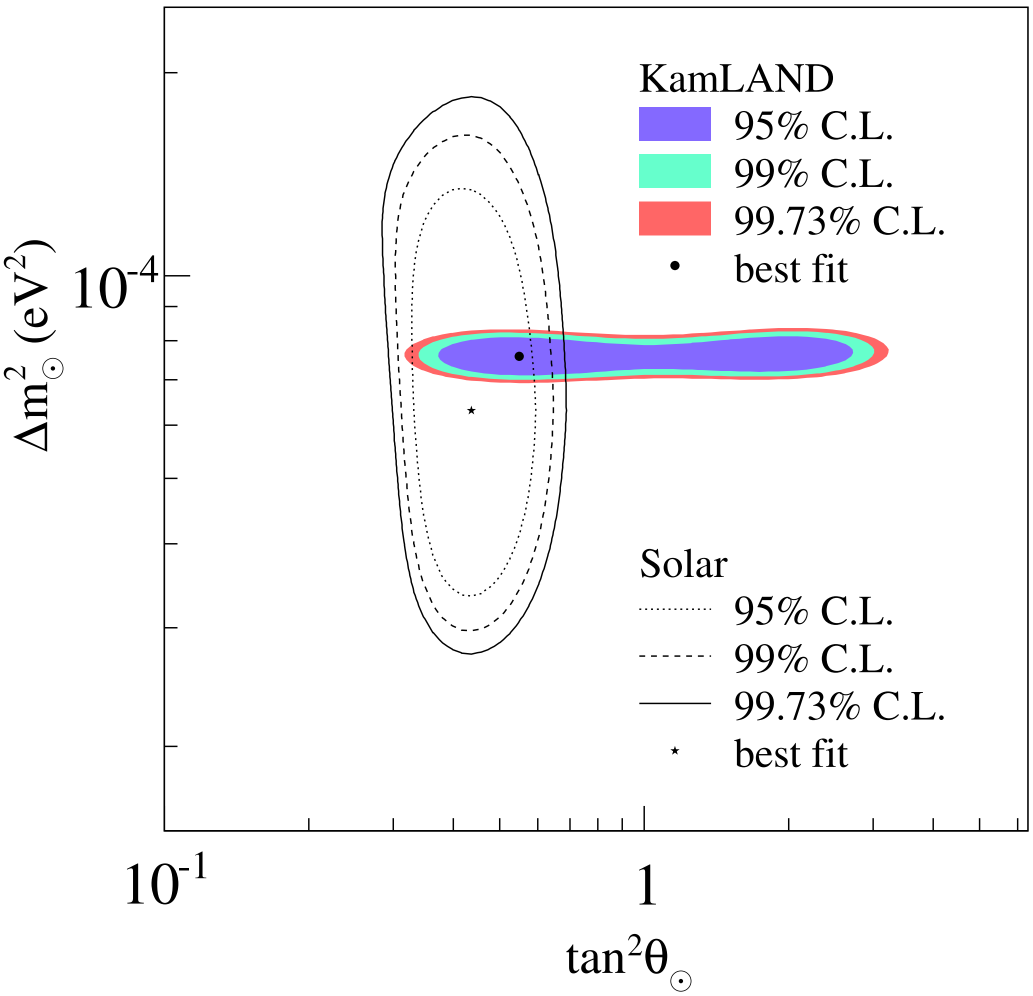

Figure 2.14: Measurements of the solar neutrino parameters from SNO and radiochemical experiments (black lines) and KamLAND (coloured regions). The measurements are consistent with each other and provide complementary information. Taken from [63].

of energy, KamLAND achieves much higher precision in its measurement of the mass splitting than the solar experiments which see only an overall rate. However, KamLAND does suffer from an ambiguity in the quadrant of $\theta_\odot$ which is not shared by solar experiments. The complementarity of these experiments is illustrated by the contours in Figure 2.14 and a combined fit to all experiments [63] gives:

$$\Delta m^2_\odot = 7.59^{+0.21}_{-0.21} \times 10^{-5}\mathrm{eV}^2 \tag{2.43}$$

$$\tan^2\theta_\odot = 0.47^{+0.06}_{-0.05} \, . \tag{2.44}$$

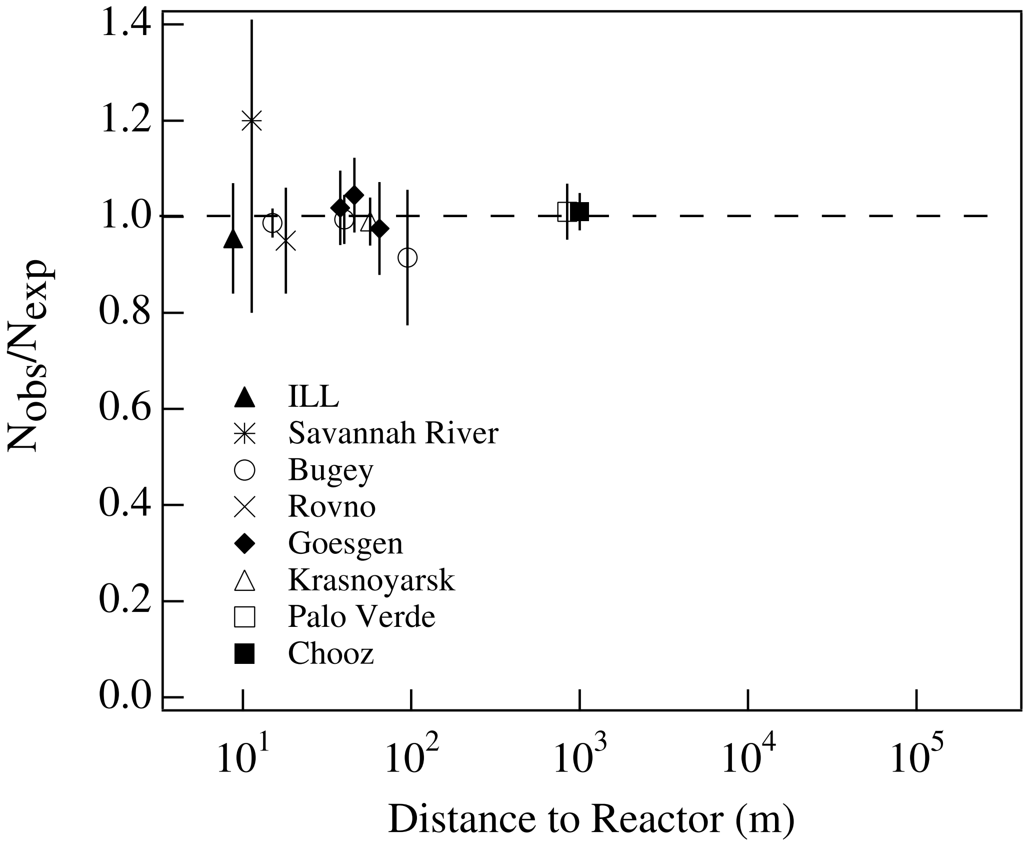

Figure 2.15: Summary of results from reactor neutrino experiments at various baselines, compared to expectations. No experiment has yet seen evidence of a $\bar{\nu}_e$ deficit. Taken from [64].

2.7 The last mixing angle

2.7.1 Reactor experiments

An alternate source of man-made neutrinos, utilized in experiments much earlier than accelerator neutrinos, comes from nuclear reactors. A large number of experiments have been performed, at varying distances from the reactor core, measuring the flux of electron antineutrinos. As shown in Figure 2.15 none of them have observed any significant deviation from expectations.

With reactor neutrinos around 3 MeV and the previously determined atmospheric mass splitting (2.41), the first maximum of the oscillation function (corresponding to maximal $\bar{\nu}_e$ disappearance) is expected to happen around 1 km. The solar oscillations observed by KamLAND occur over a much longer baseline and are not significant here.

The absence of any large deficit at this baseline means that the value of the remaining mixing angle θ_{13}, which governs the magnitude of the atmospheric oscillation component in $P_{\bar{e}\bar{e}}$, must be small.

The best limit on θ_{13} from a reactor experiment comes from CHOOZ [65], being at the correct baseline to observe any oscillation signal, and obtaining a better sensitivity than Palo Verde [66], which suffered a significantly larger cosmic muon flux due to its shallower depth.

The CHOOZ detector consisted of a cylindrical vessel containing 5 tons of gadolinium-doped liquid scintillator surrounded by photomultiplier tubes. The signal of a neutrino interaction, via inverse beta decay, was prompt photons from the annihilation of the positron, followed by gamma rays from capture of the neutron onto gadolinium. This detector was located 1 km from the Chooz nuclear power station shielded by a 300 m.w.e. rock overburden. Data was taken over 450 days in 1997 and 1998.

The error on the final result comes equally from statistical and systematic error, and the limit on the mixing angle depends on the mass splitting, but for realistic values ($\Delta m^2_{\mathrm{atm}} \gtrsim 2 \times 10^{-3}\,\mathrm{eV}^2$) CHOOZ limits $\sin^2 2\theta_{13} \lesssim 0.1$ at 90% confidence level.

The small value of this mixing angle provides additional evidence that the large ν_μ deficit seen in atmospheric neutrinos cannot be due to $\nu_\mu \leftrightarrow \nu_e$ oscillations, leaving $\nu_\mu \leftrightarrow \nu_\tau$ as the only alternative within the standard theory of oscillations.

An upgraded experiment, Double Chooz [67], with a new far detector, and the addition of a near detector to reduce uncertainties in the reactor flux, will be able to set tighter limits on θ_{13}, or potentially obtain evidence for a nonzero value.

2.7.2 MINOS

Whilst by no means designed for the identification of ν_e interactions, the MINOS experiment also has sensitivity to θ_{13} and has searched for $\nu_\mu \leftrightarrow \nu_e$ oscillations in the NuMI beam over the atmospheric baseline. The analysis observes a small excess of five ν_e-like events over the expectation of 49 [68]. Unlike CHOOZ, the oscillation probability in MINOS is dependent on δ and the neutrino mass hierarchy. The oscillations follow the form of Equation 2.26, plus a small contribution from matter effects, which is responsible for the dependence on the sign of $\Delta m^2_{\mathrm{atm}}$. Due to the observed excess, the best-fit θ_{13} is non-zero for all values of δ or the hierarchy, although a value of zero is not excluded. Despite this excess the MINOS limits, shown in Figure 2.16, are competitive with those of CHOOZ and, assuming normal hierarchy, are superior for most values of δ.

2.8 Overview of the neutrino sector

All the experimental evidence presented so far is consistent with the sole process involved in neutrino disappearance being neutrino oscillations, as described in Section 2.4.1, between three flavours. A particularly simple form of the mixing matrix is known as tri-bimaximal mixing [70]

$$U = \begin{bmatrix} \sqrt{\frac{2}{3}} & \sqrt{\frac{1}{3}} & 0 \\ \sqrt{\frac{1}{6}} & \sqrt{\frac{1}{3}} & \sqrt{\frac{1}{2}} \\ \sqrt{\frac{1}{6}} & \sqrt{\frac{1}{3}} & \sqrt{\frac{1}{2}} \end{bmatrix} , \qquad (2.45)$$

and is so far compatible with all observations. These matrix elements imply the following mixing angles:

$$\theta_{12} = 35.3^\circ \qquad \theta_{23} = 45^\circ \qquad \theta_{13} = 0 \qquad \delta = 0 \; . \qquad (2.46)$$

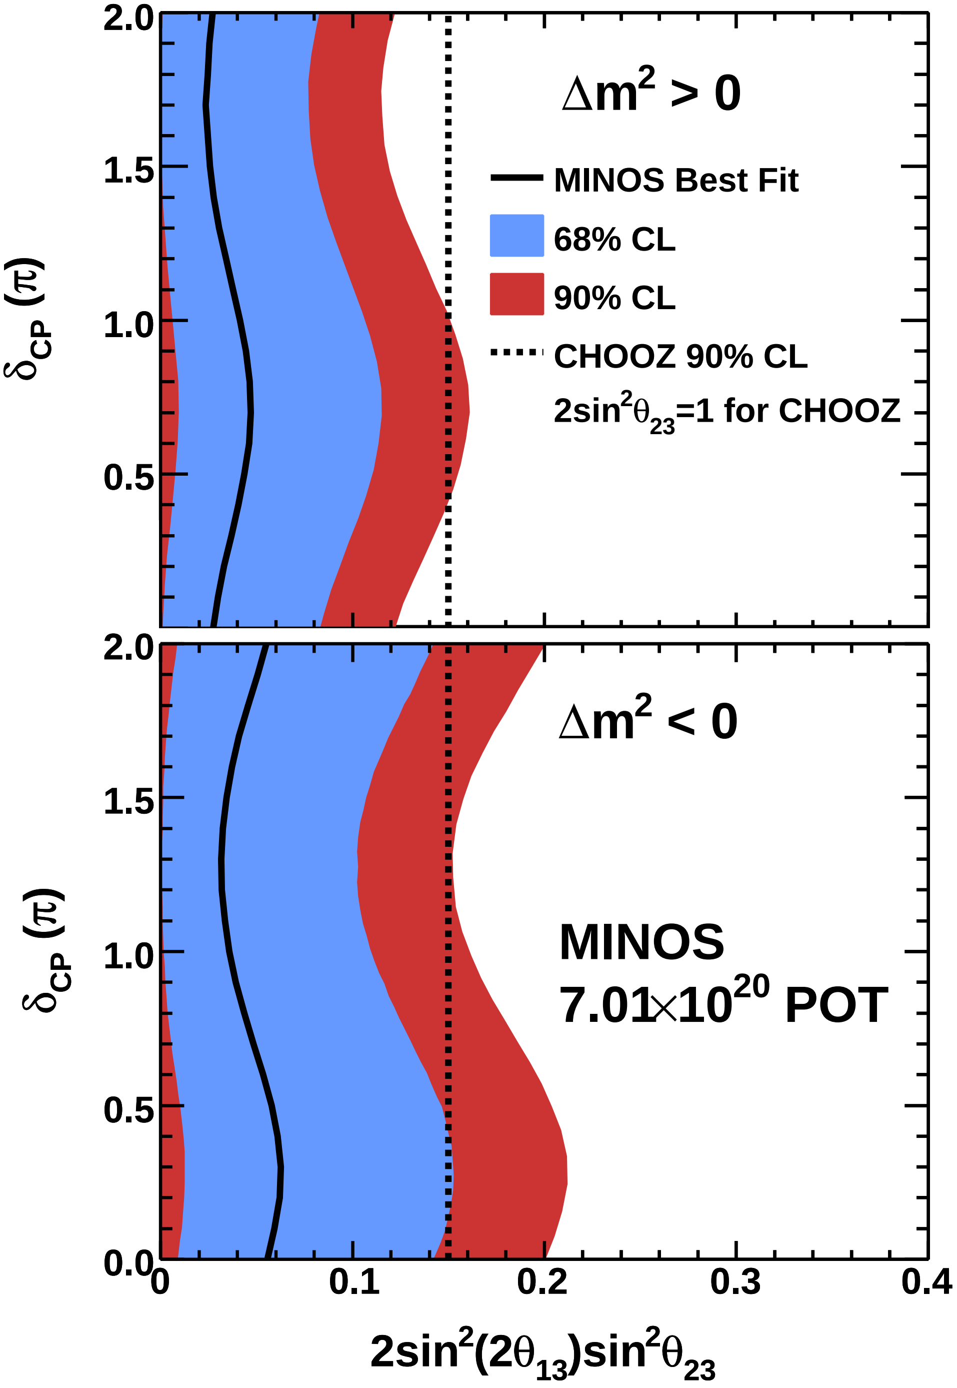

Figure 2.16: MINOS limits on $\nu_\mu \leftrightarrow \nu_e$ oscillations from [68]. The limits vary as a function of the $\mathcal{CP}$-violating phase $\delta_{\rm CP}$ and according to whether the hierarchy is normal (top) or inverted (bottom). The best fit at each value of δ is represented by the black lines, and the 68% and 90% confidence regions by blue and red shading respectively. The dotted line indicates the 90% CHOOZ limit [65], evaluated according to the MINOS best-fit atmospheric parameters in [69].

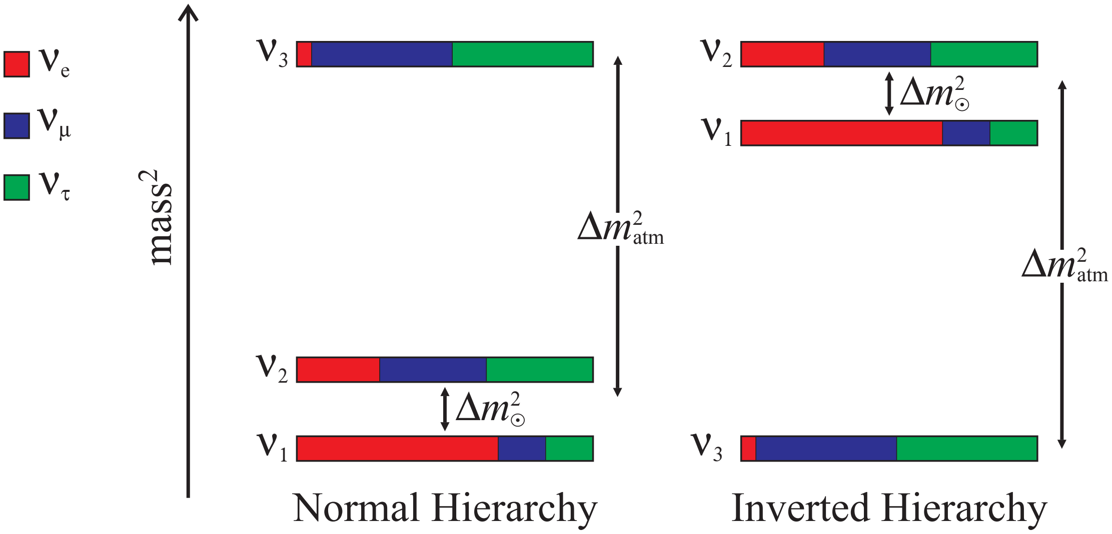

Figure 2.17: Schematic of the two possible orderings of the neutrino mass states. Each state is shown with its approximate flavour composition, assuming small but non-zero θ_{13}.

θ_{13} is set to zero here, but many theorists predict the true value to be not far below the current experimental limits, and global fits of neutrino data tend to yield non-zero θ_{13} near the present limits, but not at particularly high significance levels [71, 72].

In the absence of a definitive observation of non-zero θ_{13}, the effect of δ is also unobservable, as described in Section 2.4.3, and we have essentially no information about this phase[6].

The two mass splittings are fairly well-determined:

$$\Delta m^2_\odot \sim 8 \times 10^{-5}\,\text{eV}^2 \qquad\qquad |\Delta m^2_{\text{atm}}| \sim 2.5 \times 10^{-3}\,\text{eV}^2\ . \tag{2.47}$$

The large difference in scales, along with the small value of θ_{13}, is the reason all individual experimental results so far are well described by a two-flavour formalism.

The sign of $\Delta m^2_\odot$ and hence the ordering of the ν_1 and ν_2 mass states is determined by the necessity of invoking matter effects in the sun for the description

---

[6] Global fits tend to find large values of δ, again at low significance.

of solar neutrino data (Section 2.6.1). On the other hand, no current experiment which probes atmospheric oscillations expects a large contribution from matter effects. The sign of $\Delta m^2_{\mathrm{atm}}$ is thus still undetermined. The scenario in which ν_3 is heavier than ν_1 and ν_2 is known as “normal hierarchy”, and the scenario where it is lighter is referred to as “inverted hierarchy”. These two possibilities are indicated schematically in Figure 2.17, which also indicates the approximate flavour composition ($|U_{\alpha i}|^2$) of the three states, assuming a small but non-zero value for θ_{13}.

2.9 The LSND anomaly and MiniBooNE

A single discrepancy threatens this picture of the neutrino sector, and holds out the possibility of new physics. The LSND experiment at Los Alamos [73] consisted of a roughly cylindrical vessel containing 167 tons of doped mineral oil scintillator. The 1220 photomultiplier tubes observing this volume detected both scintillation and Čerenkov light. This detector was sited 30 m downstream of the neutrino production, where an 800 MeV proton beam was incident on a water (later high-Z metal) target. π^+ coming to rest in the target and decaying produced ν_μ, ν_e and $\bar{\nu}_\mu$, but no $\bar{\nu}_e$. Any oscillation of $\bar{\nu}_\mu$ to $\bar{\nu}_e$ could thus be detected via the reaction $\bar{\nu}_e + p \to e^+ + n$, by detecting both the positron and the photon from neutron capture $n + p \to d + \gamma$, with very little background from $\bar{\nu}_e$ contamination in the beam.

An excess of 88 $\bar{\nu}_e$ events was observed [74], at a significance of 3.8σ, implying a transition probability for $\bar{\nu}_\mu \to \bar{\nu}_e$ of around 0.3%. Interpreting this result in terms of neutrino oscillations, and given neutrino energies up to about 50 MeV, requires a mass-splitting $\mathcal{O}(1\,\mathrm{eV}^2)$. The allowed parameter space is shown by the blue regions in Figure 2.18. If the mass splitting is $\gtrsim 1\,\mathrm{eV}^2$ then the mixing angle must be very small to explain the low oscillation probability. Alternatively the mixing may be large, and the mass splitting lower, indicating that the oscillations have not yet fully developed by the LSND baseline. In any case Δm^2 is required to be larger than $10^{-2}\,\mathrm{eV}^2$.

This mass splitting is incompatible with the previously determined solar and atmospheric oscillation L/E regimes. In a three-neutrino scheme there are only two independent mass splittings. The LSND result could therefore possibly point to the existence of a fourth neutrino state, which would be required to be sterile in order to be consistent with the Z-width measurements from LEP (Figure 2.3). Even adding a sterile neutrino, global analyses still find significant tension fitting all known results; adding a second or third additional neutrino allows all of the data to be fitted, at the expense of a large number of free parameters. The LSND observations and their incompatibility with the standard oscillation theory are known as the "LSND anomaly".

The Bugey [75] and KARMEN [76] experiments were able to exclude parts of the LSND parameter space, but significant regions remained allowed.

The MiniBooNE experiment at Fermilab was constructed specifically to confirm or reject the LSND anomaly. The MiniBooNE detector consists of a spherical tank holding 800 tons of pure mineral oil, observed by 1520 PMTs. It is situated 541 m downstream of the muon neutrino beam produced by 8 GeV protons from Fermilab's Booster striking a beryllium target inside a magnetic focusing horn. The neutrino spectrum peaks at 700 MeV, which means MiniBooNE is sensitive to the same L/E range as LSND, but with different experimental characteristics, hopefully shielding it from any possible systematic problems with the LSND result.

Initially MiniBooNE ran with a muon neutrino beam, to take advantage of the larger cross-sections compared to antineutrinos, testing the LSND result under the assumption that $P_{\mu e} = P_{\bar{\mu}\bar{e}}$. This running successfully ruled out the whole of the LSND 90% allowed region [77] (left half of Figure 2.18). There was however an unexplained 3.0σ excess of electron-like events below 475 MeV [79], which cannot

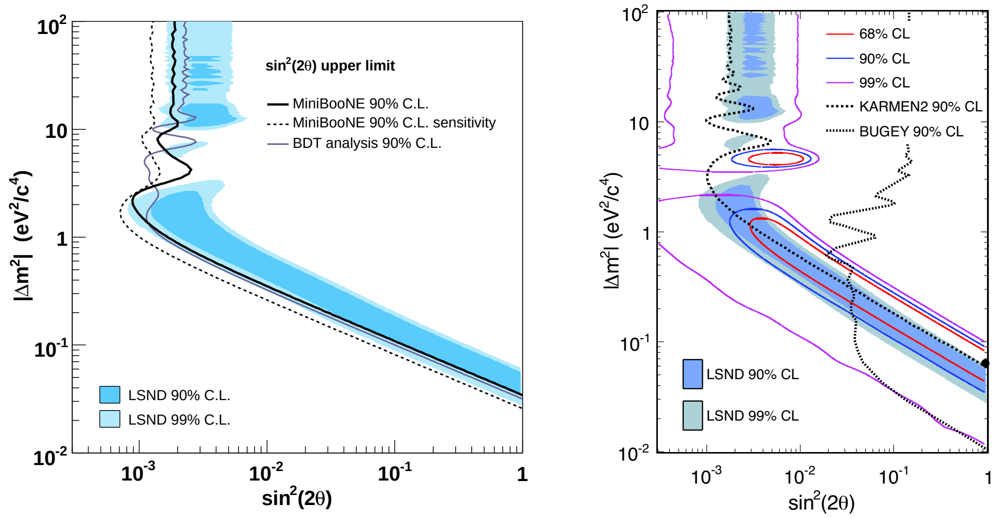

Figure 2.18: MiniBooNE confidence limits for $\nu_\mu \rightarrow \nu_e$ oscillations (left) and $\bar{\nu}_\mu \rightarrow \bar{\nu}_e$ oscillations (right) using events in the range 475–3000 MeV. Figures from [77] and [78]. In each case the contours are overlayed on the LSND $\bar{\nu}_\mu \rightarrow \bar{\nu}_e$ result [74] (blue). Regions to the right of black lines are excluded, regions within the red, blue and purple contours are allowed.

be explained by a two-neutrino oscillation model. Due to concerns about possible poorly modelled backgrounds, this region was excluded from the analysis before the dataset was unblinded.

MiniBooNE then ran in antineutrino mode, in order to make a more direct test of LSND. In this mode MiniBooNE observe an excess of $\bar{\nu}_e$-like events, consistent with LSND oscillations, and no evidence of a low-energy excess as seen in neutrino mode [78]. Interpreting this $\bar{\nu}_e$ excess as an oscillation signal produces the contours shown in the right half of Figure 2.18, consistent with the LSND allowed region, and excluding null oscillations at high probability. Additional running in antineutrino mode is expected to approximately double the available statistics.

The correct explanation for the LSND and MiniBooNE observations, particularly the difference between neutrinos and antineutrinos, is as yet unknown. Possibilities include $\mathcal{CP}$-violating oscillations involving new sterile neutrino states, unexpectedly high $\bar{\nu}_e$ cross-sections, or unanticipated systematic effects (no systematic in MiniBooNE can explain the observed excess with less than a 3σ shift).

We can only hope that time and additional data clarify the situation, although the origin of the LSND anomaly has now remained unexplained for 14 years.

2.10 Direct mass searches and neutrinoless double beta decay

2.10.1 The β-endpoint

Whilst neutrino oscillation experiments can measure the (squared) mass differences between the neutrino states, they are not sensitive to the absolute masses. Although, of course, the measured splittings mean that at least one neutrino must be more massive than $\sqrt{\Delta m^2_\odot} \sim 9 \times 10^{-3}\,\mathrm{eV}^2$ and another more massive than $\sqrt{\Delta m^2_{\mathrm{atm}}} \sim 5 \times 10^{-2}\,\mathrm{eV}^2$. One method of investigating the absolute mass scale is via the study of β decays.

Traditionally, assuming the neutrino to be massless, the Q-value of a decay

$$Q_\beta = M_i - M_f - m_e \tag{2.48}$$

is the maximum kinetic energy that can be imparted to the decay electron. Introducing a neutrino mass (and disregarding for the moment the effects of neutrino mixing, so that the electron neutrino can be taken to have a definite mass) the maximum electron energy is now

$$T_{\mathrm{max}} = Q_\beta - m_{\nu_e} \ . \tag{2.49}$$

In principle, this shift in the end-point of the β spectrum from expectations allows the neutrino mass to be measured. In practice, due to the low rate of decays near the end-point and uncertainties in its expected position, experiments instead look at the departure from the expected behaviour which occurs just before the end-point. The phase space factor in Fermi's Golden Rule gives a differential rate

$$\frac{\mathrm{d}\Gamma}{\mathrm{d}T} \propto E_e p_e E_\nu p_\nu \; . \tag{2.50}$$

Substituting for the neutrino momentum

$$p_\nu = \sqrt{E_\nu^2 - m_{\nu_e}^2} = \sqrt{(Q_\beta - T)^2 - m_{\nu_e}^2} \tag{2.51}$$

the differential decay rate becomes

$$\frac{\mathrm{d}\Gamma}{\mathrm{d}T} \propto E_e p_e (Q_\beta - T)\sqrt{(Q_\beta - T)^2 - m_{\nu_e}^2} \; . \tag{2.52}$$

When the neutrino mass is non-zero, the rate falls off faster than would be otherwise expected, in order to reach zero early at $T = Q_\beta - m_{\nu_e}$. The data is traditionally interpreted in terms of the Kurie function, where all the matrix elements, and the electron part of the phase-space factor, are divided out:

$$\begin{aligned} K(T) &\propto \sqrt{\frac{\mathrm{d}\Gamma/\mathrm{d}T}{E_e p_e}} \\ &= \sqrt{(Q_\beta - T)\sqrt{(Q_\beta - T)^2 - m_{\nu_e}^2}} \; . \end{aligned} \tag{2.53}$$

In the absence of neutrino mass this function forms a straight line $K(T) = Q_\beta - T$ and departures from linearity in this curve can be interpreted in terms of the neutrino mass.

The Mainz [80] and Troitzk [81] experiments have observed the beta spectrum of tritium (^{3}H) decay with sensitive spectrometers and used this technique to extract upper bounds on the neutrino mass

$$m_{\nu_e} < 2.3\,\mathrm{eV} \qquad (95\%\ \mathrm{C.L.}) \tag{2.54}$$

$$m_{\nu_e} < 2.5\,\mathrm{eV} \qquad (95\%\ \mathrm{C.L.}) \tag{2.55}$$

respectively. These limits are more than five orders of magnitude lower than the next lightest known particle, the electron at 511 keV.

The discussion so far has assumed the existence of a well-defined electron neutrino mass m_{ν_e}. Insofar as such a treatment is valid, what is actually measured is an incoherent sum over all the neutrino mass states that couple to the electron, an effective "beta-decay mass":

$$m_\beta^2 = \sum_i |U_{ei}|^2 m_i^2 \ . \tag{2.56}$$

Combined with Δm^2 and mixing-angle measurements from oscillation experiments, in certain circumstances a measurement of m_β can even determine the hierarchy. Because U_{e3} is small, m_β is primarily determined by m_1 and m_2. If m_β is found to be lower than about $5 \times 10^{-2}\,\mathrm{eV}^2$ then the inverted hierarchy would be ruled out, since the mass m_3 would then need to be negative to match the observed atmospheric splitting.

In principle, the definition (2.56) is only an approximation, and with sufficiently high experimental resolution individual features should be observed in the Kurie plot for each neutrino state individually, thus determining their masses directly. In practice, the required experimental resolution seems unlikely ever to be achieved.

2.10.2 Neutrinoless double-beta decay

In addition to the usual beta decay of nuclei

$$\mathcal{N}(A, Z) \rightarrow \mathcal{N}(A, Z+1) + e^- + \bar{\nu}_e \ , \tag{2.57}$$

a double-beta process

$$\mathcal{N}(A, Z) \rightarrow \mathcal{N}(A, Z+2) + 2e^- + 2\bar{\nu}_e \tag{2.58}$$

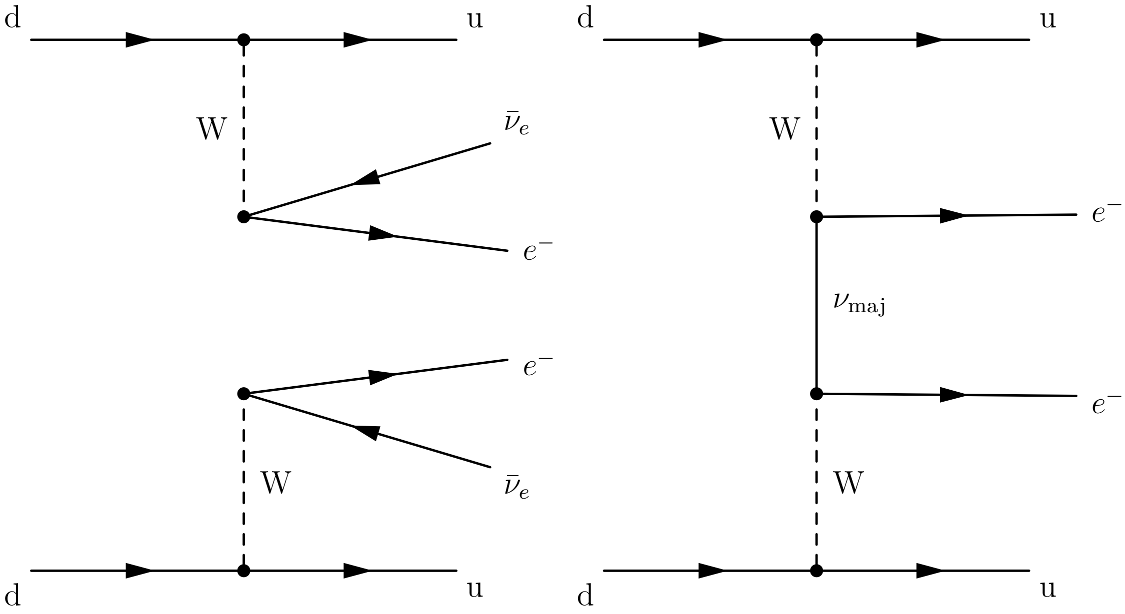

Figure 2.19: Left: Feynman diagram for double-beta decay, leading to a final state which includes two electron antineutrinos. Right: Neutrinoless double-beta decay, the Majorana neutrino forms an internal line. This diagram violates lepton number conservation and is thus not allowed by the Standard Model.

is also possible via a second-order diagram (left side of Figure 2.19). When single-beta decay is energetically possible then double-beta decay is unobservable due to its much lower rate. However, for some even-even nuclei single beta decay is energetically impossible, whilst double-beta decay is allowed, due to the pairing force between identical nucleons. The four-body final state of double-beta decay means that the energy distribution of the outgoing electrons form a continuous spectrum.

An alternate decay method "neutrinoless double-beta decay" is

$$\mathcal{N}(A,Z) \rightarrow \mathcal{N}(A,Z+2) + 2e^- \ , \tag{2.59}$$

as illustrated in the right half of Figure 2.19. Here, the final-state neutrinos are joined together as an internal line and only the electrons remain as final-state leptons. Of course this process is not allowed in the Standard Model since it violates lepton number conservation by two units. It is possible, however, that neutrinos are Majorana fermions, that is the neutrino and antineutrino are merely the left- and right-handed states of the same particle. In this case, lepton number is not conserved, and the process (2.59) is allowed. The difficulty remains that the internal neutrino needs opposite helicity at each vertex and so for massless neutrinos this process would again be disallowed. Fortunately, we know that neutrinos do indeed have mass, and the probability of the necessary helicity change is proportional to the mass. Taking into account neutrino mixing, the rate of neutrinoless double-beta decay depends on an effective neutrino mass

$$m_{2\beta} = \sum_i U_{ei}^2 m_i \ . \tag{2.60}$$

Although the presence or absence of the final state neutrinos is not directly detectable, in the neutrinoless process the two electrons carry away the full Q-value of the decay which should appear as a peak in the energy spectrum against the continuous background from conventional double-beta decay.

Several experiments have searched for the neutrinoless double-beta decay of various elements, the detection of which would confirm the Majorana nature of the neutrino. The current best limit comes from the Heidelberg-Moscow experiment [82], which sets a limit on the half-life of ^{76}Ge via neutrinoless double-beta decay of

$$\tau_{1/2}^{0\nu}(\,^{76}\mathrm{Ge}) > 1.9 \times 10^{25}\,\mathrm{years} \qquad (90\%\ \mathrm{C.L.})\ . \tag{2.61}$$

This may be converted to a limit of $m_{2\beta} < 0.35\,\mathrm{eV}$ (90% C.L.), although there are large uncertainties on the nuclear matrix elements required to perform this conversion.

A claim exists [83], based on data from the same experiment, of a discovery of $\tau_{1/2}^{0\nu}(\,^{76}\mathrm{Ge}) = 1.19^{+1.00}_{-0.17}$ years. However, the analysis has been widely criticized [84, 85, 86] and the result is not commonly accepted.

For similar reasons as in beta-endpoint experiments, discovery of a sufficiently small $m_{2\beta}$ would be able to rule out the inverted hierarchy scenario.

2.11 Astrophysical neutrinos

2.11.1 Supernova 1987A

In 1987 a type II supernova was observed in the Large Magellanic Cloud, the first supernova visible to the naked eye since 1604. This supernova is known as SN1987A, being the first discovered that year. Supernova explosions are expected to produce very large fluxes of electron neutrinos, with energies around 10 MeV. Three neutrino detectors operating at the time did indeed see an excess of events a few hours before the optical discovery of SN1987A. These events constitute the only conclusive observations of neutrinos originating outside of the solar system.

The Kamiokande-II collaboration, upon examining their data, found 11 likely signal events around 10 MeV within a 12 s window of each other [87]. The IMB detector recorded eight events with energies between 20 and 40 MeV within a 6 s time interval [88]. Finally, the Baksan liquid scintillator detector found a cluster of five events between 10 and 25 MeV within a 10 s time window[7] [89].

Unfortunately, accurate timing information was not considered a priority and there thus exists a one minute uncertainty in the absolute time of the Kamiokande and Baksan events. Within this range the timings are consistent with all three detectors observing the same ~ 10 s burst of neutrinos simultaneously.

A joint analysis of all the experimental signals [90] finds them to be consistent with standard models of supernova explosions. The emitted neutrinos are expected to cover a range of energies, and so travel at different speeds. The initial short neutrino pulse will reach the Earth spread out in time, with a greater divergence indicating larger neutrino masses. Using the best knowledge of the duration of

---

[7] The background rate is quite high, but in joint analyses with the Kamiokande and IMB data these events are found to be consistent with the supernova expectations.

the neutrino-production phase of the supernova, the analysis in [90] finds a limit on the electron neutrino mass[8]

$$m_{\mathrm{SN}} < 5.7\,\mathrm{eV} \qquad (95\%\ \mathrm{C.L.})\ . \tag{2.62}$$

The expected rate of supernovae within the Milky Way is a few per century and neutrino detectors have advanced considerably since 1987. The next galactic supernova is expected to yield $\mathcal{O}(10^4)$ events and is eagerly awaited.

2.11.2 Cosmological limits

The observed matter density of the universe can be used to set a direct limit on absolute neutrino masses [91, 92]. From the measured temperature of the cosmic microwave background, and requiring equilibrium conditions in the early universe, the relic neutrino background is found to be at about 2 K and to have a number density $\sim 100\,\mathrm{cm}^{-3}$. The WMAP measurement of the matter component of the universe is around 14% of the critical density required for a flat universe [93]. The contribution from all the neutrino masses is thus limited:

$$\sum_i m_i \lesssim 13\,\mathrm{eV}\ . \tag{2.63}$$

This limit assumes the neutrinos are relativistic at decoupling, and thus does not include any contribution from possible heavy sterile neutrino states.

Fits to the details of structure formation in the early universe obtain much more stringent bounds on the neutrino masses, competitive with the beta-endpoint and neutrinoless double-beta decay limits. There is, however, quite a large degree of uncertainty based on the exact technique and data-sets included. Most analyses

---

[8] Due to neutrino mixing this is again an effective mass, like m_β and $m_{2\beta}$.

[93, 94, 95] obtain limits:

$$\sum_i m_i \lesssim 0.5 - 1\,\mathrm{eV}\ . \tag{2.64}$$

Chapter 3

The MINOS experiment

This chapter gives an overview of the MINOS (Main Injector Neutrino Oscillation Search) experiment. The NuMI neutrino beam is described in Section 3.2 and the MINOS detectors in Section 3.3. Sections 3.4 and 3.5 describe the Monte Carlo simulation and reconstruction software used in data analysis. Full descriptions of the NuMI beamline are given in [96, 97] and of the detectors in [98].

3.1 Overview and physics goals

MINOS, in common with other long-baseline accelerator neutrino experiments (K2K before it, OPERA, and T2K and NOνA currently under construction), is based on a two-detector paradigm. The “near detector” (ND) measures the properties of the neutrino beam approximately 1 km downstream of its production, allowing comparison with the spectrum observed in the “far detector” (FD), 735 km from the origin of the beam. The layout is indicated schematically in Figure 3.1. This approach allows for cancellation of systematic uncertainties in the beam flux and in cross-sections and reconstruction effects in the detectors. For this reason, the detectors are designed to be as similar as possible within cost considerations.

The primary measurement of the MINOS experiment, and the one presented

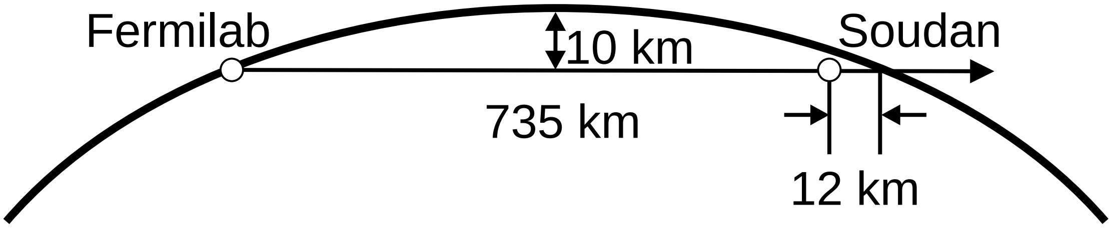

Figure 3.1: Sketch of the path of the NuMI neutrino beam from its production at Fermilab, and measurement by the near detector there, to the far detector at the Soudan mine, 735 km away.

in this thesis, is of the atmospheric oscillation parameters Δm^2_{atm} and $\sin^2 2\theta_{\text{atm}}$. The baseline and the energy of the beam are optimized so that upon reaching the far detector the primarily ν_μ beam is at the peak of its first oscillation into ν_τ. The oscillation signal is an energy-dependent deficit of ν_μ at the far detector[1]. The values of the oscillation parameters can be found by comparing the position and depth of the dip against the unoscillated prediction based on the near detector data and the Monte Carlo simulation.

The NuMI beam can also be run in antineutrino mode, allowing the same measurement to be made with muon-antineutrinos. Although not optimized to distinguish ν_e events, MINOS has obtained a competitive limit on the sub-dominant $\nu_\mu \leftrightarrow \nu_e$ oscillations (see Section 2.7.2). MINOS has also searched for sterile neutrinos via the neutral-current event rate [99], is sensitive to atmospheric neutrinos [58, 59] and cosmic rays [100, 101, 102], has measured neutrino cross-sections using the near detector [103], and has set limits on Lorentz violation [104, 105].

3.2 The NuMI beam

The neutrino beam for MINOS is provided by the Neutrinos at the Main Injector (NuMI) beamline at Fermilab. 120 GeV protons are extracted from the Main

---

[1] Very few ν_τ undergo charged-current interactions in MINOS. Due to the mass of the tau lepton, the incident neutrino needs more than $m_\tau + m_\tau^2/2m_n \approx 3.5$ GeV to be above threshold, and even then the phase-space factor is small.

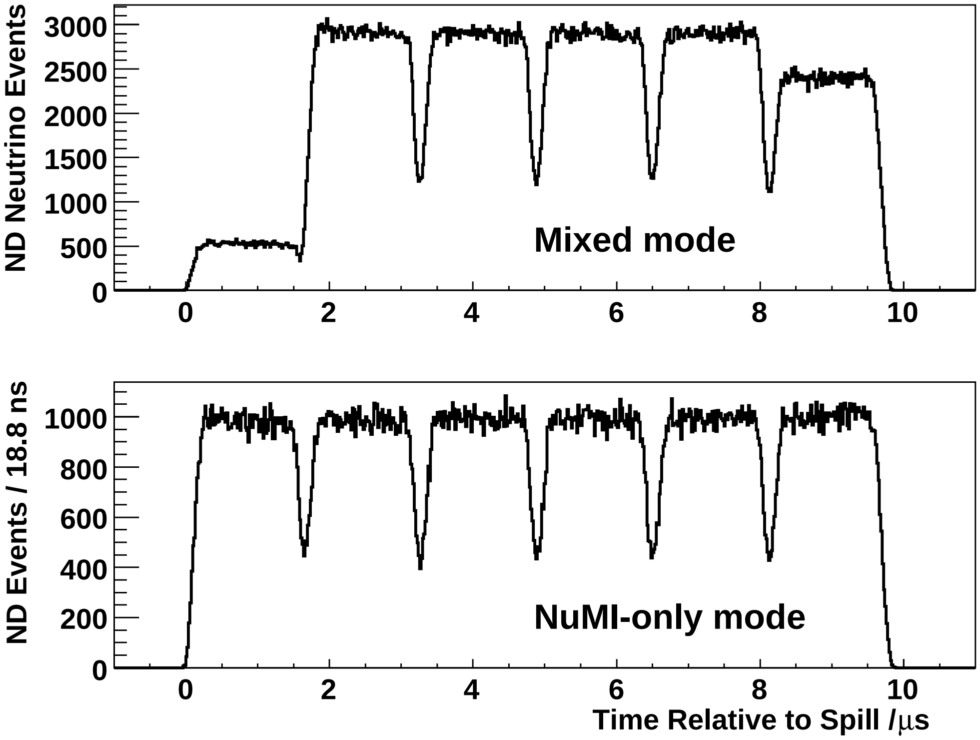

Figure 3.2: Timing structure of NuMI beam spills, in mixed mode (top) and NuMI-only mode (bottom). In mixed mode, the first batch of each spill is taken to produce antiprotons for the Tevatron, except during the transfer of antiprotons from the accumulator to the recycler. All but the last batch are formed from a combination of two Booster batches via a process known as slipstacking [106].

Injector and strike a 940 mm long graphite target. These “spills” occur every 2.2 s, last for 10 μs, and vary in intensity, containing between about 10^{13} and 3.5×10^{13} protons. Protons-on-target, or POT, are the unit in which MINOS beam exposure is measured. Each spill contains timing substructure, consisting of five or six batches (see Figure 3.2). The highest NuMI intensities are obtained when the Tevatron is not running.

The protons striking the target produce primarily pions and kaons. Positively charged pions and kaons are focused by two magnetic horns, separated from each other by 10 m. These hadrons then enter a 675 m decay pipe where most decay to ν_μ, with a small fraction to ν_e. Just before the aluminium and steel beam dump, the beam is observed by a hadron monitor. The remaining muons are

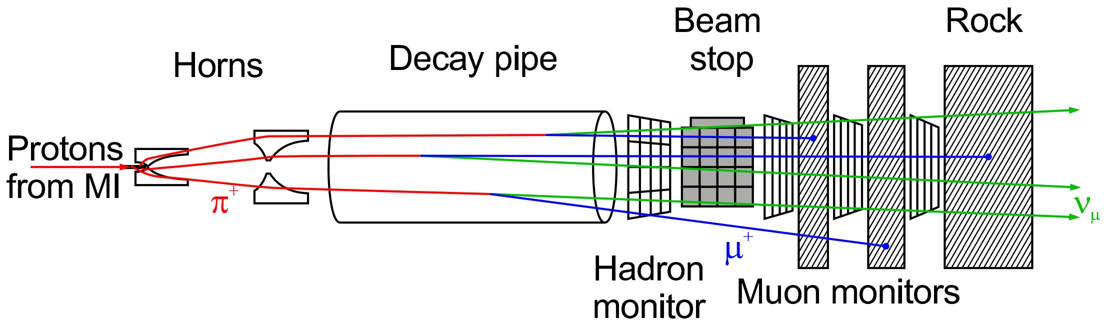

Figure 3.3: Schematic view of the NuMI beamline. Protons from the main injector form pions (red) upon striking the target. These are focused and decay to muons (blue), which are absorbed by the rock, and neutrinos (green) directed towards the MINOS detectors.

Name	Start	End	FD exposure	Comments
Run I	May 2005	Feb 2006	1.27×10^{20} POT	
pHE	Jun 2006	Aug 2006	1.53×10^{19} POT	High energy beam configuration
Run II	Sep 2006	Jul 2007	1.94×10^{20} POT	
Run III	Oct 2007	Jun 2009	3.88×10^{20} POT	Helium added to decay pipe
Run IV	Sep 2009	Mar 2010	1.71×10^{20} POT	Antineutrino beam configuration

Table 3.1: MINOS run periods used in this thesis. See text for details.

then absorbed by 300 m of rock, instrumented with three muon monitors at the upstream end. A schematic view of the neutrino beam production is shown in Figure 3.3. In the near detector Monte Carlo, interactions from beam neutrinos have relative rates: 92.9% ν_μ, 5.8% $\bar{\nu}_\mu$ and 1.3% $\nu_e + \bar{\nu}_e$.

The majority of MINOS running has been in the "low energy" beam configuration, as this maximizes the neutrino flux at energies around the atmospheric oscillation dip. In this mode the target is fully inserted into the first horn, and the current in the horns is 185 kA. Alternative neutrino fluxes can be obtained by altering the target position and horn current. These runs, and runs with the horns off entirely, are primarily useful for investigating the accuracy of the beam simulation and tuning the Monte Carlo to match (see Section 4.2). The neutrino fluxes produced by three different beam configurations are shown in Figure 3.4.

Table 3.1 shows the various MINOS run periods of significance to this thesis. Runs I and II were taken under very similar conditions, although using a different

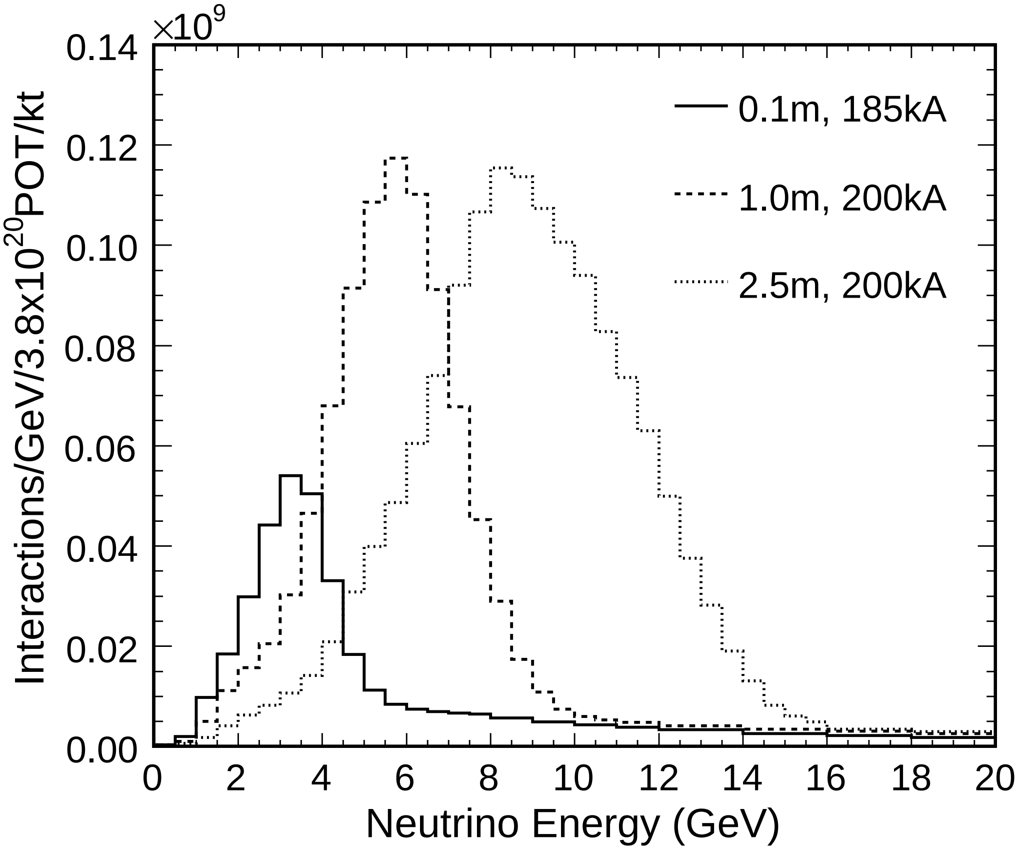

Figure 3.4: The true energy distribution of near detector events from three different NuMI configurations. The target position (distance upstream of a nominal position) and horn current of each configuration are shown in the caption. "0.1 m, 185 kA" is the "low-energy" configuration in which most MINOS data has been taken. "2.5 m, 200 kA" is the "pseudo-high-energy" configuration.

target and slightly different target positions. In between these runs (using the Run I target) a fairly large quantity of data was taken in the "pseudo-high-energy" (pHE) configuration. The higher peak energy of this running is primarily useful for testing alternative models of neutrino disappearance, which predict deviations from the standard oscillation form at high energies as well as low. Before Run III, safety concerns about a corroded and radiation-damaged inspection window led the previously evacuated decay pipe to be filled with helium at ∼1 atm. This causes a slight attenuation of the neutrino flux, especially in the beam peak. Finally, for Run IV, the current in the focusing horns was flipped from the usual "forward horn current" configuration to "reverse horn current". This configuration focuses

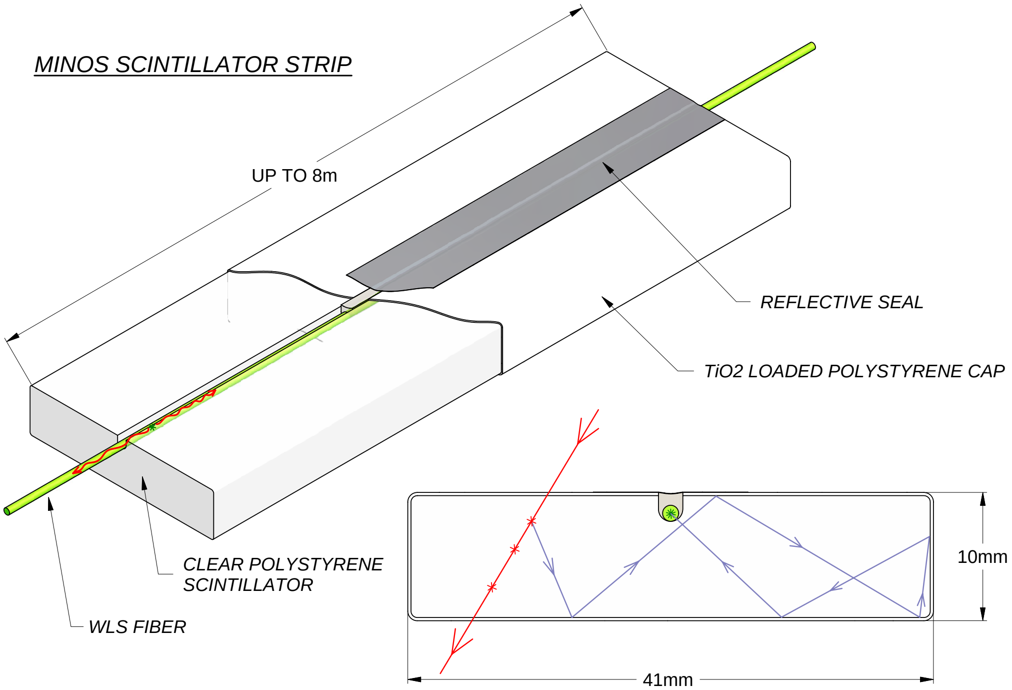

Figure 3.5: Construction of the MINOS scintillator strips. Each strip is covered in a TiO_2 coating that ensures the scintillation light is collected by the embedded wavelength-shifting fibre (inset).

negative pions and kaons, and thus produces an antineutrino-enriched beam.

3.3 The MINOS detectors

The MINOS near and far detectors are as similar as possible in order to allow cancellation of systematic errors between them. The general design is presented here, with information specific to the individual detectors in the following sections.

Both detectors are steel-scintillator tracking calorimeters. The design is optimized for neutrino interactions in the few GeV range. The bulk of each detector consists of 2.54 cm (one inch) thick steel planes. Plastic scintillator strips, with cross-section 4.1×1 cm, are affixed to each steel plane. The planes are separated by an air gap leading to a total pitch of 5.94 cm. The scintillator strips are formed from extruded polystyrene doped with the organic scintillators PPO and

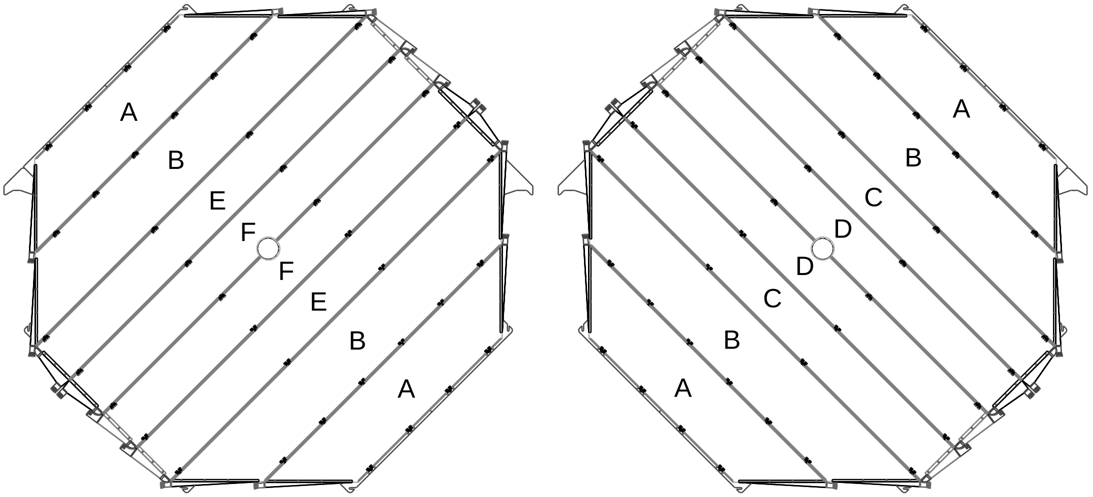

Figure 3.6: Arrangement of the scintillator modules in the far detector. The left-hand diagram shows an odd-numbered or U-plane and the right-hand diagram an even-numbered or V-plane. The A and B modules comprise 28 scintillator strips, the others have 20.

POPOP at a concentration of 1%. The strips have a diffusely reflective titanium dioxide (TiO_2) coating, ensuring that scintillation light is collected by the embedded wavelength-shifting fibre. The fibre carries the light out of the detector to one pixel of a multi-anode photomultiplier tube (PMT). The construction of the MINOS scintillator strips is illustrated in Figure 3.5.

The strips are assembled into modules and mounted to the steel planes at a 45° angle. The planes are thus read out from their diagonal edges. Alternate planes have their strips mounted at right-angles to each other, allowing events to be observed in "U" and "V" views. Figure 3.6 shows the construction of U and V planes in the far detector, and Figure 3.7 illustrates the alternation of the two views with the steel planes and air gap.

Both detectors are toroidally magnetized to around 1 T by a coil running longitudinally through them, with the return leg outside of the detector. The curvature of particles in this magnetic field allows their charge to be determined and provides an estimate of their energy, which is necessary when a particle exits the detector

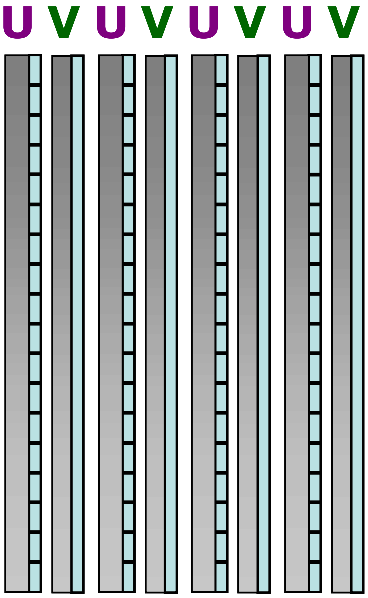

Figure 3.7: Illustration of the alternation of steel plane, scintillator and air gap in MINOS. Alternate scintillator planes are U-type or V-type and thus viewed end- or side-on in this diagram.

so that the more accurate estimate from range cannot be used. The magnetic field also helps to contain muons in the detector, focusing negative particles (or positive particles during the Run IV antineutrino running) towards the coil. During Runs III and IV the near detector field was periodically reversed in order to study the wrong-sign component of the beam.

Calibration of the detectors, to ensure uniformity both across the volume of the detector and across time, is achieved using cosmic ray muons, which provide a flux of particles with known energies and well-understood energy deposition. The calibration is observed to be correlated with the cavern temperature, and the scintillator material has degraded over time. This effect is shown in Figure 3.8. The detectors are also equipped with a light injection system. Ultraviolet LEDs periodically illuminate the fibres, allowing calibration of the response of each PMT pixel to a known light level.

The absolute calibration of the detectors was determined using a special calibration detector exposed to test beams at CERN [107, 108]. This was a small-scale

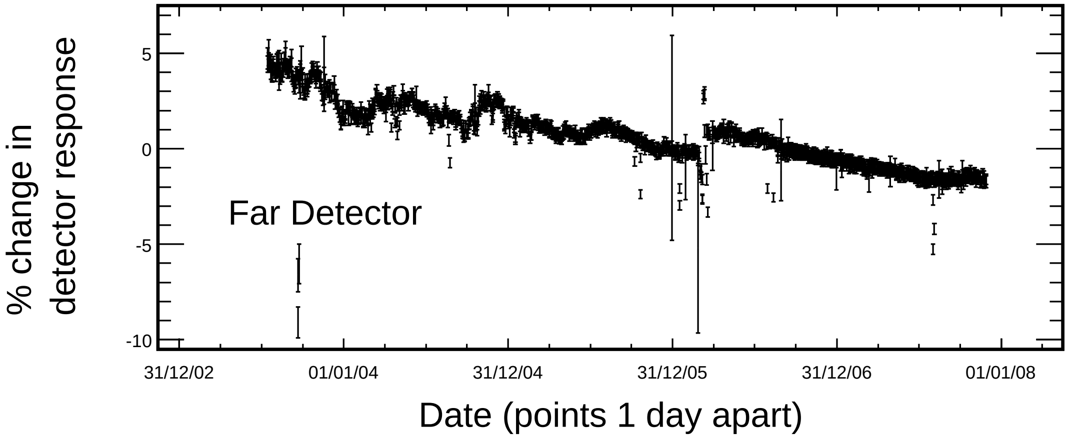

Figure 3.8: Median response of the far detector to cosmic rays as a function of time over several years. The fractional change in response is plotted compared to an arbitrary reference point. The light yield is seen to degrade by around 2% per year, due to aging of the detector components. The short-term variations are well correlated with temperature variations in the cavern. The discontinuous step in 2006 is due to a retuning of the PMT high voltage supplies.

prototype of the MINOS detector technology, consisting of 60 1×1 m unmagnetized planes, equipped with both near-detector and far-detector electronics. Beams of known energy and particle composition allowed the response of the scintillator and PMTs as a function of energy to be accurately determined.

3.3.1 The near detector

The MINOS near detector is located at Fermilab, 1040 m downstream of the NuMI target, in a cavern 100 m underground which provides 225 metres-water-equivalent (m.w.e.) of cosmic ray shielding. The layout of the cavern and a photograph of the near detector are shown in Figure 3.9. The detector consists of 282 planes and weighs 980 tons. The planes are approximately $4\,\mathrm{m} \times 3\,\mathrm{m}$ in a shape described as a "squashed octagon".

At the near detector the NuMI beam is still very intense, providing a high rate of interactions in the detector, and is concentrated on the beam spot (point A in Figure 3.9). It is therefore only necessary to partially instrument the near de-

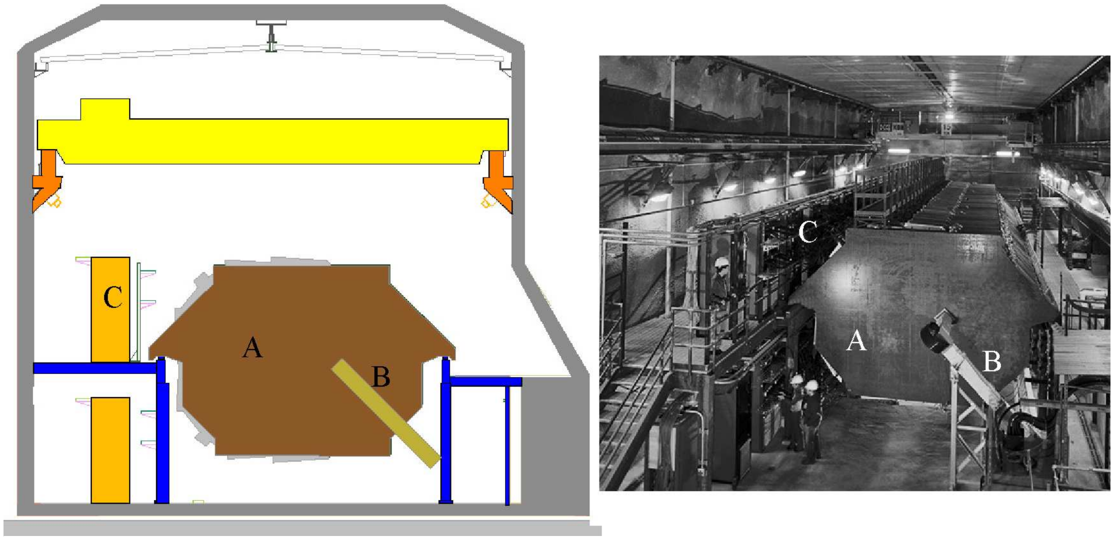

Figure 3.9: Layout of the near detector cavern. The schematic and photograph are both looking towards Soudan (downstream). The NuMI beam strikes the detector around the *A* label. *B* indicates the magnetic coil and *C* an electronics rack. Since this photograph was taken, the MINERνA detector has been constructed upstream of MINOS, obscuring much of this view.

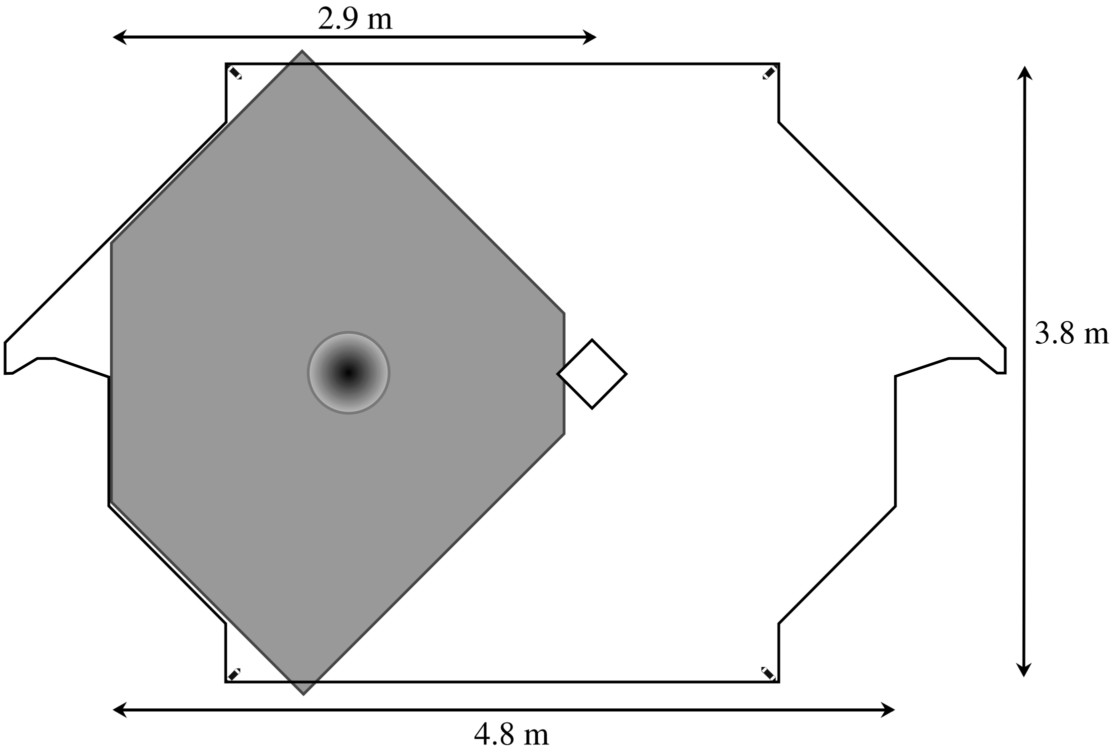

Figure 3.10: Geometry of the near detector. The diamond shape is the coil hole, and the circles indicate the beam centre. The “ears” are used to suspend the planes. In partially instrumented planes only the shaded region is covered by scintillator.

tector. In the upstream 120 planes of the detector, the “calorimeter”, the region around the beam spot is fully instrumented, but the remainder of the detector cross-section is only covered by scintillator in every fifth plane. The geometry of the fully-instrumented region is illustrated in Figure 3.10. In the remaining downstream region of the detector, the “spectrometer”, every fifth plane is fully instrumented and the remaining planes are uninstrumented. These partially instrumented regions allow muons that escape the fully instrumented region to continue to be tracked.

The strips in the near detector are read out at the western end (the side away from the magnet coil-return), being silvered at the other end. Each fibre is directed to one of the 64 pixels of a Hamamatsu R5900-00-M64 PMT [109]. The electronics digitize the PMT signals continuously at a frequency of 53 MHz. The buffers are large enough to hold the output from an entire beam spill. All activity during the 18.8 μs spill gate is read out. Additional activity-based triggers allow readout of out-of-spill events such as cosmic ray muons or atmospheric neutrinos.

3.3.2 The far detector

The MINOS far detector is located in a disused iron mine in Soudan, northern Minnesota, 735 km from the beam production point. The specially excavated cavern is 705 m underground, providing 2070 m.w.e. of shielding. The layout of the cavern and a photograph of the far detector are shown in Figure 3.11. The detector consists of 486 octagonal planes, 8 m across, and weighs 5,400 tons. It is divided into two “supermodules” of 249 and 237 planes, independently magnetized, separated by a 1.15 m gap.

Due to the divergence of the NuMI beam, the event rate in the far detector is much lower than in the near detector, leading to different decisions and cost optimizations in the design.

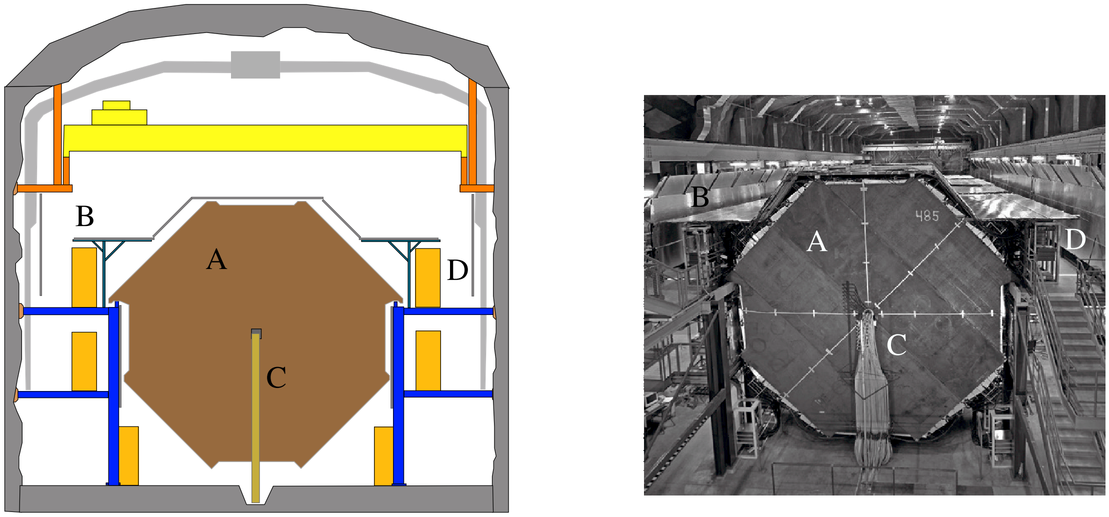

Figure 3.11: Layout of the far detector cavern. The schematic and photograph are both looking back towards Fermilab (upstream). A labels the most downstream steel plane, B the cosmic ray veto shield, C the magnetic coil and D an electronics rack.

Due to the greater attenuation along the length of the strips, some of which are 8 m long, they are read out at both ends. The outputs are attached to one of the 16 pixels of a Hamamatsu R5900-00-M16 PMT [110]. The much lower event rate, and dual-ended readout, allow an eightfold reduction in PMTs by multiplexing eight fibres onto each individual pixel. The eight fibres come from strips separated by about 1 m and the multiplexing pattern is different at each end, allowing the introduced ambiguity to be resolved at reconstruction-time [111]. PMT signals are digitized whenever the output exceeds $\frac{1}{3}$ of that expected from a single photo-electron. Due to noise in the WLS fibres [112], an additional requirement is made for two such signals to be detected within 400 ns in a group of 36 PMTs. The spill trigger is received at the far detector from Fermilab over the internet. Data is buffered sufficiently long to allow the spill timestamps to arrive. Activity-based triggers, as in the near detector, allow recording of out-of-spill events and of beam events in the case of an internet outage. Additionally, fake spill triggers are also used at the far detector to provide an estimate of the cosmic ray and detector noise backgrounds expected in true spills.

3.4 The MINOS Monte Carlo simulation

The MINOS Monte Carlo simulation begins with a simulation of the neutrino beam from NuMI. The FLUGG package combines FLUKA [113] for simulation of the proton interactions in the target, and the propagation and re-interaction of the produced particles, with a detailed description of the target and decay pipe geometry using GEANT4 [114]. The output of this step is a collection of neutrinos at the point of the decay of their parent particle labelled with their flavour, energy and momentum. The neutrino flux is simulated separately for each run period, taking into account the known shifts in the target position, and the effect of the decay pipe helium from Run III onward.

The interactions of these neutrinos in each detector are simulated using NEUGEN3 [115] for the interaction and GEANT3 [116] to propagate the resulting particles. The development of hadronic showers is modelled using the GCALOR package [117], which was selected because it displayed the best agreement with data from the calibration detector. The MINOS `PhotonTransport` package is then used to convert the simulated energy deposits to the expected PMT signals. The calibration constants used at this step are chosen randomly from the date range of the run period being simulated.

In the near detector, the effects of multiple interactions from the same beam spill are taken into account by overlaying the individually simulated events. The Monte Carlo for each run period is simulated at ten different beam intensities, chosen by dividing the intensity distribution in the MINOS data from that period into ten equal quantiles and finding the mean intensities [118]. “Rock muon” events from the interactions of neutrinos in the rock of the cavern walls, whose products then penetrate the detector, are also overlayed at this stage.

The greatly decreased event rate in the far detector means that the probability of two neutrino interactions occurring in one spill is completely negligible, and there is no need to perform the overlays. The effect of neutrino oscillations is taken into account at analysis-time by reweighting the Monte Carlo events, and by the use of special "tau-files", in which the composition of the beam is altered by interchanging ν_μ and ν_τ. These files model a transition probability $P_{\mu\tau} = 1$, and linear combination with the usual Monte Carlo with $P_{\mu\mu} = 1$ allows the construction of a spectrum for any $\nu_\mu \leftrightarrow \nu_\tau$ oscillation probability which includes the small contribution from ν_τ that do interact in the far detector.

3.5 Event reconstruction in MINOS

The first step of the MINOS reconstruction is designed to deal with the high interaction rate in the near detector. Recorded hits are divided by timing and topology cuts into "slices", each of which is intended to contain activity from only one event.

The next step is to identify the muon tracks that characterize charged-current events. First, hits in each view that appear to lie on a straight line are collected into "clusters". These clusters are then joined up to form candidate tracks. The candidate tracks in the U- and V-views are then compared and the best matches taken to form 3D tracks. The algorithm is described in detail in [119]. The track finder is fairly aggressive, and frequently identifies tracks within hadronic showers. Some of these may be real muon tracks, some are tracks from other particles in the shower, and many are simply "drawn on" to unrelated hits by the algorithm. The focus of the particle identification step of the analysis is predominantly to distinguish these tracks from true muon tracks.

Once a track candidate is identified, a Kalman filter [120] algorithm is used to estimate the true path of the muon, as it follows a helical trajectory in the magnetic field. At this stage, additional hits that were not located by the original track finding step, but which are close to the reconstructed muon trajectory, are added to the track. For each track, the track fitter outputs an estimate of the muon energy, both from the amount of material it traverses before stopping, and from its curvature in the magnetic field. It also outputs the reconstructed charge of the particle from its curvature[2].

The remaining hits in the slice are then also formed into clusters, and the clusters grouped into showers. Only hits with an energy deposit above two photoelectrons are considered, due to poorly understood data/MC disagreements at low pulse-heights, likely caused by fibre noise and crosstalk between PMT pixels. The energy of the shower is estimated by a calorimetric approach, summing the calibrated energy deposits of all the included hits.

The "event builder" stage then takes the collection of tracks and showers and associates them into events, using a long series of cuts tuned to avoid various observed reconstruction pathologies [121]. An event may have any number of tracks and showers but, where appropriate, one of each may be defined as the primary track and shower. The interaction vertex is also estimated at this stage, usually assigned to the first hit in the primary track.

---

[2] Technically the momentum and charge from curvature are combined in a single q/p number. The analysis assumes particles with unit charge.

Chapter 4

The charged-current analysis

The three components necessary to perform an analysis of ν_μ disappearance are: a method to obtain a pure sample of ν_μ charged-current events, a prediction of the expected far detector spectrum taking into account information gained from the near detector, and the ability to find the oscillation parameters which best describe the data, taking into account systematic uncertainties. These are described in Sections 4.3, 4.4 and 4.7 respectively.

Two additional components of the analysis, division of the far detector spectrum into bins of energy resolution, and inclusion of an additional sample of events from outside the fiducial volume, both improve the statistical power of the result and are described in Sections 4.5 and 4.6.

4.1 Fiducial volume

The primary analysis dataset consists of events whose primary vertex is contained within the fiducial volume. The cuts were chosen with the aim that all selected events should have their hadronic shower fully contained within the detector. In the far detector, the fiducial volume needs to be as large as possible, to maximize the number of selected events and thus the statistical power of the analysis. In the near detector, events are not at such a premium and the fiducial volume can be more conservative.

In the far detector, the vertex of an event must be within $\sqrt{14}\,\mathrm{m} \approx 3.74\,\mathrm{m}$ of the detector centre-line, but outside of 40 cm. These cuts remove events too close to the detector edge[1] or to the magnetic coil. The vertex must also not lie in the first three planes of either supermodule – vetoing events that originate outside the detector and enter through the front face – or the last ten planes of the first supermodule or the last 22 planes of the second supermodule – vetoing events where hadronic activity is likely to be lost out of the back of the detector. The far detector fiducial volume is marked in Figure 4.19.

In the near detector the radial cuts follow the axis of the beam, which, due to the curvature of the earth, must be angled 3° downwards to direct it towards Soudan. The event vertex is required to be within 80 cm of the neutrino beam centre-line. Events with their vertex before the 14th plane or beyond the 68th are vetoed. This volume ends 3 m upstream of the spectrometer region, ensuring shower containment within the calorimeter.

4.2 Flux reweighting

Comparisons of near detector data with the Monte Carlo energy spectra show fairly significant disagreements, especially on the falling edge of the peak. An effect due to neutrino cross-section uncertainties, or a feature of the reconstruction or detector calibration, would be expected to manifest in a similar energy range in different beam configurations. Instead, while similarly-sized discrepancies are observed in these spectra, the energy dependence varies with the beam configuration. This indicates that the disagreements are due primarily to a mismodelling of the beam flux.

---

[1] The distance of this cylinder from the detector edge is never less than 35 cm.

A reweighting procedure is thus adopted for simulated neutrino events, based on a reweighting of the momentum distribution of the parent hadrons, with parameters based on a fit to near detector data. The fit procedure is described in detail in [122]. This reweighting is not technically necessary for the analysis, since the extrapolation procedure described in Section 4.4 successfully corrects the far detector prediction for beam mismodelling[2], but the consistent use of these weights improves data/MC agreement for other studies and simulates the effect of the full extrapolation when developing analysis components.

The fit proceeds by parameterizing the (p_z, p_T) distribution of hadrons produced from the NuMI target in the beam simulation and allowing the parameters to vary. The best parameters are obtained from a fit to the observed near detector spectra. Spectra from runs taken in different beam configurations are used, since these each probe different regions of (p_z, p_T) space. Additional fit parameters account for focusing uncertainties in the magnetic horns (the horn current and distribution) and for the effects of radiation damage to the target observed in Runs II and III. Parameters are also included to account for detector effects: the absolute energy scale, antineutrino cross-sections, and the fraction of neutral-current background. These are necessary to prevent the beam being incorrectly reweighted in an attempt to account for Monte Carlo disagreements due to unknown detector response and neutrino cross-sections. Penalty terms are included in the fit to keep the parameters within their known error ranges, including a constraint on the π^+/π^- ratio from the NA49 experiment [123].

Figure 4.1 shows the improvement in near detector data/Monte Carlo agreement obtained by applying the best fit beam weights. The fit makes a correction to the shape of the focusing peak, and a large increase to the flux in the high energy tail.

---

[2] Predictions made with and without the flux weights applied are found to be virtually identical.

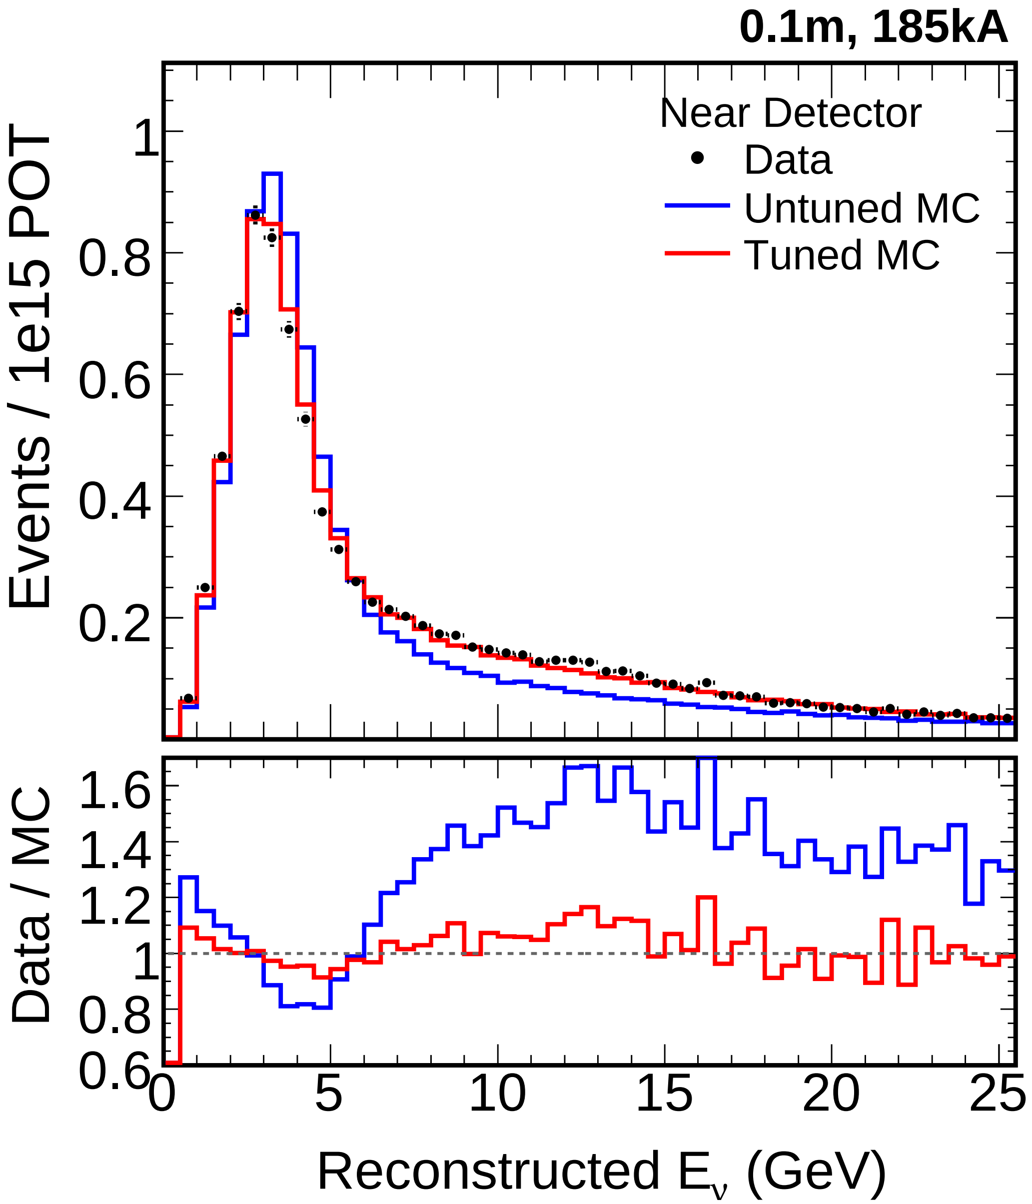

Figure 4.1: Comparison of the simulated Run III near detector charged-current energy spectrum with data (points), before (blue) and after (red) the flux tuning described in Section 4.2 is applied. The bottom panel shows the data/MC ratio before and after tuning.

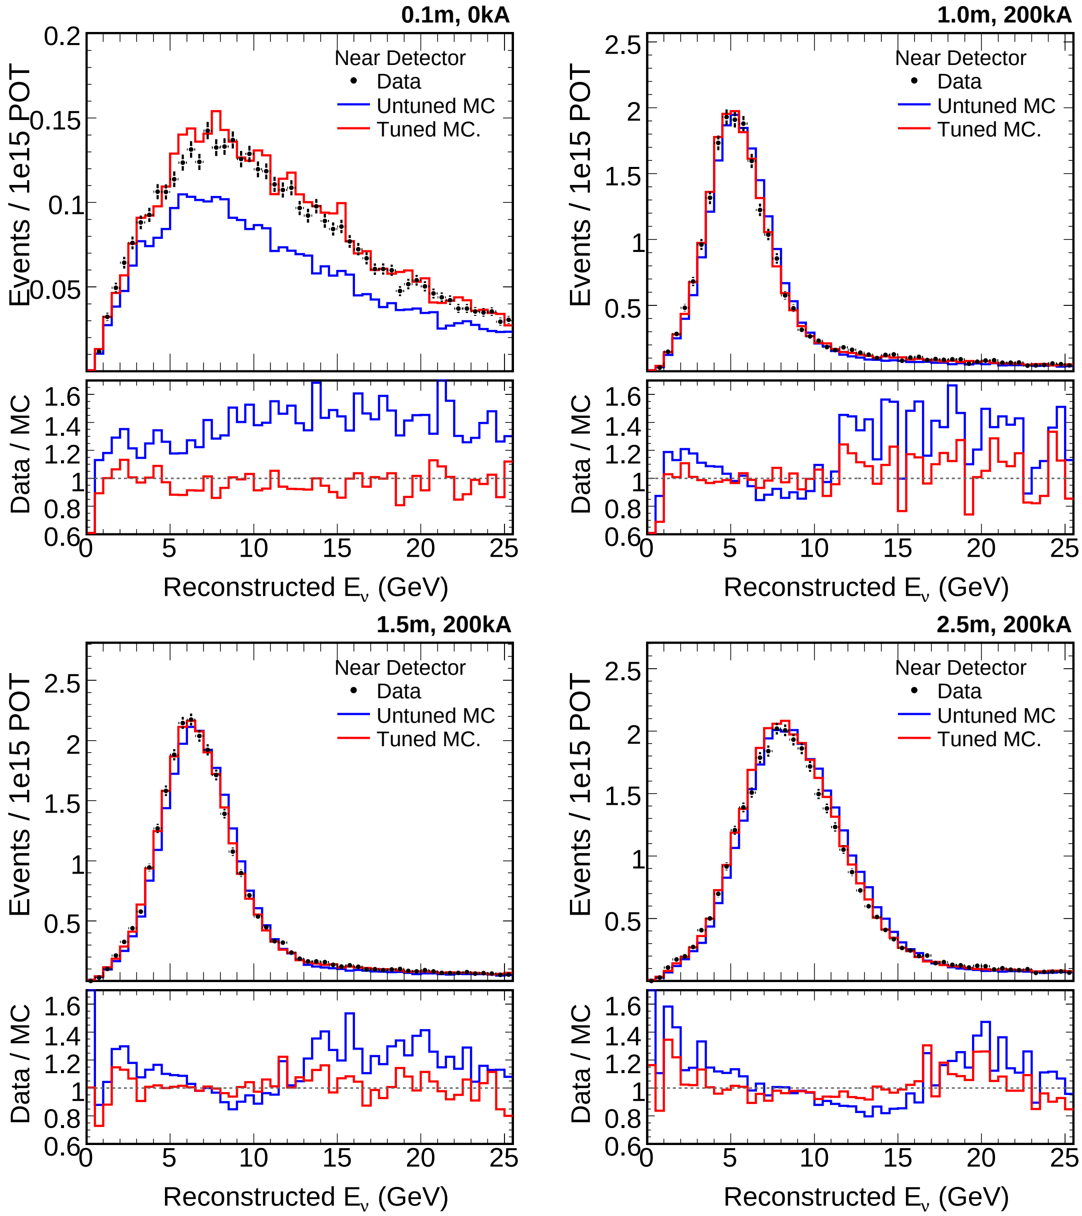

Figure 4.2: Comparison of untuned (blue) and tuned (red) simulated spectra with near detector data (points) for four non-standard beam configurations. In the top-left panel: the target is in the usual low-energy position, but the current in the focusing horns is off. The other three panels show horn-on running with the target 1, 1.5, and 2.5 m back from the first horn. In each case the flux tuning improves agreement with the observed data.

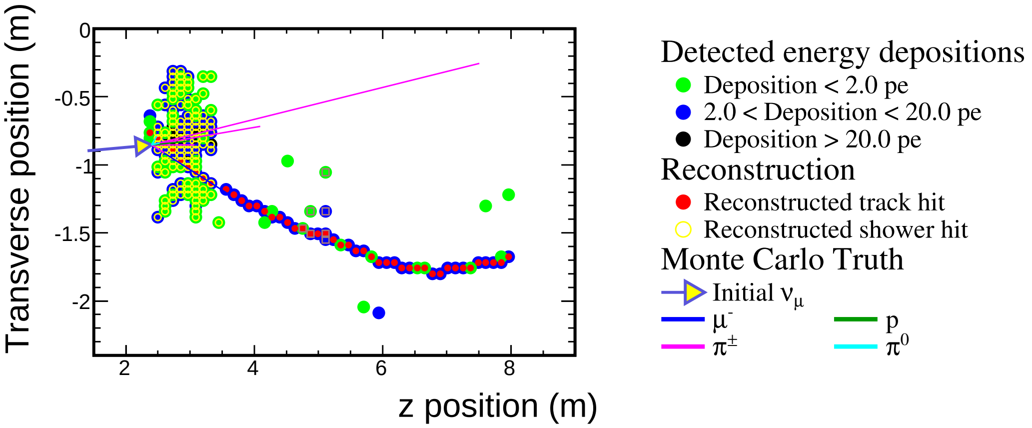

Figure 4.3: An illustrative Monte Carlo charged-current ν_μ event. The circles mark hit strips, the colour scheme indicates the energy deposited in photoelectrons as shown in the legend. The true interaction point and the momenta and identity of the initially produced particles are marked by lines. Hits reconstructed into a track are marked in red, and into a shower in yellow. The reconstruction has correctly identified the long muon track characteristic of charged-current events. A large amount of hadronic activity is also visible.

The agreement with data is similarly improved for other beam configurations (Figure 4.2). The reweighting procedure also provides an estimate of the error envelope on the Monte Carlo energy spectra, reflecting the remaining freedom in the fit parameters.

4.3 Particle identification

Since the couplings of all neutrino flavours with the Z^0 are identical, neutral-current events carry no information about $\nu_\mu \leftrightarrow \nu_\tau$ oscillations. It is therefore necessary for a study of ν_μ disappearance to identify a pure sample of charged-current events, without sacrificing efficiency.

Figures 4.3 and 4.4 illustrate the important properties of charged- and neutral-current events in MINOS[3]. The unique feature of a charged-current event is the charged muon track. In Figure 4.3 this track is very obvious, in events with a lower

---

[3] The rate of ν_e-CC events is sufficiently low that they are ignored in this discussion, they are similar in appearance to NC events.

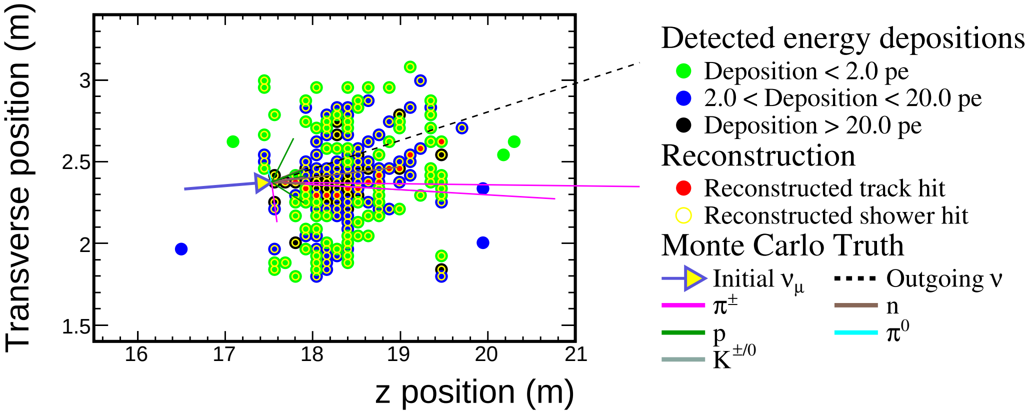

Figure 4.4: An illustrative Monte Carlo neutral-current event. Details of the event display are as in Figure 4.3. The hits marked in red are believed, incorrectly, by the reconstruction to constitute a muon track.

neutrino energy or a larger inelasticity (y-value) the track may be much shorter. The PID described here requires all selected events to have a reconstructed track. However, as Figure 4.4 illustrates, not all events with reconstructed tracks are due to charged-current interactions. The goal of the PID algorithm in MINOS is to evaluate candidate tracks and evaluate the likelihood that they are left by a muon or that they are due to some other particle or to no one individual particle.

The primary CC selector used is unchanged from that used for the previous charged-current analysis [69]. A full description can be found in [124]. A preselection requires a candidate event to have a reconstructed track of 10 planes or more. The four input variables that provide separation between neutral-current and charged-current events are shown in Figure 4.5. They are:

- Number of muon scintillator planes – Muons leave long tracks as they smoothly lose energy. In general, hadronic particles interact strongly in the detector and travel a shorter distance.
- Mean energy deposited per strip – Muons deposit a near-constant amount of charge (one MIP) in each plane. The multiple particles of hadronic showers give a higher mean energy deposition. The first 30% of planes after the track

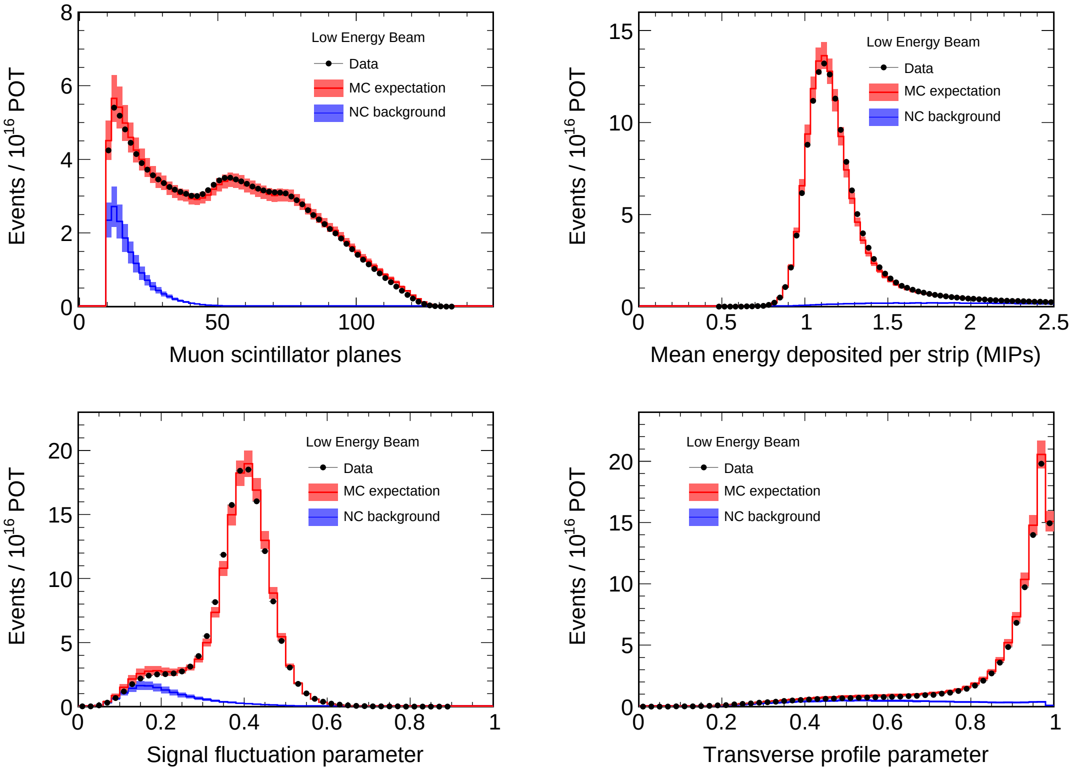

Figure 4.5: Distributions of the four input variables to the primary CC selection. The red histograms show the Monte Carlo prediction, with the systematic error associated with the beam flux indicated by shading. The expected distribution of the neutral-current background is shown in blue. Each variable shows strong identifying power. The black points show the distributions in near detector data. Deviations from the Monte Carlo predictions are within expectation.

vertex are excluded in the calculation of the mean.

- Signal fluctuation parameter – Tracks due to hadronic activity have larger fluctuations in their energy deposits than muons. All hits within 4 strips of the reconstructed track, after the first 30% of planes, are sorted by pulse-height. The variable is calculated as the mean of the lowest pulse-height strips divided by the mean of the highest strips. A larger value indicates more uniform energy deposits.

- Transverse profile parameter – Typical muon tracks deposit energy in only one strip per plane, whereas showers are wider. The variable is formed as

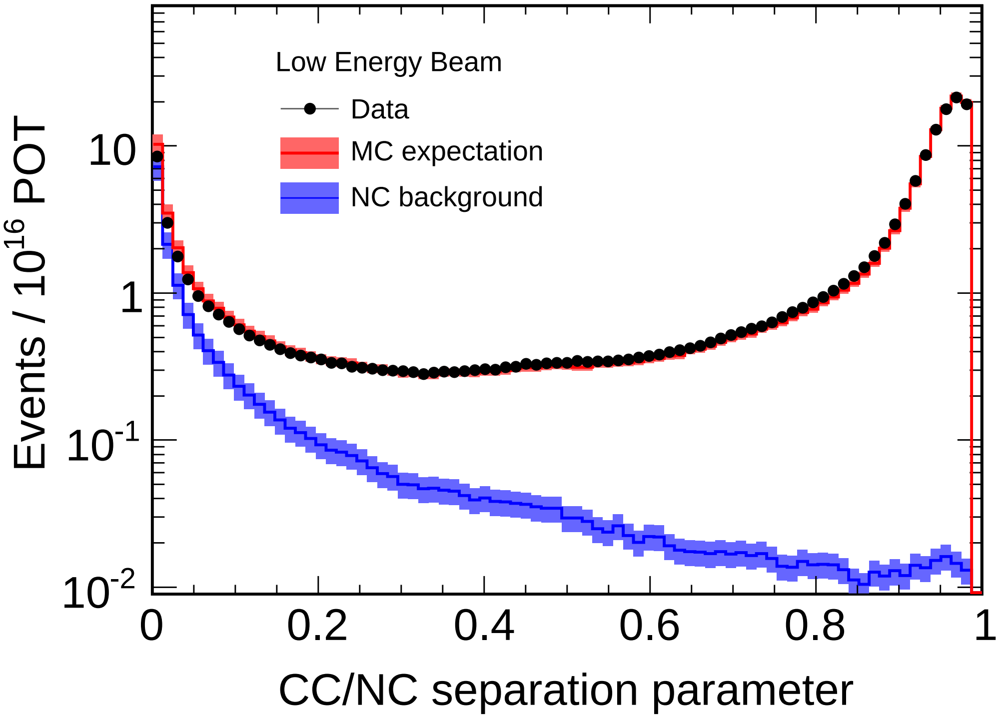

Figure 4.6: Output of the kNN constructed from the variables shown in Figure 4.5. The majority of NC events are found at low PID values. The data/MC agreement of the kNN output is good.

> the fraction of energy within a 4 strip window of the track that is assigned to the reconstructed track. The 50% of track planes closest to the vertex are excluded.

Each individual variable exhibits a strong separation between signal and background. The information provided by all four variables is combined by means of a k-Nearest-Neighbours (kNN) technique. A training sample of events is formed from Monte Carlo and for each event to be classified the k (in this case 80) training events nearest in the parameter space are found. The PID output is then simply the fraction of these events that were due to a ν_μ-CC interaction. See Figure 4.7 for an illustration. kNNs are discussed further in Section 5.2. The distribution of the kNN output (R) is shown in Figure 4.6. Optimizing for the best combination of efficiency and purity $\epsilon \times p$ [125] the best cut position is found to be $R > 0.3$.

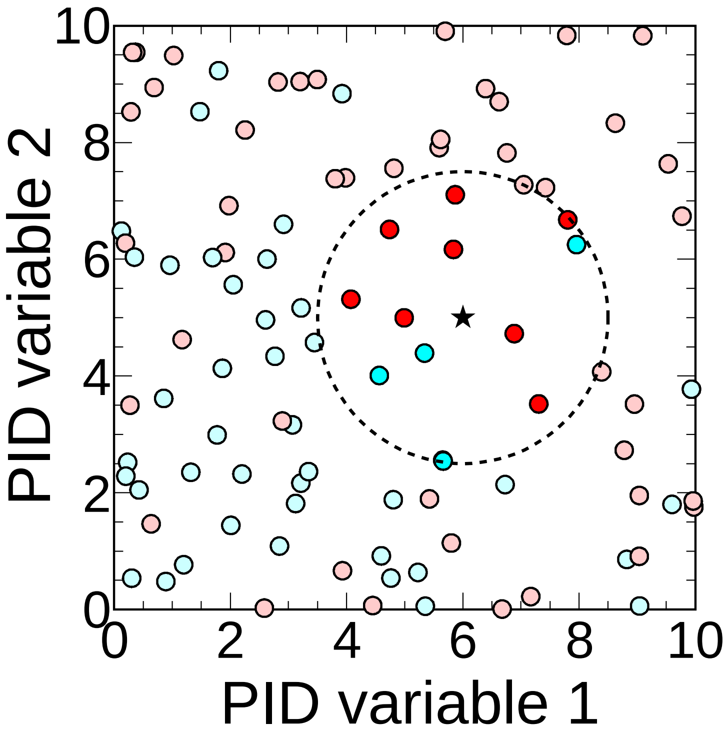

Figure 4.7: Cartoon of the use of a kNN for particle identification. In this two-dimensional space, the red signal events correlate with larger values of the variables than blue background events. The identity of the event marked by the star is estimated by examining its nearest neighbours, marked by more intense colouring. In this case $k = 12$ and the value of the PID output variable is $\frac{2}{3}$.

For the present analysis, an additional PID was developed to improve selection efficiency at low energies. This PID is constructed in a similar fashion to the primary selector, using a four-dimensional kNN. A full description can be found in [126] and [127]. As this selector focuses on low energy events the requirement for a minimum number of track planes is lifted. The input variables are shown in Figure 4.8.

- Number of muon scintillator planes – Defined as for the equivalent variable in R.
- End pulse height – A hadronic track that terminates in a nuclear interaction may lead to a large energy deposit at the end of the track, which does not occur for muons. This variable consists of the total energy deposit in the

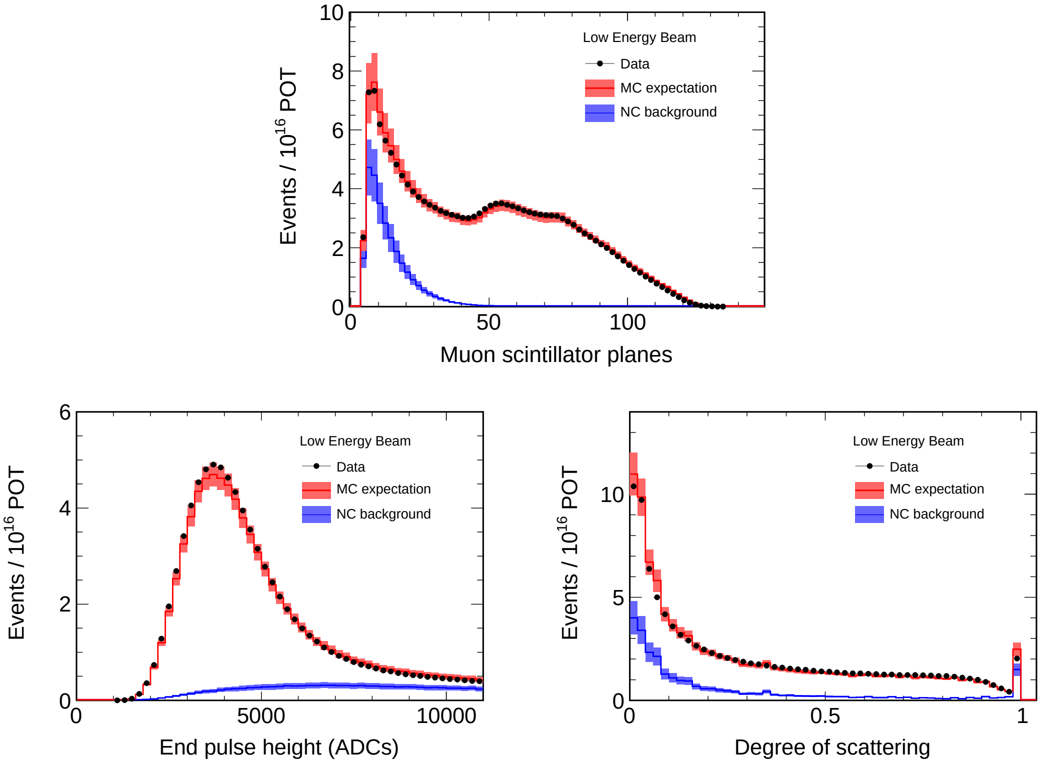

Figure 4.8: Distributions of the input variables to the auxiliary CC selection kNN. Conventions are as in Figure 4.5.

last five planes of the track.

- Degree of scattering – Muon tracks are relatively smooth and straight, whereas a hadronic track may undergo more scattering. A track constructed incorrectly from unrelated hits is also likely to be less straight. The variable is calculated from a combination of the Pearson (correlation) coefficients of the track in each view. The coefficients calculated in the U and V views are used as independent input variables.

Figure 4.9 shows the the distribution of the kNN output (J) for these variables.

The final PID is constructed as a logical-OR between these two selectors. That is, an event that passes either is considered selected. With the addition of this new selector, the cut positions of both R and J need to be re-optimized. In addition to oscillation sensitivity, the discriminatory power against neutrino decay and

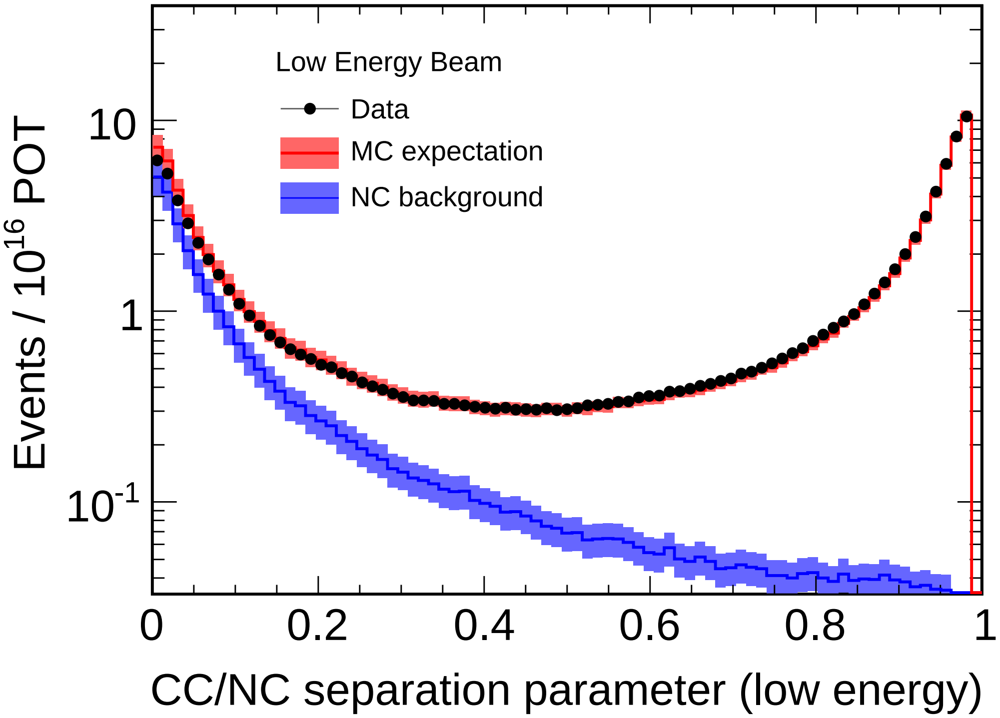

Figure 4.9: Output of the kNN constructed from the variables shown in Figure 4.8.

decoherence scenarios was also considered in the optimization [126]. The final selection criterion obtained was somewhat looser than the optimal selection with R alone:

$$(R > 0.25) \vee (J > 0.5)\ . \tag{4.1}$$

The selection efficiency and level of contamination achieved by this overall selector in far detector Monte Carlo is shown in Figure 4.10 as a function of reconstructed energy. Above 1 GeV the selection performs very strongly.

4.4 Beam matrix extrapolation

In order to take advantage of the two-detector design of MINOS, it is necessary to have some method of adjusting the far detector prediction based on information from the near detector (referred to as “extrapolation”). The “beam matrix”

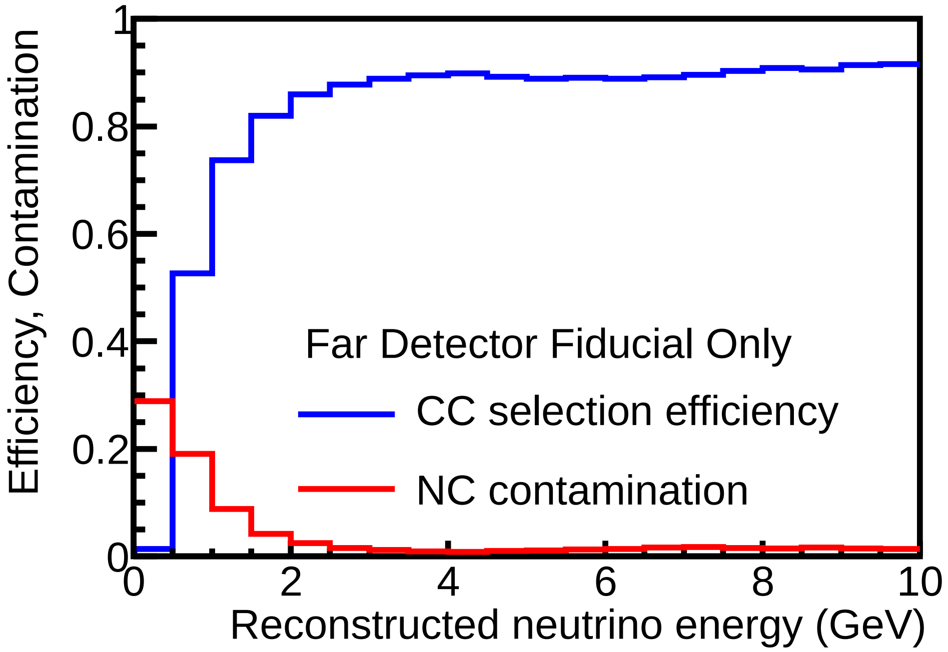

Figure 4.10: Efficiency (blue) and contamination (red) achieved by the overall charged-current selection at the far detector as a function of reconstructed energy. Above 1 GeV the selection achieves high efficiency and low contamination.

method adopted is identical to that used in the previous charged-current analysis[4], and is described in [128].

The neutrino flux seen by the two detectors is not identical, purely for geometric reasons. As illustrated by Figure 4.11, due to its proximity, the near detector subtends a much greater angle than the far detector at the point of pion decay. The effect is that neutrinos from lower energy pions, decaying at larger angles, can still give rise to interactions in the near detector but will likely miss the far detector entirely. Figure 4.12 shows this same effect in terms of neutrino energy spectra. The range of pions that give rise to neutrinos of certain energies in the near detector contribute a somewhat different spectrum of neutrinos which are destined to intersect the far detector.

---

[4] Except for the minor complication introduced by resolution binning, described in Section 4.5.

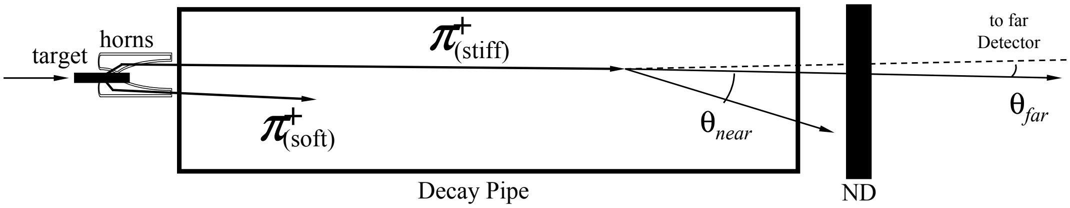

Figure 4.11: Sketch of the geometry of pion decay in the NuMI decay pipe. Due to its proximity to the decay pipe the near detector subtends a much larger angle than the far detector. Neutrinos emitted at large angles may miss the far detector while still intersecting the near detector. Relatively “softer” (lower energy) pions thus contribute a larger fraction of the observed near detector neutrino interactions.

The core of the extrapolation is a near to far transfer matrix, or “beam matrix”. This encodes the necessary geometric and kinematic information to transform an observed near detector flux into a predicted far detector flux. For each parent decay in the beam simulation, a random interaction vertex is picked in the near detector. The decay kinematics then determine the neutrino energy this geometry corresponds to. The same procedure is carried out for the far detector, although picking an interaction point is unnecessary since all points in the far detector are virtually straight down the z-axis. This procedure associates a neutrino energy at the near detector to one in the far detector via their shared parent. Repeating this procedure for all the simulated parents, and taking into account the angular dependence of the decay cross-sections, produces a beam matrix as shown in Figure 4.13. This procedure is carried out separately for each run period to take into account differences in the beam properties, and independently for neutrinos and antineutrinos to account for the differing parent particles. There is, of course, a strong correlation between neutrino energies in the two detectors, meaning the matrix is nearly diagonal, but each near detector energy does correspond to a distribution of far detector energies.

The near detector does not directly measure neutrino flux. Likewise the far detector prediction is not required in the form of a flux, but as a visible energy

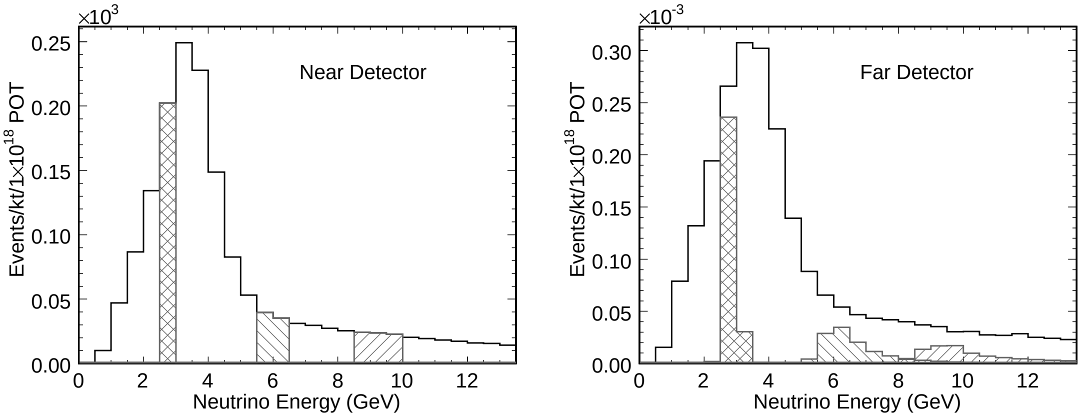

Figure 4.12: Comparison of simulated near and far detector energy spectra in true neutrino energy. The shaded regions in each histogram correspond to the same set of parent meson decay positions and momenta. The difference in geometrical acceptance between the two detectors leads to the observed smearing and difference between the shapes of the two spectra. The neutrino energy in the beam peak is slightly higher in the far detector, since neutrinos from the decay of high energy parents are emitted in a more forward direction and thus more likely to be directed towards the near detector. The difference in the y-axis scales by a factor of approximately 10^6 reflects the divergence of the neutrino beam between the two sites.

distribution. A further series of steps are thus required to complete the extrapolation procedure. Additional matrices, derived from Monte Carlo simulations, remove neutral-current contamination from the observed near detector spectrum, convert it to true neutrino energy, and correct for reconstruction efficiency. The near detector flux thus obtained is then extrapolated to the far detector via the beam matrix and the equivalent matrices are applied in reverse to generate a far detector prediction in reconstructed energy. While the far detector prediction is expressed in neutrino energy the neutrino oscillation probabilities may be applied. The full procedure is illustrated in Figure 4.14.

This extrapolation procedure is designed to use the near detector spectrum as a measurement of the neutrino flux, and thus correct systematic errors due to uncertainty in the beam simulation. Due to the nearly direct transfer of the near detector observation to the far detector prediction (the beam matrix is near-

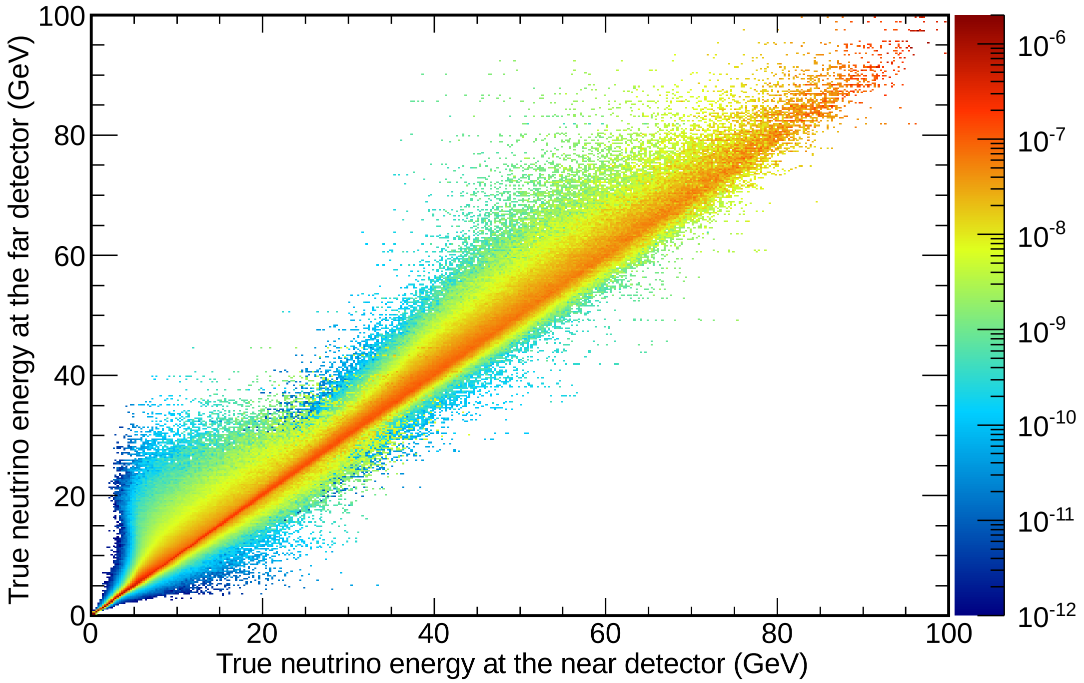

Figure 4.13: The beam matrix calculated from Run III Monte Carlo. Application of this matrix converts a neutrino flux at the near detector into a predicted neutrino flux at the far detector. This is essentially a two-dimensional version of Figure 4.12. See main text for more details.

diagonal) it also serves to propagate reconstruction or cross-section effects from one detector to the other, thus serving to, at least partially, account for these factors in its prediction as well.

4.5 Resolution binning

The better the energy resolution of the far detector spectrum to be fitted for oscillations, the more clearly the oscillation dip will be defined. A more sharply defined dip allows less freedom for the fit to vary oscillation parameters and thus leads to tighter constraints on the Δm^2 and $\sin^2 2\theta$.

The topic of directly improving the resolution of reconstructed events is covered in Chapter 5. This section describes the complementary procedure of dividing

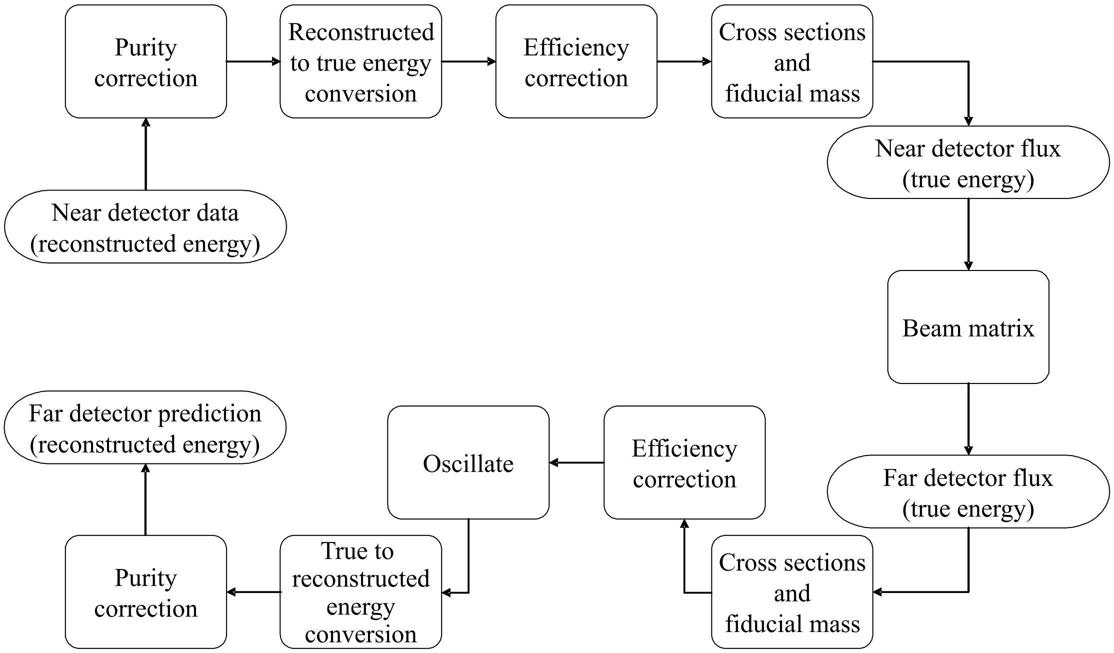

Figure 4.14: Flowchart of the near detector to far detector extrapolation procedure. The observed near detector energy spectrum is background-subtracted based on the Monte Carlo background expectation. The spectrum is then converted to a true energy distribution, and selection and reconstruction efficiencies taken into account, again using Monte Carlo. Dividing out cross-sections and the near detector mass provides an estimate of the true ν_μ flux at the near detector. This is then transformed into a far detector flux prediction using the beam matrix, and the corresponding corrections all applied in reverse to finally produce a prediction of the energy spectrum observed in the far detector.

reconstructed events according to an estimate of their energy resolution.

The reconstructed energy of an event is formed by taking the sum of the reconstructed track and shower energies. The resolution estimate is likewise formed from a combination of the track and shower resolution estimates. The Monte Carlo sample of selected truly-CC events was used to construct $E_{\text{reco}} - E_{\text{true}}$ distributions for the track and shower energies. A gaussian fit was made to each distribution and the fitted standard deviation used as the resolution estimate. For showers, and tracks whose energy is measured from range, this procedure was repeated for each bin of reconstructed energy, and the results parametrized with a functional form consisting of the sum in quadrature of terms linear in the energy and in the square-root of the energy, and for showers an additional constant term. For the

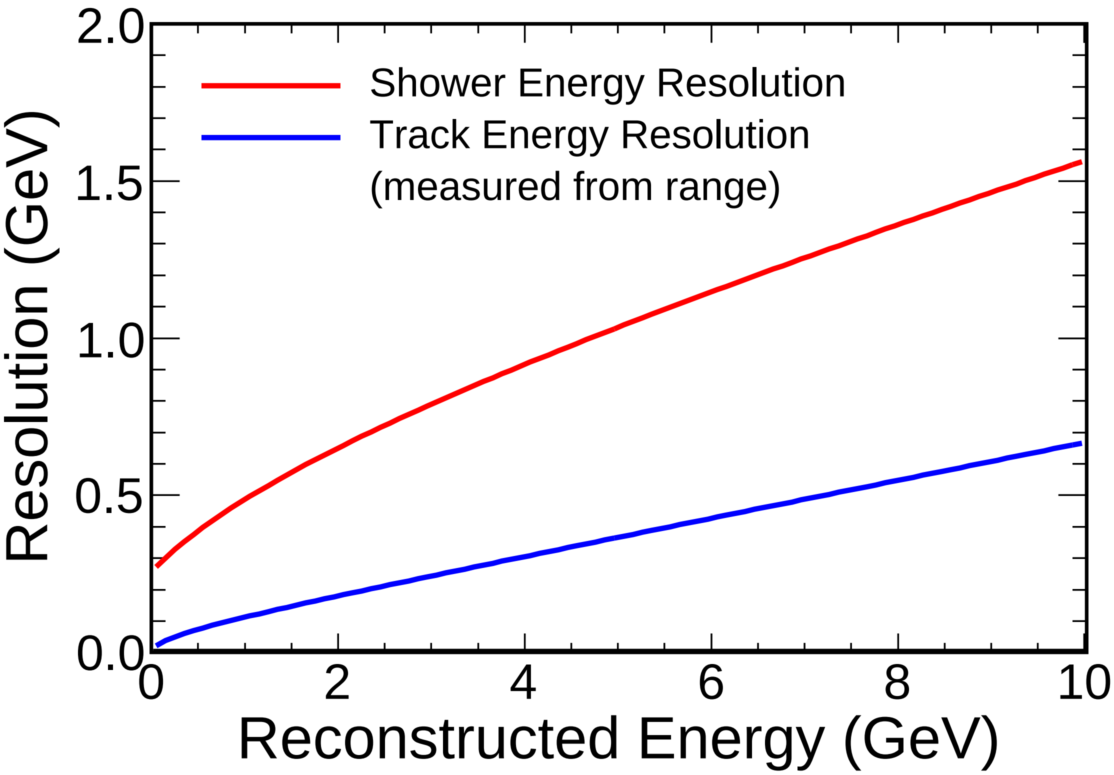

Figure 4.15: Track and shower energy resolution functions as described by equations (4.2).

case of tracks whose energy is measured from curvature, the independent variable was the product $p^2_{\text{curv}}\sigma_{q/p}$, where p_{curv} is the momentum estimate from the tracker, and $\sigma_{q/p}$ is its estimate of the error on the charge : momentum ratio. This fitting procedure is described in more detail in [129]. The resolution functions determined by this procedure are:

$$\begin{aligned}\sigma^{\text{range}}_{\text{trk}} &= 6.9\%\ E_{\text{range}} \oplus \ \ 5.1\%\ \sqrt{E_{\text{range}}} \\ \sigma^{\text{curv}}_{\text{trk}} &= \ \ p^2_{\text{curv}}\sigma_{q/p} \ \ \oplus\ 134\%\ \sqrt{p^2_{\text{curv}}\sigma_{q/p}} \\ \sigma_{\text{shw}} &= \ 8.6\%\ E_{\text{shw}} \ \oplus 40.4\%\ \sqrt{E_{\text{shw}}} \qquad \oplus\ 275\,\text{MeV}\ ,\end{aligned} \tag{4.2}$$

where the units of energy are GeV. Figure 4.15 shows the form of the shower and range-measured track functions. The overall resolution estimate for an event is then formed from the sum in quadrature of the track and shower resolution

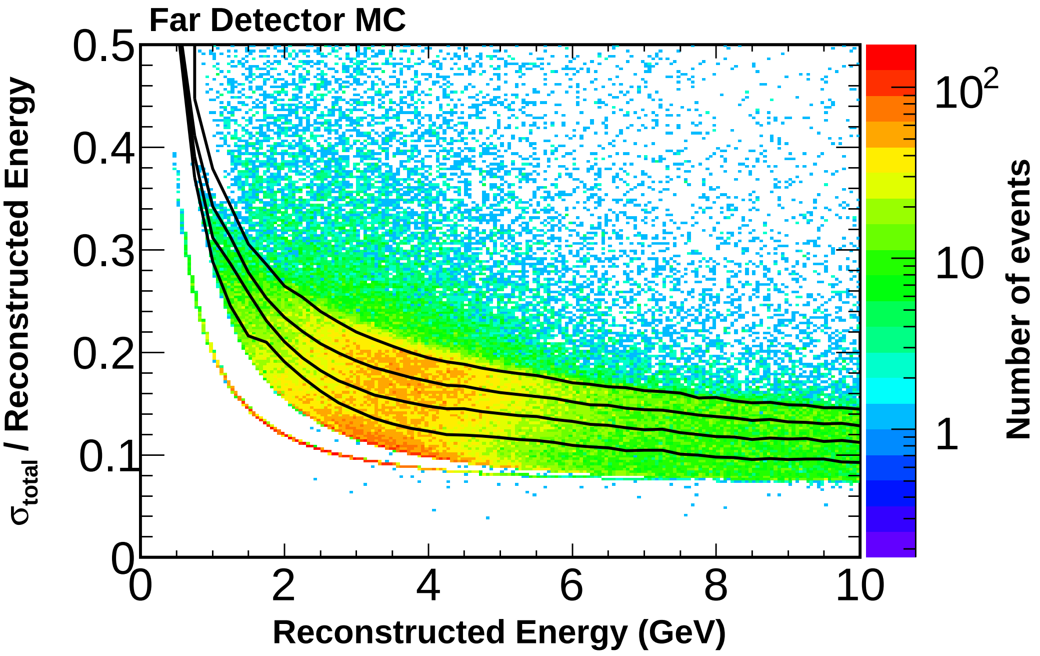

Figure 4.16: The distribution of energy resolution σ_{tot} as a function of reconstructed energy. The black lines mark the quantile boundaries. At each energy 20% of simulated events fall in each quantile.

components:

$$\sigma_{\text{tot}} = \sigma_{\text{trk}} \ \oplus \ \sigma_{\text{shw}} \tag{4.3}$$

Using this parametrization, it is possible to divide the energy spectrum of far detector events into bins of differing resolution. The choice of binning made is to have five resolution quantiles, each with equal statistics in the unoscillated case. Figure 4.16 shows the distribution of estimated energy resolution against total reconstructed energy for unoscillated far detector Monte Carlo. The black lines represent the cuts used to separate the events. The cut positions vary as a function of reconstructed energy, and are defined so as to divide the event sample into five equal components at each value of reconstructed energy. The unoscillated far detector predictions for each resolution bin are thus identical by construction. Separate cuts are constructed for each individual run period. This procedure is described in [130] and [131].

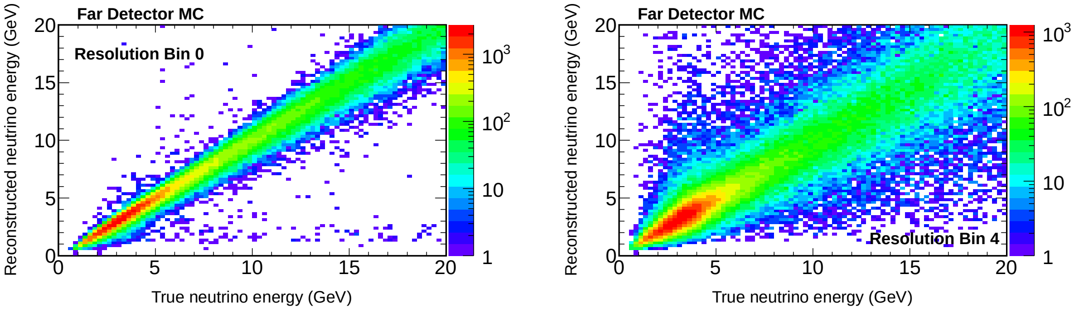

Figure 4.17: Distribution of reconstructed energy versus true energy for the best 20% of events (left) and the worst 20% (right).

Events in which the muon track is reconstructed to have positive charge are excluded from the resolution binning scheme and allocated to a separate spectrum – making six in all. Many of these events are indeed from antineutrino interactions, but at low energies charge misidentification is common, and this sample also includes a large fraction of neutrino events.

The near detector spectrum is divided according to reconstructed muon charge, but not by estimated event resolution. The extrapolation procedure described in Section 4.4 is carried out independently for each of the six subsamples, starting from either the positive or negative near detector spectrum as appropriate, reflecting the different characteristics of the events making up each, producing six far detector predictions for comparison with far detector data for each run period. Figure 4.17 shows the reconstructed-versus-true energy matrices employed in the extrapolation procedure for the highest and lowest resolution quantiles. The large difference in the accuracy of energy reconstruction between the samples is clear.

Dividing the far detector energy spectrum in this way leads to superior oscillation parameter sensitivity. The highest resolution events most clearly define the oscillation dip, and are allowed to do so uncontaminated by events with poorer energy resolution, which would tend to smear out the oscillation signal. By maintaining the poorer events in their own spectra, the information they provide about

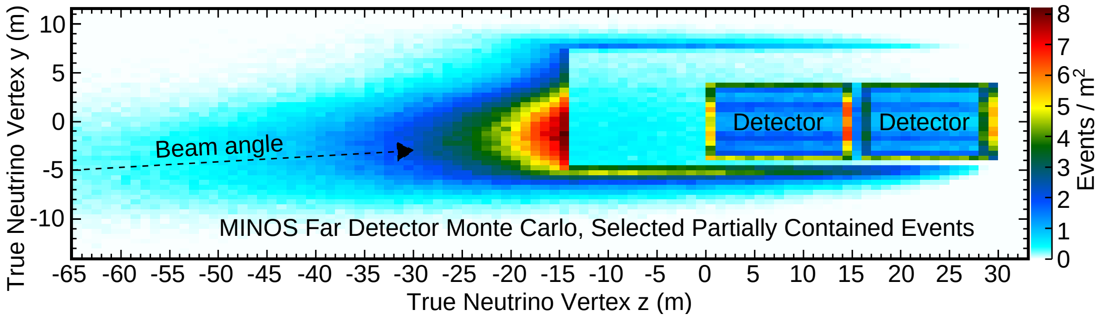

Figure 4.18: Distribution of true interaction vertices of Monte Carlo events selected into the rock and antifiducial sample. The structure of the far detector supermodules and the cavern are clearly visible. The majority of interactions occur in the rock of the upstream cavern wall.

the oscillation parameters is also fully utilized. A side-effect of the resolution binning is that neutral-current background events are mostly relegated to the two worst resolution bins. The higher resolution spectra thus also draw from a purer sample of charged-current events, again allowing a more precise determination of the oscillation parameters.

4.6 Rock and antifiducial events

In addition to neutrino interactions within the fiducial volume, the MINOS far detector also records events originating from neutrino interactions in the outer regions of the detector or in the rock of the surrounding cavern walls. Figure 4.18 shows the distribution of true interaction vertices for these events. The largest contribution comes from interactions in the upstream wall of the far detector cavern. In previous analyses, these events have been excluded from the sample. Events originating in the detector but outside the fiducial volume ("antifiducial events") were excluded because shower containment is not guaranteed, and because they are difficult to distinguish from events originating in the surrounding rock. Events originating in the cavern walls ("rock events") were rejected because the shower almost certainly does not reach the detector, and some unknown portion of the track energy is also dissipated in the rock. For this analysis events from these classes were included. Whilst the energy resolution of these events is poor, they comprise a sample with comparable statistics to the fiducial sample, and provide a valuable additional contribution to the oscillation measurement. Full details of the treatment of these events are given in [132] and summarized below.

Due to the uncertain or nonexistent containment of hadronic showers in the sample of rock and antifiducial (RAF) events, the energy estimator used, and the variable in which the energy spectra are binned, is chosen as simply the visible track energy (from range or curvature as in the fiducial analysis). As well as neglecting the energy in the hadronic shower, this variable is also unable to account for any energy lost by the muon as it travels through the rock of the cavern walls. However, the visible muon energy is the maximum amount of information available about the true neutrino energy, and some correlation remains. In case of particles from a hadronic shower penetrating the detector, the selection of muon-like tracks is made using the same PID variables, R and J, described in Section 4.3.

While both classes of event are hampered by the absence of shower information, those from interactions within the antifiducial volume have much better energy resolution than events whose true vertex is in the surrounding rock. In a rock event the muon has lost an unknown fraction of its energy in transit to the detector. Splitting the RAF sample according to energy resolution is advantageous for the same reasons as presented in Section 4.5. The RAF sample is therefore divided into six spectra based on the detector region in which the first track hit is recorded. These regions are illustrated in Figure 4.19. The types of events recorded in each region are determined by geometrical considerations. The rate of rock muons depends on the surface area exposed by each region to events originating in the rock, whilst the detector mass contained in each region governs the rate of neutrino interactions within them. Events detected in the region of the detector front face are predominantly rock events, those in the gap and back regions are

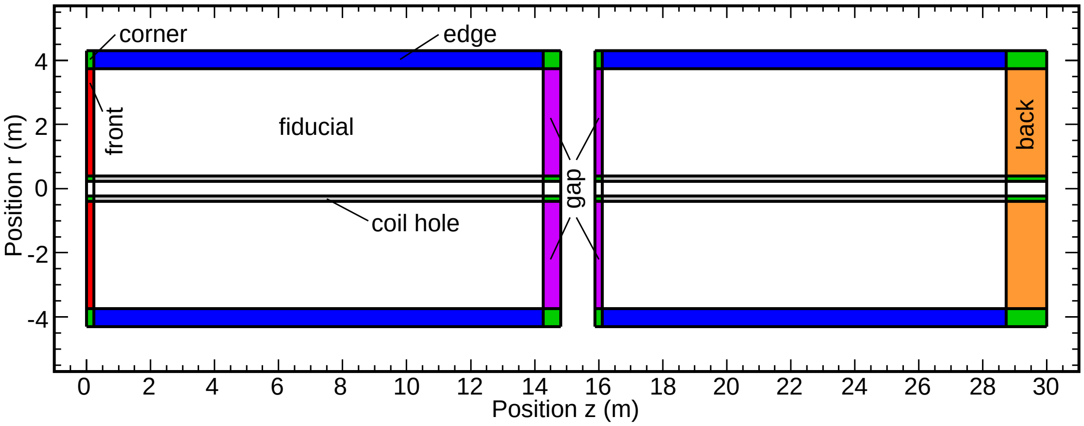

Figure 4.19: Scale diagram of the various far detector regions as considered in the analysis. Events with reconstructed vertices in the coloured regions are placed in the corresponding rock and antifiducial spectrum. Events with vertices in the fiducial volume go into the main sample, divided according to charge and resolution.

predominantly antifiducial interactions. The detector edge regions have both a large mass and a large surface area exposed to the rock, they thus collect a mixture of events. The coil hole region is not used in this analysis. Edge events are further separated into rock-like and detector-like samples based on consideration of the precise sequence of hits at the start of the track and of the geometry of the scintillator strips at the detector edge.

With an event selection and binning scheme tuned to give maximum oscillation sensitivity, the final necessary component of the RAF analysis is an extrapolation technique in order to take advantage of the information provided by the near detector. Attempting to use rock and antifiducial events from the near detector is technically complex, and is not in general an effective source of information about the corresponding far detector spectra. The geometry of the near detector and its cavern differ from the far detector, as does the geometrical relationship to the beam. The near detector cavern and the mass distribution within it are also not well modelled in the Monte Carlo. The composition and thus density of the rock also differs at the two sites. For these reasons, the near detector fiducial spectrum is used as the basis of the extrapolation procedure. The simulated near detector spectrum, and that actually observed in data, are put through the extrapolation procedure described in Section 4.4 and illustrated in Figure 4.14. The procedure is halted when the far detector flux prediction is obtained, and the ratio between the flux predicted from data and that from Monte Carlo is applied as a correction factor to the far detector RAF Monte Carlo as a function of true neutrino energy. Whilst different in conception to the extrapolation of the fiducial sample this procedure is effectively equivalent, the various correction factors are simply handled implicitly by the Monte Carlo instead of being broken out explicitly into separate histograms.

The RAF analysis obtains good sensitivity to Δm^2, comparable to the fiducial analysis, but very poor sensitivity to $\sin^2 2\theta$. Considering that Δm^2 governs the position of the oscillation dip and $\sin^2 2\theta$ its magnitude, it may seem counter-intuitive that an analysis with poor energy resolution and good statistics behaves in this way. However, due to the position of the oscillation dip on the rising edge of the NuMI beam peak, variations in Δm^2 affect the expected event rate considerably. Determination of the depth of the oscillation dip, meanwhile, requires good energy resolution to probe the oscillation maximum clearly.

4.7 Systematics and fitter

Comparison of the far detector prediction with the observed energy spectrum in data is performed using a binned maximum likelihood method. Specifically, the best fit oscillation parameters are found at the minimum of the log-likelihood ratio for Poisson-distributed data [46]:

$$\begin{aligned} \chi^2 &= 2\sum_i e_i - o_i + o_i \ln\left(\frac{o_i}{e_i}\right) \qquad (4.4) \\ e_i &= e_i(\sin^2 2\theta, \Delta m^2) \ , \end{aligned}$$

C.L.	1D	2D
68%	1.00	2.30
90%	2.71	4.61

Table 4.1: $\Delta\chi^2$ values denoting the 68% and 90% confidence levels in one- and two-dimensional parameter spaces. Excerpted from [46].

where the summation over i includes each bin of energy for every spectrum under consideration (six fiducial and six antifiducial for each of the four run periods[5]). o_i represents the number of data events observed in each bin and e_i the extrapolated prediction. The expectation in each bin is a function of Δm^2 and $\sin^2 2\theta$, calculated by reweighting each truly charged-current event as a function of true energy, according to Equation (2.17). The small effect of ν_τ appearance is taken into account by adding a spectrum of similarly reweighted tau events to the prediction.

Confidence levels may be drawn in the two-dimensional $(\sin^2 2\theta, \Delta m^2)$ space based on the $\Delta\chi^2$ from the best fit point, also known as the "up-value". The $\Delta\chi^2$ values required to cover the true oscillation parameters with 68% or 90% probability are given in Table 4.1. The two-dimensional likelihood surface and its resulting contours may be converted into a one-dimensional $\Delta\chi^2$ projection and corresponding single parameter error via the technique of marginalization. For example, at each value of Δm^2 all possible values of $\sin^2 2\theta$ are scanned, and the lowest χ^2 found represents that value of Δm^2 in the one-dimensional projection. The best fit Δm^2 is found at the minimum of this curve and will correspond, by construction, with the Δm^2 value of the best-fit point. The errors on Δm^2 are found by inspecting the $\Delta\chi^2$ from the minimum, this time using the tabulated one-dimensional up-values.

Sensitivity contours and predictions of the one-dimensional parameter errors can be constructed with the use of "fake data". The full Monte Carlo sample is taken, and reweighted according to some input values of the oscillation parameters,

---

[5] The positively-charged spectrum for the pHE run period is omitted due to the very low predicted event count.

before being scaled down to an exposure equivalent to the data sample. The correct operation of the fitting procedure is confirmed by using this spectrum in place of data and ensuring that the best-fit point obtained matches the input values. The contours and single-parameter errors obtained from this procedure represent the sensitivity of the experiment at the given input parameters.

4.7.1 Treatment of systematic errors

The magnitudes of a large number of possible sources of systematic bias have been estimated and their impact on the oscillation fit calculated. For each systematic studied, a method was developed to simulate the effect of a $\pm 1\sigma$ shift in Monte Carlo. The full list of parameters and the details of how the shifts were input is given in [133].

The impact of each systematic shift on the oscillation fit is evaluated as the fit bias that would be introduced to each oscillation parameter if the data were to have an unmodelled $\pm 1\sigma$ shift. For each systematic, Monte Carlo events were modified to reflect the effect, producing fake data samples. These were input to the extrapolation and fit procedure in place of data, to be compared with an unshifted Monte Carlo sample. For systematics affecting both the near and far detectors, both were shifted correspondingly. The output best-fit point was then compared to the input oscillation values, the result reflecting the bias introduced by the systematic even in the presence of the near to far extrapolation. Figure 4.20 shows a “star plot” summarizing the impact of the various sources of error considered. The centre of the plot represents the input oscillation parameters, and each arm the deviation from these true values introduced by a $\pm 1\sigma$ shift in a single systematic. The four systematics with the greatest impact on the oscillation fit are:

- Absolute hadronic energy scale – uncertainty in the energy scale of hadronic

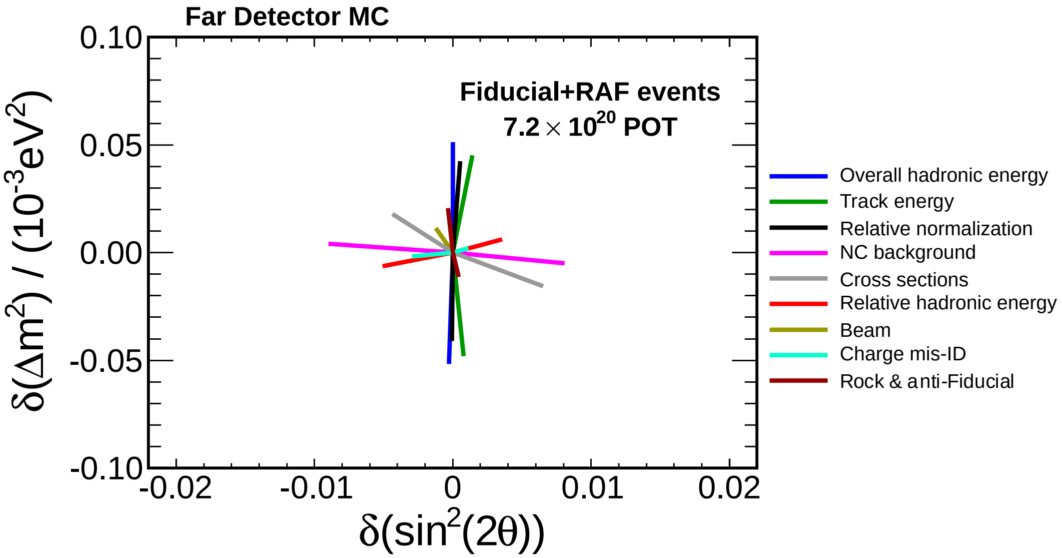

Figure 4.20: Shifts induced in an oscillation fit to Monte Carlo by the major categories of systematic error. The centre-point represents the oscillations input into the fake data. Each arm of the star indicates the difference to the best fit value obtained by a statistics-only fit to fake data based on the same oscillation parameters but including a $\pm 1\sigma$ shift in one systematic. The top four systematics in the list, representing the largest shifts in Δm^2 and $\sin^2 2\theta$, are included as nuisance parameters in the final fit.

showers due to factors common to both detectors. The derivation of the magnitude of this uncertainty is presented in Section 5.6.

- Absolute track energy scale – uncertainty in the energy scale of muon tracks due to factors common to both detectors. A 2% adjustment is made to energies measured by range, with a fully correlated 3% shift for tracks whose energies are measured from curvature. The range-based error is calculated from the combination of known uncertainties in the detector simulation and particle propagation, with the curvature error including an additional component from observed differences in the range and curvature measurements for individual tracks [134].

- Relative normalization – uncertainty in the expected number of neutrino events from factors not shared between the two detectors. A 1.54% scale factor is applied to the far detector spectrum. The uncertainty arises from a combination of relative uncertainties on fiducial mass, live-time and reconstruction efficiency between the two detectors [135].

- Neutral-current background – uncertainty in the efficiency of the charged-current selection due to mismodelling of hadronic showers and neutral-current cross-sections. A 20% scale factor is applied to the level of misidentified neutral-current events in both detectors. This value comes from a comparison of the near detector PID distribution (R) between data and Monte Carlo, and of data and Monte Carlo charged-current events re-reconstructed with the muon track removed [136, 127].

The other systematic shifts that were investigated, and found to have a smaller effect on the fit were [133]:

- Cross sections – variations of the various neutrino and antineutrino charged-current cross-sections. The contributions of the different contributing systematics are summed in quadrature for display on the star plot.
- Relative hadronic energy – uncorrelated error in the shower energy scale between the detectors. A 1.9% shift in the near detector, or a 1.1% shift in the far detector.
- Beam – the bin-by-bin error on the flux prediction, as determined by the beam reweighting procedure.
- Charge mis-ID – a 40% scale on the amount of $\bar{\nu}_\mu$ charged-current background selected into the negatively-charged event samples.
- Rock and anti-fiducial – a 0.9% normalization scale on events originating in the far detector rock, a 1% scale on the non-DIS cross-sections of events in the rock, and the uncertainty in the separation of rock-like and detector-like

events along the detector edges. These effects are summed in quadrature for the star plot.

The effect of each of these four major sources of systematic uncertainty is included in the fit, and in the calculation of the confidence limits, by allowing the fit freedom to adjust the parameters in question. The log-likelihood formula in (4.4) becomes:

$$\begin{aligned} \chi^2 &= 2\sum_i e_i - o_i + o_i \ln\left(\frac{o_i}{e_i}\right) + \sum_j \frac{s_j^2}{\sigma_j^2} \\ e_i &= e_i(\sin^2 2\theta, \Delta m^2, \vec{s}) \ , \end{aligned} \tag{4.5}$$

where the expectation in each bin e_i is now also a function of the four systematic shifts s_j. The final term in the χ^2 formula acts to constrain the systematic terms within about their estimated range of variation. The σ_j are the 1σ error ranges described above. In the absence of any preference from the data, the systematic parameters will fit best at their nominal values. In the case of evidence of a systematic shift in the data, the best-fit shift is a trade-off against our prior knowledge of the parameter. Fitting for the systematic parameters leads to a better fit to the data. The χ^2 value is marginalized over the systematic parameters $\vec{s}$ independently at all points in the $(\Delta m^2, \sin^2 2\theta)$ plane, thus reducing the height of the $\Delta\chi^2$ surface, and increasing the size of the confidence contours. This captures the increased uncertainty in the estimation of the oscillation parameters which is introduced by these sources of systematic error.

4.7.2 Systematics interpolation and fitter

Performing the oscillation fit exactly as described above would be very computationally intensive. A straightforward implementation of the track and shower energy systematics entails adjusting the energies of all Monte Carlo events in both detectors, rebuilding the histograms required for the extrapolation and recalculating the far detector predictions. This must happen many hundreds of times for each point in $(\Delta m^2, \sin^2 2\theta)$ space as the systematic parameters are searched over.

The approach adopted instead involves interpolation between representative templates. The inputs to the fitter are reconstructed-versus-true spectra as produced by the extrapolation, one for each event class (charged-current ν_μ, $\bar{\nu}_\mu$, ν_τ and $\bar{\nu}_\tau$, and neutral-current) for each resolution bin or detector region, and for each run period. These templates are produced at the nominal point, with no systematic shifts applied, and at $\pm 1\sigma$ and $\pm 2\sigma$ shifts of each parameter, yielding $4n + 1 = 17$ systematic configurations in total. Spectra with multiple shifts applied and at intermediate values are obtained by interpolation. The MINUIT package [137] is used for the minimization over the systematic parameters and avoids the necessity of a brute-force search over a four-dimensional space, but requires the input χ^2 function to be doubly-differentiable in order to reliably form the covariance matrix the algorithm depends on. A simple linear interpolation scheme between the templates would not achieve that property, exhibiting unwanted behaviour at the boundaries between each segment. The interpolation scheme used is constructed as a weighted average of the different linear interpolations between neighbouring templates. It may be described in terms of the weight given in the prediction to any one template, as a function of the input systematic shift. Figure 4.21 illustrates these weighting functions. At each integer value of the systematic shift the exact spectrum known from that template is reproduced, with intermediate values being influenced more heavily by the nearby templates. The interpolation scheme and fitter in general are described in [129]. Spectra with multiple shifted parameters were generated both by this method and by the full event-by-event procedure and the discrepancies shown to be small.

The oscillation fit and construction of contours proceeds as follows. A scan is made over a grid in $(\Delta m^2, \sin^2 2\theta)$ space. At each point, every template is

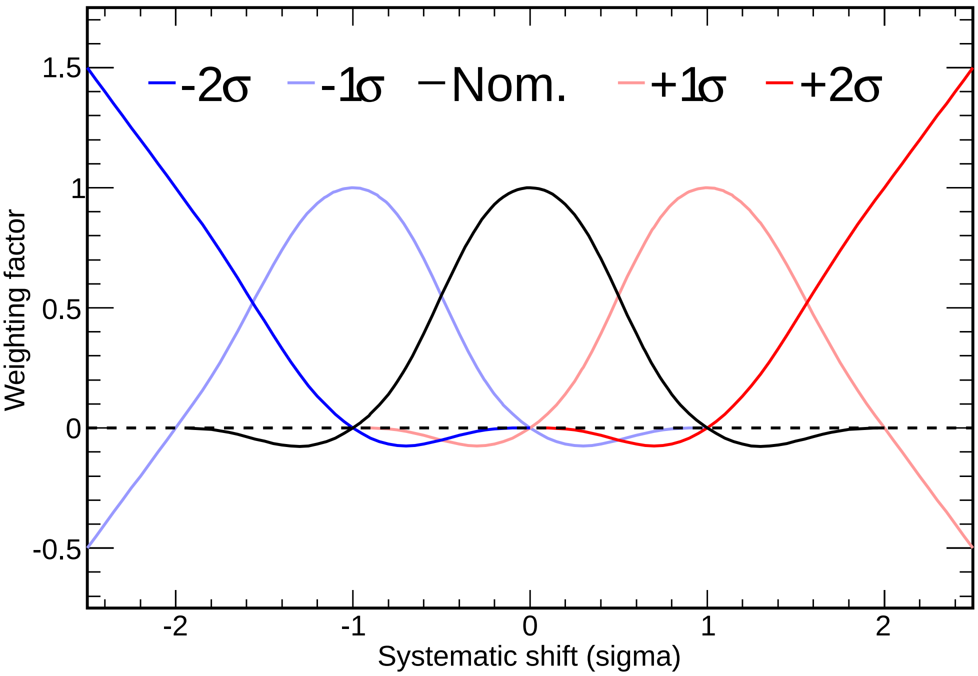

Figure 4.21: Weight assigned to each of the five template histograms, representing $0, \pm 1, \pm 2\sigma$ shifts in one systematic parameter in order to interpolate between them as a function of the input shift. At the central value for each template the weights for all other templates are zero, thus reproducing the known shifted spectra.

reweighted according to the two-flavour survival probability (2.17) at these parameters[6]. The templates are then reduced down to reconstructed energy. MINUIT is allowed to search over the four systematic parameters, with the prediction at each shifted point generated by interpolation between the oscillated templates. The minimum χ^2 found is taken as the value for these oscillation parameters, and the procedure repeats at the next point on the grid. The grid point with the lowest χ^2 is taken as the best-fit point and confidence levels are drawn according to the up values in Table 4.1.

---

[6] The fitter makes use of the linearity of the probability in $\sin^2 2\theta$ to more quickly generate predictions at the same Δm^2. The oscillations are precalculated at $\sin^2 2\theta = \{0, 1\}$. In each bin the weight is taken as the average of the weights for all Monte Carlo events populating this bin in the nominal sample. This procedure produces a more accurate representation of the oscillations than simply using the formula evaluated at the bin centre.

Chapter 5

The kNN energy estimator

This chapter presents a new technique, developed for this thesis, for the estimation of the energy content of hadronic showers in MINOS. The energy resolution of low-energy showers is significantly improved by using this new estimator in place of the usual calorimetric estimator.

The sensitivity of the MINOS charged-current analysis to the oscillation parameters Δm^2 and $\sin^2 2\theta$ is dependent on the ability to pick out the features of the characteristic oscillation dip. A more sharply resolved dip allows the oscillation fit less freedom in matching its depth ($\sin^2 2\theta$) and position (Δm^2), and results in a smaller allowed region of the parameter space. The technique of dividing events by their estimated energy resolution was discussed in Section 4.5; this chapter focuses on directly improving the resolution of the individual events. Resolution binning can of course then be applied to the result, yielding further improvements.

Sections 5.2 to 5.5 describe the details of the kNN algorithm and the procedure followed to construct an optimized estimator for use in the MINOS charged-current analysis. Section 5.6 then describes the evaluation and combination of the various systematic uncertainties that affect the hadronic energy estimate.

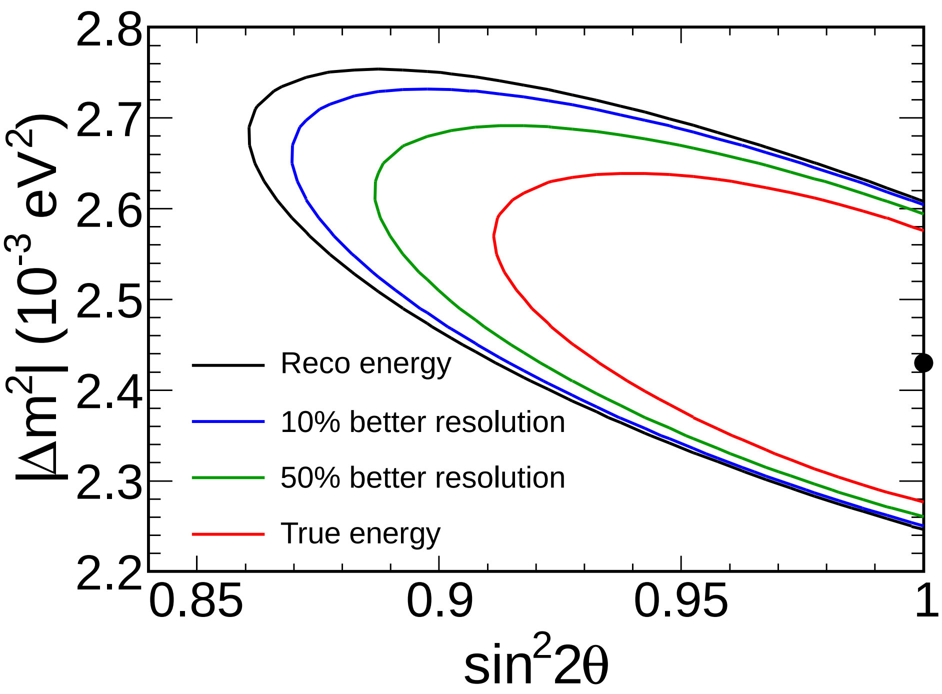

Figure 5.1: Comparison of the oscillation sensitivity obtained using the standard energy reconstruction (black), or the true neutrino energy (red). The blue and green contours show the sensitivity gains achieved by hypothetical 10% and 50% improvements in energy reconstruction, as described in the main text.

5.1 Motivation

Except where noted, all sensitivities in this chapter are based on 7×10^{20} POT of data[1] taken in the low energy, forward horn current, beam configuration. Only fiducial events are included, and resolution binning is not performed. The input oscillation parameters are $\Delta m^2_{\text{atm}} = 2.43 \times 10^{-3}\,\text{eV}^2$, $\sin^2 2\theta_{\text{atm}} = 1$, the best fit from the 2008 charged-current analysis [69]. Contours are drawn at the 90% confidence level.

Figure 5.1 illustrates the scale of the possible sensitivity gains. For the red contour, the usual reconstructed energy of each event has been replaced by the true neutrino energy from the Monte Carlo. Access to this “perfect” energy estimator

---

[1] The kNN energy estimator was developed in part before the final POT count for the analysis was determined. This exposure approximately reflects the quantity of data that was taken in the low energy beam configuration.

gives a very substantial improvement in oscillation sensitivity. This sensitivity is, of course, unobtainable in practice. Much of the necessary information is irretrievably lost, for example in the uncertain quantity of energy deposited in the detector steel. Figure 5.1 also includes two other scenarios, which show that even a more realistic improvement in energy resolution can have a significant impact on oscillation sensitivity. In each case, the reconstructed energy of simulated events has been adjusted so that the difference from the true neutrino energy is reduced[2]:

$$E_{\text{adjust}} = E_{\text{true}} + \frac{E_{\text{reco}} - E_{\text{true}}}{f} \ . \tag{5.1}$$

The "10%" and "50%" contours in Figure 5.1 are generated with $f = 1.1$ and $f = 1.5$ respectively.

The resolution functions (4.2), plotted in Figure 4.15, show that, at any given energy, the energy resolution of showers is much worse than that of tracks. Since, for an average event, the energy is split roughly equally between track and shower, the energy resolution is dominated by the shower contribution. Improvements in track energy resolution will therefore have very little effect on the oscillation sensitivity, whilst improvements in shower resolution will be reflected directly.

5.2 The kNN algorithm

The approach taken here to attempt to improve the energy resolution of MINOS showers is to adopt a multivariate technique, making use of a broader range of information about the shower, and the event as a whole. The suggestion to use a kNN was inspired by work on the charged-current PID, and the technique showed promise in initial investigations [138].

The algorithm begins by taking a training sample of Monte Carlo events. These events are distributed across the run periods in the correct proportions, so they

---

[2] Here, E stands for the full neutrino energy, not just the shower component.

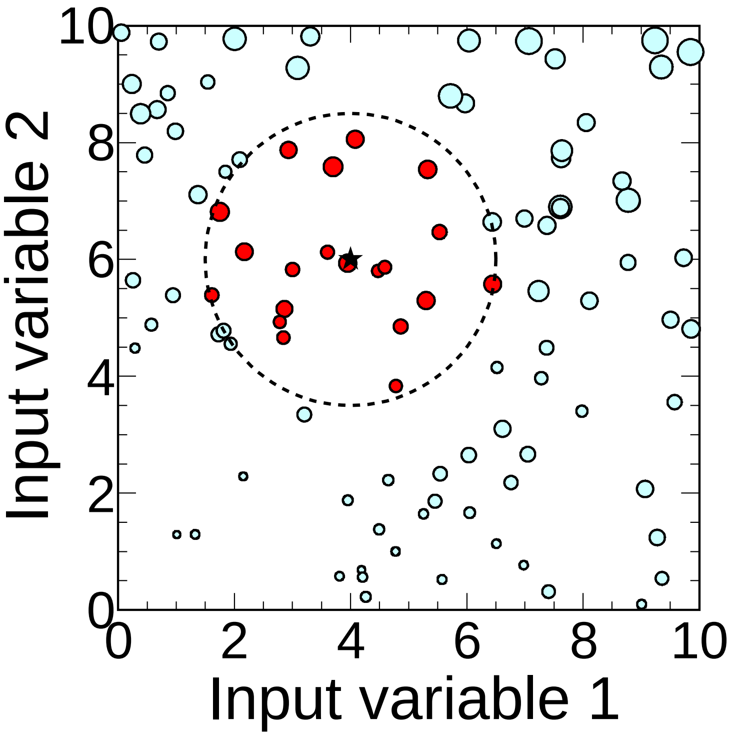

Figure 5.2: Cartoon of the use of a kNN for energy estimation. In this two-dimensional space, higher energy training events (represented by larger circles) correlate with larger values of the variables. The energy of the event marked by the star is estimated by examining its k (here 20) nearest neighbours, highlighted in red. The energy estimate is the mean of the true energies of these neighbouring events.

accurately reflect the properties of the data, and they are required to pass the charged-current selection. Separate samples are used for the near and far detectors, drawn from the respective Monte Carlo, to allow for any inter-detector differences.

For a D-dimensional kNN, D variables characterizing the properties of an event, and correlating with true shower energy, are stored for each training point, along with the true shower energy that was simulated by the Monte Carlo. In order to estimate the shower energy of a candidate event, the k closest training events, in the D-dimensional space formed by the input variables, are found. The output energy estimate is then the mean of the true shower energies of these training events. Figure 5.2 illustrates this procedure.

The distance measure, or metric, used to define which are the most nearby events is simply the D-dimensional Euclidean metric. The squared distance Δs^2 between a trial point $\vec{y}$ and a training point $\vec{x}$ is

$$\Delta s^2 = \sum_{i=1}^{D} \frac{(y_i - x_i)^2}{\sigma_i^2} \ . \tag{5.2}$$

Because the input variables are not guaranteed to be in compatible units, they must be normalized in the distance calculation. Otherwise, the variable with the largest numerical range would dominate over the others. The values of the input variables are thus compared in terms of their standard deviations (σ_i) in the training sample.

Other choices are possible for the construction of the metric, including a form varying throughout the training space. Similarly, the energy estimate could be constructed in some way other than a simple mean, perhaps weighting more distant training events less heavily. These, however, are the simplest choices, and have no adjustable parameters, yielding the smallest search space in which to find an optimal kNN. Note that in the limit of large training sample statistics, the precise definition of the metric becomes irrelevant as all chosen points will be "very nearby". The only choices remaining are in the set of variables to use as inputs, including their total number, and in the number of neighbours, k, to include in the energy calculation.

A naïve implementation of the kNN algorithm requires $\mathcal{O}(n)$ time to find the closest events from an n-element training set. Using a training set of over a million events to estimate the energies for a similarly sized set of test points would be very resource-intensive. Substantially faster neighbour searches are achieved by storing the training points in a kd-tree structure [139]. This is a multi-dimensional analogue of a balanced binary tree. Figure 5.3 illustrates the space partition construction. The structure consists of a set of nested axis-aligned boxes. Each box is associated with one training point and, except at the leaf level, two child boxes

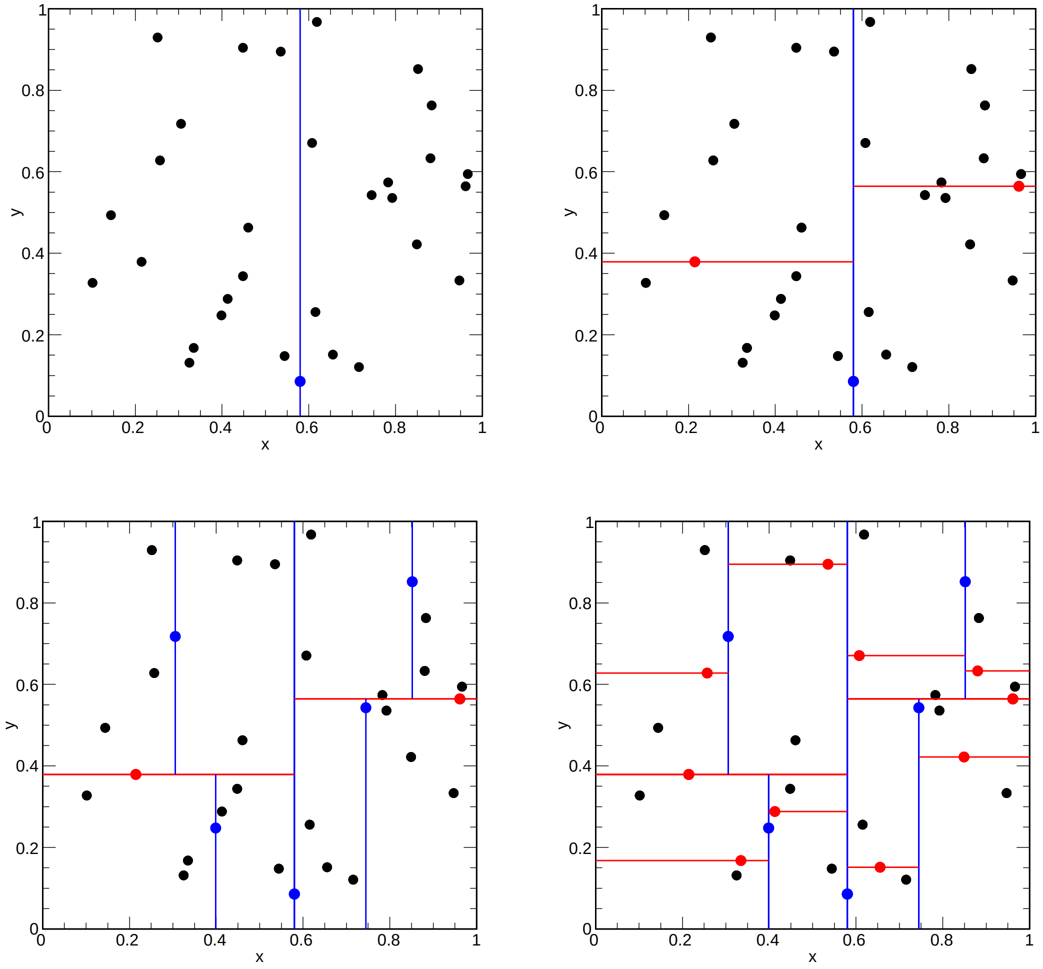

Figure 5.3: Partition of two-dimensional points by a kd-tree. The points are first divided along the horizontal axis (marked in blue) and then subdivided along the vertical axis (marked in red). This procedure repeats until no points share the same region. Each boundary passes through exactly one point, dividing the remaining points as equally as possible (*i.e.* the median for an odd number of points).

containing those points “left” and “right” of this point. The dimension along which the boxes are subdivided advances at each layer of the tree. The point at which each box is split is the median point along this dimension, dividing the remaining points equally between each child. The nearest neighbour search proceeds as a breadth-first traversal of the tree. There is no need for the search to descend into boxes whose entire volume is further from the test point than the worst of the k candidate neighbours so far identified. For randomly distributed data, the neighbour search takes $\mathcal{O}(\log n)$ time, a significant improvement in performance and scalability.

5.3 Selection of variables

With N candidate input variables, there are

$$^{N}\mathrm{C}_{D} = \frac{N!}{D!(N-D)!} \tag{5.3}$$

possible choices of D-dimensional kNN. Seven input variables were considered as potential components of the new estimator. With this number of candidates, a brute-force search over all combinations remains tractable (there are 127 possibilities to test[3]). The variables considered were:

1. The calorimetric energy in the primary shower.
2. The sum of the calorimetric energy in the primary and secondary showers.
3. The "deweighted" energy of the primary shower.
4. The sum of the deweighted energies of all showers within 1 m of the track vertex.
5. The number of hit strips in the primary shower.
6. The number of detector planes with activity from the primary shower.
7. The number of planes between the start of the shower and the plane with the maximum energy deposition.

The width of the shower, defined in terms of the variance of the strip U-V positions $\sigma_U^2 + \sigma_V^2$, was considered as a variable and found to have little power.

---

[3] Including or not including each variable, requiring at least one variable to be included: $2^7 - 1 = 127$.

For each of the variables which is defined as an energy, this is the fully calibrated energy. That is, the values are those produced at a late stage of reconstruction, after variations between channels, and over the life of the experiment, have been calibrated away.

"Deweighted" energy refers to an energy estimator [119] in which the energy deposited in each strip is raised to a power between zero and one before they are summed. The power used varies with the size of the shower, so that for large showers the estimate approaches the calorimetric energy, and for small showers the estimator tends towards counting strips. The estimator is described in detail in [119], but the large improvements reported there appear to have been negated by the two-photoelectron cut since introduced into the reconstruction.

For each combination of variables, the shower energies for the full far detector Monte Carlo set were calculated using the corresponding kNN. In the case of analyzing Monte Carlo using a training file derived from the same Monte Carlo sample, it is necessary to discard the very nearest neighbour found - this corresponds to the trial event itself[4]. From this data, it is possible to calculate an oscillation sensitivity assuming the conditions stated at the start of this chapter. This sensitivity then needs to be converted into a single figure of merit, for the purposes of comparing different kNN choices. The figure of merit chosen was the 1σ error achieved on Δm^2, assuming maximal mixing. This definition allows a shortcut to be taken. Instead of calculating the full $\Delta\chi^2$ curve, and locating $\Delta\chi^2 = 1$, it is sufficient to sample a single Δm^2 value and use the parabolic shape of the $\Delta\chi^2$ curve, and the fact that the minimum χ^2 is zero by construction, to derive the 1σ Δm^2 error. The actual figure used is the $\Delta\chi^2$ between one spectrum constructed at $\Delta m^2 = 2.43 \times 10^{-3}\,\mathrm{eV}^2$, $\sin^2 2\theta = 1$ and another at $\Delta m^2 = 2.70 \times 10^{-3}\,\mathrm{eV}^2$. This value is monotonically related to the error on Δm^2.

---

[4] In production, this selection is implemented as a requirement that the true energy of neighbours differs by more than 10 keV from the trial event.

The kNN with the best figure of merit was three-dimensional, consisting of:

- The calorimetric energy in the first two showers.
- The deweighted energy within 1 m of the track vertex[5].
- The number of planes in the primary shower.

Using a larger number of input variables might be expected to improve the energy reconstruction due to the availability of additional information. However, a higher-dimensional training space also leads to a sparser distribution of training points. A variable that provides only a small amount of extra information may only serve to "confuse" the algorithm, causing the resulting neighbour selection to be somewhat randomized.

In order to avoid possible problems in the kd-tree algorithm caused by the discrete nature of the number of planes variable, a deterministically seeded pseudo-random gaussian jitter of magnitude 10^{-5} is added to this variable in the training set.

In response to concerns that the second variable might pick up showers from unrelated events, a timing cut is applied in the near detector. Figure 5.4 shows the time difference between shower hits and the reconstructed track vertex. Showers associated with tracks leave a peak at small time difference, independent activity gives a flat background. On the basis of this plot the timing cut is set at $-25\,\mathrm{ns} < \Delta t < +75\,\mathrm{ns}$, a window 1% of the length of the beam spill.

The same set of three variables is used for energy estimation in both the near and far detectors, although based on training samples drawn from the respective Monte Carlos.

Having selected the input variables to the kNN, the final remaining step is to optimize the number of neighbours used. Increasing the number of neighbours

---

[5] Attempting to replace this variable with calorimetric energy near the track vertex yielded a decrease in sensitivity.

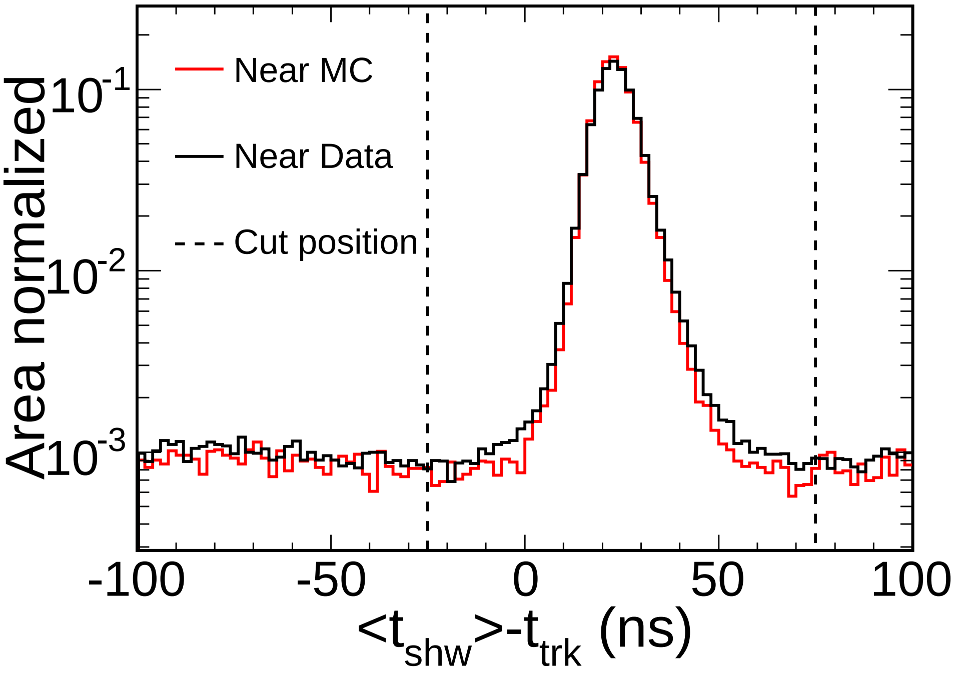

Figure 5.4: Distribution of $\langle t_{\mathrm{shw}}\rangle - t_{\mathrm{trk}}$ in data (black) and Monte Carlo (red). This distribution is formed from all pairs of reconstructed tracks and shower in a single spill. $\langle t_{\mathrm{shw}}\rangle$ is the median time of all hits in the shower, t_{trk} is the time assigned to the track vertex by the track fitter. The peak is due to tracks and showers from the same neutrino interaction, whilst the flat background is due to coincidences between independent events. The dashed lines show the timing cut applied to ensure that only showers associated with the same event as the track are included.

examined provides more information to the energy estimation and decreases the effect of statistical fluctuations. However, with increasing number of neighbours the algorithm begins to sample parts of phase space which bear less similarity to the input event under consideration. The optimum number of neighbours to use is a trade-off between these two considerations, and a function of the size and properties of the training sample in use. Figure 5.5 shows the figure of merit, as defined above, for the kNN with the three selected variables, obtained as a function of the number of neighbours. The expected behaviour is observed, with a rapid rise, followed by a gradual decline. On the basis of this plot, the number of neighbours to use was chosen to be 400 in each detector. Of course, the optimum

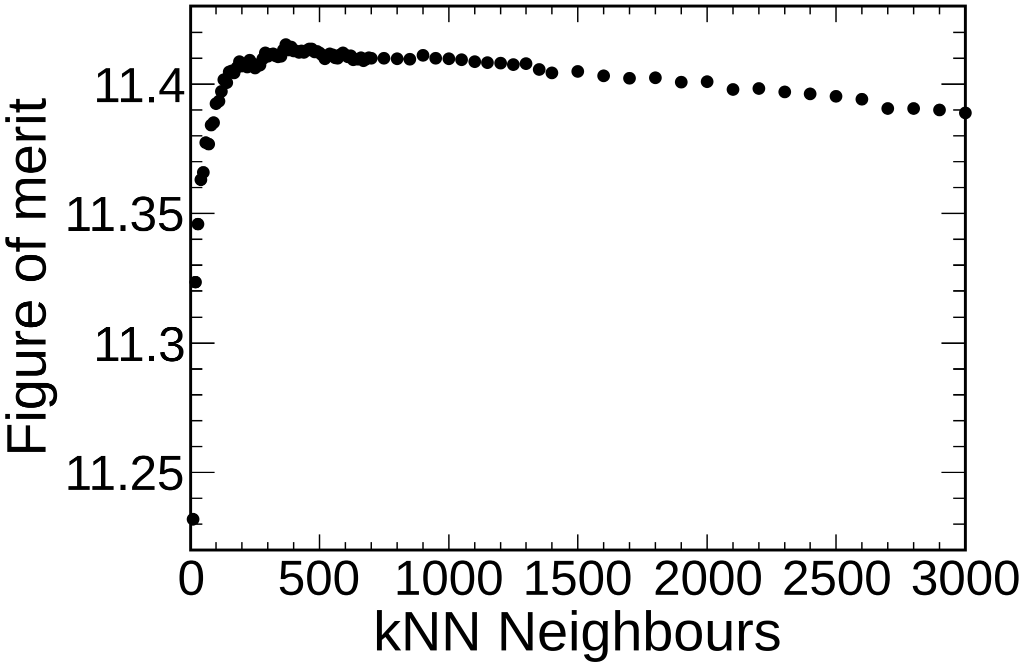

Figure 5.5: Figure of merit for the shower energy kNN as a function of the number of neighbours used (k). The figure of merit is defined in the main text. Performance of the kNN improves rapidly, before levelling off around 300-500 neighbours, and then gradually declining. The finalized kNN uses 400 neighbours.

number of neighbours is also a function of the size of the training sample. The training sample in use, consisting of the full charged-current selected Monte Carlo sample available, contains around $1,200,000$ events.

The kNN energy estimation algorithm also allows an estimate of the shower energy resolution of individual events to be calculated, by taking the standard deviation of the true energies of the neighbours used in the mean. This estimate could be used for the resolution binning described in Section 4.5. However, investigation of this variable showed that it performed no better than the parameterization of Equation (4.2), and so the simpler approach was retained.

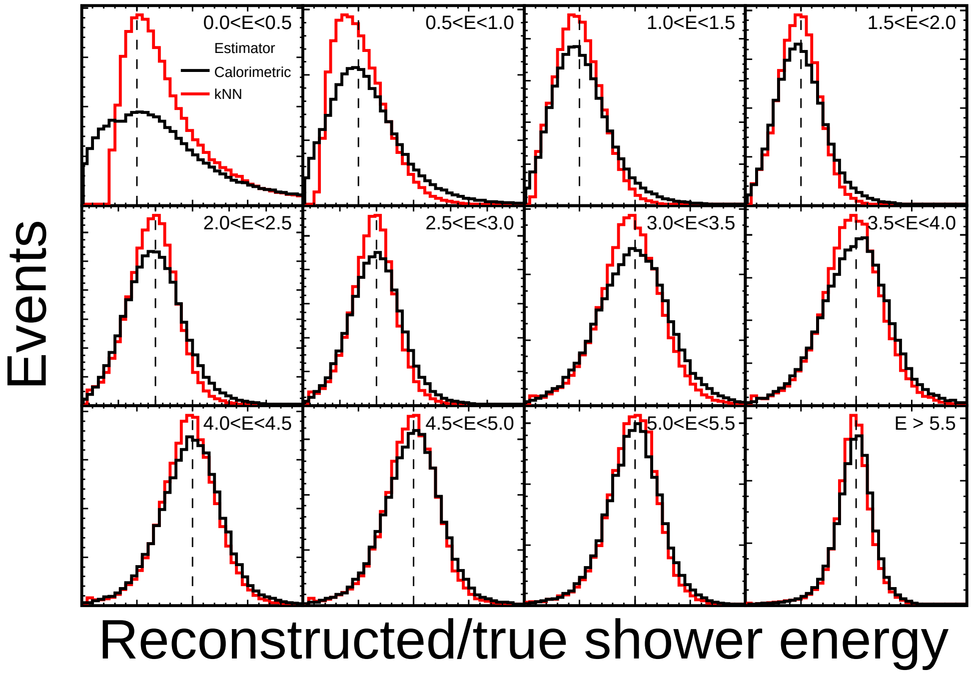

Figure 5.6: Distribution of reconstructed over true shower energy for different ranges of true shower energy (indicated in the corner of each panel). The calorimetric estimator is shown in black, and the kNN estimator in red. In each panel, the dashed line indicates the position of $E_{\text{reco}}/E_{\text{true}} = 1$.

5.4 Sensitivity improvement

Figure 5.6 shows the distribution of reconstructed over true shower energy for the standard calorimetric estimator, and for the kNN estimator, constructed as described above. The behaviour of the peak in each case is displayed as a function of the true shower energy. At high shower energies, the kNN and calorimetric energies behave similarly. However, at lower shower energies, as the resolution of both estimators falls, the kNN energy becomes noticeably superior, displaying a much sharper peak, indicating better energy resolution for individual events.

One way of quantifying the improvement in each event is via the adjustment factor it corresponds to in the artificially improved Monte Carlo of Section 5.1 (Figure 5.1 and Equation (5.1)). That is, the amount closer the kNN energy estimate is to the truth, as a fraction of the true energy. This can be expressed in

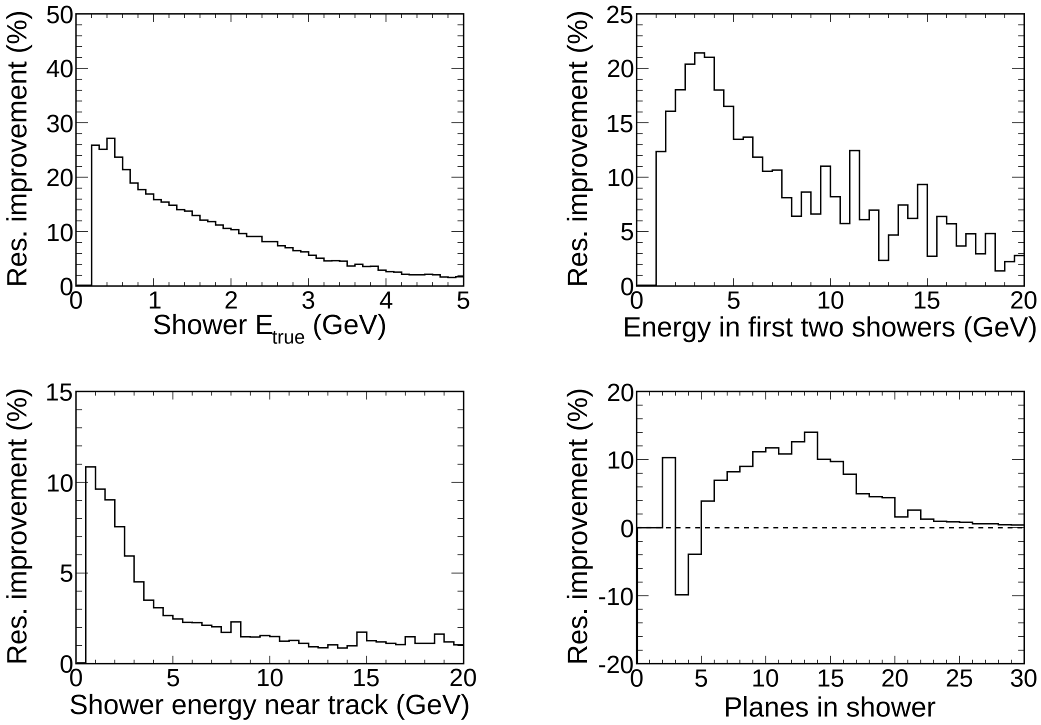

Figure 5.7: Improvement of kNN shower energy over calorimetric energy as defined in Equation (5.4). In the top-left is the improvement as a function of the true shower energy. In the other panels, this value is displayed as a function of the input variables.

two equivalent forms:

$$\begin{aligned} R &= \frac{|E_{\text{calo}} - E_{\text{true}}| - |E_{\text{kNN}} - E_{\text{true}}|}{E_{\text{true}}} \\ &\equiv \left|\frac{E_{\text{calo}}}{E_{\text{true}}} - 1\right| - \left|\frac{E_{\text{kNN}}}{E_{\text{true}}} - 1\right| \ . \end{aligned} \tag{5.4}$$

The second form makes clear the relationship of this measure to the widths of the distributions in Figure 5.6. In this equation, and onwards, the variable E represents the hadronic, or shower, component of the full neutrino energy, unless otherwise noted. The average value of R gives the overall improvement in a collection of events of interest. The top-left panel of Figure 5.7 shows this average improvement as a function of the true shower energy. The improvement is small at high energies, but rises as high as 25% at the lowest energies. The other panels

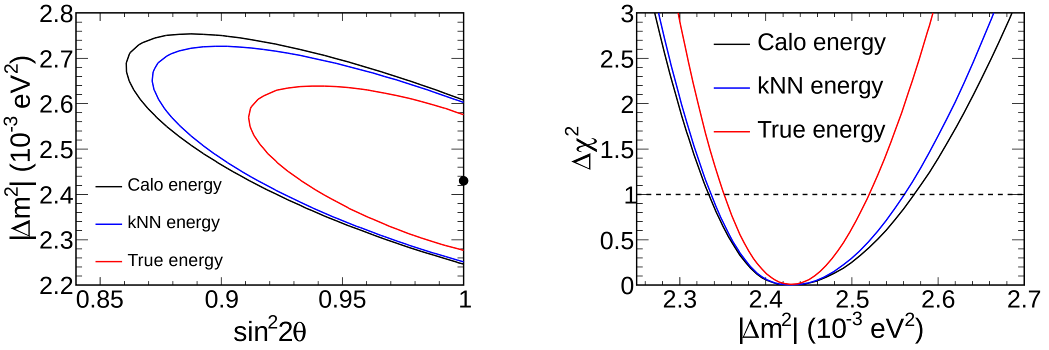

Figure 5.8: Improvement in oscillation sensitivity achieved by the kNN energy estimator (blue), compared to the calorimetric estimator (black), plus the sensitivity from perfect reconstruction (red) for scale. On the left, 90% confidence contours in $(\sin^2 2\theta, \Delta m^2)$. On the right, the projection onto the Δm^2 axis. The dashed line at $\Delta\chi^2 = 1$ indicates the 68% confidence level.

show how the improvement varies as a function of the input variables. That is, what type of event gain the greatest improvement. For all types of event (except 3- and 4-plane events) R is positive, so on average the kNN estimator makes a better estimate than the calorimetric estimator. In general, low energy events benefit the most.

Of course, the ultimate test of the performance of the estimator is in the improvement in oscillation sensitivity that it achieves. Figure 5.8 shows sensitivity contours equivalent to those in Figure 5.1, with the addition of the contour obtained when using the kNN energy estimator. The projection onto the Δm^2 axis is also shown. One way to quantify this improvement is in terms of the amount of data taking that would be required to achieve it from increased statistics alone. Using the 1σ errors derived from the Δm^2 projection, and taking sensitivity to improve proportionally to $\sqrt{\text{POT}}$, the improvement expected from calorimetric energy to kNN energy is found to be equivalent to 10^{20} additional protons-on-target, *i.e.* around five months of data taking.

5.5 Energy corrections

Ideally, an energy estimator is unbiased. That is, the mean value of $E_{\text{reco}}/E_{\text{true}}$ is unity. For the neutrino energy estimate as a whole, a biased estimate is not actually harmful. Information about the nature of the bias is reflected in the Monte Carlo simulation. In the analysis, a biased estimator is equivalent simply to a remapping of the energy axis, and has no systematic effect on the quality of the result[6]. However, when estimates are summed, as is the case with track and shower energy in MINOS, a bias in one estimator can lead to an overall sensitivity degradation. The track energy estimator has been tuned to be essentially unbiased. Summing this with a biased shower estimate, in a proportion depending on the inelasticity of the event, leads to a smearing of the overall estimate, a decrease in energy resolution, and a consequent decrease in overall sensitivity.

Figure 5.9 shows the distribution of $E_{\text{kNN}}/E_{\text{true}}$ for showers using the kNN estimator, as a function of true shower energy. At high energies the distribution is narrow, and the mean is one. Below about 5 GeV, however, there are deviations in the mean, up to 15%. The existence and qualitative form of this feature can easily be explained in terms of the properties of the training sample. The training events are drawn directly from the standard Monte Carlo, including the features of the simulated neutrino energy spectrum. There are thus more training events in the beam peak, around 3 GeV, than elsewhere. This provides a prior expectation to the kNN, biasing nearby energy estimates, such that events below the beam peak tend to be reconstructed with higher energies, and vice versa.

Bias in the standard calorimetric estimator is dealt with by the application of an energy correction function to map the biased raw output to an unbiased final estimate [140, 141]. A similar approach is followed here. The top plot in Figure 5.10 shows the means from Figure 5.9, on a logarithmic scale to highlight

---

[6] So long as the mapping is monotonic.

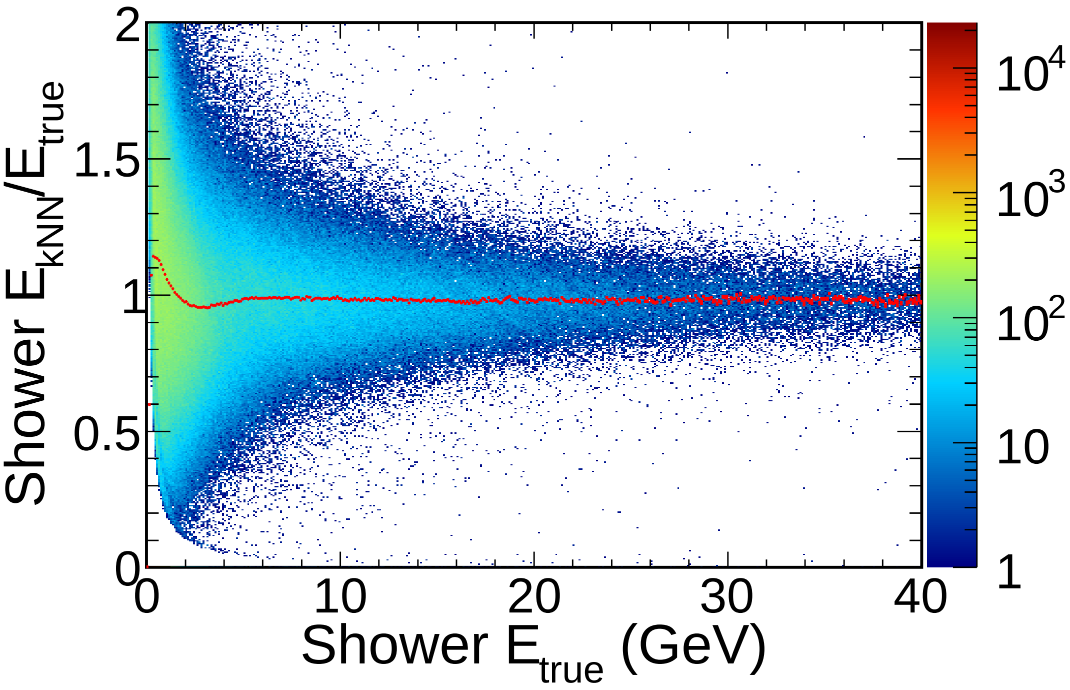

Figure 5.9: Distribution of $E_{\mathrm{kNN}}/E_{\mathrm{true}}$ in far detector Monte Carlo, as a function of E_{true}. The red points are placed at the mean of the distribution for each value of E_{true}. Below around 5 GeV the estimator displays a bias of up to 15%.

the structure at low energy. The aim of the energy corrections is to flatten this distribution so that $\left\langle \frac{E_{\mathrm{kNN}}}{E_{\mathrm{true}}} \right\rangle = 1$ for all values of E_{true}[7]. The plot also shows a 5th order polynomial in $\log(E_{\mathrm{true}})$ fit to the points. This serves as a first approximation to the desired correction function. The function does not successfully capture the full detail of the variation. More seriously, the correction has been determined as a function of E_{true}, whereas what is required is a function in E_{kNN}. Since E_{kNN} is already a reasonable estimate of E_{true}, the polynomial correction is applied as a function of $\log(E_{\mathrm{kNN}})$ as an approximation. After this first approximate correction, the new means are found, and the procedure iterated. In all, three correction functions are applied sequentially. The result is shown in the bottom plot of Figure 5.10. In the crucial region $500\,\mathrm{MeV} < E < 4\,\mathrm{GeV}$ the deviations

---

[7] An alternative aim would be to flatten the distribution of $\left\langle \frac{E_{\mathrm{kNN}}}{E_{\mathrm{true}}} \right\rangle$ as a function of E_{kNN}. Doing this is found not to lead to the same sensitivity improvement.

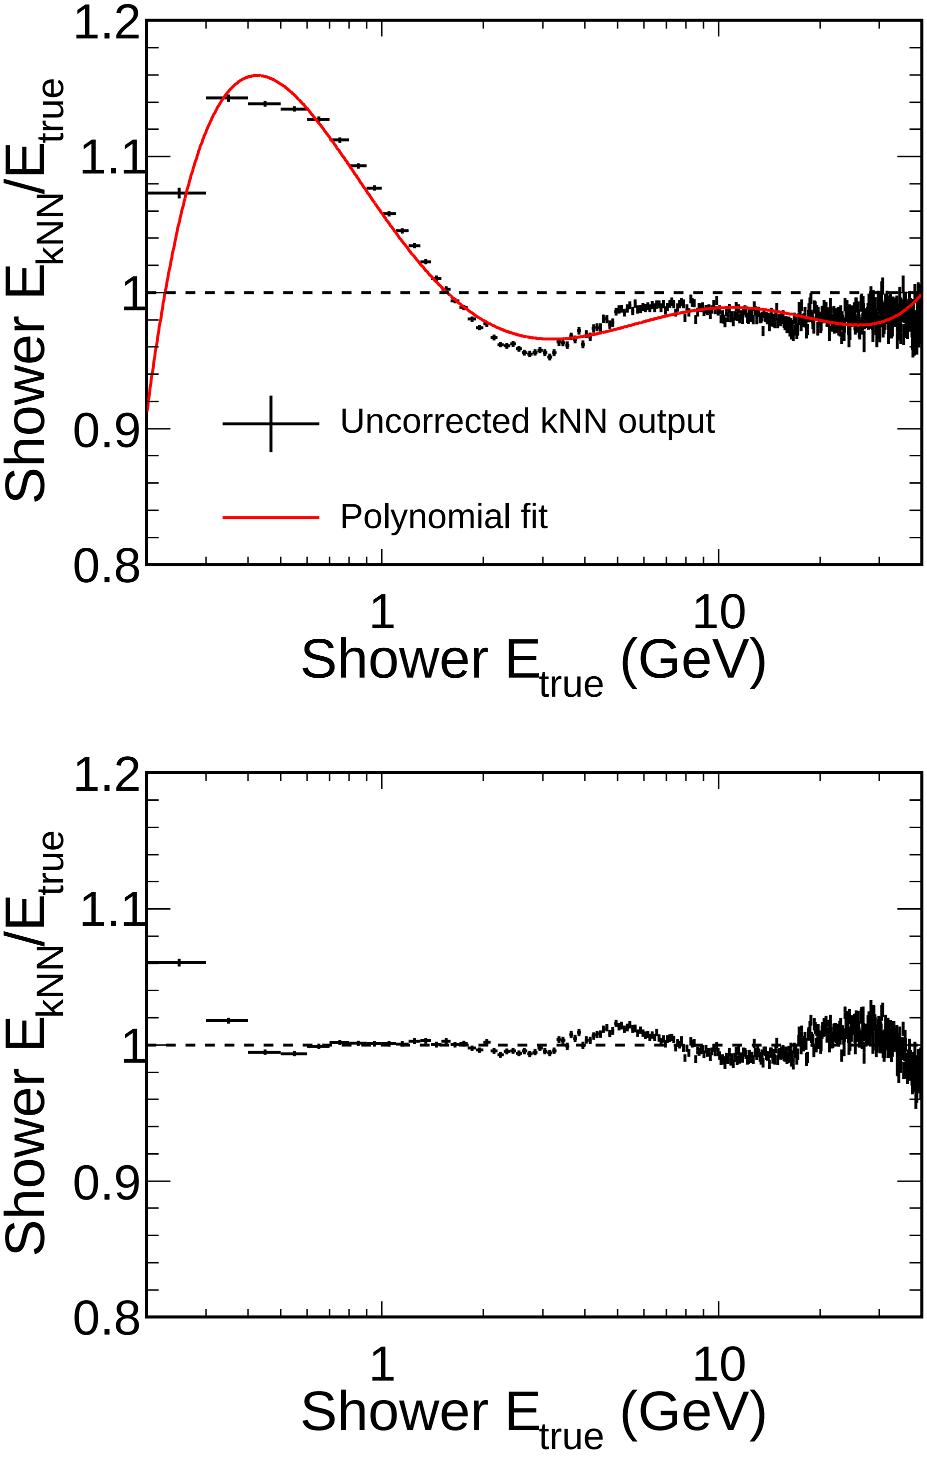

Figure 5.10: Mean value of $E_{\mathrm{kNN}}/E_{\mathrm{true}}$ as a function of E_{true}. In the top plot, the black points are the means from Figure 5.9. Note the logarithmic x-scale. The red curve is a polynomial fit to the points, the first step of the energy correction procedure described in the text. The bottom plot shows the means, calculated by the same procedure, after the full energy correction has been applied.

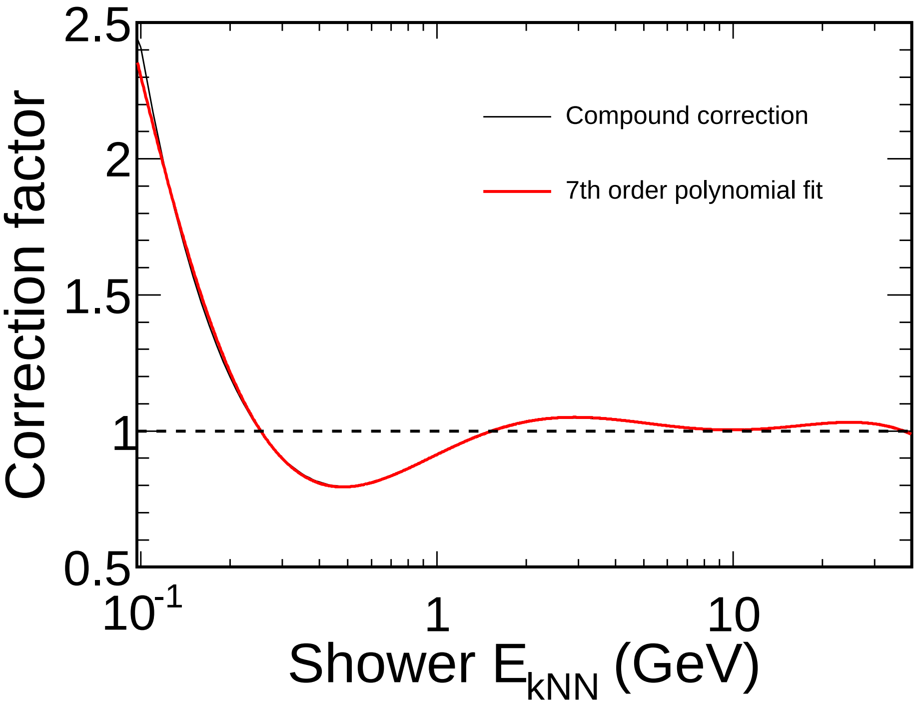

Figure 5.11: The correction factor applied to the kNN energy. In black, the correction calculated as a compounding of three sequential corrections. In red, the 7th order polynomial which is actually used. The black line lies underneath the red line for almost all values of E_{kNN}.

have been reduced to the sub-percent level.

This iterative correction procedure seems somewhat inelegant, corresponding in principle to a 125th order polynomial correction. The correction applied in practice is in fact a 7th order polynomial, which successfully reproduces all the features of the full correction. Figure 5.11 shows the final correction function $f(E_{\mathrm{kNN}})$. To guard against divergence of the function at high energies, the value of the correction factor is held constant at all energies above 40 GeV. A similar mechanism at low energies proved unnecessary. The correction procedure is thus:

$$E_{\mathrm{kNN}} \mapsto f(\max[E_{\mathrm{kNN}}, 40\,\mathrm{GeV}])E_{\mathrm{kNN}}\ . \tag{5.5}$$

The same procedure of fitting the $\left\langle \frac{E_{\mathrm{kNN}}}{E_{\mathrm{true}}} \right\rangle$ distribution is followed to obtain a similar correction function for use in the near detector.

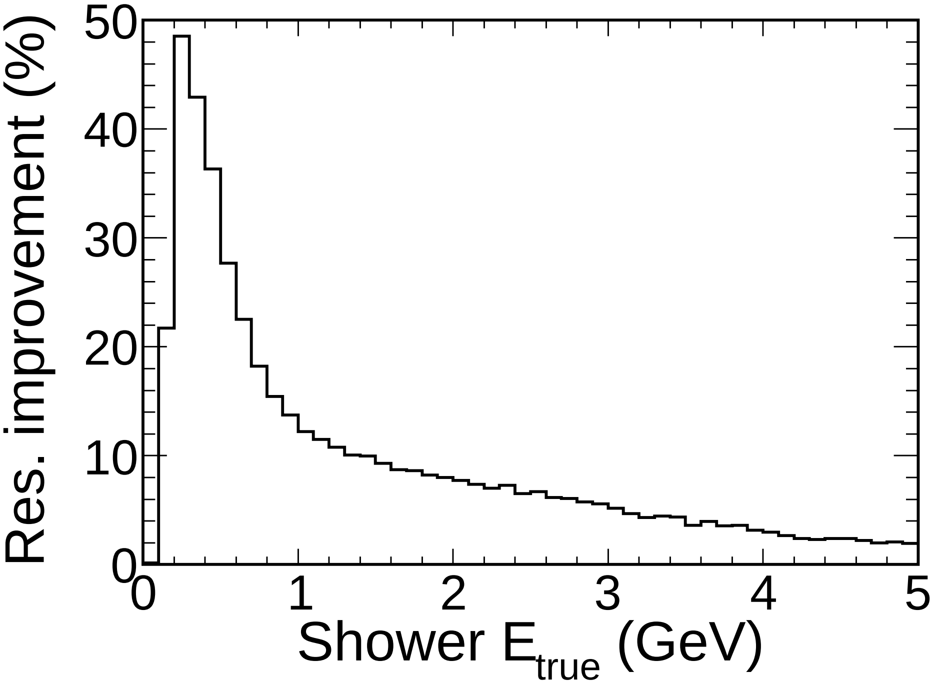

Figure 5.12: Improvement of kNN shower energy over calorimetric energy as defined in Equation (5.4), after energy corrections have been applied.

With the energy correction function applied, the improvement over calorimetric energy, as measured by the R variable (Equation (5.4)), is significantly increased at the lowest energies, reaching nearly 50%, as shown in Figure 5.12. As expected, this leads to a further small improvement in oscillation sensitivity, as shown in Figure 5.13.

5.6 Evaluation of the systematic error

For this new shower energy estimator to be safely used in an oscillation analysis, it is necessary to evaluate the degree of systematic error introduced into the hadronic energy measurement, and the impact on the overall oscillation fit. There is some concern that the kNN estimator may be more sensitive to mismodelling, as it uses more detailed information about the pattern of hit strips in an event.

The evaluation of the systematic error on the kNN shower energy follows a

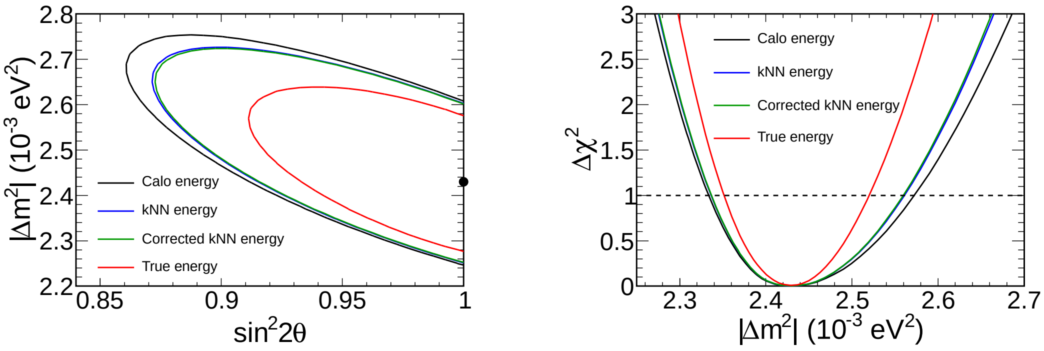

Figure 5.13: Improvement in oscillation sensitivity achieved by the corrected kNN estimator. Black, blue and red lines are as in Figure 5.8, the green lines show the small additional improvement from applying energy corrections.

similar approach to that used for the evaluation of the systematic error on the calorimetric energy for previous charged-current [142, 143] and ν_e [144] analyses.

The uncertainty in the calibration of the response of the MINOS detectors to hadronic showers originating from single particles has been well determined by test beam measurements at CalDet and by studies of cosmic ray events in the two detectors. A significant source of error also arises from the unknown hadronic particle content from the initial neutrino interaction. In the MINOS Monte Carlo simulation, there is uncertainty both in the modelling of the hadronization process in the initial interaction (described in Section 5.6.1) and in the subsequent rescattering of these particles inside the target nucleus (described in Section 5.6.2). The approach taken to evaluate this systematic is to consider the effect of adjusting each uncertain aspect of the simulation. These areas of uncertainty may be simple tunable parameters, or assumptions built into the models themselves.

Sections 5.6.3 and 5.6.4 describe the two different approaches that are employed to evaluate the various effects arising from modelling of the initial shower content. Section 5.6.5 presents the full list of the effects considered. The following sections describe additional systematic effects, notably the uncertainty in single particle response in Section 5.6.8. The overall systematic uncertainty obtained by combining all effects is presented in Section 5.6.11.

5.6.1 Neutrino interactions and the AGKY model

Neutrinos interact with the MINOS detectors in one of three ways. Quasielastic interactions do not give rise to a hadronic shower (the outgoing proton is essentially undetectable) and so create no uncertainty in the energy scale of the hadronic reconstruction. Resonance interactions make a substantial contribution to the interaction cross-section in the $\mathcal{O}(1\,\text{GeV})$ energy regime relevant to MINOS oscillation analyses. They are simulated according to the Rein-Sehgal model [145]. Although the rates of these interactions suffer from some uncertainty (which is included as a component the cross-section systematic in analyses), the final state particles from the initial interaction are well known. For example, in the case of the dominant $\Delta(1232)$ resonance, the branching ratio for the decay $\Delta \rightarrow N\pi$ is greater than 99% [46].

Deep inelastic scattering also makes a major contribution to the cross-section. In this case, the particle content of the final state is not so easy to determine. The initial hadronization is performed by the "AGKY model", transitioning to PYTHIA/JETSET [146] at higher energies. This model is described in [147] and [148]. The algorithm is, briefly:

- Compute the average charged hadron and total hadron multiplicities.
- Generate the actual hadron multiplicity.
- Generate a single final state baryon on the basis of quark model arguments.
- Generate mesons to make up the correct multiplicity, taking into account charge conservation and the available invariant mass.
- Generate the baryon four-momentum.
- Generate a phase-space decay of the remaining hadronic system.

The various probabilities and distributions required: for the multiplicities, meson production probabilities, and distribution of the baryon momentum, are formed using data from bubble chamber experiments.

5.6.2 Intranuclear rescattering with INTRANUKE

The majority of the detector mass in MINOS is in the form of steel. This means that in most neutrino interactions, the hadrons produced have to escape from an iron nucleus, with a significant chance of re-interaction, before they can be seen in the detector. This re-scattering process is simulated by a semi-classical intranuclear cascade model, INTRANUKE [149].

There are four processes by which a hadronic particle may re-interact:

- Elastic scattering – The particle is deflected by in an interaction, but retains its identity.
- Inelastic scattering – The particle is scattered, losing energy but retaining its identity.
- Charge exchange – The hadron is transformed from charged to neutral or vice versa during an inelastic scattering process.
- Absorption – The hadron does not escape the nucleus, transferring its energy to the constituent nucleons.

Initially the particles from the hadronization process are in the form of quarks, not yet bound into hadrons, which travel relatively freely through the nucleus. This phenomenon is implemented in INTRANUKE by allowing particles to travel freely for a “formation time” $\tau = 0.342\,\mathrm{fm}/c$ before becoming candidates for interaction.

The probability for a particle, travelling along a trajectory s, through nuclear matter of density $\rho(r)$, with total interaction cross-section $\sigma(E)$ to interact is given by:

$$P = 1 - \exp\left(- \int_s \rho(r)\sigma(E) \left| \frac{\mathrm{d}\vec{r}}{\mathrm{d}s} \right| \mathrm{d}s \right) \ . \tag{5.6}$$

This is approximated as a sequence of $\Delta s = 0.1\,\mathrm{fm}$ steps:

$$P \approx 1 - \exp\left(-\sigma \sum_i \rho(r_i) \Delta s \right) \ . \tag{5.7}$$

Once an interaction is deemed to have occurred, the exact nature is determined from energy-dependent branching ratios to the various "fates". Additional particles are created as required by the form of the interaction. These secondary particles are not themselves subject to reinteraction in the nucleus, since the cross-sections used were determined from data on $\pi + {}^{56}\mathrm{Fe}$ and $p + {}^{56}\mathrm{Fe}$ collisions, so these additional interactions have already been taken into account.

5.6.3 INTRANUKE reweighting

Sufficient information is contained in the Monte Carlo output events to enable the effect of varying INTRANUKE parameters to be studied using a reweighting scheme [150]. The identity and reinteraction type, if any, of all the initial hadronization products is stored. The probability for an event to have been simulated as it was is the product of the probabilities of all the initial hadrons to have reinteracted as they did, a function of the cross-sections and branching ratios. This probability can be calculated from the Monte Carlo record, as can the probability with differing constants. To imitate the effect that a change in INTRANUKE parameters would have, each simulated event is reweighted by the ratio of these two probabilities. When adjusting the branching ratio to one fate, the branching ratios to the other fates are automatically adjusted to maintain the unitarity condition that the sum of the probabilities of all fates is unity.

The formation time variable causes a unique problem for the reweighting scheme. Above a certain energy, 2.8 GeV for a pion with the default parameters, the free step is sufficiently long that it will always take the particle clear of the nucleus, leaving no chance for reinteraction. Attempting to simulate a decreased formation time via reweighting, some of these particles should now have a chance to reinteract. However, since no such reinteracting events have been simulated, there are no suitable events available to reweight. When this occurs, the weight is left set to 1.

5.6.4 Special Monte Carlo samples

Other possible systematic effects, either in the initial hadronization, or model uncertainties in rescattering which are not expressible as simple changes of parameters, cannot be studied via a reweighting scheme. To evaluate these effects, as well as two of the largest INTRANUKE effects, special Monte Carlo samples were generated and reconstructed. In all, thirteen such samples were generated, some representing a discrete model change, others, in pairs, reflecting $\pm 1\sigma$ shifts in model input parameters. The samples were generated simulating the near detector, with the Run III beam flux (this being the configuration that the largest quantity of data was taken in). To save generation and reconstruction time, only events with their vertex in the fiducial volume were generated, and overlaying was not performed, so each "spill" consists of a single event. An additional, unshifted, nominal sample was produced under the same conditions for comparison. After application of the charged-current selection, the nominal sample consisted of about 217,000 events.

5.6.5 Nuclear modelling systematics considered

The full set of hadronization, and intranuclear rescattering systematics considered are listed below, grouped according to the technique by which their contribution was evaluated.

Evaluated from special Monte Carlo samples:

- Pion/nucleon absorption – The probability for absorption of pions and nucleons is adjusted within the bounds allowed by external data (±30% and ±20% respectively) [151]. A correlated shift in these values is implemented as $\pm 1\sigma$ special MC samples, and independent shifts are available as INTRANUKE weights.

- Baryon x_F selection – Based on the scenario of a neutrino striking a constituent quark of a nucleon, the final state baryon is most likely to form from the remaining two quarks, *i.e.* in the opposite direction to the momentum transfer, the backwards hemisphere. This leaves the pions to be formed in the forward hemisphere and thus to have higher energies when boosting back to the lab frame. The baryon x_F distribution has been measured in bubble chamber data, confirming this effect, and the AGKY model draws from this distribution [148]. The model change in this sample is, conservatively, to undo this behaviour and simply generate the hadron momenta as a phase space decay in the centre of mass [144].

- Formation zone – The $0.342\,\mathrm{fm}/c$ formation time described above is based on measurements in the SCAT bubble chamber [152]. For these samples the formation time was adjusted by ±50% ($\tau = 0.171, 0.513\,\mathrm{fm}/c$), reflecting the uncertainty on this measurement. The effects of formation time uncertainty can also be studied via the INTRANUKE reweighting procedure, but as described above, this is expected to suffer from some problems.

  An improved formation zone model, in which interactions of newly created particles turn on gradually, instead of with a step function as at present, is currently under development [153]. When available, simulations using this model should, at the least, be incorporated as an additional systematic sample. Previous work [143] has shown the effect of this modification to be similar in magnitude to the adjustment of τ described above.

- INTRANUKE assumptions: absorption – INTRANUKE simulates the pion absorption process according to the Ransome model. The energy of the absorbed pion is most often distributed between 3 or 4 nucleons [154]. As a conservative estimate, this sample doubles this number of nucleons.

- INTRANUKE assumptions: de Broglie ring – At low pion energies, the quantum mechanical nature of the interactions becomes more significant as the de Broglie wavelength $\lambda = h/p$ increases. This is taken into account in an empirical manner, increasing the reinteraction probability by artificially increasing the size of the nucleus. This adjustment was tuned by matching the output to data on the total cross-section of pion-iron scattering, and cross-checked for consistency with neutrino data [155]. By default the best fit of a 0.5λ increase is used [156]. For these samples the constant of proportionality is varied by the larger ± 0.6 uncertainty obtained from the neutrino data. To get good agreement with this data, it was also necessary to modify the pion absorption cross-section by $+0.5\sigma$. This change is also included in both samples.

- Charged/neutral particle correlations – The AGKY model first selects the total multiplicity and then divides this between charged and neutral particles. Measurements show that the charged and neutral multiplicities are in fact independent. This sample uses an updated model where the multiplicities are chosen independently [144].

- π^0 probability – The relative probability for AGKY to produce a π^0 pair instead of $\pi^+\pi^-$ is 0.30. For these samples the probability is varied by

±30% to 0.21 or 0.39 [144].

- Two-body decays – All decays to two-body hadronic final states, both in the AGKY model and in the simulation of resonant interactions, are made isotropically in the centre of mass. These samples consist of two extreme (unphysical) modifications. In one sample, the particles are produced perpendicular to the direction of momentum transfer, in the other they are produced parallel [144].

Evaluated by INTRANUKE reweighting:

The 1σ errors quoted on the INTRANUKE parameters are tabulated in [143]. A detailed description of the construction and evolution of the model is given in [149].

- Pion charge exchange – The branching ratio for charge exchange is varied by ±50%.
- Pion elastic scattering – The branching ratio for elastic scattering is varied by ±10%.
- Pion inelastic scattering – Similarly, the inelastic scattering branching ratio is varied by ±40%.
- Pion secondary pion production – The fraction of pion interactions creating a second pion is scaled by ±20%.
- Nucleon secondary pion production – Likewise, the fraction of nucleon interactions that create a pion is scaled by ±20%.
- Pion cross-section – The total pion interaction cross-section is adjusted by ±10%.

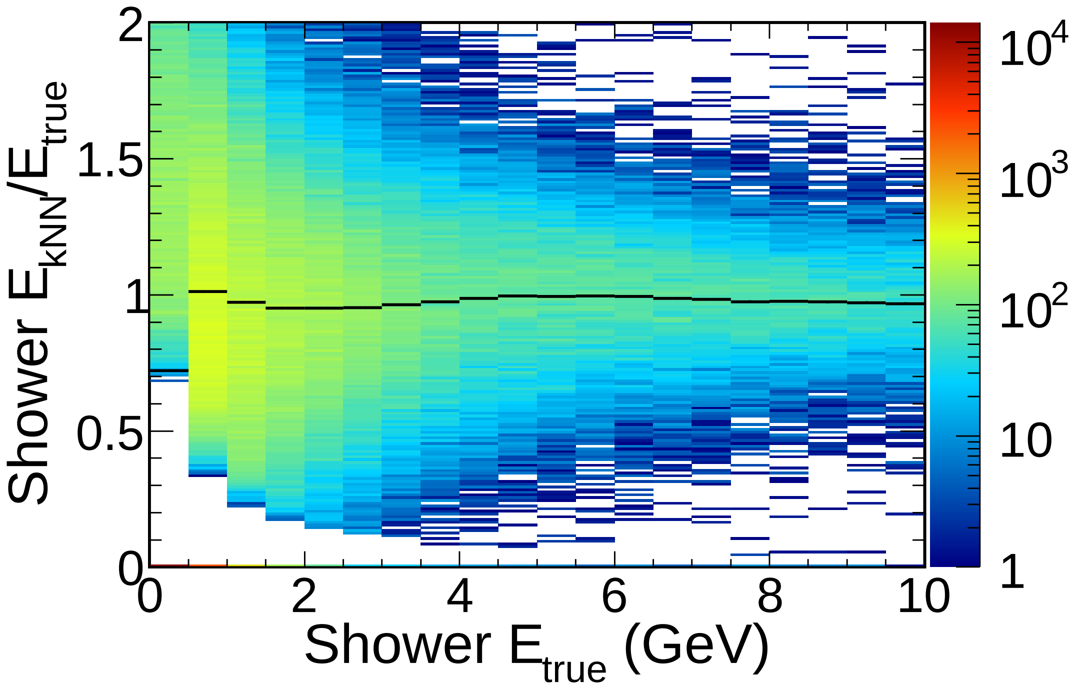

Figure 5.14: Example distribution of $E_{\rm kNN}/E_{\rm true}$ for a systematically shifted Monte Carlo sample. The black marker in each bin indicates the mean.

- Nucleon cross-section – Similarly, the total nucleon interaction cross-section is adjusted by ±15%.

5.6.6 Evaluation of individual error contributions

The evaluation of the systematic error introduced by each model uncertainty begins by selecting events from the shifted Monte Carlo which pass the charged-current selection. For these events, an $E_{\rm kNN}/E_{\rm true}$ vs. $E_{\rm true}$ histogram is constructed, as for the construction of the energy corrections (Section 5.5 and Figure 5.9). The statistics in these samples are significantly lower than the full Monte Carlo, so a coarser 500 MeV binning is used. An example distribution is shown in Figure 5.14. The mean is evaluated in each true energy bin, including the uncertainty $\sigma/\sqrt{N}$. The same procedure is applied to the nominal sample, and the ratio of the means is taken, to leave the bias introduced by the systematic shift, with any component due to the special generation circumstances of the Monte Carlo removed.

For systematic components evaluated by INTRANUKE reweighting, the equivalent procedure is followed for the full Monte Carlo sample, with one weight factor applied to each event, divided by the case with no weights applied. The uncertainty on the errors evaluated for these systematics is much lower, due to the higher statistics of the full Monte Carlo sample. Additionally, the two distributions being divided are correlated, so the statistical uncertainties are overestimated.

Figure 5.15 shows the shifts resulting from all of the uncertainties described in Section 5.6.5. In some cases, the systematic consists of a parameter adjusted by $\pm 1\sigma$ in which case an approximately symmetric error band is expected. In other cases, a discrete model change results in a one-sided error.

An alternative, arguably more correct, approach to evaluating the effect of the INTRANUKE weights is to pick all the parameters simultaneously from Gaussian distributions and use the average spread from a large sample of such variations. This approach was found to give very similar results for the overall INTRANUKE-induced systematic. The treatment with discrete $\pm 1\sigma$ components was adopted, since it separates out the contribution from each component, and for consistency with the treatment of the special Monte Carlo samples.

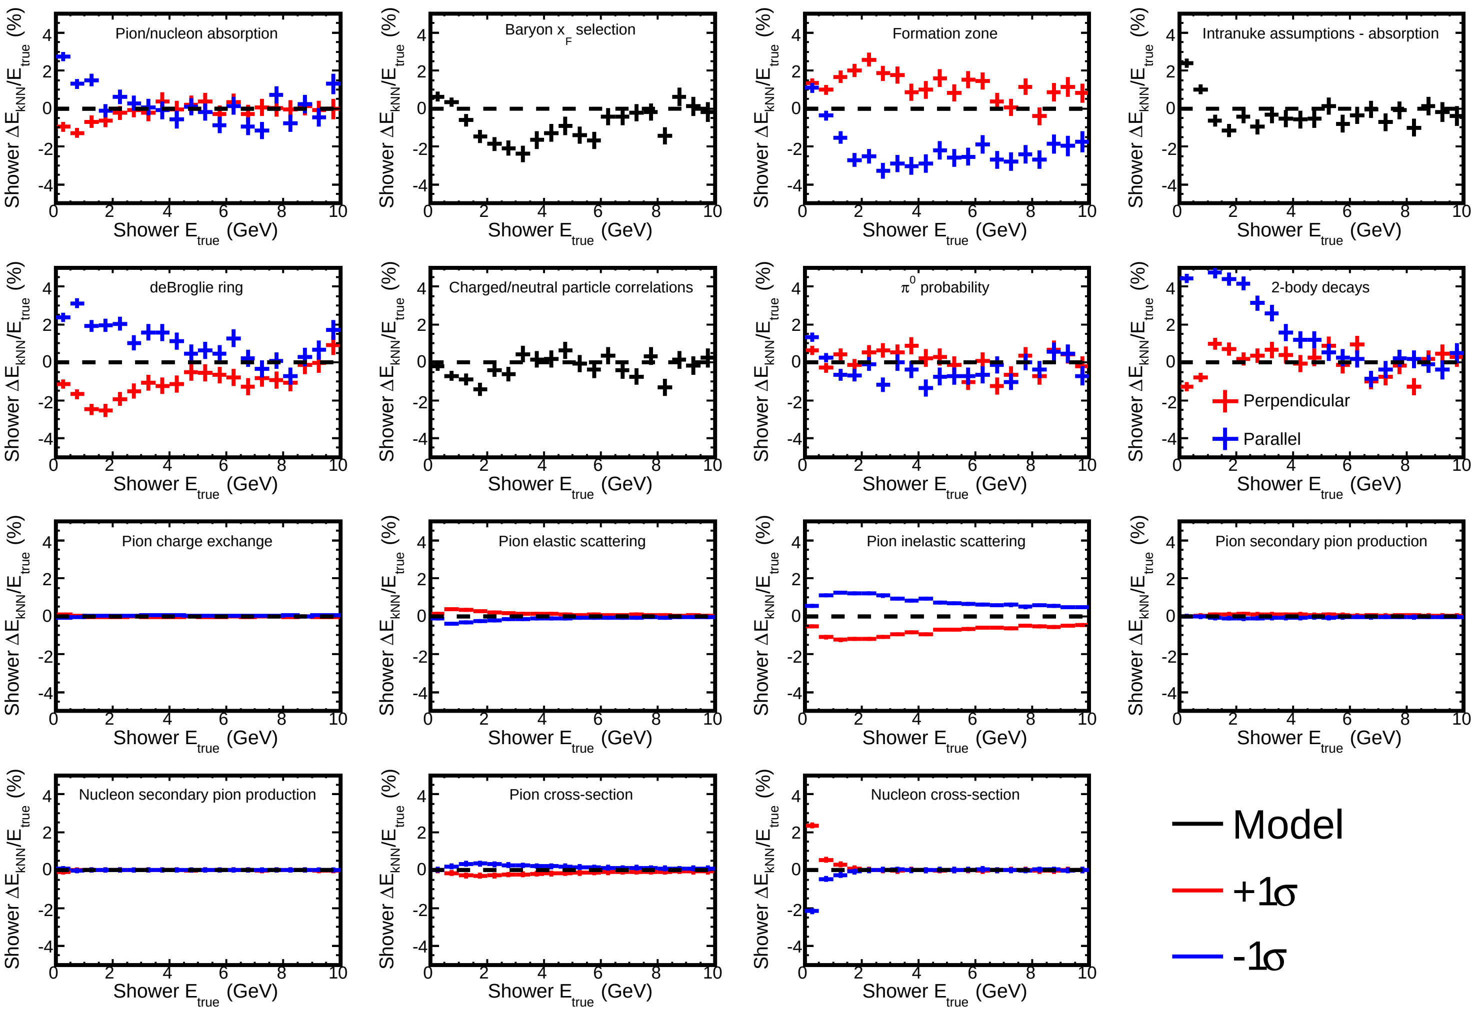

Figure 5.15: Systematic shifts induced by the 15 items listed in Section 5.6.5 as a function of true shower energy, including uncertainty on the estimate. In cases with a $\pm 1\sigma$ parameter change, the two cases are shown by red and blue points. In cases with a discrete model change, the shift is shown by black points. Note that the "2-body decays" case actually consists of two independent model changes, to perpendicular or parallel decays.

5.6.7 Crosschecks

The consistency of the shifted Monte Carlo and INTRANUKE reweighting estimates can be checked, since two of the Monte Carlo samples correspond to equivalent INTRANUKE weights. One sample implements a shift of both the pion and nucleon absorption cross-sections. These are implemented as independent weighting factors. The top plot of Figure 5.16 shows a comparison of these two approaches, where the two INTRANUKE results have been linearly summed to represent a correlated shift. Agreement between the two estimates is reasonable. The bottom plot shows the equivalent comparison for the formation time estimates. Agreement between the two methods here is very good. There is some evidence of the problems applying a reweighting scheme for the negative shift at higher energies, as described in Section 5.6.3. In both cases, the estimate from the dedicated Monte Carlo samples is taken for use in the determination of the total error.

5.6.8 Calibration and single particle response

The response of the MINOS detectors to the passage of single particles (protons, pions and muons) was determined from exposure of CalDet to test beams at CERN [157]. The systematic error associated with such energy measurements comes from uncertainty in the Monte Carlo modelling of the energy deposit of these particles, and from uncertainty in the absolute calibration of the detector channels. This error is quoted in two components: "absolute calibration" is the full systematic error of the measurement, taken to be fully correlated between the two detectors; "relative calibration" is the additional uncorrelated component due to differences between the two detectors, in practice this is implemented as a shift applied in the far detector, but not the near. Based on the largest data/MC discrepancies seen at CalDet, these values are taken as 5.7% and 2.1% in the MINOS analyses [158].

To determine the effect of these uncertainties on the kNN shower energy es-

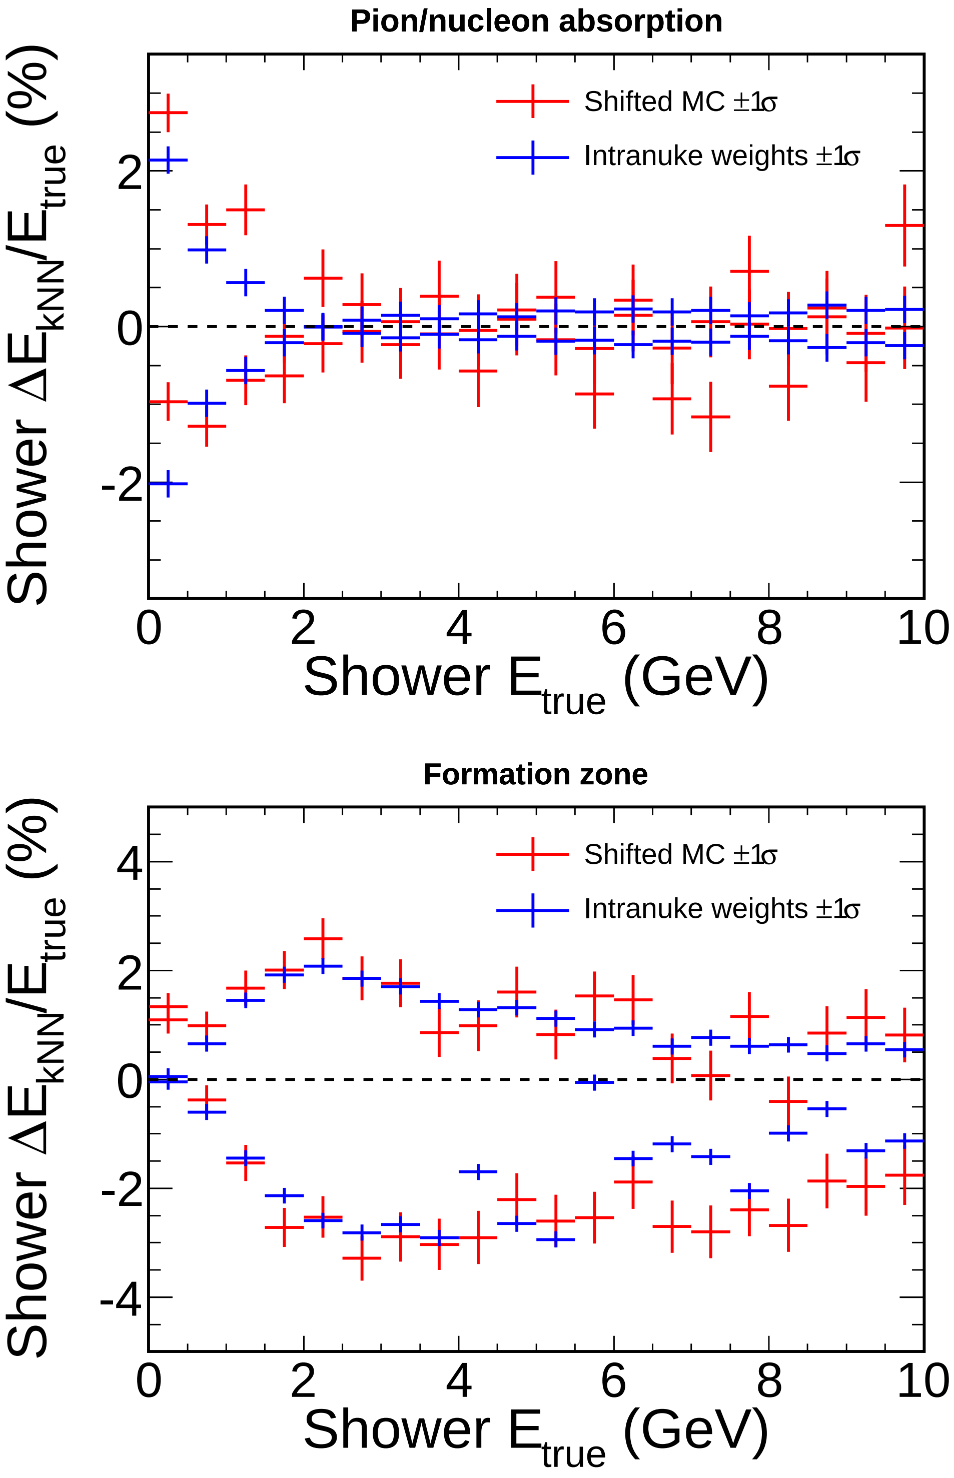

Figure 5.16: Comparison of systematic error estimates from shifted Monte Carlo samples in red and from INTRANUKE reweighting in blue. Top: correlated shift of pion and nucleon absorption cross-sections. Bottom: adjustment of the formation time parameter.

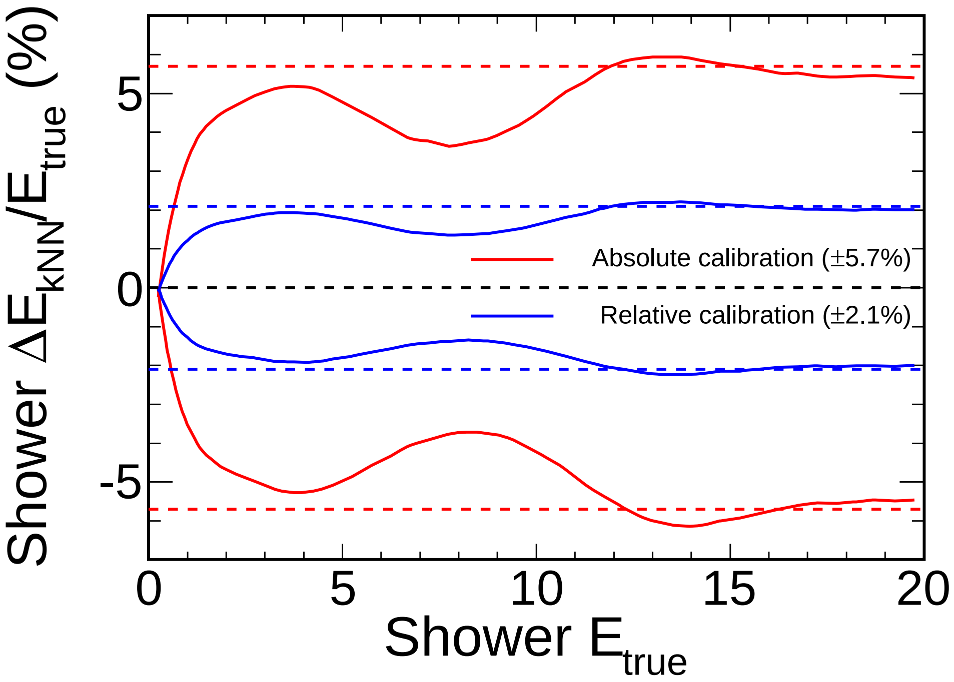

Figure 5.17: Effect of shifting the calorimetric input variables of the kNN by ±2.1% (blue) and ±5.7% (red). The level of uncertainty taken for the calorimetric estimator is shown by the dashed lines, for comparison.

timate, it is necessary to create Monte Carlo samples that reflect such a shift in calibration. Two of the kNN variables are directly affected: the energy in the first two showers, and within 1 m of the vertex. The third variable, length of the primary shower, is unaffected by the scale of the detector response. The samples are generated simply by taking the nominal Monte Carlo and scaling these two variables by the appropriate amount. These samples are then processed through the standard kNN procedure, using the usual unshifted training file. The same procedure is then followed as for the hadronization and rescattering uncertainties. Figure 5.17 shows the effect of ±2.1% and ±5.7% shifts on the kNN output. The uncertainty on the shape is very low, since the Monte Carlo sample is large and the shifted and nominal events are strongly correlated. For the purposes of combination later, the uncertainty is taken as zero. The effect is nearly symmetric for $\pm 1\sigma$ shifts, as might be expected. At high shower energies, where the kNN behaves similarly to the calorimetric estimator, the full error in the calorimetric components is reflected in the results. At lower energies, where the kNN relies more heavily on the number of planes, it is seen to be affected to a somewhat lesser extent by calibration errors.

5.6.9 Gain calibration

Whilst the effects of an overall calibration scale are dealt with in the previous section, there is another point in the calibration chain which can affect the final energy scale. This is the question of the PMT gains: how much charge corresponds to a certain number of photons. In MINOS this is measured by the light-injection procedure. For each strip, LEDs frequently ($\mathcal{O}(0.1\,\mathrm{Hz})$) inject a constant amount of light, and the resultant charge is measured. The uncertainty on this procedure is estimated to be around 5% [159].

Whilst the direct (linear) effect of poorly known gains is corrected out by the rest of the calibration procedure, various threshold effects remain. Most obviously, the two photo-electron cut made at reconstruction time. If the measured gains are too low, the Monte Carlo will incorrectly throw some hits away as they fail to pass threshold. Likewise, high gains will cause some hits to be recorded that would not have been in reality. The kNN estimator is likely more sensitive to these effects than the calorimetric estimator. For example, a difference of a single hit could cause the length of the shower to change, or a small secondary shower not to be reconstructed.

The impact of the gains uncertainty was investigated by taking a Monte Carlo sample (Run III, far detector) and passing it through the full reconstruction[8] with gains set either 5% high or low. The standard procedure is then followed to arrive at the systematic error estimate. Figure 5.18 shows the result. By construction,

---

[8] The output of the Monte Carlo is a series of energy deposits. It is actually the first step of the “reconstruction” that converts these to light levels and then to charges.

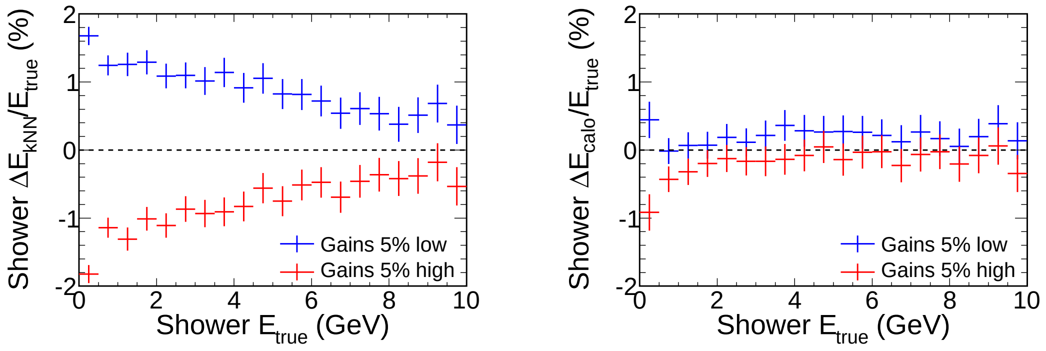

Figure 5.18: The effect of a ±5% shift in the PMT gains in Monte Carlo on kNN energy (left) and calorimetric energy (right).

altering the gain factor in the Monte Carlo does not alter ADC/energy. But with higher gains, the same amount of deposited energy must correspond to less simulated light and a greater chance of the hit failing some threshold. For this reason, the energy estimates for high gains are systematically low, and vice versa. The same effect, calculated for calorimetric energy, is shown for comparison. The impact on the kNN estimator is indeed significantly larger, but not unreasonably so.

5.6.10 Intensity effects

It is conceivable that the kNN estimator might be vulnerable to effects of the beam intensity in the near detector. For example, that activity associated with another event from the same beam spill, the number of which increases with intensity, may be included in the estimate, leading to an intensity-dependent energy bias.

The varying beam intensity is taken into account in the Monte Carlo, as described in Section 3.4. Figure 5.19 shows the shift in $E_{\mathrm{kNN}}/E_{\mathrm{true}}$ for the lowest and highest intensity Monte Carlo samples. Because the intensity of the NuMI beam has generally increased with time, the largest available samples are drawn from Run I (1.24×10^{13} POT/spill) and Run III (3.24×10^{13} POT/spill). The "nominal" dataset used as the basis for comparison is the full Monte Carlo sample

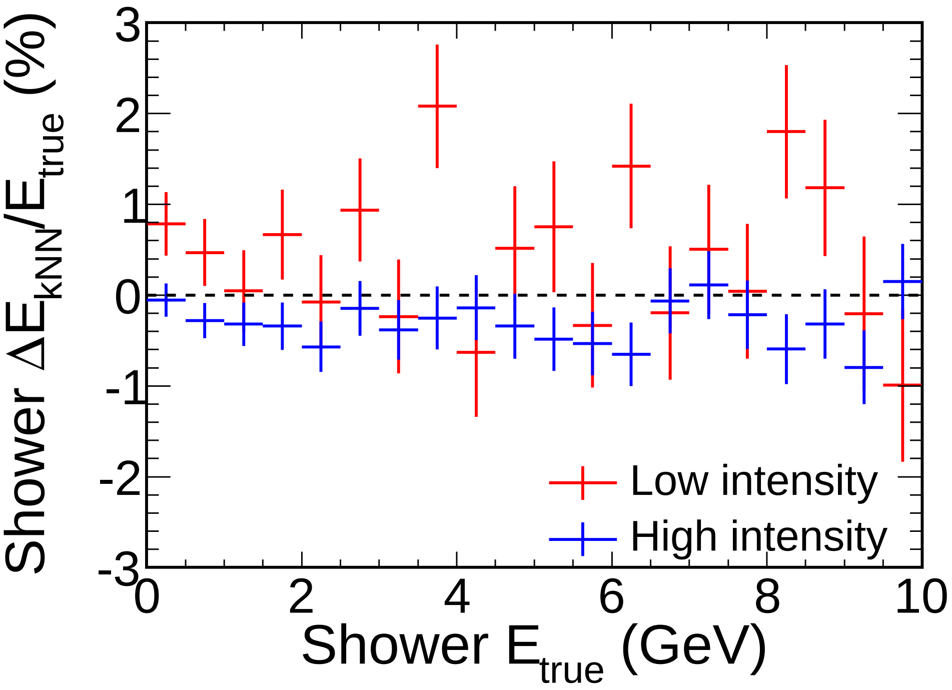

Figure 5.19: Shift of kNN energy with low intensity MC (red) and high intensity MC (blue) relative to the full MC dataset.

over all run periods and intensity conditions. There is evidence of a small ($< 1\%$) effect, but this is in any case an overestimate, approximately corresponding to a case where beam intensity variations are not simulated at all in Monte Carlo. The calorimetric shower energy displays similar behaviour.

Figure 5.20 shows the distribution of E_{kNN} as a function of beam spill intensity in near detector data. There is no visible evidence that the mean energy varies in any way as a function of intensity. Similarly, there is no visible intensity dependence in the Monte Carlo. There is some evidence that the energies are systematically higher in Monte Carlo, but at a level that can be comfortably covered by the uncertainties described in previous sections.

On this evidence, intensity effects are not considered in the construction of the overall kNN systematic error.

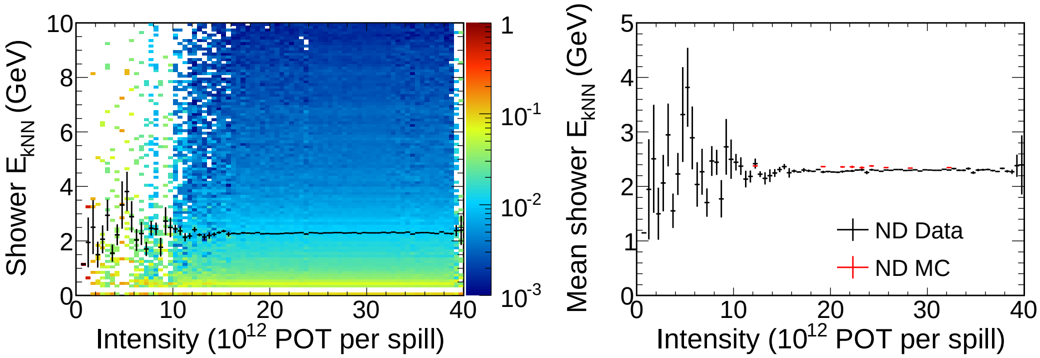

Figure 5.20: Left: distribution of E_{kNN} in near detector data as a function of beam intensity. Right: the mean value of E_{kNN} as a function of beam intensity in data, and at the discrete points simulated in Monte Carlo.

5.6.11 Combination and overall systematic

To calculate the overall error on the absolute kNN shower energy scale, the effects from all the sources considered above are combined. Systematics that were evaluated as $\pm 1\sigma$ shifts are converted to a symmetric, or one-sided, error by averaging the two cases:

$$S_{\pm 1\sigma}(E_{\mathrm{true}}) = \frac{1}{2}\left(|S_{+1\sigma}(E_{\mathrm{true}})| + |S_{-1\sigma}(E_{\mathrm{true}})|\right) \ . \tag{5.8}$$

This procedure is also followed for the "two-body decays" systematic – averaging the effect from the two different dramatic models changes made.

As mentioned in Section 5.6.5, an additional component representing a change to the implementation of the formation time ought to be included. This component is expected to be of similar magnitude to, but uncorrelated with, the already evaluated formation time systematic. For that reason, this effect is scaled by a factor of $\sqrt{2}$, to represent both effects, before the combination is performed:

$$S_{\mathrm{form}}(E_{\mathrm{true}}) \mapsto \sqrt{2} S_{\mathrm{form}}(E_{\mathrm{true}}) \ . \tag{5.9}$$

The full systematic estimate is then the sum in quadrature of all the AGKY and

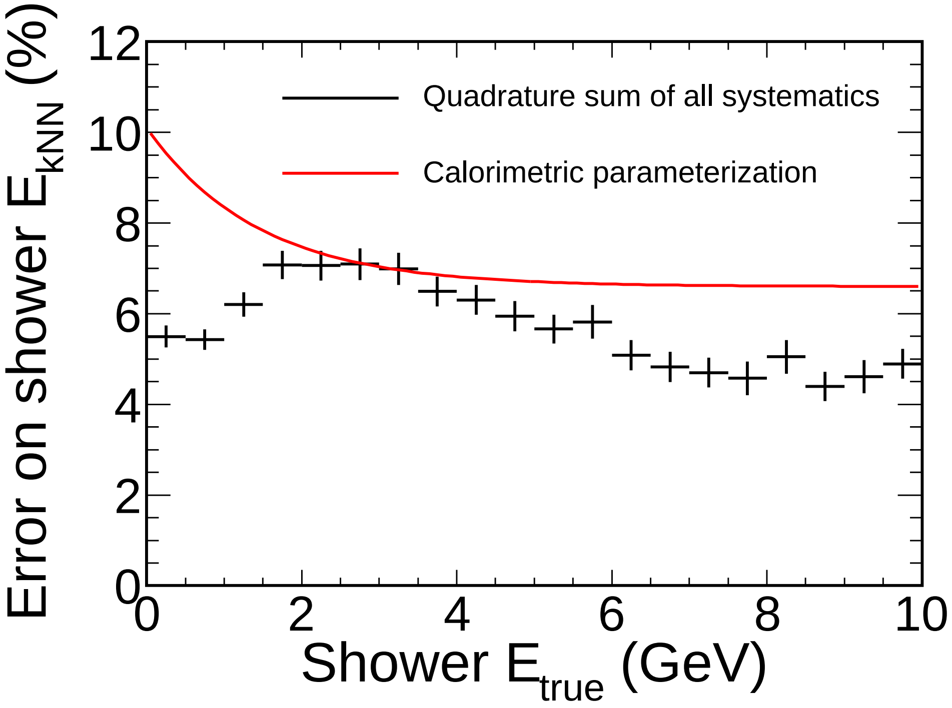

Figure 5.21: Total estimated error on the kNN shower energy estimator from all sources (black points). Also shown, for comparison, is the parameterization previously used for the error on the calorimetric estimator (red curve).

INTRANUKE systematics, and the calibration and gain errors:

$$S_{\mathrm{tot}}(E_{\mathrm{true}}) = \sqrt{\left(\sum_{i}^{\mathrm{had}} S_i^2(E_{\mathrm{true}})\right) + S_{\mathrm{cal}}^2(E_{\mathrm{true}}) + S_{\mathrm{gain}}^2(E_{\mathrm{true}})}\ . \tag{5.10}$$

The uncertainties on the error evaluations are propagated correctly through this combination.

Figure 5.21 shows the resulting systematic estimate as a function of true shower energy. For previous analyses, the error on the calorimetric energy, evaluated in a similar manner, was parameterized with a simple functional form for use in the oscillation fit [160]:

$$\sigma_{\mathrm{abs}} = 6.6\% + 3.5\% \times e^{E_{\mathrm{shw}}/1.44\,\mathrm{GeV}}\ . \tag{5.11}$$

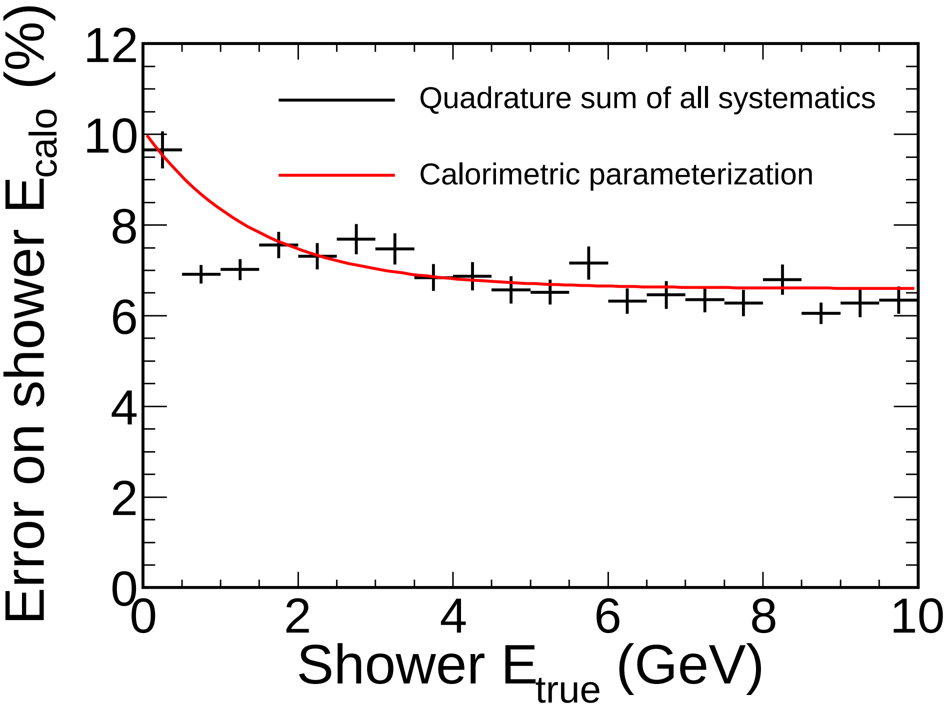

Figure 5.22: Total estimated error on the calorimetric shower energy estimator from all sources (black points), for comparison with Figure 5.21.

This curve is also shown in Figure 5.21 for comparison. The evaluated kNN systematic error does not exceed this level at any energy. The majority of the improvement at low energy is due to the shape of the calibration systematic, as compared to the flat 5.7% used for the calorimetric error.

To cross-check the systematic evaluation procedure, the evaluation of the various components was repeated for the calorimetric estimator, using a flat 5.7% absolute calibration error. Figure 5.22 shows the calorimetric error evaluated this way, again with the previous parameterization for comparison. The agreement is very good, validating the procedure against previous estimates.

Due to time constraints, the error on the shower energy used in the oscillation fit was taken as the old parameterization, Equation (5.11). This is something of an overestimate. In future, a new parameterization specific to the kNN energy could be used (for example, something like a flat 7%).

5.6.12 An alternative error estimate

An alternative route to estimating the systematic on the kNN energy, used as a cross-check, is to take a top-down approach, based on observed data/MC discrepancies, in contrast to the main bottom-up, first principles, evaluation presented above.

The idea is to start from the observed disagreements between near detector data and Monte Carlo in the kNN input variables. The Monte Carlo should then be adjusted so that it matches what is observed in data, regardless of the source of the discrepancies. Passing this modified Monte Carlo through the unmodified kNN will yield a shifted $E_{\mathrm{kNN}}/E_{\mathrm{true}}$ distribution, which can act as an error estimate.

The most obvious way to modify the distribution of a single variable to match what is observed in data is via a reweighting scheme. There is a unique choice of weights to apply to individual events such that the overall distributions match. Specifically, ordering the data and Monte Carlo events by their value for the variable, and forming cumulative distributions, each Monte Carlo event in turn should have its weight adjusted so that the cumulative distribution at that point matches. Unfortunately, when more than one (in this case three) variables need to be made to match, there is no consistent set of weights which will completely achieve this goal.

A better approach is to modify the values of the variables themselves in the Monte Carlo events, leaving the weights unaltered. The advantage is that this can be done independently for all three variables and achieve agreement in all of them simultaneously. Figure 5.23 illustrates the two methods of adjusting events to make a cumulative distribution match. If there were an equal number of data and Monte Carlo events (and no beam reweighting), the algorithm would simply be to sort all the events by the value of one variable, and copy the data values into the Monte Carlo event in corresponding place in sequence, and then repeat

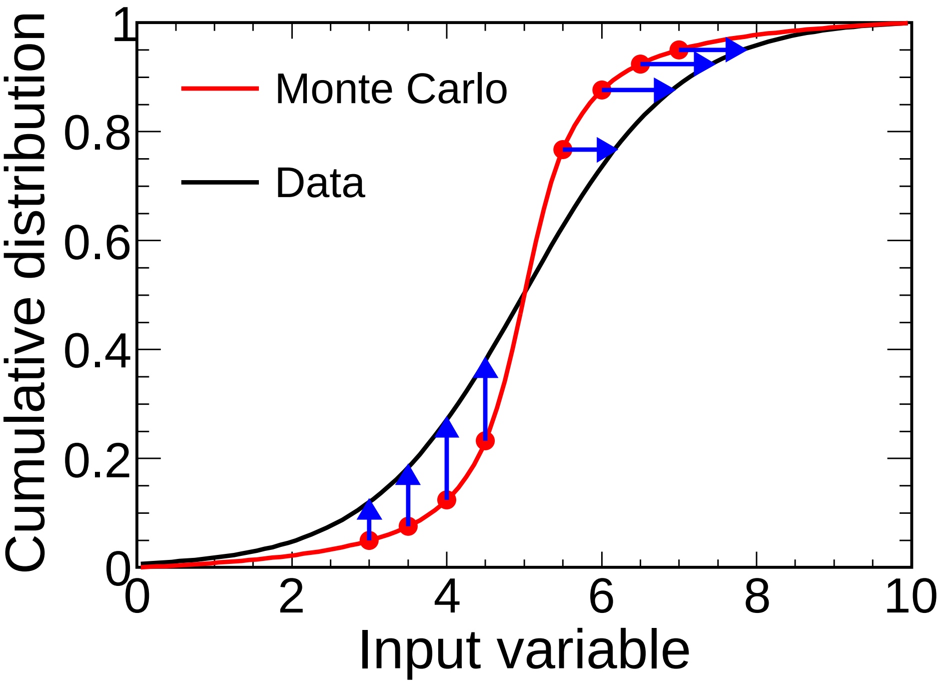

Figure 5.23: Cartoon of two methods to distort one distribution into another. The black and red curves represent (normalized) cumulative distributions in data and Monte Carlo respectively. In the left half, the blue arrows represent the approach of applying a weight to individual events to make the Monte Carlo match the data. In the right half, the arrows represent modifying the value of the variable in individual Monte Carlo events to achieve agreement.

for the other two variables. In practice, the weighted average from the two data events with the closest cumulative weight is used.

Figure 5.24 shows the resultant shift in E_{kNN} from this method, compared with the conventional estimate from Section 5.6.11. The statistical error on this estimate is small, due to the large Monte Carlo sample and correlation between the shifted and unshifted events, and the resulting curve is smooth. The deviation expected in the kNN energy estimate, as determined by this method, is well covered at all energies by the error band estimated previously.

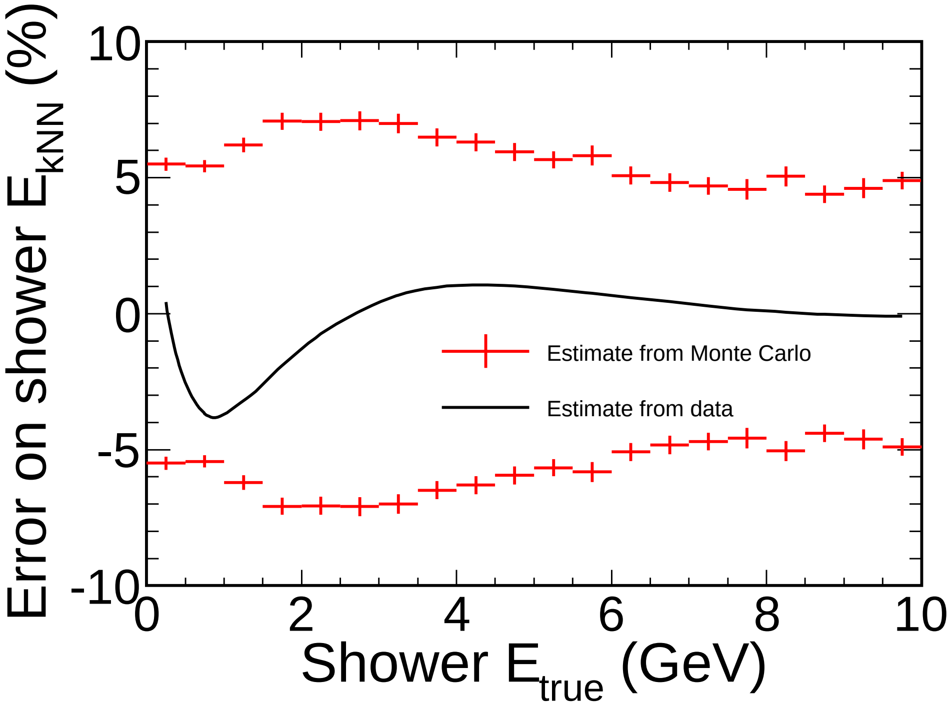

Figure 5.24: The shift to the E_{kNN} distribution when Monte Carlo events are adjusted to match the distributions observed in data (black curve). The red points are the conventionally estimated error (Figure 5.21) for comparison.

5.7 Conclusion

The kNN energy estimator presented in this chapter achieves substantially improved shower energy reconstruction over the previous calorimetric estimator. This leads directly to an improvement in the MINOS oscillation sensitivity, approximately equivalent to the sensitivity gain that would be made with an additional 10^{20} POT of data.

Various sources of systematic error have been considered, and crosschecked. The final estimate of the systematic error on the shower energy scale is slightly smaller than the error on the calorimetric estimator. This error, while still among those with the largest impact, is acceptable for the oscillation analysis.

When combined with other analysis improvements made since the previous (2008) analysis, the overall sensitivity improvement is substantial. Figure 5.25

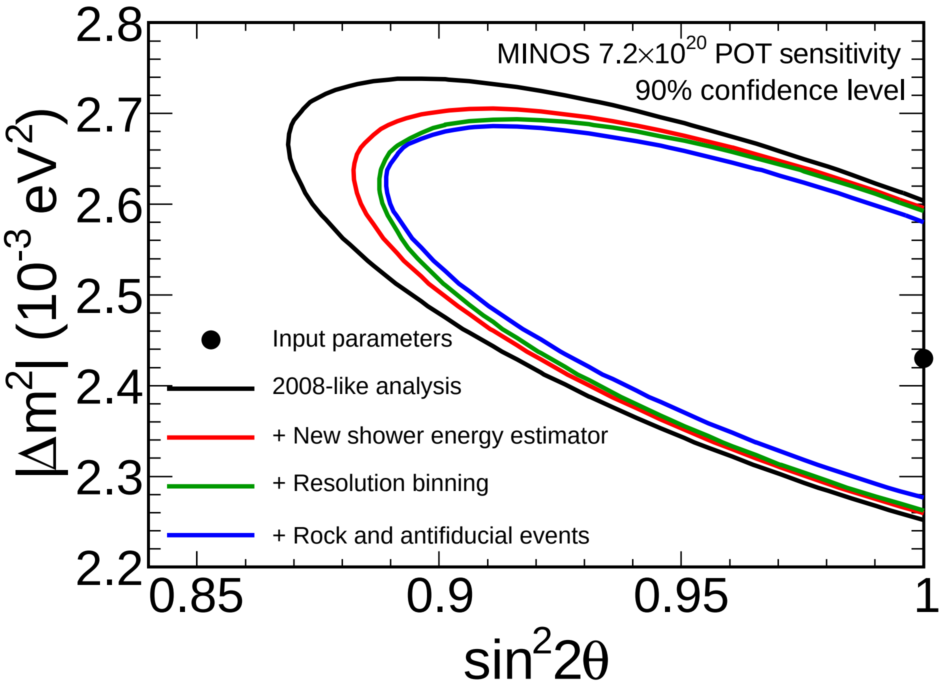

Figure 5.25: MINOS sensitivity to neutrino oscillation parameters with 7.2×10^{20} POT of data. The black curve shows the expected 90% confidence contour before any analysis improvements are made. The red curve shows the new sensitivity when using the kNN shower energy. The green and blue curves show the cumulative improvement towards the final sensitivity obtained by additionally dividing events by bins of resolution, and by including the rock and antifiducial event sample.

shows the cumulative effect on the sensitivity of the analysis as successive improvements are added.

Chapter 6

Results of the MINOS ν_μ and $\bar{\nu}_\mu$ disappearance analyses

This chapter presents, in Section 6.1, the results of the charged-current ν_μ-disappearance analysis described by earlier chapters of this thesis. Section 6.2 then presents the results of the similar analysis performed on data from the subsequent period of antineutrino-mode running. Finally, Section 6.3 briefly discusses the significance and interpretation of these results.

6.1 ν_μ disappearance

This section presents the results of the analysis described in Chapter 4.

Applying the charged-current selection of Section 4.3, and tuning the beam Monte Carlo as described in Section 4.2, near detector data and Monte Carlo spectra are obtained, as shown in Figure 6.1. The data and Monte Carlo spectra are in good agreement, for both positively and negatively charged tracks. The distribution of the $\bar{\nu}_\mu$ background illustrates the decision to include events with positive charge reconstruction in the analyzed data sample. Their statistical contribution is non-negligible and, in the most significant sub-5 GeV region, about

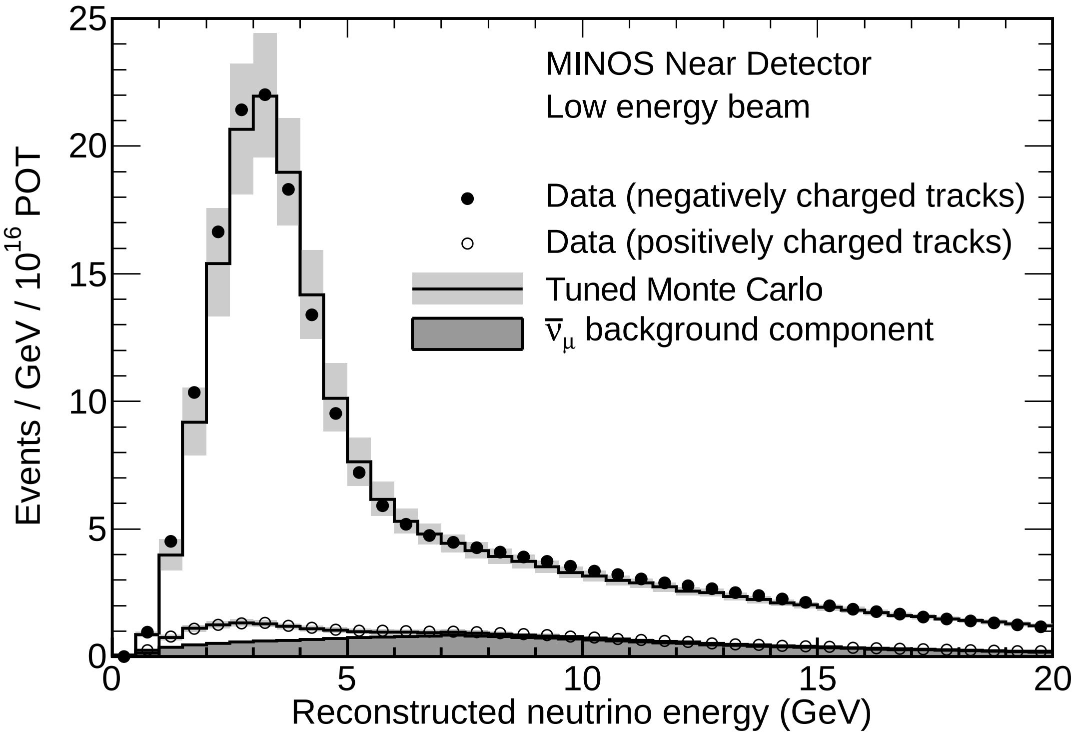

Figure 6.1: The energy spectrum of charged-current selected near detector events, summed over the three low energy run periods. The solid points show the observed spectrum of events with negatively charged tracks. The black line shows the (tuned) Monte Carlo expectation, with 1σ error band. The open points, and associated line show the data and expectation for events with positively charged tracks. The shaded region indicates the component of the positively-charged spectrum expected to be due to $\bar{\nu}_\mu$ interactions.

half of this sample is composed of mis-identified ν_μ events. Including events of both charge-signs in the analysis also helps to mitigate any systematic effect in the charge-sign determination.

The remaining data/MC discrepancies contribute, via the beam matrix technique described in Section 4.4, to the far detector prediction. They also modify the prediction for the rock and antifiducial sample, as described in Section 4.6.

The analysis was carried out in a “blind” fashion. That is, to prevent unintentional bias, the selection criteria, extrapolation technique, fitting procedure and so forth were decided on and finalized before inspecting the far detector data.

Figure 6.2 shows the spectrum obtained when the far detector data was un-

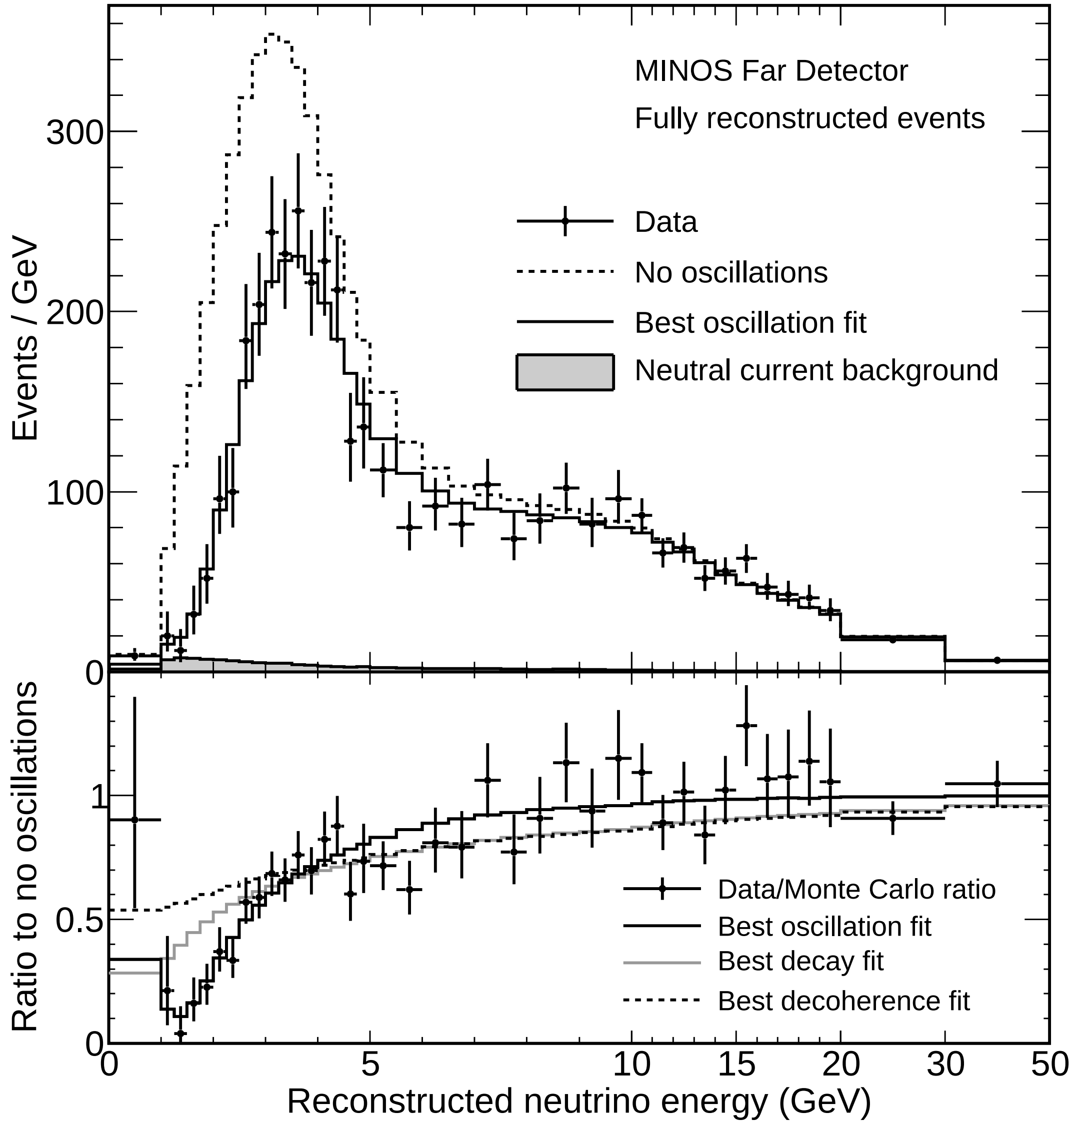

Figure 6.2: The spectrum of charged-current selected events in the far detector (top panel). All runs, resolution bins, and reconstructed charges are summed. The dashed line shows the predicted far detector spectrum in the absence of neutrino oscillations. The points are the observed data spectrum, and the solid line shows the prediction at the best-fit oscillation parameters. The shaded area indicates the neutral-current component expected from Monte Carlo. The bottom panel shows the ratio of the data to the unoscillated prediction (points). The estimated neutral-current background is subtracted from both spectra before the ratio is formed. Also shown are the ratio of the best oscillation fit to the unoscillated case (solid line), and the best fits assuming pure neutrino decay or decoherence scenarios (grey and dashed lines respectively). The x-axis is non-uniform above 10 GeV in order to compactly display the content of the high-energy tail.

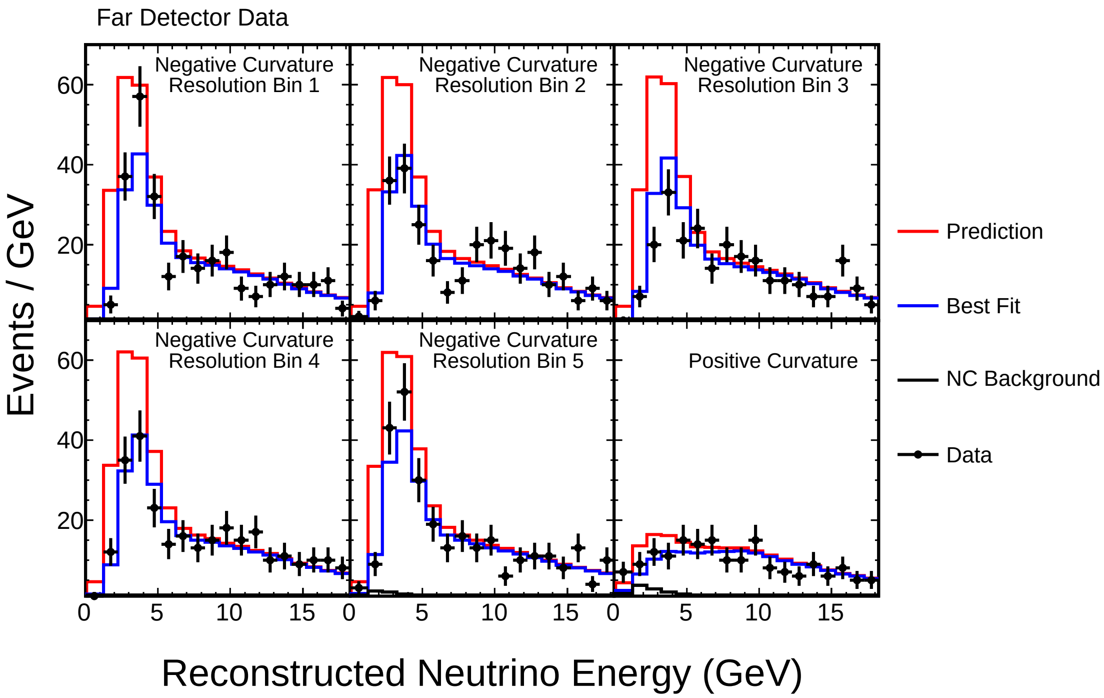

Figure 6.3: Far detector energy spectra of observed data (black points) compared to the unoscillated (red) and best-fit (blue) predictions for each resolution bin. The spectrum of events in the best resolution quantile is shown first, in the top-left panel, followed by the lower resolution bins, and ending with the spectrum of events having positive reconstructed charge.

blinded, compared to the prediction in the case of no oscillations. In all, 1986 events were observed, against an unoscillated prediction of 2451. This numerical deficit alone provides overwhelming evidence of ν_μ disappearance. Figure 6.3 breaks the same data down by resolution bin.

Figure 6.4 shows the spectrum of visible muon energy for observed far detector rock and antifiducial events, again compared to the unoscillated prediction. In this sample, a total of 2017 events were obtained, compared to an unoscillated prediction of 2206.

The event deficits are concentrated at low energies, as expected for neutrino oscillations driven by the atmospheric mass splitting. Figure 6.2 shows the ratio of events observed to the unoscillated prediction as a function of reconstructed energy. The data displays a convincing oscillation dip form, with some evidence

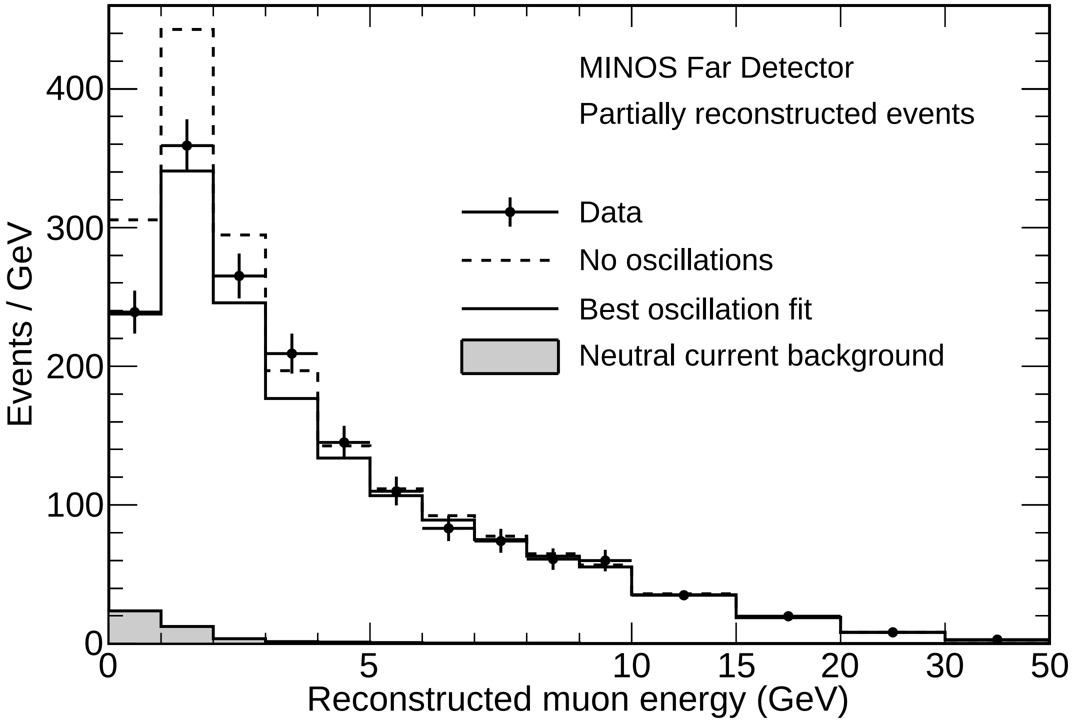

Figure 6.4: The spectrum of selected far detector rock and antifiducial events. All runs and detector regions are summed. The dashed line shows the far detector prediction in the absence of oscillations. The points show the observed data, and the solid line the prediction at the best-fit oscillation parameters. The neutral-current background prediction is shown by the shaded region. Above 10 GeV the scale is compressed, in order to focus on the oscillation region.

of a subsequent rise in the lowest two bins.

While Figures 6.2 and 6.4 show the overall fiducial and RAF spectra for simplicity, the actual oscillation analysis divides the events into separate spectra by run, and by estimated energy resolution in the case of fiducial events (see Section 4.5), or by detector region for RAF events (see Section 4.6).

Performing a maximum-likelihood fit to the model of two-flavour $\nu_\mu \leftrightarrow \nu_\tau$ oscillations, and taking into account the largest systematics as nuisance parameters, as described in Section 4.7, the best-fit oscillation parameters are found to be:

$$|\Delta m^2| = 2.32^{+0.12}_{-0.08} \times 10^{-3}\mathrm{eV}^2 \tag{6.1}$$

$$\sin^2 2\theta > 0.90 \text{ (90\% C.L.)} . \tag{6.2}$$

The best-fit value of $\sin^2 2\theta$ is at maximal mixing[1]. Expressing the result as a measurement, instead of a limit, gives $\sin^2 2\theta = 1.00_{-0.06}$. The quoted errors include both the statistical and systematic contributions. None of the systematic parameters had a best-fit value outside of its 1σ range. Figures 6.2 and 6.4 include the Monte Carlo prediction with oscillations applied at these best-fit parameters. This oscillated prediction is seen to agree well with the observed data.

Also shown in the ratio plot of Figure 6.2, are the best fits to the pure-decay and pure-decoherence models described in Section 2.4.5. It is clear that these models have difficulty reproducing the shape of the deficit at low energies, having bad agreement with the data in this region. Comparing the χ^2 of these fits with the best oscillation fit, these models are disfavoured at 7σ and 9σ respectively. Work is underway to fit the mixed models of Section 2.4.5, in which a small admixture of these effects coexists with standard oscillations [161].

Scanning over the oscillation parameter space, and forming the $\Delta\chi^2$ to the best-fit point, gives the two-dimensional allowed region shown in Figure 6.5. The results of the previous MINOS analysis [69], and of two alternative Super-Kamiokande analyses [54, 55], are shown for comparison. This result is consistent with the earlier measurements, and achieves significantly tighter limits on Δm^2. The Super-Kamiokande results retain an advantage in their measurement of $\sin^2 2\theta$.

These results have been published by the collaboration in [162].

6.2 $\bar{\nu}_\mu$ disappearance

After taking data for the ν_μ analysis in Runs I-III, MINOS then took 1.71×10^{20} POT in the "reversed horn current" (RHC) configuration. In this mode, the current in the magnetic horns is reversed, leading them to focus negative pions and yielding an antineutrino-enriched beam.

---

[1] In fact, if the fit is not constrained to the physical region, the best-fit value for $\sin^2 2\theta$ is very slightly in the unphysical region, at 1.001.

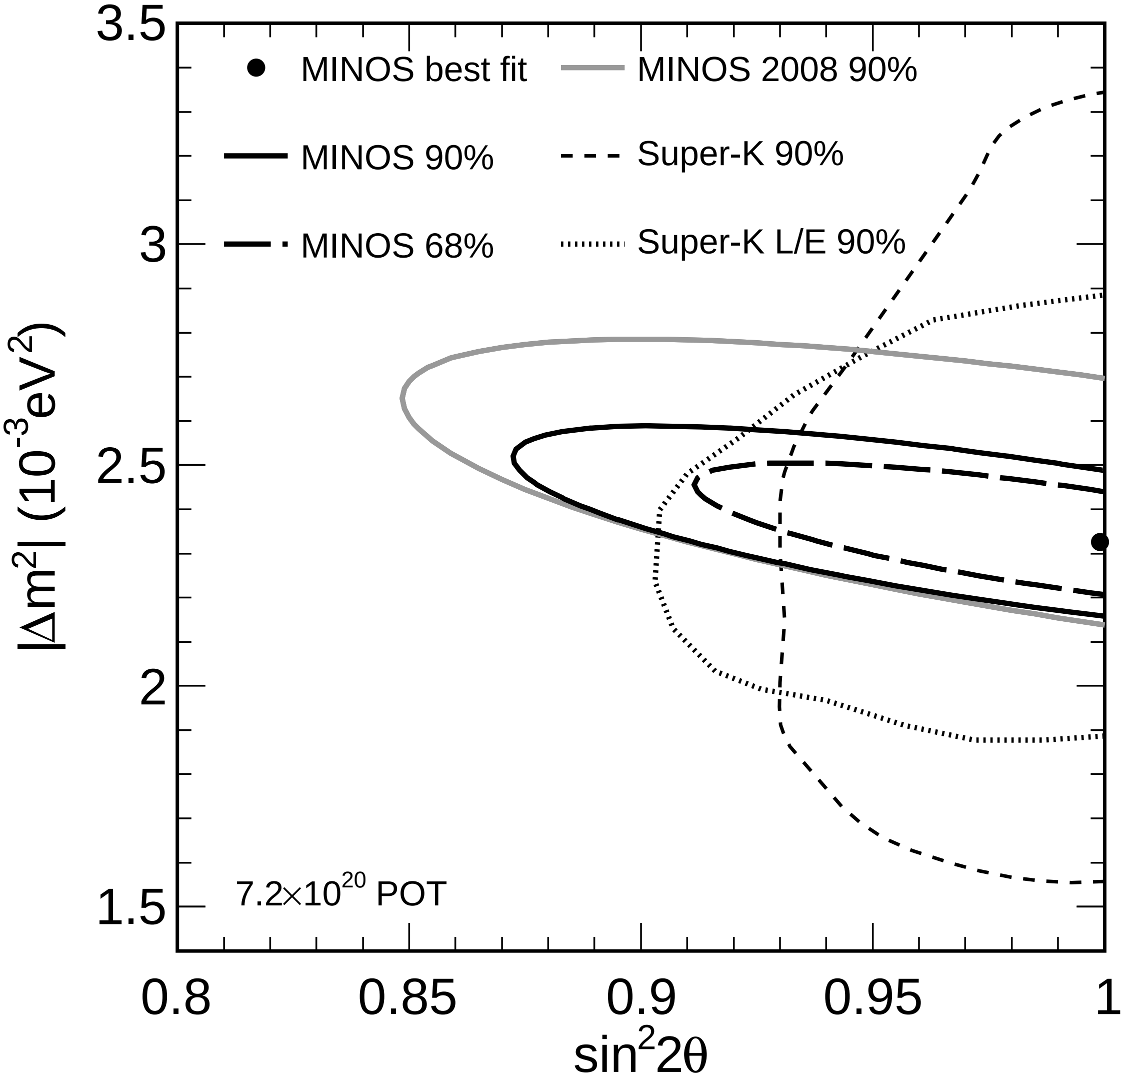

Figure 6.5: Best fit oscillation parameters for the ν_μ charged-current analysis, along with 68% (dashed black) and 90% (solid black) confidence contours. Also shown, for comparison, are the 90% contours from the previous MINOS results (grey), and from Super-Kamiokande's zenith angle [54] and L/E [55] analyses (dashed and dotted lines respectively).

The decision to run in this mode, instead of collecting further neutrino-mode data, which would benefit the ν_μ-disappearance and ν_e-appearance analyses, was in part motivated by the results of an analysis of the antineutrino component of the usual forward horn current beam [128, 163]. The analysis was heavily statistically limited, due to the low proportion of antineutrino "contamination" in the usual beam mode. Nevertheless, a clear deficit was observed: 42 events against an unoscillated expectation of 65. This in itself is unsurprising, what was a surprise was that this is also a deficit compared to a prediction based on the best-fit FHC neutrino-oscillation parameters (58 expected events). Detailed investigations found no flaws in the analysis. Interpreting this deficit in terms of neutrino oscillations, the best-fit point is at a much higher mass-splitting than the neutrino result, and at sub-maximal mixing. The poor statistics meant that the significance of this discrepancy was not particularly high. The decision was therefore taken that Run IV should be in the RHC configuration, where the data would have a large impact on this measurement, compared with the incremental improvement further neutrino data would make.

In most respects, the antineutrino analysis was carried out in a very similar fashion to the analysis described in Chapter 4. Due to the relative cross-sections, both at production and detection, the neutrino contamination in the antineutrino beam is much larger than the antineutrino contamination of the neutrino beam. Since the analysis is also designed to investigate possible differences between neutrinos and antineutrinos, a charge-sign selection $q/p > 0$ is applied. Because of time constraints, and the degree to which the measurement is statistically limited, the resolution binning of Section 4.5, the rock and antifiducial sample of Section 4.6, and the kNN shower energy estimator of Chapter 5, were not used in the analysis. For future analyses, it is intended to include the kNN energy.

Figure 6.6 shows the near detector energy spectra for both positively and negatively charged events. As with the ν_μ analysis, there is good data/MC agreement in both samples. The remaining disagreement in the positively-charged sample is used to adjust the far detector prediction, via the usual beam matrix method. As with the ν_μ analysis, this analysis was performed in a blinded fashion, with all aspects finalized before the far detector data was revealed.

Figure 6.7 shows the unoscillated far detector prediction, along with the observed far detector data. As expected, there is a deficit: 155 events were expected, and 97 observed. Taking the ratio of observed to expected (Figure 6.8) reveals the

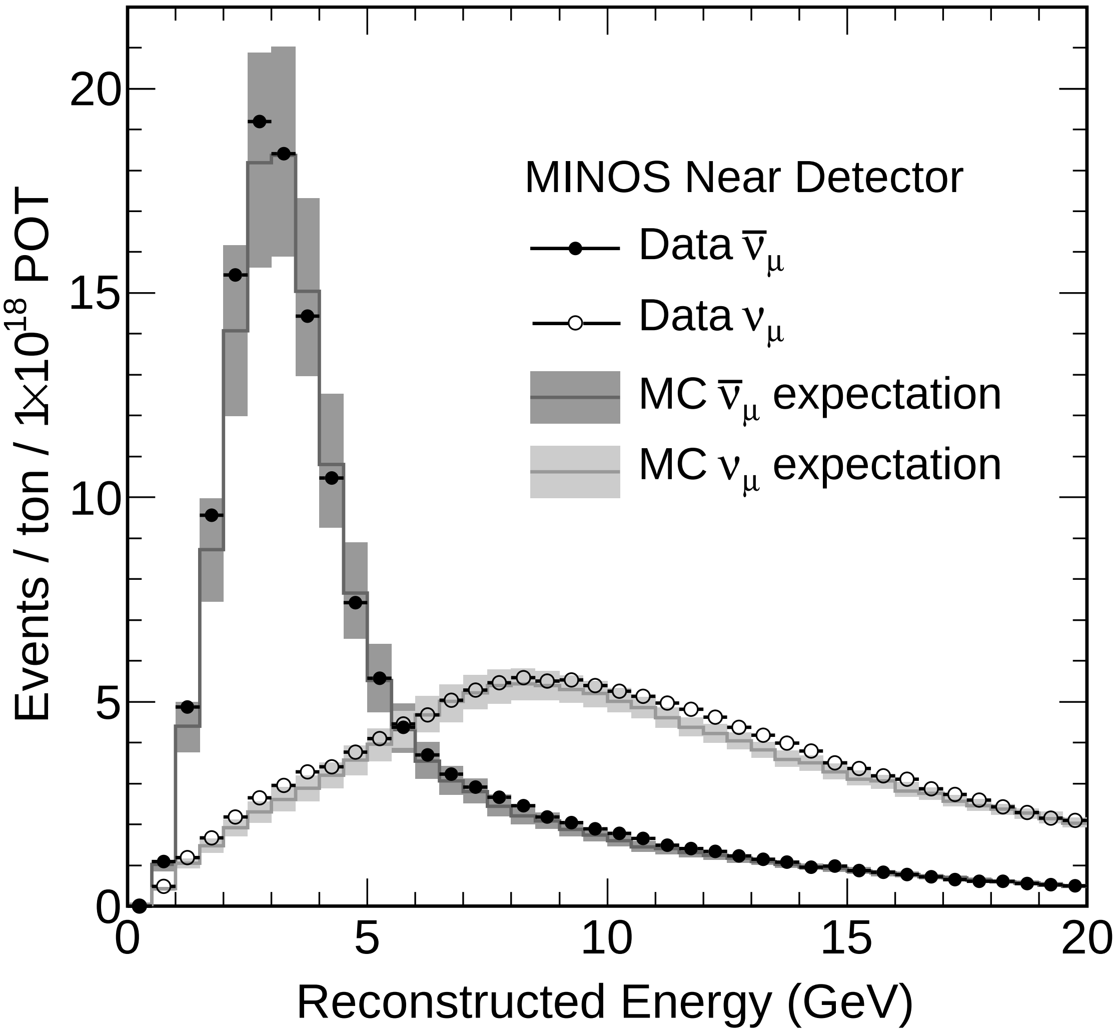

Figure 6.6: Near detector energy spectra from Run IV. The solid and open points represent the observed spectra of events with positively and negatively charged tracks respectively. The associated lines show the tuned Monte Carlo expectation for each case, including a 1σ error band.

usual oscillation dip structure.

This deficit is however, somewhat larger than would be expected from the best-fit neutrino oscillation parameters (Equations (6.1) and (6.2)). Using these values, the expected number of events is 111.

Performing a binned log-likelihood fit to the energy spectrum, the best fit

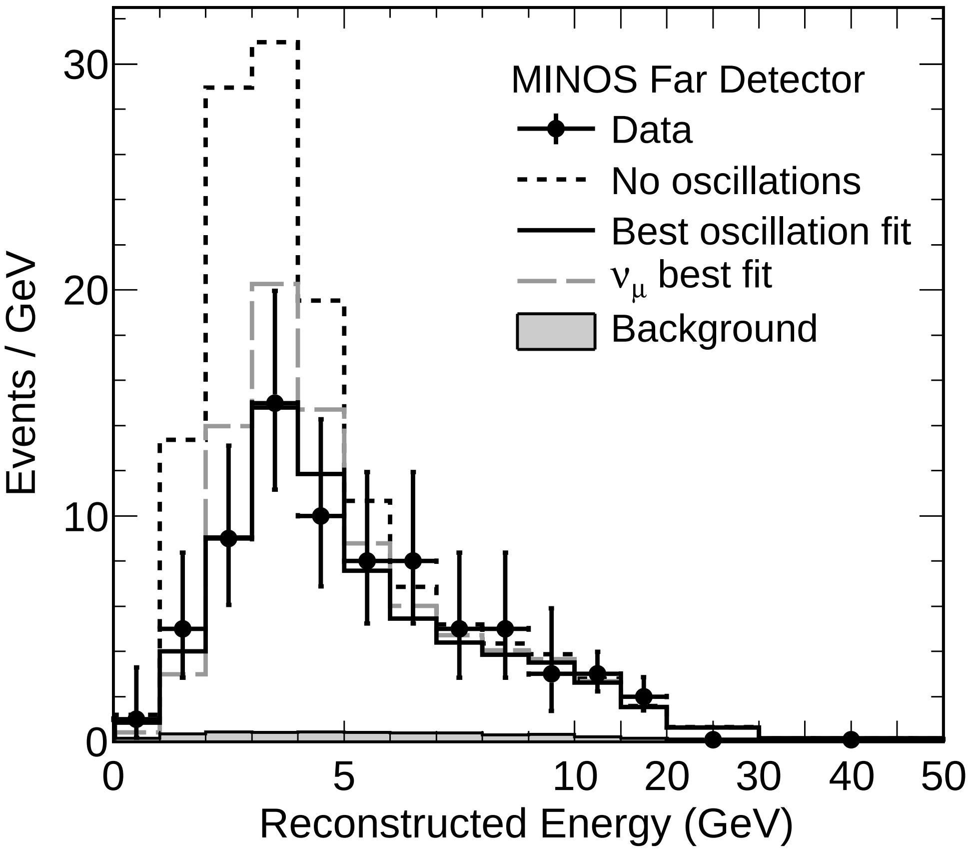

Figure 6.7: The spectrum of selected far detector events from the Run IV antineutrino analysis. The dashed black line shows the unoscillated prediction. The points show the observed data, and the best oscillation fit is shown by the solid line. The filled region shows the predicted background from neutral-current and neutrino-induced charged-current events. The grey dashed line shows the prediction assuming the best-fit oscillation parameters from the neutrino analysis.

antineutrino oscillation parameters are:

$$|\Delta\bar{m}^2| = \left(3.36^{+0.46}_{-0.40}(\text{stat.}) \pm 0.06(\text{syst.})\right) \times 10^{-3}\text{eV}^2 \tag{6.3}$$

$$\sin^2 2\bar{\theta} = 0.86^{+0.11}_{-0.12}(\text{stat.}) \pm 0.01(\text{syst.})\ . \tag{6.4}$$

The Monte Carlo predictions based on these best-fit parameters are shown in Figures 6.7 and 6.8 by the solid black lines. These oscillation parameters provide a good fit to the observed shape of the deficit. These values differ significantly from those measured in the neutrino-mode running. The far detector prediction

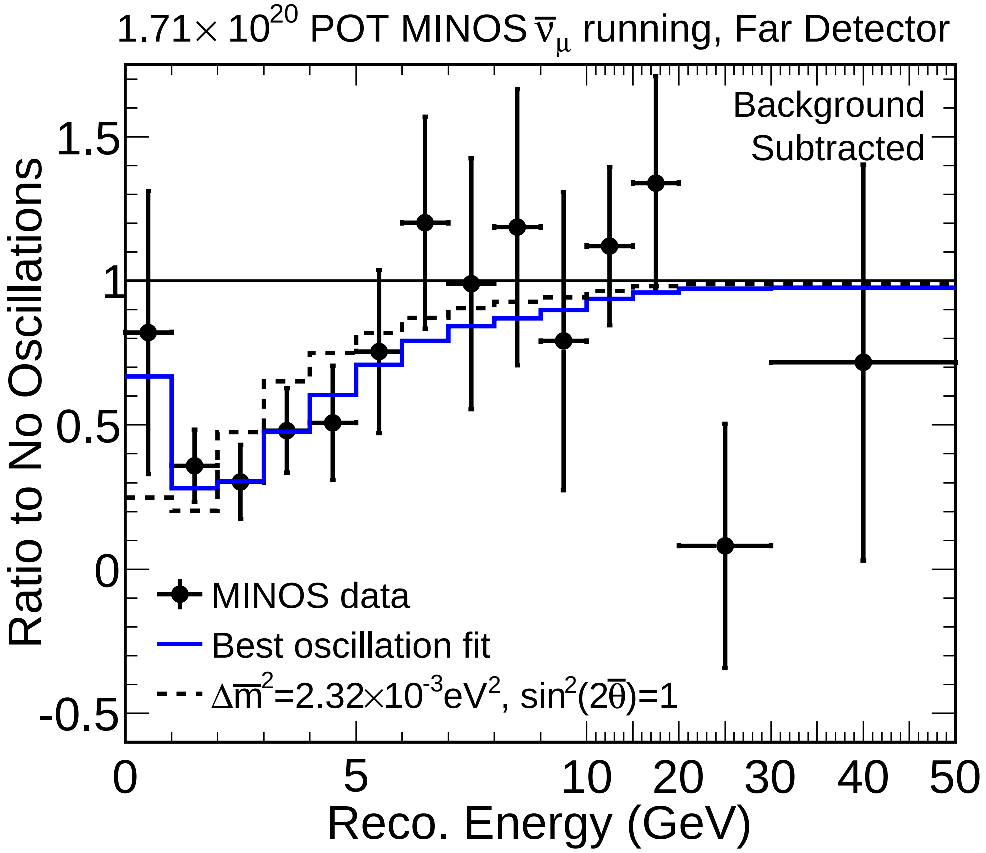

Figure 6.8: Ratio of antineutrino data (points) and best fit prediction (black line) to the unoscillated prediction. Estimated backgrounds were subtracted before the ratio was taken. For comparison, the dashed grey line represents the expected ratio on the basis of the fiducial-only neutrino best fit.

with those parameters is also shown for comparison.

In order to investigate the performance of, and any possible bias in, charge reconstruction in MINOS, the near detector has been operated, at various times, with the magnetic field reversed from the usual direction. During Run III, the field was reversed for multiple periods of around two weeks each, totalling 1.00×10^{20} POT of beam exposure. This configuration, focusing positive muons whilst defocusing negative muons, is known as “FHC+”, in contrast to the usual FHC−. Similarly, during Run VII, a 13 day period (1.09×10^{19} POT) of antineutrino data was taken in the reversed-field (RHC−) configuration. The amount of data that

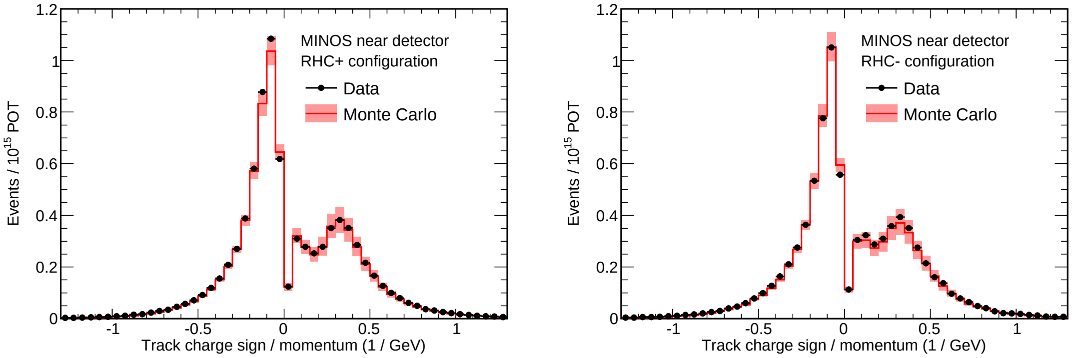

Figure 6.9: The near detector charge/momentum ratio for events passing all other selection cuts. Left: with the magnetic field oriented to focus positively charged particles. Right: with the field oriented to defocus positive particles. In each case the Monte Carlo expectation is shown by the red line, and the observed data by the black points. Data/Monte Carlo agreement is good in both configurations.

could be taken in this configuration was unfortunately limited by MINERνA's preference for conventional field direction in MINOS, allowing them to continue to track muons that exit the back of the MINERνA detector and enter MINOS.

Since the only difference between FHC+/− or RHC+/− is the direction of the magnetic field, any major difference in data/MC agreement would have to be due to mismodelling of the effects of detector geometry on charge determination. Figure 6.9 shows the reconstructed track charge/momentum ratio for the RHC+ and RHC− configurations. The distributions are very similar, and good data/MC agreement is observed in each case. Figure 6.10 shows the same distribution in the far detector. Both neutrino and antineutrino samples ($q/p < 0$ and $q/p > 0$) are consistent, within statistical errors, with their expected distributions.

Calculating correct confidence contours for a low-statistics measurement such as this takes additional care. In the case of a high-statistics data sample, such as the ν_μ analysis, the constant $\Delta\chi^2$ definitions of Table 4.1, which assume the best-fit point is distributed around the true value in a Gaussian fashion, are sufficient. In a lower statistics case, where the confidence contour is expected to extend over a wide range of the parameter space, and the fit to be more heavily constrained

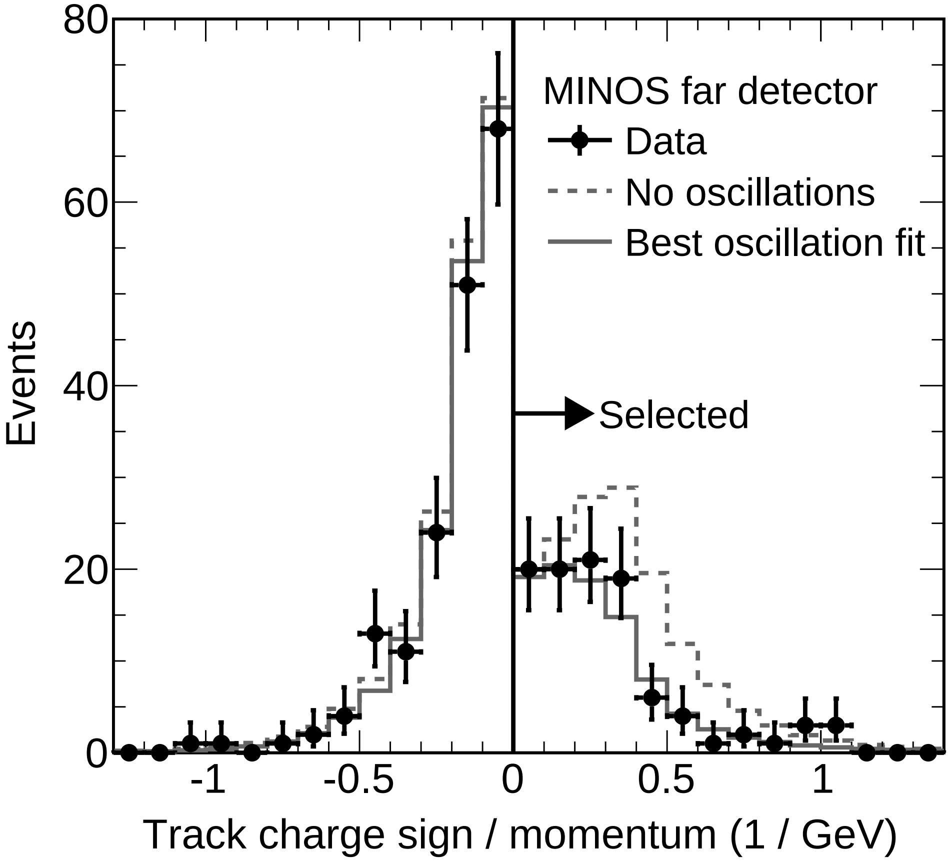

Figure 6.10: The far detector track charge/momentum ratio for events passing all other selection cuts. The observed data is shown by the points, the dashed line represents the prediction in the absence of oscillations, and the solid line shows the expectation at the best-fit oscillation parameters. The negatively-charged neutrino events to the left of the plot are almost unaffected by oscillations due to their higher average energy.

by the physical boundaries, it is necessary to take a more empirical approach to ensure the correct statistical properties.

The $\bar{\nu}_\mu$ analysis makes use of the Feldman-Cousins technique [164] in the construction of its confidence contours. The desired property is that of correct "coverage". That is, if the experiment were repeated multiple times, the true oscillation parameters should be enclosed by the 90% contours in 90% of the trials. At each point in the $(\sin^2 2\bar{\theta}, \Delta\bar{m}^2)$ parameter space, a large number of mock experiments are created. These consist of Monte Carlo, oscillated with these true parameters, with Poisson fluctuations applied to emulate the appearance of low-statistics data.

Each of these samples is then input to the usual maximum-likelihood fit, and the best-fit point that would have been obtained in each trial is determined. The $\Delta\chi^2$ between the best-fit and true oscillation parameters provides the contour level definition that would have been required for the true point to be enclosed by the contour in this case. The value of this $\Delta\chi^2$ which would have caused 90% of these experiments to have covered the true point is then recorded at each point in the $(\sin^2 2\bar{\theta}, \Delta\bar{m}^2)$ space. In this way, a map of critical $\Delta\chi^2$ values is built up across the whole parameter space. When the final fit to the real data is performed, the $\Delta\chi^2$ as a function of the oscillation parameters is compared to this map. All those points where the $\Delta\chi^2$ exceeds the value at the corresponding point in the map are outside of the final confidence contour, and those that fall below the critical value at that point are within the contour. This method guarantees the correct coverage properties, even in the presence of large departures from gaussian behaviour. In addition, it provides a simple way to include the effect of systematic errors, without the need to introduce additional fit parameters as in Section 4.7. For each fake experiment, in the construction of the critical $\Delta\chi^2$ map, in addition to the introduction of Poisson fluctuations, random values of the systematic shifts are also chosen, with their assumed distributions, and used to adjust the generated spectra. In this way, the constituent experiments fully reflect the degree of variation expected in the real data.

Figure 6.11 shows the 68% and 90% confidence contours for the $\bar{\nu}_\mu$ analysis, determined by this technique. Also shown are the ν_μ contours from Figure 6.5, and the results of a global fit for the antineutrino parameters, based on data from other experiments [165] (which is dominated by the Super-Kamiokande atmospheric neutrino analysis), for comparison. As in the ν_μ analysis, this measurement is the world's most precise in $\Delta\bar{m}^2$, whilst achieving somewhat lower sensitivity in $\sin^2 2\bar{\theta}$. This plot gives some indication of the significance of the discrepancy between the two measurements, neither best-fit point is within the other measurement's 90%

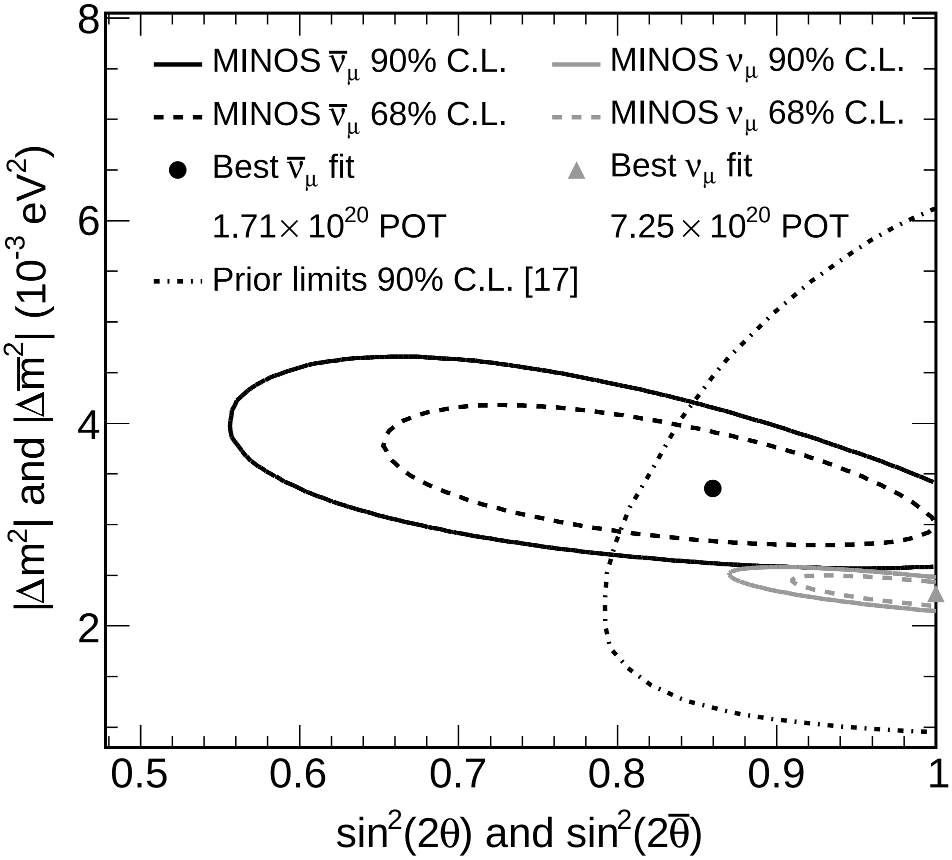

Figure 6.11: Best fit oscillation parameters for the antineutrino charged-current analysis, along with 68% (dashed black) and 90% (solid black) confidence contours. Also shown for comparison, in grey, are the best fit parameters and contours from the neutrino analysis (Figure 6.5). The dotted line shows the results of a previous global fit for the antineutrino parameters (without MINOS data) [165].

confidence contour. The probability of such a discrepancy occurring by chance has been evaluated to be at the few percent level.

The results of this analysis are also currently being prepared for publication [166].

6.3 Discussion

The statistical significance of this observed discrepancy between the oscillations of neutrinos and antineutrinos is still sufficiently low that there is a substantial chance the entire effect is simply a statistical fluctuation, which will be resolved by future data.

However, especially considering the apparent different behaviour of neutrinos and antineutrinos seen by MiniBooNE (Section 2.9), the possibilities are intriguing.

As argued in Section 2.4.1, the survival probabilities for neutrinos and antineutrinos in vacuum are guaranteed to be identical by $\mathcal{CPT}$ symmetry. Any departure from this identity would be a signal of $\mathcal{CPT}$-violation.

Accepting the discrepancy as a real effect, we may be saved from the introduction of so extreme a hypothesis by the fact that the MINOS experiment does not take place in a vacuum. The neutrino beam passes through the Earth's crust, composed of matter, and interacts in matter detectors. The symmetry between the two measurements is therefore not complete. Perhaps there is some mechanism that can generate a neutrino-antineutrino difference in these circumstances?

The phenomenon of matter effects (discussed in Section 2.4.4) provides an example of such a mechanism. However, the impact of matter effects on a ν_μ- or $\bar{\nu}_\mu$-disappearance experiment over the MINOS baseline is completely negligible.

One possibility that has been proposed and investigated [167, 168], is to introduce a new, non-standard, interaction between neutrinos and matter. This interaction serves to enhance or inhibit the conversion between ν_μ and ν_τ, and leads to an apparent adjustment of the underlying oscillation parameters, in a similar fashion to conventional matter effects. This possibility is also under investigation within the MINOS collaboration [169], which will allow official fit parameters and confidence limits to be produced under the assumptions of such a model.

Chapter 7

Summary and outlook

MINOS has now been running for almost six years. The detectors have performed very reliably over this period, and the power of the NuMI beam has continually increased. Over 10^{21} protons-on-target have been delivered since the start of operations.

With the data accumulated in the first four run periods (to March 2010), MINOS has made the most precise measurements of the atmospheric mass splitting for neutrinos $\Delta m^2_{\mathrm{atm}}$ [162], and for antineutrinos $\Delta\bar{m}^2_{\mathrm{atm}}$ [166]. The corresponding measurements of $\sin^2 2\theta$ and $\sin^2 2\bar{\theta}$ are consistent with the Super-Kamiokande best-fit to maximal mixing.

Searching for ν_e appearance in the first three neutrino-mode runs, MINOS sets limits on θ_{13} [68], competitive with the previous limits from the CHOOZ experiment [65].

This thesis focused on the details of the ν_μ disappearance analysis. Chapter 4 gave an overview of the full analysis: how charged-current events are effectively selected, and how the energy spectrum observed in the near detector is used in the construction of the far detector expectation. This chapter also described two improvements that were new for this analysis: binning events by their estimated energy resolution, and a new sample of events originating from outside the fiducial

Name	Start	End	FD exposure	Comments
Run V	Mar 2010	May 2010	4.59×10^{19} POT	
Run VI	May 2010	Jul 2010	6.16×10^{19} POT	
Run VII	Nov 2010	Mar 2011	1.24×10^{20} POT	Antineutrino beam
Run VIII+	Mar 2011	Mar 2012	$\sim 3 \times 10^{20}$ POT	Until NOνA upgrade

Table 7.1: Unanalyzed and future MINOS run periods. See text for details.

volume.

Chapter 5 described in detail another improvement that was developed for this round of analysis. This new technique was motivated by the sensitivity improvement potentially available from better energy resolution, particularly of the shower component of an event. A wide variety of potential estimators based on a kNN technique were tested, and the one yielding the largest performance improvement selected. The raw kNN output was corrected to ensure that the final estimator was unbiased. The final sensitivity improvement achieved was comparable to taking an additional 10^{20} POT of neutrino data. Evaluating all potential sources of systematic error, the final uncertainty on the hadronic energy scale introduced by this new technique was found to be no larger than with the previous approach, and the kNN energy was adopted for the analysis.

Chapter 6 presented the results of this analysis. The best-fit point and the confidence contours are consistent with previous measurements of the atmospheric oscillation parameters, and provide the most precise measurement of Δm^2 to date. Exotic explanations of the ν_μ deficit: neutrino decay and decoherence, were disfavoured at high significance. This chapter also presented the results of the similar analysis performed on the Run IV antineutrino data. This is the only dedicated analysis to date of antineutrino oscillations over the atmospheric baseline, and improves on previous world limits for $\Delta\bar{m}^2_{\mathrm{atm}}$. The tension between the oscillation parameters measured from the neutrino and antineutrino data samples was unanticipated. Such a discrepancy cannot be explained under the standard picture of neutrino oscillations.

The most direct approach to investigating this puzzle is simply to take more antineutrino data, which MINOS is currently doing with "Run VII" (see Table 7.1). This run is expected to approximately double the current antineutrino dataset. If the observed effect is real, inclusion of this data should somewhat increase the statistical significance. If the current results are merely a statistical fluctuation, the analysis should show some reversion towards the ν_μ best-fit. In either case the evidence is unfortunately unlikely to be conclusive.

Within MINOS, taking antineutrino data in place of neutrino data significantly hurts the ν_e analysis, and obviously adds no further contribution to the ν_μ analysis. Increasing the $\bar{\nu}_\mu$ statistics, however, has been set as the top priority.

External constraints also impact on the ability of MINOS to take as much antineutrino data as desired. The MINERνA experiment [170] is situated upstream of the near detector, and is designed to make precision measurements of neutrino cross-sections, which will be of help in future oscillation experiments. They prioritize neutrino-mode running, and are given equal weight by the laboratory. The current compromise position is that Run VII will end, and NuMI will revert to neutrino-mode, sometime in May 2011.

The time constraint on all of these decisions is the planned NuMI shutdown, from March 2012 to February 2013, to upgrade the beam for the NOνA experiment. Once the shutdown ends, although the beam power should be greatly increased, the beamline will be operated in something similar to the medium energy configuration, and it will in fact be physically impossible to return to the low energy configuration. The vast majority of the medium energy beam flux is outside of the oscillation region of interest to MINOS and MINERνA[1]. At this point, it is likely that the MINOS far detector will cease operations.

In addition to the analysis of the $\bar{\nu}_\mu$ data currently being collected, a ν_e-

---

[1] The NOνA far detector is being constructed at an off-axis angle of 0.8°. At this angle the beam peak once again falls in the few-GeV oscillation region.

appearance analysis of the expanded neutrino dataset, incorporating several analysis improvements, is planned. A unified analysis of the complete MINOS dataset is also a possibility.

Within the framework of standard neutrino oscillations, there are three parameters that remain undetermined. These are: the mixing angle θ_{13}, the $\mathcal{CP}$-violating phase δ, and the mass hierarchy. Three reactor experiments, searching for θ_{13} via ν_e disappearance, are currently under construction: Double-Chooz [67], Daya Bay [171], and RENO [172]. The initial goal of the two long-baseline experiments, T2K [173] and NOνA [174], is to search for θ_{13} via ν_e appearance in their ν_μ beams. Of these, T2K has recently begun operations, and NOνA is under construction. If the value of θ_{13} is large, near the current limit from CHOOZ and MINOS, then one of these experiments will discover it in the next few years. If the value of θ_{13} is small, or even zero, then these experiments will set stringent new limits. The long-baseline experiments will also make greatly improved measurements of the atmospheric parameters, for both neutrinos and antineutrinos, confirming or denying the presence of the effect hinted at by MINOS within about a year of running [168]. Assuming a sufficiently large value of θ_{13}, NOνA will have some sensitivity to δ and the hierarchy by comparing the results of neutrino and antineutrino data taking, taking into account the expected matter effects. In the best case, they can both determine the hierarchy and make a crude measurement of δ. In the worst case, the two effects cancel and they can merely rule out some combinations of possibilities.

In the case where θ_{13} is discovered, but found to be small, or where the signals from δ and the hierarchy are ambiguous, the proposed LBNE experiment could investigate further. This experiment would consist of a new neutrino beam from Fermilab to the Homestake mine in South Dakota. The far detector would be a 100 kt scale water-Čerenkov or multi-kiloton Liquid Argon detector.

With a robust programme to measure the last remaining pieces of the oscillation puzzle, a couple of tantalizing anomalies, and the increasing sensitivity of non-oscillation experiments to the absolute neutrino mass scale, the next decade in neutrino physics looks set to be as exciting as the previous two have been.

Bibliography

[1] CDF Collaboration: F. Abe *et al.*, Observation of Top Quark Production in $p\bar{p}$ Collisions with the Collider Detector at Fermilab, Phys. Rev. Lett. **74**, 2626 (1995).

[2] D0 Collaboration: S. Abachi *et al.*, Search for High Mass Top Quark Production in $p\bar{p}$ Collisions at $\sqrt{s} = 1.8$ TeV, Phys. Rev. Lett. **74**, 2422 (1995).

[3] W. Pauli, On the Earlier and more recent history of the neutrino, Camb. Monogr. Part. Phys. Nucl. Phys. Cosmol. **14**, 1 (2000).

[4] J. Chadwick, The Existence of a Neutron, Proceedings of the Royal Society of London **A136**, 692 (1932).

[5] E. Fermi, Versuch einer Theorie der β-Strahlen. I., Zeitschrift für Physik **88**, 161 (1934).

[6] W. Heisenberg, Zur Theorie der "Schauer" in der Höhenstrahlung, Zeitschrift für Physik **101**, 533 (1936).

[7] H. Bethe and R. Peierls, The Neutrino, Nature **133**, 532 (1934).

[8] F. Reines and C. Cowan, Detection of the Free Neutrino, Phys. Rev. **92**, 830 (1953).

[9] 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).

[10] T. D. Lee and C. N. Yang, Question of Parity Conservation in Weak Interactions, Phys. Rev. **104**, 254 (1956).

[11] C. S. Wu, E. Ambler, R. W. Hayward, D. D. Hoppes, and R. P. Hudson, Experimental Test of Parity Conservation in Beta Decay, Phys. Rev. **105**, 1413 (1957).

[12] M. Goldhaber, L. Grodzins, and A. W. Sunyar, Helicity of Neutrinos, Phys. Rev. **109**, 1015 (1958).

[13] H. Yukawa, On the Interaction of Elementary Particles, Proc. Phys. Math. Soc. Japan **17**, 48 (1935).

[14] S. L. Glashow, Partial Symmetries of Weak Interactions, Nucl. Phys. **22**, 579 (1961).

[15] A. Salam and J. C. Ward, Electromagnetic and weak interactions, Phys. Lett. **13**, 168 (1964).

[16] S. Weinberg, A Model of Leptons, Phys. Rev. Lett. **19**, 1264 (1967).

[17] Gargamelle Neutrino Collaboration: F. J. Hasert *et al.*, Search for elastic muon neutrino electron scattering, Phys. Lett. **B46**, 121 (1973).

[18] 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).

[19] 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).

[20] UA1 Collaboration: G. Arnison *et al.*, Experimental observation of lepton pairs of invariant mass around 95 GeV/c^2 at the CERN SPS collider, Phys. Lett. **B126**, 398 (1983).

[21] UA2 Collaboration: R. Banner *et al.*, Observation of Single Isolated Electrons of High Transverse Momentum in Events with Missing Transverse Energy at the CERN $\overline{p}p$ Collider, Phys. Lett. **B122**, 476 (1983).

[22] 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).

[23] E. J. Konopinski and H. M. Mahmoud, The Universal Fermi Interaction, Phys. Rev. **92**, 1045 (1953).

[24] G. Danby *et al.*, Observation of High-Energy Neutrino Reactions and the Existence of Two Kinds of Neutrinos, Phys. Rev. Lett. **9**, 36 (1962).

[25] M. L. Perl *et al.*, Evidence for Anomalous Lepton Production in $e^+ - e^-$ Annihilation, Phys. Rev. Lett. **35**, 1489 (1975).

[26] Precision electroweak measurements on the Z resonance, Phys. Rept. **427**, 257 (2006).

[27] DONUT Collaboration: K. Kodama *et al.*, Observation of Tau Neutrino Interactions, Phys. Lett. **B504**, 218 (2001).

[28] DONUT Collaboration: K. Kodama *et al.*, Final tau-neutrino results from the DONuT experiment, Phys. Rev. D **78**, 052002 (2008).

[29] W.-M. Yao *et al.*, Review of Particle Physics, J. Phys. G **33**, 1 (2006).

[30] R. Davis, D. S. Harmer, and K. C. Hoffman, Search for neutrinos from the sun, Phys. Rev. Lett. **20**, 1205 (1968).

[31] B. T. Cleveland *et al.*, Measurement of the Solar Electron Neutrino Flux with the Homestake Chlorine Detector, The Astrophysical Journal **496**, 505 (1998).

[32] Kamiokande-II Collaboration: K. S. Hirata *et al.*, Observation of ^{8}B solar neutrinos in the Kamiokande-II detector, Phys. Rev. Lett. **63**, 16 (1989).

[33] A. I. Abazov *et al.*, Search for neutrinos from the Sun using the reaction ^{71}Ga(ν_e,e^-)^{71}Ge, Phys. Rev. Lett. **67**, 3332 (1991).

[34] GALLEX Collaboration: P. Anselmann *et al.*, Solar neutrinos observed by GALLEX at Gran Sasso, Phys. Lett. **B285**, 376 (1992).

[35] Kamiokande Collaboration: K. S. Hirata *et al.*, Experimental study of the atmospheric neutrino flux, Phys. Lett. **B205**, 416 (1988).

[36] IMB Collaboration: R. Becker-Szendy *et al.*, Electron- and muon-neutrino content of the atmospheric flux, Phys. Rev. D **46**, 3720 (1992).

[37] W. A. Mann, Atmospheric neutrinos and the oscillations bonanza, Int. J. Mod. Phys. **A15S1**, 229 (2000).

[38] NUSEX Collaboration: M. Aglietta *et al.*, Experimental Study of Atmospheric Neutrino Flux in the NUSEX Experiment, Europhys. Lett. **8**, 611 (1989).

[39] Fréjus Collaboration: K. Daum *et al.*, Determination of the atmospheric neutrino spectra with the Fréjus detector, Z. Phys. **C66**, 417 (1995).

[40] Soudan2 Collaboration: W. Allison *et al.*, Measurement of the atmospheric neutrino flavour composition in Soudan 2, Phys. Lett. **B391**, 491 (1997).

[41] B. Pontecorvo, Inverse beta processes and nonconservation of lepton charge, Sov. Phys. JETP **7**, 172 (1958).

[42] Z. Maki, M. Nakagawa, and S. Sakata, Remarks on the unified model of elementary particles, Prog. Theor. Phys. **28**, 870 (1962).

[43] J.-M. Levy, On ultra-relativistic approximations, unobservable phases and other hand-waving in the derivation of the neutrino oscillation length, hep-ph/0901.0408.

[44] E. K. Akhmedov and A. Y. Smirnov, Paradoxes of neutrino oscillations, Phys. Atom. Nucl. **72**, 1363 (2009).

[45] C. Giunti and C. W. Kim, *Fundamentals of Neutrino Physics and Astrophysics* (Oxford University Press, 2007), First edition.

[46] K. Nakamura *et al.*, Review of Particle Physics, J. Phys. G **37**, 075021 (2010).

[47] S. J. Parke, CP violation in the Neutrino Sector, hep-ph/0807.3311 (2008).

[48] L. Wolfenstein, Neutrino oscillations in matter, Phys. Rev. **D17**, 2369 (1978).

[49] S. P. Mikheev and A. Y. Smirnov, Resonance enhancement of oscillations in matter and solar neutrino spectroscopy, Sov. J. Nucl. Phys. **42**, 913 (1985).

[50] V. Barger, J. G. Learned, S. Pakvasa, and T. J. Weiler, Neutrino Decay as an Explanation of Atmospheric Neutrino Observations, Phys. Rev. Lett. **82**, 2640 (1999).

[51] G. L. Fogli, E. Lisi, A. Marrone, and D. Montanino, Status of Atmospheric Neutrino $\nu_\mu \to \nu_\tau$ Oscillations and Decoherence after the first K2K Spectral Data, Phys. Rev. **D67**, 093006 (2003).

[52] Super-Kamiokande Collaboration: Y. Fukuda *et al.*, Evidence for oscillation of atmospheric neutrinos, Phys. Rev. Lett. **81**, 1562 (1998).

[53] T. Kajita, for the Super-Kamiokande and Kamiokande Collaborations, Atmospheric neutrino results from Super-Kamiokande and Kamiokande: Evidence for ν_μ oscillations, Nucl. Phys. Proc. Suppl. **77**, 123 (1999).

[54] Super-Kamiokande Collaboration: Y. Ashie *et al.*, A measurement of atmospheric neutrino oscillation parameters by Super-Kamiokande I, Phys. Rev. **D71**, 112005 (2005).

[55] Super-Kamiokande Collaboration: Y. Ashie *et al.*, Evidence for an oscillatory signature in atmospheric neutrino oscillations, Phys. Rev. Lett. **93**, 101801 (2004).

[56] 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).

[57] MACRO Collaboration: S. Ahlen *et al.*, Atmospheric neutrino flux measurement using upgoing muons, Phys. Lett. **B357**, 481 (1995).

[58] MINOS Collaboration: P. Adamson *et al.*, Charge-separated atmospheric neutrino-induced muons in the MINOS far detector, Phys. Rev. **D75**, 092003 (2007).

[59] A. Blake, Observations of Atmospheric Neutrinos and Antineutrinos in the MINOS Detector (paper draft), MINOS-doc-6926v2 (2010).

[60] K2K Collaboration: M. H. Ahn *et al.*, Measurement of neutrino oscillation by the K2K experiment, Phys. Rev. **D74**, 072003 (2006).

[61] OPERA Collaboration: N. Agafonova *et al.*, Observation of a first ν_τ candidate in the OPERA experiment in the CNGS beam, Phys. Lett. **B691**, 138 (2010).

[62] SNO Collaboration: B. Aharmim *et al.*, Electron Energy Spectra, Fluxes, and Day-Night Asymmetries of ^{8}B Solar Neutrinos from the 391-Day Salt Phase SNO Data Set, Phys. Rev. **C72**, 055502 (2005).

[63] KamLAND Collaboration: S. Abe *et al.*, Precision Measurement of Neutrino Oscillation Parameters with KamLAND, Phys. Rev. Lett. **100**, 221803 (2008).

[64] KamLAND Collaboration: K. Eguchi *et al.*, First results from KamLAND: Evidence for reactor anti-neutrino disappearance, Phys. Rev. Lett. **90**, 021802 (2003).

[65] 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).

[66] F. Boehm *et al.*, Final results from the Palo Verde neutrino oscillation experiment, Phys. Rev. **D64**, 112001 (2001).

[67] Double Chooz Collaboration: F. Ardellier *et al.*, Double Chooz: A search for the neutrino mixing angle θ_{13}, hep-ex/0606025 (2006).

[68] MINOS Collaboration: P. Adamson *et al.*, New constraints on muon-neutrino to electron-neutrino transitions in MINOS, Phys. Rev. **D82**, 051102 (2010).

[69] MINOS Collaboration: P. Adamson *et al.*, Measurement of Neutrino Oscillations with the MINOS Detectors in the NuMI Beam, Phys. Rev. Lett. **101**, 131802 (2008).

[70] P. F. Harrison, D. H. Perkins, and W. G. Scott, Tri-bimaximal mixing and the neutrino oscillation data, Phys. Lett. **B530**, 167 (2002).

[71] G. L. Fogli, E. Lisi, A. Marrone, A. Palazzo, and A. M. Rotunno, Hints of $\theta_{13} > 0$ from Global Neutrino Data Analysis, Phys. Rev. Lett. **101**, 141801 (2008).

[72] M. C. Gonzalez-Garcia, M. Maltoni, and J. Salvado, Updated global fit to three neutrino mixing: status of the hints of $\theta_{13} > 0$, JHEP **04**, 056 (2010).

[73] C. Athanassopoulos *et al.*, The liquid scintillator neutrino detector and LAMPF neutrino source, Nucl. Instrum. Meth. **A388**, 149 (1997).

[74] LSND Collaboration: A. Aguilar *et al.*, Evidence for neutrino oscillations from the observation of $\overline{\nu}_e$ appearance in a $\overline{\nu}_\mu$ beam, Phys. Rev. **D64**, 112007 (2001).

[75] Y. Declais *et al.*, Search for neutrino oscillations at 15-meters, 40-meters, and 95-meters from a nuclear power reactor at Bugey, Nucl. Phys. **B434**, 503 (1995).

[76] KARMEN Collaboration: B. Armbruster *et al.*, Upper limits for neutrino oscillations muon-antineutrino to electron-antineutrino from muon decay at rest., Phys. Rev. **D65**, 112001 (2002).

[77] MiniBooNE Collaboration: A. A. Aguilar-Arevalo *et al.*, A Search for electron neutrino appearance at the $\Delta m^2 \sim 1$ eV^2 scale, Phys. Rev. Lett. **98**, 231801 (2007).

[78] MiniBooNE Collaboration: A. A. Aguilar-Arevalo *et al.*, Event Excess in the MiniBooNE Search for $\bar{\nu}_\mu \to \bar{\nu}_e$ Oscillations, Phys. Rev. Lett. **105**, 181801 (2010).

[79] MiniBooNE Collaboration: A. A. Aguilar-Arevalo *et al.*, Unexplained Excess of Electronlike Events from a 1-GeV Neutrino Beam, Phys. Rev. Lett. **102**, 101802 (2009).

[80] C. Kraus *et al.*, Final Results from phase II of the Mainz Neutrino Mass Search in Tritium β Decay, Eur. Phys. J. **C40**, 447 (2005).

[81] B. M. Lobashev *et al.*, Direct search for the mass of neutrino and anomaly in the tritium beta-spectrum, Phys. Lett. **B460**, 227 (1999).

[82] H. V. Klapdor-Kleingrothaus *et al.*, Latest Results from the Heidelberg-Moscow Double Beta Decay Experiment, Eur. Phys. J. **A12**, 147 (2001).

[83] H. V. Klapdor-Kleingrothaus, I. V. Krivosheina, A. Dietz, and O. Chkvorets, Search for neutrinoless double beta decay with enriched ^{76}Ge in Gran Sasso 1990-2003, Phys. Lett. **B586**, 198 (2004).

[84] F. Feruglio, A. Strumia, and F. Vissani, Neutrino oscillations and signals in β and $0\nu2\beta$ experiments, Nucl. Phys. **B637**, 345 (2002).

[85] C. E. Aalseth *et al.*, Comment on “Evidence for Neutrinoless Double Beta Decay”, Mod. Phys. Lett. **A17**, 1475 (2002).

[86] A. M. Bakalyarov, A. Y. Balysh, S. T. Belyaev, V. I. Lebedev, and S. V. Zhukov, Results of the experiment on investigation of Germanium-76 double beta decay, Phys. Part. Nucl. Lett. **2**, 77 (2005).

[87] K. S. Hirata *et al.*, Observation in the Kamiokande-II detector of the neutrino burst from supernova SN1987A, Phys. Rev. **D38**, 448 (1988).

[88] IMB Collaboration: J. C. Van Der Velde *et al.*, Neutrinos from SN1987A in the IMB detector, Nucl. Instrum. Meth. **A264**, 28 (1988).

[89] A. E. Chudakov, Y. S. Elensky, and S. P. Mikheev, Characteristics of the neutrino emission from supernova SN1987A, JETP Lett. **46**, 373 (1987).

[90] T. J. Loredo and D. Q. Lamb, Bayesian analysis of neutrinos observed from supernova SN 1987A, Phys. Rev. **D65**, 063002 (2002).

[91] S. S. Gershtein and Y. B. Zeldovich, Rest mass of muonic neutrino and cosmology, JETP Lett. **4**, 120 (1966).

[92] R. Cowsik and J. McClelland, An Upper Limit on the Neutrino Rest Mass, Phys. Rev. Lett. **29**, 669 (1972).

[93] WMAP Collaboration: D. N. Spergel *et al.*, First Year Wilkinson Microwave Anisotropy Probe (WMAP) Observations: Determination of Cosmological Parameters, Astrophys. J. Suppl. **148**, 175 (2003).

[94] SDSS Collaboration: M. Tegmark *et al.*, Cosmological parameters from SDSS and WMAP, Phys. Rev. **D69**, 103501 (2004).

[95] SDSS Collaboration: U. Seljak *et al.*, Cosmological parameter analysis including SDSS Ly-alpha forest and galaxy bias: Constraints on the primordial spectrum of fluctuations, neutrino mass, and dark energy, Phys. Rev. **D71**, 103515 (2005).

[96] J. Hylen *et al.*, NuMI Technical Design Handbook, Internal NuMI report (2003).

[97] S. E. Kopp, The NuMI neutrino beam at Fermilab, physics/0508001 (2005).

[98] MINOS Collaboration: D. G. Michael *et al.*, The Magnetized Steel and Scintillator Calorimeters of the MINOS Experiment, Nucl. Instrum. Meth. **A569**, 190 (2008).

[99] MINOS Collaboration: P. Adamson *et al.*, Search for sterile neutrino mixing in the MINOS long- baseline experiment, Phys. Rev. **D81**, 052004 (2010).

[100] MINOS Collaboration: P. Adamson *et al.*, Measurement of the atmospheric muon charge ratio at TeV energies with MINOS, Phys. Rev. **D76**, 052003 (2007).

[101] MINOS Collaboration: P. Adamson *et al.*, Observation of muon intensity variations by season with the MINOS far detector, Phys. Rev. **D81**, 012001 (2010).

[102] MINOS Collaboration: S. Osprey *et al.*, Sudden stratospheric warmings seen in MINOS deep underground muon data, Geophys. Res. Lett. **36**, L05809 (2009).

[103] P. Adamson *et al.*, Neutrino and antineutrino inclusive charged-current cross section measurements with the MINOS near detector, Phys. Rev. **D81**, 072002 (2010).

[104] MINOS Collaboration: P. Adamson *et al.*, Testing Lorentz Invariance and CPT Conservation with NuMI Neutrinos in the MINOS Near Detector, Phys. Rev. Lett. **101**, 151601 (2008).

[105] MINOS Collaboration: P. Adamson *et al.*, Search for Lorentz Invariance and CPT Violation with the MINOS Far Detector, Phys. Rev. Lett. **105**, 151601 (2010).

[106] S. Shekhar *et al.*, Slip Stacking in the Fermilab Main Injector, DPF/DPB Summer Study: New Direction for High Energy Physics, Snowmass, Colorado (1996).

[107] A. Cabrera *et al.*, Comparisons of the MINOS near and far detector readout systems at a test beam, Nucl. Instrum. Meth. **A609**, 106 (2009).

[108] P. Adamson *et al.*, Calibration of the MINOS Near and Far Detector Readout Systems, To be submitted to Nucl. Instrum. Meth.

[109] N. Tagg *et al.*, Performance of Hamamatsu 64-anode photomultipliers for use with wavelength-shifting optical fibres, Nucl. Instrum. Meth. **A539**, 668 (2005).

[110] K. Lang *et al.*, Characterization of 1600 Hamamatsu 16-anode photomultipliers for the MINOS Far detector, Nucl. Instrum. Meth. **A545**, 852 (2005).

[111] B. J. Rebel, *Neutrino induced muons in the MINOS far detector*, PhD thesis, Indiana University, 2004.

[112] S. Avvakumov *et al.*, Spontaneous light emission from fibers in MINOS, Nucl. Instrum. Meth. **A545**, 145 (2005).

[113] 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).

[114] S. Agostinelli *et al.*, GEANT4 – a simulation toolkit, Nucl. Instrum. and Meth. **A506**, 250 (2003).

[115] H. Gallagher, The NEUGEN neutrino event generator, Nucl. Phys. Proc. Suppl. **112**, 188 (2002).

[116] GEANT—Detector description and simulation tool, CERN Program Library, Long Writeup, W5013, `http://wwwasd.web.cern.ch/wwwasd/geant/index.html`.

[117] C. Zeitnitz and T. A. Gabriel, The GEANT - CALOR interface and benchmark calculations of ZEUS test calorimeters, Nucl. Instrum. Meth. **A349**, 106 (1994).

[118] J. Coelho and G. Pawloski, Incorporating Intensity Profile in Monte Carlo - Note, MINOS-doc-6032 (2009).

[119] J. S. Marshall, *A Study of Muon Neutrino Disappearance with the MINOS Detectors and the NuMI Neutrino Beam*, DPhil thesis, University of Cambridge, 2008.

[120] G. Welch and G. Bishop, An Introduction to the Kalman Filter, Technical Report TR 95-041 (2006).

[121] C. Smith, SS Event Builder Algorithm, MINOS-doc-2821 (2007).

[122] M. Dorman, Beam Fit Position Paper, MINOS-doc-7146 (2010).

[123] NA49 Collaboration: C. Alt *et al.*, Inclusive production of charged pions in p + C collisions at 158-GeV/c beam momentum, Eur. Phys. J. **C49**, 897 (2007).

[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] B. List, Why and When to Optimize Efficiency Times Purity, ETH Zurich internal note 2002.

[126] J. Ratchford, A k-Nearest Neighbor based Particle Identification for the 2010 Charged Current Analysis, MINOS-doc-7116 (2010).

[127] J. Ratchford, PhD thesis, University of Texas at Austin, 2011.

[128] J. J. Evans, *Measuring antineutrino oscillations with the MINOS experiment*, DPhil thesis, University of Oxford, 2008.

[129] J. Mitchell, DPhil thesis, University of Cambridge, 2011.

[130] S. Coleman, Extrapolating with antineutrinos and resolution bins, MINOS-doc-6874 (2010).

[131] S. J. Coleman, *A measurement of neutrino oscillations with muon neutrinos in the MINOS experiment*, PhD thesis, College of William and Mary, 2011.

[132] M. Strait, *A Measurement of Oscillation Parameters using Antifiducial Charged Current Events in MINOS*, PhD thesis, University of Minnesota, 2010.

[133] J. Evans, Systematic uncertainties for the 2010 CC analysis, MINOS-doc-7145 (2010).

[134] R. Hatcher *et al.*, Range Curvature Task Force Position Paper, MINOS-doc-3134 (2007).

[135] P. A. Rodrigues, *A sterile neutrino search for the MINOS experiment*, DPhil thesis, University of Oxford, 2010.

[136] J. Ratchford, Position Paper: MRCC motivated NC background, MINOS-doc-7115 (2010).

[137] F. James and M. Roos, MINUIT - A System for function minimization and analysis of the parameter errors and correlations, Comput. Phys. Commun. **10**, 343 (1975).

[138] L. Loiacono, kNN Shower Energy Estimation, MINOS-doc-6340 (2009).

[139] `http://en.wikipedia.org/wiki/kd-tree`.

[140] M. Ishitsuka, Shower energy tuning for CedarPhyV03/DaikonV04, MINOS-doc-5131 (2008).

[141] X. Huang, CC shower energy tuning update, MINOS-doc-6170 (2009).

[142] M. Kordosky, H. Gallagher, and S. Dytman, Estimated Uncertainty on the Shower Energy Scale in ν_μ-CC Events, MINOS-doc-3362 (2007).

[143] M. Kordosky and H. Gallagher, Shower Energy Scale Uncertainty For the Run I+II CC Analysis, MINOS-doc-4287 (2008).

[144] J. Boehm, H. Gallagher, and T. Yang, Hadronization Model Uncertainties for the ν_e Analysis, MINOS-doc-5392 (2008).

[145] D. Rein and L. M. Sehgal, Neutrino Excitation of Baryon Resonances and Single Pion Production, Annals Phys. **133**, 79 (1981).

[146] T. Sjostrand, S. Mrenna, and P. Z. Skands, PYTHIA 6.4 Physics and Manual, JHEP **05**, 026 (2006).

[147] T. Yang *et al.*, A Hadronization Model for Few-GeV Neutrino Interactions, Eur. Phys. J. **C63**, 1 (2009).

[148] T. Yang, *A Study of Muon Nuetrino to Electron Neutrino Oscillations in the MINOS Experiment*, PhD thesis, Stanford University, 2009.

[149] S. Dytman and H. Gallagher, INTRANUKE Advancements for NEUGEN3 hA model, MINOS-doc-3484 (2007).

[150] M. Kordosky, A procedure to re-weight events to account for uncertainties in final state interactions, MINOS-doc-3449 (2007).

[151] M. Kordosky, H. Gallagher, and S. Dytman, Estimated Uncertainty on the Shower Energy Scale in ν_μ-CC Events., MINOS-doc-3362 (2007).

[152] D. S. Baranov *et al.*, Evaluation Of Hadron Formation Length In Neutrino Induced Interactions. (In Russian), Yad. Fiz. **40**, 1454 (1984).

[153] M. Kordosky, Improvements to neugen's formation time treatment, MINOS-doc-3539 (2007).

[154] R. D. Ransome, Pion absorption and re-scattering, Nucl. Phys. Proc. Suppl. **139**, 208 (2005).

[155] H. Gallagher *et al.*, Comparing Intranuclear Rescattering Models to ν_μ–D and ν_μ–Ne Data for An Atmospheric Neutrino Flux, MINOS-doc-4112 (2008).

[156] S. Dytman and H. Gallagher, Changes to the Determination of the Mean Free Path of Hadrons in Nuclear Matter in INTRANUKE, MINOS-doc-4358 (2008).

[157] M. A. Kordosky, *Hadronic interactions in the MINOS detectors*, PhD thesis, University of Texas at Austin, 2004.

[158] The MINOS Calibration Group, Calibration Position Paper, MINOS-doc-3137 (2007).

[159] N. Tagg, Gain Calibration Systematics, MINOS-doc-4865 (2008).

[160] J. Evans, Absolute (calorimetric) shower energy uncertainty, MINOS-doc-7173 (2010).

[161] J. Coelho, PhD thesis, Universidade Estadual de Campinas, 2011.

[162] MINOS Collaboration: P. Adamson *et al.*, Measurement of the neutrino mass splitting and flavor mixing by MINOS, Phys. Rev. Lett. **106**, 181801 (2011).

[163] J. Hartnell, NuMI Muon Anti-Neutrinos in MINOS, Fermilab Joint Experimental-Theoretical Seminar (2009).

[164] G. J. Feldman and R. D. Cousins, A Unified Approach to the Classical Statistical Analysis of Small Signals, Phys. Rev. **D57**, 3873 (1998).

[165] M. C. Gonzalez-Garcia and M. Maltoni, Phenomenology with Massive Neutrinos, Phys. Rept. **460**, 1 (2008).

[166] P. Adamson *et al.*, First direct observation of muon antineutrino disappearance, Submitted to Phys. Rev. Lett. (2011).

[167] W. A. Mann, D. Cherdack, W. Musial, and T. Kafka, Apparent Multiple Δm^2_{32} in $\bar{\nu}_\mu$ and ν_μ Survival Oscillations from Non-Standard Interaction Matter Effect, Phys. Rev. **D82**, 113010 (2010).

[168] J. Kopp, P. A. N. Machado, and S. J. Parke, Interpretation of MINOS data in terms of non-standard neutrino interactions, Phys. Rev. **D82**, 113002 (2010).

[169] Z. Isvan, PhD thesis, University of Pittsburgh, 2011.

[170] MINERνA Collaboration: D. Drakoulakos *et al.*, Proposal to perform a high-statistics neutrino scattering experiment using a fine-grained detector in the NuMI beam, (2004), hep-ex/0405002.

[171] 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).

[172] 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).

[173] Y. Itow *et al.*, The JHF-Kamioka neutrino project, (2001), hep-ex/0106019.

[174] NOνA Collaboration: D. S. Ayres *et al.*, NOνA proposal to build a 30-kiloton off-axis detector to study neutrino oscillations in the Fermilab NuMI beamline, (2004), hep-ex/0503053.