Probing Future DUNE Sensitivity to Oscillation Parameters with Atmospheric Neutrinos

Master's Thesis

Department of Physics and Astronomy

Stony Brook University

Ian Segal-Gould

September 12, 2026

Advisor: Prof. Ciro Riccio

Abstract

Neutrino oscillations are one of the few confirmed signatures of physics beyond the Standard Model, which predicts neutrinos to be massless. The Deep Underground Neutrino Experiment (DUNE) is an international experiment designed to resolve open questions in neutrino physics, including the CP-violating phase, the neutrino mass ordering, and the octant of the atmospheric mixing angle. While DUNE's primary science program uses a long baseline neutrino beam, its large underground liquid argon far detector modules are also sensitive to atmospheric neutrinos, which provide a complementary probe of oscillation physics through their broad range of baselines and energies, as well as their sensitivity to matter effects as they traverse the Earth's interior.

This thesis presents a sensitivity study of DUNE's atmospheric neutrino oscillation program using the GUNDAM fitting framework, within a semi-frequentist statistical approach. An optimization of the oscillation probability grid is performed, establishing that a 300×300 grid in true neutrino energy and true cosine of the zenith angle is sufficient for reliable oscillation weight extraction.

Sensitivity results are presented through one dimensional profile likelihood scans, two dimensional profiled contour scans, and a mass ordering determination study, all constructed from simulated datasets at a 400 kt·yr exposure corresponding to ten years of data collection with the full four module DUNE far detector. The one dimensional scans show that Δm^2_{32} is tightly constrained beyond the external NuFit 5.2 prior, $\sin^2\theta_{23}$ exhibits an asymmetric profile reflecting partial octant resolution, and δ_{CP} is broadly unconstrained by the atmospheric sample alone. A mass ordering sensitivity study finds that DUNE atmospheric neutrinos can reach 3 to 4 sigma significance depending on the true ordering and external constraints, providing an independent complement to the beam based mass ordering determination.

Acknowledgements

First and foremost I am thankful to my mentors throughout my time as a master's student at Stony Brook University. In particular, I am thankful for the support and guidance of my advisor Prof. Ciro Riccio, and the post-doc I've had the honor to work with and learn from: Dr. Uladzislava Yevarouskaya. Thank you to the following professors, for their mentorship and patience with teaching me physics: Professor M.C. Gonzalez-Garcia, Professor Clark McGrew, Professor Paul Goldbart, and Professor Katherine Hall from my undergraduate institution, who acted as both my academic advisor and instructor for numerous classes while at the University of Connecticut. I am thankful for the friends I have made during my time at Stony Brook University. While the number of people who have made a difference in my life is longer than I can share here, I at least want to acknowledge some of the main individuals who I've had the pleasure of calling my friends. Béla Arwen and Ian Johnson, thank you for listening to my thoughts on neutrinos and their place in the cosmos. I would also like to thank Carrie Cox, Gloria Tejedor García, Edoardo Buonocore, Melissa Rasmussen, and Jacob Smith. I would like to thank my family, notably my sister Arielle and my grandmother, Rochelle. While living on Long Island, I've been able to see them regularly and they've reminded me of where I come from, supporting me and assuring me I would realize whatever I set my mind to. I would also like to thank my mother Laura, and my brother Noah, who have both heard me out when I needed to vent. I wish I could mention every single person who has made a difference in my life these past two years, but unfortunately I cannot. So if I didn't mention you, I still appreciate our interactions. I would also like to thank a very special being in my life, my cat Jackson. Although Jackson is not mathematically inclined or able to read, he provided invaluable companionship throughout my research and thesis writing process.

Contents

Acknowledgements **ii**

1 Neutrino Physics **3**
1.1 Introduction to Neutrinos 3
1.1.1 What Neutrinos Are 3
1.1.2 History of Neutrino Discovery 4
1.1.3 The Solar Neutrino Problem 6
1.1.4 The Atmospheric Neutrino Problem 6
1.1.5 Discovery of Neutrino Oscillations 7
1.2 Neutrino Oscillations 8
1.2.1 Vacuum Oscillations 8
1.2.2 Matter Effects 10
1.2.3 CP Violation 12
1.3 Open Questions in Neutrino Physics 12
1.4 Atmospheric Neutrinos 13
1.4.1 Cosmic Ray Primaries 14
1.4.2 Hadronic Interactions and Meson Production 15
1.4.3 Decay Chains and the Flavor Ratio 15
1.4.4 Energy Dependence of the Flavor Ratio 15
1.4.5 Production Height and Zenith Angle Dependence 16
1.4.6 Atmospheric Neutrino Flux Modeling 19
1.4.7 Oscillation Signatures in Atmospheric Neutrinos 19

2 The DUNE Experiment **21**
2.1 Physics Goals 22
2.1.1 CP Violation 22
2.1.2 Mass Ordering Determination 22
2.1.3 Precision Oscillation Measurements 22
2.1.4 BSM Searches, Supernova Neutrinos, and Nucleon Decay 22
2.2 LBNF Beam 23
2.2.1 Fermilab Accelerator Complex 23
2.2.2 Beam Production and Energy Spectrum 23
2.2.3 The 1300 km Baseline 23
2.3 Far Detector 25
2.3.1 LArTPC Operating Principle 25
2.4 Near Detector Complex 26
2.4.1 ND-LAr 26
2.4.2 TMS 27

2.4.3 SAND . . . 28
2.5 Event Reconstruction . . . 28
2.5.1 Particle Identification . . . 28
2.5.2 Interaction Classification . . . 29
2.6 Atmospheric Neutrinos at DUNE Far Detector . . . 30

3 The GUNDAM Framework and Analysis 31
3.1 Overview . . . 31
3.2 Statistical Framework . . . 32
3.2.1 Systematic Uncertainties . . . 32
3.2.2 The Extended Binned Likelihood . . . 32
3.2.3 Parameter Estimation and Profiling . . . 33
3.2.4 The Asimov Fit . . . 34
3.3 The GUNDAM Framework . . . 35
3.3.1 Overview and Design . . . 36
3.3.2 Oscillation Probability Calculation . . . 38
3.3.3 Minimization and Fitting . . . 38
3.4 Analysis Configuration . . . 38
3.4.1 Event Selection and Analysis Binning . . . 40
3.4.2 Oscillation Parameters Configuration in GUNDAM . . . 40
3.5 Production Height Studies . . . 41
3.5.1 Energy Dependence . . . 42
3.5.2 Zenith Dependence . . . 43
3.5.3 Summary . . . 47
3.6 Asimov Fit . . . 47

4 Oscillation Probability Calculation Optimization 49
4.1 Oscillation Grid Configuration . . . 49
4.1.1 Interpolation Error Tolerance . . . 49
4.1.2 Three-Flavor Period Counting . . . 51
4.1.3 Grid Construction Methods . . . 52
4.2 Grid Resolution Studies . . . 53
4.2.1 Impact of Grid Resolution on Event Rate . . . 54
4.2.2 Impact of Grid Resolution on Oscillation Parameters . . . 54
4.2.3 Quantitative Summary . . . 56
4.2.4 Conclusion . . . 58

5 DUNE Sensitivity to Atmospheric Neutrino Oscillations 60
5.1 One-Dimensional Profile Likelihood Scans . . . 60
5.1.1 Prior-Dominated Parameters . . . 61
5.1.2 Atmospheric Neutrino Oscillation Parameters . . . 62
5.1.3 Summary . . . 63
5.2 Two-Dimensional Contour Scans . . . 64
5.2.1 $\sin^2\theta_{23}$ vs Δm^2_{32} . . . 64
5.2.2 Contours involving δ_{CP} . . . 64
5.3 Mass Ordering Sensitivity . . . 68
5.3.1 Method . . . 68
5.3.2 Results . . . 68
5.3.3 MSW Resonance and Octant Competition . . . 69
5.3.4 Comparison . . . 70

5.3.5 Summary 70

Bibliography **80**

List of Tables

1.1 Principal decay modes contributing to atmospheric neutrino production. Branching ratios from [53]. 15
1.2 Characteristic baselines, energies, and resulting L/E coverage of atmospheric neutrinos. 20

3.1 Oscillation parameters injected in the simulated datasets used in this analysis, taken from the NuFit 5.2 global fit with and without Super-Kamiokande atmospheric data [34, 35], for both normal (NO) and inverted (IO) ordering. The with SK configuration is the baseline used in Sections 5.1 and 5.2; both configurations are used in the mass ordering study of Section 5.3. For IO, $\Delta m^2_{3\ell} \equiv \Delta m^2_{32}$ following the NuFit convention. 41
3.2 Baseline prior constraints on oscillation parameters under both NuFit 5.2 configurations used in this analysis. Central values follow the corresponding NuFit 5.2 fit [35]; Gaussian 1σ widths are from the covariance matrix of each fit. The with SK configuration is used for the profile likelihood scans and two dimensional contour scans in Chapter 5; the mass ordering sensitivity study in Section 5.3 is performed under both configurations. 41

4.1 Validation of the interpolation error bound: number of equally-spaced points per period required to achieve a given maximum error tolerance on $\sin^2(x)$. 51
4.2 Oscillation probability grid configurations tested in this study. All grids span $E \in [0.1, 100]$ GeV and $\cos\theta_z \in [-1, 0]$. 53

5.1 Grid ranges for the three two dimensional profiled likelihood scans. Each scan uses a 50×50 uniform grid (2500 points). . . . 64

List of Figures

1.1 Standard Model 3

1.2 The continuous energy spectrum of electrons emitted in β-decay. Image credit: H. Paul, CC BY-SA 4.0 [56]. 4

1.3 Zenith angle distributions of μ-like and e-like events for sub-GeV and multi-GeV data sets measured by Super-Kamiokande. The hatched region shows the no oscillation expectation; the bold line shows the best fit $\nu_\mu \leftrightarrow \nu_\tau$ oscillation. The μ-like deficit for upward-going events ($\cos\Theta < 0$) provided definitive evidence for atmospheric neutrino oscillations. Figure from [38]. 7

1.4 Schematic of atmospheric neutrino production. A primary cosmic ray interacts in the upper atmosphere, producing pions. Charged pions decay to muons and ν_μ, and the muons subsequently decay to electrons, ν_e, and ν_μ. Figure from Ref. [48]. 13

1.5 The cosmic ray energy spectrum, showing flux multiplied by E^2 versus energy. Individual components (H, He, O, Fe) are shown at lower energies, with the all-particle spectrum extending to the highest energies. The knee ($\sim 10^6$ GeV), second knee, and ankle features are labeled. The red box indicates the approximate primary cosmic ray energy range responsible for producing the atmospheric neutrino flux relevant to DUNE far detector reconstruction studies, which covers neutrino energies from 100 MeV to 10 GeV [32]. Figure from the Cosmic Ray Database (CRDB) [10]. 14

1.6 Calculated $(\nu_\mu + \bar{\nu}_\mu)/(\nu_e + \bar{\nu}_e)$ ratio of the atmospheric neutrino flux as a function of neutrino energy, using three different flux model calculations. Figure from [48]. 17

1.7 Geometry of atmospheric neutrino production. A neutrino produced at height h above the Earth's surface travels a baseline L to reach the underground detector, depending on the zenith angle θ. 17

1.8 Geometry of an atmospheric neutrino trajectory through the Earth. The zenith angle at production θ_{in} and at detection θ_{out} satisfies $\theta_{\text{in}} = \pi - \theta_{\text{out}}$, reflecting the up-down symmetry of atmospheric neutrino production. Figure from [48]. 18

1.9 Ratio of observed to expected atmospheric ν_μ events as a function of L/E. The dip around $L/E \sim 500$ km/GeV corresponds to the first oscillation minimum. Figure from [14]. 20

2.1 Schematic of the LBNF beamline and near detector complex at Fermilab. Figure from [4]. 23

2.2 The LBNF neutrino flux as a function of energy at the far detector location. Figure from [6]. 24

2.3 The $\nu_\mu \to \nu_e$ appearance probability at a baseline of 1300 km as a function of neutrino energy for neutrinos (left) and antineutrinos (right), for several values of δ_{CP} assuming normal ordering. The first oscillation maximum near 2.5 GeV coincides with the LBNF beam peak energy. Figure from [28]. 25

2.4 Schematic of the DUNE near detector hall in the on-axis configuration (left) and with ND-LAr and ND-GAr in an off-axis configuration (right), illustrating the DUNE-PRISM concept. SAND remains on the beam axis in both configurations. The muon spectrometer shown here is ND-GAr, which will replace TMS in Phase II. Figure from [29]. 26

2.5 Left: the full ND-LAr detector, composed of a 7×5 array of individual modules. Right: a single $1 \times 1 \times 3$ m^3 ND-LAr module, divided into two optically isolated LArTPCs by a central cathode. Ionization electrons drift toward pixelated charge readout anode planes on opposite faces, while scintillation light is collected by light readout modules on the remaining two faces. Figure from [29]. 27

2.6 Examples of track, shower, and Michel electron topologies in a simulated ProtoDUNE-SP event, shown as CNN (Convolutional Neural Network) input patches. The track topology shows a long, straight energy deposition characteristic of a muon, while the shower topology shows the branching cascade structure of an electromagnetic shower. Figure from [30]. 29

3.1 Baseline consistency check: one dimensional profile likelihood scans for all six oscillation parameters using the 300×300 oscillation probability grid at 400 kt·yr exposure. At each scan point the parameter of interest is held fixed at the scan value and the remaining five oscillation parameters are profiled over by Minuit2. The pink band shows the NuFit 5.2 1σ range and the red triangle marks the NuFit 5.2 best fit. Horizontal dashed lines at $\Delta\chi^2 = 1$, 4, and 9 correspond to the 1σ, 2σ, and 3σ intervals by Wilks' theorem. All six parameters recover clean curves centered on truth, with the atmospheric sector parameters ($\sin^2\theta_{23}$, Δm^2_{32}, δ_{CP}) showing local sensitivity from the atmospheric neutrino sample and the solar sector parameters ($\sin^2\theta_{12}$, $\sin^2\theta_{13}$, Δm^2_{21}) dominated by their external priors. The y-axis is labeled χ^2 but denotes $\Delta\chi^2$ relative to the scan minimum. 34

3.2 Software architecture of the GUNDAM framework, showing the six core libraries and their principal class hierarchies. The `FitterEngine` delegates to a `MinimizerBase` backend and drives the `Propagator` for spectrum prediction. The `StatisticalInference` library provides the likelihood evaluation through `LikelihoodInterface`. Event data is managed by the `SamplesManager`, while oscillation and systematic parameters are organized by the `ParametersManager` and applied to events through the `DialDictionary` reweighting engine. Class names reflect the current GUNDAM source code. . 37

3.3 Analysis workflow used in this thesis. Configuration files in the `gudi-anoa-sen25a` repository specify the binned Monte Carlo inputs for the ν_e- and ν_μ-selected samples together with the oscillation and systematic parameter priors. Oscillation probabilities are precomputed by TabulatedNuOscillator on a two dimensional grid in true neutrino energy and $\cos\theta_z$. A single `gundamFitter` call reweights the Monte Carlo events by the flux, interaction cross section, and oscillation probability, compares the resulting binned predictions to the simulated dataset, and returns the best fit parameter values. The simulated dataset is itself produced by running the reweighting pipeline once at the injected truth values. Profile likelihood scans are produced by external shell scripts that rerun the fit many times, each time holding the parameter of interest fixed at a different value; the collected results form the $\Delta\chi^2$ profiles used in the sensitivity studies. . . . 39

3.4 Reconstructed energy spectra for the ν_μ (left) and ν_e (right) selections comparing fixed production height and height averaged predictions. The lower panels show the relative difference. The ν_μ deficit is broad and systematic, peaking around 1.5% in the sub GeV region. The ν_e differences are negligible. 43

3.5 True neutrino energy spectra for the ν_μ (left) and ν_e (right) selections comparing fixed production height and height averaged predictions. The shape of the ν_μ deficit is similar to the reconstructed energy case, with the addition of a small positive excursion at the lowest energies. The ν_e differences remain negligible. 44

3.6 Event rates as a function of true $\cos\theta_z$ for the ν_μ (left) and ν_e (right) selections. The lower panels show the relative difference between height averaged and fixed height predictions. The ν_μ effect is largest near the horizon ($\cos\theta_z \approx 0$ to 0.3), where the production height variation is a significant fraction of the total path length. 45

3.7 Event rates as a function of reconstructed $\cos\theta_z$ for the ν_μ (left) and ν_e (right) selections. The zenith angle reconstruction smears the sharp near horizon feature seen in the true $\cos\theta_z$ distribution but the overall pattern is preserved. The ν_e differences remain negligible. 46

3.8 Postfit parameter uncertainties from the Asimov Hesse fit. For the solar sector parameters $\sin^2\theta_{12}$, $\sin^2\theta_{13}$, and Δm^2_{21}, which carry NuFIT 5.2 Gaussian priors, the postfit errors are shown normalized to the prior σ (red band); ratios close to unity indicate that the atmospheric sample does not improve on the external constraint. For the parameters of interest $\sin^2\theta_{23}$, Δm^2_{32}, and δ_{CP}, which enter the fit with flat priors, the Hesse errors reflect the local curvature of the likelihood at the Asimov minimum and are shown for diagnostic purposes only; the physical sensitivity to these parameters is given by the profile likelihood scans in Section 5.1. 48

4.1 Muon neutrino survival probability $P(\nu_\mu \to \nu_\mu)$ as a function of true neutrino energy and true zenith angle, computed with CUDAProb3 at 3000×1000 (E, θ_z) resolution. The arrows on the left indicate neutrino direction: $\cos\theta_z = +1$ corresponds to downward-going neutrinos (short atmospheric path), $\cos\theta_z = 0$ to horizontal trajectories, and $\cos\theta_z = -1$ to upward-going neutrinos that traverse the full Earth diameter. Rapid oscillations at low energies demand fine grid spacing to avoid interpolation artifacts. 50

4.2 Piecewise-linear interpolation of $\sin^2(x)$ sampled at 4 points per period (left) and 8 points per period (right). Green markers show the sample locations, the blue curve shows the true function, and the red shaded band shows the interpolation error. At 4 points per period the linear interpolant visibly misses the peaks and troughs, while at 8 points per period the error drops below the 8% tolerance criterion of Eq. 4.3. 51

4.3 Predicted event rate spectra for the ν_e-selected (left) and ν_μ-selected (right) samples, shown as a function of reconstructed energy (top) and reconstructed $\cos\theta_Z$ (bottom), for five oscillation grid resolutions. The ν_e selection uses calorimetric electron energy (`recoEnuEcalo`) and electron zenith angle (`recoEle_thetaZ`), while the ν_μ selection uses calorimetric lepton energy (`recoEnuLepCalo`) and muon zenith angle (`recoMu_thetaZ`). Lower panels show the fractional deviation of each grid relative to the 4500×300 reference. In reconstructed space, all grids agree at the sub-percent level across most of the kinematic range; residual differences appear only at the highest reconstructed energies where event counts are small. 55

4.4 Two-dimensional maps of the fractional event rate residual $\Delta N/N$ [%] relative to the 4500×300 reference grid, shown in bins of true neutrino energy and $\cos\theta_Z$ for the ν_e selection (top) and ν_μ selection (bottom). The leftmost column shows the absolute event distribution from the reference grid. Residuals are concentrated in the upward-going, multi-GeV region where oscillation structure is most rapid. 56

4.5 Profile likelihood scan curves ($\Delta\chi^2$ vs. parameter value) for all six oscillation parameters, overlaid for the five grid resolutions. At each scan point the parameter of interest is held fixed at the scan value and the remaining parameters are profiled over. Lower panels show the residual $\Delta(\Delta\chi^2)$ of each grid relative to the 4500×300 reference. The atmospheric parameters ($\sin^2\theta_{23}$, Δm^2_{32}, δ_{CP}) show genuine sensitivity from the data, while the solar parameters and $\sin^2\theta_{13}$ are prior-dominated. All grids produce indistinguishable curves on the scale of the full $\Delta\chi^2$ range. 57

4.6 Quantitative summary of grid resolution effects on the profile likelihood comparison, computed relative to the 4500×300 reference grid. Left: fractional shift in the width at $\Delta\chi^2 = 1$, Δw_G (Equation 4.4). Right: maximum absolute pointwise $\Delta\chi^2$ deviation R_G (Equation 4.5), plotted on a logarithmic vertical axis. The 150×300 grid (pink diamonds) exhibits width shifts of 3.1% for Δm^2_{32} and 1.6% for $\sin^2\theta_{23}$ and a pointwise deviation of $R_G = 57.4$ for $\sin^2\theta_{23}$, confirming it is insufficiently resolved for a reliable sensitivity analysis. All configurations at 300×300 and above show width shifts at or below 2.4% and are discussed in the text. Width shifts for δ_{CP} are not reported because the likelihood curve does not cross $\Delta\chi^2 = 1$ within the scan range. . 58

5.1 One dimensional profile likelihood scans for the three prior-dominated oscillation parameters: $\sin^2\theta_{12}$, $\sin^2\theta_{13}$, and Δm^2_{21}. At each scan point the parameter of interest is held fixed at the scan value and the remaining five oscillation parameters are profiled over by Minuit2. The pink band shows the NuFit 5.2 1σ range and the red triangle marks the NuFit 5.2 best fit. The horizontal dashed lines at $\Delta\chi^2 = 1$, 4, and 9 correspond to the 1σ, 2σ, and 3σ intervals by Wilks' theorem. All three parameters recover clean parabolas consistent with their input prior widths, indicating that the atmospheric neutrino sample does not add significant constraint beyond the external priors. 61

5.2 One dimensional profile likelihood scans for the three parameters to which DUNE atmospheric neutrinos are most sensitive: $\sin^2\theta_{23}$, Δm^2_{32}, and δ_{CP}. At each scan point the parameter of interest is held fixed at the scan value and the remaining five oscillation parameters are profiled over. The pink band shows the NuFit 5.2 1σ range and the red triangle marks the NuFit 5.2 best fit. The $\sin^2\theta_{23}$ curve exhibits an asymmetric structure reflecting residual octant sensitivity. The Δm^2_{32} curve is tightly peaked, with the $\Delta\chi^2 = 1$ interval narrower than the NuFit prior band. The δ_{CP} curve is highly asymmetric: $\Delta\chi^2$ rises sharply toward positive δ_{CP} and remains below unity across the scan range on the negative side, so no two sided 1σ interval can be extracted. The full allowed regions including correlations between these parameters are presented in Section 5.2. 63

5.3 Two dimensional profiled $\Delta\chi^2$ contour for $\sin^2\theta_{23}$ versus Δm^2_{32}. Contours are drawn at $\Delta\chi^2 = 2.30$, 6.18, and 11.83 (corresponding to 1σ, 2σ, 3σ for 2 DOF). The red star marks the profiled best fit; the green triangle marks the injected NuFit 5.2 truth. Side panels show the one dimensional profile likelihoods obtained by minimizing $\Delta\chi^2$ along the orthogonal axis. 65
5.4 Two-dimensional profiled $\Delta\chi^2$ contour for δ_{CP} versus $\sin^2\theta_{23}$. Contours are drawn at $\Delta\chi^2 = 2.30$, 6.18, and 11.83 (1σ, 2σ, 3σ for 2 DOF). The red star marks the profiled best fit; the green triangle marks the injected NuFit 5.2 truth. Side panels show the one-dimensional profile likelihoods. 66
5.5 Two dimensional profiled $\Delta\chi^2$ contour for δ_{CP} versus Δm^2_{32}. Contours are drawn at $\Delta\chi^2 = 2.30$, 6.18, and 11.83 (1σ, 2σ, 3σ for 2 DOF). The red star marks the profiled best fit; the green triangle marks the injected NuFit 5.2 truth. Side panels show the one-dimensional profile likelihoods. 67
5.6 Mass ordering sensitivity as a function of exposure for DUNE atmospheric neutrinos, using NuFit 5.2 oscillation parameters with SK atmospheric data. The blue curve assumes normal ordering is true, and the red curve assumes inverted ordering is true. . . . 69
5.7 Mass ordering sensitivity as a function of exposure for DUNE atmospheric neutrinos, using NuFit 5.2 oscillation parameters without SK atmospheric data. Both orderings reach higher significance compared to the with-SK configuration, and the relative ordering of the curves is reversed. 70
5.8 Comparison of mass ordering sensitivity between the two NuFit 5.2 configurations. The upper panel shows all four sensitivity curves (NO and IO truth, with and without SK atmospheric data). The lower panel shows $\Delta\sigma = \sigma_{\mathrm{noSK}} - \sigma_{\mathrm{wSK}}$ for each true ordering. 71

Introduction

Neutrinos are at the forefront of modern physics. Within the last two decades, our understanding of the fundamental nature of the universe has been shaped by these elusive particles. According to the Standard Model of Particle Physics, neutrinos should be massless. However, experimental observations from the turn of the 21st century, including Super-Kamiokande (SK) in Japan, definitively showed that neutrinos oscillate between flavors, a phenomenon that requires neutrinos to have mass. This result provided clear evidence of physics beyond the Standard Model (BSM), and the open questions that follow from it such as role of neutrinos in the matter-antimatter asymmetry in our Universe, and the origin of neutrino masses represent some of the most concrete steps currently available for probing BSM physics.

Several of these open questions drive the current generation of neutrino experiments. Resolving the mass ordering of the neutrino mass eigenstates, whether $m_3 > m_1$ (normal ordering) or $m_3 < m_1$ (inverted ordering), may improve our understanding of the generational structure of the Standard Model, constrain BSM models, and inform the interpretation of complementary experiments. In particular, neutrinoless double-beta decay ($0\nu\beta\beta$) searches, which probe the possible Majorana nature of neutrinos, are themselves sensitive to the mass ordering. Agreement between oscillation-based and $0\nu\beta\beta$-based determinations would strengthen confidence in both results, while a disagreement could point to new physics in either sector. The CP-violating phase δ_{CP} in the leptonic mixing matrix is also of significant experimental interest: if sufficiently large, CP violation in the lepton sector could drive leptogenesis, a mechanism that may explain the observed matter-antimatter asymmetry in the universe. Finally, the question of whether the atmospheric mixing angle θ_{23} lies in the upper octant, lower octant, or is exactly maximal remains open. Maximal mixing would hint at an underlying symmetry that is not yet understood.

The Deep Underground Neutrino Experiment (DUNE) is a next-generation, US-led international long-baseline neutrino oscillation experiment designed to address these questions. DUNE will use an intense neutrino beam 2.4 MW (by 2030)[4] produced at Fermilab, directed over a baseline of approximately 1300 km to cutting-edge liquid argon time projection chamber (LArTPC) far detector modules located at the Sanford Underground Research Facility (SURF) in Lead, South Dakota. Although the beam program is DUNE's primary oscillation measurement tool, atmospheric neutrinos offer a complementary probe. Their broad, continuous L/E coverage, spanning roughly 10^{-1} to 10^5 km/GeV, provides independent sensitivity to the mass ordering, the atmospheric mixing angle, and matter effects. Upward-going atmospheric neutrinos traverse the Earth's mantle and core, sampling varying density profiles that enhance matter-driven oscillation effects. Atmospheric neutrinos also provide simultaneous access to both neutrinos and antineutrinos without requiring beam polarity switching. Importantly, atmospheric neutrino data collection can begin as soon as the far detector modules are filled with liquid argon, independently of the beam commissioning timeline.

This work represents one of the first dedicated studies of DUNE's atmospheric neutrino sensitivity using the GUNDAM fitting framework. The oscillation probability grid convergence study presented in Chapter 4 is, to the author's knowledge, the first rigorous and systematic evaluation of the grid resolution requirements for this analysis, confirming that the 300×300 configuration adopted by the working group is sufficient for reliable parameter extraction. The analysis also explores the treatment of atmospheric neutrino production height within GUNDAM, comparing fixed and averaged configurations and their impact on the oscillation probability calculation. Sensitivity results are presented through one-dimensional profile likelihoods, two-dimensional contour scans, and a mass ordering determination study, all constructed from simulated datasets at 400 kt·yr exposure.

The thesis is organized as follows. Chapter 1 reviews neutrino oscillation physics and atmospheric neutrino production, which underpins the analysis in later chapters. Chapter 2 describes the DUNE experiment, including the beamline, far detector technology, and atmospheric neutrino detection capabilities. Chapter 3 presents the statistical methodology, the GUNDAM fitting framework, the analysis configuration, and studies of the neutrino production height parameterization. Chapter 4 develops the oscillation probability grid optimization, including convergence studies that determine the resolution requirements for the analysis. Chapter 5 presents the sensitivity results, including profile likelihoods, two-dimensional contours, and the mass ordering study.

Chapter 1

Neutrino Physics

1.1 Introduction to Neutrinos

1.1.1 What Neutrinos Are

The Standard Model

The standard model (SM) of particle physics describes three of the four known fundamental interactions (electromagnetic, weak, and strong forces), as well as the fundamental particles of nature.

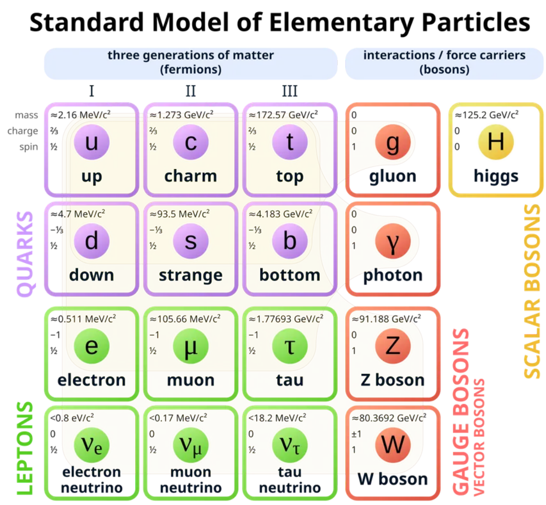

Figure 1.1: Standard Model

The SM describes elementary particles as either fermions or bosons. Fermions are the building blocks of matter, consisting of quarks and leptons with half integer spin. Bosons have integer spin and mediate the fundamental interactions: photons for electromagnetism, W and Z bosons for weak interactions, and gluons for the strong interaction. The Higgs boson is responsible for giving particles mass, and each fundamental particle has its corresponding antiparticle, with opposite charge and parity.

While the SM has great success in describing a variety of physical phenomena, it fails to explain the matter-antimatter asymmetry of the universe and the fundamental nature of dark matter. Furthermore, according to the SM, neutrinos are predicted to be massless. Experiments at the turn of the century including Super-Kamiokande (SK) [38] and Sudbury Neutrino Observatory (SNO) [9] definitively showed that neutrinos oscillate between flavors as they propagate through space. This phenomenon is due to the fact that neutrinos produced with a given flavor are themselves superpositions of distinct neutrino mass eigenstates (what is commonly referred to as mixing of flavor and mass). As the neutrino propagates through space, it travels in a basis known as the mass basis, in which these distinct mass eigenstates travel at different rates, allowing for the eventually detected neutrino to become a different flavor than what is started out as.

While the absolute mass of individual neutrinos is currently unknown, they are at least 500,000 times smaller than the next smallest fundamental particle, the electron, which has a mass of $m_e \approx 9.109 \times 10^{-31}$ kg (0.511 MeV/c^2) [64].

Properties and Flavors

Neutrinos are neutral, nearly massless, spin-$\frac{1}{2}$ particles that interact only via the weak force and gravity. They are the most abundant massive particle in the Universe and belong to the lepton family of the SM. Neutrinos have three active flavors–electron (ν_e), muon (ν_μ), and tau (ν_τ)–defined by the charged lepton produced simultaneously during weak interactions.

1.1.2 History of Neutrino Discovery

Pauli's Postulate (1930): Neutrino Hypothesis

In 1930, in a desperate attempt to explain the continuous β-decay spectrum (Figure 1.2) [56, 57], physicist Wolfgang Pauli hypothesized a neutral particle of negligible mass.

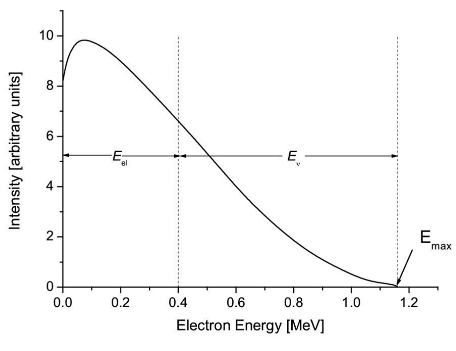

Figure 1.2: The continuous energy spectrum of electrons emitted in β-decay. Image credit: H. Paul, CC BY-SA 4.0 [56].

Pauli was convinced the particle would never be sensibly measured, famously stating,

> "I have done a terrible thing. I have postulated a particle that cannot be detected." [58]

While Pauli initially called the particle a neutron, Enrico Fermi later coined the name neutrino while developing his theory of weak interactions [36]. When Bethe and Peierls calculated the neutrino interaction cross section to be extremely small (10^{-44} cm^2 at $E_\nu = 2$ MeV) [19], detection seemed nearly impossible.

Cowan & Reines (1956): Electron Antineutrino Discovery

In 1956, C. Cowan and F. Reines made the first direct experimental confirmation of the neutrino [23] by studying inverse beta decay:

$$\bar{\nu}_e + p \rightarrow e^+ + n \tag{1.1}$$

Using electron antineutrinos from a nuclear reactor, they observed interactions in water containing dissolved cadmium chloride ($CdCl_2$). The positron annihilated with an electron releasing two gamma rays, and the neutron was captured by cadmium, producing another gamma ray, providing a conclusive signature. This led to the 1995 Nobel Prize in Physics.

Lederman, Schwartz, Steinberger (1962): Muon Neutrino Discovery

In 1962, Lederman, Schwartz, and Steinberger discovered the muon neutrino [25]. Using the Alternating Gradient Synchrotron (AGS) at Brookhaven National Laboratory, they produced a shower of pi mesons whose decays yielded muons and neutrinos. Only the neutrinos could pass through a 5,000 ton steel wall to reach a spark chamber detector. The neutrinos interacted with aluminum plates producing muon trails, definitively providing evidence of a second neutrino flavor and earning the 1988 Nobel Prize in Physics.

DONUT (2000): Tau Neutrino Discovery

The last active flavor was discovered by the DONUT (Direct Observation of the Nu Tau) collaboration [49]. Using neutrinos produced from a high energy proton beam striking a tungsten target, the experiment detected the characteristic tau lepton signature from tau neutrino interactions with iron nuclei, providing direct evidence of ν_τ.

LEP Z-Width Measurement: Constraint on Maximum Active Neutrino Flavors

The CERN LEP experiment constrained the number of active neutrino flavors by measuring the Z-boson decay width [60]. After subtracting the visible decay width from the total, the remaining invisible width, corresponding to undetectable neutrino-antineutrino pairs, was consistent with exactly three active neutrino flavors. This measurement does not rule out additional "sterile neutrinos" that do not couple to the Z boson.

1.1.3 The Solar Neutrino Problem

In the late 1960s, Raymond Davis, Jr. and John N. Bahcall conducted the Homestake experiment in Lead, South Dakota [26], aiming to count neutrinos produced in the Sun via nuclear fusion. They observed only about one third of the expected neutrinos, introducing the *solar neutrino problem.* For the discovery of solar neutrinos, Davis was awarded the 2002 Nobel Prize in Physics [54].

In response to the solar neutrino problem, theorist Bruno Pontecorvo proposed *neutrino oscillations*, in which neutrinos transform between flavors during propagation [59]. Multiple experiments confirmed the deficit, including the gallium-based GALLEX [42] and SAGE [1], and the water-based Kamiokande [44].

A theoretical framework for understanding the deficit involves the Wolfenstein matter potential [63]. When neutrinos propagate through matter, they undergo coherent forward scattering off particles in the medium. This process involves no momentum transfer and no deflection of the neutrino; rather, the collective weak interaction with the medium produces an effective potential that modifies the neutrino's phase during propagation, analogous to a refractive index for light in glass. Crucially, this potential is flavor dependent: electron neutrinos experience an additional contribution from charged current scattering off electrons (via W boson exchange), an interaction not available to ν_μ or ν_τ, which only interact through the flavor universal neutral current. This flavor dependent potential modifies the effective mixing angle and mass-squared splitting in matter. The Mikheyev-Smirnov-Wolfenstein (MSW) effect [51] describes how, at a particular electron density, the effective mixing angle reaches $\pi/4$ regardless of the vacuum mixing angle, producing a resonance in which efficient flavor conversion occurs. Solar neutrinos, produced in the dense core and propagating outward through a decreasing density profile, can cross this resonance and undergo significant ν_e disappearance, providing a compelling explanation for the observed flux suppression.

1.1.4 The Atmospheric Neutrino Problem

In the 1980s, a deficit in atmospheric muon neutrinos was observed relative to predictions. Cosmic ray interactions in the atmosphere produce pions and kaons, which decay into muons and neutrinos, with an expected ν_μ-to-ν_e ratio of approximately 2:1 assuming no new physics; the origin of this ratio from pion and kaon decay chains is discussed in Section 1.4.

The Kamiokande and Irvine-Michigan-Brookhaven (IMB) detectors, both water-based Cherenkov detectors originally developed for proton decay searches, were among the first to characterize this anomaly. Their large instrumented volumes (necessary to compensate for small neutrino cross sections) naturally lent themselves to detecting atmospheric neutrinos [18]. Using the double ratio $R = (\mu/e)_{\text{data}}/(\mu/e)_{\text{MC}}$, both experiments found $R \approx 0.6$, indicating fewer muon neutrinos than expected [20, 43].

In 1998, Super-Kamiokande observed a striking zenith angle dependence of the muon neutrino flux (Figure 1.3) [38]. Upward-going neutrinos, traveling approximately 13,000 km through the Earth, showed roughly half the expected flux, while downward-going neutrinos, produced about 15 km above the detector, arrived at the expected rate. Electron neutrinos showed no such asymmetry.

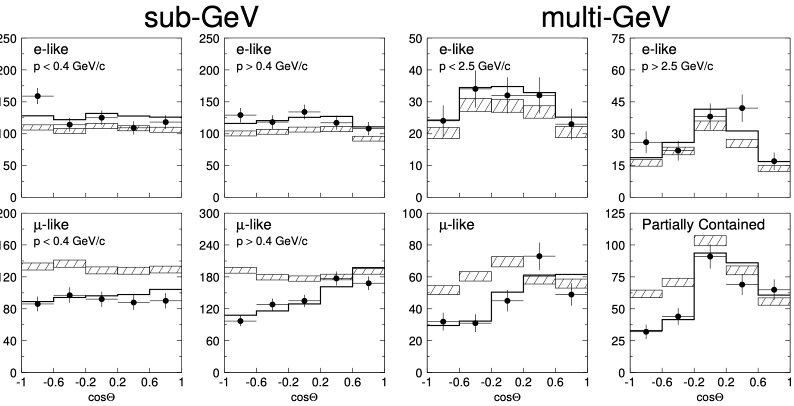

Figure 1.3: Zenith angle distributions of μ-like and e-like events for sub-GeV and multi-GeV data sets measured by Super-Kamiokande. The hatched region shows the no oscillation expectation; the bold line shows the best fit $\nu_\mu \leftrightarrow \nu_\tau$ oscillation. The μ-like deficit for upward-going events ($\cos\Theta < 0$) provided definitive evidence for atmospheric neutrino oscillations. Figure from [38].

1.1.5 Discovery of Neutrino Oscillations

The Super-Kamiokande (SK) zenith angle measurement provided the first definitive evidence for neutrino oscillations [38]. The zenith angle serves as a proxy for the neutrino propagation baseline length L. Under the oscillation hypothesis, a neutrino's survival probability depends on L/E: short baseline (downward-going) neutrinos have insufficient distance to oscillate, while long baseline (upward-going) neutrinos have. The disappearance of only muon neutrinos, following the expected L/E dependence, pointed to $\nu_\mu \to \nu_\tau$ oscillations. The data was consistent with maximal mixing ($\sin^2 2\theta_{23} \approx 1$) and $\Delta m^2_{32} \approx 2.5 \times 10^{-3}$ eV^2, providing the first compelling evidence that neutrinos have mass.

The solar neutrino problem was independently resolved by the Sudbury Neutrino Observatory (SNO) [9]. SNO used heavy water (D_2O) to study neutrinos through neutral current (NC), charged current (CC), and elastic scattering (ES) channels. The NC measurement of the total neutrino flux was consistent with solar model predictions, while the CC measurement, sensitive only to ν_e, showed only about one third of the expected total. This provided direct evidence for neutrino flavor transformation: the solar ν_e flux was not disappearing, but converting into ν_μ and ν_τ through the MSW effect described in Section 1.1.3.

These two independent results, from different neutrino sources and different L/E regimes, converged on the same conclusion: neutrinos are massive particles that undergo flavor oscillations described by the Pontecorvo-Maki-Nakagawa-Sakata (PMNS) mixing matrix. Takaaki Kajita (Super-Kamiokande) and Arthur B. McDonald (SNO) shared the 2015 Nobel Prize in Physics for the discovery of neutrino oscillations [55]. With neutrino oscillations firmly established, several fundamental questions about the oscillation framework remain unresolved and are discussed in Section 1.3.

1.2 Neutrino Oscillations

Neutrino oscillations occur because neutrino flavor eigenstates $(\nu_e, \nu_\mu, \nu_\tau)$ produced and absorbed in weak interactions are coherent linear superpositions of the neutrino mass eigenstates (ν_1, ν_2, ν_3), related by the Pontecorvo–Maki–Nakagawa–Sakata (PMNS) leptonic mixing matrix (1.5).

Because the mass eigenstates have different masses, they accumulate different phases as they propagate. When the neutrino is detected via a weak interaction, interference between these mass eigenstate components makes the probability of observing a given flavor vary sinusoidally with $\frac{L}{E}$, where L is the neutrino's propagation length and E is the neutrino energy, with frequencies set by the mass squared splittings Δm^2_{ij}.

The flavor eigenstates,

$$(\nu_e, \nu_\mu, \nu_\tau)\,, \tag{1.2}$$

are related to the mass eigenstates,

$$(\nu_1, \nu_2, \nu_3)\,, \tag{1.3}$$

through the PMNS matrix:

$$\nu_\alpha = \sum_{i=1}^{3} U_{\alpha i}\nu_i, \tag{1.4}$$

where $\alpha = e, \mu, \tau$ labels flavor and i labels mass eigenstates.

The PMNS matrix is parameterized by three mixing angles $(\theta_{12}, \theta_{23}, \theta_{13})$ and a CP-violating phase δ_{CP}:

$$U = \begin{pmatrix} 1 & 0 & 0 \\ 0 & c_{23} & s_{23} \\ 0 & -s_{23} & c_{23} \end{pmatrix} \begin{pmatrix} c_{13} & 0 & s_{13}e^{-i\delta_{\text{CP}}} \\ 0 & 1 & 0 \\ -s_{13}e^{i\delta_{\text{CP}}} & 0 & c_{13} \end{pmatrix} \begin{pmatrix} c_{12} & s_{12} & 0 \\ -s_{12} & c_{12} & 0 \\ 0 & 0 & 1 \end{pmatrix}, \tag{1.5}$$

where $c_{ij} = \cos\theta_{ij}$ and $s_{ij} = \sin\theta_{ij}$.

1.2.1 Vacuum Oscillations

In vacuum, the time evolution of a neutrino mass eigenstate with energy E_i is given by

$$\nu_i(t) = e^{-iE_i t}\nu_i(0). \tag{1.6}$$

For relativistic neutrinos, the energy can be approximated as

$$E_i \simeq E + \frac{m_i^2}{2E}, \tag{1.7}$$

where E is the neutrino energy and m_i is the mass of the ith eigenstate.

The general three-flavor probability for a neutrino produced in flavor state ν_α to be detected as flavor ν_β after traveling a distance L is

$$P_{\alpha\to\beta} = \left|\sum_i U_{\beta i}U^*_{\alpha i}\exp\left(-i\frac{\Delta m^2_{i1}L}{2E}\right)\right|^2, \tag{1.8}$$

where $\Delta m^2_{ij} = m^2_i - m^2_j$. As a pedagogical limit, if only two flavors mix with a single mixing angle θ and a single mass-squared splitting Δm^2, Equation 1.8 reduces to

$$P_{\alpha\to\beta} = \sin^2(2\theta)\,\sin^2\left(\frac{\Delta m^2 L}{4E}\right), \tag{1.9}$$

which illustrates the characteristic L/E dependence of neutrino oscillations. The realistic three-flavor case, relevant for DUNE, requires expanding Equation 1.8 channel by channel.

For the channels relevant to atmospheric neutrino oscillations, the general probability in Equation 1.8 can be expanded in terms of the PMNS matrix elements and the oscillation phases $\Delta_{ij} \equiv \Delta m^2_{ij} L/4E$. Throughout this section, the "atmospheric limit" refers to the regime relevant to DUNE, in which $\Delta_{21} \ll 1$ and $\Delta_{31} \approx \Delta_{32}$. For each channel we give the exact three-flavor expression first, followed by its atmospheric limit approximation.

ν_μ **survival probability:** For the disappearance channel $\nu_\mu \to \nu_\mu$, the imaginary terms in Equation 1.8 vanish and the exact three-flavor probability reduces to [41]

$$P(\nu_\mu \to \nu_\mu) = 1 - 4|U_{\mu 1}|^2|U_{\mu 2}|^2 \sin^2\Delta_{21} - 4|U_{\mu 1}|^2|U_{\mu 3}|^2 \sin^2\Delta_{31} - 4|U_{\mu 2}|^2|U_{\mu 3}|^2 \sin^2\Delta_{32}\,. \tag{1.10}$$

In the atmospheric limit, the first term is negligible and the second and third combine to give the atmospheric limit survival probability

$$P(\nu_\mu \to \nu_\mu) \approx 1 - 4|U_{\mu 3}|^2\big(1 - |U_{\mu 3}|^2\big)\sin^2\left(\frac{\Delta m^2_{32} L}{4E}\right), \tag{1.11}$$

where $|U_{\mu 3}|^2 = \sin^2\theta_{23}\cos^2\theta_{13}$. The effective mixing factor $4\sin^2\theta_{23}\cos^2\theta_{13}(1-\sin^2\theta_{23}\cos^2\theta_{13})$ reduces to $\sin^2 2\theta_{23}$ in the limit $\theta_{13} \to 0$, recovering the two flavor form of Equation 1.9. Because θ_{13} is small ($\sin^2\theta_{13} \approx 0.022$ from Table 3.1), the correction is at the percent level, and atmospheric ν_μ disappearance remains primarily sensitive to θ_{23} and Δm^2_{32}.

$\nu_\mu \to \nu_e$ **appearance probability.** The appearance channel provides sensitivity to θ_{13}, $\delta_{\rm CP}$, and the mass ordering. For $\alpha \neq \beta$, the exact three-flavor vacuum probability contains both CP-conserving and CP-violating terms [41]:

$$\begin{aligned} P(\nu_\mu \to \nu_e) = &-4\sum_{i>j} \mathrm{Re}\big(U^*_{\mu i}\,U_{ei}\,U_{\mu j}\,U^*_{ej}\big)\sin^2\Delta_{ij} \\ &+ 2\sum_{i>j} \mathrm{Im}\big(U^*_{\mu i}\,U_{ei}\,U_{\mu j}\,U^*_{ej}\big)\sin 2\Delta_{ij}\,. \end{aligned} \tag{1.12}$$

The first sum is CP-even (identical for neutrinos and antineutrinos) while the second is CP-odd (changes sign under $\nu \to \bar\nu$). In the atmospheric limit ($\Delta_{21} \to 0$), the leading three-flavor vacuum appearance probability simplifies to

$$P(\nu_\mu \to \nu_e) \approx \sin^2\theta_{23}\,\sin^2 2\theta_{13}\,\sin^2\left(\frac{\Delta m^2_{32} L}{4E}\right). \tag{1.13}$$

This leading term scales as $\sin^2\theta_{23}$ rather than $\sin^2 2\theta_{23}$, making the appearance channel sensitive to the θ_{23} octant, in contrast to the disappearance channel (Equation 1.11).

All CP-violating effects in Equation 1.12 are proportional to the Jarlskog invariant [41],

$$J = c_{12}\, s_{12}\, c_{23}\, s_{23}\, c_{13}^2\, s_{13}\, \sin\delta_{\mathrm{CP}}\,, \tag{1.14}$$

which is the unique parametrization-independent measure of CP violation in the PMNS matrix and vanishes if any mixing angle is zero or if $\delta_{\mathrm{CP}} = 0$ or π. The CP asymmetry between neutrino and antineutrino appearance channels is [41]

$$P(\nu_\mu \to \nu_e) - P(\bar{\nu}_\mu \to \bar{\nu}_e) = 16\, J\, \sin\Delta_{21}\, \sin\Delta_{31}\, \sin\Delta_{32}\,, \tag{1.15}$$

which requires all three mass splittings to contribute simultaneously and is therefore an intrinsically three flavor effect.

$\nu_\mu \to \nu_\tau$ appearance. By unitarity of the PMNS matrix, the $\nu_\mu \to \nu_\tau$ transition probability is fully determined by the other two channels:

$$P(\nu_\mu \to \nu_\tau) = 1 - P(\nu_\mu \to \nu_\mu) - P(\nu_\mu \to \nu_e)\,. \tag{1.16}$$

The DUNE far detector does not reconstruct τ leptons directly, so ν_τ appearance manifests as a depletion of the ν_μ sample. The unitarity relation ensures that fits to the ν_μ and ν_e selected samples implicitly account for this channel.

1.2.2 Matter Effects

When neutrinos propagate through matter, coherent forward scattering on electrons modifies their effective Hamiltonian. Electron neutrinos experience an additional charged-current potential:

$$V_{\mathrm{CC}} = \sqrt{2} G_F N_e, \tag{1.17}$$

where G_F is the Fermi constant and N_e is the electron number density.

Resonant enhancement of oscillations occurs when

$$\sqrt{2} G_F N_e \simeq \frac{\Delta m^2 \cos 2\theta}{2E}. \tag{1.18}$$

For atmospheric neutrinos traversing the Earth, matter effects become significant for multi-GeV energies and long baselines, providing sensitivity to the neutrino mass ordering. Since the matter potential changes sign for antineutrinos, the MSW effect introduces an asymmetry between neutrino and antineutrino oscillations that depends on the sign of Δm^2_{31}. The complexity of the exact three-flavor oscillation probabilities in matter motivates the use of numerical approaches, as discussed in Section 3.3.

Approximate appearance probability in matter. For neutrinos propagating through matter of constant electron number density n_e, an analytic approximation to the $\nu_\mu \to \nu_e$ appearance probability, valid to second order in the small parameters $\alpha \equiv \Delta m^2_{21}/\Delta m^2_{31}$ and $\sin\theta_{13}$, was derived by Cervera et al. [22] and Freund [37]:

$$
\begin{aligned}
P(\overset{(-)}{\nu}_\mu \to \overset{(-)}{\nu}_e) \approx{}& \sin^2\theta_{23}\,\frac{\sin^2 2\theta_{13}}{(\hat{A}-1)^2}\,\sin^2\big[(\hat{A}-1)\Delta_{31}\big] \\
&\mp\,\alpha\,\frac{J_0\sin\delta_{\mathrm{CP}}}{\hat{A}(1-\hat{A})}\,\sin\Delta_{31}\,\sin(\hat{A}\Delta_{31})\,\sin\big[(1-\hat{A})\Delta_{31}\big] \\
&+\,\alpha\,\frac{J_0\cos\delta_{\mathrm{CP}}}{\hat{A}(1-\hat{A})}\,\cos\Delta_{31}\,\sin(\hat{A}\Delta_{31})\,\sin\big[(1-\hat{A})\Delta_{31}\big] \\
&+\,\alpha^2\cos^2\theta_{23}\,\frac{\sin^2 2\theta_{12}}{\hat{A}^2}\,\sin^2(\hat{A}\Delta_{31})\,,
\end{aligned}
\tag{1.19}
$$

where the upper sign $(-)$ applies to neutrinos and the lower sign $(+)$ to antineutrinos, and the auxiliary quantities are defined as

$$
\alpha = \Delta m^2_{21}/\Delta m^2_{31}\,, \qquad \Delta_{ij} = \Delta m^2_{ij}\,L/4E\,, \tag{1.20}
$$

$$
\hat{A} = 2\sqrt{2}\,G_F\,n_e\,E\,/\,\Delta m^2_{31}\,, \qquad J_0 = \sin 2\theta_{12}\,\sin 2\theta_{13}\,\sin 2\theta_{23}\,\cos\theta_{13}\,. \tag{1.21}
$$

For antineutrinos, the matter parameter changes sign: $\hat{A} \to -\hat{A}$.

The four terms have distinct physical roles. The first, proportional to $\sin^2\theta_{23}\sin^2 2\theta_{13}$, is the dominant θ_{13}-driven appearance term and is resonantly enhanced when $\hat{A} \to 1$. This MSW resonance occurs for neutrinos if $\Delta m^2_{31} > 0$ (normal ordering) and for antineutrinos if $\Delta m^2_{31} < 0$ (inverted ordering), providing sensitivity to the mass ordering (Section 5.3). The second term, proportional to $\sin\delta_{\mathrm{CP}}$, is the CP-violating interference between the atmospheric and solar oscillation amplitudes. The third, proportional to $\cos\delta_{\mathrm{CP}}$, is a CP-conserving interference term. The fourth is the pure solar contribution, suppressed by $\alpha^2 \approx 10^{-3}$.

An important limitation of Equation 1.19 is that it assumes a constant matter density along the neutrino trajectory. This is a reasonable approximation for long-baseline beam experiments such as T2K [3] and NOvA [7], where neutrinos travel through the Earth's crust at roughly constant density over baselines of a few hundred kilometers. For atmospheric neutrinos, this approximation breaks down: upward-going neutrinos traverse the mantle and potentially the core, encountering density variations from $\sim$3 g/cm^3 near the surface to $\sim$13 g/cm^3 at the inner core, as described by the PREM density profile [33]. No closed-form analytic solution exists for propagation through a continuously varying density profile [37].

The numerical approach used in this analysis handles the varying density exactly. The oscillation probability engine CUDAProb3 [61] divides the Earth into concentric shells of constant density following the PREM model. Within each shell, the three-flavor evolution equation

$$
i\frac{d}{dx}\begin{pmatrix}\nu_e\\ \nu_\mu\\ \nu_\tau\end{pmatrix} = \left[\frac{1}{2E}\,U\,\mathbf{M}^2\,U^\dagger + \mathbf{V}(x)\right]\begin{pmatrix}\nu_e\\ \nu_\mu\\ \nu_\tau\end{pmatrix}, \tag{1.22}
$$

where $\mathbf{M}^2 = \mathrm{diag}(0, \Delta m^2_{21}, \Delta m^2_{31})$ and $\mathbf{V}(x) = \mathrm{diag}(\sqrt{2}\,G_F\,n_e(x), 0, 0)$, is solved exactly by matrix exponentiation at the constant density of that shell. The total propagation matrix is obtained by multiplying the individual shell matrices in sequence. All oscillation probabilities presented in this thesis are computed using this method, as described in Section 3.3.

Despite this limitation, the analytic form is useful for identifying which parameters enter at leading order and where the matter resonance occurs. The qualitative dependence on θ_{13}, θ_{23}, δ_{CP}, $\hat{A}$, and the mass ordering carries over to the full numerical calculation, even when the constant-density approximation itself does not hold.

1.2.3 CP Violation

The CP-violating phase δ_{CP} further enriches the structure of neutrino oscillations in matter. In three flavor oscillations, CP-violating effects arise from interference between multiple oscillation frequencies and are intrinsically entangled with matter-induced asymmetries.

1.3 Open Questions in Neutrino Physics

Several major open questions motivate the next generation of neutrino experiments, including DUNE.

Mass Ordering Hierarchy. Three neutrino mass eigenstates (ν_1, ν_2, ν_3) are established and the mass-squared splittings Δm^2_{ij} have been measured, but the sign of Δm^2_{32} remains unknown. In *normal ordering* (NO), $m_1 < m_2 < m_3$; in *inverted ordering* (IO), $m_3 < m_1 < m_2$. Determining the ordering has implications for neutrinoless double beta decay searches, cosmological models, and supernova neutrino physics.

CP Violation. The PMNS matrix contains a complex phase δ_{CP}. If non-zero, neutrinos and antineutrinos oscillate at different rates, implying CP violation in the leptonic sector. Such violation could contribute to explaining the matter–antimatter asymmetry through leptogenesis [39]. While δ_{CP} is not yet well constrained, DUNE is specifically designed to measure this parameter.

The $\boldsymbol{\theta_{23}}$ Octant. Experiments including Super-K [38], T2K [3], and NOvA [7] have established that $\sin^2 2\theta_{23}$ is close to unity. However, exactly maximal mixing ($\theta_{23} = 45°$) would suggest an underlying symmetry such as μ–τ symmetry, while a deviation would require flavor models to account for it. DUNE and Hyper-Kamiokande are designed to resolve the octant [2, 4].

Dirac vs. Majorana Nature. An open question beyond the reach of oscillation experiments: are neutrinos their own antiparticles? If so, they are *Majorana* fermions; if not, *Dirac* fermions. Neutrinoless double beta decay ($0\nu\beta\beta$) searches can resolve this [27]. If neutrinos are Majorana particles, the seesaw mechanism would provide a natural explanation for their small masses [52, 65].

Absolute Neutrino Mass Neutrino oscillation experiments only measure the mass squared splittings Δm^2_{ij}, where i, j denote the neutrino mass eigenstates. In this sense, oscillation experiments are blind to the overall mass scale; all three neutrino masses could be nearly degenerate or strongly hierarchical. There are three main avenues by which physicists attempt to constrain the absolute neutrino mass: direct kinematic measurements such as the KATRIN experiment, which constrains the effective electron antineutrino mass from the tritium beta decay endpoint; cosmological bounds on the sum of neutrino masses from observations of the Cosmic Microwave Background (CMB) and large scale structure; and neutrinoless double beta decay ($0\nu\beta\beta$) experiments, which probe the effective Majorana mass $m_{\beta\beta}$. Note that the effective Majorana mass $m_{\beta\beta} = \left|\sum_i U_{ei}^2 m_i\right|$ is a coherent sum over the mass eigenstates, in which the Majorana phases can lead to partial or complete cancellation; it is therefore distinct from the absolute mass scale and its interpretation depends on knowledge of the mixing parameters.

While DUNE cannot answer all of these questions, it is uniquely positioned to probe the mass ordering, δ_{CP}, and the θ_{23} octant through its long-baseline beam and atmospheric neutrino programs, as discussed in Chapter 2.

1.4 Atmospheric Neutrinos

Atmospheric neutrinos are produced by interactions of cosmic rays with atomic nuclei in Earth's atmosphere. High energy cosmic ray particles, predominantly protons, collide with nitrogen and oxygen nuclei in the upper atmosphere, initiating hadronic cascades that produce mesons. The subsequent decays of pions and kaons generate the atmospheric neutrino flux.

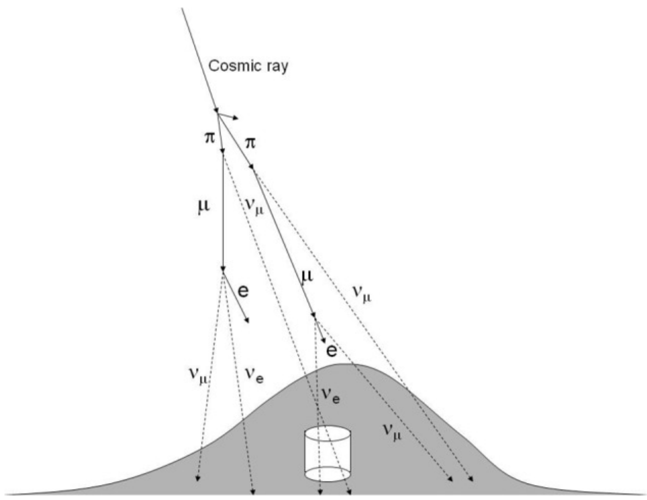

Figure 1.4: Schematic of atmospheric neutrino production. A primary cosmic ray interacts in the upper atmosphere, producing pions. Charged pions decay to muons and ν_μ, and the muons subsequently decay to electrons, ν_e, and ν_μ. Figure from Ref. [48].

1.4.1 Cosmic Ray Primaries

The primary cosmic ray flux is composed of approximately 90% protons, 9% helium nuclei, and 1% heavier nuclei [8, 53]. The energy spectrum spans approximately eleven orders of magnitude, from $\sim 10^9$ eV to beyond 10^{20} eV. Below $\sim$ 100 TeV, the spectrum is measured directly by balloon based and space based experiments including AMS-02, PAMELA, CALET, and DAMPE [53]. At higher energies, ground-based observatories such as the Pierre Auger Observatory measure extensive air showers [53].

The differential flux follows an approximate power law,

$$\frac{d\Phi}{dE} \propto E^{-\gamma}, \tag{1.23}$$

with spectral index $\gamma \approx 2.7$ below the "knee" at $\sim$ 3 PeV [8]. This steeply falling spectrum is inherited by the atmospheric neutrino flux, which also decreases rapidly with energy [48], so lower-energy cosmic rays dominate atmospheric neutrino production. For DUNE, the relevant primary cosmic ray energies range from approximately 1 GeV to 1 TeV, corresponding to neutrino energies from sub-GeV to $\sim$ 100 GeV.

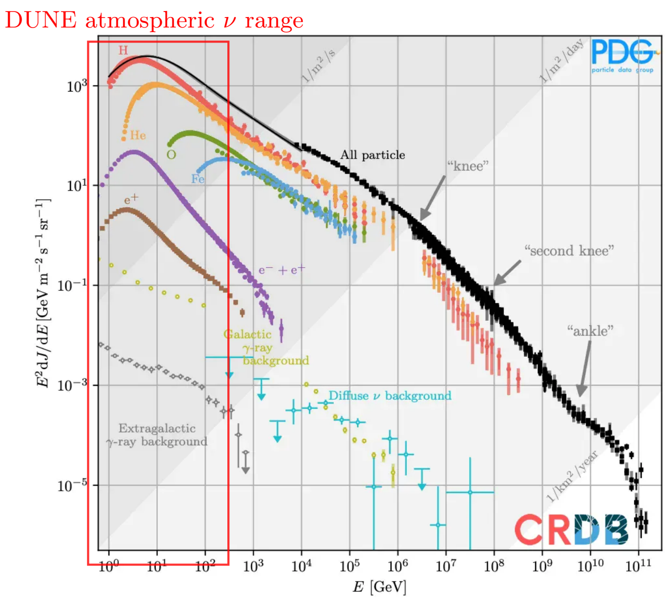

Figure 1.5: The cosmic ray energy spectrum, showing flux multiplied by E^2 versus energy. Individual components (H, He, O, Fe) are shown at lower energies, with the all-particle spectrum extending to the highest energies. The knee ($\sim 10^6$ GeV), second knee, and ankle features are labeled. The red box indicates the approximate primary cosmic ray energy range responsible for producing the atmospheric neutrino flux relevant to DUNE far detector reconstruction studies, which covers neutrino energies from 100 MeV to 10 GeV [32]. Figure from the Cosmic Ray Database (CRDB) [10].

1.4.2 Hadronic Interactions and Meson Production

When cosmic ray protons collide with atmospheric nuclei, mesons are produced. The first interaction typically occurs at an altitude of $\sim$ 15 km [21]:

$$p + \mathrm{N} \rightarrow \pi^{\pm}, \pi^0 + X, \tag{1.24}$$

where N represents an atmospheric nucleus (nitrogen or oxygen) and X denotes additional hadronic products. Kaon production requires higher energies and is therefore rarer. The collision produces roughly equal numbers of π^+, π^-, and π^0, but only the charged pions matter for neutrino production since $\pi^0 \rightarrow \gamma + \gamma$.

1.4.3 Decay Chains and the Flavor Ratio

The dominant pion decay mode (branching ratio 99.99% [53]) is leptonic:

$$\pi^+ \rightarrow \mu^+ + \nu_\mu, \qquad \pi^- \rightarrow \mu^- + \bar{\nu}_\mu \tag{1.25}$$

The resulting muons then decay with a branching ratio of nearly 100% [53]:

$$\mu^+ \rightarrow e^+ + \nu_e + \bar{\nu}_\mu, \qquad \mu^- \rightarrow e^- + \bar{\nu}_e + \nu_\mu \tag{1.26}$$

At higher energies, pions are more likely to interact with air nuclei before decaying. Above the pion critical energy ($\epsilon_\pi \approx 115$ GeV), kaon decays become the dominant source of atmospheric neutrinos [40].

Table 1.1: Principal decay modes contributing to atmospheric neutrino production. Branching ratios from [53].

Decay	Products	Branching Ratio	Neutrino Type
$\pi^+ \rightarrow \mu^+\nu_\mu$	μ^+, ν_μ	99.99%	ν_μ
$\mu^+ \rightarrow e^+\nu_e\bar{\nu}_\mu$	e^+, ν_e, $\bar{\nu}_\mu$	100%	ν_e, $\bar{\nu}_\mu$
$K^+ \rightarrow \mu^+\nu_\mu$	μ^+, ν_μ	63.6%	ν_μ
$K^+ \rightarrow \pi^0 e^+\nu_e$	π^0, e^+, ν_e	5.1%	ν_e

The charge conjugate decays of π^-, μ^-, and K^- proceed with identical branching ratios (to the precision shown above) and produce the corresponding antineutrinos, contributing the $\bar{\nu}_\mu$ and $\bar{\nu}_e$ components to the atmospheric flux.

Counting the neutrinos produced in a complete $\pi \rightarrow \mu \rightarrow e$ decay chain yields two muon flavor neutrinos and one electron flavor neutrino, giving the expected flavor ratio:

$$\frac{(\nu_\mu + \bar{\nu}_\mu)}{(\nu_e + \bar{\nu}_e)} \approx 2 \tag{1.27}$$

As discussed in Section 1.1, this expected ratio was not observed experimentally, establishing the atmospheric neutrino anomaly [38].

1.4.4 Energy Dependence of the Flavor Ratio

In the muon rest frame, the mean lifetime is $\tau_0 \approx 2.197$ μs [53], so even traveling at the speed of light a muon would cover only $c\tau_0 \approx 660$ m before decaying, far shorter than the $\sim$ 15 km path to the surface, and one might naively conclude that essentially no muons should reach the ground.

Time dilation resolves this apparent issue. In the Earth's frame, a muon with energy E_μ and rest mass $m_\mu \approx 105.66$ MeV/c^2 [53] has a Lorentz factor $\gamma = E_\mu/(m_\mu c^2)$, and its lifetime is dilated to $\gamma\tau_0$. The corresponding mean decay length becomes $L_\mu = \gamma\beta c\tau_0 \approx \gamma c\tau_0$ (for $\beta \approx 1$), or equivalently

$$L_\mu \approx 6.2 \text{ km} \times \left(\frac{E_\mu}{\text{GeV}}\right). \tag{1.28}$$

For a muon with $E_\mu = 1$ GeV, $\gamma \approx 9.5$ and $L_\mu \approx 6.2$ km, so a substantial fraction survive the trip to the surface. At $E_\mu = 10$ GeV, $L_\mu \approx 62$ km, so the decay length far exceeds the atmospheric thickness and most muons reach the ground before decaying.

Since most electron neutrinos come from muon decay, high energy muons that reach the ground do not contribute their ν_e and second ν_μ to the atmospheric flux. Therefore, at high energies ($\gtrsim$ 100 GeV), the flavor ratio increases above 2, while at low energies the 2:1 ratio is preserved [40].

Figure 1.6 shows the calculated flavor ratio as a function of energy from three independent flux calculations. The ratio is approximately 2 below 1 GeV, as expected from the decay chain in Section 1.4.3. Above this energy, the ratio increases as muons increasingly reach the ground before decaying. The close agreement between models demonstrates that this ratio is well constrained, with uncertainties of only a few percent in the GeV range [16, 17, 46].

1.4.5 Production Height and Zenith Angle Dependence

Atmospheric neutrinos have a production height distribution that depends on energy and zenith angle. For GeV-scale neutrinos, the production height peaks at around 15–20 km [45]. As energy increases, mesons penetrate deeper into the atmosphere before decaying, shifting the mean production height downward and broadening the distribution [45]. At even higher energies, muon and pion decays become negligible and kaon decays dominate.

For a spherical Earth with radius $R_\oplus$, the baseline is given by:

$$L = \sqrt{(R_\oplus + h)^2 - R_\oplus^2 \sin^2\theta} - R_\oplus \cos\theta \tag{1.29}$$

where θ_z is the zenith angle, defined such that $\cos\theta_z = 1$ corresponds to downward-going and $\cos\theta_z = -1$ to upward-going neutrinos. While $L \approx h$ for vertically down-going neutrinos, small variations in h can produce large changes in L for near horizontal trajectories, since the trajectory is essentially tangent to the Earth's surface.

If neutrinos did not oscillate, the flux would be symmetric: $\Phi(\theta) = \Phi(\pi - \theta)$, since cosmic rays arrive isotropically and the atmosphere is uniform on average. Instead, $N_{up}(E) \neq N_{down}(E)$ because upward-going neutrinos travel baselines up to $\sim$12,700 km, sufficient for significant oscillation, while downward-going neutrinos travel only $\sim$15 km. The observation of this asymmetry by Super-Kamiokande provided definitive evidence for neutrino oscillations [38].

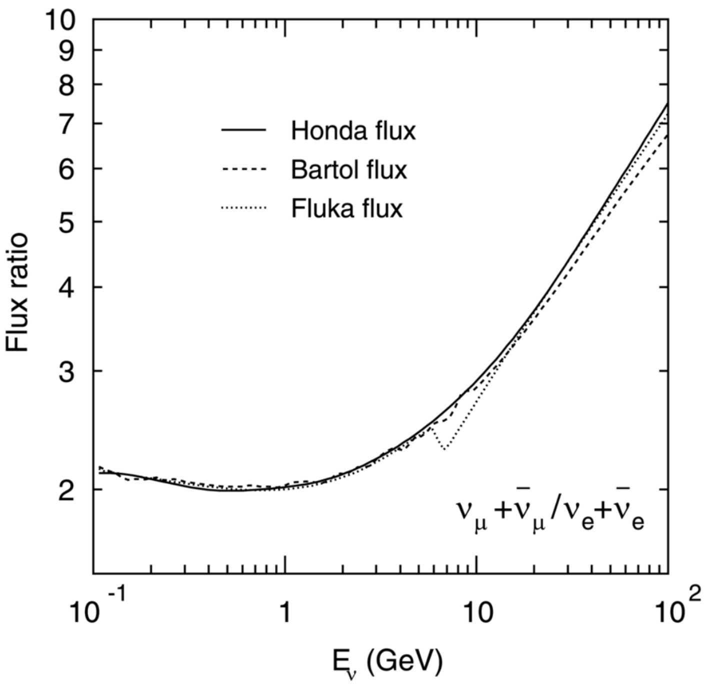

Figure 1.6: Calculated $(\nu_\mu + \bar{\nu}_\mu)/(\nu_e + \bar{\nu}_e)$ ratio of the atmospheric neutrino flux as a function of neutrino energy, using three different flux model calculations. Figure from [48].

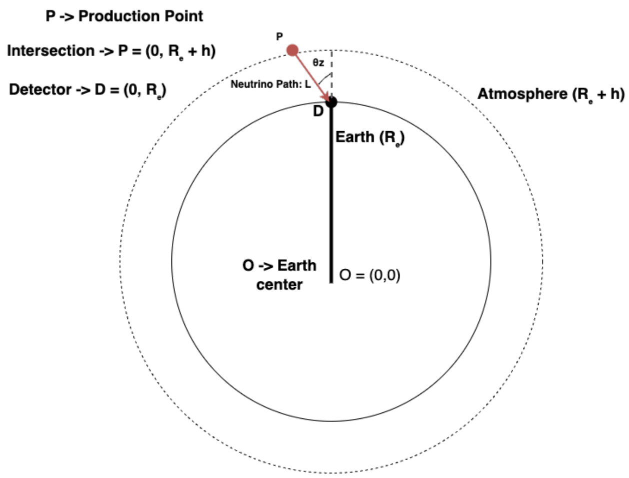

Figure 1.7: Geometry of atmospheric neutrino production. A neutrino produced at height h above the Earth's surface travels a baseline L to reach the underground detector, depending on the zenith angle θ.

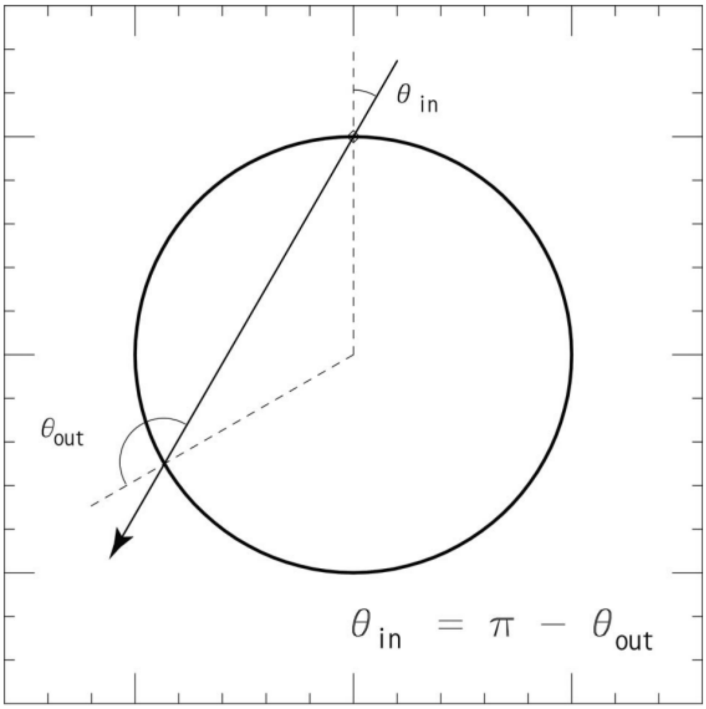

Figure 1.8: Geometry of an atmospheric neutrino trajectory through the Earth. The zenith angle at production θ_{in} and at detection θ_{out} satisfies $\theta_{\text{in}} = \pi - \theta_{\text{out}}$, reflecting the up-down symmetry of atmospheric neutrino production. Figure from [48].

1.4.6 Atmospheric Neutrino Flux Modeling

Accurate neutrino flux modeling requires knowledge of the incoming cosmic ray spectrum, hadronic interaction physics ($p + N \to \pi, K$, etc.), and meson decay kinematics. The dominant uncertainty comes from the hadronic cross sections, particularly for forward meson production, where secondary pions and kaons emerge at small angles relative to the incoming cosmic ray. This regime is difficult to measure at accelerator experiments for three reasons: the forward direction coincides with the beam pipe, so there is a geometric blind spot where detectors cannot be placed; particle identification between pions, kaons, and protons becomes increasingly difficult at high forward momentum; and most existing data are taken on proton or carbon targets rather than the nitrogen and oxygen nuclei that make up the atmosphere, requiring nuclear extrapolations. Forward production dominates the atmospheric neutrino flux because forward going mesons carry the largest fraction of the parent cosmic ray's energy, so uncertainties in this regime propagate directly into the predicted flux.

Several Monte Carlo flux calculations exist, including Honda, FLUKA, and Bartol, each employing different hadronic interaction models (DPMJET, FLUKA's internal model, and TARGET, respectively). For this analysis we use the Honda flux [45], the standard reference in oscillation physics experiments, which combines the DPMJET-III hadronic interaction model (above 32 GeV) and the JAM nuclear interaction model (below 32 GeV) with the NRLMSISE-00 atmospheric model, tuned against atmospheric muon measurements. The calculation accounts for all combinations of projectile target collisions: while most of the flux comes from p+N and p+O, helium and heavier nuclei also contribute. For the primary cosmic ray energies relevant to DUNE's atmospheric neutrino sample (approximately 1 to a few hundred GeV, producing neutrinos in the 100 MeV to 10 GeV range [32]), Honda is the standard reference used by Super-Kamiokande, IceCube, and DUNE.

1.4.7 Oscillation Signatures in Atmospheric Neutrinos

As noted in Section 1.2.1, in the two-flavor approximation, the oscillation probability is

$$P_{\alpha\to\beta} = \sin^2(2\theta)\,\sin^2\left(1.27\,\frac{\Delta m^2\,(\mathrm{eV}^2)\,L\,(\mathrm{km})}{E_\nu\,(\mathrm{GeV})}\right), \tag{1.30}$$

[41]. The oscillation probability depends on the dimensionless phase $\frac{\Delta m^2 L}{E}$, a general property of neutrino oscillations.

To leading order, atmospheric neutrino oscillations are well described by two flavor ν_μ disappearance driven by Δm^2_{32} and θ_{23}, with three flavor and matter effects providing subleading corrections. As shown in Table 1.2, atmospheric neutrinos span baselines of $L =$ 15-12,700 km and energies $E =$ 0.1-100 GeV [41].

This broad L/E coverage allows direct observation of the oscillation probability's dependence on L/E. In 1998, Super-Kamiokande demonstrated this by measuring the zenith angle dependence of atmospheric ν_μ events [38]: downward going neutrinos ($L \sim 15$ km) showed no deficit, while upward going neutrinos ($L \sim 10{,}000$ km) exhibited a clear deficit. The ν_e events showed no corresponding deficit, indicating the dominant channel is $\nu_\mu \to \nu_\tau$.

Because upward going neutrinos traverse the Earth's mantle and core, they undergo coherent forward scattering off electrons and nucleons. While ν_e can

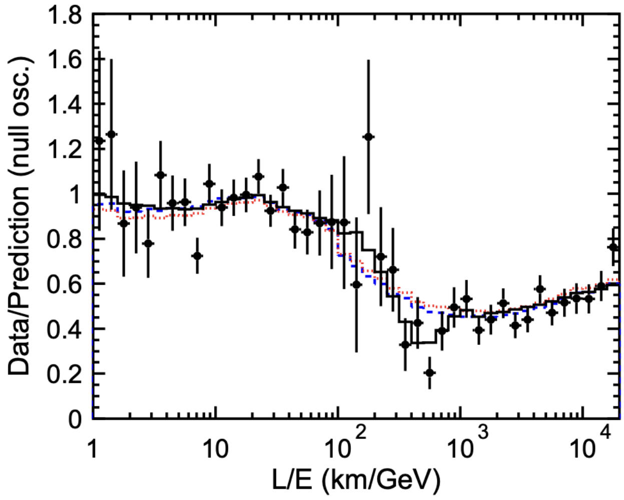

Figure 1.9: Ratio of observed to expected atmospheric ν_μ events as a function of L/E. The dip around $L/E \sim 500$ km/GeV corresponds to the first oscillation minimum. Figure from [14].

Table 1.2: Characteristic baselines, energies, and resulting L/E coverage of atmospheric neutrinos.

Quantity	Typical range	Physical origin
Baseline L	$\sim$ 15–12,700 km	Down-going to up-going trajectories
Energy E	$\sim$ 0.1–100 GeV	Atmospheric production spectrum
L/E	$\sim 10^{-1}$–10^{5} km/GeV	Derived from L and E ranges

scatter via both charged current (W exchange) and neutral current (Z exchange), ν_μ and ν_τ scatter only via neutral current. This generates an effective potential experienced only by ν_e, as discussed in Section 1.2.2.

The DUNE atmospheric neutrino program complements beam-based measurements by providing independent sensitivity to the neutrino mass ordering and δ_{CP}. In the following chapter, we describe the DUNE experiment in detail.

Chapter 2

The DUNE Experiment

The Deep Underground Neutrino Experiment (DUNE) is an international experiment hosted by Fermilab and the U.S. Department of Energy. The experiment will consist of a far detector module located underground at the Sanford Underground Research Facility (SURF) in Lead, South Dakota, approximately 1300 km from Fermilab, as well as a near detector module located at Fermilab in Batavia, Illinois. As the next generation of experiment designed to address fundamental unresolved questions in modern neutrino physics, DUNE is a large international effort consisting of over 1500 collaborators from more than 200 institutions [4].

DUNE will be the largest long-baseline neutrino experiment ever built, with a baseline distance of 1300 km from Fermilab to SURF. While DUNE is currently under construction and has experienced numerous delays and cost growth, Phase I beam operations are expected to begin in the early 2030s [4]. The experiment, given its scale, has numerous benchmarked phases, making it effectively a staged experiment.

Phase I, currently under construction, involves:

- Two far detector modules: FD1 (horizontal drift) and FD2 (vertical drift), with cryostats being installed at SURF as of 2026–2027,
- A subset of the near detector complex,
- A 1.2 MW beam from the PIP-II upgraded accelerator at Fermilab.

Phase II is planned but not yet fully funded or deployed. This phase consists of a third and fourth far detector module, the full near detector suite, and an upgrade of the neutrino beam to 2.4 MW. While there is currently no firm completion date, design decisions are expected within the next few years. The full four-module far detector is the ultimate goal, but meaningful physics, including the mass ordering determination and CP violation searches, can be extracted with just the first two modules.

DUNE seeks to address: precision measurement of the CP-violating leptonic mixing phase δ_{CP}, the neutrino mass eigenstate ordering (hierarchy), precision measurements of oscillation parameters such as θ_{23} (the atmospheric mixing angle) and its octant, beyond-the-Standard-Model physics searches, supernova neutrino detection, and proton decay searches. The DUNE experiment will integrate the world's most intense neutrino beam with a deep underground, state-of-the-art far detector using novel liquid argon detection technology.

In this chapter, we outline the physics goals, the Long-Baseline Neutrino Facility (LBNF) beam, the near detector complex, the far detector, event reconstruction, and the atmospheric neutrino capabilities of DUNE.

2.1 Physics Goals

2.1.1 CP Violation

The search for leptonic CP violation is one of the central goals of the DUNE physics program. CP violation in the lepton sector is of particular interest because it may help explain the observed matter-antimatter asymmetry of the universe through leptogenesis [4]. DUNE's sensitivity to δ_{CP} emerges from comparing $\nu_\mu \to \nu_e$ and $\bar{\nu}_\mu \to \bar{\nu}_e$ appearance rates in the long-baseline beam; a difference between these rates beyond what is expected from matter effects alone would indicate CP violation.

The 2023 Particle Physics Project Prioritization Panel (P5) identified DUNE Phase II as the "definitive long-baseline neutrino oscillation experiment of its kind," recommending early implementation of an accelerator upgrade to enable 2.1 MW beam operation, together with the deployment of a third far detector module and an upgraded near detector complex [13]. Under this plan, DUNE is expected to reach the Phase I design exposure of 120 kt·MW·yr by the mid-2030s and 600 kt·MW·yr by the mid-2040s [13]. Detailed sensitivity projections based on a full simulation and reconstruction of the DUNE far and near detectors are presented in Ref. [5]: DUNE has the potential to observe CP violation in the neutrino sector at 3σ (5σ) after an exposure of 5 (10) years, for 50% of all δ_{CP} values [5].

2.1.2 Mass Ordering Determination

One of the primary goals of DUNE is to determine the neutrino mass ordering, determining the sign of $\Delta m^2_{31} \equiv m^2_3 - m^2_1$. As discussed in Chapter 1, the long 1300 km baseline and the resulting matter effects provide DUNE with strong sensitivity to the mass ordering for any value of δ_{CP}.

2.1.3 Precision Oscillation Measurements

Another main goal of DUNE is the precision measurement of the mixing angle θ_{23} and determination of which octant the mixing angle is in [4]. DUNE is also designed to test the three flavor oscillation paradigm itself, for example through searches for non-unitarity of the PMNS matrix, sterile neutrinos, and non-standard neutrino interactions[4].

2.1.4 BSM Searches, Supernova Neutrinos, and Nucleon Decay

DUNE is also uniquely positioned to search for supernova neutrino bursts (SNBs) and for proton decay, a form of nucleon decay as of yet unobserved. The large underground liquid argon far detector provides an ideal environment for these searches due to its massive fiducial volume and low background rates at depth.

2.2 LBNF Beam

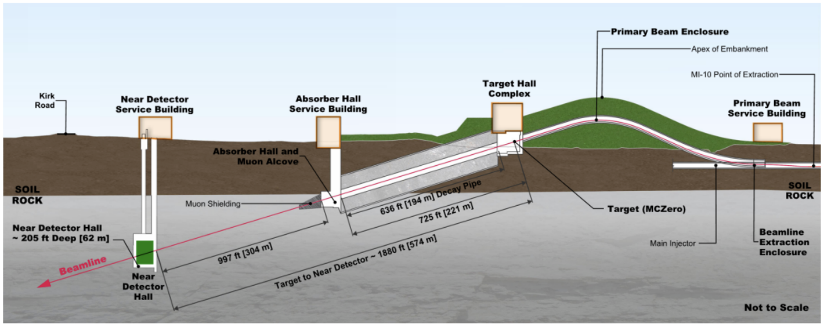

Figure 2.1: Schematic of the LBNF beamline and near detector complex at Fermilab. Figure from [4].

2.2.1 Fermilab Accelerator Complex

The Fermilab accelerator complex is a world-class particle beam facility. All sensitivities discussed in this thesis, excluding atmospheric neutrino results (as atmospheric neutrinos are not generated via the neutrino beam), are calculated assuming a 1.2 MW neutrino beam with a corresponding rate of protons on target (POT) per year on the order of 1.1×10^{21} POT, assuming a 56% LBNF beam uptime [4].

The PIP-II upgrade of the Fermilab accelerator complex will deliver between 1.0 MW and 1.2 MW of proton beam power from Fermilab's Main Injector in the energy range of 60–120 GeV at the start of DUNE operations, and provide a platform for extending beam power to greater than 2 MW [4]. There is a further planned upgrade that would bring the accelerator complex beam to provide up to 2.4 MW of beam power.

2.2.2 Beam Production and Energy Spectrum

The protons from the Main Injector strike a graphite target, producing pions and kaons. These secondary particles are focused by magnetic horns into a decay pipe, where they decay primarily into ν_μ (or $\bar{\nu}_\mu$ in antineutrino beam mode) [4].

While the energy spectrum peaks around 2.5 GeV, the neutrino beam covers a broad range of energies, allowing DUNE to be sensitive to both the first and second oscillation maxima [4].

2.2.3 The 1300 km Baseline

The long baseline of 1300 km is a physics driven choice: the first oscillation maximum occurs at approximately 2.5 GeV for Δm^2_{32}, matching the beam's

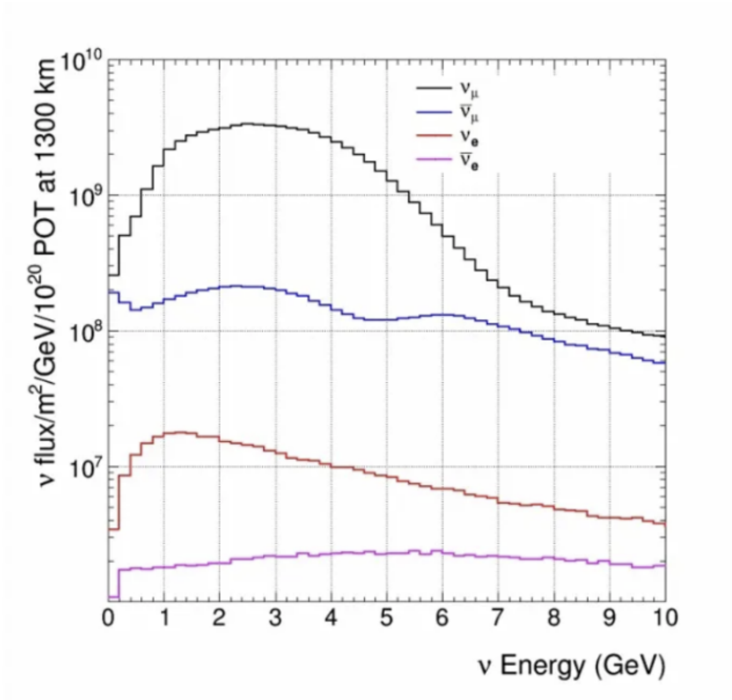

Figure 2.2: The LBNF neutrino flux as a function of energy at the far detector location. Figure from [6].

peak energy. Figure 2.3 shows the $\nu_\mu \to \nu_e$ appearance probability at this baseline as a function of neutrino energy. Given this baseline and the fact that the beam traverses through the Earth from Fermilab (Batavia, IL) to SURF (Lead, SD), matter effects become significant enough to enable robust sensitivity to both the mass ordering and δ_{CP} [28].

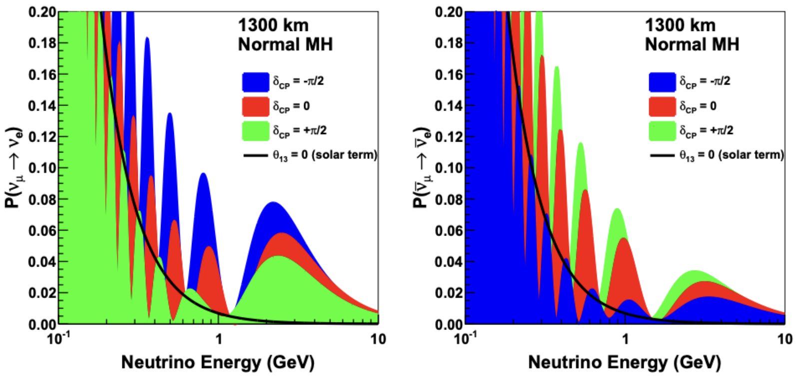

Figure 2.3: The $\nu_\mu \to \nu_e$ appearance probability at a baseline of 1300 km as a function of neutrino energy for neutrinos (left) and antineutrinos (right), for several values of δ_{CP} assuming normal ordering. The first oscillation maximum near 2.5 GeV coincides with the LBNF beam peak energy. Figure from [28].

2.3 Far Detector

The DUNE far detector is located at the Sanford Underground Research Facility (SURF) in Lead, South Dakota, at a depth of 4850 feet (approximately 1.5 km) underground. The rock overburden provides roughly 4300 meters water equivalent of shielding, reducing the cosmic ray muon rate to a level that enables rare event searches and atmospheric neutrino measurements [4].

The far detector consists of four separate LArTPC modules, each containing a fiducial mass of approximately 10 kt of liquid argon, for a total fiducial mass of 40 kt. The large mass is essential for accumulating sufficient statistics from the low-rate neutrino beam at a baseline of 1285 km, as well as for detecting atmospheric neutrinos and potential nucleon decay or supernova neutrino burst events [4].

2.3.1 LArTPC Operating Principle

Liquid argon time projection chambers detect charged particles through the ionization they produce as they traverse the liquid argon volume. When a charged particle passes through the argon, it ionizes argon atoms along its trajectory, liberating electrons. A uniform electric field, maintained by a cathode at high voltage and an anode at the opposite side of the detector, drifts these ionization electrons toward the anode. Scintillation light is also produced promptly during the ionization process and is detected by a photon detection system, providing the event start time (t_0).

At the anode, the drifting electrons induce signals on sensing elements. In the single-phase (horizontal drift) design adopted for the first far detector module (FD1), these consist of three planes of sense wires oriented at different angles. The first two wire planes (induction planes) record bipolar signals as charge drifts past, while the third plane (collection plane) collects the charge, producing a unipolar signal. By combining the wire signals with the drift time measured relative to t_0, the detector reconstructs three-dimensional images of particle trajectories with millimeter-scale spatial resolution [4].

The second far detector module (FD2) employs a vertical drift design, in which the electric field is oriented vertically and the charge is collected on horizontal charge-readout planes at the top and bottom of the module. This design allows for a longer maximum drift distance and a simplified detector construction [4].

2.4 Near Detector Complex

The DUNE near detector (ND) complex is located at Fermilab, approximately 574 m downstream of the neutrino production target. Its primary role is to constrain the neutrino flux, cross-section, and detector response systematic uncertainties that would otherwise limit the sensitivity of oscillation measurements at the far detector. As shown in Figure 2.4, the ND complex consists of three detector systems: ND-LAr, the Temporary Muon Spectrometer (TMS), and the System for on-Axis Neutrino Detection (SAND) [29].

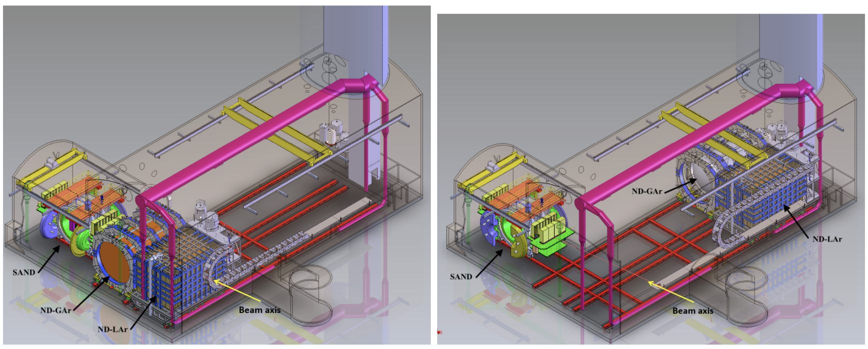

Figure 2.4: Schematic of the DUNE near detector hall in the on-axis configuration (left) and with ND-LAr and ND-GAr in an off-axis configuration (right), illustrating the DUNE-PRISM concept. SAND remains on the beam axis in both configurations. The muon spectrometer shown here is ND-GAr, which will replace TMS in Phase II. Figure from [29].

2.4.1 ND-LAr

ND-LAr is a modular liquid argon time projection chamber (LArTPC) that serves as the primary detector for measuring neutrino interactions in the near detector complex. Because it uses the same target nucleus (argon) as the far detector, cross section and detector response systematic uncertainties are highly correlated between the two detectors, and ND-LAr measurements provide strong constraints on the combined flux and interaction model used in the oscillation analysis [29]. As shown in Figure 2.5, ND-LAr consists of 35 optically isolated modules arranged in a 7×5 array, each module approximately $1 \times 1 \times 3$ m^3 and divided into two LArTPCs by a central cathode. ND-LAr employs a pixelated charge readout rather than the projective wire readout used in the far detector, providing direct three dimensional imaging of charge deposits. This avoids the wire plane matching ambiguities ("ghost tracks") that arise when multiple particles traverse a wire based LArTPC simultaneously, which is particularly important at the near detector site, where the high beam intensity produces many overlapping neutrino interactions per spill [29].

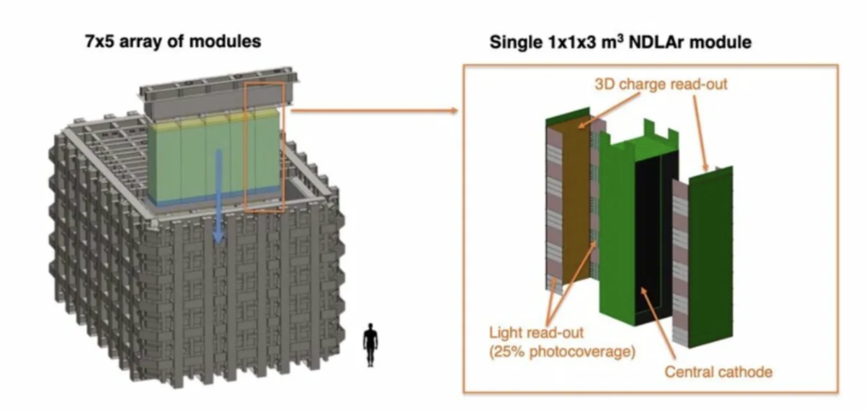

Figure 2.5: Left: the full ND-LAr detector, composed of a 7×5 array of individual modules. Right: a single $1 \times 1 \times 3$ m^3 ND-LAr module, divided into two optically isolated LArTPCs by a central cathode. Ionization electrons drift toward pixelated charge readout anode planes on opposite faces, while scintillation light is collected by light readout modules on the remaining two faces. Figure from [29].

A key feature of the ND-LAr design is the DUNE Precision Reaction Independent Spectrum Measurement (DUNE-PRISM) capability. By moving the detector off the beam axis, DUNE-PRISM samples different neutrino energy spectra at different off axis positions. This allows the near detector to effectively map out the relationship between true and reconstructed neutrino energy, reducing flux and cross section modeling uncertainties [29].

2.4.2 TMS

The Temporary Muon Spectrometer (TMS) is a magnetized detector located directly downstream of ND-LAr. It consists of alternating layers of steel absorber plates and scintillator strip detector planes. The primary purpose of TMS is to measure the charge sign and momentum of muons that exit ND-LAr, enabling separation of ν_μ and $\bar{\nu}_\mu$ charged-current interactions. This charge identification is essential for the CP violation measurement, which requires independent constraints on neutrino and antineutrino cross sections [29].

2.4.3 SAND

The System for on-Axis Neutrino Detection (SAND) remains on-axis at all times, monitoring the neutrino beam flux continuously even as ND-LAr and TMS move off-axis for PRISM measurements. SAND reuses the superconducting magnet from the KLOE experiment at LNF and incorporates a tracking system and electromagnetic calorimeter. By providing a persistent on-axis flux measurement, SAND ensures that any time-dependent variations in the beam are tracked and accounted for in the oscillation analysis [29].

2.5 Event Reconstruction

2.5.1 Particle Identification

Particle interactions in the DUNE far detector produce two fundamental topologies: tracks and showers. The energy loss per unit length, dE/dx, characterizes the energy deposited by a charged particle as it travels through the LArTPC. Different particles have characteristic dE/dx profiles, making it the main particle identification tool [4, 12]. This dE/dx based particle identification is critical for distinguishing ν_e and ν_μ charged current interactions, which is essential for the oscillation measurements discussed in Chapter 5.

Muons pass through the liquid argon as minimum ionizing particles (MIPs), with dE/dx roughly constant at $\sim$2 MeV/cm, typically exiting the detector before slowing enough to produce a Bragg peak. The Bragg peak is a sharp spike in energy deposition occurring immediately before a particle comes to rest. As the particle slows, it spends more time near each argon atom, transferring more energy per unit length. Protons exhibit this feature prominently, with dE/dx increasing dramatically in the last few centimeters of their track. Charged pions resemble muons at high energies but can undergo hadronic interactions with argon nuclei, producing a burst of secondary particles and a characteristic kink in the track. If a pion stops without interacting, it produces a Bragg peak similar to a proton, making pion identification challenging.

When electrons enter liquid argon, they immediately begin to lose energy by emitting bremsstrahlung photons (high-energy photons radiated by the decelerating electron as it interacts with the electric fields of argon nuclei). These photons then produce e^+e^- pairs, which radiate more bremsstrahlung photons, amplifying the cascade. At the beginning of this shower, before any branching occurs, one observes a single electron track with dE/dx at the MIP level of $\sim$2 MeV/cm, giving an important tagging characteristic for electron identification. Critically, the shower is completely connected to the interaction vertex with no empty space. Photons, on the other hand, are electrically neutral and travel invisibly for several centimeters before pair-producing into an e^+e^- pair and initiating a shower. Because the photon shower begins with two particles rather than one, the dE/dx at the shower start is $\sim$4 MeV/cm (double MIP). These two features, the presence or absence of a gap, and the single- vs. double-MIP dE/dx at the shower start, together provide the primary handles for e/γ

separation in a LArTPC [4]. This separation is important for rejecting neutral current π^0 backgrounds, in which a $\pi^0 \to \gamma\gamma$ decay can mimic a ν_e interaction.

The Pandora multi-algorithm pattern recognition framework [50] is the primary event reconstruction software used by LArTPC detectors in experiments such as MicroBooNE, ProtoDUNE-SP, SBND, ICARUS, and DUNE. Pandora reconstructs neutrino interaction events by building 3D tracks and showers from raw wire signals through a chain of clustering and matching algorithms [31]. CNN-based classifiers complement this by providing hit-level track/shower discrimination [30].

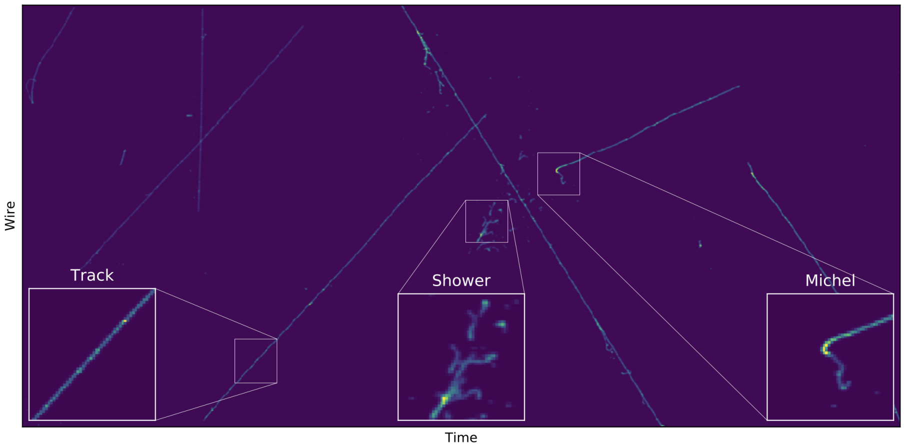

Figure 2.6: Examples of track, shower, and Michel electron topologies in a simulated ProtoDUNE-SP event, shown as CNN (Convolutional Neural Network) input patches. The track topology shows a long, straight energy deposition characteristic of a muon, while the shower topology shows the branching cascade structure of an electromagnetic shower. Figure from [30].

2.5.2 Interaction Classification

The DUNE oscillation analysis depends on correctly sorting events into key categories. In ν_μ charged-current (CC) interactions, a neutrino exchanges a W boson with an argon nucleus, producing a muon and hadronic activity at the interaction vertex. This is the disappearance channel. In ν_e CC interactions, an electron is produced instead, yielding an electromagnetic shower plus hadronic activity at the vertex. This is the appearance channel. In neutral-current (NC) interactions, a neutrino exchanges a Z boson, producing only hadronic activity with no charged lepton, so the neutrino flavor cannot be determined. NC interactions in which a π^0 is produced constitute the main misidentification background: the π^0 decays immediately into two photons, and the resulting showers can mimic a ν_e CC event, particularly when one photon is low-energy or the two showers overlap. The e/γ separation capability described above, the conversion gap and single- vs. double-MIP dE/dx, is the primary tool for rejecting this background [4].

2.6 Atmospheric Neutrinos at DUNE Far Detector

The DUNE far detector is well suited to observe atmospheric neutrinos. The full 40 kt fiducial mass of the far detector modules provides sufficient statistics despite the relatively low atmospheric neutrino flux. Assuming neutrino oscillations, the expected event rates for a 40 kt detector are approximately 1600 fully contained electron-like events per year, 2400 fully contained muon-like events per year, and 790 partially contained muon-like events per year [4]. Fully contained events are those in which the final-state lepton is tracked until it stops or reaches the boundary of the active volume. Partially contained events are those in which a final-state muon exits the detector before stopping. The deep underground location of the far detector at SURF, at a depth of 4850 ft, sufficiently suppresses the cosmic ray muon background to enable a high-purity atmospheric neutrino sample [4]. Furthermore, atmospheric data collection can begin as soon as the detector modules are filled with liquid argon, independent of the beam commissioning timeline. The LArTPC imaging and calorimetry described in Sections 2.3.1 and 2.5.1 enable lepton identification and energy reconstruction, while the particle identification techniques discussed in Section 2.5.1 allow for discrimination between ν_e-selected and ν_μ-selected event samples based on the distinct topologies they exhibit in the detector.

One advantage of atmospheric neutrinos relative to beam neutrinos is the broad and continuous coverage in L/E, ranging from $\sim 10^{-1}$ to $\sim 10^5$ km/GeV (see Table 1.2), compared to the single fixed baseline of 1300 km provided by the LBNF beam. Up-going atmospheric neutrinos traverse the Earth's mantle and core, experiencing MSW matter effects (Section 1.2.2) that depend on the sign of Δm^2_{31}. The MSW resonance enhances the oscillation probability for electron neutrinos at energies of a few GeV if the mass ordering is normal ($\Delta m^2_{32} > 0$), and for electron antineutrinos if the mass ordering is inverted ($\Delta m^2_{32} < 0$) [4]. This asymmetry between the neutrino and antineutrino channels provides sensitivity to the neutrino mass ordering that is independent of the beam program. Because both neutrinos and antineutrinos arrive simultaneously in the atmospheric flux, no beam polarity switching is required. Additionally, the broad range of baselines and energies provides complementary sensitivity to the θ_{23} octant and to Δm^2_{32} beyond what the beam measurement alone can achieve.

For the analysis presented in Chapter 3, events are classified into ν_e-selected and ν_μ-selected samples using the reconstruction and particle identification described in Section 2.5. These samples are binned in reconstructed energy and $\cos\theta_Z$ into 260 analysis bins, with the binning scheme detailed in Section 3.4.1. The sensitivity to oscillation parameters is extracted via profile likelihood methods using the GUNDAM fitting framework. The statistical methods, fitting framework, and full analysis configuration are presented in Chapter 3.

Chapter 3

The GUNDAM Framework and Analysis

3.1 Overview

Chapters 1 and 2 developed the physics of neutrino oscillations and the DUNE experimental apparatus. This chapter combines them into an analysis: a procedure for extracting oscillation parameters from DUNE's atmospheric neutrino sample.

An oscillation analysis estimates the values of the PMNS oscillation parameters $\sin^2\theta_{12}$, $\sin^2\theta_{13}$, $\sin^2\theta_{23}$, Δm^2_{21}, Δm^2_{32}, and δ_{CP} by comparing a predicted event distribution to an observed one. In this analysis, events are binned in reconstructed neutrino energy and reconstructed cosine of the zenith angle. The prediction is assembled from the atmospheric neutrino flux (Section 1.4.6), the three flavor oscillation probabilities including matter effects (Section 1.2.2), the neutrino argon cross sections described by GENIE, and the LArTPC detector response (Section 2.5). The extracted parameter values are those for which the prediction best matches the data.

The parameters of the fit divide into two groups. Parameters of interest are the physics quantities the analysis aims to measure: in this thesis, $\sin^2\theta_{23}$, Δm^2_{32}, and δ_{CP}. The sign of Δm^2_{32}, which determines the neutrino mass ordering, is also a parameter of interest and is extracted through a dedicated hypothesis test in Section 5.3. Nuisance parameters affect the prediction but are not themselves the object of the measurement: the remaining oscillation parameters, which are constrained more tightly by other experiments, and the systematic parameters representing flux, cross section, and detector uncertainties. Nuisance parameters are profiled during the fit.

A fully frequentist analysis treats all parameters as fixed but unknown and constructs confidence intervals from the frequency with which a test statistic would exceed a given threshold under repeated experiments. A fully Bayesian analysis treats parameters as random variables with prior probability distributions and reports posterior distributions after conditioning on the data. This thesis adopts a semi-frequentist approach, standard in neutrino oscillation analyses at T2K [3], NOvA [7], and DUNE [4]: the parameters of interest enter the likelihood without priors, while nuisance parameters are constrained by

Gaussian prior penalty terms derived from external measurements. This avoids specifying priors on δ_{CP} and the mass ordering, which is contested, while still incorporating external constraints on the other oscillation parameters.

3.2 Statistical Framework

3.2.1 Systematic Uncertainties

Systematic uncertainties in a DUNE atmospheric neutrino analysis fall into three broad categories. Flux uncertainties arise from the modeling of the atmospheric neutrino flux, including the primary cosmic ray spectrum, hadronic interaction models, and the resulting neutrino energy and angular distributions as computed by Honda et al. [46]. Cross section uncertainties reflect the limited knowledge of neutrino nucleus interaction rates and final state topologies, as modeled by generators such as GENIE [11]. Detector uncertainties account for the response of the liquid argon TPC, including energy scale, particle identification efficiency, and reconstruction biases.

The sensitivity studies presented in this thesis do not include systematic uncertainties and therefore represent a statistical only sensitivity. The GUNDAM framework supports their inclusion through prior covariance matrices, and a comprehensive treatment incorporating the full suite of DUNE systematic uncertainties is deferred to future work within the collaboration. The prior configuration used in this analysis is detailed in Section 3.4.2.

3.2.2 The Extended Binned Likelihood

The likelihood $\mathcal{L}$ for the DUNE far detector is described by a Poisson distribution, since atmospheric neutrino data consists of event counts grouped into bins of reconstructed energy and zenith angle. A Poisson distribution gives the probability of observing N_i^{obs} events in the i-th bin given a prediction of N_i^{pred} events in that bin. The likelihood is

$$\mathcal{L} = \prod_{i=1}^{N_{\mathrm{bins}}} \frac{\left(N_i^{\mathrm{pred}}\right)^{N_i^{\mathrm{obs}}} e^{-N_i^{\mathrm{pred}}}}{N_i^{\mathrm{obs}}!}, \tag{3.1}$$

where N_i^{pred} depends on a set of physics and nuisance parameters.

When performing statistical analysis, one typically considers the likelihood ratio rather than the likelihood itself. The likelihood ratio normalizes Eq. (3.1) to the case where the prediction equals the observation in every bin, canceling the factorials:

$$\mathcal{L}_{\mathrm{ratio}} = \prod_{i=1}^{N_{\mathrm{bins}}} \frac{\left(N_i^{\mathrm{pred}}\right)^{N_i^{\mathrm{obs}}} e^{-N_i^{\mathrm{pred}}}}{\left(N_i^{\mathrm{obs}}\right)^{N_i^{\mathrm{obs}}} e^{-N_i^{\mathrm{obs}}}}. \tag{3.2}$$

Taking $-2\ln\mathcal{L}_{\mathrm{ratio}}$ converts the product into a sum, yielding the $\Delta\chi^2$ test statistic:

$$\Delta\chi^2_{\mathrm{stat}} = 2\sum_{i=1}^{N_{\mathrm{bins}}} \left(N_i^{\mathrm{pred}} - N_i^{\mathrm{obs}} + N_i^{\mathrm{obs}} \ln \frac{N_i^{\mathrm{obs}}}{N_i^{\mathrm{pred}}} \right). \tag{3.3}$$

The total test statistic adds a Gaussian penalty term for the nuisance parameter priors introduced in Section 3.1: $\Delta\chi^2_{\text{total}} = \Delta\chi^2_{\text{stat}} + \Delta\chi^2_{\text{sys}}$, where

$$\Delta\chi^2_{\text{sys}} = \sum_{p}^{N_{\text{par}}} \sum_{q}^{N_{\text{par}}} \Delta\theta_p \, (V_\theta^{-1})_{pq} \, \Delta\theta_q \,, \tag{3.4}$$

with $\vec{\theta}$ the set of nuisance parameters, V_θ the prior covariance matrix encoding their uncertainties and correlations, and $\Delta\theta_p = \theta_p - \theta_p^{\text{nom}}$ the deviation of parameter p from its nominal value.

In practice, Monte Carlo predictions carry their own statistical uncertainty due to finite sample sizes. The Barlow-Beeston method [15] extends the likelihood to account for this by introducing per bin scaling parameters that are solved analytically during the fit. The specific implementation used in this analysis is described in Section 3.3.1.

3.2.3 Parameter Estimation and Profiling

The oscillation parameters are estimated by finding the values that minimize $\Delta\chi^2_{\text{total}}$. This is a maximum likelihood estimation (MLE) problem: the parameter values that minimize $\Delta\chi^2$ are the same values that maximize the likelihood.

In practice, the minimization is performed using Minuit2 [47], a general purpose function minimization package. After convergence, the HESSE algorithm computes an estimate of the parameter uncertainties from the inverse of the second derivative matrix (Hessian) of $\Delta\chi^2$ at the minimum. These Hesse errors assume the likelihood surface is approximately parabolic near the minimum, corresponding to Gaussian distributed parameter uncertainties.

All sensitivity results in this thesis use profile likelihood scans. For each fixed value of the parameter θ_k, the $\Delta\chi^2$ is minimized over all remaining parameters $\vec{\eta}$. The resulting one dimensional curve,

$$\Delta\chi^2(\theta_k) = \min_{\vec{\eta}} \; \Delta\chi^2(\theta_k, \vec{\eta}) - \Delta\chi^2(\hat{\theta}_k, \hat{\vec{\eta}}) \,, \tag{3.5}$$

gives the profiled $\Delta\chi^2$ as a function of the parameter of interest, where $\hat{\theta}_k$ and $\hat{\vec{\eta}}$ are the global best fit values. By Wilks' theorem [62], this quantity follows a χ^2 distribution with one degree of freedom, so that confidence intervals correspond to $\Delta\chi^2 = 1.00$ (1σ), 4.00 (2σ), and 9.00 (3σ).

The same profiling procedure extends to two parameters of interest. In a two dimensional profile scan, both parameters are fixed on a grid of values and $\Delta\chi^2$ is minimized over all remaining parameters at each grid point. The resulting surface is projected as a contour plot, with confidence regions defined by the two degree of freedom thresholds: $\Delta\chi^2 = 2.30$ (1σ), 6.18 (2σ), and 11.83 (3σ). The two dimensional contour results presented in Section 5.2 use this profiled approach, with Minuit minimizing over the remaining oscillation parameters at each grid point.

As a consistency check on the baseline configuration used throughout this thesis, Figure 3.1 shows one dimensional profile likelihood scans for all six oscillation parameters on the 300×300 oscillation probability grid at 400 kt·yr exposure. All six parameters produce clean curves with minima sitting at the injected truth values, and the expected sensitivity hierarchy is reproduced: the atmospheric parameters ($\sin^2\theta_{23}$, Δm^2_{32}, δ_{CP}) show local sensitivity from the atmospheric neutrino sample, while the solar sector parameters ($\sin^2\theta_{12}$, $\sin^2\theta_{13}$, Δm^2_{21}) remain dominated by their external priors. This baseline configuration is the starting point for the formal grid convergence study presented in Chapter 4 and for the sensitivity results presented in Chapter 5.

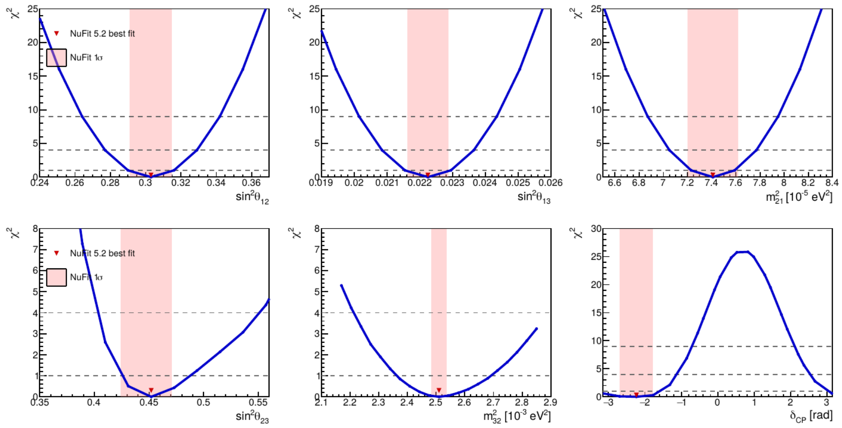

Figure 3.1: Baseline consistency check: one dimensional profile likelihood scans for all six oscillation parameters using the 300×300 oscillation probability grid at 400 kt·yr exposure. At each scan point the parameter of interest is held fixed at the scan value and the remaining five oscillation parameters are profiled over by Minuit2. The pink band shows the NuFit 5.2 1σ range and the red triangle marks the NuFit 5.2 best fit. Horizontal dashed lines at $\Delta\chi^2 = 1$, 4, and 9 correspond to the 1σ, 2σ, and 3σ intervals by Wilks' theorem. All six parameters recover clean curves centered on truth, with the atmospheric sector parameters ($\sin^2\theta_{23}$, Δm^2_{32}, δ_{CP}) showing local sensitivity from the atmospheric neutrino sample and the solar sector parameters ($\sin^2\theta_{12}$, $\sin^2\theta_{13}$, Δm^2_{21}) dominated by their external priors. The y-axis is labeled χ^2 but denotes $\Delta\chi^2$ relative to the scan minimum.

The specific implementation of profiling within the GUNDAM framework, including the minimization procedure and the construction of $\Delta\chi^2$ profiles, is described in Section 3.3.3. The one dimensional profile likelihood scans and two dimensional profiled contours for the baseline configuration are presented in Chapter 5.

3.2.4 The Asimov Fit

The sensitivity studies presented in this thesis evaluate the expected performance of DUNE before any data have been collected. This requires a well defined procedure for generating synthetic data that yields representative results. The simulated dataset, introduced by Cowan et al. [24], serves this purpose.

An Asimov fit is constructed by setting the observed bin counts exactly equal to the predicted expectation under a chosen set of true parameter values:

$$N_i^{\mathrm{obs}} = N_i^{\mathrm{pred}}(\vec{\theta}_{\mathrm{true}}) \tag{3.6}$$

for every analysis bin i. Because the data and prediction are identical at the true values, the global minimum of the test statistic returns $\Delta\chi^2 = 0$ at $\vec{\theta} = \vec{\theta}_{\text{true}}$ by construction. This removes statistical fluctuations entirely and isolates the sensitivity arising from the physics content of the analysis: the oscillation structure in the predicted spectra, the analysis binning, and the prior constraints on the parameters.

The key property of the Asimov dataset is that it yields the median expected sensitivity [24]. The test statistic $\Delta\chi^2$ is itself a random variable, because its value depends on the Poisson fluctuations in a given dataset. In a frequentist framework, one could generate many pseudo experiments drawn from Poisson statistics at the true parameter values, fit each one, and build a distribution of test statistic values. This distribution is generally skewed and bounded from below, so the mean and the median are not the same. The median is the natural summary of typical sensitivity, since half of the pseudo experiments would give a stronger result and half would give a weaker one. The result obtained from the Asimov dataset corresponds to this median. It provides a representative summary of the expected experimental sensitivity from a single fit, rather than requiring thousands of pseudo experiments to be generated and individually fitted.

This property makes the Asimov approach the standard tool for sensitivity projections in neutrino oscillation experiments, including DUNE [4], T2K [3], and NOvA [7]. It is well suited for studies requiring large numbers of fits, such as the grid convergence study in Chapter 4 (where six grid configurations are each scanned over six parameters), the two dimensional contour scans in Section 5.2 (where 2500 fits are performed per parameter pair), and the mass ordering sensitivity study in Section 5.3 (where fits are repeated at eleven exposures under four configurations).

The Asimov approach has a known limitation. Because it replaces statistical fluctuations with the expectation value, it does not capture the spread of possible experimental outcomes. In regions of parameter space where the likelihood surface is non parabolic or where discrete degeneracies exist (such as the θ_{23} octant degeneracy), the median sensitivity may not represent the full distribution of outcomes. Characterizing this spread would require toy Monte Carlo studies, which are beyond the scope of this thesis.

For this analysis, the Asimov truth values are taken from the NuFit 5.2 global fit (normal ordering, with SK atmospheric data) [35] as listed in Table 3.1. The exposure is 400 kt·yr, corresponding to ten years of data collection with the full four module DUNE far detector. As a validation of this procedure, Section 3.6 presents an Asimov fit at the nominal parameter values with Hesse error estimation enabled (Figure 3.8). The fit recovers the injected truth values exactly, and the post fit Hesse errors are compared to the NuFit 5.2 prior widths, confirming that the fitting framework is functioning correctly before proceeding to the sensitivity results in Chapter 5.

3.3 The GUNDAM Framework

GUNDAM (Generalized and Unified Neutrino Data Analysis Methods) is a binned likelihood fitting framework used across neutrino oscillation experiments, including both long baseline and atmospheric analyses. Its purpose is to compare a predicted event spectrum to an observed or Asimov dataset and extract the parameter values that best describe the data. The core framework is written in C++ and built on ROOT, and is configuration driven: the samples, parameters, and fit settings are all specified through YAML or JSON files without modifying the source code.

3.3.1 Overview and Design

The framework is organized into six core libraries, shown in Figure 3.2. The `FitterEngine` is the main driver of the fit and delegates minimization to a `MinimizerBase` backend, which can be configured for gradient based optimization (via Minuit2), Markov Chain Monte Carlo sampling, or other algorithms through a common interface. The `FitterEngine` drives the `Propagator`, which turns a set of parameter values into a predicted spectrum by applying oscillation probabilities and systematic effects to the input events. For likelihood evaluation, the `FitterEngine` interfaces with the `StatisticalInference` library through the `LikelihoodInterface`.

Event data is managed by the `SamplesManager`, where each `Sample` owns its `Event` objects and the `Histogram` that summarizes them. Parameters are organized by the `ParametersManager` into `ParameterSet` groups. Each parameter is applied to events through the `DialDictionary`, which serves as a reweighting engine: `DialCollection` objects hold the `DialBase` implementations (normalizations, tabulated weights, splines, and so on) that define how an event weight responds to a change in a given parameter.

The likelihood evaluated by the `StatisticalInference` library is a Barlow-Beeston log likelihood [15], which extends the standard Poisson likelihood to account for finite Monte Carlo statistics in each analysis bin. For a bin with observed count n_i and predicted count μ_i, GUNDAM introduces an auxiliary scaling parameter β_i that adjusts the prediction to $\hat{\mu}_i = \beta_i \, \mu_i$. The per bin contribution to $-2\ln\mathcal{L}$ is

$$-2\ln\mathcal{L}_i = 2\big(\beta_i \, \mu_i - n_i\big) + 2\, n_i \, \ln\frac{n_i}{\beta_i \, \mu_i} + \frac{(\beta_i - 1)^2}{\sigma^2_{\mathrm{rel},i}} \,, \tag{3.7}$$

where $\sigma^2_{\mathrm{rel},i} = \sigma^2_{\mathrm{MC},i}/\mu_i^2$ is the relative variance of the Monte Carlo prediction in bin i. The first two terms correspond to the Poisson log likelihood (up to a data only constant), while the final term is a Gaussian penalty that constrains β_i to remain close to unity. For each bin, β_i is solved analytically from the quadratic

$$\beta_i = \frac{-b_i + \sqrt{b_i^2 + c_i}}{2} \,, \qquad b_i = \mu_i \, \sigma^2_{\mathrm{rel},i} - 1 \,, \qquad c_i = 4\, n_i \, \sigma^2_{\mathrm{rel},i} \,. \tag{3.8}$$

The total test statistic also includes Gaussian penalty terms for systematic parameters:

$$-2\ln\mathcal{L} = \sum_i^{N_{\mathrm{bins}}} \big(-2\ln\mathcal{L}_i\big) + \sum_j^{N_{\mathrm{syst}}} \left(\frac{p_j - p_j^0}{\sigma_j}\right)^2 , \tag{3.9}$$

where p_j^0 and σ_j are the prior central value and width of systematic parameter p_j.

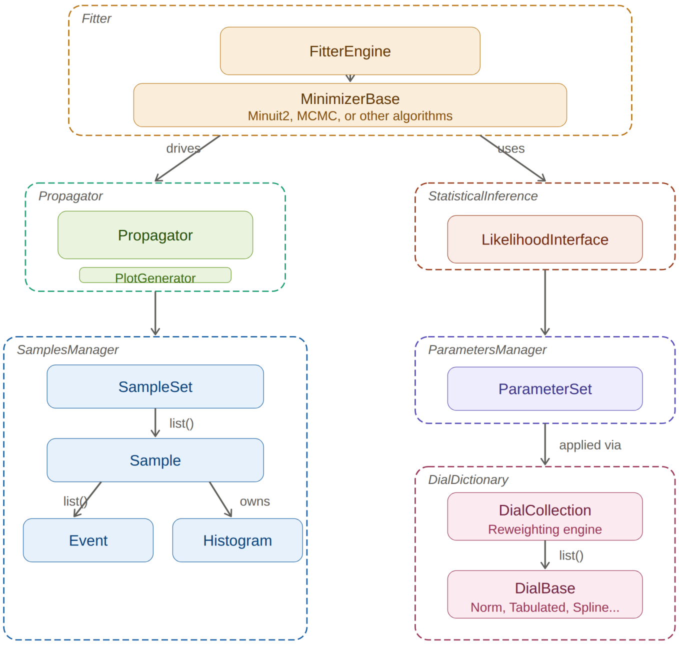

Figure 3.2: Software architecture of the GUNDAM framework, showing the six core libraries and their principal class hierarchies. The `FitterEngine` delegates to a `MinimizerBase` backend and drives the `Propagator` for spectrum prediction. The `StatisticalInference` library provides the likelihood evaluation through `LikelihoodInterface`. Event data is managed by the `SamplesManager`, while oscillation and systematic parameters are organized by the `ParametersManager` and applied to events through the `DialDictionary` reweighting engine. Class names reflect the current GUNDAM source code.

3.3.2 Oscillation Probability Calculation

Oscillation probabilities are computed by NuOscillator [61], which pre computes three flavour transition probabilities on a two dimensional grid in true neutrino energy E and cosine of the zenith angle $\cos\theta_z$. During the fit, probabilities at arbitrary $(E, \cos\theta_z)$ are obtained by interpolation from this table. The fidelity of the interpolation depends on the grid resolution, motivating the optimisation studies presented in Chapter 4.

Matter effects are incorporated using the Preliminary Reference Earth Model (PREM) [33], which provides a radially symmetric density profile of the Earth. Neutrino production heights in the atmosphere are averaged over the distribution discussed in Section 3.5.

3.3.3 Minimization and Fitting

Oscillation parameters are estimated by minimising Eq. (3.9) with Minuit2 [47]. Confidence regions are constructed from the profile log likelihood ratio

$$\Delta\chi^2(\boldsymbol{\theta}) = -2\ln\frac{\mathcal{L}(\boldsymbol{\theta},\, \hat{\hat{\boldsymbol{\eta}}})}{\mathcal{L}(\hat{\boldsymbol{\theta}},\, \hat{\boldsymbol{\eta}})}\,, \tag{3.10}$$

where $\hat{\boldsymbol{\theta}}$ and $\hat{\boldsymbol{\eta}}$ denote the global best fit values of the parameters of interest and nuisance parameters, respectively, and $\hat{\hat{\boldsymbol{\eta}}}$ is the conditional best fit of the nuisance parameters at fixed $\boldsymbol{\theta}$. By Wilks' theorem [62], $\Delta\chi^2$ follows a χ^2 distribution with degrees of freedom equal to the number of parameters of interest. The relevant thresholds are:

- One parameter: $\Delta\chi^2 < 1.00$ (1σ), < 4.00 (2σ), < 9.00 (3σ).
- Two parameters: $\Delta\chi^2 < 2.30$ (1σ), < 6.18 (2σ), < 11.83 (3σ).

3.4 Analysis Configuration

The sensitivity studies in this thesis are carried out with Gundam by performing a semi-frequentist fit using the Barlow-Beeston likelihood. GUNDAM is configured through a dedicated set of configuration files maintained in the `gudi-anoa-sen25a` repository. The configuration includes the event selections and analysis binning (Section 3.4.1) and the oscillation parameters entering the fit together with their associated priors and uncertainties (Section 3.4.2). The oscillation probability grid and the precalculation of oscillation probabilities are developed separately in Chapter 4.

The overall workflow is summarized in Figure 3.3. Generated Monte Carlo events are reweighted by the expected flux, interaction cross section, and oscillation probability before entering the fit. The Asimov fit then returns a best fit point. A profile likelihood scan is performed by an external driver that relaunches the fit at fixed values of the parameters of interest. Together, the pieces described in this section define what a single point in the $\Delta\chi^2$ profiles shown in Chapter 5 represents.

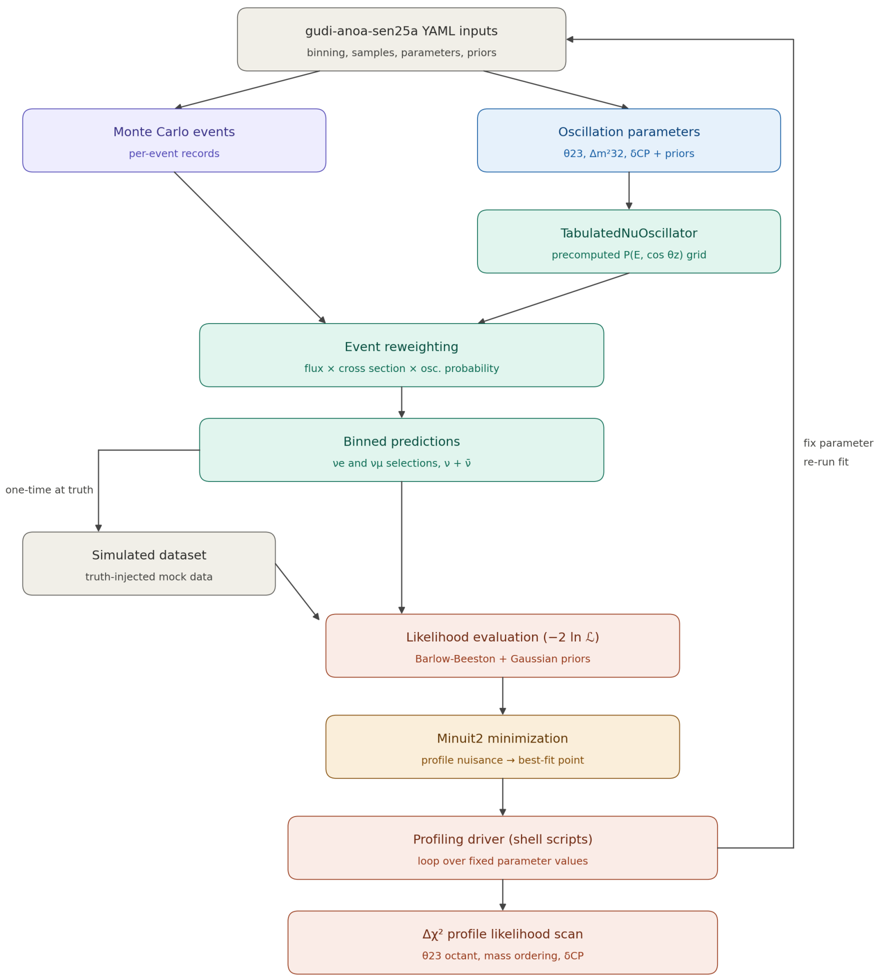

Figure 3.3: Analysis workflow used in this thesis. Configuration files in the `gudi-anoa-sen25a` repository specify the binned Monte Carlo inputs for the ν_e- and ν_μ-selected samples together with the oscillation and systematic parameter priors. Oscillation probabilities are precomputed by TabulatedNuOscillator on a two dimensional grid in true neutrino energy and $\cos\theta_z$. A single `gundamFitter` call reweights the Monte Carlo events by the flux, interaction cross section, and oscillation probability, compares the resulting binned predictions to the simulated dataset, and returns the best fit parameter values. The simulated dataset is itself produced by running the reweighting pipeline once at the injected truth values. Profile likelihood scans are produced by external shell scripts that rerun the fit many times, each time holding the parameter of interest fixed at a different value; the collected results form the $\Delta\chi^2$ profiles used in the sensitivity studies.

3.4.1 Event Selection and Analysis Binning

Events are classified into two charged current selections based on the reconstructed final state lepton: a ν_e selection and a ν_μ selection. The selection criteria are described in Section 2.5.

Each sample is binned in two dimensions: reconstructed neutrino energy and the cosine of the reconstructed zenith angle, $\cos\theta_z \in [-1, +1]$, using 10 uniform bins of width 0.2. The energy binning is sample dependent and uses non uniform bin widths chosen to match the detector resolution and the oscillation structure at each energy scale:

- ν_e **selection**: 16 reconstructed energy bins spanning 0–10 000 GeV, with edges at 0, 0.07, 0.1, 0.15, 0.22, 0.33, 0.47, 0.7, 1.1, 1.7, 2.9, 5, 9, 16, 32, 48, and 10 000 GeV.

- ν_μ **selection**: 10 reconstructed energy bins spanning 0–10 000 GeV, with edges at 0, 0.12, 0.22, 0.37, 0.63, 1.5, 4.0, 11, 19, 32, and 10 000 GeV.

This gives $16 \times 10 = 160$ bins for the ν_e sample and $10 \times 10 = 100$ bins for the ν_μ sample, for a total of 260 analysis bins. The final bin in energy acts as an overflow bin. We emphasize the distinction between the two binnings used in this analysis: the analysis bins defined above are in *reconstructed* neutrino energy and zenith angle, whereas the finer oscillation probability grid discussed in Chapter 4 is in *true* neutrino energy and zenith angle. All studies in this work use simulated datasets [24], in which the observed bin counts are set exactly equal to the predicted expectation at the true parameter values. The Monte Carlo samples correspond to an exposure of 400 kt·yr, equivalent to ten years of data collection with the full four module DUNE far detector (4×10 kt fiducial mass).

3.4.2 Oscillation Parameters Configuration in GUNDAM

The oscillation parameters used in the Asimov fits are taken from the NuFit 5.2 global fit [34, 35]. Two NuFit 5.2 configurations are used in this thesis. The fit that includes Super-Kamiokande atmospheric data is used as the baseline, and its injected truth values drive the profile likelihood scans of Section 5.1 and the two dimensional contour scans of Section 5.2. The mass ordering sensitivity study in Section 5.3 is run under both configurations: with SK atmospheric data and without it. These two NuFit fits return different best fit values for several oscillation parameters and different prior widths, and comparing them is central to the mass ordering interpretation in that section. Table 3.1 lists the injected truth values for both configurations and both orderings, and Table 3.2 lists the prior constraints for both configurations.

Table 3.1: Oscillation parameters injected in the simulated datasets used in this analysis, taken from the NuFit 5.2 global fit with and without Super-Kamiokande atmospheric data [34, 35], for both normal (NO) and inverted (IO) ordering. The with SK configuration is the baseline used in Sections 5.1 and 5.2; both configurations are used in the mass ordering study of Section 5.3. For IO, $\Delta m^2_{3\ell} \equiv \Delta m^2_{32}$ following the NuFit convention.

	With SK atm		Without SK atm		
Parameter	NO	IO	NO	IO	Unit
$\sin^2\theta_{12}$	0.303	0.303	0.303	0.303	—
$\sin^2\theta_{13}$	0.02225	0.02223	0.02203	0.02219	—
$\sin^2\theta_{23}$	0.451	0.569	0.572	0.578	—
Δm^2_{21}	7.41×10^{-5}	7.41×10^{-5}	7.41×10^{-5}	7.41×10^{-5}	eV^2
Δm^2_{32}	$+2.507 \times 10^{-3}$	-2.486×10^{-3}	$+2.511 \times 10^{-3}$	-2.498×10^{-3}	eV^2
δ_{CP}	-2.233 rad	-1.466 rad	-2.845 rad	-1.292 rad	rad

Table 3.2 lists the prior constraints applied to each oscillation parameter under both NuFit configurations. The three nuisance oscillation parameters ($\sin^2\theta_{12}$, $\sin^2\theta_{13}$, and Δm^2_{21}) receive Gaussian priors with central values and 1σ widths from the corresponding NuFit 5.2 fit [35]; the three parameters of interest ($\sin^2\theta_{23}$, Δm^2_{32}, and δ_{CP}) are assigned flat priors and are constrained entirely by the atmospheric neutrino data.

Table 3.2: Baseline prior constraints on oscillation parameters under both NuFit 5.2 configurations used in this analysis. Central values follow the corresponding NuFit 5.2 fit [35]; Gaussian 1σ widths are from the covariance matrix of each fit. The with SK configuration is used for the profile likelihood scans and two dimensional contour scans in Chapter 5; the mass ordering sensitivity study in Section 5.3 is performed under both configurations.

		Central value $\pm$ 1σ width	
Parameter	Prior type	With SK atm	Without SK atm
$\sin^2\theta_{12}$	Gaussian	0.303 ± 0.013	0.303 ± 0.012
$\sin^2\theta_{13}$	Gaussian	0.02225 ± 0.00070	0.02203 ± 0.00062
$\sin^2\theta_{23}$	Flat	—	—
Δm^2_{21} [10^{-5} eV^2]	Gaussian	7.41 ± 0.18	7.41 ± 0.21
Δm^2_{32}	Flat	—	—
δ_{CP}	Flat	—	—

This baseline configuration does not include systematic uncertainties on neutrino flux normalizations, interaction cross sections, or detector response. The fits therefore represent a statistical only sensitivity. A comprehensive treatment incorporating the full suite of DUNE systematic uncertainties is deferred to future work within the collaboration.

3.5 Production Height Studies

As discussed in Section 1.4, atmospheric neutrinos are produced at varying altitudes in the atmosphere, typically ranging from around 10 to 30 km depending on the energy and zenith angle of the parent meson. The production height enters the oscillation calculation through the neutrino path length: a neutrino produced higher in the atmosphere travels a slightly longer distance to the detector than one produced lower. Since oscillation probabilities depend on the ratio L/E, the treatment of production height can in principle affect the predicted event rates and the sensitivity to oscillation parameters.

There are two approaches to handling this in the analysis. The first is to fix the production height at a single representative value, taken here to be 20 km. In this case, the oscillation probability for each event is computed using a path length determined entirely by the zenith angle and the fixed altitude. The second approach averages the oscillation probability over a distribution of production heights for each flavor, binned in energy, zenith angle, and altitude in 2.5 km height bins. The averaging replaces the oscillation probability at each point in the energy zenith grid with a weighted sum over the height distribution.

The fixed height approach is computationally less expensive, since it avoids the additional loop over height bins when tabulating the oscillation probabilities. However, it is important to verify that this simplification does not introduce a significant bias in the predicted event rates. To quantify the effect, both configurations were run using the same 300×300 oscillation probability grid, the same NuFit 5.2 normal ordering parameters, and the same Monte Carlo samples. A simulated dataset was generated for each, and the predicted event distributions were compared at the pre fit level.

The total difference between the two approaches for the ν_μ selection is -369.3 weighted events, corresponding to a relative change of -1.294%. The deficit is systematic and negative across nearly all bins, with the largest differences appearing in the oscillation maximum region. This is expected: height averaging effectively smears the distribution of path lengths for a given zenith angle, which smooths the oscillation pattern and reduces the event rate at the peaks.

For the ν_e selection, the total difference is only -3.6 weighted events, or -0.014%, roughly 100 times smaller than the ν_μ effect. The physical explanation is that ν_e at these energies originate primarily from the $\pi \to \mu \to e$ decay chain. Because the intermediate muon travels on the order of kilometers before decaying, the effective production point of the ν_e is spread out over a range of altitudes regardless of where the pion was produced. This dilutes the sensitivity to the initial production height.

3.5.1 Energy Dependence

Figures 3.4 and 3.5 show the event rate comparison as a function of reconstructed and true neutrino energy, respectively. In both cases, the lower panels show the relative difference defined as

$$\frac{N_{\text{avg}} - N_{\text{fixed}}}{N_{\text{fixed}}} \times 100\%. \tag{3.11}$$

For ν_μ, the height averaged prediction is systematically lower across the full energy range in both reconstructed and true energy, with the relative deficit reaching roughly 1.5% in the sub GeV region where the oscillation probability varies most rapidly with L/E. The true energy distribution shows a small positive excursion at the lowest energies ($E \lesssim 0.2$ GeV), where the oscillation length becomes comparable to the height variation itself. For ν_e, the relative differences are at the level of 0.02–0.04% and show no coherent structure, consistent with the expectation that the ν_e sample is insensitive to production height.

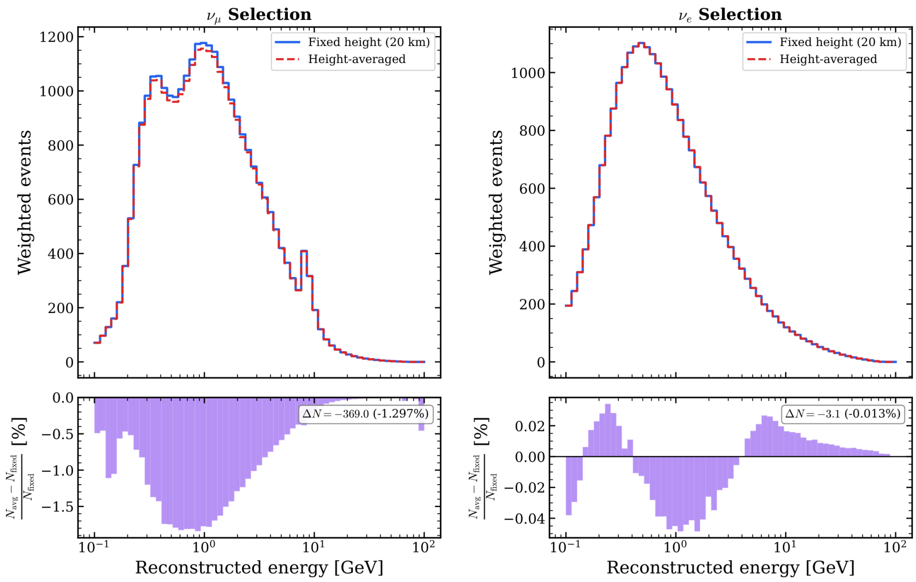

Figure 3.4: Reconstructed energy spectra for the ν_μ (left) and ν_e (right) selections comparing fixed production height and height averaged predictions. The lower panels show the relative difference. The ν_μ deficit is broad and systematic, peaking around 1.5% in the sub GeV region. The ν_e differences are negligible.

3.5.2 Zenith Dependence

Figures 3.6 and 3.7 show the event rates binned in true and reconstructed $\cos\theta_z$, respectively, where $\cos\theta_z = -1$ corresponds to up going neutrinos and $\cos\theta_z = +1$ to down going.

The ν_μ selection shows a clear pattern in both variables: the production height effect is largest for near horizontal and slightly up going directions ($\cos\theta_z \approx$ 0 to 0.3), reaching a relative difference of roughly 4% in true $\cos\theta_z$ and 2.5% in reconstructed $\cos\theta_z$. The smearing from zenith angle reconstruction softens the peak but preserves the overall shape. For steeply up going neutrinos ($\cos\theta_z < -0.5$), the effect is smaller because the total path length through the Earth is much larger than the 10 to 30 km height variation, making the fractional change in L negligible. For steeply down going neutrinos, the oscillation baseline is short enough that the oscillation probability is small regardless of the exact height.

The ν_e selection shows no significant dependence on zenith angle in either variable, with relative differences below 0.1% in all bins.

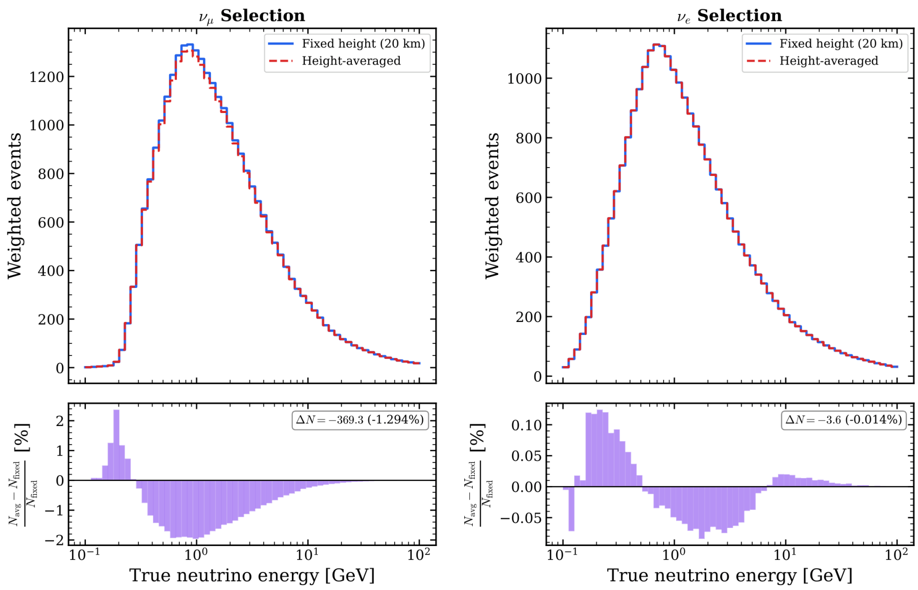

Figure 3.5: True neutrino energy spectra for the ν_μ (left) and ν_e (right) selections comparing fixed production height and height averaged predictions. The shape of the ν_μ deficit is similar to the reconstructed energy case, with the addition of a small positive excursion at the lowest energies. The ν_e differences remain negligible.

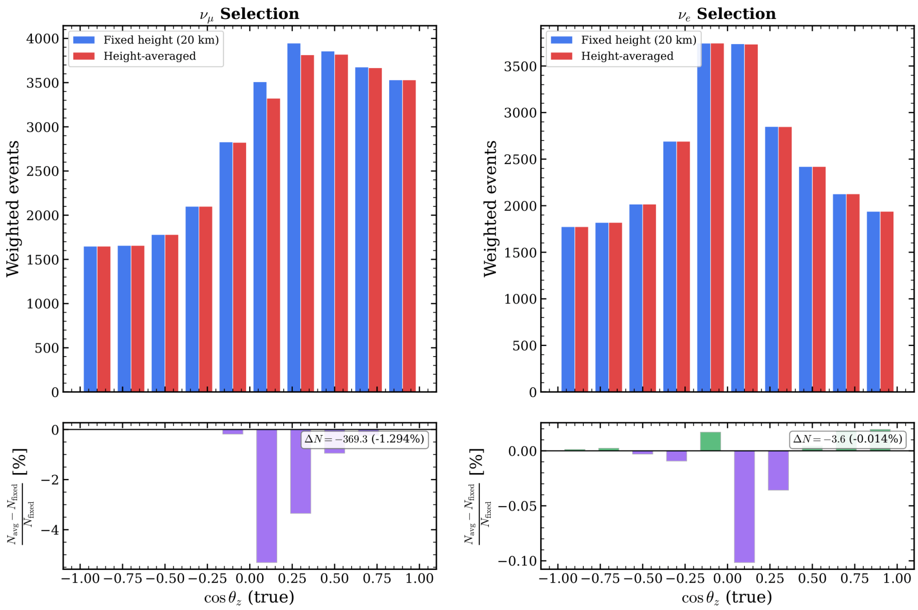

Figure 3.6: Event rates as a function of true $\cos\theta_z$ for the ν_μ (left) and ν_e (right) selections. The lower panels show the relative difference between height averaged and fixed height predictions. The ν_μ effect is largest near the horizon ($\cos\theta_z \approx 0$ to 0.3), where the production height variation is a significant fraction of the total path length.

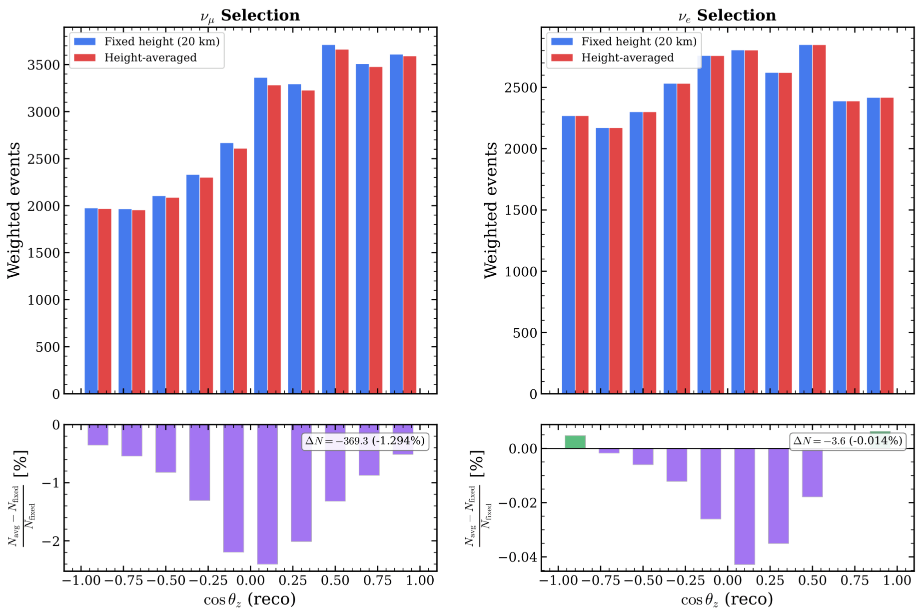

Figure 3.7: Event rates as a function of reconstructed $\cos\theta_z$ for the ν_μ (left) and ν_e (right) selections. The zenith angle reconstruction smears the sharp near horizon feature seen in the true $\cos\theta_z$ distribution but the overall pattern is preserved. The ν_e differences remain negligible.

3.5.3 Summary

The fixed production height approximation introduces a bias of at most 1.3% in the ν_μ event rate and a negligible 0.014% in the ν_e event rate. All subsequent sensitivity studies presented in this thesis, including the grid convergence studies (Section 4.2), δ_{CP} profile likelihood scans, two dimensional confidence contours, mass ordering sensitivity, and θ_{23} octant sensitivity, use the fixed production height of 20 km.

3.6 Asimov Fit

A standard validation of the fitting configuration is performed by running a single Asimov fit at the nominal parameter values with Hesse error estimation enabled. In an Asimov fit, the observed event counts are set exactly equal to the predicted counts under the nominal parameter values, such that the best fit point is guaranteed to return the injection values with a likelihood at the minimum of zero. This provides a clean test that the minimizer, oscillation probability engine, and systematic parameter handling are all functioning correctly.

The post fit parameter errors returned by the Hesse algorithm are shown in Figure 3.8. The solar sector parameters $\sin^2\theta_{12}$, $\sin^2\theta_{13}$, and Δm^2_{21} carry NuFIT 5.2 Gaussian priors, and their post fit errors are shown normalized to the prior σ: a ratio near unity indicates the atmospheric sample adds no information beyond the external prior, while a ratio below unity would indicate the data is actively constraining that parameter. The parameters of interest $\sin^2\theta_{23}$, Δm^2_{32}, and δ_{CP} enter the fit with flat priors, so the Hesse errors reported for them reflect the local curvature of the likelihood at the Asimov minimum rather than a physical constraint; the sensitivity to these parameters is quantified by the $\Delta\chi^2$ profiles in Section 5.1 rather than by Hesse.

For the parameters of interest, the Hesse errors are compared against the 1σ intervals extracted from the profile likelihood scans (Section 5.1), defined by $\Delta\chi^2 = 1$ from the scan minimum. Agreement between the two indicates the likelihood is approximately parabolic near the minimum and the Gaussian approximation used by Hesse is valid; disagreement indicates a non Gaussian likelihood, in which case the scan interval is used as the reported uncertainty.

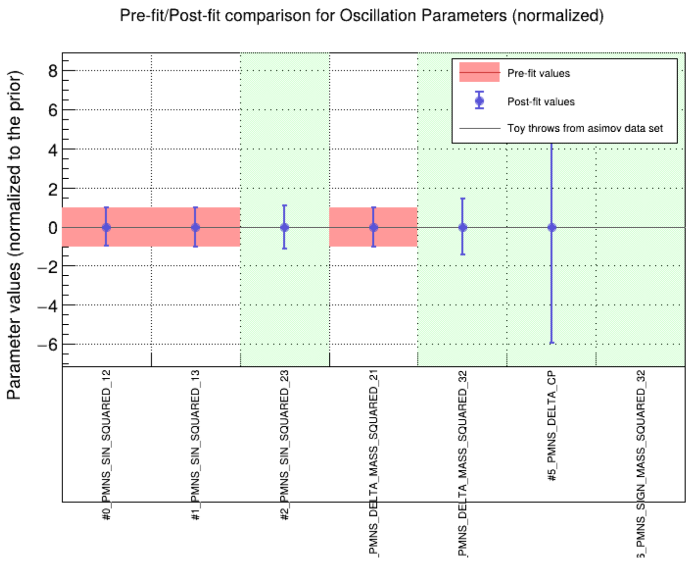

Figure 3.8: Postfit parameter uncertainties from the Asimov Hesse fit. For the solar sector parameters $\sin^2\theta_{12}$, $\sin^2\theta_{13}$, and Δm^2_{21}, which carry NuFIT 5.2 Gaussian priors, the postfit errors are shown normalized to the prior σ (red band); ratios close to unity indicate that the atmospheric sample does not improve on the external constraint. For the parameters of interest $\sin^2\theta_{23}$, Δm^2_{32}, and δ_{CP}, which enter the fit with flat priors, the Hesse errors reflect the local curvature of the likelihood at the Asimov minimum and are shown for diagnostic purposes only; the physical sensitivity to these parameters is given by the profile likelihood scans in Section 5.1.

Chapter 4

Oscillation Probability Calculation Optimization

4.1 Oscillation Grid Configuration

The GUNDAM framework computes oscillation probabilities on a two-dimensional grid in true neutrino energy E_ν and true zenith angle $\cos\theta_z$, using the `NuOscillator` interface to CUDAProb3 for three-flavor calculations including matter effects via the PREM Earth density profile. During fitting, probabilities at arbitrary $(E, \cos\theta_z)$ values are obtained by piecewise-linear interpolation between grid points. The accuracy of this interpolation depends critically on the grid resolution: if the spacing is too coarse relative to the oscillation structure, interpolation errors distort the predicted probabilities and could bias the extracted oscillation parameters. Figure 4.1 shows the $\nu_\mu \to \nu_\mu$ survival probability evaluated at high resolution (3000×1000). The oscillation structure is strongly energy dependent: at low energies ($E \lesssim 1$ GeV) the probability oscillates rapidly, while at high energies the oscillations are slow and well resolved by any reasonable grid. This motivates a grid that is fine at low energies and coarser at high energies.

4.1.1 Interpolation Error Tolerance

During fitting, GUNDAM does not recompute the oscillation probability from scratch for every Monte Carlo event. Instead, the probability is precomputed on a fixed grid in true neutrino energy E and true zenith angle $\cos\theta_z$ at the start of each fit iteration, and the value for each individual event is obtained by interpolation between neighboring grid nodes. The accuracy of this interpolation step sets a floor on how faithfully the predicted event rate tracks the true oscillation probability, and therefore places a requirement on how finely the grid must sample the oscillation structure. This subsection derives a sufficiency criterion for the grid spacing, expressed as a minimum number of points per oscillation period. The derivation below models the interpolation as piecewise-linear on a pure sinusoid, which serves as a conservative analytic proxy rather than a description of the actual interpolation scheme used by `TabulatedNuOscillator`;

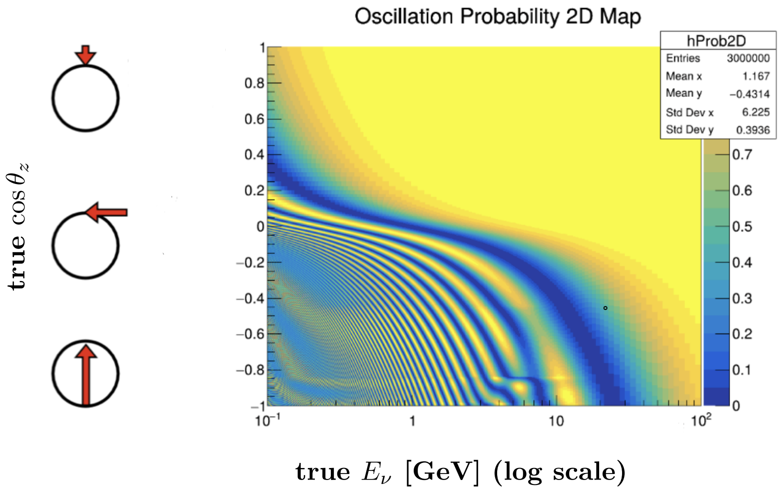

Figure 4.1: Muon neutrino survival probability $P(\nu_\mu \to \nu_\mu)$ as a function of true neutrino energy and true zenith angle, computed with CUDAProb3 at $3000{\times}1000$ (E, θ_z) resolution. The arrows on the left indicate neutrino direction: $\cos\theta_z = +1$ corresponds to downward-going neutrinos (short atmospheric path), $\cos\theta_z = 0$ to horizontal trajectories, and $\cos\theta_z = -1$ to upward-going neutrinos that traverse the full Earth diameter. Rapid oscillations at low energies demand fine grid spacing to avoid interpolation artifacts.

the empirical grid convergence study in Section 4.2 provides the direct test of grid adequacy.

The maximum error of piecewise-linear interpolation of a function $f(x)$ with uniform spacing h is bounded by

$$\epsilon_{\max} = \frac{h^2}{8} \max|f''(x)|\,. \tag{4.1}$$

For a sinusoidal oscillation $f(x) = \sin(x)$, where $\epsilon_{\max}$ is the maximum allowed error tolerance and $\max|f''| = 1$ and $h = 2\pi/N$ for N points per period, this gives

$$\epsilon = \frac{\pi^2}{2\,N^2}\,. \tag{4.2}$$

Inverting for the number of points per period required to achieve a given tolerance:

$$\boxed{N(\epsilon) = \left\lceil \frac{\pi}{\sqrt{2\,\epsilon}} \right\rceil}\,. \tag{4.3}$$

Setting $\epsilon \approx 0.077$ (corresponding to a maximum interpolation error of $\sim$8% on the oscillation probability) yields $N \approx 8$ points per period. Figure 4.2 shows this directly for $\sin^2(x)$: at 4 points per period the interpolation residuals are large, while at 8 points per period the error drops to $\epsilon \lesssim 8\%$. This criterion was also validated numerically by interpolating $\sin^2(x)$ at various resolutions and comparing to the exact function, with the results summarized in Table 4.1.

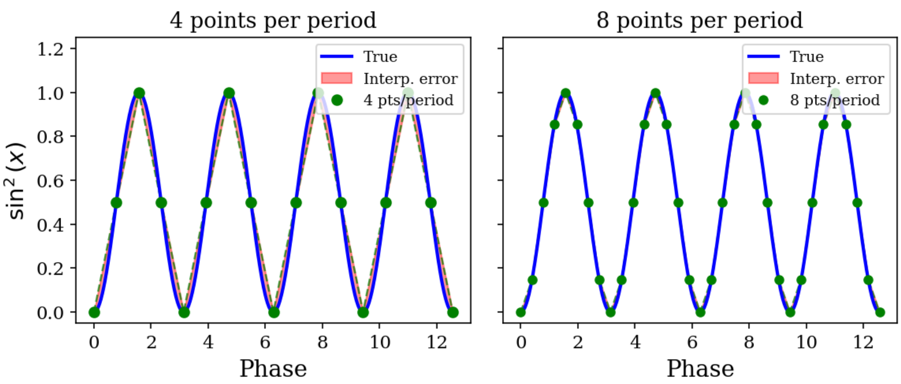

Figure 4.2: Piecewise-linear interpolation of $\sin^2(x)$ sampled at 4 points per period (left) and 8 points per period (right). Green markers show the sample locations, the blue curve shows the true function, and the red shaded band shows the interpolation error. At 4 points per period the linear interpolant visibly misses the peaks and troughs, while at 8 points per period the error drops below the 8% tolerance criterion of Eq. 4.3.

Table 4.1: Validation of the interpolation error bound: number of equally-spaced points per period required to achieve a given maximum error tolerance on $\sin^2(x)$.

Tolerance ϵ	Points per period
3.0×10^{-2}	13
4.5×10^{-2}	11
6.7×10^{-2}	9
$\mathbf{7.7 \times 10^{-2}}$	$\sim \mathbf{8}$

4.1.2 Three-Flavor Period Counting

Setting the grid resolution requires a concrete number: how many oscillation periods does $P(\nu_\mu \to \nu_\mu)$ contain across the analysis energy range? Multiplying that count by the eight-points-per-period sufficiency criterion derived in Section 4.1.1 fixes the required number of energy bins. The count must be made on the full three-flavor probability rather than on an analytic two-flavor form, because the three-flavor calculation includes solar oscillations driven by Δm^2_{21}, MSW resonances in the Earth's mantle and core, and interference between the atmospheric and solar amplitudes. These contributions introduce structure on shorter effective scales than the dominant atmospheric term alone, and an analytic estimate built only from Δm^2_{32} and a representative baseline would undercount the true number of cycles.

The counting was performed directly on the simulated probability curve. A one-dimensional slice of $P(\nu_\mu \to \nu_\mu)$ was computed with CUDAProb3 and the PREM Earth density profile at $\cos\theta_z = -1$, corresponding to vertically upgoing neutrinos that traverse the full Earth diameter (L = 12,742 km). This zenith angle is the worst case for grid resolution: the long baseline compresses the largest number of oscillation cycles into the analysis energy window, so a grid that adequately samples $\cos\theta_z = -1$ samples every shorter baseline at least as well. The probability was evaluated at high resolution and stored as a one-dimensional histogram of P versus E, with the sampling chosen fine enough that the counting itself was not limited by resolution of the input curve.

Each full oscillation period corresponds to one local minimum of $P(\nu_\mu \to \nu_\mu)$, where the disappearance is maximal. The minima were located using SciPy's `find_peaks` routine applied to $1-P(\nu_\mu \to \nu_\mu)$: `find_peaks` identifies local maxima, and the maxima of $1-P$ are exactly the minima of P. Before peak-finding, the probability curve was smoothed with a Savitzky-Golay filter (window length 7, polynomial order 3) to suppress high-frequency numerical noise from the PREM density steps without broadening the underlying oscillation features, and a prominence threshold of 0.04 was applied to reject spurious fluctuations while retaining every physical oscillation minimum. The procedure returned **approximately 115 oscillation periods** across the full energy range.

Section 4.1.1 established that eight points per oscillation period are sufficient to keep the piecewise linear interpolation error below $\sim 8\%$. This criterion was chosen conservatively but kept as small as possible to limit the computational cost of the fit, which scales with the number of grid points. Combining the 115 counted periods with the eight points per period criterion gives $8\times115+1 = 921$ energy bins.

The oscillation phase $\phi = 1.267\,\Delta m^2\,L/E$ is linear in $1/E$, so full oscillation periods are uniformly spaced in that variable. Placing bin edges at equal intervals in $1/E$ therefore puts the same number of grid points inside every oscillation period across the full energy range, which is the condition under which the eight points per period criterion applies uniformly. The grid of 921 energy bins was generated using a direct `numpy.linspace` construction between $1/E_{\mathrm{max}}$ and $1/E_{\mathrm{min}}$, followed by inversion to recover energies. No adaptive step logic or curve driven edge placement is used.

This grid, with 300 uniform bins in $\cos\theta_z$ held fixed across all configurations, is referred to throughout this chapter as the 921×300 grid. It represents the minimum resolution that satisfies the analytic sampling criterion for the three flavor probability at the worst case baseline, and it serves as one of the reference configurations against which coarser grids are compared in Section 4.2.

4.1.3 Grid Construction Methods

Two methods were used to construct the grids tested in this study.

Pure uniform-in-$1/E$. The 921×300 grid was generated using a new grid-construction method developed for this analysis, motivated directly by the period counting argument of Section 4.1.2. Given 115 oscillation periods across the full energy range and an eight points per period target, 922 energy bin edges were placed at equally spaced values of $1/E$ between $1/E_{\mathrm{max}}$ and $1/E_{\mathrm{min}}$ and then inverted to obtain energies. No adaptive logic is applied at any stage, so the grid contains relatively coarse spacing at high energy where the oscillation phase evolves slowly. This is by design, not a defect: because the phase is linear in $1/E$, coarse E spacing at high energies corresponds to the same phase interval as fine E spacing at low energies, and the points per period density is preserved across the full range.

The `makeInverseEnergy` algorithm. The 150×300, 300×300, 1500×300, and 4500×300 grids were generated using the `makeInverseEnergy` function provided in the GUNDAM oscillation tools, which is the default grid construction algorithm shipped with `gundamInputs`. This algorithm starts from uniform in $1/E$ spacing but applies a `fractionLimit` correction: if the ratio of consecutive bin widths in $\log E$ exceeds a threshold, additional bins are inserted to prevent overly coarse spacing at high energies. The result is a hybrid scheme that follows $1/E$ at low energies and transitions toward logarithmic spacing at high energies. The behavior of the correction depends on the total bin count. At 150 and 300 bins the `fractionLimit` threshold is triggered and the high energy portion of the grid is pulled toward logarithmic spacing. At 1500 bins and above, the threshold is never reached and the algorithm returns a grid that is effectively indistinguishable from pure uniform in $1/E$.

All grids use 300 uniformly distributed bins in $\cos\theta_z$.

Table 4.2: Oscillation probability grid configurations tested in this study. All grids span $E \in [0.1, 100]$ GeV and $\cos\theta_z \in [-1, 0]$.

Grid ($N_E \times N_{\cos\theta}$)	N_E	Method	Effective spacing
150×300	150	`makeInverseEnergy`	Hybrid $1/E + \log E$
300×300	300	`makeInverseEnergy`	Hybrid $1/E + \log E$
921×300	921	Pure $1/E$	Uniform in $1/E$
1500×300	1500	`makeInverseEnergy`	$\approx$ pure $1/E$
4500×300	4500	`makeInverseEnergy`	$\approx$ pure $1/E$

The five grids span more than an order of magnitude in energy resolution. Only the 921×300 grid has its bin count set by the sampling criterion derived in this chapter: it is the smallest grid at which uniform in $1/E$ spacing places eight points in every one of the 115 counted oscillation periods. The other four grids come from different motivations. The 300×300 grid is the baseline configuration used by GUNDAM and by the DUNE atmospheric analysis for this work, and is the grid whose adequacy this study is ultimately meant to assess. The 1500×300 and 4500×300 grids serve as high resolution references against which coarser configurations are compared; at these bin counts the `makeInverseEnergy` correction is never triggered, so both are effectively pure uniform in $1/E$. The 150×300 grid is included to establish the lower bound of the convergence regime, below which grid effects become apparent. The convergence of oscillation parameters across these five grids is examined in Section 4.2.

4.2 Grid Resolution Studies

The oscillation grid configuration described in Section 4.1 proposes a 300×300 grid as the baseline for this analysis. Before proceeding to physics extraction, it is essential to verify that this grid provides sufficient resolution for the oscillation probability calculation, and that the choice of grid does not bias the extracted oscillation parameters. This section presents a systematic convergence study across the five configurations listed in Table 4.2, using the 4500×300 grid (1.35×10^6 nodes) as the reference against which coarser grids are evaluated.

4.2.1 Impact of Grid Resolution on Event Rate

The first test of grid adequacy is whether the predicted event rates are stable across grid resolutions. Figure 4.3 overlays the predicted event spectra from the five grids as a function of reconstructed neutrino energy (top row) and $\cos\theta_Z$ (bottom row), for both the ν_e-selected (left) and ν_μ-selected (right) samples. The ratio panels show the fractional deviation of each grid relative to the 4500×300 reference.

For the ν_e selection, all grids agree at the sub-percent level across the full energy range. The ν_μ selection shows slightly larger deviations, particularly in the 150×300 grid at high energies where the coarser energy binning under resolves the rapid oscillation structure. Even for the ν_μ sample the fractional differences remain below roughly 2% for all grids and below 1% for the 300×300, 921×300, and 1500×300 configurations. The $\cos\theta_Z$ projections (bottom row) show even better agreement, as the zenith binning is held fixed at 300 bins across all configurations.

To understand where in the analysis phase space the grid resolution effects are concentrated, Figure 4.4 presents two dimensional residual maps of the fractional event rate difference $\Delta N/N$ between each coarser grid and the 4500×300 reference, binned in true energy and $\cos\theta_Z$. The leftmost column shows the absolute event distribution from the reference grid for context.

The residual structure is concentrated in the region of upward-going neutrinos ($\cos\theta_Z < -0.2$) at energies between 1 and 30 GeV, precisely where the atmospheric oscillation probability undergoes the most rapid variation. The ν_μ sample (bottom row) shows systematically larger residuals than the ν_e sample (top row), consistent with the ν_μ disappearance channel containing faster oscillation structure that demands finer energy resolution. The 300×300, 921×300, and 1500×300 grids show comparably small residuals across the full phase space, while the 150×300 grid exhibits the largest deviations, particularly in the upward-going, multi-GeV region where its resolution falls below the threshold needed to resolve individual oscillation cycles.

4.2.2 Impact of Grid Resolution on Oscillation Parameters

The most stringent test of grid adequacy is whether the choice of oscillation grid resolution affects the extracted physics: the $\Delta\chi^2$ profile likelihood curves and their $\Delta\chi^2 = 1$ crossing widths. For each of the five grid configurations, a 21-point profile likelihood scan was performed for all six oscillation parameters. At each scan point the parameter of interest was held fixed at its scan value and the remaining parameters were varied in the fit. (Table 3.1).

Figure 4.5 overlays the $\Delta\chi^2$ curves for all six parameters across the five grid resolutions. The atmospheric sensitive parameters, $\sin^2\theta_{23}$, Δm^2_{32}, and δ_{CP}, produce well-defined minima near the injected truth values, reflecting genuine sensitivity from the atmospheric neutrino sample. The solar parameters ($\sin^2\theta_{12}$

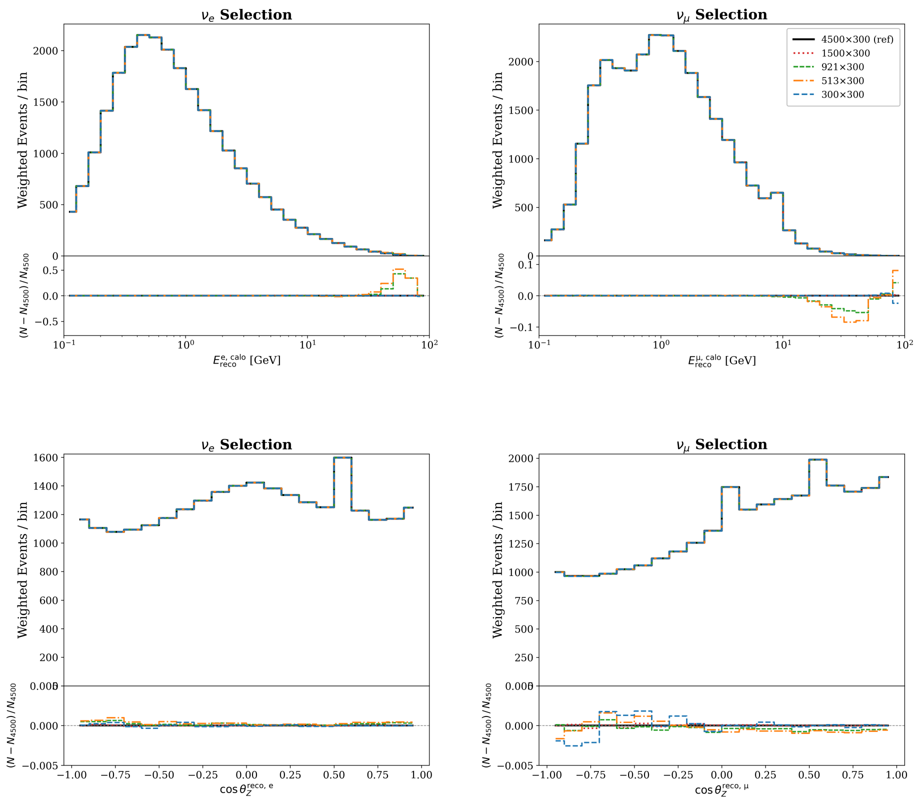

Figure 4.3: Predicted event rate spectra for the ν_e-selected (left) and ν_μ-selected (right) samples, shown as a function of reconstructed energy (top) and reconstructed $\cos\theta_Z$ (bottom), for five oscillation grid resolutions. The ν_e selection uses calorimetric electron energy (`recoEnuEcalo`) and electron zenith angle (`recoEle_thetaZ`), while the ν_μ selection uses calorimetric lepton energy (`recoEnuLepCalo`) and muon zenith angle (`recoMu_thetaZ`). Lower panels show the fractional deviation of each grid relative to the 4500×300 reference. In reconstructed space, all grids agree at the sub-percent level across most of the kinematic range; residual differences appear only at the highest reconstructed energies where event counts are small.

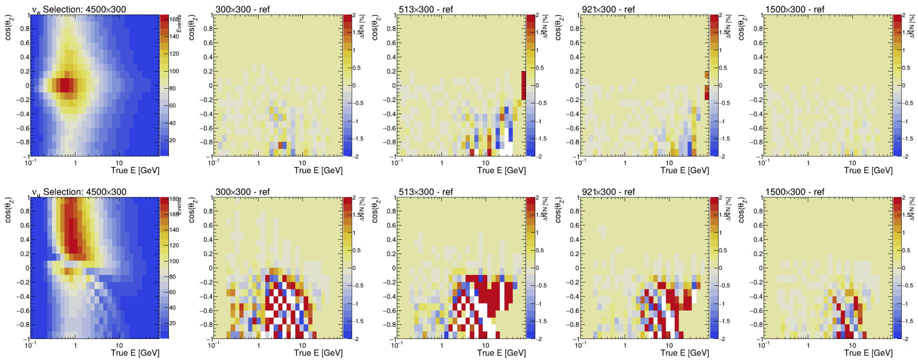

Figure 4.4: Two-dimensional maps of the fractional event rate residual $\Delta N/N$ [%] relative to the 4500×300 reference grid, shown in bins of true neutrino energy and $\cos\theta_Z$ for the ν_e selection (top) and ν_μ selection (bottom). The leftmost column shows the absolute event distribution from the reference grid. Residuals are concentrated in the upward-going, multi-GeV region where oscillation structure is most rapid.

and Δm^2_{21}) and $\sin^2\theta_{13}$ return curves dominated by their Gaussian prior constraints, confirming that the atmospheric data at this exposure contribute negligible additional information on these parameters. For all six parameters, the curves from the five grids are visually indistinguishable on the scale of the full $\Delta\chi^2$ range.

The residual panels in Figure 4.5 reveal the fine structure of the grid differences. The largest residuals appear in $\sin^2\theta_{23}$ for the 150×300 grid, where the curve deviates from the reference by tens of units in $\Delta\chi^2$ far from the minimum, and in δ_{CP} for the same grid, where the residual reaches of order unity. For all other grid and parameter combinations the residuals remain well below the dynamic range of the curves near the $\Delta\chi^2 = 1$ crossing region that determines the extracted width.

4.2.3 Quantitative Summary

To quantify the residual grid dependence, Figure 4.6 summarizes two metrics for each parameter and grid configuration, both computed relative to the 4500×300 reference. The left panel shows the fractional shift in the width at $\Delta\chi^2 = 1$,

$$\Delta w_G = \frac{|w_G - w_{\mathrm{ref}}|}{w_{\mathrm{ref}}} \times 100\%, \tag{4.4}$$

where w_G is the distance between the two points at which the curve for grid G crosses $\Delta\chi^2 = 1$, determined by linear interpolation. The right panel shows the maximum absolute pointwise deviation,

$$R_G = \max_i \left|\Delta\chi^2_G(x_i) - \Delta\chi^2_{\mathrm{ref}}(x_i)\right|, \tag{4.5}$$

where the maximum is taken over all scan points x_i shared between grid G and the reference. These two metrics are complementary. The width shift Δw_G

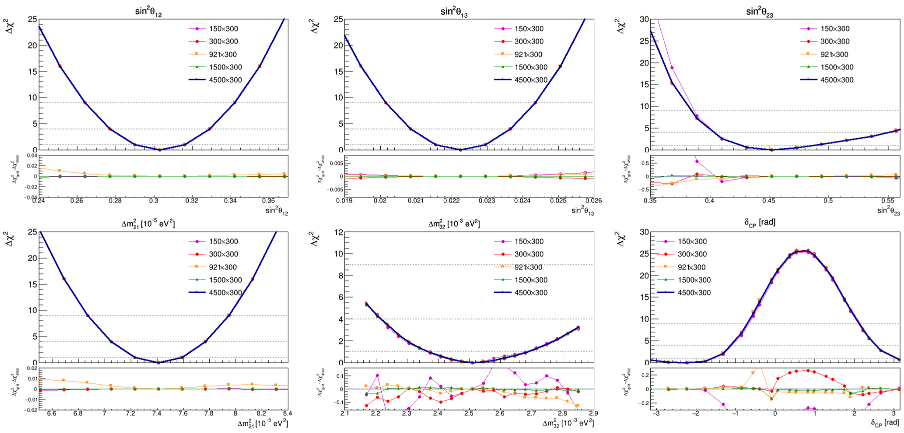

Figure 4.5: Profile likelihood scan curves ($\Delta\chi^2$ vs. parameter value) for all six oscillation parameters, overlaid for the five grid resolutions. At each scan point the parameter of interest is held fixed at the scan value and the remaining parameters are profiled over. Lower panels show the residual $\Delta(\Delta\chi^2)$ of each grid relative to the 4500×300 reference. The atmospheric parameters ($\sin^2\theta_{23}$, Δm^2_{32}, δ_{CP}) show genuine sensitivity from the data, while the solar parameters and $\sin^2\theta_{13}$ are prior-dominated. All grids produce indistinguishable curves on the scale of the full $\Delta\chi^2$ range.

captures whether the 1σ uncertainty that would be quoted from a given grid changes with resolution, while R_G captures whether the curve shape deviates anywhere along the scan, including regions far from the $\Delta\chi^2 = 1$ crossing where a large pointwise disagreement has no effect on the reported uncertainty. For δ_{CP}, the $\Delta\chi^2 = 1$ width is not reported because the likelihood curve does not cross unity within the scan range at this exposure; only R_G is quoted.

The 150×300 grid is included to establish the lower bound of the convergence regime, and its behavior confirms the role of that bound. The width shifts at this resolution reach 3.1% for Δm^2_{32} and 1.6% for $\sin^2\theta_{23}$, while all other parameters stay below 0.02%. The pointwise deviations are far more dramatic: $R_G = 57.4$ for $\sin^2\theta_{23}$, 0.74 for δ_{CP}, and 0.22 for Δm^2_{32}. Taken together, these results confirm that the 150×300 grid is insufficiently resolved for a reliable sensitivity analysis.

For all configurations at 300×300 and above the picture changes qualitatively. The prior dominated parameters ($\sin^2\theta_{12}$, $\sin^2\theta_{13}$, Δm^2_{21}) show width shifts at or below 0.1% and R_G values below 0.11 across every grid, reflecting that the atmospheric sample provides negligible information on these parameters and the likelihood surface is set almost entirely by the external priors. For δ_{CP}, R_G falls from 0.26 at 300×300 to 0.14 at 1500×300.

The atmospheric sensitive parameter most sensitive to grid resolution is Δm^2_{32}. At 300×300 the width shift is 2.4%, falling to 2.0% at 921×300 and to 0.09% at 1500×300. The modest improvement from 300 to 921, despite the three-fold increase in energy bins, indicates that the finite-resolution effect on the Δm^2_{32} width near its minimum is not primarily set by the density of points per oscillation period in the energy grid. Only when the grid is refined to

1500×300 does the Δm^2_{32} width converge to the reference at the sub-0.1% level. On $\sin^2\theta_{23}$, the width shifts remain modest across all grids (1.2% at 300×300, 0.5% at 921×300), while the R_G values are larger and not monotonic with grid size: the 921×300 grid yields $R_G = 8.3$ against $R_G = 1.3$ for 300×300. Because these pointwise deviations occur in regions of the scan far from the $\Delta\chi^2 = 1$ crossing, they do not affect the extracted 1σ uncertainty.

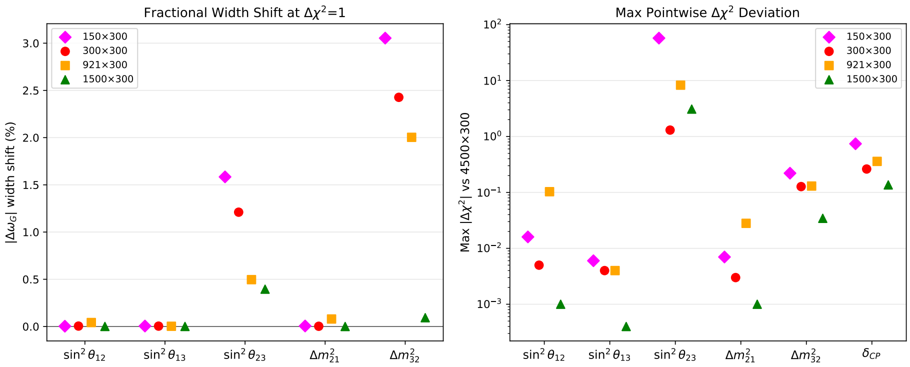

Figure 4.6: Quantitative summary of grid resolution effects on the profile likelihood comparison, computed relative to the 4500×300 reference grid. Left: fractional shift in the width at $\Delta\chi^2 = 1$, Δw_G (Equation 4.4). Right: maximum absolute pointwise $\Delta\chi^2$ deviation R_G (Equation 4.5), plotted on a logarithmic vertical axis. The 150×300 grid (pink diamonds) exhibits width shifts of 3.1% for Δm^2_{32} and 1.6% for $\sin^2\theta_{23}$ and a pointwise deviation of $R_G = 57.4$ for $\sin^2\theta_{23}$, confirming it is insufficiently resolved for a reliable sensitivity analysis. All configurations at 300×300 and above show width shifts at or below 2.4% and are discussed in the text. Width shifts for δ_{CP} are not reported because the likelihood curve does not cross $\Delta\chi^2 = 1$ within the scan range.

4.2.4 Conclusion

On the basis of these comparisons, the 300×300 grid (generated with the `makeInverseEnergy` algorithm) is adopted as the baseline configuration for all subsequent analyses in this work. Across the four parameters for which the $\Delta\chi^2 = 1$ width is well defined, the 300×300 grid agrees with the 4500×300 reference to better than 0.01% on the prior-dominated parameters $\sin^2\theta_{12}$, $\sin^2\theta_{13}$, and Δm^2_{21}, to 1.2% on $\sin^2\theta_{23}$, and to 2.4% on Δm^2_{32}. The largest of these, the 2.4% shift on Δm^2_{32}, is small compared to the roughly 1 to 2% projected 1σ precision on Δm^2_{32} from DUNE atmospheric data at the exposure considered here, and translates to a change in the reported uncertainty that is well below the experimental systematic budget. The pure uniform in $1/E$ 921×300 grid, designed around the three-flavor period counting criterion derived in Section 4.1.2, produces a width shift of 2.0% on Δm^2_{32}, only marginally smaller than the baseline, indicating that the residual width effect at 300×300 is not primarily driven by the energy grid spacing. Only the 1500×300 grid reduces the Δm^2_{32} width shift below 0.1%, and at a threefold cost in oscillation nodes relative to the baseline. The 300×300 configuration is therefore sufficient for the sensitivity studies presented in the remainder of this thesis.

Chapter 5

DUNE Sensitivity to Atmospheric Neutrino Oscillations

In this chapter, we present sensitivity results using the 300×300 $(E_\nu, \cos\theta_z)$ baseline grid validated in Chapter 4, with simulated data generated at the NuFit 5.2 best fit values (normal ordering, with Super-Kamiokande atmospheric data; Table 3.1). All results represent a statistical only sensitivity; systematic uncertainties on flux, cross sections, and detector response are deferred to future work within the collaboration. Section 5.1 presents one dimensional profile likelihood scans for all six oscillation parameters. Section 5.2 maps the correlations between the atmospheric sector parameters through two dimensional profiled likelihood contour scans. Section 5.3 evaluates the mass ordering sensitivity as a function of exposure under two configurations of external constraints.

5.1 One-Dimensional Profile Likelihood Scans

For each of the six oscillation parameters, we perform a one dimensional profile likelihood scan by varying the parameter of interest over a range of $\pm 10\sigma$ around its NuFit 5.2 central value in 21 equally spaced steps, where σ denotes the NuFit 5.2 1σ width. At each scan point the parameter of interest is held fixed at its scan value and the remaining five oscillation parameters are profiled: Minuit2 minimizes the negative log likelihood over them at each point. The test statistic is $\Delta\chi^2$, defined relative to the global minimum of $-2\ln\mathcal{L}$ across the scan. By Wilks' theorem the resulting curves follow a χ^2 distribution with one degree of freedom, so the $\Delta\chi^2 = 1$, 4, and 9 crossings correspond to the 1σ, 2σ, and 3σ confidence intervals on the scanned parameter. The two dimensional profiled contours presented in Section 5.2 extend this procedure to pairs of parameters, simultaneously profiling the remaining four.

The results divide into two categories. Three parameters, $\sin^2\theta_{12}$, $\sin^2\theta_{13}$, and Δm^2_{21}, are prior dominated: the atmospheric neutrino sample at DUNE does not significantly improve upon the precision set by the external NuFit constraints. The remaining three, $\sin^2\theta_{23}$, Δm^2_{32}, and δ_{CP}, are parameters to which DUNE atmospheric neutrinos are expected to contribute sensitivity beyond the priors, though as discussed below the degree of constraint varies considerably among them.

5.1.1 Prior-Dominated Parameters

Figure 5.1 shows the profile likelihood scans for $\sin^2\theta_{12}$, $\sin^2\theta_{13}$, and Δm^2_{21}. All three curves are clean, symmetric parabolas whose $\Delta\chi^2 = 1$ crossings coincide with the NuFit 5.2 prior band (pink shading), and the minima sit at the injected truth values to within numerical precision.

The parameters $\sin^2\theta_{12}$ and Δm^2_{21} govern oscillations at the solar Δm^2_{21} scale. Because Δm^2_{21} is roughly 30 times smaller than Δm^2_{32}, the associated oscillation phase $\Delta m^2_{21}L/4E$ remains small even for the longest atmospheric baselines. Thus, atmospheric neutrino data provide negligible constraining power on these parameters relative to existing global fits.

Their influence on the atmospheric event rate is indirect, entering through sub-leading terms in the three-flavor oscillation probability, and is far too weak to compete with the precise constraints from KamLAND and solar neutrino experiments encoded in the NuFit priors. Similarly, $\sin^2\theta_{13}$ is constrained to sub-percent precision by the reactor experiments Daya Bay, Double Chooz, and RENO. While θ_{13} does affect the atmospheric $\nu_\mu \to \nu_e$ appearance channel (Equation 1.13), the prior is sufficiently tight that the atmospheric data cannot meaningfully narrow it further.

These scans also serve as a consistency check on the fit configuration: clean parabolas centered on the injected truth values confirm that the likelihood evaluation, parameter injection, and profiling machinery are functioning correctly. Any deviation from this behavior for the prior dominated parameters would point to a problem in the fit configuration rather than a physical effect.

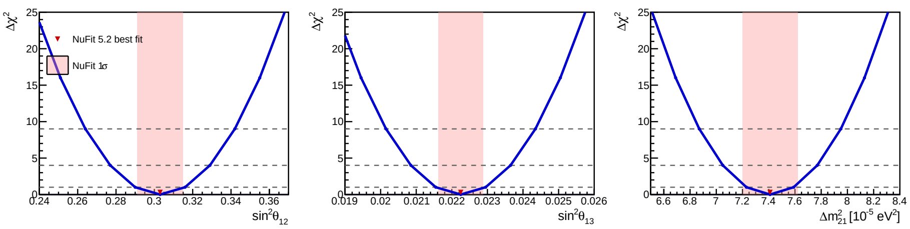

Figure 5.1: One dimensional profile likelihood scans for the three prior-dominated oscillation parameters: $\sin^2\theta_{12}$, $\sin^2\theta_{13}$, and Δm^2_{21}. At each scan point the parameter of interest is held fixed at the scan value and the remaining five oscillation parameters are profiled over by Minuit2. The pink band shows the NuFit 5.2 1σ range and the red triangle marks the NuFit 5.2 best fit. The horizontal dashed lines at $\Delta\chi^2 = 1$, 4, and 9 correspond to the 1σ, 2σ, and 3σ intervals by Wilks' theorem. All three parameters recover clean parabolas consistent with their input prior widths, indicating that the atmospheric neutrino sample does not add significant constraint beyond the external priors.

5.1.2 Atmospheric Neutrino Oscillation Parameters

Figure 5.2 shows the profile likelihood scans for the three parameters to which DUNE atmospheric neutrinos are most sensitive. As with the prior dominated scans, the non-scanned oscillation parameters are profiled over at each scan point, so the resulting $\Delta\chi^2$ curves give confidence intervals on the scanned parameter after accounting for the best fit response of the remaining parameters. The full allowed regions that account for the correlations between $\sin^2\theta_{23}$, Δm^2_{32}, and δ_{CP} are presented as profiled two dimensional contours in Section 5.2.

$\sin^2\theta_{23}$: The mixing angle θ_{23} governs the dominant ν_μ disappearance channel. In the two flavor approximation, the survival probability depends on $\sin^2 2\theta_{23}$ (Equation 1.9), which is symmetric about $\theta_{23} = \pi/4$ and therefore insensitive to the octant. However, the full three flavor oscillation probabilities, particularly the $\nu_\mu \to \nu_e$ appearance channel whose leading term scales as $\sin^2\theta_{23}$ (Equation 1.13), partially break this degeneracy. Matter effects in upward going atmospheric neutrinos further distinguish the two octants.

The resulting curve exhibits a visible asymmetry, with $\Delta\chi^2$ rising more steeply on the lower octant side than on the upper octant side. The $\Delta\chi^2 = 1$ crossings sit at approximately $\sin^2\theta_{23} \approx 0.425$ and 0.475, giving a 1σ interval of roughly 0.05 that is comparable in width to the NuFit 5.2 prior band in this region. The broader rise on the upper octant side reflects the residual octant degeneracy, partially but not fully broken by matter effects at this exposure.

Δm^2_{32}: The atmospheric mass splitting sets the frequency of ν_μ disappearance oscillations and is imprinted directly on the position of the first oscillation dip in the L/E distribution (Figure 1.9). The profile scan for Δm^2_{32} is tightly peaked around the injected truth, with $\Delta\chi^2 = 1$ crossings at approximately $\Delta m^2_{32} \approx$ 2.48 and 2.54×10^{-3} eV^2. The resulting 1σ interval of roughly 0.06×10^{-3} eV^2 is narrower than the NuFit 5.2 prior band, indicating that the atmospheric sample at this exposure provides additional constraint on Δm^2_{32} beyond the NuFit prior. Of the six oscillation parameters, this is the one to which the atmospheric sample is most sensitive, as the dip position can be precisely located given sufficient statistics and energy resolution. The curve is nearly symmetric and parabolic near the minimum, consistent with Gaussian behavior in the regime close to the truth value.

δ_{CP}: The CP-violating phase enters the atmospheric oscillation probabilities through the sub-leading $\nu_\mu \to \nu_e$ appearance channel (Equation 1.19), where the CP-odd interference term is proportional to $\sin\delta_{\mathrm{CP}}$. Unlike long baseline beam experiments, where the fixed baseline and peaked energy spectrum allow a direct rate comparison between neutrino and antineutrino modes, atmospheric neutrinos sample a continuous distribution of L/E values. Any CP-violating effect is diluted by the broad energy spectrum and the integration over production heights and zenith angles.

The δ_{CP} profile scan is dramatically broader than the other two sensitive parameters and is not symmetric about the minimum. Moving to the right of the injected truth (toward $\delta_{\mathrm{CP}} = 0$), the likelihood rises sharply and peaks at $\Delta\chi^2 \approx 26$ near $\delta_{\mathrm{CP}} \approx +0.63$ rad before falling again toward $+\pi$. Moving to the left (toward $-\pi$), $\Delta\chi^2$ remains essentially flat and does not cross unity within the scan range. A well defined two-sided $\Delta\chi^2 = 1$ interval therefore cannot be extracted from this scan. Consistent with this, the post-fit Hesse uncertainty on δ_{CP} is approximately six times the NuFit 5.2 1σ width (Figure 3.8), reflecting the shallow curvature of the profile likelihood near the minimum rather than a constraint competing with NuFit. The atmospheric sample alone does not meaningfully constrain δ_{CP} at this exposure. This is consistent with the expectation that atmospheric δ_{CP} sensitivity is limited compared to dedicated beam measurements. The primary value of the atmospheric channel for δ_{CP} lies in its complementarity to beam results: because the atmospheric measurement uses different baselines, energies, and systematics, it provides an independent cross check when eventually combined with the DUNE beam analysis.

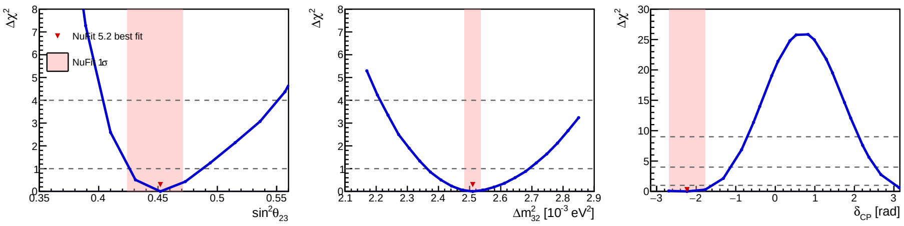

Figure 5.2: One dimensional profile likelihood scans for the three parameters to which DUNE atmospheric neutrinos are most sensitive: $\sin^2\theta_{23}$, Δm^2_{32}, and δ_{CP}. At each scan point the parameter of interest is held fixed at the scan value and the remaining five oscillation parameters are profiled over. The pink band shows the NuFit 5.2 1σ range and the red triangle marks the NuFit 5.2 best fit. The $\sin^2\theta_{23}$ curve exhibits an asymmetric structure reflecting residual octant sensitivity. The Δm^2_{32} curve is tightly peaked, with the $\Delta\chi^2 = 1$ interval narrower than the NuFit prior band. The δ_{CP} curve is highly asymmetric: $\Delta\chi^2$ rises sharply toward positive δ_{CP} and remains below unity across the scan range on the negative side, so no two sided 1σ interval can be extracted. The full allowed regions including correlations between these parameters are presented in Section 5.2.

5.1.3 Summary

The one dimensional profile likelihood scans confirm the expected sensitivity hierarchy for DUNE atmospheric neutrinos. The solar parameters and θ_{13} are entirely prior dominated, serving as consistency checks on the fit configuration. The atmospheric mass splitting Δm^2_{32} is the parameter to which the atmospheric sample is most sensitive, benefiting from the direct imprint of the oscillation frequency on the L/E distribution. The mixing angle θ_{23} shows genuine atmospheric sensitivity with an asymmetric curve reflecting partial octant resolution. The CP phase δ_{CP} is not meaningfully constrained by the atmospheric sample alone, though it will provide an independent measurement when combined with the beam program.

The two dimensional profiled contour scans presented in Section 5.2 extend this analysis to pairs of atmospheric sector parameters, profiling over the remaining four at each grid point to map the correlated allowed regions for $\sin^2\theta_{23}$, Δm^2_{32}, and $\delta_{\rm CP}$.

5.2 Two-Dimensional Contour Scans

To map the correlations between oscillation parameters and quantify the allowed regions in parameter space, three two dimensional profiled likelihood scans were performed over pairs of ($\delta_{\rm CP}$, $\sin^2\theta_{23}$), ($\delta_{\rm CP}$, Δm^2_{32}), and ($\sin^2\theta_{23}$, Δm^2_{32}). Each scan evaluates a 50×50 grid (2500 points), where the two scan parameters are fixed at each grid point and all remaining oscillation parameters are profiled (minimized) using the GUNDAM fitter. The grid ranges are summarized in Table 5.1. Confidence regions are drawn at the two degree of freedom thresholds $\Delta\chi^2 = 2.30$, 6.18, and 11.83, corresponding to 1σ, 2σ, and 3σ, respectively. One dimensional profile likelihoods are obtained by projecting the minimum $\Delta\chi^2$ onto each axis.

Table 5.1: Grid ranges for the three two dimensional profiled likelihood scans. Each scan uses a 50×50 uniform grid (2500 points).

Scan	Parameter 1	Range	Parameter 2	Range
$\delta_{\rm CP}$ vs $\sin^2\theta_{23}$	$\delta_{\rm CP}$	$[-\pi,\ \pi]$	$\sin^2\theta_{23}$	[0.41, 0.59]
$\delta_{\rm CP}$ vs Δm^2_{32}	$\delta_{\rm CP}$	$[-\pi,\ \pi]$	Δm^2_{32}	$[2.30,\ 2.70]\times 10^{-3}$ eV^2
$\sin^2\theta_{23}$ vs Δm^2_{32}	$\sin^2\theta_{23}$	[0.40, 0.60]	Δm^2_{32}	$[2.20,\ 2.80]\times 10^{-3}$ eV^2

5.2.1 $\sin^2\theta_{23}$ vs Δm^2_{32}

Figure 5.3 shows the two dimensional profiled $\Delta\chi^2$ surface for $\sin^2\theta_{23}$ versus Δm^2_{32}. Of the 2500 grid points, 1604 fits converged successfully, providing sufficient coverage to draw the 1σ, 2σ, and 3σ contours. The profiled best-fit values ($\sin^2\theta_{23} = 0.4531$, $\Delta m^2_{32} = 0.002506$ eV^2) are in close agreement with the injected NuFit 5.2 truth ($\sin^2\theta_{23} = 0.451$, $\Delta m^2_{32} = 0.00251$ eV^2), confirming that the Asimov fit correctly recovers the input parameters for these two quantities.

The 1σ contour is closed and compact, indicating strong constraining power. The shape is elongated toward higher $\sin^2\theta_{23}$, reflecting an asymmetry between the lower and upper octants: atmospheric neutrino disappearance is governed by $\sin^2 2\theta_{23}$, which is symmetric about maximal mixing ($\sin^2\theta_{23} = 0.5$), so values above 0.5 produce a mirror degeneracy. The injected truth lies in the lower octant ($\sin^2\theta_{23} = 0.451$), and the tighter constraint on the lower side reflects DUNE's ability to partially resolve this octant ambiguity through matter effects and the broad energy coverage of atmospheric neutrinos. The Δm^2_{32} constraint is tight, with the 1σ band spanning roughly $[2.44,\ 2.56]\times 10^{-3}$ eV^2, consistent with the precision expected from the many oscillation periods sampled by atmospheric neutrinos traversing the Earth.

5.2.2 Contours involving $\delta_{\rm CP}$

The two remaining scans probe the sensitivity to the CP-violating phase $\delta_{\rm CP}$ in combination with $\sin^2\theta_{23}$ and Δm^2_{32}. These pairings are motivated by known

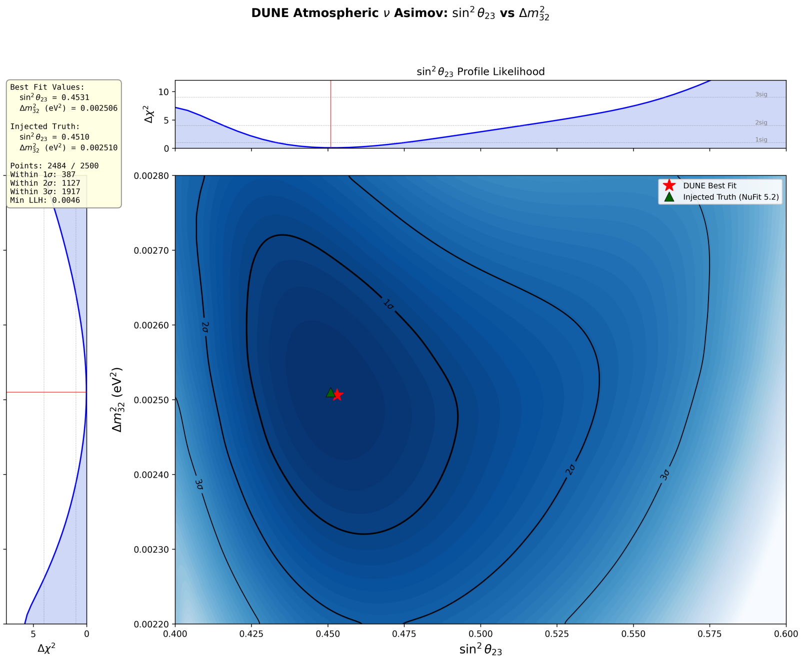

Figure 5.3: Two dimensional profiled $\Delta\chi^2$ contour for $\sin^2\theta_{23}$ versus Δm^2_{32}. Contours are drawn at $\Delta\chi^2 = 2.30$, 6.18, and 11.83 (corresponding to 1σ, 2σ, 3σ for 2 DOF). The red star marks the profiled best fit; the green triangle marks the injected NuFit 5.2 truth. Side panels show the one dimensional profile likelihoods obtained by minimizing $\Delta\chi^2$ along the orthogonal axis.

correlations in the oscillation probabilities. The sub-dominant $\nu_\mu \to \nu_e$ appearance probability couples δ_{CP} and $\sin^2\theta_{23}$ through the leading term, which is proportional to $\sin^2\theta_{23}\sin^2 2\theta_{13} f(\delta_{\text{CP}})$: a shift in θ_{23} can be partially compensated by a shift in δ_{CP} to produce a similar appearance rate, giving rise to the well-known octant and δ_{CP} degeneracy. Similarly, δ_{CP} and Δm^2_{32} are correlated because the CP interference term is proportional to $\sin\delta_{\text{CP}}$ modulated by sinusoidal factors of $\Delta m^2_{32}L/4E$, so shifts in either parameter produce correlated changes in the predicted energy spectrum. These correlations are hidden in one dimensional profile likelihoods and can only be assessed through two dimensional scans.

Figure 5.4 shows the δ_{CP} versus $\sin^2\theta_{23}$ contour and Figure 5.5 shows the δ_{CP} versus Δm^2_{32} contour. In both cases, the $\sin^2\theta_{23}$ and Δm^2_{32} best-fit values lie close to truth, and the profiled best fit for δ_{CP} is consistent with the injected value, with the minimum occurring at $\delta_{\text{CP}} \approx -2.24$ rad compared to the injected NuFit 5.2 value of -2.233 rad. This small offset reflects a near-degeneracy in the likelihood surface rather than a bias in the fit.

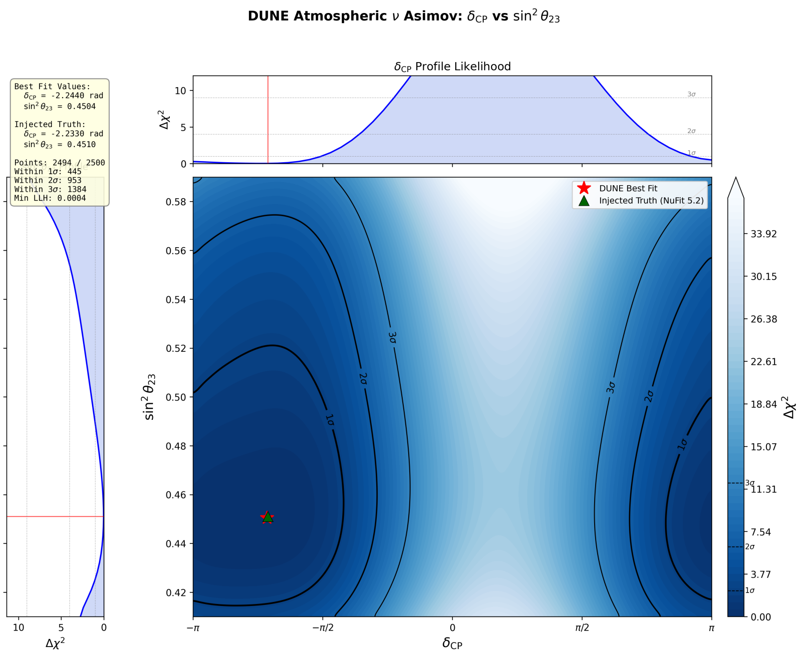

Figure 5.4: Two-dimensional profiled $\Delta\chi^2$ contour for δ_{CP} versus $\sin^2\theta_{23}$. Contours are drawn at $\Delta\chi^2 = 2.30$, 6.18, and 11.83 (1σ, 2σ, 3σ for 2 DOF). The red star marks the profiled best fit; the green triangle marks the injected NuFit 5.2 truth. Side panels show the one-dimensional profile likelihoods.

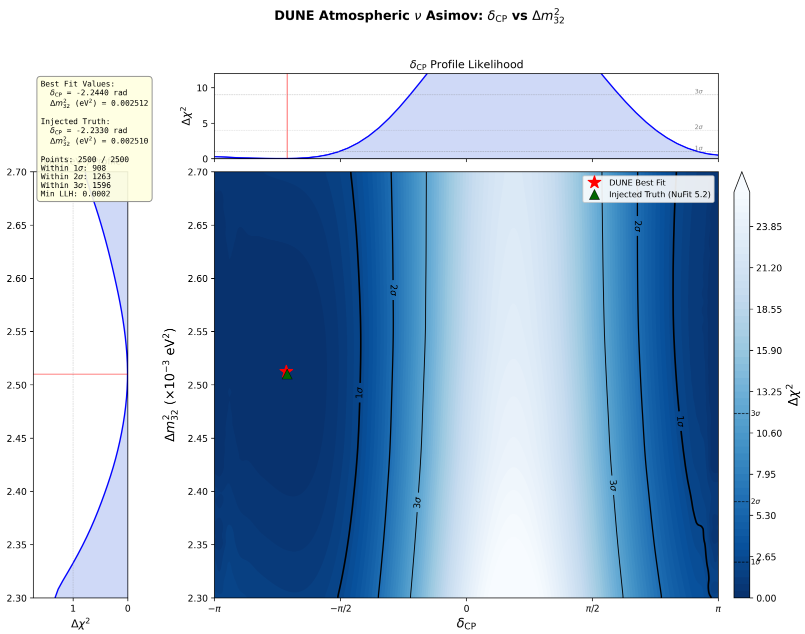

Figure 5.5: Two dimensional profiled $\Delta\chi^2$ contour for $\delta_{\rm CP}$ versus Δm^2_{32}. Contours are drawn at $\Delta\chi^2 = 2.30$, 6.18, and 11.83 (1σ, 2σ, 3σ for 2 DOF). The red star marks the profiled best fit; the green triangle marks the injected NuFit 5.2 truth. Side panels show the one-dimensional profile likelihoods.

5.3 Mass Ordering Sensitivity

The neutrino mass ordering refers to whether the third mass eigenstate ν_3 is heavier (normal ordering, NO) or lighter (inverted ordering, IO) than ν_1 and ν_2. As discussed in Section 1.2.2, the sign of Δm^2_{31} determines whether the MSW resonance condition is met for neutrinos or antineutrinos as they propagate through the Earth. For normal ordering ($\Delta m^2_{31} > 0$), the resonance occurs in the neutrino channel, while for inverted ordering ($\Delta m^2_{31} < 0$), it occurs in the antineutrino channel. Since the neutrino interaction cross section and atmospheric flux are both larger than their antineutrino counterparts, the resonance enhancement produces a larger observable effect when it falls in the neutrino channel. DUNE's atmospheric neutrino sample can therefore provide sensitivity to the mass ordering that is complementary to the beam-based measurement.

5.3.1 Method

The mass ordering sensitivity is evaluated using an Asimov hypothesis test. For a given true ordering, a simulated dataset is generated using the NuFit 5.2 [35] best-fit oscillation parameters for that ordering. Importantly, the best-fit values differ between the two orderings: in particular, the NuFit 5.2 fit with SK atmospheric data places NO in the lower octant ($\sin^2\theta_{23} = 0.451$) and IO in the upper octant ($\sin^2\theta_{23} = 0.569$), while the fit without SK data places both orderings in the upper octant ($\sin^2\theta_{23} \approx 0.57$). As discussed below, this difference plays a significant role in the results.

Each Asimov dataset is fit twice: once under the same-ordering hypothesis (matching the injection) and once under the opposite-ordering hypothesis. The ordering is enforced in the fit through the GUNDAM parameter `PMNS_SIGN_MASS_SQUARED_32`, whose sign determines whether the oscillation probability engine uses $\Delta m^2_{32} > 0$ (normal ordering) or $\Delta m^2_{32} < 0$ (inverted ordering); its magnitude carries no physical meaning. This parameter is held fixed throughout the fit, while the remaining oscillation parameters are profiled within a single ordering hypothesis under their prior constraints. The sensitivity is defined as

$$\sqrt{\Delta\chi^2_{\mathrm{MO}}} = \sqrt{\chi^2_{\mathrm{opposite}} - \chi^2_{\mathrm{same}}}, \tag{5.1}$$

where χ^2_{same} and $\chi^2_{\mathrm{opposite}}$ are the best fit negative log likelihoods under the same-ordering and opposite ordering hypotheses, respectively. This procedure is repeated at eleven exposure values ranging from 40 to 2000 kt·yr (corresponding to 1 to 50 years of running with a 40 kt far detector module). The analysis includes only oscillation related systematic uncertainties; flux, cross-section, and detector response systematics are not included.

The study is performed under two configurations of the NuFit 5.2 global fit: one that includes Super-Kamiokande atmospheric data and one that excludes it. The injected truth values, prior central values, and 1σ widths for both configurations are listed in Table 3.2.

5.3.2 Results

Figure 5.6 shows the mass ordering sensitivity as a function of exposure using NuFit 5.2 with SK atmospheric data included. Assuming IO is the true ordering, the sensitivity reaches approximately 3.2σ at 2000 kt·yr. Assuming NO is the true ordering, the sensitivity is lower, reaching approximately 2.0σ at the same exposure.

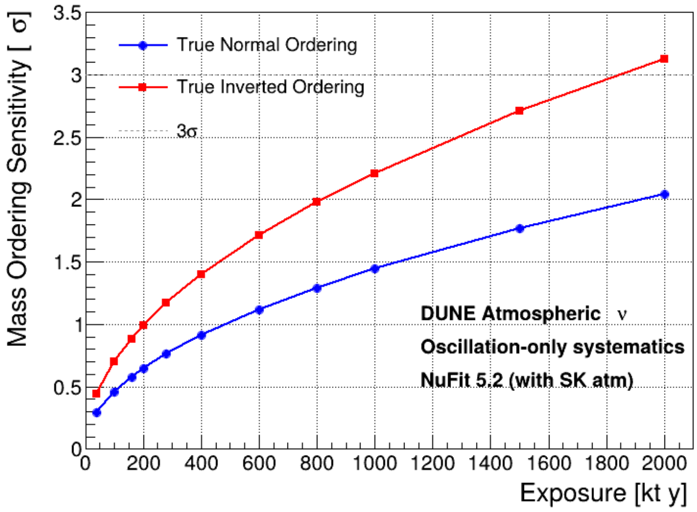

Figure 5.6: Mass ordering sensitivity as a function of exposure for DUNE atmospheric neutrinos, using NuFit 5.2 oscillation parameters with SK atmospheric data. The blue curve assumes normal ordering is true, and the red curve assumes inverted ordering is true.

Figure 5.7 shows the same study using NuFit 5.2 without SK atmospheric data. The sensitivities are notably higher: approximately 4.3σ for true NO and 3.9σ for true IO at 2000 kt·yr. The ordering of the curves also reverses, with NO truth now giving higher sensitivity than IO truth.

5.3.3 MSW Resonance and Octant Competition

The reversal of the sensitivity curves between the two configurations can be understood as a competition between two physical effects.

The first is the MSW matter resonance. For normal ordering, the resonance falls in the neutrino channel, which has higher statistics due to larger flux and cross section. This gives NO truth a natural advantage in terms of the observable signal. For inverted ordering, the resonance falls in the antineutrino channel, where the statistics are lower. In the absence of other effects, one would therefore expect NO truth to yield higher sensitivity than IO truth.

The second effect is the θ_{23} octant. The appearance probability $P(\nu_\mu \to \nu_e)$ scales with $\sin^2\theta_{23}$: a value in the upper octant ($\sin^2\theta_{23} > 0.5$) produces more ν_e appearance events than a value in the lower octant, effectively providing a "signal boost" for the mass ordering measurement.

Without SK atmospheric data, the NuFit 5.2 best fit places both orderings in the upper octant ($\sin^2\theta_{23} \approx 0.57$). In this case the octant effect is similar for

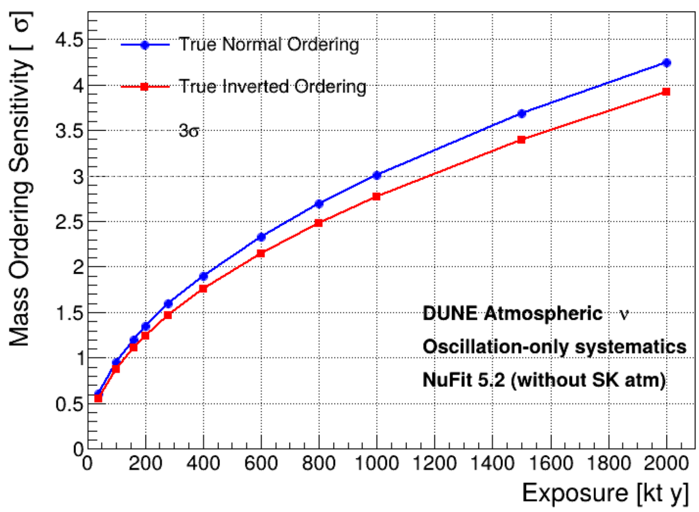

Figure 5.7: Mass ordering sensitivity as a function of exposure for DUNE atmospheric neutrinos, using NuFit 5.2 oscillation parameters without SK atmospheric data. Both orderings reach higher significance compared to the with-SK configuration, and the relative ordering of the curves is reversed.

both orderings and the MSW resonance dominates, yielding higher sensitivity for NO truth. This is the expected behavior.

With SK atmospheric data, however, the NuFit 5.2 best fit shifts the NO value into the lower octant ($\sin^2\theta_{23} = 0.451$) while keeping IO in the upper octant ($\sin^2\theta_{23} = 0.569$). The upper octant signal boost for IO truth compensates for its lack of MSW resonance in the neutrino channel, and the lower octant suppression for NO truth weakens its MSW advantage. The net result is that IO truth achieves higher sensitivity than NO truth when SK constraints are included.

Additionally, the overall sensitivity is lower in the with-SK configuration because the tighter external constraints on the oscillation parameters give the minimizer more freedom to partially absorb the wrong-ordering hypothesis by adjusting the profiled parameters, reducing the $\Delta\chi^2$.

5.3.4 Comparison

Figure 5.8 combines all four sensitivity curves and includes a lower panel showing the difference in significance between the two NuFit configurations. For both true orderings, removing the SK constraints increases the sensitivity by roughly 0.5 to 1.5σ across the exposure range, with the difference growing at higher exposures.

5.3.5 Summary

DUNE's atmospheric neutrino sample provides meaningful sensitivity to the neutrino mass ordering, reaching up to 3-4σ depending on the true ordering

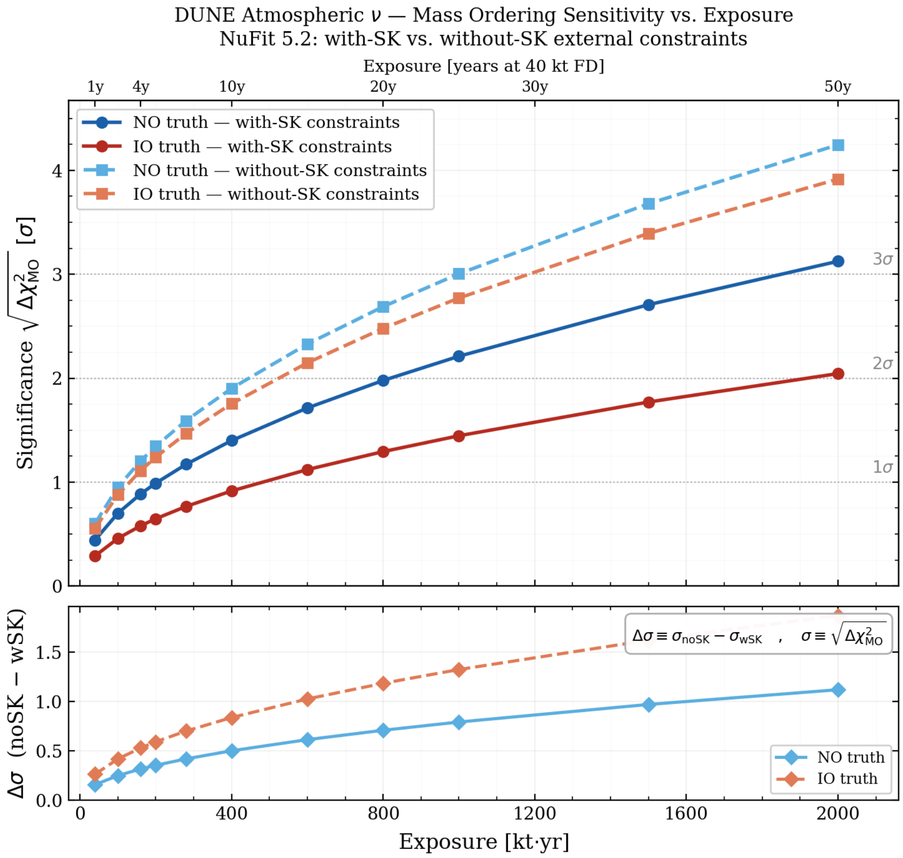

Figure 5.8: Comparison of mass ordering sensitivity between the two NuFit 5.2 configurations. The upper panel shows all four sensitivity curves (NO and IO truth, with and without SK atmospheric data). The lower panel shows $\Delta\sigma = \sigma_{\rm noSK} - \sigma_{\rm wSK}$ for each true ordering.

and the external constraints assumed. The sensitivity depends not only on the MSW matter effect but also on the θ_{23} octant assumed for each ordering, illustrating the interplay between these two physical effects. The analysis presented here includes only oscillation related systematics, so these results represent an upper bound on the atmospheric only contribution. In practice, the atmospheric measurement will be combined with DUNE's beam based mass ordering analysis to strengthen the overall determination.

Conclusion

As the next generation of long baseline neutrino oscillation experiment, DUNE is positioned to address several of the open questions in modern particle physics, from the neutrino mass ordering to the extent of CP violation in the leptonic sector. DUNE's atmospheric neutrino program is complementary to its long baseline beam program, providing independent sensitivity to oscillation parameters such as $\sin^2\theta_{23}$ and Δm^2_{32}. Because of the wide range of L/E coverage spanning roughly 10^{-1} to 10^5 km/GeV, atmospheric neutrinos are sensitive to multiple oscillation regimes simultaneously, allowing the full shape of the oscillation curves to constrain the atmospheric sector parameters rather than the narrow L/E window accessible to a fixed baseline beam experiment. Atmospheric neutrinos also carry a different systematic budget than beam neutrinos for the mass ordering measurement: while the beam determination relies on robust modeling of beam flux and cross sections, atmospheric mass ordering sensitivity is primarily set by the zenith angle shape of the sample. Furthermore, matter effects vary continuously with baseline, and atmospheric neutrinos probe everything from short downward going trajectories to upward going neutrinos that traverse the core of the Earth, providing a path to the mass ordering that is independent of the beam program.

In this thesis, we performed Asimov sensitivity studies using the GUNDAM fitting framework at a 400 kt·yr exposure, corresponding to ten years of data collection with the full four module DUNE far detector. The injected truth oscillation parameters were taken from the NuFit 5.2 global fit with Super-Kamiokande atmospheric data (Table 3.1). The analysis configuration uses a common binning of 260 reconstructed level analysis bins, with 160 bins in the ν_e selection and 100 bins in the ν_μ selection, each binned in reconstructed energy and in 10 uniform bins of the cosine of the reconstructed zenith angle. We use the Barlow-Beeston likelihood, in which per bin Monte Carlo statistical uncertainties are absorbed into analytically profiled scaling factors. The only free parameters in the fit are the six oscillation parameters, and Minuit2 is used to minimize the negative log likelihood over them. For our 1D profile likelihood scans, the scanned parameter was held fixed at each grid point while Minuit2 profiled over the remaining oscillation parameters, yielding the minimum $\Delta\chi^2$ at each point. For our 2D profile likelihood scans, the two scanned parameters were fixed at each grid point and Minuit2 profiled over the remaining oscillation parameters, yielding the $\Delta\chi^2$ surface from which confidence intervals were extracted. The scope of this thesis is a statistics only sensitivity study, with no flux, cross section, or detector systematic uncertainties included in the fit. This choice isolates the statistical power of the DUNE atmospheric neutrino sample at 400 kt·yr and provides an upper bound on the achievable sensitivity. A full systematic treatment for the DUNE atmospheric sample is an active area of development within the collaboration and is left to future work.

Our convergence study established that the 300×300 (true energy $\times$ true zenith) oscillation probability grid used as the GUNDAM default is sufficient to capture atmospheric oscillation features. Against the 4500×300 reference, the 300×300 grid agrees to better than 0.01% in the $\Delta\chi^2 = 1$ width on the prior dominated parameters $\sin^2\theta_{12}$, $\sin^2\theta_{13}$, and Δm^2_{21}, to 1.2% on $\sin^2\theta_{23}$, and to 2.4% on Δm^2_{32}, the largest of the five and still small compared to the projected precision on Δm^2_{32} from DUNE atmospheric data at this exposure. A coarser 150×300 grid was shown to fail this test, with the recovered Δm^2_{32} and $\sin^2\theta_{23}$ profile widths biased by 3.1% and 1.6% respectively and with pointwise $\Delta\chi^2$ deviations reaching 57 units on $\sin^2\theta_{23}$. Convergence was confirmed through comparisons of charged current event rates, 2D residuals, and profile likelihood scans across five grid resolutions. We further established that fixing the atmospheric neutrino production height at 20 km, rather than averaging over the full height distribution, biases the predicted event rate by only -1.294% in the ν_μ sample and -0.014% in the ν_e sample, justifying the use of the fixed height approximation in all subsequent fits.

The 1D profile likelihood scans show that the parameters DUNE atmospheric neutrinos meaningfully constrain are Δm^2_{32} and $\sin^2\theta_{23}$. Δm^2_{32} is the most tightly constrained of the six oscillation parameters, with the $\Delta\chi^2 = 1$ crossing well within the NuFit 5.2 prior band, reflecting the direct imprint of the oscillation frequency on the L/E distribution. $\sin^2\theta_{23}$ exhibits a visible asymmetry in its scan, reflecting partial but incomplete resolution of the octant degeneracy through three flavor and matter effects. δ_{CP} receives only a weak, broad constraint in this fit, with a post fit Hesse error roughly six times the NuFit 5.2 prior width (Figure 3.8). This apparent sensitivity is itself partly an artifact of the statistics only configuration: flux, cross section, and detector systematics would absorb much of it in a realistic fit. An independent measurement of δ_{CP} from atmospheric data alone is therefore not realistic at this exposure, and the value of the atmospheric channel for δ_{CP} lies in its eventual combination with the DUNE beam program. The solar parameters Δm^2_{21} and $\sin^2\theta_{12}$, along with $\sin^2\theta_{13}$, are entirely prior dominated in the GUNDAM atmospheric fit configuration, and the recovery of clean parabolas centered on truth for these parameters serves as a consistency check on the fit machinery.

The 2D profiled likelihood contours map the correlations between the atmospheric sector parameters. The $\sin^2\theta_{23}$ versus Δm^2_{32} contour closes at 1σ around the injected truth, with the profiled best fit recovering $\sin^2\theta_{23} = 0.4531$ and $\Delta m^2_{32} = 2.506 \times 10^{-3}$ eV^2, in close agreement with the NuFit 5.2 injection. The δ_{CP} versus $\sin^2\theta_{23}$ and δ_{CP} versus Δm^2_{32} contours both confirm that δ_{CP} is poorly localized by the atmospheric sample alone: the profiled best fit sits within a fraction of a radian of the injected NuFit 5.2 truth, but the contours extend across most of the scanned δ_{CP} range.

The mass ordering sensitivity reaches 3 to 4σ at 2000 kt·yr, depending on the true ordering and the choice of external constraints. Using NuFit 5.2 priors that include Super-Kamiokande atmospheric data, DUNE atmospheric neutrinos achieve approximately 3.2σ assuming true inverted ordering and 2.0σ assuming true normal ordering. Removing the SK constraints from the priors raises both significances and reverses their relative ordering, yielding 4.3σ for true NO and 3.9σ for true IO. This reversal reflects a competition between two physical effects: the MSW resonance, which favors NO truth because the matter enhancement falls in the higher statistics neutrino channel, and the θ_{23} octant assumed for each ordering, which favors IO truth in the with-SK configuration because NuFit 5.2 places NO in the lower octant and IO in the upper octant. The interplay of these two effects, rather than either alone, sets DUNE's atmospheric mass ordering reach.

The results presented in this thesis should be read as an upper bound on the achievable sensitivity of DUNE's atmospheric neutrino program. A full systematic treatment incorporating Honda flux uncertainties, GENIE cross section uncertainties, and LArTPC detector response uncertainties will degrade these statistics only sensitivities by an amount that remains to be quantified within the collaboration. Several extensions of this work are natural next steps. The grid convergence study in Chapter 4 used uniform binning in true $\cos\theta_z$, but a non uniform zenith binning aligned with the core mantle boundary may improve sensitivity in the energy and baseline regime where the MSW resonance enhances the oscillation probability. A dedicated octant sensitivity study, complementary to the mass ordering work in Section 5.3, would quantify how much of the residual octant asymmetry seen in the $\sin^2\theta_{23}$ versus Δm^2_{32} contour can be resolved at higher exposures. The injected truth values can also be updated to NuFit 6.1 once those constraints are available, and the analysis rerun for direct comparison with the most recent global fit. Finally, the principal value of the DUNE atmospheric neutrino program lies in its complementarity with the long baseline beam measurement: atmospheric neutrinos provide independent sensitivity to Δm^2_{32} and θ_{23}, an independent route to the mass ordering through Earth traversing matter effects, and a δ_{CP} cross check that does not share the beam program's flux and cross section systematics. Because atmospheric data collection can begin as soon as the first far detector modules are filled with liquid argon, well before beam commissioning is complete, the analysis framework, oscillation grid, and fitting procedures developed in this thesis position the DUNE atmospheric neutrino program to deliver its first oscillation results during the experiment's earliest physics capable phase.

Bibliography

[1] Abdurashitov, J. N. et al. (1999). Measurement of the solar neutrino capture rate with gallium metal. *Phys. Rev. C*, 60:055801.

[2] Abe, K. et al. (2018). Hyper-Kamiokande Design Report. *arXiv preprint*.

[3] Abe, K. et al. (2023). Measurements of neutrino oscillation parameters from the T2K experiment using 3.6×10^{21} protons on target. *Eur. Phys. J. C*, 83:782.

[4] Abi, B. et al. (2020a). Deep Underground Neutrino Experiment (DUNE), Far Detector Technical Design Report, Volume II: DUNE Physics. *arXiv preprint*.

[5] Abi, B. et al. (2020b). Long-baseline neutrino oscillation physics potential of the DUNE experiment. *Eur. Phys. J. C*, 80(10):978.

[6] Acciarri, R. et al. (2016). Long-Baseline Neutrino Facility (LBNF) and Deep Underground Neutrino Experiment (DUNE) Conceptual Design Report Volume 2.

[7] Acero, M. A. et al. (2022). Improved measurement of neutrino oscillation parameters by the NOvA experiment. *Phys. Rev. D*, 106:032004.

[8] Aguilar, M. et al. (2015). Precision Measurement of the Proton Flux in Primary Cosmic Rays from Rigidity 1 GV to 1.8 TV with the Alpha Magnetic Spectrometer on the International Space Station. *Phys. Rev. Lett.*, 114:171103.

[9] Ahmad, Q. R. et al. (2002). Direct Evidence for Neutrino Flavor Transformation from Neutral-Current Interactions in SNO. *Phys. Rev. Lett.*, 89:011301.

[10] Alvarez-Muñiz, J., Cao, Z., Katz, U., Mertsch, P., and Spiering, C. (2024). Cosmic Rays. *Phys. Rev. D*, 110:030001. Chapter 30 in the Review of Particle Physics.

[11] Andreopoulos, C., Barry, C., Dytman, S., Gallagher, H., Golan, T., Hatcher, R., Perdue, G., and Yarba, J. (2015). The genie neutrino monte carlo generator: Physics and user manual.

[12] ArgoNeuT Collaboration (2014). Calorimetry for low-energy electrons using charge and light in liquid argon. *Phys. Rev. D*, 89:112003.

[13] Asai, S. et al. (2024). Exploring the quantum universe: Pathways to innovation and discovery in particle physics. Report of the 2023 Particle Physics Project Prioritization Panel (P5).

[14] Ashie, Y. et al. (2004). Evidence for an oscillatory signature in atmospheric neutrino oscillation. *Phys. Rev. Lett.*, 93:101801.

[15] Barlow, R. J. and Beeston, C. (1993). Fitting using finite Monte Carlo samples. *Comput. Phys. Commun.*, 77:219–228.

[16] Barr, G. D., Gaisser, T. K., Lipari, P., Robbins, S., and Stanev, T. (2004). Three-dimensional calculation of atmospheric neutrinos. *Phys. Rev. D*, 70:023006.

[17] Battistoni, G., Ferrari, A., Montaruli, T., and Sala, P. R. (2003). The FLUKA atmospheric neutrino flux calculation. *Astropart. Phys.*, 19:269–290.

[18] Becker-Szendy, R. et al. (1992). Electron- and muon-neutrino content of the atmospheric flux. *Phys. Rev. D*, 46:3720–3724.

[19] Bethe, H. and Peierls, R. (1934). The 'neutrino'. *Nature*, 133:532.

[20] Casper, D. et al. (1991). Measurement of atmospheric neutrino composition with the IMB-3 detector. *Phys. Rev. Lett.*, 66:2561–2564.

[21] Cecchini, S. and Spurio, M. (2012). Atmospheric muons: experimental aspects. *Geoscientific Instrumentation, Methods and Data Systems*, 1:185–196.

[22] Cervera, A., Donini, A., Gavela, M. B., Gomez Cadenas, J. J., Hernandez, P., Mena, O., and Rigolin, S. (2000). Golden measurements at a neutrino factory. *Nucl. Phys. B*, 579:17–55. Erratum: Nucl. Phys. B 593 (2001) 731.

[23] Cowan, C. L., Reines, F., Harrison, F. B., Kruse, H. W., and McGuire, A. D. (1956). Detection of the Free Neutrino: a Confirmation. *Science*, 124(3212):103–104.

[24] Cowan, G., Cranmer, K., Gross, E., and Vitells, O. (2011). Asymptotic formulae for likelihood-based tests of new physics. *Eur. Phys. J. C*, 71:1554.

[25] Danby, G., Gaillard, J.-M., Goulianos, K., Lederman, L. M., Mistry, N., Schwartz, M., and Steinberger, J. (1962). Observation of High-Energy Neutrino Reactions and the Existence of Two Kinds of Neutrinos. *Phys. Rev. Lett.*, 9(1):36–44.

[26] Davis, Jr., R., Harmer, D. S., and Hoffman, K. C. (1968). Search for Neutrinos from the Sun. *Phys. Rev. Lett.*, 20(21):1205–1209.

[27] Dolinski, M. J., Poon, A. W. P., and Rodejohann, W. (2019). Neutrinoless Double-Beta Decay: Status and Prospects. *Ann. Rev. Nucl. Part. Sci.*, 69:219–251.

[28] DUNE Collaboration (2016). Long-Baseline Neutrino Facility (LBNF) and Deep Underground Neutrino Experiment (DUNE): Conceptual Design Report, Volume 2: The Physics Program for DUNE at LBNF.

[29] DUNE Collaboration (2021). Deep Underground Neutrino Experiment (DUNE) Near Detector Conceptual Design Report. *Instruments*, 5(4):31.

[30] DUNE Collaboration (2022). Separation of track- and shower-like energy deposits in ProtoDUNE-SP using a convolutional neural network. *Eur. Phys. J. C*, 82:903.

[31] DUNE Collaboration (2023). Reconstruction of interactions in the ProtoDUNE-SP detector with Pandora. *Eur. Phys. J. C*, 83:618.

[32] DUNE Collaboration (2026). Reconstruction of atmospheric neutrinos in DUNE's horizontal-drift far-detector module. *arXiv preprint*.

[33] Dziewonski, A. M. and Anderson, D. L. (1981). Preliminary reference Earth model. *Phys. Earth Planet. Inter.*, 25:297–356.

[34] Esteban, I., Gonzalez-Garcia, M. C., Maltoni, M., Schwetz, T., and Zhou, A. (2020). The fate of hints: updated global analysis of three-flavor neutrino oscillations. *JHEP*, 09:178.

[35] Esteban, I., Gonzalez-Garcia, M. C., Maltoni, M., Schwetz, T., and Zhou, A. (2022). NuFIT 5.2 (2022). `http://www.nu-fit.org`.

[36] Fermi, E. (1934). Versuch einer Theorie der β-Strahlen. I. *Z. Phys.*, 88:161–177.

[37] Freund, M. (2001). Analytic approximations for three neutrino oscillation parameters and probabilities in matter. *Phys. Rev. D*, 64:053003.

[38] Fukuda, Y. et al. (1998). Evidence for Oscillation of Atmospheric Neutrinos. *Phys. Rev. Lett.*, 81:1562–1567.

[39] Fukugita, M. and Yanagida, T. (1986). Baryogenesis Without Grand Unification. *Phys. Lett. B*, 174:45–47.

[40] Gaisser, T. K. and Honda, M. (2002). Flux of atmospheric neutrinos. *Ann. Rev. Nucl. Part. Sci.*, 52:153–199.

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

[42] Hampel, W. et al. (1999). GALLEX solar neutrino observations: Results for GALLEX IV. *Phys. Lett. B*, 447:127–133.

[43] Hirata, K. S. et al. (1988). Experimental study of the atmospheric neutrino flux. *Phys. Lett. B*, 205:416–420.

[44] Hirata, K. S. et al. (1989). Observation of ^{8}B Solar Neutrinos in the Kamiokande-II Detector. *Phys. Rev. Lett.*, 63:16.

[45] Honda, M., Athar, M. S., Kajita, T., Kasahara, K., and Midorikawa, S. (2015). Atmospheric neutrino flux calculation using the NRLMSISE-00 atmospheric model. *Phys. Rev. D*, 92:023004.

[46] Honda, M., Kajita, T., Kasahara, K., and Midorikawa, S. (2004). New calculation of the atmospheric neutrino flux in a three-dimensional scheme. *Phys. Rev. D*, 70:043008.

[47] James, F. and Roos, M. (1975). Minuit: A System for Function Minimization and Analysis of the Parameter Errors and Correlations. *Comput. Phys. Commun.*, 10:343–367.

[48] Kajita, T. (2010). Atmospheric neutrinos and discovery of neutrino oscillations. *Proc. Jpn. Acad. Ser. B*, 86:303–321.

[49] Kodama, K. et al. (2001). Observation of tau neutrino interactions. *Phys. Lett. B*, 504:218–224.

[50] Marshall, J. S. and Thomson, M. A. (2015). The Pandora Software Development Kit for Pattern Recognition. *Eur. Phys. J. C*, 75:439.

[51] Mikheyev, S. P. and Smirnov, A. Y. (1985). Resonance amplification of oscillations in matter and spectroscopy of solar neutrinos. *Sov. J. Nucl. Phys.*, 42:913–917. [Yad. Fiz. 42, 1441 (1985)].

[52] Minkowski, P. (1977). $\mu \to e\gamma$ at a rate of one out of 10^9 muon decays? *Phys. Lett. B*, 67:421–428.

[53] Navas, S. et al. (2024). Review of Particle Physics. *Phys. Rev. D*, 110:030001.

[54] Nobel2002 (2002). The Nobel Prize in Physics 2002. `https://www.nobelprize.org/prizes/physics/2002/summary/`. Awarded to Raymond Davis Jr. and Masatoshi Koshiba.

[55] Nobel2015 (2015). The Nobel Prize in Physics 2015. `https://www.nobelprize.org/prizes/physics/2015/summary/`. Awarded to Takaaki Kajita and Arthur B. McDonald.

[56] Paul, H. (2014). Beta decay spectrum. Wikimedia Commons, CC BY-SA 4.0.

[57] Pauli, W. (1930). Open letter to the Gauverein meeting in Tübingen. Reprinted in W. Pauli, *Wissenschaftlicher Briefwechsel*, Vol. II, Springer, 1985.

[58] PauliQuote (1930). Wolfgang Pauli Quotes. `https://libquotes.com/wolfgang-pauli`. Historical attribution; exact phrasing may vary.

[59] Pontecorvo, B. (1957). Mesonium and Antimesonium. *Sov. Phys. JETP*, 6:429. Originally published in Zh. Eksp. Teor. Fiz. 33, 549 (1957).

[60] Schael, S. et al. (2006). Precision electroweak measurements on the Z resonance. *Phys. Rept.*, 427:257–454.

[61] Sherwood, P. et al. (2024). NuOscillator: A neutrino oscillation probability engine. `https://github.com/dbarrow257/NuOscillator`.

[62] Wilks, S. S. (1938). The Large-Sample Distribution of the Likelihood Ratio for Testing Composite Hypotheses. *Ann. Math. Statist.*, 9(1):60–62.

[63] Wolfenstein, L. (1978). Neutrino oscillations in matter. *Phys. Rev. D*, 17:2369–2374.

[64] Workman, R. L. et al. (2022). Review of Particle Physics.

[65] Yanagida, T. (1979). Horizontal gauge symmetry and masses of neutrinos. In *Proceedings of the Workshop on the Baryon Number of the Universe and Unified Theories, Tsukuba, Japan*, pages 95–99.