Imperial College London

Doctoral Thesis

First Bayesian Neutrino Oscillation Sensitivities of DUNE

Author:

Liban Warsame

Supervisors:

Dr. Patrick Dunne

Prof. Alex Tapper

A thesis submitted in fulfilment of the requirements

for the degree of Doctor of Philosophy

in the

High Energy Physics

Department of Physics

February 25, 2025

Abstract

The Deep Underground Neutrino Experiment (DUNE) is a next-generation long-baseline neutrino oscillation experiment which aims to make world-leading precision measurements. DUNE will produce an intense beam of neutrinos at the Fermi National Accelerator Laboratory (Fermilab) directed at the Sandford Underground Research Facility (SURF) in South Dakota, providing a 1285 km baseline. Neutrino interactions will be measured by a near detector facility at Fermilab and a far detector facility at SURF. The neutrino beam will be primarily composed of either muon neutrinos, ν_μ or antimuon neutrinos, $\bar{\nu}_\mu$. As a consequence of neutrino oscillation, DUNE will observe the disappearance of ν_μ and $\bar{\nu}_\mu$, along with the appearance of electron neutrinos, ν_e or antielectron neutrinos, $\bar{\nu}_e$. The observation of these neutrino oscillation channels provides DUNE with sensitivity to neutrino oscillation parameters, specifically θ_{23}, θ_{13}, Δm^2_{32} and δ_{CP}, as well as the neutrino mass ordering.

To ensure that the experiment will achieve its physics goals, sensitivity studies are performed based on predictions of the neutrino event distributions that DUNE will observe. The results of these studies will also motivate and justify design choices that have an impact on sensitivity. In addition, DUNE will measure an order of magnitude more neutrino events than current long-baseline experiments. As a result, DUNE's sensitivity will be limited by systematic uncertainty, rather than statistical uncertainty. It is therefore imperative that these are modelled correctly in these studies to ensure accurately predicted sensitivities. Performing a Bayesian analysis allows for a more efficient method of studying the behaviour of systematic uncertainty parameters.

This thesis presents the first Bayesian neutrino oscillation sensitivities of DUNE. The Markov Chain Monte Carlo methods and Bayesian techniques used in this analysis will be described, as well as the implementation of systematic uncertainties which surpass that of previous DUNE sensitivity studies. The predicted sensitivities to the oscillation parameters will be shown, as well as the impact of systematic uncertainties. The deficiencies in the current systematic model are also discussed, with a demonstration of the level of complexity required in future analyses.

Declaration

The analyses presented in this thesis are my own work as a member of the DUNE collaboration. The work contributed by others is properly cited and referenced.

Chapters 1 and 2 provide experimental background and an overview of neutrino physics. This work contributed by previous researchers is properly cited throughout. Chapter 3 introduces the Deep Underground Neutrino Experiment (DUNE) and cites technical and conceptual reports produced by the DUNE collaboration.

Chapter 4 describes the statistical techniques used in this thesis. These techniques are commonly used in the scientific community, as well as in industry. The citations for the techniques have been provided. Chapter 5 describes the advancement of a neutrino oscillation calculator designed for parallel CPU computations. The work carried out by the creator of this software has been properly cited, along with other relevant work in the field of oscillation probability calculation.

Chapter 6 describes and presents the results of the far detector sensitivity studies. The Monte Carlo simulation, flux model and cross-section model uncertainties were developed by the DUNE Collaboration. However, both the cross-section and flux systematics were implemented in a novel way compared with previous DUNE sensitivity studies, which is my own work. The detector systematic uncertainties were implemented by myself, following a similar model to previous studies but again implemented in a novel way. The Markov Chain Monte Carlo framework used to produce the sensitivities in this thesis was written by members of the T2K, SK and DUNE collaborations over a number of years. I have added improvements to this framework throughout my PhD. The production of physics plots and sensitivities in Chapter 6 is wholly my own work.

Chapter 7 describes and presents the results of the joint far and near detector sensitivity studies. Similarly to Chapter 6, the near-detector Monte Carlo simulation was produced by members of the DUNE collaboration. The final implementation of the detector systematics

for the near detector was also used in the previous sensitivity study. The study showing the constraining power of the ND and the effect degeneracy has on systematic constraints was wholly my own work, as well as all physics plots and sensitivities.

Acknowledgements

Despite the challenges, this PhD has been a wonderful experience, through what I've learned, what I've achieved and the people I've met along the way. The work in this thesis would not have been possible without the help, advice and support of many.

First and foremost I would like to thank Patrick Dunne for being an incredible supervisor. The help he's given me over the past 4 years has been invaluable. He never fails to solve a problem and always ensures your thoughts and feelings are considered. Thanks for answering all of my seemingly random questions. A special shoutout for reading this thesis, chapter by chapter, while on paternity leave. I wish nothing but the best for your new family.

A sincere thanks to the rest of the Imperial College London long-baseline group and the high-energy physics group in general. Thanks to Ed Atkin, for sitting through all the compilation errors in MaCh3 and rants about ROOT plotting scripts. To my fellow students, I've had countless engaging discussions, fun socials and laughs. Special thanks to Marvin, Jay, Naseem, Anežka, Rohan, Charlotte, Tiago and Charlie (the postdoc amongst us). Life in the Imperial group has been merry.

Being a member of the DUNE collaboration has been an eye-opening experience. To work alongside so many brilliant researchers has been a privilege. None of this work would have been possible without the likes of Luke Pickering, Callum Wilkinson and Chris Marshall. Thanks for always responding to my Slack questions, I imagine I'll still be referring to those documented insights for years to come (as I do now).

To all my friends who have been there for me in the tough times and the great ones. I can finally say you no longer have to hear me moan about writing this thesis. To mention some special names, Manjosh, Niall, Zsolt, Kiran, Karolina, Robert, Rowan, Kajal and Hamza. Thanks for keeping me sane and I look forward to many more years of friendship.

Finally, to my family, especially my mum, dad and brother. They've been the greatest pillar of support. I wouldn't have been in this position without them.

Contents

List of Figures

List of Tables

List of Abbreviations

APA Anode Plane Assembly.

CC Charged Current.

CNN Convolutional Neural Network.

CP Charge-Parity.

CP Charge-Parity.

CPA Cathode Plane Assembly.

CRP Charge Readout Planes.

CVN Convolutional Visual Network.

DIS Deep Inelastic Scattering.

DP Dual-Phase.

DUNE Deep Underground Neutrino Experiment.

ECAL Electromagnetic Calorimeter.

FD Far Detector.

FHC Forward Horn Current.

FSI Final State Interactions.

GAr Gaseous Argon.

GEM Gaseous Electron Multiplier.

GENIE Generates Events for Neutrino Interaction Experiment.

HPD Highest Posterior Density.

HPgTPC High-Pressure Gas Time Projection Chamber.

IMB Irvine-Michigan-Brookhaven.

IO Inverted Ordering.

LAr Liquid Argon.

LBNF Long-Baseline Neutrino Facility.

LHC Large Hadron Collider.

MC Monte Carlo.

MCMC Markov Chain Monte Carlo.

MSW Mikheyev-Smirnov-Wolfenstein.

MWPC Multi-wire Proportional Chamber.

NC Neutral Current.

ND Near Detector.

NO Normal Ordering.

PCB Printed Circuit Board.

PIP Proton Improvement Plan.

PMNS Pontecorvo-Maki-Nakagawa-Sakata.

PRISM Precision Reaction-Independent Spectrum Measurement.

QE Quasielastic.

RES Resonant.

RHC Reverse Horn Current.

RPA Random Phase Approximation.

SAND System for on-Axis Neutrino Detection.

STT Straw Tube Trackers.

SURF Sandford Underground Research Facility.

SiPM Silicon Photomultipliers.

TDR Technical Design Report.

TMS Temporary Muon Spectrometer.

TPC Time Projection Chamber.

WLS Wavelength-Shifting.

Chapter 1

Introduction

The universe as we observe it is overwhelmingly composed of matter, with only a negligible presence of antimatter. This asymmetry between matter and antimatter is one of the most profound mysteries in modern physics. According to the Big Bang theory, the universe should have produced matter and antimatter in equal amounts. The standard model of physics does not currently explain this asymmetry. Charge-parity (CP) violating processes are those which occur at different rates for matter and antimatter and are necessary to explain the asymmetry. CP violation has already been observed in the quark sector [1] but is orders of magnitude too small to provide an explanation.

Neutrinos are one of the most abundant particles in the universe, but very rarely interact, making them a challenging particle to study. A groundbreaking discovery in the late 20th century revealed that neutrinos have mass and can oscillate between their flavours. The discovery of neutrino oscillation opened a new door that could help to explain the matter-antimatter asymmetry. The question was asked: Do neutrino oscillations violate CP symmetry? There are some indications of CP violation in the neutrino sector from current generation experiments, but no definitive evidence. A new generation of neutrino experiments is in development, intending to conclusively determine whether neutrino oscillations violate CP symmetry.

The Deep Underground Neutrino Experiment (DUNE) is a long-baseline neutrino exper-

iment which aims to precisely measure the parameters in the model that describes neutrino oscillations. In doing so, DUNE also aims to definitively determine whether neutrino oscillations violate CP symmetry. DUNE will do this by producing the world's most powerful neutrino beam at Illinois and firing the neutrinos 1285 km west towards South Dakota. Using several detectors, both close and far from the neutrino production site, neutrino interactions will be collected and studied to determine the oscillation parameter values. At the time of writing, DUNE is under construction and will begin to collect events from the neutrino beam at the beginning of the next decade.

Although the experiment is not running yet, it is important to determine how well we expect DUNE to perform in measuring the neutrino oscillation parameters. We need to confirm that DUNE will achieve its physics goals, as well as help motivate and justify design choices for the neutrino beam and detectors. DUNE will collect an order of magnitude more neutrino events than current long-baseline experiments, and so DUNE's performance will be determined by the size of the systematic uncertainty. A frequentist sensitivity study of DUNE was performed for the technical design report (TDR) [2]. This thesis will present the first Bayesian sensitivity study of DUNE. It uses the same inputs, i.e. event simulation, as the TDR analysis, with a more comprehensive implementation of the systematic uncertainty. Given that DUNE's sensitivity will be limited by systematic uncertainty, it is vital that these are implemented correctly in analysis frameworks to produce accurate predictions of DUNE's performance. In addition, taking a Bayesian approach allows for a more efficient method to study the behaviour of systematic uncertainty parameters, as well as new quantities that are not explicitly included in the oscillation model.

This thesis begins with an overview of neutrino physics in Chapter 2. The discovery of neutrinos and the history of neutrino physics are briefly reviewed. The theory of neutrino interactions in the Standard Model is described, as well as neutrino-nucleon interactions which are especially relevant in long-baseline experiments. Finally, the theory of neutrino oscillation is covered, including its discovery and our current best knowledge of the parameters that

characterise the oscillation.

Chapter 3 describes DUNE, covering each component that comprises the experiment. Initially, the production, properties and characteristics of the neutrino beam are described. Following this, an overview of the principal technology used in DUNE, liquid argon time projection chambers, is provided. The near detector complex and the far detector are then described, as well as a brief overview of DUNE's physics potential.

The analysis described in this chapter uses Markov Chain Monte Carlo and Bayesian inference to make estimations of neutrino oscillation parameters. Chapter 4 describes these concepts, starting with the basics of Bayesian statistics. A brief description of Monte Carlo methods is provided, followed by a more detailed overview of Markov Chain Monte Carlo. Finally, the methods for interring parameter estimates and uncertainties from the full posterior distribution are described.

Calculating oscillation probabilities is a computationally costly component of a neutrino oscillation analysis. The analysis in this thesis uses an oscillation calculator designed to efficiently parallelize the operations on CPU resources. Chapter 5 begins with a description of oscillation probability calculation and the difficulties involved. The calculator used in this analysis is then outlined and the performance of the calculator is presented as well as the resulting motivation for its use in this analysis.

The techniques described in Chapter 4 are used to perform a sensitivity study of the DUNE far detector and are presented in Chapter 6. Firstly, the event selection criteria are described, followed by the resulting Monte Carlo prediction of the event distributions expected given an assumed exposure. A detailed description of the systematic uncertainties considered in this analysis is provided, as well as the effect these uncertainties have on the event distributions. A brief description of the validations that were performed on the framework is provided. The resulting oscillation parameter sensitivities are then presented, from studies with and without the systematic uncertainties. The effect of the systematic uncertainties on the sensitivities is illustrated, as well as the effect of changing the prior information, specifically introducing a

constraint on the θ_{13} mixing angle from reactor neutrino experiments.

The near detector in long-baseline experiments collects neutrino events before significant oscillations occur, allowing us to establish more stringent constraints on the systematic uncertainties. Chapter 7 presents a joint far and near detector sensitivity study. Initially, the near detector design used in this analysis and event selection criteria are described. This is followed by the Monte Carlo predictions and the systematic uncertainties associated with the near detector samples. Due to its proximity to the neutrino beam source, the near detector collects significantly more events than the far detector. As well as this, the unprecedented intensity of DUNE's neutrino beam will result in the collection of far more near-detector events than current long-baseline experiments. Although this is generally desirable, the sheer number of events can make analyses susceptible to over-constraining systematic uncertainties. This occurs if the uncertainty models implemented in the analysis do not accurately represent the real uncertainties in the experiment. This issue is discussed, and a study is presented illustrating the necessity of including the true degeneracy between parameters in the systematic models. Finally, the results of the joint far and near detector sensitivity study are presented, including comparisons with the far detector study.

The thesis concludes with Chapter 8, reiterating the key results from the sensitivity studies. The next steps required to produce more realistic sensitivities for DUNE are also discussed.

Chapter 2

Neutrino Physics

The neutrino is thought to be the most abundant particle in the universe after the photon (i.e. light) and yet one of the least understood in our model of particle physics. Their properties make them extremely difficult to detect, and hence they earned the appropriate nickname the "Ghost Particle". In the nearly 100 years since their existence was first postulated, significant knowledge has been gained about neutrinos. They have three distinct flavours and one known property of theirs is their tendency to oscillate between them. Neutrino oscillation is the phenomenon this thesis aims to investigate.

This chapter begins with a brief history lesson, outlining the discovery of the neutrino and some of its key properties in Section 2.1. Section 2.2 will discuss neutrino interactions in the Standard Model of particle physics, including how the interaction terms developed over time, and introduce some relevant Standard Model theory to do so. Section 2.3 will discuss the phenomenon of neutrino oscillation, including its discovery, the theory behind it and our current knowledge of the parameters that govern it.

2.1 Discovery of Neutrinos

The story of neutrinos began in 1914 when James Chadwick made accurate measurements of electrons released in beta decay [3]. Unlike alpha or gamma decays, in which the resulting alpha and gamma particles have very narrow energy distributions, he found that the kinetic energy of the electrons followed a continuous spectrum, shown in Figure 2.1. Under the assumption that, like alpha and gamma, beta decay follows a two-body problem (i.e. a nucleus decays into another nucleus and an electron), a distribution of energies would violate the law of conservation of energy. Some members of the physics community rejected the result, and some even agreed with the violation of the law [1]. In 1930, Wolfgang Pauli wrote an infamous letter (addressed, in German but translating to, "Dear Radioactive Ladies and Gentlemen") [4], proposing a neutral, light, particle which he called the 'neutron' that was also emitted during beta decay. After Chadwick discovered the actual neutron in 1932, Enrico Fermi named the mysterious particle 'neutrino', a term which translates literally from Italian as the 'little neutral one'.

Enrico Fermi developed a theory of beta decay in 1933, consisting of the decay of a neutron into a proton, an electron and a neutrino [5]:

$$n \rightarrow p + e^{-} + \bar{\nu} \tag{2.1}$$

The theory solved the issue of the continuous energy spectrum of the electron since the available energy from the decay was now shared between the electron and the neutrino. Fermi's theory correctly predicted almost all of the observed characteristics of beta decay, resulting in a consensus that the neutrino did indeed exist.

Shortly after, Bethe and Peierels suggested that if a neutrino can be created from beta decay, then it can also be captured [6]. They proposed that a neutrino could hit a nucleus,

[1] Niels Bohr, considered one of the founders of quantum theory, suggested that energy was only conserved on average, but could be violated in any given decay.

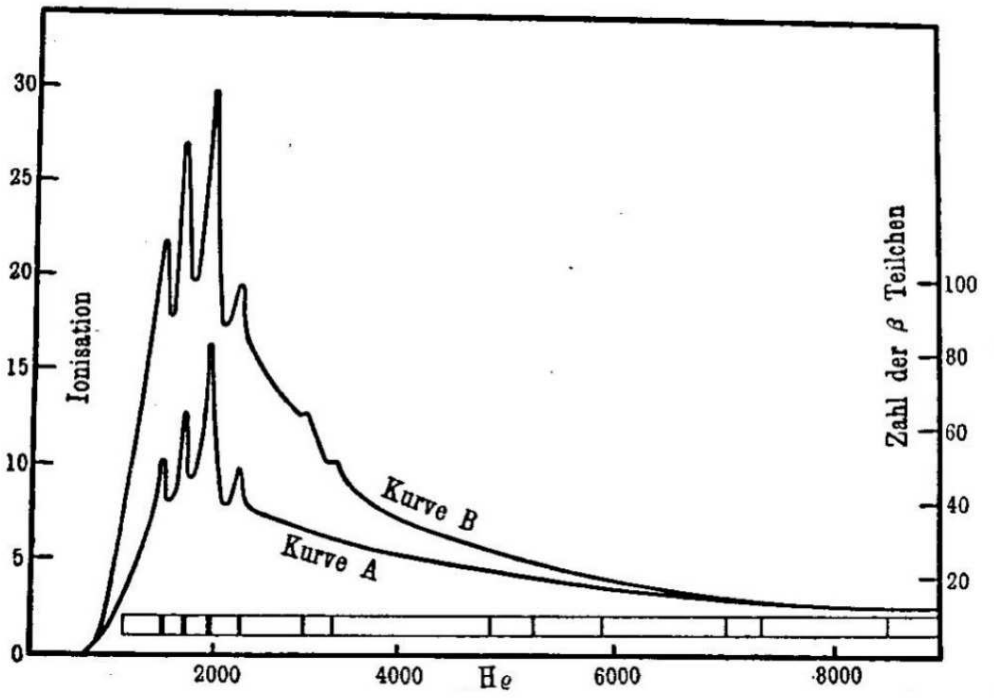

Figure 2.1: Number of β particles as a function of energy found by Chadwick. A continuous spectrum can be seen in Curve A, which uses a Geiger counter, and Curve B, which uses an ionization chamber [3].

producing an electron/positron and changing the charge of the nucleus by 1, e.g.

$$\bar{\nu} + p \rightarrow n + e^{+} \tag{2.2}$$

However, using Fermi's theory, they estimated the interaction cross-section (i.e. the probability of an interaction/process to take place) and concluded that observing the neutrino was impractical. This statement held until 1956, 26 years after Pauli first postulated its existence when the first observation of a neutrino was confirmed by Reines, Cowan, et al. [7]. The development of fission reactors provided an intense source of neutrinos, counterbalancing the unlikeliness of the interaction. By using tanks and liquid scintillators, electron antineutrinos emitted from a nuclear reactor were detected. When the neutrinos interacted with the water via inverse beta decay (Eq 2.2), the resulting e^{+} would annihilate producing two gamma rays

that could be detected in the liquid scintillator. The emitted neutron could also be detected via n capture on cadmium added to the water, resulting in a delayed gamma ray. The occurrence of both signals provided a clear event signature.

The experiment was initially deployed in 1953 at the site of a nuclear reactor in Hanford, Washington [8]. However, both cosmic and fission-induced backgrounds prevented a definitive observation. The experiment was relocated to the Savannah River Plant, South Carolina, allowing the detector to be placed underground which provided significant shielding from the cosmic background [7, 9]. The experiment was also redesigned to have a larger target volume and better signal-background separation. The neutrino interaction signal was 20 times greater than the accidental background, allowing for a confirmation of the existence of neutrinos in 1956.

At the same time, an experiment was conducted at the Brookhaven reactor, New York, by Davis et al., which attempted to measure antineutrinos via interactions with Chlorine:

$$\bar{\nu} + \mathrm{Cl} \rightarrow e^{-} + \mathrm{Ar} \tag{2.3}$$

Tanks filled with tetrachloromethane were placed close to the reactor, however, no significant excess of Argon was found. This result, in conjunction with the observation of antineutrinos from positrons by Reines and Cowan, initiated the concept of lepton number conservation i.e. an antineutrino cannot produce an electron. The interaction in Eq 2.3 would require a neutrino, not an antineutrino, a process called reverse electron capture.

The discovery of the muon, μ, in 1936 [10, 11] raised a new question: Are neutrinos associated with electrons identical to those produced alongside muons? In 1959, Pontecorvo argued that they were different particles [12] and just 3 years later this was confirmed by Lederman, Schwartz, Steinberger, et al. at Brookhaven National Laboratory [13]. A beam of protons was fired at a Beryllium target producing secondary pions and kaons. At the time, it was known

that the pions decayed into muons and neutrinos, $\pi \rightarrow \mu\nu$, but it was not certain that these neutrinos were different from those involved in beta decay. After observing the interactions of the emitted neutrinos in a spark chamber, they found an excess of muons compared with electrons. Had the neutrinos been the same, it was expected that an equal number of muons and electrons would have been observed. Hence, the experiment simultaneously proved the existence of the ν_μ and its flavour independence with the ν_e. Subsequent experiments at CERN further validated the "two-neutrino hypothesis" with a substantial increase in recorded neutrino interactions [14, 15].

For decades following the discovery of the muon, the question lingered: Were there additional, heavier siblings of the electron? Between 1974 and 1977, a series of experiments conducted by Perl et al used the (at the time) new e^+e^- collider called SPEAR at SLAC resulting in the discovery of an even heavier lepton called the tau, τ [16]. Based on the symmetry of lepton and neutrino flavours observed thus far, there was strong speculation regarding the existence of a corresponding third neutrino. The DONUT experiment at Fermilab, Illinois, was built specifically to observe the postulated ν_τ and did so in 2000 [17]. The interactions of neutrinos in the detector were initially identified by the sudden appearance of tracks (as neutrinos themselves leave no discernible trace as they travel). Of those, ν_τ events were identified by a characteristic "kink" in the track after ~ 2 mm, which indicated the decay of a τ as shown in Figure 2.2. The common leptonic decays of a τ^- is $\tau^- \rightarrow \mu^-\bar{\nu}_\mu\nu_\tau$ or $\tau^- \rightarrow e^-\bar{\nu}_e\nu_\tau$, of which only the emitted lepton is observable resulting in the single kinked track. The result was made from only four ν_τ events, but since the particle was expected to exist the discovery was generally accepted.

Having found two further generations of leptons after the electron, many questioned whether there was a fourth. Some postulated that this pattern continued indefinitely for higher masses and energies. However, in 1989, the ALEPH detector, using the Large Electron Positron (LEP) collider ruled out the possibility of a fourth neutrino flavour at 98% C.L. via measurements of the Z-boson decay width [18]. A combined fit was also performed with four experiments at LEP

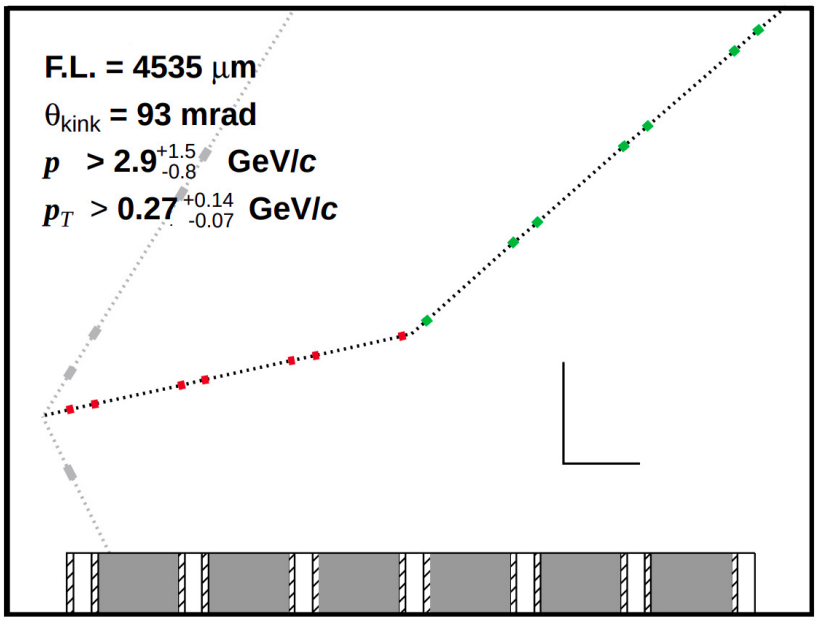

Figure 2.2: A ν_τ event at DONUT, showing the track of a τ (red-dotted) which decays into a lepton (green-dotted) causing a characteristic "kink" [16].

and the SLD experiment at the e^+e^- colliding ring at SLAC, resulting in a $N_\nu = 2.984 \pm 0.008$ [19]. This result is consistent with the existence of three light neutrino species. However, this does not rule out further species of neutrinos that do not couple to the Z (i.e. sterile neutrinos) or are too heavy to be produced in Z decay (i.e. heavy neutrinos).

2.2 Neutrino Interactions

2.2.1 Neutrino Interactions in the Standard Model

After the postulation of the neutrino by Pauli, Fermi developed a theory of weak interactions to describe beta decays, known as Fermi's theory. In Fermi's theory, beta decay was a point-like interaction of four fermions as shown Figure 2.3. The matrix element for the interaction in Fermi's theory was given by

$$\mathcal{M}_F = \frac{\mathrm{G}_F}{\sqrt{2}}\, g_{\mu\nu} \Big[\bar{\psi}_3 \gamma^\mu \psi_1\Big] \Big[\bar{\psi}_4 \gamma^\mu \psi_2\Big] \tag{2.4}$$

where G_F is the Fermi constant, $g_{\mu\nu}$ is the diagonal Minkowski metric tensor, ψ are the fermion fields and γ^μ are the gamma matrices. A description of ψ, $g_{\mu\nu}$ and γ^μ can be found in [20]. G_F denotes the coupling factor of the interaction. Fermi's theory largely drew from electromagnetic interactions and provides fairly reliable calculations at low energies. According to the theory, the cross-section for the interaction increases linearly with energy (forever...), highlighting that modifications were necessary, especially at higher energies.

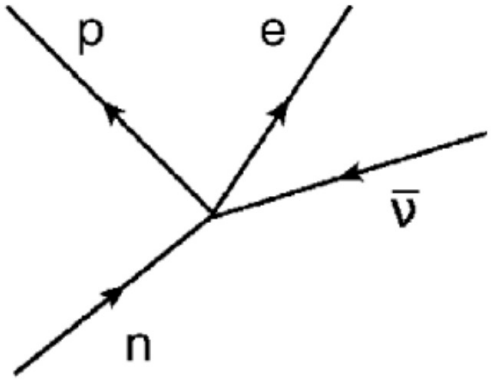

Figure 2.3: Feynman diagram of a point-like Fermi interaction.

In 1957, Wu found experimental evidence that beta decay violated parity conservation [21]. Parity conservation refers to a symmetry whereby the laws of physics remain unchanged under an inversion of the spatial coordinates (i.e. $\vec{x} \rightarrow -\vec{x}$). By observing the decays of Co^{60} in a magnetic field, it was noticed that the direction of the emitted electron was almost always opposite to the nuclear spin of the ^{60}Co. If parity were conserved, the electrons would be emitted isotropically. Parity was believed to be a fundamental symmetry in nature and holds for both electromagnetic and strong interactions, so the discovery shocked the physics community. The form of the interaction in Eq 2.4 does not violate parity, and hence providing further evidence modification to the theory was necessary.

Soon after, Feynman and Gell-Mann introduced a new form of interaction which maximally violated parity [22]. The general idea behind the new formalisation was to find a combination of gamma matrices which contained this violation. The derived form was given by

$$\frac{\mathrm{G}_F}{\sqrt{2}}\, g_{\mu\nu}\Big[\bar{\psi}_3\gamma^\mu\left(1-\gamma^5\right)\psi_1\Big]\Big[\bar{\psi}_4\gamma^\mu\left(1-\gamma^5\right)\psi_2\Big] \tag{2.5}$$

and is widely referred to as the "V - A" structure, since $\psi\gamma^\mu\bar{\psi}$ is a vector (occurring in electromagnetic interactions) and $\psi\gamma^\mu\gamma^5\bar{\psi}$ is an axial vector. It then follows that the weak interaction only couples to left-handed chiral particles (i.e. the direction of motion is opposite to that of the spin) since the left and right-handed chiral projection operators are defined as

$$P_L = \frac{1}{2}\left(1-\gamma^5\right); \qquad\qquad P_R = \frac{1}{2}\left(1+\gamma^5\right) \tag{2.6}$$

When applied to particles, the chiral projection operators return the corresponding projection, i.e. $P_L\psi = \psi_L$. For anti-particles, they return the opposite projection, i.e. $P_L\bar{\psi} = \bar{\psi}_R$. Hence, the weak force only couples to right-handed anti-particles. This suggests that, although parity is violated, the combination of charge and parity (CP) is still conserved. A CP operation on a left-handed neutrino returns a right-handed antineutrino. However, Cronin, Fitch, et al. found that CP is also violated by weak interactions [23]. The topic of CP violation will be discussed in more detail in Section 2.3.2.

As quantum electrodynamics, the theory governing electromagnetic interactions achieved success, efforts were made to formulate an analogous theory for weak interactions. This resulted in the development of a unified theory, called the electroweak theory, by Sheldon Glashow, Steven Weinberg, and Abdus Salam [24, 25]. The theory predicted that massive gauge bosons were necessary to mediate the weak force. The weak force now has two formulations: charged-current (CC) through the exchange of $W^\pm$ boson and neutral-current (NC) through the exchange of a Z^0 boson. Feynman diagrams for the two interactions can be seen in Figure 2.5. Unlike $W^\pm$ bosons, Z^0 bosons can couple to both left-handed and right-handed particles and

antiparticles. The introduction of the propagator term in the weak interaction, $\frac{1}{M_{W,Z}^2 - q^2}$ where $M_{W,Z}$ is the mass of the boson and q is the momentum transfer, suppresses the total cross-section at high energies. At low energies, $M_{W,Z}^2 >> q^2$, we return back to Fermi's theory.

2.2.2 Neutrino-Nucleon Interactions

Neutrino interactions on nucleons have a much larger cross-section than interactions on free particles but their description provides more theoretical challenges. These interactions can be further categorised. Firstly, all neutrino-nucleon interactions can be placed into two categories: elastic and inelastic. Elastic interactions dominate at low q^2 and cause the struck nucleon to recoil. In CC interactions, there is a change of charge since a charged lepton is produced. Due to the transfer of mass to the lepton, elastic CC interactions are referred to as quasielastic (QE).

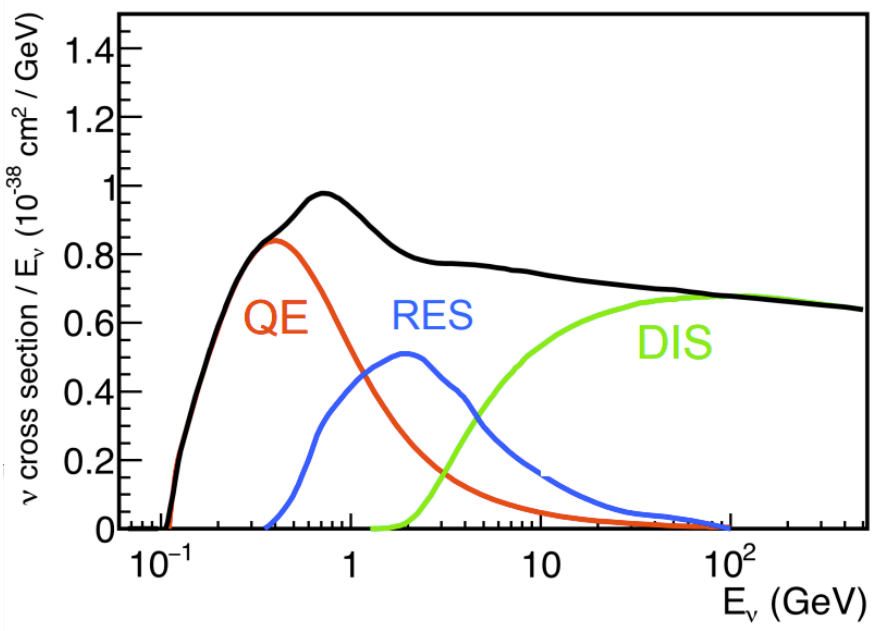

Figure 2.4: Total muon neutrino CC interaction cross-section on nucleon, and the three major components: QE, resonance (RES) and DIS, as a function of neutrino energy [26]

Inelastic interactions are dominated by two further groups: resonant (RES) and deep inelastic scattering (DIS). RES interactions dominate inelastic interactions at low q^2 and result in the excitation of the struck nucleon into a baryonic resonance, e.g. Δ. This excited state then

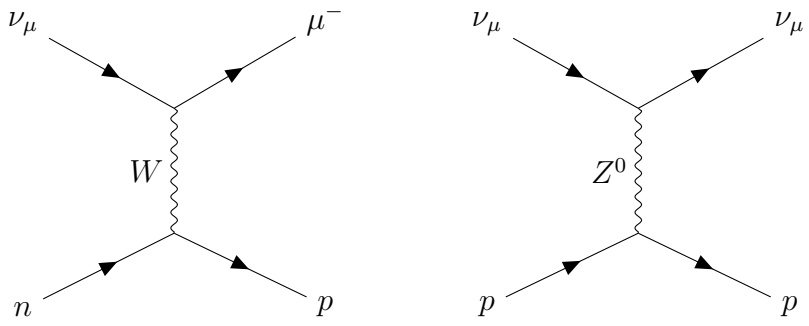

(a) CCQE scattering **(b)** NC elastic scattering

Figure 2.5: Example Feynman diagrams showing (a) CCQE scattering and (b) NC elastic scattering of a ν_μ.

decays, emitting further particles from interaction e.g. a pion. At higher q^2, DIS interactions dominate and are when the neutrino scatters directly off of a constituent quark. This often results in the fragmentation of the incident nucleon. Figure 2.4 shows the total muon neutrino CC interaction cross section as a function of neutrino energy, E_ν, and also the cross-section for the major three components: QE, RES, DIS.

NC Elastic and CC Quasi-Elastic Scattering

Figure 2.5 shows example Feynman diagrams for both CCQE and NC elastic scattering of a ν_μ on a neutron. At very low neutrino energies, NC elastic interactions are the only nucleon interactions available, since a threshold of energy is required in CCQE scattering to produce mass of the charged lepton. For a ν_μ with E_ν ¡ 1 GeV, CCQE scattering is the dominant type of interaction.

CCQE interactions are particularly valuable in neutrino physics. The two-body nature of the interaction allows the kinematics to be completely reconstructed, specifically the initial neutrino energy, assuming the initial nucleon is at rest. Accurately reconstructing the neutrino energy is vital for neutrino oscillation measurements. In fact, for a nucleon at rest (or if E_ν is large enough such that this is a valid approximation) E_ν can be calculated with just the momentum, p_l, and angle relative to the incoming neutrino, θ_l, of the charged lepton:

$$E_\nu = \frac{m_p^2 + 2m_n \mathrm{E}_l - m_n^2 - m_l^2}{2\left(m_n - \mathrm{E}_l + p_l \cos\theta_l\right)} \tag{2.7}$$

where $m_{p,n,l}$ are the masses of the proton, neutron and lepton and E_l is the energy of the lepton. This can be derived directly from mass and energy conservation. For nucleons bound within a nucleus, Eq 2.7 must be amended slightly by introducing the binding energy.

Resonance Production

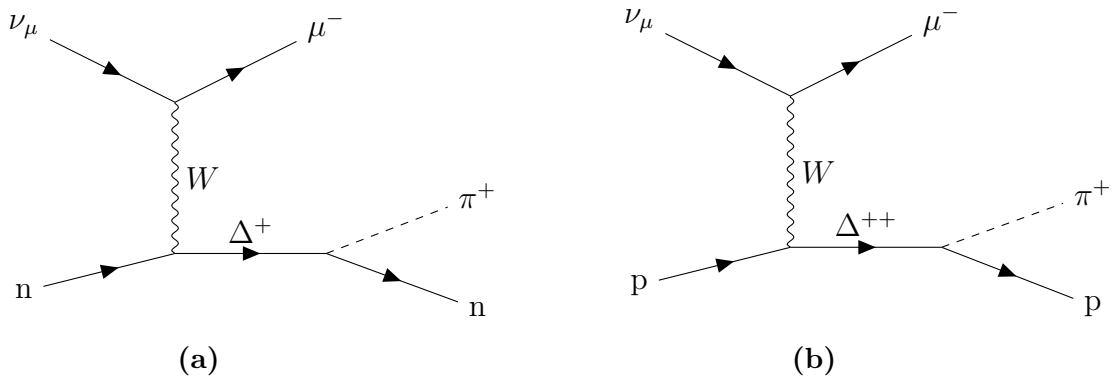

Figure 2.6: Example Feynman diagram showing ν_μ CC resonant scattering off of a (a) neutron and (b) proton producing a single pion (CC1π).

For higher neutrino energies, where larger q^2 is possible, neutrinos can undergo inelastic interactions. The incoming neutrino "knocks" the target nucleon into a baryonic resonance, e.g. Δ. The resonant state then decays back to a nucleon while also emitting a meson or photon. Most often just a single pion is emitted, which is shown in Figure 2.6. The cross-section for RES interactions peaks in the energy region between QE and DIS i.e. 0.5 ¡ E_ν ¡ 10 GeV.

Deep Inelastic Scattering

As E_ν increases further, q^2 becomes sufficiently large that the nucleon structure can be resolved by the incoming neutrino. The neutrino can scatter directly off of the quarks within the nucleon, including the "sea" of quarks and antiquarks that are constantly being created and annihilated. DIS interactions are the dominant process for E_ν ¿ 10 GeV. The DIS cross-section increases

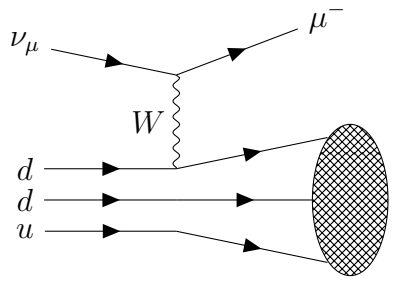

Figure 2.7: Example Feynman diagram showing ν_μ CC DIS interaction.

linearly until E_ν $\sim$M$_{W,Z}$. A Feynman diagram of a CC DIS interaction is shown in Figure 2.7.

DIS interactions result in the fragmentation of the struck nucleon. In a simplistic view, the struck quark recoils resulting in the creation of quark-antiquark pairs, a process known as hadronization. In an experiment, a jet of hadrons is observed.

Nuclear Effects and Final-State Interactions

Neutrino interactions are significantly more complex in real experiments since the nucleon is contained within a nucleus. The initial state of these nucleons, which are constantly moving and changing momentum within the nuclear potential, affects the kinematics and cross-section of neutrino interactions, especially at lower energies where Q^2 is smaller. Models like the relativistic Fermi-gas [27] have traditionally been used to describe these interactions, but newer experiments at lower neutrino energies are exploring alternative models, such as spectral functions [28]. Additionally, the nuclear potential limits the final-state kinematics of nucleon-producing interactions through an effect called "Pauli-blocking", which requires the final-state nucleon's momentum to exceed the Fermi momentum. This occurs because the final-state nucleon cannot occupy a quantum state of an existing nucleon due to the Pauli exclusion principle.

Furthermore, neutrino interactions are not confined to individual nucleons but can involve correlated nucleon pairs, alpha particles, or other quasi-bound states. After the interaction, final-state particles must propagate out of the nucleus, where they can undergo strong interactions with other nucleons, known as final-state interactions (FSI). These FSIs can significantly alter the momentum, direction, type, and number of particles, with pions and nucleons poten-

tially being absorbed or generating additional particles. Consequently, the particles detected outside the nucleus can differ greatly from those initially created at the interaction vertex.

2.3 Neutrino Oscillations

This thesis focuses on neutrino oscillation, the phenomenon whereby a neutrino created with a given flavour (ν_e, ν_μ or ν_τ) can later be observed in a different flavour. This section will first describe the initial evidence for neutrino oscillation in Section 2.3.1. The theory of neutrino oscillation will be described in Section 2.3.2 and a review of our current knowledge of neutrino oscillation parameters in Section 2.3.3.

2.3.1 Evidence for Neutrino Oscillation

The Solar Neutrino Puzzle

In 1968, Davis, Harmer and Hoffman devised an experiment to measure neutrinos from the sun at the Homestake Mine, South Dakota [29]. Specifically, the neutrinos produced by the decay of ^{8}B: $^8\mathrm{B} \rightarrow {}^8\mathrm{Be} + e^+ + \nu_e$. They planned to use the same method of detection as Davis' previous experiment roughly 10 years prior; use Cl to observe the interaction in Eq 2.3. However, since this time they aimed to measure neutrinos and not antineutrinos, it would actually work! The results of the experiment showed a significant deficit, roughly 7 times smaller, in the predicted flux of neutrinos. This result became known as "the solar neutrino puzzle". The consensus at the time was that this was an issue with the predicted flux of neutrinos which came from the Standard Solar Model.

This result was validated by the Kamiokande Collaboration in 1989 [30]. The experiment was initially intended to measure proton decay using a large water Chereknov detector. However, they produced a measurement of the flux of solar ν_e by observing ν_e - e elastic scattering events ($\nu_e + e^- \rightarrow \nu_e + e^-$). They similarly observed a deficit, of roughly $\frac{1}{2}$, compared with the Standard Solar Model prediction, as shown in Figure 2.8. Given that this result was still

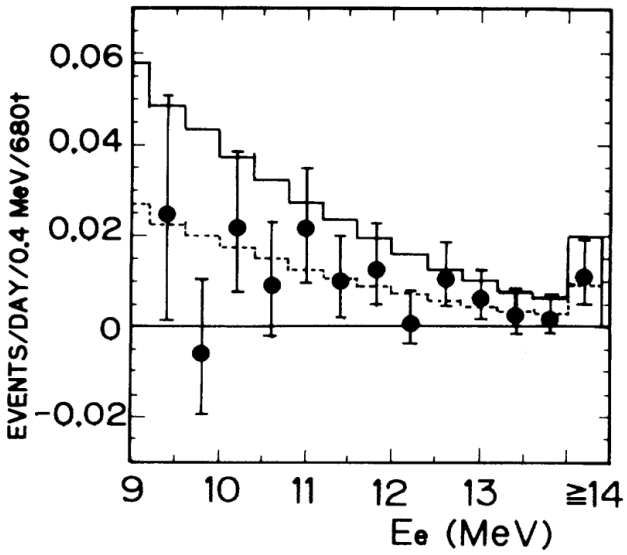

Figure 2.8: Energy distribution of solar neutrino events seen by the Kamiokande detector. The histogram shows the prediction by the Standard Solar Model. The dotted line shows the best fit to the data [30].

dependent on the same solar model, this result still was not clear evidence that something was missing from our picture of the neutrino.

The solar neutrino puzzle continued to be investigated into the early 1990s. The GALLEX [31] and SAGE [32] experiments both made measurements of the solar ν_e flux using a similar radiochemical method to Davis, except using Gallium rather than Chlorine: $\nu_e + {}^{71}\mathrm{Ga} \rightarrow e^- + {}^{71}\mathrm{Ge}$. The threshold energy for this interaction is much lower than that for Cl, allowing lower-energy neutrinos to be observed. This allowed neutrinos produced from the proton-proton chain, one of the nuclear fusion chain reactions that occur in the sun, to be detected, which at the time was a more understood process. Nevertheless, both experiments saw a significant deficit once again. Although still reliant on solar models, these results provided stronger evidence for the disappearance of ν_e.

The confirmation of the disappearance of solar neutrinos came from the SNO collaboration in 2002 [33]. Using a spherical water Cherenkov detector filled with heavy water, D_2O, SNO measured the flux of solar neutrinos, once again predominantly from the decay of ${}^8\mathrm{B}$. However, unlike previous experiments, SNO was able to measure both the ν_e flux and the total ν_x flux.

The ν_e flux was measured via the CC interaction off of deuterium

$$\nu_e + d \rightarrow p + p + e^- \tag{2.8}$$

which is not available to ν_μ or ν_τ. The total flux of neutrinos can be measured via NC interactions

$$\nu_x + d \rightarrow \nu_x + p + n \tag{2.9}$$

$$\nu_x + e^- \rightarrow \nu_x + e^- \tag{2.10}$$

which is available to all three neutrino flavours. They found that the total neutrino flux was consistent with Standard Solar Model predictions, however, they found the flux of ν_e to be $\sim \frac{1}{3}$ of the total flux. Only ν_e are produced by the sun, and so this gave fairly conclusive evidence that ν_e were changing flavour to ν_μ and ν_τ.

Atmospheric Neutrinos

In 1986, both the Irvine-Michigan-Brookhaven (IMB) [34] and Kamiokande [35] experiments were searching for proton decay, of which atmospheric neutrinos were the largest background. Both experiments made measurements of the rate of atmospheric neutrinos. Atmospheric neutrinos are primarily produced by the decay of pions and muons in the atmosphere, e.g.

$$\begin{aligned} \pi^\pm \rightarrow \mu^\pm + \nu_\mu/\bar{\nu}_\mu \\ \hookrightarrow e^\pm + \nu_e/\bar{\nu}_e + \bar{\nu}_\mu/\nu_\mu \end{aligned} \tag{2.11}$$

Given that Eq 2.11 is the sufficiently dominant process, this would result in an expected ratio of $(\nu_\mu/\bar{\nu}_\mu) : (\nu_e/\bar{\nu}_e) = 2 : 1$. However, both experiments measured a lower rate of ν_μ than

expected.

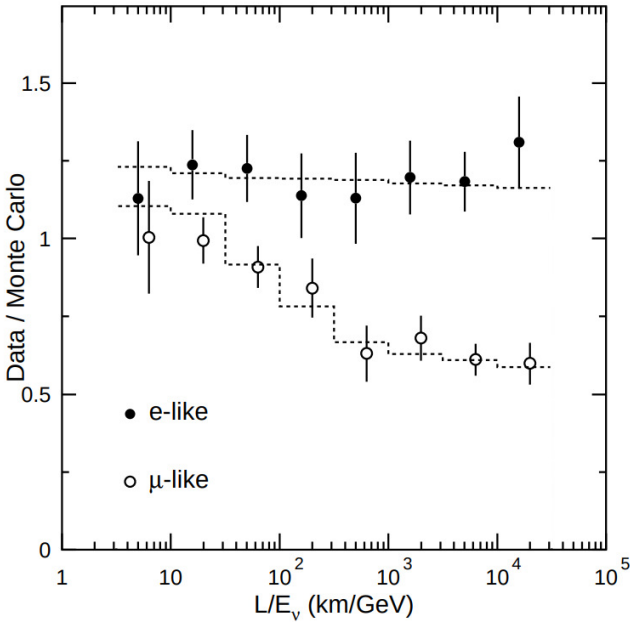

Figure 2.9: The ratio of the number of observed events to Monte Carlo prediction in the absence of neutrino oscillation against the reconstructed L/E_ν in Super-Kamiokande. The dashed lines show the expected shape of the distribution when accounting for $\nu_\mu \leftrightarrow \nu_\tau$ in a two-flavour oscillation model at $\Delta m^2 = 2.2x10^{-3}eV^2$ and $\sin^2 2\theta = 1$ [35].

In 1998, the successor to Kamiokande, Super-Kamiokande, made a far more definitive measurement of the disappearance of atmospheric ν_μ [36]. The Super-Kamiokande detector is a 50 kt water Cherenkov detector with good reconstruction on the angle of the incoming neutrino. By measuring the incoming angle, the ν_μ events could be separated into two samples: up-going and down-going. The down-going neutrinos will have been produced in the atmosphere directly above the detector and will have travelled $\sim 15km. The up-going neutrinos will have travelled the whole way t$ $13000km. Under the hypothesis that there is no oscillation, and provided the cosmic rays which produce the neut$ $going and down-going neutrino events. On the contrary, Super-Kamiokande observed a deficit in up-$ $going neutrinos compared with down-going neutrinos, which proved that the observed atmospheric \nu_\mu$ flux was dependent on the distance the neutrino propagated. The oscillation of atmospheric ν_μ became evident when plotting the ratio of observed events to prediction against distance, L, over E_ν, as shown in Figure 2.9. For muon-like events, there is a clear deficit as L/E_ν increases, but not for the electron-like events. This was clear evidence of the oscillation of atmospheric ν_μ.

Reactor Neutrinos

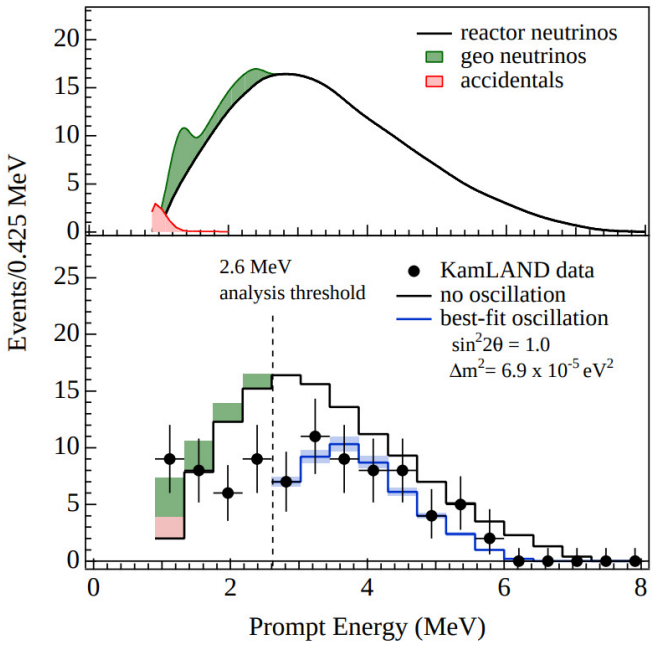

Figure 2.10: Upper Panel: Predicted reactor $\bar{\nu}_e$ flux as a function of energy, showing background contributions e.g. geo-neutrinos (from decays in the Earth). Lower panel: Observed $\bar{\nu}_e$ flux as a function of energy. The black histogram shows the predicted flux given no oscillation. The blue histogram shows the best fit using two-flavour oscillation [37]

As well as solar and atmospheric neutrinos, strong evidence was found for the oscillation of $\bar{\nu}_e$ from nuclear reactors. The KamLAND experiment measured $\bar{\nu}_e$ from 53 nuclear reactors in Japan, with the vast majority of the flux coming from 26 reactors with a distance range of 138-214 km [37]. The relatively narrow band of distances for the majority of the flux allows KamLAND to observe the effect of oscillation in the flux as a function of energy. The detector consists of a balloon filled with liquid scintillator, in which $\bar{\nu}_e$ can undergo inverse beta decay, as in Eq 2.2. The emitted positron produces scintillation light as it passes through the liquid before it annihilates. The emitted neutron is then captured by hydrogen, $\sim$ 200 μs later, the combination of which produces a unique signature. A significant deficit was observed, as shown in Figure 2.10, consistent with the findings of the solar and atmospheric neutrino experiments.

Both the Super-Kamiokande and SNO experiments were accredited with the discovery of neutrino oscillation and, as a result, shared the 2015 Nobel Prize for Physics [38]. Since then, several constraints have been made on the oscillation parameters in the model (discussed

in Section 2.3.2) from a wide range of modern neutrino experiments, including long-baseline experiments. Their results and the current state of knowledge of neutrino oscillation parameters will be discussed in Section 2.3.3.

2.3.2 Theory of Neutrino Oscillation

Neutrinos interact via the weak force, which couples to the flavour of the neutrino. In charged-current weak interactions, mediated by a W boson, a neutrino of lepton flavour ν_l will be coupled to the corresponding lepton l. Hence, the flavour eigenstates undergo interactions which are detectable experimentally.

The confirmation of neutrino oscillation, outlined in Section 2.3.1, proves that neutrinos are not massless particles. This means that as well as flavour eigenstates, mass eigenstates also exist, each having a well-defined mass. The sets of eigenstates also cannot be equivalent, another consequence of the observation of oscillation. We can therefore describe the flavour eigenstates, $|v_\alpha\rangle$ for $\alpha = e, \mu, \tau$, as a superposition of the mass eigenstates, $|v_i\rangle$ for $i = 1, 2, 3$,

$$|v_\alpha\rangle = \sum_i U^*_{\alpha i} |v_i\rangle \tag{2.12}$$

where U is the unitary matrix that describes the transformations between the flavour and mass eigenstates, known as the Pontecorvo-Maki-Nakagawa-Sakata (PMNS) matrix [20]. Given three flavour eigenstates and three mass eigenstates, the PMNS matrix can be written as

$$\mathrm{U} = \begin{pmatrix} \mathrm{U}_{e1} & \mathrm{U}_{e2} & \mathrm{U}_{e3} \\ \mathrm{U}_{\mu 1} & \mathrm{U}_{\mu 2} & \mathrm{U}_{\mu 3} \\ \mathrm{U}_{\tau 1} & \mathrm{U}_{\tau 2} & \mathrm{U}_{\tau 3} \end{pmatrix} \tag{2.13}$$

The PMNS matrix can be decomposed into three rotation matrices and a diagonal matrix,

and parameterised in terms of three rotation angles and three complex phases [39]

$$U = \underbrace{\begin{pmatrix} 1 & 0 & 0 \\ 0 & c_{23} & s_{23} \\ 0 & -s_{23} & c_{23} \end{pmatrix}}_{\text{Atmospheric}} \underbrace{\begin{pmatrix} c_{13} & 0 & s_{13}e^{-i\delta_{CP}} \\ 0 & 1 & 0 \\ -s_{13}e^{i\delta_{CP}} & 0 & c_{13} \end{pmatrix}}_{\text{Cross-mixing}} \underbrace{\begin{pmatrix} c_{12} & s_{12} & 0 \\ -s_{12} & c_{12} & 0 \\ 0 & 0 & 1 \end{pmatrix}}_{\text{Solar}} \underbrace{\begin{pmatrix} e^{i\alpha_1/2} & 0 & 0 \\ 0 & e^{i\alpha_2/2} & 0 \\ 0 & 0 & 1 \end{pmatrix}}_{\text{Majorana}} \tag{2.14}$$

where $c_{lm} = \cos\theta_{lm}, s_{lm} = \sin\theta_{lm}$, θ_{lm} is the rotation angle for the relevant rotation and δ_{CP} is the charge-parity violation phase. α_1 and α_2 are complex phases which are necessary only if neutrinos are Majorana particles. This decomposition separates the mixing angles appropriately into the experiments which are most sensitive to them.

The first matrix contains only θ_{23} and the third matrix contains only θ_{12}. They are often referred to as the 'atmospheric' and 'solar' parameters respectively since they have major contributions to the oscillation of those neutrinos. For this reason, both the solar and atmospheric neutrino oscillations can be approximated to a two-flavour model with the respective flavours involved, reducing the three-flavour model to only the relevant parameter.

The second matrix, known as the 'cross mixing' matrix contains the mixing angle θ_{13}, known as the 'reactor' parameter, and δ_{CP}. If δ_{CP} is non-zero, U becomes a complex matrix which, as shown below, will lead to probability differences between $\upsilon_\alpha \longrightarrow \upsilon_\beta$ and $\bar{\upsilon}_\alpha \longrightarrow \bar{\upsilon}_\beta$. This result is of high interest in the scientific community as it would suggest an asymmetry between matter and anti-matter. Accurate measurements of this parameter will determine how much charge-parity violation is generated by neutrino oscillations which could help explain the matter dominance in the universe.

The 'Majorana' matrix is only necessary if neutrinos have a Majorana mass, meaning that they are their own antiparticle. If true, two extra complex phases need to be introduced to redefine the phase of Majorana fields. Otherwise, the right-hand side matrix in Eq 2.14 can be

omitted. However, the mass nature of neutrinos cannot be concluded in oscillation experiments as the extra complex phases do not affect the oscillation probability. This requires the observation of other phenomena such as neutrino-less double beta decay, which is currently under study (MAJORANA [40], GERDA [41], SNO+ [42], KamLAND-Zen [43] and SuperNEMO [44], among others).

The unitarity of the PMNS matrix is a consequence of the Standard Model. In the Standard Model, due to lepton-number conservation, a neutrino of flavour ν_α can only interact to create a lepton of flavour l_α. This means that the flavour states $|v_\alpha\rangle$ must be orthogonal, resulting in the mixing matrix being unitary. If we also assume the mass eigenstates are orthogonal, we can write

$$\langle \nu_\alpha \mid \nu_\beta \rangle = \left(\sum U^*_{\alpha i} |\nu_i\rangle \right)^\dagger \left(\sum U^*_{\beta j} |\nu_j\rangle \right) \tag{2.15}$$

$$\delta_{\alpha\beta} = \sum_i \sum_j U^*_{\beta j} U_{\alpha i} \langle \nu_i \mid \nu_j \rangle \tag{2.16}$$

$$\delta_{\alpha\beta} = \sum_i U^*_{\beta i} U_{\alpha i} \tag{2.17}$$

where $\delta_{\alpha\beta}$ is the Kronecker delta ($\delta_{\alpha\beta} = 1$ if $\alpha = \beta$, $\delta_{\alpha\beta} = 0$ if $\alpha \neq \beta$). Since the PMNS matrix is unitary, we can invert Eq 2.12 to express the mass eigenstates as a superposition of the flavour eigenstates:

$$|v_i\rangle = \sum_\alpha U_{\alpha i} |v_\alpha\rangle \tag{2.18}$$

The fraction of flavour α in a mass eigenstate ν_i is equal to $|U_{\alpha i}|^2$. This is equivalent to the probability that a ν_i will interact to produce a lepton of flavour α.

We now aim to calculate the probability of a neutrino produced in flavour α being de-

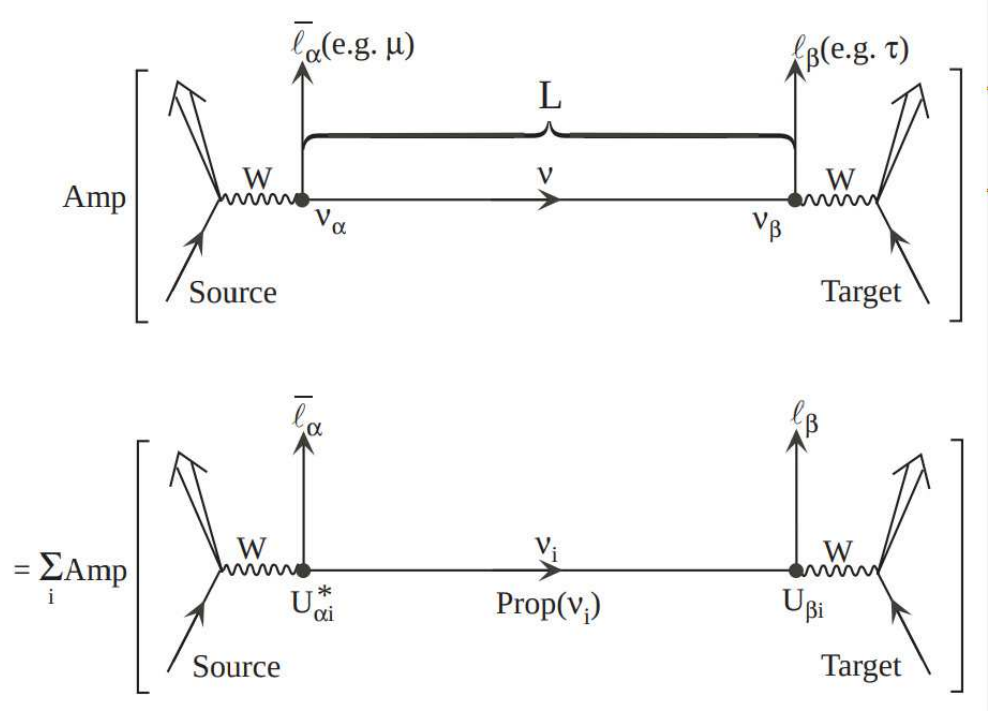

Figure 2.11: A simple diagram illustrating the oscillation of a neutrino flavour eigenstate (upper diagram) and the contributions of each mass eigenstate to the calculation (lower). "Amp" denotes the amplitude of the oscillation and "Prop" denotes the propagator. Figure taken from [39].

tected in flavour β after propagating some distance, L. A schematic of the problem is shown in Figure 2.11, illustrating that a neutrino has changed flavour if it was produced alongside a lepton of flavour α and later interacts to produce a lepton of flavour β. It also shows that the amplitude of the oscillation is a sum of the amplitudes of each mass eigenstate:

$$\text{Amp}\,(\nu_\alpha \to \nu_\beta) = \sum_i U^*_{\alpha i}\,\text{Prop}\,(\nu_i)\,U_{\beta i} \tag{2.19}$$

where $\text{Amp}\,(\nu_\alpha \to \nu_\beta)$ denotes the amplitude of the oscillation and $\text{Prop}\,(\nu_i)$ denotes the propagator. We now must calculate what $\text{Prop}\,(\nu_i)$ is. We start by considering the time evolution of a neutrino state, in its rest frame, given by the Schrödinger equation:

$$i\frac{d}{d\tau}\left|\nu_i(\tau)\right\rangle = \mathcal{H}\left|\nu_i\right\rangle \tag{2.20}$$

where $\mathcal{H}$ is the Hamiltonian operator and τ is the propagation time. Since the neutrino mass eigenstate is also an eigenstate of the Hamiltonian

$$\mathcal{H}\left|\nu_i\right\rangle = m_i\left|\nu_i\right\rangle \tag{2.21}$$

where m_i is the mass of the eigenstate i. By equating Eq 2.20 and Eq 2.21, we find that the solution to the Schrödinger equation is a plane wave

$$\left|\nu_i\left(\tau\right)\right\rangle = e^{-im_i\tau}\left|\nu_i\right\rangle \tag{2.22}$$

Now we must convert this into the laboratory frame. By Lorentz invariance,

$$m_i\tau = E_i t - p_i L \tag{2.23}$$

where, in the laboratory frame, E_i is the energy of the mass eigenstate, t is the time, p_i is the momentum of the mass eigenstate and L is the distance propagated. The momentum of the mass eigenstate is related to its energy and mass via

$$p_i = \sqrt{E_i^2 - m_i^2} \approxeq E_i - \frac{m_i^2}{2E_i} \tag{2.24}$$

where, since neutrinos are extremely light i.e. $m_i^2 \ll E_i^2$, we can approximate using a binomial expansion. Placing this back into Eq 2.23, we find

$$m_i\tau \cong E\,(t-L) + \frac{m_i^2}{2E}L \tag{2.25}$$

where we have assumed that the energy of each mass eigenstate, for a given flavour state, is equal (i.e. $E_i = E_j$). The term $E(t-L)$ is the same for all eigenstates and will become irrelevant when calculating the oscillation probability. Substituting Eq 2.25 into Eq 2.22, we find a final expression for our propagator:

$$\text{Prop}\,(\nu_i) = exp\left[-i\frac{m_i^2}{2E}L\right] \tag{2.26}$$

From Eq 2.19, we can write the amplitude for a neutrino to change flavour as

$$\text{Amp}\,(\nu_\alpha \to \nu_\beta) = \sum_i U_{\alpha i}^* U_{\beta i} e^{-im_i^2\frac{L}{2E}} \tag{2.27}$$

From the amplitude, we can calculate the probability of a $\nu_\alpha \to \nu_\beta$ oscillation:

$$P\,(\nu_\alpha \to \nu_\beta) = |\,\text{Amp}\,(\nu_\alpha \to \nu_\beta)\,|^2 = \sum_i \sum_j U_{\alpha i}^* U_{\beta i} U_{\alpha j} U_{\beta j}^* e^{-i\Delta m_{ij}^2\frac{L}{2E}} \tag{2.28}$$

where $\Delta m_{ij}^2 = m_i^2 - m_j^2$ and is referred to as the mass-squared difference. This result can be simplified by splitting Eq 2.28 into its real and imaginary components and utilising the unitarity of U, leaving the following more convenient expression:

$$P\left(\nu_\alpha \to \nu_\beta\right) = \delta_{\alpha\beta} - 4\Re\sum_{i>j} U^*_{\alpha i}U_{\beta i}U_{\alpha j}U^*_{\beta j}\sin^2\left(\frac{\Delta m^2_{ij}}{4E}L\right) + 2\Im\sum_{i>j} U^*_{\alpha i}U_{\beta i}U_{\alpha j}U^*_{\beta j}\sin\left(\frac{\Delta m^2_{ij}}{2E}L\right) \tag{2.29}$$

The first term in Eq 2.29 corresponds to no oscillations. The following term corresponds to oscillations, with the real components of the PMNS matrix controlling the amplitude of oscillation and the terms inside the sine function controlling the phase. The final term in the expression accounts for the impact of CP violation within the PMNS matrix.

Accelerator-based neutrino experiments, such as DUNE, generally study neutrino oscillation by using a ν_μ dominated neutrino beam and measuring the ν_μ disappearance, $P\left(\nu_\mu \to \nu_\mu\right)$, and the ν_e appearance, $P\left(\nu_\mu \to \nu_e\right)$. These probabilities can also be approximated to

$$P\left(v_\mu \to v_\mu\right) \approx 1 - \cos^4\theta_{13}\sin^2 2\theta_{23}\sin^2\left(1.27\frac{\Delta m^2_{32}L}{E}\right) \tag{2.30}$$

$$P\left(v_\mu \to v_e\right) \approx \sin^2 2\theta_{13}\sin^2\theta_{23}\sin^2\left(1.27\frac{\Delta m^2_{31}L}{E}\right) \tag{2.31}$$

where the factor 1.27 comes when the units of L and E and the values and units of $\hbar$, c are included. From Eq 2.30 and Eq 2.31, we can see that L and E are the controllable experimental variables that must be optimised to maximise the oscillation probability. If we consider the ν_μ disappearance, the important mass-squared splitting for the oscillation phase is Δm^2_{32}, which is $\approx 0.0025\,\mathrm{eV}^2$. Accelerator experiments typically produce neutrinos with energies of $\mathcal{O}(1\,\mathrm{GeV})$. Since we know that

$$1.27\frac{\Delta m^2_{32}L}{E} \approx \frac{\pi}{2} \tag{2.32}$$

maximises the ν_μ disappearance, we can estimate the rough baseline needed to achieve this:

$$L \;=\; \frac{\pi\, E}{2.54\, \Delta m^2_{32}} \;\approx\; 400\,\text{km} \tag{2.33}$$

Accelerator experiments with higher energies will require longer baselines (i.e. DUNE).

If we refer back to the atmospheric oscillation discovery, discussed in Section 2.3.1, we can use Eq 2.30 to understand the observations. The downgoing ν_μ neutrinos, which only travelled $\sim$ 15 km, would have a disappearance probability of ~ 0.001. Hence, the baseline is much too short for those oscillations to have started. The up-going neutrinos, which travelled a much longer distance, $\sim$ 13000 km, would have a disappearance probability of ~ 0.5. This explains the observed data in Figure 2.9.

Another interesting characteristic can be found in Eq 2.30. The θ_{23} mixing angle appears in the disappearance probability, to leading order, in the term $\sin^2 2\theta_{23}$. Hence, maximal mixing is defined as $\sin^2 2\theta_{23} = 1$ and hence $\theta_{23} = \frac{\pi}{4}$. However, if θ_{23} is not maximal, of which the current data suggests it's close but not quite (see Section 2.3.3), there is a symmetry in the values of $\sin^2 2\theta_{23}$. Values of $\theta_{23} < \frac{\pi}{4}$ and $\theta_{23} > \frac{\pi}{4}$ (or $\sin^2\theta_{23} < 0.5$ and $\sin^2\theta_{23} > 0.5$), referred to as the lower octant and upper octant respectively, return the same disappearance probability. This degeneracy can be broken by appearance channels, i.e. Eq 2.31, or by the effect of higher order terms of disappearance probability.

Anti-neutrino Oscillation

So far we've derived neutrino oscillation probabilities, but we haven't explicitly checked what happens to anti-neutrinos. To do this, we first adjust Eq 2.18

$$|\bar{\nu}_i\rangle = \sum_\alpha U^*_{\alpha i}\, |\bar{\nu}_\alpha\rangle \tag{2.34}$$

where the only change is the conjugate on the PMNS element. The same derivation can be applied, which results in the following expression of the oscillation probability

$$P\left(\bar{\nu}_\alpha \to \bar{\nu}_\beta\right) = \delta_{\alpha\beta} - 4\Re \sum_{i>j} U^*_{\alpha i} U_{\beta i} U_{\alpha j} U^*_{\beta j} \sin^2\left(\frac{\Delta m^2_{ij}}{4E} L\right) - 2\Im \sum_{i>j} U^*_{\alpha i} U_{\beta i} U_{\alpha j} U^*_{\beta j} \sin\left(\frac{\Delta m^2_{ij}}{2E} L\right) \tag{2.35}$$

where the only major change is the sign in front of the imaginary term. This difference can give rise to CP violation.

Charge-Parity Violation

In Section 2.2.1, parity violation was discussed and shown to occur in weak interactions. The initial theory was that, although parity is violated in weak interactions, the combination of both charge and parity operators would leave the interaction properties the same. In the neutrino sector, a CP operation on a left-handed neutrino would return a right-handed antineutrino. Hence, CP-violation in the neutrino sector would manifest as a difference between the physics of neutrinos and antineutrinos. In the PMNS matrix, the CP violating phase factor, δ_{CP}, causes a difference between the oscillations of neutrinos and antineutrinos (if $\delta_{CP} \neq 0$ or $\pm\pi$).

The existence of CP violation in the neutrino sector is an exciting prospect as it may help to explain the matter-antimatter asymmetry that is observed in the universe. The universal matter-antimatter asymmetry observation requires there to have been some responsible process during the evolution of the universe since equal parts of matter and antimatter were created in the Big Bang. Sakharov proposed a set of three necessary conditions that must be satisfied to produce a matter-antimatter asymmetry [45]:

1. Baryon number violation

2. C and CP violation

3. Out of equilibrium interactions

CP violation has already been observed in the quark sector [1], however, the magnitude is not enough to explain the observed asymmetry.

The Jarlskog invariant, J_{CP}, is a quantity that indicates the magnitude of CP violation [46]. It is invariant under changes to the parameterisation of the mixing matrix that governs the process, i.e. the PMNS matrix. When applied to the PMNS matrix, J_{CP} takes the form of

$$J_{CP} = \frac{1}{8} \cos\theta_{13} \sin(2\theta_{13}) \sin(2\theta_{12}) \sin(2\theta_{23}) \sin\delta_{CP} \tag{2.36}$$

Mass Ordering

In the Standard Model, neutrinos are assumed to have no mass since, without the existence of a right-handed chiral state, there is no way to construct the Dirac mass term. However, as is quite evident from Eq 2.29, for neutrino oscillation to occur the mass-splittings must be non-zero.

Given three neutrino eigenstates, there exists two independent mass splitting: Δm_{32} (referred to as the atmospheric mass splitting) and Δm_{21} (referred to as the solar mass splitting). The third splitting is defined by the relationship: $\Delta m_{32} + \Delta m_{21} + \Delta m_{31} = 0$. Generally, neutrino oscillation only has direct sensitivity to the squared mass difference, meaning there is ambiguity in the ordering of the mass states. Solar neutrino experiments were able to deduce that the $m_2^2 > m_1^2$ through the effect that matter in the sun has on the rate of oscillations (this effect is discussed further in the following section). That leaves two possible hierarchies, normal ordering (NO): $m_3^2 > m_2^2 > m_1^2$ and inverted ordering (IO): $m_2^2 > m_1^2 > m_3^2$, as shown in Figure 2.12.

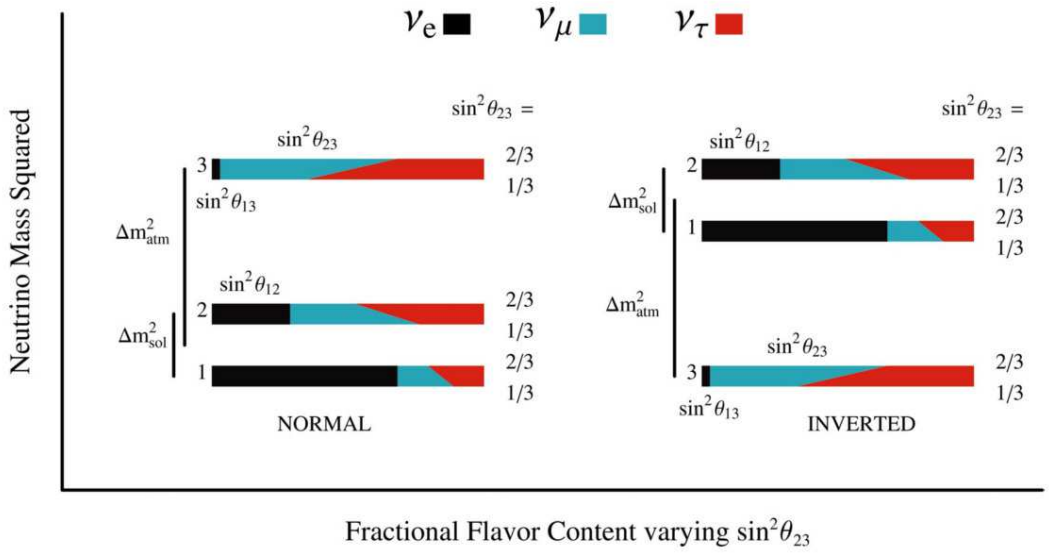

Figure 2.12: Illustration of the two possible orderings of the mass eigenstates. The flavour compositions of each mass state, given by the current best knowledge of the mixing angles (under both $\sin^2\theta_{23} = 2/3$ and $\sin^2\theta_{23} = 1/3$ hypotheses), are denoted by the colours. Image taken from [47].

Matter Effects

Thus far, we've only considered neutrino oscillations in a vacuum. In reality, the neutrinos that are measured in experiments travel through matter. As neutrinos propagate through matter, an additional potential is introduced from the particles within that medium causing the mass eigenstates to change. This is mainly a result of coherent scattering (scattering which leaves the neutrino state unchanged), which can either be CC or NC, as shown in Figure 2.13.

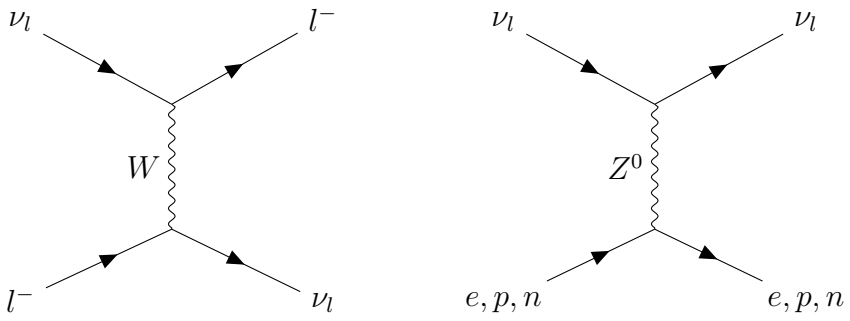

(a) CC coherent scattering **(b)** NC coherent scattering

Figure 2.13: Example Feynman diagrams showing (a) CC and (b) NC coherent scattering in matter.

Since NC coherent scattering is independent of the incoming neutrino flavour, it leaves the

oscillation probability unchanged. On the contrary, the CC scattering only affects ν_e since electrons are the only charged lepton present in the Earth's matter. The contributing term to the Hamiltonian and how this affects the neutrino oscillation is not given here but can be found in several neutrino textbooks such as [48]. In the sun, the rapidly changing density of matter gives rise to resonances, a phenomenon known as the Mikheyev-Smirnov-Wolfenstein (MSW) effect [49]. In the Earth's crust, through which neutrinos in long-baseline experiments propagate, the density of matter can be approximated as constant.

Furthermore, the sign of the extra potential introduced by the matter effect is reversed for antineutrinos. This is inherently CP-asymmetric and has the effect of mimicking the CP-violation which is caused by the PMNS matrix. It is therefore vital that these effects are taken into account when inferring the values of oscillation parameters from data.

The matter effects also allow for differentiation between the two mass hierarchies since the effect is also dependent on the sign of Δm^2_{ij}. Due to the long baseline and intense neutrino beam at DUNE, the mass hierarchy will be completely resolved within the first few years of running [2]. This is essential for the true effects of CP violation in the PMNS matrix to be measured. In this analysis, matter effects are fully taken into account while assuming a constant matter density of 2.848 g/cm^3 [50].

2.3.3 Current Knowledge of Oscillation Parameters

Neutrino oscillation parameters have been measured by a wide range of experiments using accelerator (both long-baseline and short-baseline), reactor, atmospheric and solar neutrinos. The current best estimates of the neutrino oscillation parameters are presented here, as published in the 2023 update of the 2022 Review of Particle Physics by the Particle Data Group [51]. These results were performed by NuFit which provides a global neutrino oscillation measurement [52]. Specifically, these are the NuFit 5.2 results (without Super-K Atmospheric constraints).

- **θ_{23}:**

  The current best-estimate of $\sin^2\theta_{23}$:

$$\begin{aligned} \sin^2\theta_{23} &= 0.572^{+0.018}_{-0.023} \quad \text{(NO)} \\ &= 0.578^{+0.016}_{-0.021} \quad \text{(IO)} \end{aligned}$$

where, since there is dependence on the mass ordering, values are presented for both NO and IO. The global sensitivity of θ_{23} comes from the disappearance channel (Eq 2.30) of long-baseline experiments, of which the major contributors are T2K [53], NOvA [54] and MINOS [55]. There is no significant preference for the octant of θ_{23}. Often results are published as a 2D contour of $\Delta m^2_{32} - \sin^2\theta_{23}$ since both parameters significantly contribute to the disappearance probability. Figure 2.14 shows the 90% contours for T2K, NOvA, Super-K, MINOS and IceCube experiments superposed on each other.

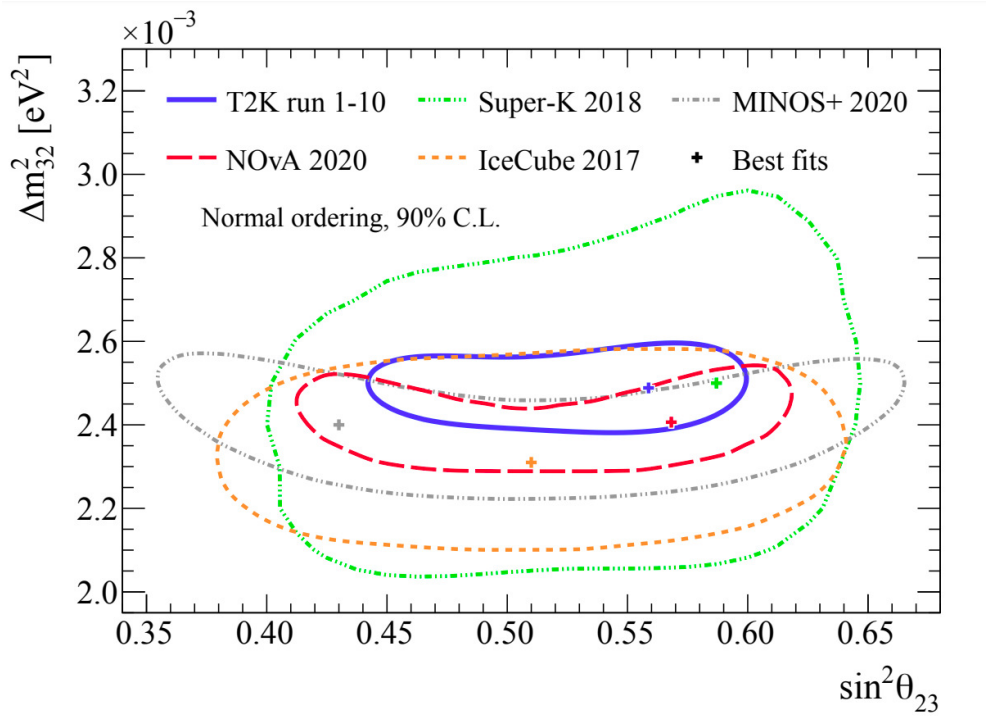

Figure 2.14: The 90% confidence level region in $\Delta m^2_{32} - \sin^2\theta_{23}$ and the best-fit points from T2K, NOvA, Super-K, MINOS and IceCube experiments. Figure from [56]

- **Δm^2_{32}:**

  The current best-estimate of $|\Delta m^2_{32}|$:

$$|\Delta m^2_{32}| = 2.437^{+0.028}_{-0.027} \times 10^{-3}\,\text{eV}^2 \quad \text{(NO)}$$
$$= 2.498^{+0.032}_{-0.025} \times 10^{-3}\,\text{eV}^2 \quad \text{(IO)}$$

where, since there is dependence on the mass ordering, values are presented for both NO and IO. The modulus is present since the sign of Δm^2_{32} is still unknown which is equivalent referring to the mass ordering. This global fit analysis has a slight preference for NO.

- **θ_{12}:**

  The current best-estimate of $\sin^2\theta_{12}$:

$$\sin^2\theta_{12} = 0.303^{+0.012}_{-0.011}$$

  where a constraint on the value of θ_{13} has been used. The current global sensitivity is dominated by solar neutrino experiments (e.g. SNO [33], BOREXINO [57], Super-K [58]).

- **Δm^2_{21}:**

  The current best-estimate of $|\Delta m^2_{21}|$:

$$|\Delta m^2_{21}| = 7.41^{+0.21}_{-0.20} \times 10^{-5}\,\text{eV}^2$$

  where a constraint on the value of θ_{13} has been used. The current global sensitivity is dominated by long-baseline reactor experiments, i.e. KamLAND [37, 59]. Results of $\tan^2\theta_{12}$ and Δm^2_{21} often presented together as a 2D contour. Figure 2.15 shows the 2D contours from KamLAND and solar neutrino experiments. It can be seen that the solar neutrino experiments have much better precision on the θ_{12}, whereas KamLAND has

much better precision on Δm^2_{21}.

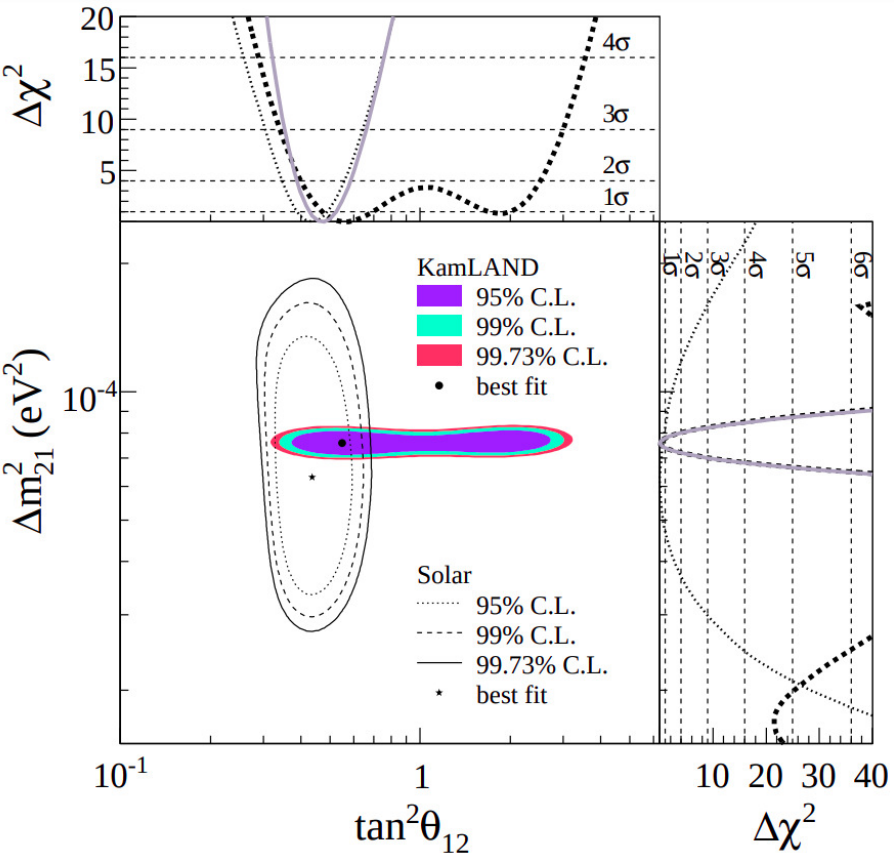

Figure 2.15: The 95%, 99%, 99.73% confidence level regions for $\Delta m^2_{21} - \tan^2\theta_{12}$ from KamLAND (colours) and solar experiments (lines). The side panels show the 1D $\Delta\chi^2$ profiles for KamLAND (dashed), solar experiments (dotted) and the combination of both (solid). Figure from [59]

- ***θ_{13}*:**

  The current best-estimate of $\sin^2\theta_{13}$:

$$\begin{aligned}\sin^2\theta_{13} &= 2.2203^{+0.056}_{-0.059} \times 10^{-2} \quad \text{(NO)} \\ &= 2.2219^{+0.060}_{-0.057} \times 10^{-2} \quad \text{(IO)}\end{aligned}$$

  where the choice of mass ordering has a slight effect on the best estimate. The current global sensitivity is dominated by the Daya Bay short-baseline reactor experiment[60], with contributions from similar experiments (e.g. Double Chooz [61] and RENO [62]).

θ_{13} is considerably smaller than the other mixing angles and, before the results of precision reactor experiments, was thought to be zero for some time.

- ***δ_{CP}*:**

  The current best-estimate of δ_{CP}:

$$\begin{aligned} \delta_{CP} &= \left(197^{+42}_{-25}\right)^\circ \text{ or } -2.84^{+0.73}_{-0.44} \text{ rad (NO)} \\ &= \left(286^{+27}_{-32}\right)^\circ \text{ or } -1.29^{+0.47}_{-0.56} \text{ rad (IO)} \end{aligned}$$

  where values are given in both degrees and radians and shown for NO and IO. The choice of NO and IO has a significant impact on the best estimate. δ_{CP} is the worst constrained parameter in the PMNS matrix, with the dominant constraining power coming from the appearance channels of current long-baseline accelerator experiments, specifically T2K (although doesn't appear in the leading term shown in Eq 2.31). The 90% confidence level regions in $\delta_{CP} - \sin^2\theta_{23}$ from T2K, NOvA and Super-K are shown in Figure 2.16.

  Since the appearance probability is significantly affected by θ_{13}, a much stronger constraint on δ_{CP} is obtained when an external θ_{13} constraint is applied. Sensitivity to δ_{CP} is also achieved by collecting both ν and $\bar{\nu}$ data. However, due to the low number of events collected for $\bar{\nu}$, the appearance channel for ν still contributes more to the sensitivity to δ_{CP}.

Next-generation Neutrino Oscillation Experiments

A new era of neutrino oscillation experiments is being built to unambiguously measure neutrino oscillation parameters. More specifically, these experiments aim to extend the measurements of θ_{13}, measure the value of δ_{CP} to determine whether CP is violated in the neutrino sector and determine the mass hierarchy.

Two future long-baseline accelerator experiments, Hyper-Kamiokande (Hyper-K) [63] and

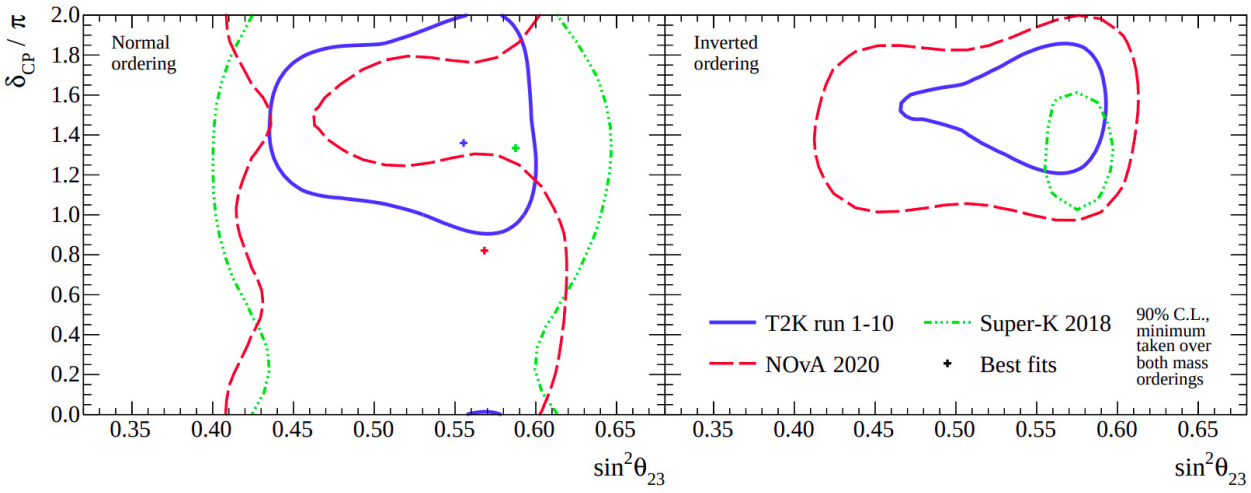

Figure 2.16: The 90% confidence level region in $\delta_{CP} - \sin^2\theta_{23}$ and best-fit points from T2K, NOvA and Super-K for both NO and IO. Figure from [56]

DUNE [2, 64–66], are currently in their construction phase. The Hyper-K detector is the successor to Super-K, using the same water Cherenkov technology but with 5 times the volume of water. DUNE will use liquid-argon technology and a neutrino beam with a wide energy range (this experiment is the topic of this thesis and is described in detail in Chapter 3). Both experiments will collect unprecedented numbers of neutrino events, with the aim of making 5σ statements on neutrino oscillation parameters, importantly δ_{CP}.

Jiangmen Underground Neutrino Observatory (JUNO) [67] is a medium baseline reactor neutrino experiment currently under construction in Southern China. Compared with Daya Bay, JUNO is much further away from the reactors and is therefore sensitive to a much larger peak of neutrino oscillation. This large distance also requires a much larger detector and better background rejection. By measuring the flux of $\bar{\nu}_e$, JUNO aims to make a high precision measurement of θ_{13} and determine the mass ordering.

This thesis focuses on DUNE, and going forward the potential of Hyper-K and JUNO to make precise oscillation measurements will not be discussed. However, given the differences between these experiments, significant opportunities exist in the future for joint analyses.

Chapter 3

The Deep Underground Neutrino Experiment

The Deep Underground Neutrino Experiment (DUNE) is a next-generation long-baseline neutrino experiment being constructed in the US [2, 64–66, 68, 69]. DUNE will be installed within the Long-Baseline Neutrino Facility (LBNF). Fermilab, Illinois, hosts the neutrino beam production facilities and the near detector complex. The far detector suite will be at the Sandford Underground Research Facility (SURF), South Dakota. The two sites provide a baseline of approximately 1300 km, as shown by the cartoon in Figure 3.1.

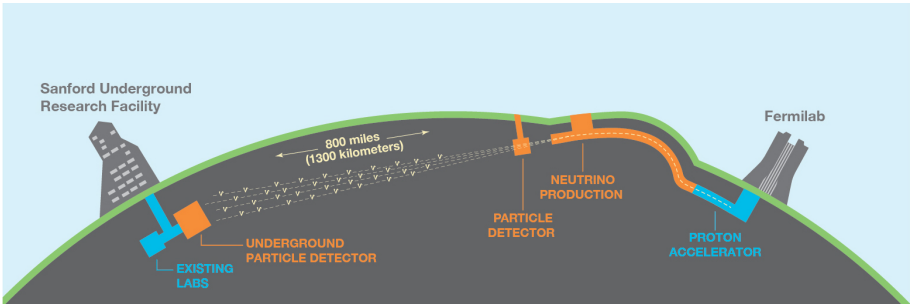

Figure 3.1: A cartoon representation of the DUNE experimental setup. The orange components need to be built, while the components in blue already exist. The figure is not to scale and is taken from [64].

This chapter will outline the components of DUNE. Section 3.1 will describe the design of the neutrino beam, Section 3.2 will describe liquid argon technology, Section 3.3 will describe

the design of the near detector, Section 3.4 will describe the design of the far detector and Section 3.5 will discuss the physics potential of the experiment.

3.1 Beam

DUNE will use the LBNF beam, taking advantage of the Proton Improvement Plan II (PIP-II) upgrade of the Fermilab accelerator complex [64]. The accelerator complex has several stages [70], starting with the ionisation of hydrogen atoms using a magnetron. The ions are incrementally accelerated to 750 keV where they enter the linear accelerator (Linac). The Linac, composed of roughly 150 m of drift tube cavities and side-coupled cavities, accelerates the H^- ions to 400 MeV. As they leave the Linac, they pass through a carbon foil which strips two electrons from each ion, producing H^+ ions (protons). The protons then enter the Booster, a circular accelerator measuring 468 m in circumference, and exit with 8 GeV. Magnets force the protons into a circular trajectory, where they gain energy after each revolution. Finally, the protons enter the 3.3 km circumference Main Injector, where they are extracted at up to 120 GeV.

By replacing the current linear accelerator, PIP-II aims to surpass the intensity of Fermilab's current high-energy neutrino beam, NuMI, producing the world's most intense neutrino beam. PIP-II plans to deliver 1.2 MW of proton beam power to the DUNE target at the start of operations and provide a platform for extending beam power to 2 MW [64].

Once the beam's ultimate trajectory is established, the protons strike a target, producing secondary mesons. These mesons are focused using magnetic horns before they can decay into muons and neutrinos. Depending on the choice of the polarity of the magnetic horns, the beam is either dominated by neutrinos, forward horn current (FHC), or antineutrinos, reverse horn current (RHC). After the decay pipe, muons and other decay products are stopped leaving a beam of neutrinos or antineutrinos. The neutrinos then propagate 570 m through the near detector and then another 1300 km to the far detector. Figure 3.2 shows a schematic of the neutrino beam.

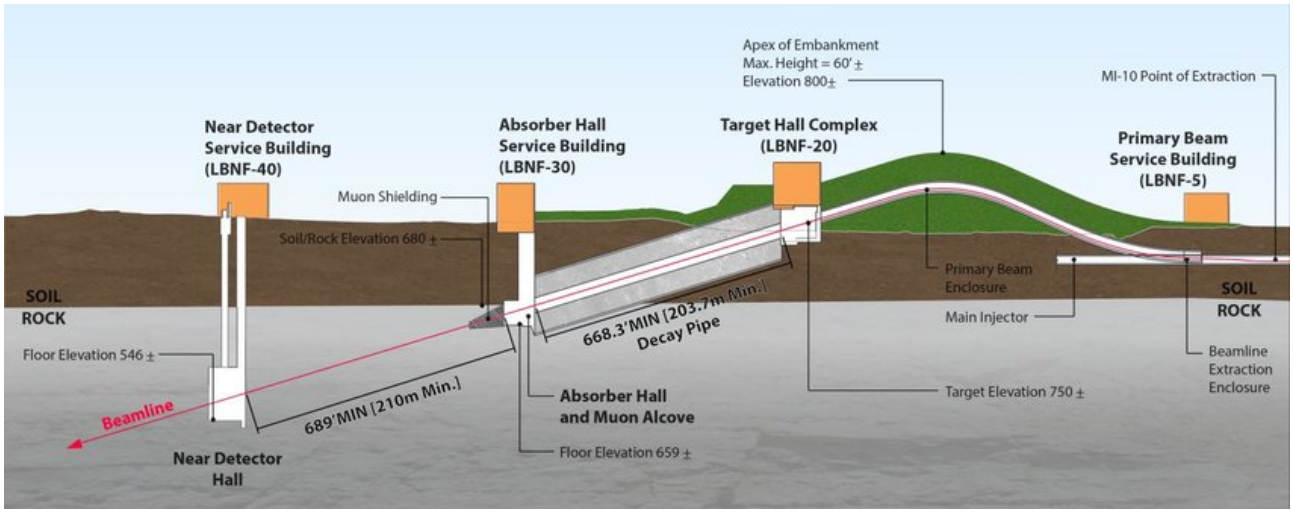

Figure 3.2: Schematic of the LBNF neutrino beamline and DUNE near detector hall [64].

Figure 3.3 shows the simulated DUNE neutrino fluxes at the near detector broken down by the initial decay particle. The vast majority of ν_μ and $\bar{\nu}_\mu$ flux are produced via the decay of charged pions:

$$\pi^+ \rightarrow \mu^+ + \nu_\mu \tag{3.1}$$

$$\pi^- \rightarrow \mu^- + \bar{\nu}_\mu \tag{3.2}$$

These channels account for 99.98% of charged pion decays. The decay of charged kaons, $K^\pm$, also produces a significant portion of the ν_μ and $\bar{\nu}_\mu$ flux, especially the right-sign flux, via the same decay products as Eq 3.1 and Eq 3.2. Right-sign and wrong-sign neutrinos refer to whether the charge of the produced lepton is the same as the signal process i.e. in FHC, ν_μ is right-sign and $\bar{\nu}_\mu$ is wrong sign and vice versa for RHC. The wrong-sign ν_μ and $\bar{\nu}_\mu$ flux is dominated by the decay wrong sign pions, via Eq 3.3 and Eq 3.4, which escape the magnetic horns. A small portion also occurs via muon and antimuon decay:

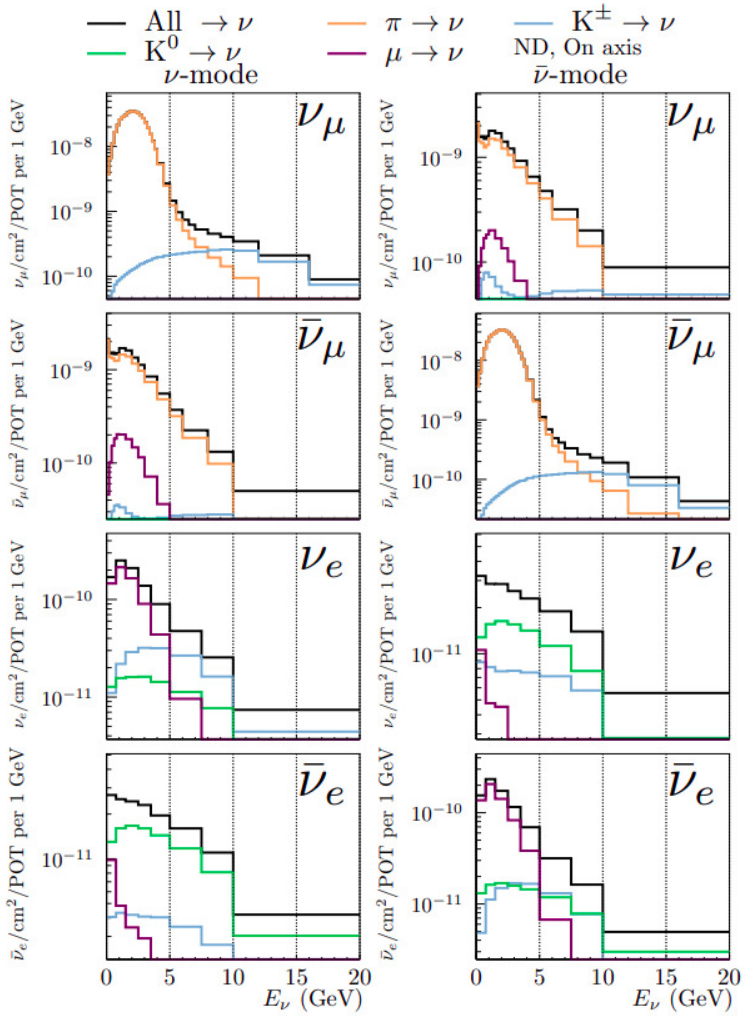

Figure 3.3: Simulated DUNE near detector neutrino on-axis fluxes broken down by initial decay particle for neutrino mode (left) and antineutrino mode (right). muon neutrino, muon antineutrino, electron neutrino and electron antineutrino fluxes are shown from top to bottom [2].

$$\mu^+ \rightarrow e^+ + \nu_e + \bar{\nu}_\mu \tag{3.3}$$

$$\mu^- \rightarrow e^- + \bar{\nu}_e + \nu_\mu \tag{3.4}$$

which, not unsurprisingly, dominates the right-sign ν_e and $\bar{\nu}_e$ flux. The ν_e and $\bar{\nu}_e$ flux in the muon beam before oscillations are called the intrinsic ν_e and $\bar{\nu}_e$ flux and is an irreducible background. The main component of the wrong-sign ν_e and $\bar{\nu}_e$ flux is produced by the decay

of neutral kaons via:

$$K^0 \rightarrow \pi^{\pm} + e^{\pm} + \bar{\nu}_e(\nu_e) \tag{3.5}$$

One might also notice from Figure 3.3 that the total flux of right-sign neutrinos is higher in neutrino mode compared with antineutrino mode. This is a result of the nature of proton-nucleus collisions at the target. Since these interactions involve positive initial states, positive secondary mesons are preferred, resulting in this asymmetry. The magnetic horn is not 100% efficient in defocusing wrong-sign secondary mesons. Hence, this asymmetry also results in a higher intensity of wrong-sign neutrinos in antineutrino mode compared with neutrino mode. Consequently, in neutrino (antineutrino) mode, the beams are 92% (90.4%) muon neutrinos (antineutrinos), with wrong-sign contamination contributing 7% (8.6%) and intrinsic electron neutrino and electron antineutrino background contributing 1% (1%) [2].

The DUNE neutrino beam will have a wide spread of energies, which is enhanced by the on-axis nature of the detectors in Section 3.3 and Section 3.4 [71]. This design choice differs from current and future experiments which have narrow energy beams due to off-axis detectors. The DUNE beam flux will contain both the first and second oscillation maxima. At a baseline of 1300 km, these occur at 2.4 and 0.8 GeV respectively. Observing a spectrum of neutrino energies provides advantages in sensitivity to oscillation parameters, as is discussed in Section 3.5.

3.2 Liquid Argon Time Projection Chamber

The liquid argon time projection chamber, LArTPC, is the principal technology used in DUNE to detect neutrinos. Since its conception, an incredible amount of research and development has led to the maturing and optimisation of the technology. The success of experiments such as ICARUS [72], ArgoNeuT [73], LArIAT [74] and MicroBooNE [75] has led to LArTPCs becoming a leading method in neutrino experiments, alongside Cherenkov detectors and scintillator

detectors. This section will provide a brief history of LArTPC technology, describe the basic operation and explain the motivation for its use in the detectors mentioned in the following sections.

3.2.1 History of Time Projection Chambers

In the early 1970s, bubble chambers and spark chambers dominated the field of tracking detectors in particle physics. However, the rapid development of semiconductor electronics allowed digital readout, rather than photographic, to become feasible. In 1974, the use of a time projection chamber in particle physics experiments was proposed by David Nygren [76]. The idea would retain the qualities of bubble chamber data, such as spatial and timing resolution, but allow digital readout. A general time projection chamber consists of a volume of gas within an electric field with an electron collection system capable of resolving position. When a particle propagates through the gas, ionisation electrons are produced and are drifted by the electric field towards the readout. Combining the two-dimensional readout position with the drift time provides three-dimensional track information. A homogeneous electric field is required to produce an undeformed projection. A magnetic field is often applied to make a momentum measurement, but also reduces transverse diffusion which helps improve position precision. Measuring the track ionisation along with the momentum provides accurate particle identification. Figure 3.4 shows Nygren's drawing of his time projection chamber concept.

Demonstrating total absorption calorimetry in liquid argon detectors [77] led to its use as the medium in time projection chambers. Carlo Rubbia proposed one of the earliest LArTPC designs for use at CERN for the detection of neutrinos [78]. Due to the much greater target mass, noble liquids are used over gas in neutrino experiments which increases the probability of neutrino interaction. Noble elements have no electronegativity, resulting in high electron mobility and low transverse and longitudinal diffusion. Furthermore, noble elements possess excellent dielectric properties, necessary due to the high electric field used to drift the ionisation electrons. Of the noble elements, Argon is strikingly abundant in the Earth's atmosphere and hence relatively inexpensive, especially important for large volume detectors. Finally, liquid

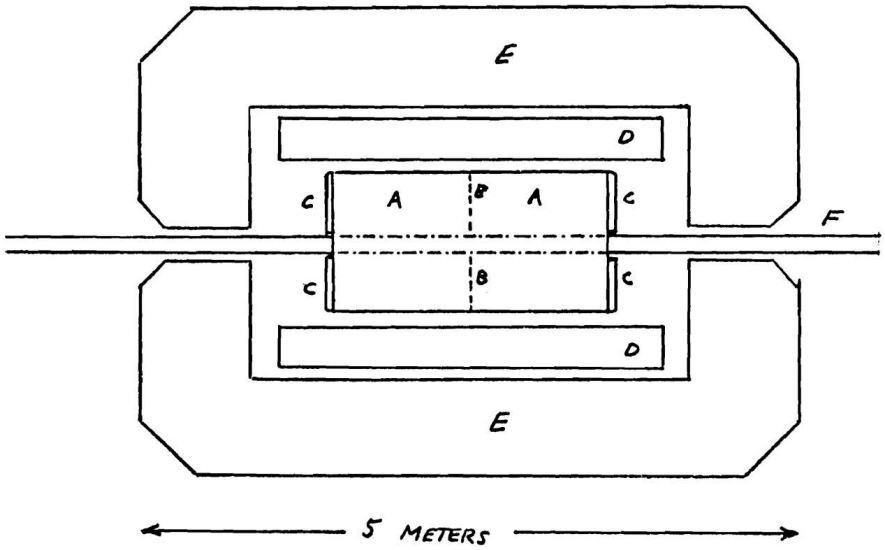

Figure 3.4: Original hand-drawn TPC design by David Nygren [76]. Sections are labelled as: methane-filled region (A), screen/foil to establish electric field (B), end-cap detector (C), superconducting solenoid (D), iron return yoke for magnetic field (E), beam vacuum pipe (F).

argon has an ionisation threshold of 23.6 ± 0.5 eV [79], meaning an LArTPC would also have a low detection threshold. These properties led Carlo Rubbia to describe liquid argon as an 'almost ideal target material' and Figure 3.5 shows his conceptual design.

3.2.2 LArTPC Operation

A LArTPC consists of at least one anode and cathode separated by a drift region, as shown in Figure 3.6. A strong uniform electric field is applied across this drift region, causing ionisation electrons produced by charged particles to drift towards the anode. Charge depositions are collected and read out providing an 'image' of the tracks of charged particles involved in the interaction.

The readout of ionisation electrons is generally done via multiple anode wire planes, also illustrated in Figure 3.6. If a detector has N anode planes, the $N-1$ planes closest to the cathode are called induction planes. These wire planes are held at lower (more negative) potentials, allowing electrons to drift past resulting in induction signals. These induction signals are bipolar since the sign of the signal is dependent on whether the electrons are moving towards or away from the plane. Figure 3.6 shows the induction planes labelled as U and V. The final plane is

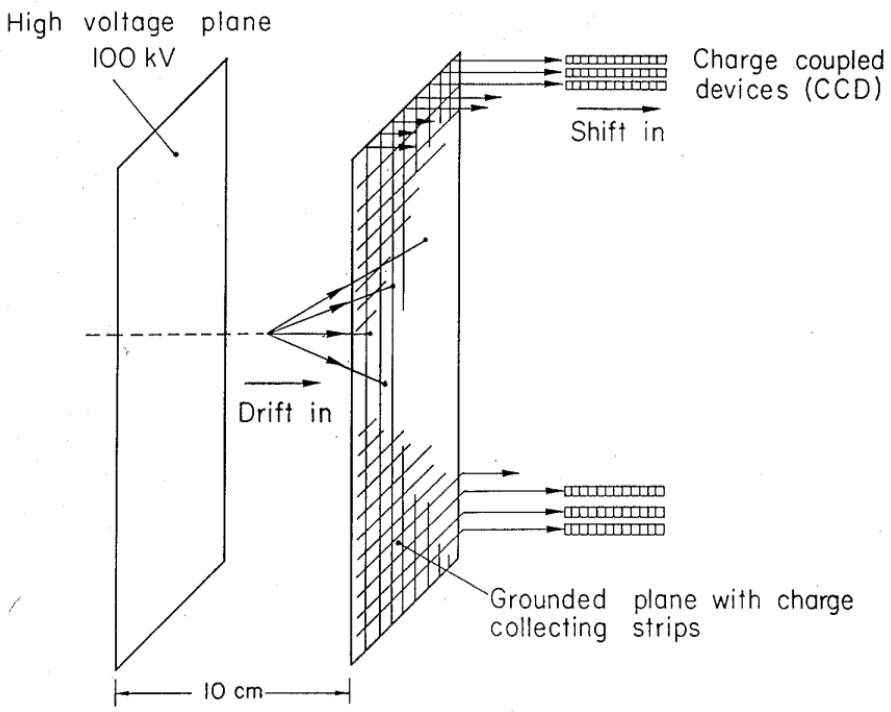

Figure 3.5: Carlo Rubbia's LArTPC detector design [78].

called the collection plane. The electric field is designed such that all field lines terminate on the collection plane. The signal caused by the collection of electrons is unipolar. A grounding layer is often included to shield the active planes from long-range induction effects. Without this plane, the induction signals appear asymmetrical.

We have just described a system for drifting ionisation electrons to a readout plane, however unsurprisingly these electrons can be lost on the way. One mechanism for this is the recombination of the ionisation electron with the atom before it can be fully separated. Although this effect harms the signal, it produces scintillation light which can be collected on a timescale of nanoseconds. This is often used to determine the interaction time, t_0. Using the observed t_0 and the electron drift velocity allows the third dimension, along the drift direction, to be resolved. In liquid argon, there are two mechanisms for the emission of scintillation light, as illustrated in Figure 3.7. The passage of ionising radiation can either cause the argon atom to become excited (top) or ionised (bottom). In either scenario, argon molecules are formed before de-excitation causes the emission of scintillation light.

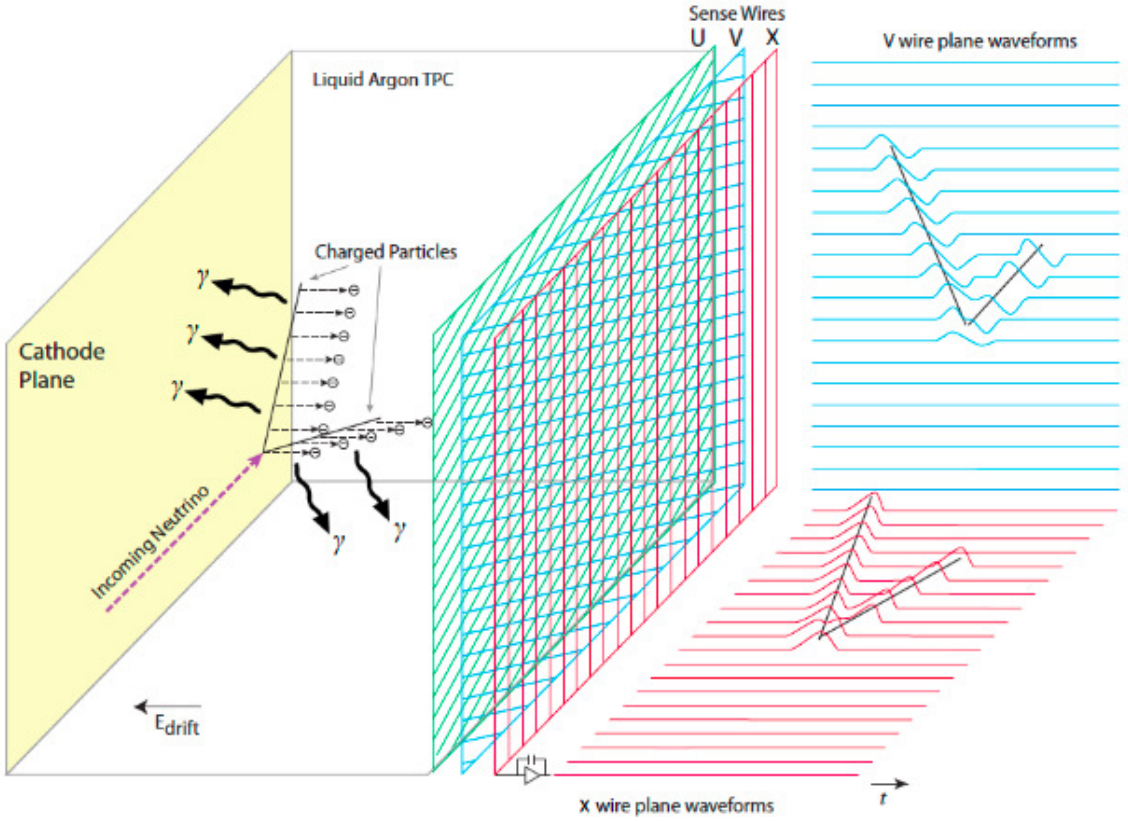

Figure 3.6: Basic operation of a LArTPC [80].

The probability of re-combination heavily depends on the strength of the electric field. The higher the voltage, the lower the probability of recombination and hence less scintillation light. However, a higher voltage also results in a better signal-to-noise ratio, better calorimetry and lower detection thresholds. Therefore the voltage becomes a trade-off.

Electronegative impurities in the liquid argon also prevent ionisation electrons from reaching the anode. Oxygen and water are common electronegative impurities. This effect is quantified by assigning an electron lifetime, τ, which is used to predict the observed charge, Q, via

$$Q(t) = Q_0 \exp -\frac{t}{\tau} \tag{3.6}$$

where t is the drift time and Q_0 is the initial ionisation charge. The electron lifetime is directly affected by the argon purity. The argon purity target for DUNE results in a τ of 3 ms

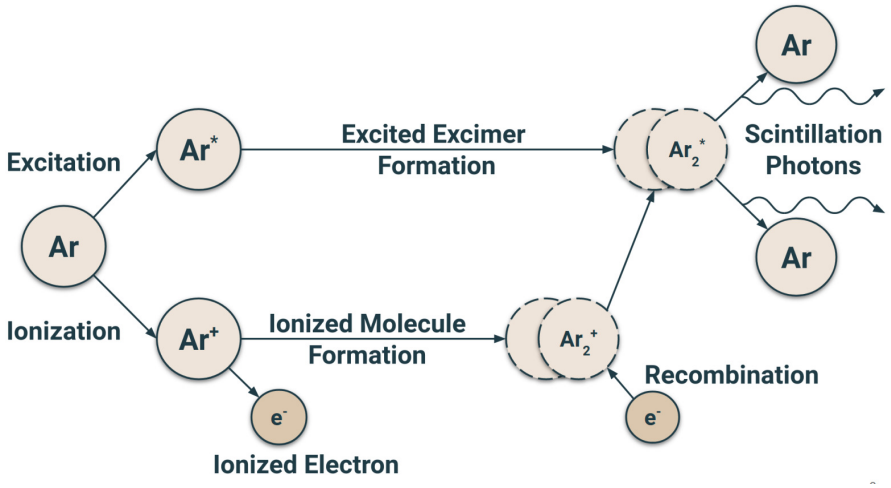

Figure 3.7: Schematic of scintillation light production in argon from [66].

[64] which is necessary for the required signal-to-noise ratio. Nitrogen contaminants also quench the scintillation photons via collisions with the excited argon atoms, de-exciting them without photon emission [81]. Ionisation electrons also often undergo interactions causing electrons to drift off of the field direction. This can change the size and apparent location of observed signals.

The above describes the operation of the default single-phase LArTPC. An alternative dual-phase technology exists which adds an extra step before the charge collection. After drifting through liquid argon, ionisation electrons then pass through a volume of gaseous argon. This amplifies the measured signal via electron avalanches, in which the ionisation electrons also ionise the gaseous argon atoms. Dual-phase technology provides two main advantages. Firstly, the charge amplification in the gaseous volume results in a better signal-to-noise ratio than for a single-phase detector. This lowers the detection threshold allowing lower energy events to be observed. Secondly, the charge amplification also allows for longer drift lengths. Hence, if both detectors have the same dimensions, the dual-phase detector would in principle have a larger fiducial volume due to fewer dead regions. Initially, DUNE considered both technologies for detector modules. More recently, DUNE has moved away from dual-phase technology.

3.3 Near Detector Complex

One of the key aims of DUNE is to extract neutrino oscillation probabilities by measuring neutrino interaction rates. Extracting these probabilities as a function of energy will allow DUNE to make precise measurements of the parameters of the PMNS matrix [20]. The near detector (ND) complex will measure the neutrino interaction rate 575m downstream of the neutrino beam source [64]. At this distance, it can be assumed no neutrino oscillations given the three neutrino paradigm. Hence, the role of the ND is to serve as the control of the experiment.

Label	Description
ND-O0	Predict the observed neutrino spectrum at FD
ND-O1	Transfer measurements to FD
ND-O2	Constrain the cross section model
ND-O3	Measure the neutrino flux
ND-O4	Obtain measurements with different fluxes
ND-O5	Monitor time variations of the neutrino beam
ND-O6	Operate in high rate environment

Table 3.1: Overarching requirements of the DUNE near detector [69]. ND-O0 represents the ultimate goal of the DUNE near detector. ND-O1 - ND-O6 outline the ingredients needed to achieve ND-O0.

In any long-baseline neutrino experiment, the observed rate of neutrino events in flavour β from a neutrino beam of flavour α as a function of the reconstructed energy, E'_ν, is given by

$$N_{\alpha\to\beta}(E'_\nu) = \int \Psi_\alpha(E_\nu) \times \sigma_\beta(E_\nu) \times \epsilon_\beta(E_\nu) \times M_\beta(E'_\nu; E_\nu) \times P_{\alpha\to\beta}(E_\nu)\, dE_\nu \tag{3.7}$$

where E_ν is the true neutrino energy, $\Psi_\alpha(E_\nu)$ is the neutrino beam flux in flavour α, $\sigma(E_\nu)$ is the neutrino cross-section for flavour β, $\epsilon(E_\nu)$ is the detection efficiency for neutrino flavour β, $M_\beta(E'_\nu; E_\nu)$ is a matrix linking the true and reconstructed energy, and $P_{\alpha\to\beta}(E_\nu)$ is the

probability of neutrino oscillation from flavour α to flavour β. This is further complicated by the numerous possible final states in a neutrino interaction, each of which has different $\sigma(E_\nu)$, $\epsilon(E_\nu)$ and $M_\beta(E'_\nu; E_\nu)$. Hence, with just one detector, one would need a perfect understanding of the flux, cross-section and detector effects to extract the oscillation probabilities. Since there is no handle to do this in just a single far detector, an ND is used to constrain these uncertainties. Given the ND has no oscillations, the event rate can be written as

$$N_\alpha^{ND}(E'_\nu) = \int \Psi_\alpha^{ND}(E_\nu) \times \sigma_\alpha(E_\nu) \times \epsilon_\alpha^{ND}(E_\nu) \times M_\alpha^{ND}(E'_\nu; E_\nu)\, dE_\nu \tag{3.8}$$

where the beam is assumed to only contain neutrinos of flavour α. This gets further complicated by the differences in these models between the near detector and the far detector. Generally, near detectors are designed to have similar efficiencies and responses as the far detector. A summary of the overarching requirements of the DUNE near detector is shown in Table 3.1.

The DUNE project will occur in two phases. The transition from Phase I to Phase II requires the construction of a more capable near-detector complex. Figure 3.8 shows the finally envisioned DUNE near detector complex within the detector hall, with the neutrino beam entering from the right-hand side. The beam initially passes through ND-LAr, a modularised LArTPC. The beam then passes through ND-GAr, a high-pressure gaseous argon TPC. ND-LAr and ND-GAr will also move to take data in positions off the beam axis, a capability called the DUNE Precision Reaction-Independent Spectrum Measurement (DUNE-PRISM) [2]. Finally, the beam will pass through a magnetised beam monitor called the System for on-Axis Neutrino Detection (SAND). ND-GAr will be added in Phase II of the experiment. In Phase I, the Temporary Muon Spectrometer, TMS, will be adjacent to ND-LAr. Section 3.3.1 - Section 3.3.4 will discuss each of these near detector components in more detail.

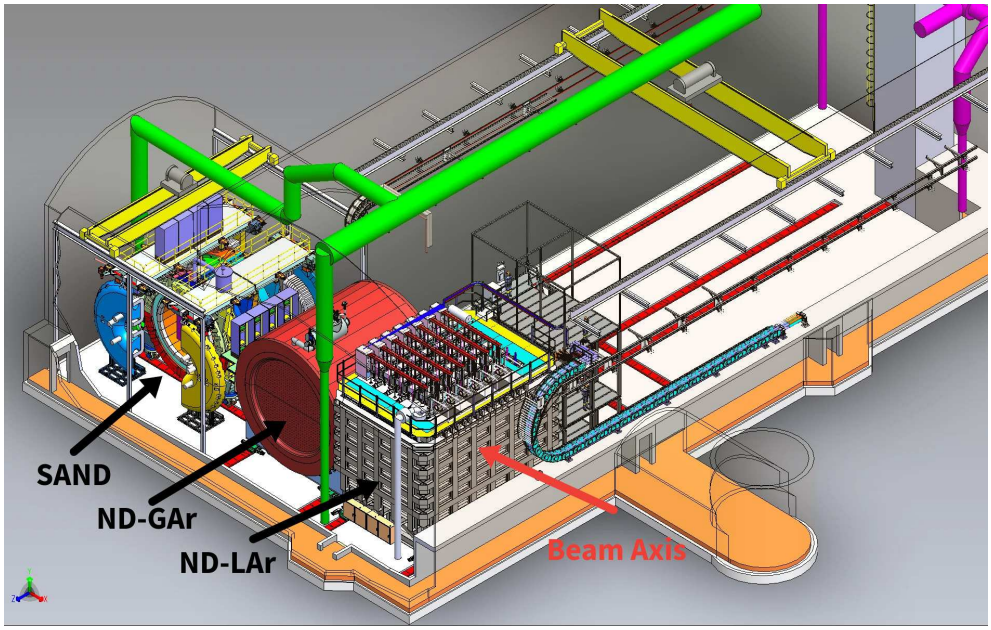

(a) On-axis

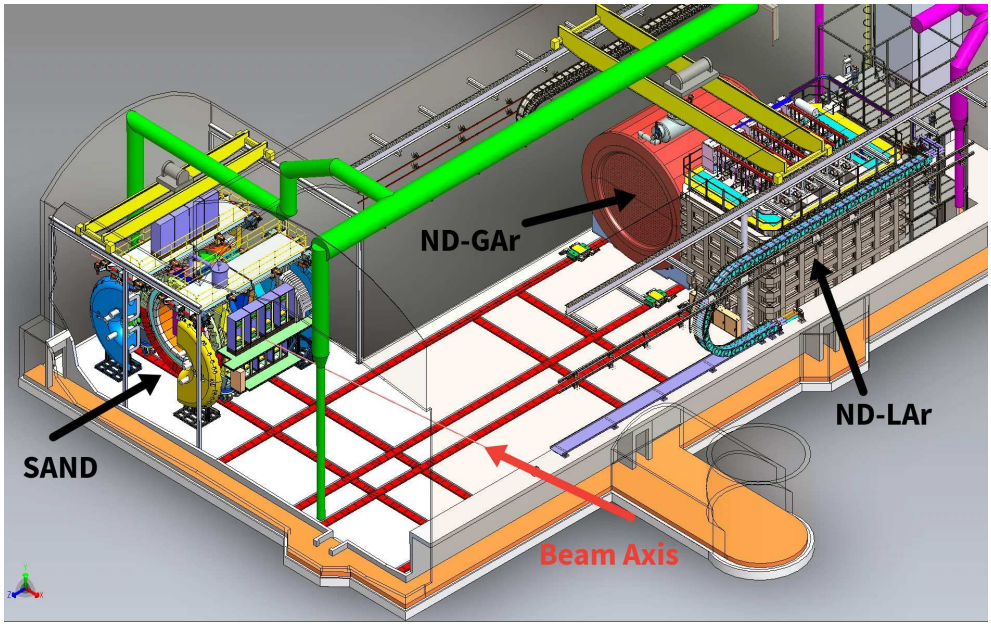

(b) Off-axis

Figure 3.8: Schematic of DUNE ND near detector complex showing composite detectors and beam direction in both (a) on-axis configuration and (b) with ND-LAr and ND-GAr in an off-axis configuration [69].

3.3.1 ND-LAr

As mentioned in Section 3.2, LArTPC is the principal technology used in DUNE. Since the far detector modules discussed in Section 3.4 use liquid argon as a target material, collecting a

sample of neutrino interactions on the same target at the near detector is vital for the reduction of cross-section and detector systematics.

The proximity of the ND to the source results in an intense neutrino flux. In each 10μs beam spill, ND-LAr will host $\mathcal{O}(50)$ neutrino interactions. Given the drift time of electrons in liquid argon across $\mathcal{O}(m)$ is $\mathcal{O}(ms)$, this would result in a significant pile-up of events. To remedy this, ND-LAr will use a novel technology called ArgonCube [82], which utilises detector modularisation. ND-LAr will consist of 35 LArTPC optically isolated modules lying within a common cryostat. With such high event rates, associating the scintillation light of an interaction with the relevant module is crucial for reconstruction. Together, the modules form a 5m (beam direction), 7m (width), and 3m (height) structure with a fiducial mass of 67 t. The dimensions of ND-LAr were driven by the need for hadronic containment, essential for the accurate reconstruction of neutrino energy. Each module consists of a cathode, field shell, light collection system and pixelated charge readout. The ArgonCube technology uses a resistive field shell over traditional field rings to minimise the dead regions caused by modularisation.

ND-LAr is required to measure signals from ionisation at the same spatial granularity or better than that in the far detector. Due to the high multiplicity of events, each LArTPC will use novel pixelated anodes capable of providing a 3D record of ionisation signals. A pixel spacing of 4 mm is chosen, corresponding to a density of 60,000 channels/m^2. Combined with a known beam spill width, the pixel readout will deliver a position accuracy of 16mm. A central cathode splits each module into two drift regions, with pixel tiles placed on either side. In addition to performance, the pixel readout design is driven by manufacturability and reliability. The latter is made important due to the difficulty of accessing the detector once the cryostat is filled with liquid argon.

As mentioned in Section 3.2, detecting scintillation light provides an absolute interaction time reference, t_0. The ND-LAr light collection system will provide a more accurate t_0 than utilising beam spill information. In addition, the light system can tag disassociated charge depositions, allowing all charge to be tagged to the correct neutrino event. Figure 3.9 shows

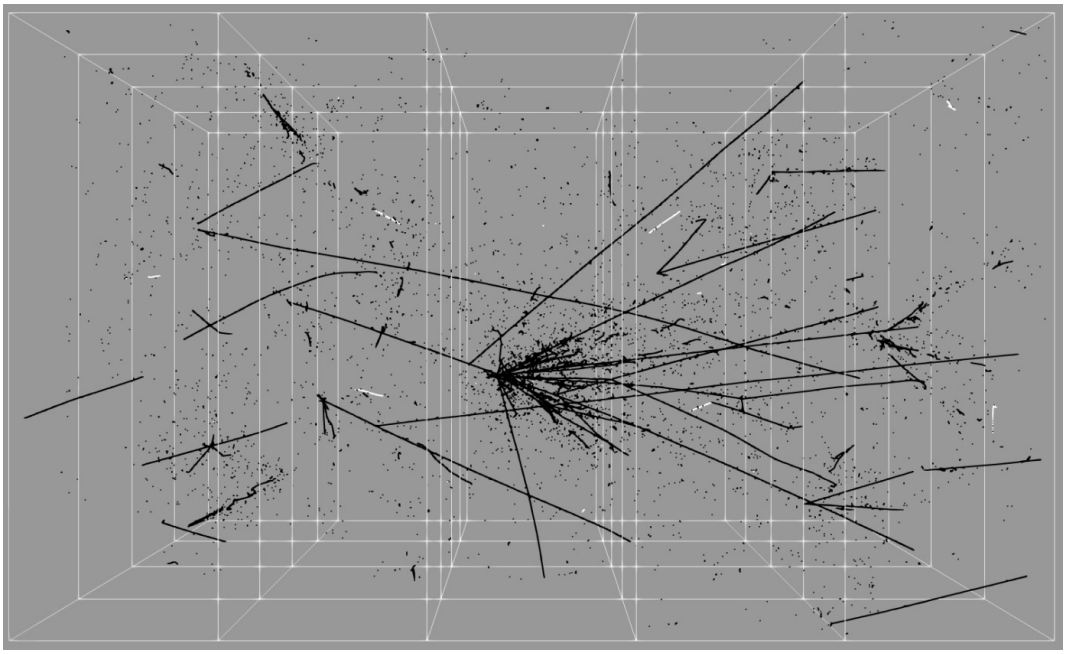

Figure 3.9: Simulated beam spill in ND-LAr. White tracks indicate fast-neutron-induced proton tracks (energy threshold of 10 MeV). All other deposits are shown in black [69].

energy deposits from a simulated beam spill in ND-LAr, highlighting the magnitude of charge deposits. Fast-neutron induced proton recoiling tracks are shown in white. One can see the difficulty of associating these with a neutrino vertex using charge alone.

For muons with momenta > 0.7 GeV, the probability of containment within ND-LAr begins to drop. If muons escape, their energy cannot be calculated calorimeterically. Furthermore, since ND-LAr is not magnetised, it is not able to determine the sign of the muons. By measuring the charge of the muon, the wong-sign background of the neutrino beam can be measured. This drives the need for a downstream spectrometer to measure these particles. In Phase I of the experiment, ND-LAr will be accompanied by TMS. Although the design is not yet confirmed, TMS will generally be a magnetised steel-scintillator sampling detector. The steel plates provide the density to help contain the muons, while the scintillator allows for track reconstruction. Magnetising the steel allows μ^- tracks to be distinguished from μ^+ tracks, hence distinguishing ν_μ and $\bar{\nu}_\mu$ interactions. In Phase II, TMS will be replaced with ND-GAr, which can measure the escaping muons as well as provide vital information to help study neutrino interaction kinematics

3.3.2 ND-GAr

The success of next-generation long-baseline experiments heavily relies on the reduction of neutrino interaction uncertainties. For DUNE, this means a better understanding of neutrino-on-argon interactions. The inclusion of a gaseous argon detector allows DUNE to study interactions at lower energies in far more detail, acting as a "microscope" for nuclear effects. This information can be used to probe and constrain parts of the interaction model.

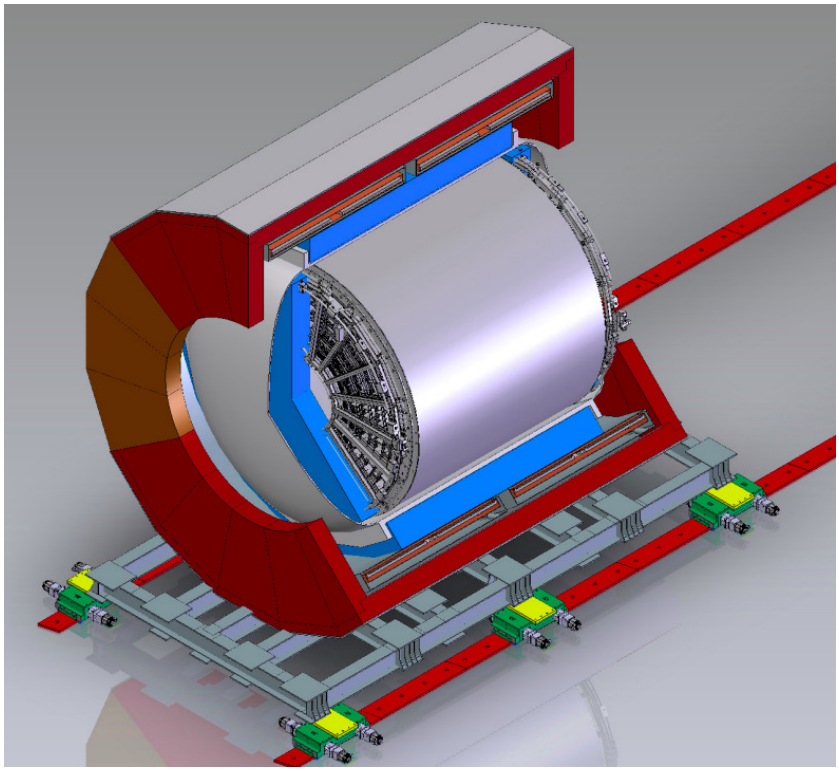

Figure 3.10: Schematic of ND-GAr detector showing the HPgTPC, the ECAL, the magnet and the return iron [69].

Figure 3.10 shows a cut-away schematic of ND-GAr. At its core, there is a high-pressure gaseous argon TPC (HPgTPC). The cylindrical TPC measures 5.2m in diameter and 5m in length. A central cathode will split the TPC into two drift regions. At the conceptual design phase, the details of the HPgTPC design were based on the ALICE TPC [83], a detector dedicated to heavy-ion physics at the Large Hadron Collider (LHC). In the ALICE TPC, the end plates of the cylinder are instrumented with multi-wire proportional chambers (MWPCs) which initiate avalanches at the anode wires. These avalanches induce signals on cathode pads

behind the MWPCs which are proportional to the initial ionisation charge. For ND-GAr, more modern amplification stages, such as gaseous electron multipliers (GEMs), are being considered. The conceptual design target gas is an argon-CH_4 mixture, 9 : 1 ratio, at a pressure of 10 bar. High pressure is required to ensure sufficient neutrino interactions. At 10 bar with this gas mixture, ND-GAr will have a fiducial mass of 1.8 t.

The HPgTPC will be surrounded by an electromagnetic calorimeter (ECAL). The principal role of the ECAL is to reconstruct photons produced from neutrino interactions. This will allow the t_0 of both tracks entering the HPgTPC and interactions occurring in the HPgTPC to be measured. Since the ECAL can measure both energy and direction of electromagnetic showers, the vertex π^0 decays within the HPgTPC can be determined. Events producing a π^0 pose as a background to ν_e events in ND-LAr, which the ECAL can reject. In addition, the ECAL can tag external background events such as rock muons and neutrons.

The HPgTPC and ECAL will be surrounded by a 0.5T magnetic field. It must run perpendicular to the particle path to cause curvature in the tracks. This allows both the sign of charged particles and their momenta to be measured. The sign measurement is important for the constraint of the wrong-sign background in the neutrino beam. The return iron yoke is required to manage the escaping field which can interact with surrounding detectors.

Since ND-GAr will have a much lower density than ND-LAr, particles will produce longer tracks. Therefore, the threshold energy for particle reconstruction in ND-GAr will be much lower. This is especially important for protons and charged pions. Due to complexities in neutrino-nucleus interactions, final-state particle kinematic distributions are extremely hard to predict. Figure 3.11 shows predictions of the low-energy proton distributions by three different neutrino generators. One can see that at kinetic energies below the liquid argon threshold, the predictions begin to diverge. Hence, the low-energy final states reconstructed in ND-GAr could inform nuclear modelling choices.

The reconstruction of low-energy pions is also vital for energy reconstruction. Unlike protons which are knocked out of the nucleus, final-state pions are produced from the incoming

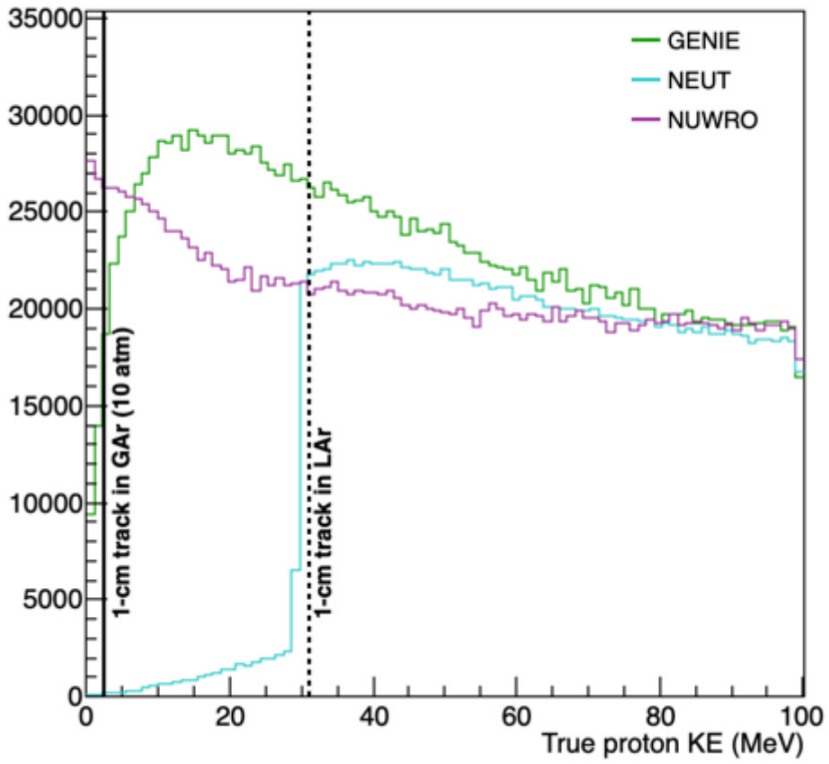

Figure 3.11: Comparison of true proton energy spectra predicted in three neutrino event generators (GENIE, NEUT and NUWRO). The vertical lines indicate the energy threshold required to produce a 1 cm track in liquid argon (dashed) and gaseous argon at 10 atm (solid) [69].

neutrino energy. Therefore, if a pion is missed in the ND-LAr, the reconstructed energy will be off by at least the pion rest mass. Since events can often have multiple final-state pions, this has a significant effect.

3.3.3 DUNE-PRISM

Neutrino-nucleus interaction modelling is the dominant systematic uncertainty in next-generation long-baseline experiments. The mapping between the observable final state and the incoming neutrino energy is not understood well enough to achieve DUNE's physics goals. Since there is currently no fully 'correct' interaction model, simply constraining the available parameters with the ND data could end up in biases. DUNE-PRISM takes advantage of the relationship between neutrino energy and observation angle to reduce the dependence of an analysis on the neutrino interaction model.

The energy of a neutrino from a pion decay can be expressed as [71]

$$E_\nu = E_\pi \left(1 - \frac{m_\mu^2}{m_\pi^2}\right) \frac{1}{1 + \theta^2 \gamma^2} \tag{3.9}$$

where E_π is the pion energy, m_π is the pion rest mass, m_μ is the muon rest mass, θ is the angle between the pion and the neutrino and γ is the Lorentz factor of the parent pion. It becomes clear that as θ increases, the neutrino energy decreases. For a neutrino beam produced from a range of pion energies, increasing θ reduces both the peak and the spread in the neutrino energy spectrum, as shown in Figure 3.12.

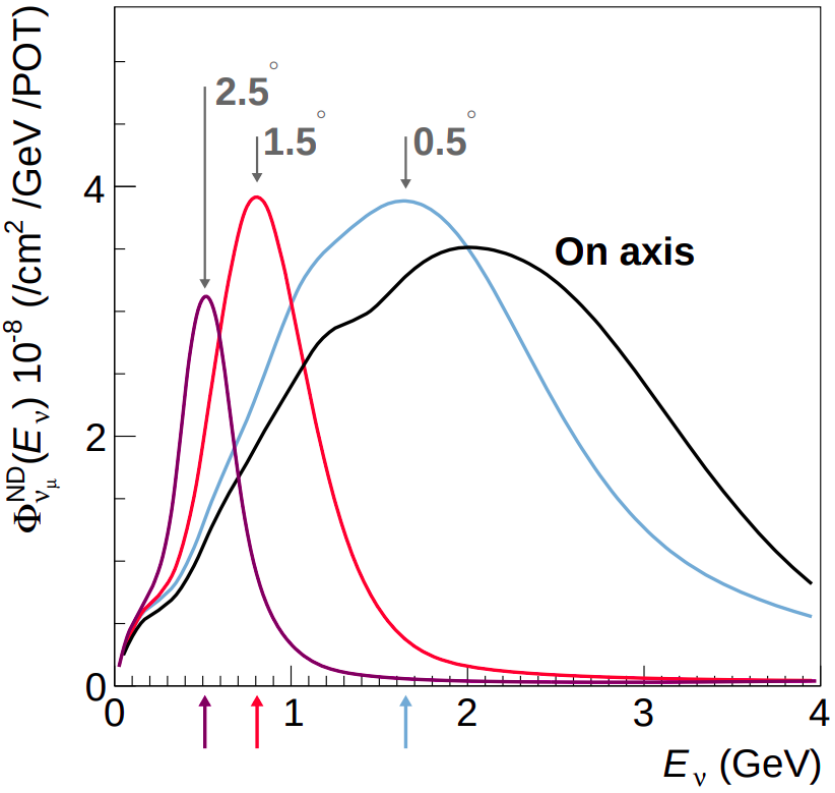

Figure 3.12: Predicted DUNE ν_μ flux at the ND for different off-axis angles. Arrows indicate peak neutrino energy [69].

This property is already widely used in current long-baseline experiments to achieve narrower energy distributions e.g. T2K (2.5° [84]) and NOvA (0.8°) [54]. Since DUNE-PRISM aims to sample the flux for a wide range of angles, ND-LAr and ND-GAr will move laterally

within the ND hall through an off-axis angle range of 0° to 3.2°. These off-axis samples can be incorporated into an analysis in two ways. Firstly, they can be compared with MC predictions, helping to identify potential biases within the neutrino interaction model that would not have been visible using just the on-axis samples. Secondly, these off-axis ND samples can be linearly combined to predict the observed FD neutrino energy spectrum. This technique, in principle, allows you to bypass the neutrino interaction model. This would be vital in the scenario that there is no theoretical neutrino interaction model which can predict the observed ND spectra with sufficient precision. However, this approach relies on an accurate prediction of the flux and any flux uncertainties snowball due to the combination of several neutrino fluxes.

3.3.4 SAND

In a long-baseline experiment, the neutrino flux must be constantly monitored so that any possible distortions that occur in the beamline are identified. These distortions are often caused by the horn positions and currents. These are most easily detected on-axis as there is a wider energy distribution. Also, for DUNE-PRISM to work correctly, the relationship between the off-axis angle and neutrino energy must be known well, which requires a stable neutrino beam. For these reasons, DUNE needs a continuous on-axis beam monitoring system which will be fulfilled by SAND.

The reference design for the SAND is illustrated in Figure 3.13. The active target uses a 3D scintillator tracker (3DST), made of several optically isolated plastic scintillator cubes and read out by three orthogonal optical fibres. This will be surrounded by low-density tracking chambers aimed at measuring the momentum and charge of outgoing particles. The design of these trackers is not decided, but the technologies being considered are TPCs, straw tube trackers (STT) or some combination of the two. An alternative design includes layers of carbon and hydrocarbon as extra targets for neutrino interaction. In theory, a sample of neutrino-hydrogen events can be extracted by taking the difference in the spectra for these two targets. This can then be used to directly constrain the neutrino-nucleon element of the interaction model without any complex nuclear effects. Exterior to the trackers will be a superconducting

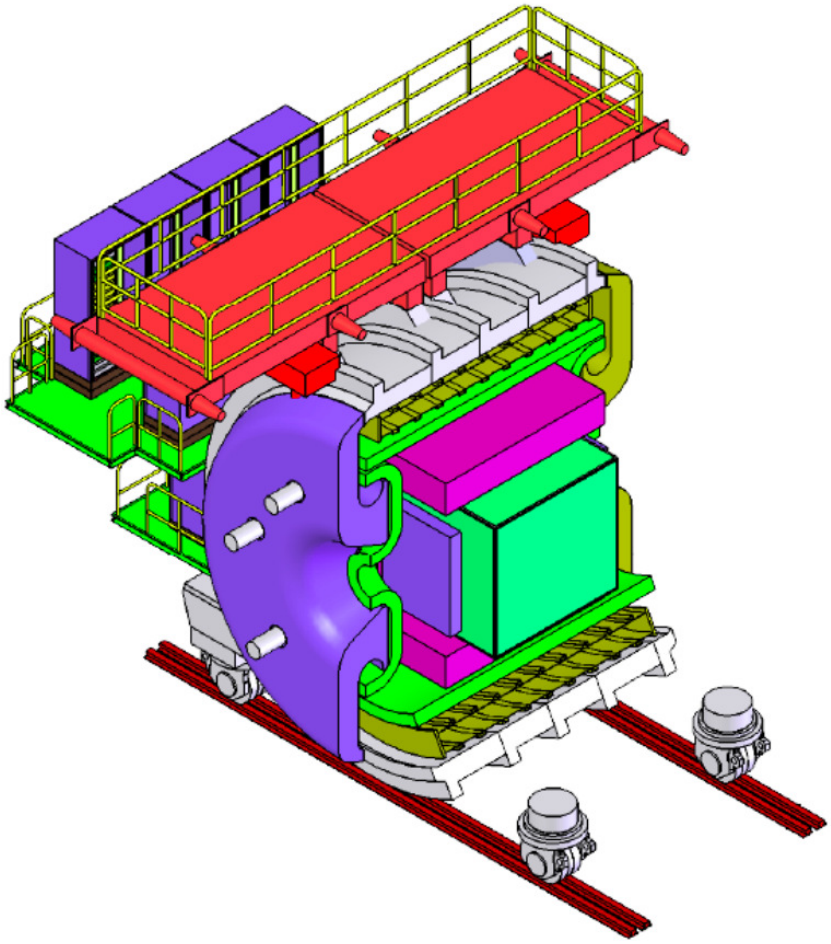

Figure 3.13: Drawing of the SAND reference design showing the central 3DST (light green), low-density tracker (magenta), ECAL (green), magnet coil (gold), and return yoke (grey) [69].

magnet and ECAL reused from the KLOE experiment [85]. The magnet can produce a 0.6T field, allowing the charge and momentum to be measured within the trackers. The KLOE ECAL is composed of lead/scintillating fibres, offering good light transmission and $\mathcal{O}(ns)$ timing.

3.4 Far Detector

The DUNE far detector (FD) will be located at the Sanford Underground Research Facility (SURF) in South Dakota. Achieving DUNE's physics program requires the observation of rare events with millimetre-scale spatial resolution. The FD also needs a many-kiloton fiducial mass to host these events in sufficient numbers. As discussed in Section 3.2, the LArTPC provides total calorimetry and high-granularity tracking and has been chosen as the principal technology

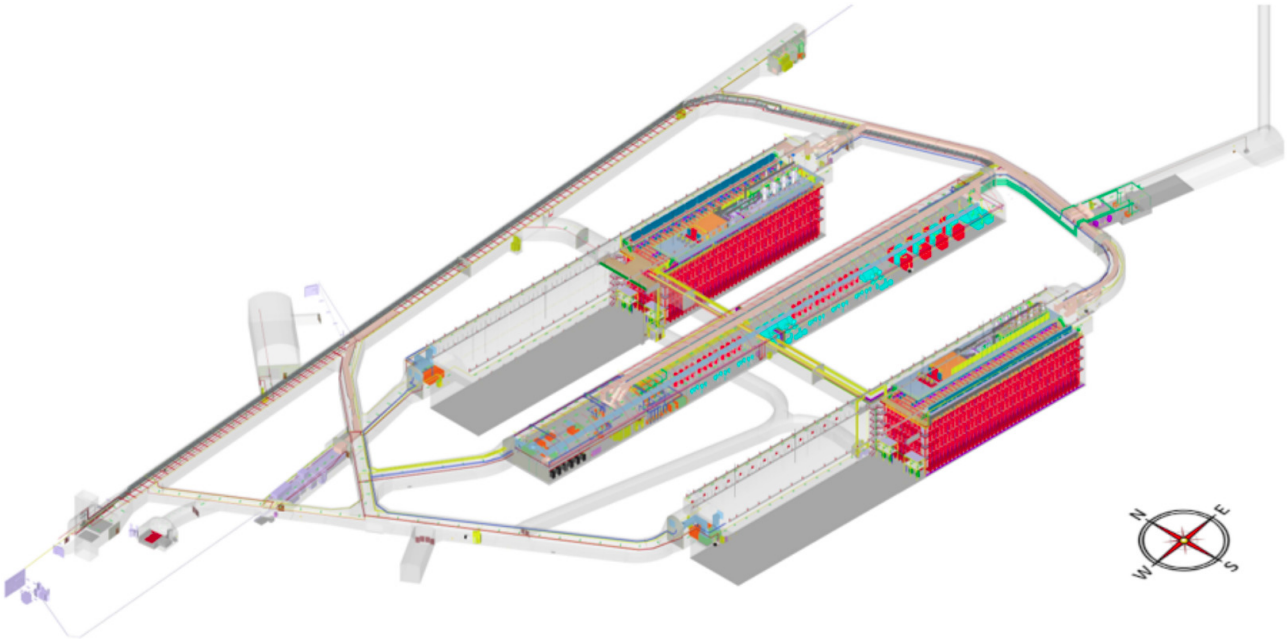

Figure 3.14: The layout of SURF underground caverns, showing the location of two detector modules at the east end of the north and south cavern. The central utility cavern will host infrastructure, utilities and upstream data acquisition [68].

in the DUNE FD. The FD will consist of four independent modules, each containing roughly 17.5kt of LAr. As of now, the specific design of two out of four modules has been chosen. The first of these modules, FD1-HD, will use horizontal drift technology, whereas the second, FD2-VD, will use novel vertical drift technology. Both of these modules will be single-phase as DUNE has moved away from dual-phase designs. Figure 3.14 shows FD1-HD and FD2-VD placed in the east end of the north and south cavern respectively at SURF. Both modules will be available in Phase I of the experiment. The third and fourth FD modules will be added in Phase II, the designs of which have not been decided. FD3 will likely be HD or VD. DUNE is welcoming designs for FD4, named the "Module of Opportunity", which can offer new physics. This section will focus on the proposed designs of FD1-HD, Section 3.4.1, and FD-VD2, Section 3.4.2.

3.4.1 Horizontal Drift Module

Horizontal drift technology has been used successfully in previous LArTPC detectors including those described in [72, 75, 86]. The ionisation electrons produced by charged particles are

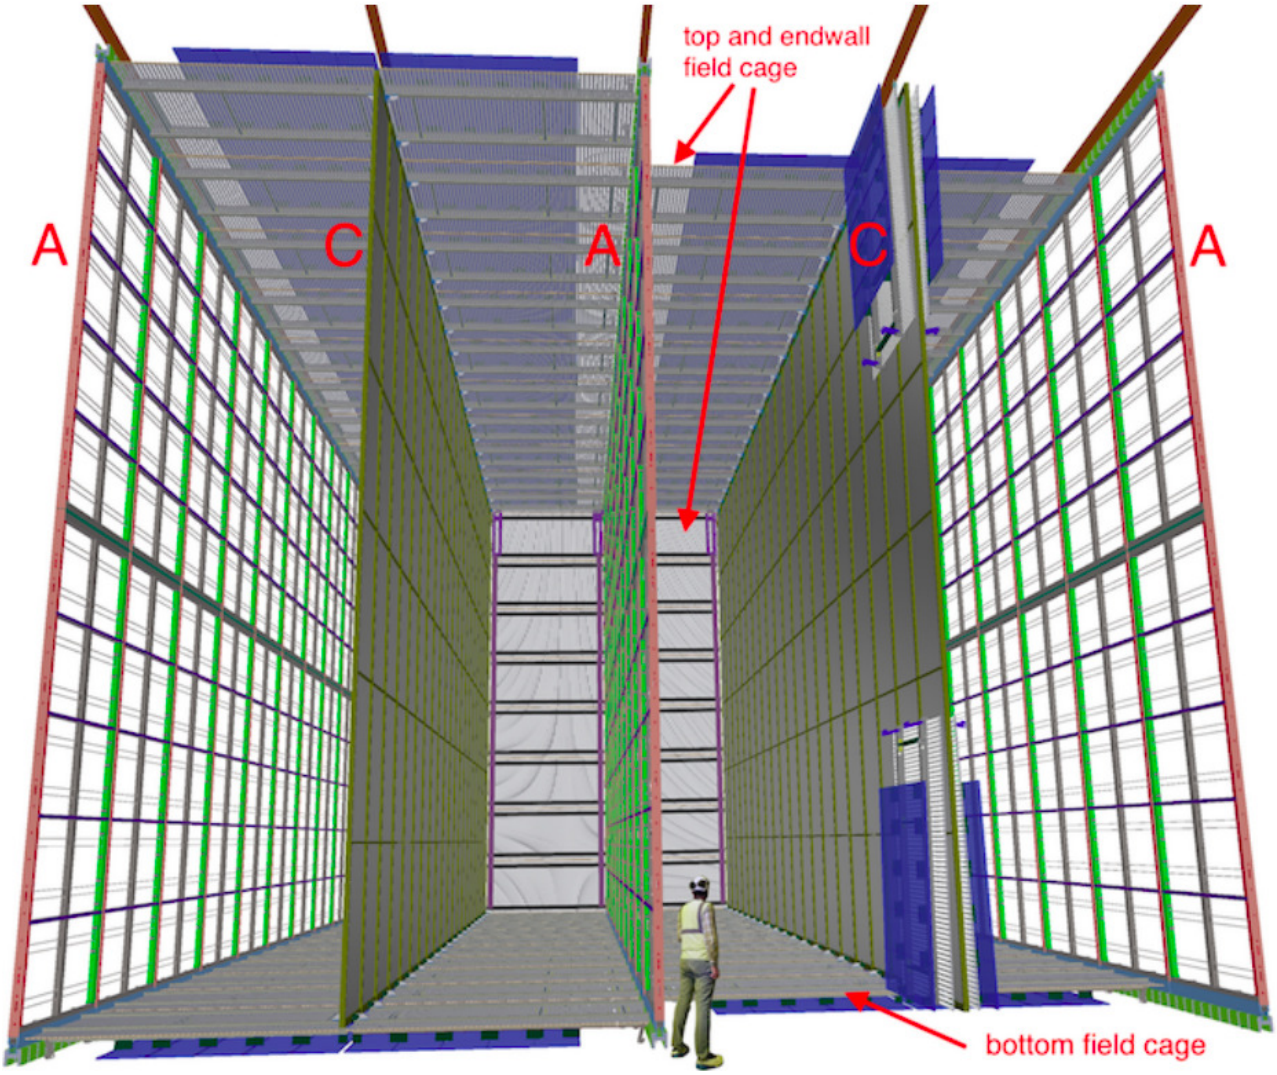

Figure 3.15: A DUNE horizontal drift module showing alternating vertical anode (A) and cathode (C) planes, as well as the field cage that surrounds the drift region. On the right-hand cathode plane, the foremost portion of the field cage is shown in its folded state [66].

drifted horizontally by vertically orientated anode and cathode planes. Figure 3.15 shows a schematic of FD1-HD. The module will have a 12 x 14 x 58.2 m active volume. The module interior consists of alternating anode and cathode walls, two cathode walls and 3 anode walls. Each of these walls is then made up of individual Anode or Cathode Plane Assemblies, APAs or CPAs.

Each APA measures 2.3m wide and 6m tall, and hence each wall consists of two rows of twenty-five APAs across the length of the module. As described in Section 3.2.2, LArTPC readout generally consists of multiple anode wire planes. The APAs in FD1-HD will consist of

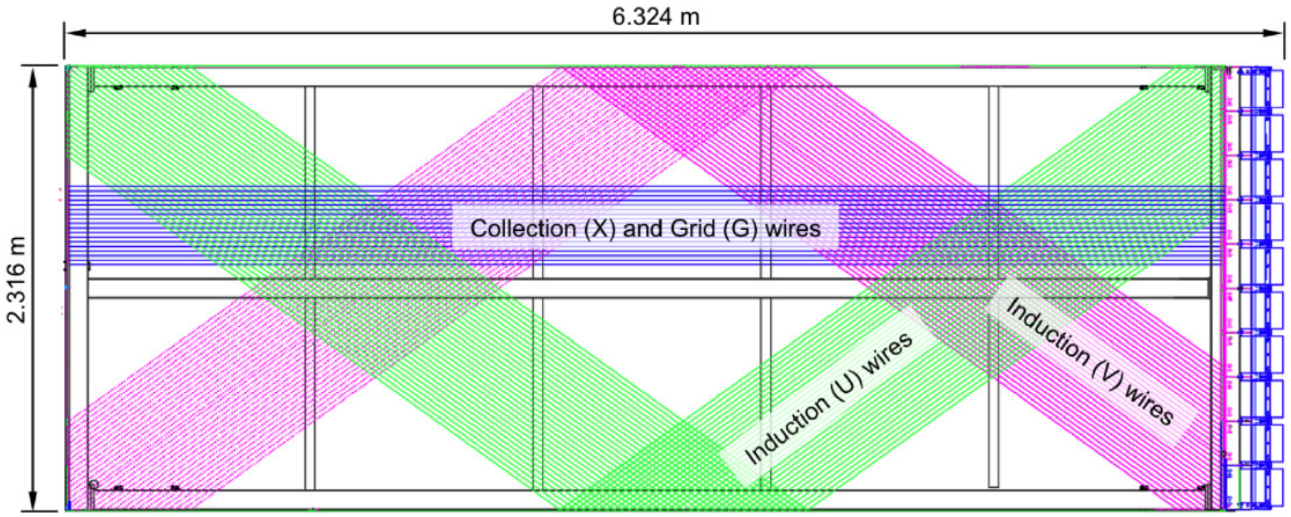

Figure 3.16: An illustration of an APA, showing the wrapping of the induction wires (U, V) in green and pink. The collection and grid wires (X, G) are shown in blue. The blue boxes on the right-hand side are the cold electronics [66].

a collection layer (X), two induction layers (V and U) and a grounding layer (G). An FD1-HD APA is shown in Figure 3.16. The X and G layers run parallel to the long axis of the APA, while the U and V layers run at $\pm 35.7°$. This angle is chosen such that each induction wire overlaps just one collection wire, reducing ambiguity on the reconstruction, while also minimising the dead space within the APA. The wire spacing determines the resolution of the APA, chosen to be 4.8 mm for X and G layers and 4.7 mm for U and V. The TPC readout electronics, referred to as "cold-electronics", are attached to the APAs and must function within the cryostat. Their job is to send out digitized signals from the APA wires. The electronics are also designed to keep the signal-to-noise ratio high, which is aided by low temperature in the cryostat.

Each cathode wall consists of three rows of fifty CPAs, held at -180 kV. Since the anode walls are held at roughly ground, this results in a 500 V/cm uniform electric field. The drift volume is surrounded by an aluminium field cage, which keeps the electric field uniform and terminates any extruding field lines.

As discussed previously, LArTPCs use scintillation photons to record a t_0 for each interaction, allowing reconstruction in the drift direction. In FD1-HD, this is achieved using arrays of X-Arapuca[1] cells which are mounted onto each APA. An X-Arapuca cell is shown in Figure 3.17.

[1]An arapuca is a South American handcrafted trap used to catch birds and other small animals, used to

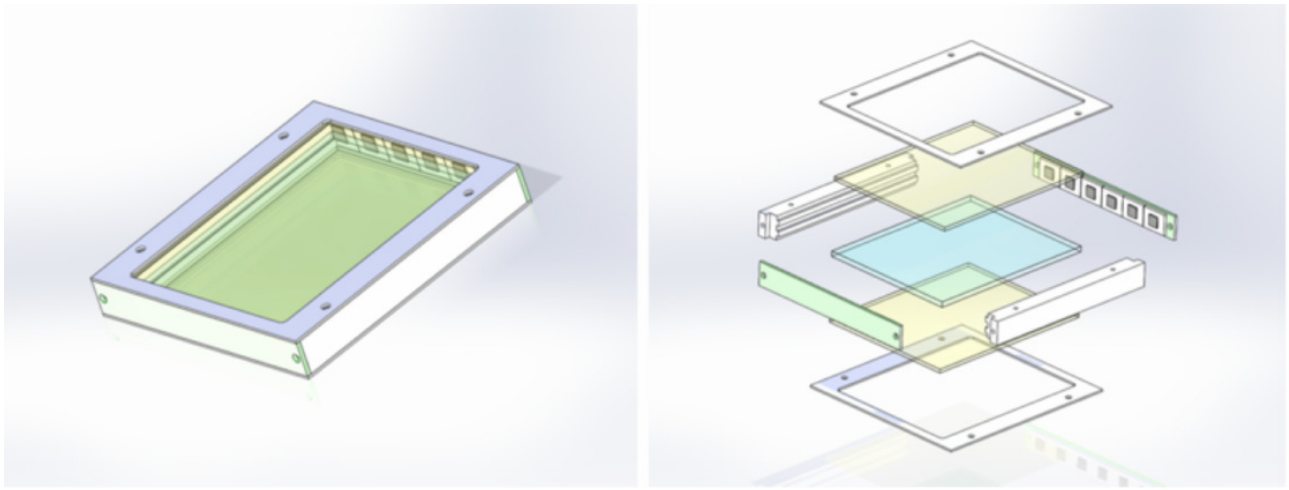

Figure 3.17: Left: an X-Arapuca cell. Right: an exploded view of the X-Arapuca cell, where the blue sheet is the wavelength-shifting plate and the yellow sheets are the dichroic filters [66].

Each cell consists of a wavelength-shifting (WLS) plate sandwiched between two dichroic filters, with silicon photomultipliers, SiPMs, at the edges. The 127 nm UV scintillation photons pass through the filter and are shifted to 430 nm photons in the visible spectrum. If the visible photons are emitted at greater than the critical angle, they are retained in the WLS plate and collected by the SiPMs. If they escape the WLS plate, they are reflected by the dichroic filter back into the plate, since the filter has an optical cutoff at 400 nm. The system effectively traps photons until they are eventually detected by the SiPMs.

3.4.2 Vertical Drift Module

The FD2-VD design emerged from the DUNE dual-phase (DP) design [87] which was tested at the European Laboratory for Particle Physics (CERN). Although DUNE moved away from the DP concept, the design of the cathode and the anode had its advantages. The FD2-VD design features a horizontal cathode place mid-height module, dividing the module into two equal drift volumes. The anodes, placed at the top and bottom of the module, are composed of perforated printed circuit boards (PCBs) rather than the wire planes in FD1-HD. Ionisation electrons will then drift vertically to one of the anode planes. FD2-VD will also use the same photon detection system as FD1-HD, with the placement re-optimised to the module geometry.

name this device which traps photons.

Figure 3.18 shows a cut-out schematic of the vertical drift concept.

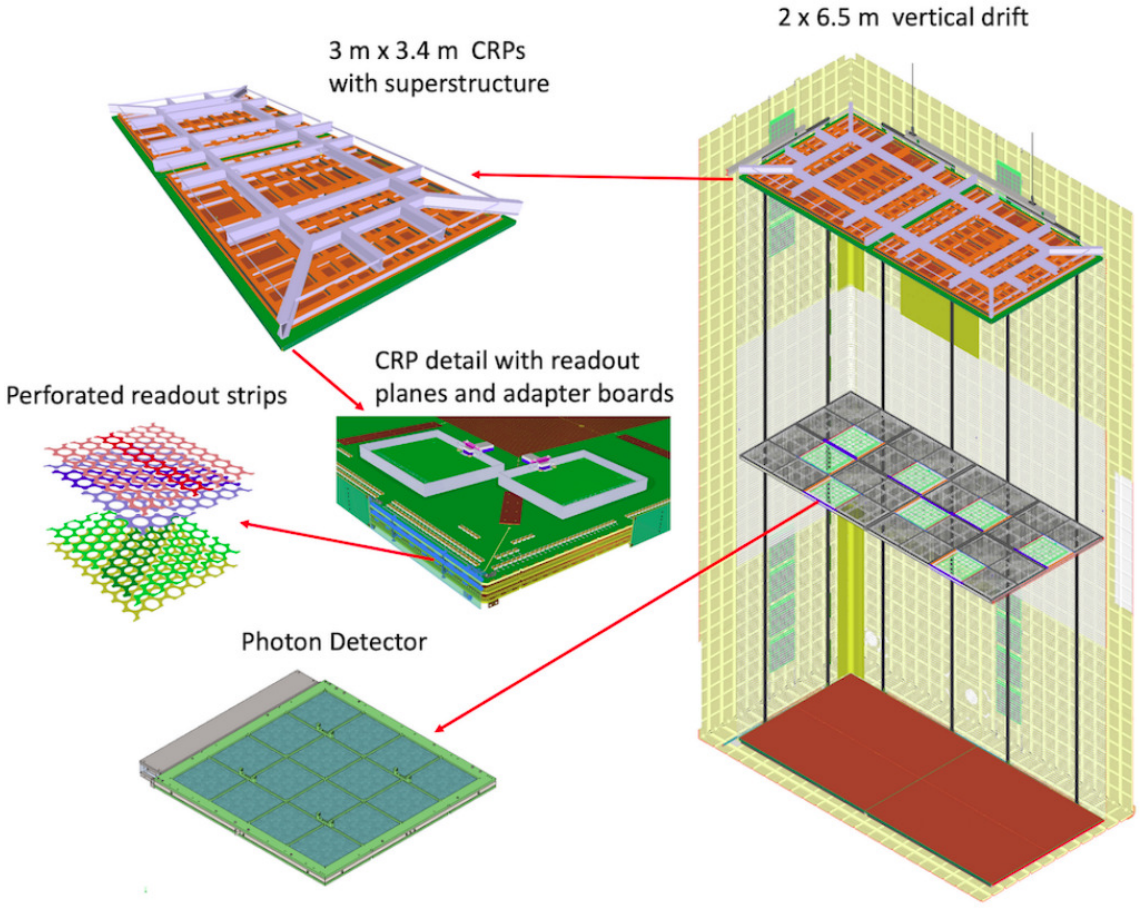

Figure 3.18: Cutaway diagram of FD2-VD design, using PCB-based readout. The cryostat wall is shown in yellow, PCB-based charge readout planes in brown at the top and bottom, the cathode in the middle in violet (with gaps for photon detectors), field cage centrally around the perimeter in white and photon detectors placed on the cathode and cryostat walls in vertical planes near the anodes [68].

Each FD2-VD anode plane consists of eighty charge readout planes (CRPs), which are composed of two double-sided perforated PCBs, connected on top of one another, and attached to a support structure. Each side of the PCBs performs the same function as the wire planes in FD1-HD. The innermost PCB has a copper guard plane facing the drift volume which acts as the shielding layer. The other side is the first induction layer, engraved with strips that act as the wires. The second PCB hosts the second induction layer on the side facing the first PCB and collection strips on the reverse side. Similarly to FD1-HD, the PCB strips are offset at an angle of 60°. To ensure complete transmission of electrons through the perforations, a roughly 1 kV voltage is applied across each PCB.

Due to the module geometry, the top and bottom drift electronics have different implementations. Since the top anode plane lies just under the LAr surface at the top of the cryostat, a combination of both cold and warm electronics will be used. The placement of the anode was chosen to exploit the total available vertical space of 13 m. The top drift electronics are also accessible through the cryostat roof. The bottom drift electronics will use the same components as FD1-HD, with the system completely submerged in LAr to maximise the signal-to-noise ratio.

Overall, the vertical drift design has a maximal active volume and extended drift distance which will help with event reconstruction. This design also features a simplified anode structure and increased modularity of the whole detector. Specifically, the modularity of the anode planes allows for cold testing without the need for large cryogenic vessels.

3.5 DUNE Physics Potential

DUNE has a wide physics program but is primarily a long-baseline neutrino experiment that aims to precisely measure the neutrino oscillation parameters. This includes the mixing angles, the mass splitting, the mass hierarchy and the CP-violating phase factor, δ_{CP}. Long-baseline experiments typically have good sensitivity to θ_{23}, Δm^2_{32}, θ_{13} and δ_{CP}. Sensitivity to θ_{12} and Δm^2_{12} is achieved by solar neutrino experiments which are at lower energies and have much longer baselines. Whether δ_{CP} is non-zero and the mass hierarchy are both unknown and of high interest in the scientific community. DUNE has great sensitivity to the mass hierarchy due to its 1300 km baseline. This can be seen from the expression for electron appearance probability, $P(\nu_\mu \rightarrow \nu_e)$, through matter with constant density which, to first order, is [88]

$$
\begin{aligned}
P(\nu_\mu \to \nu_e) \approx \sin^2(\theta_{23})\sin^2(2\theta_{13})\frac{\sin^2(\Delta_{31}-aL)}{(\Delta_{31}-aL)^2}\Delta_{31}^2 \\
+ \sin(2\theta_{23})\sin(2\theta_{13})\sin(2\theta_{12})\frac{\sin(\Delta_{31}-aL)}{(\Delta_{31}-aL)}\Delta_{31}\frac{\sin(aL)}{(aL)}\Delta_{21}\cos(\Delta_{31}+\delta_{CP}) \\
+ \cos^2(\theta_{23})\sin^2(2\theta_{12})\frac{\sin^2(aL)}{(aL)^2}\Delta_{21}^2
\end{aligned}
\tag{3.10}
$$

where $\Delta_{ab} = \Delta m^2_{ab} L/4E$, $a = G_F N_e/\sqrt{2}$, G_F is the Fermi constant, N_e is electron number density in the Earth, L is the baseline and E is the neutrino energy. For the CP conjugate process (i.e. $\bar{\nu}_\mu \to \bar{\nu}_e$), $a \to -a$ and $\delta_{CP} \to -\delta_{CP}$. It is now noticeable that the sensitivity to mass hierarchy comes from the matter effect since, without the matter effect, the leading term would be unchanged upon $\Delta_{31} \to -\Delta_{31}$. The neutrino-antineutrino asymmetry introduced by the matter effect is due to the significant imbalance in the presence of electrons and positrons in the Earth's crust. Due to the 1300 km baseline, at the energy of maximum oscillation, the asymmetry caused by the matter effect is greater than the largest possible asymmetry caused by CP violation. This allows the hierarchy to be determined by DUNE for any value of δ_{CP}.

Figure 3.19 shows $P(\nu_\mu \to \nu_e)$ as a function of energy at the DUNE baseline for different values of δ_{CP}. One can see how changing δ_{CP} changes both the phase and the amplitude of the oscillation spectrum. At lower energies, the effect on the amplitude is greater. Since DUNE is a broadband experiment, it is possible to map out the shape of the observed spectrum.

In addition, the value θ_{23} is known to be near maximal-mixing, i.e. $\theta_{23} \approx \pi/4$. However, given it's not exactly maximal, the octant of θ_{23} is unknown, i.e. $\theta_{23} < \pi/4$ or $\theta_{23} > \pi/4$. Knowing the true value of θ_{23} is important for a complete understanding of the mechanism behind neutrino mixing and masses. DUNE will also measure θ_{13}. Currently, the world-leading measurements of θ_{13} are held by reactor neutrino experiments. With a long enough exposure, DUNE will exceed the precision of these measurements.

DUNE also aims to probe physics outside of the long-baseline sector. In the event of a rare core-collapse supernova, DUNE aims to be ready to observe the incoming neutrino flux,

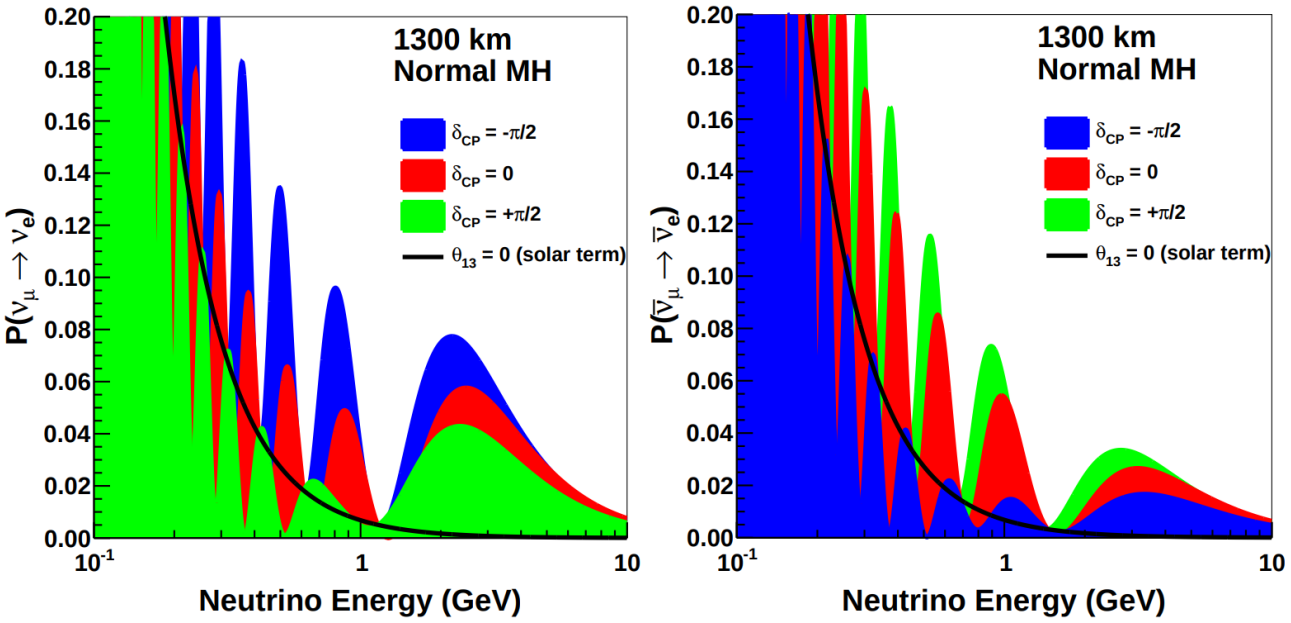

Figure 3.19: The electron (left) and anti-electron (right) neutrino appearance probability as a function of energy, at a baseline of 1300 km and under normal ordering hypothesis. Spectra are shown for $\delta_{CP} = -\pi/2$ (blue, 0 (red), $\pi/2$ (green). The black line shows the oscillation probability if θ_{13} were equal to 0 [64].

along with the rest of the neutrino community. An example of one of these events is the famous Supernova 1987A, which occurred roughly 50 kiloparsecs (kpc) from Earth. Valuable information on the mechanism of a core-collapse was uncovered from the tens of recorded neutrino events [89, 90]. It is estimated that a core-collapse supernova within a few hundred kpcs from Earth will occur every few decades. Hence, there is a reasonable chance that one will occur during the lifetime of DUNE. For this reason, it is key that at least one DUNE FD module is online at all times. This also adds further requirements to the data acquisition system, such as advanced triggers and a large storage buffer to ensure that the event is not missed.

Due to the low-energy nature of supernova neutrinos, order tens of MeV, CC interactions are prohibited for ν_μ, $\bar{\nu}_\mu$, ν_τ, $\bar{\nu}_\tau$. Typically, supernova $\bar{\nu}_e$ have been detected via inverse beta decay. However, since DUNE is using an argon target, it will have sensitivity to supernova ν_e via $\nu_e + {}^{40}\mathrm{Ar} \rightarrow e^- + {}^{40}\mathrm{K}^*$. The photon emission from the decay of the excited potassium produces a unique signature.

DUNE will also look for proton decays, despite current observations pointing towards their inherently stable nature. If this process were to occur, it would violate the baryon number, which is implicitly conserved in the Standard Model Lagrangian. However, there is no symmetry of nature which requires its conservation. DUNE will look for proton decay via the process $p \rightarrow K^{+}\bar{\nu}$. As well as proton decay, there is a wealth of Beyond-Standard-Model physics that DUNE can probe. Among these are searches for active-sterile neutrino mixing, non-unitarity of the PMNS matrix, nonstandard interactions, dark matter candidates, and heavy neutral leptons.

Chapter 4

Bayesian Statistics and Markov Chain Monte Carlo

This thesis presents an analysis using a Bayesian statistical treatment. In a Bayesian analysis, the ultimate goal is to produce the posterior distribution, which is a function of all the parameters in the model. Although there are only six oscillation parameters that are of interest, often large numbers of systematic parameters are considered in the model which predicts the data from a long-baseline experiment. Analytically constructing the posterior distribution is generally not possible for very high dimensions. In such cases, sampling methods are used to produce approximate distributions. With a large number of samples, this approximate distribution can be used to extract information about the exact posterior distribution.

This analysis uses the Markov Chain Monte Carlo (MCMC) sampling method to explore the parameter space and extract estimations of oscillation parameter values from data. MCMC is a family of Monte Carlo methods, which intelligently explores the parameter space by utilising the convergence properties of a Markov chain.

This chapter will briefly explain the theory behind Bayesian inference in Section 4.1. Monte Carlo methods in general are described in Section 4.2 and MCMC is described in Section 4.3. Section 4.4 explains how information can be extracted from the posterior distribution.

4.1 Bayesian Inference

In Bayesian statistics, probability is interpreted as the degree of belief in an event or hypothesis. This differs from the frequentist interpretation of statistics, in which probability is the relative frequency of an event or hypothesis upon "infinite" repetitions. This effective difference can be explained using the classical flip of a fair coin. A frequentist would not assign any probability to the flip of a single coin but might state that as the number of flips increases, the proportion of flips that land on tails tends to one-half. A Bayesian would assign a probability, i.e. their degree of belief, of the coin landing on either side as "fifty-fifty".

Since in the Bayesian view, neither the observables nor the model parameters are fixed, one must find a complete probability model of both to make inferences about model parameters. This is known as Bayesian inference. Let's say we're trying to make inferences about a set of model parameters, $\vec{\theta}$, from some observed data, D. The aim is to find the probability distribution of $\vec{\theta}$ given D, written as $P(\vec{\theta}|D)$. This is known as the posterior probability distribution. To obtain this, we need to first ask the opposite question: what is the likelihood of observing D given $\vec{\theta}$, written as $P(D|\vec{\theta})$? This is known as the likelihood. These are related by Bayes' theorem

$$P(\vec{\theta}|D) = \frac{P(D|\vec{\theta})P(\vec{\theta})}{\int P(D|\vec{\theta})P(\vec{\theta})d\vec{\theta}} \tag{4.1}$$

where $P(\vec{\theta})$ is our prior knowledge of the model parameters. The numerator is often called the joint probability distribution. The denominator of Eq 4.1 is the integral of the joint probability distribution over the entire parameter space and is often very difficult to evaluate. However, since this analysis aims to infer the probability distribution of the model parameters, it is sufficient to construct a distribution that is proportional to the real posterior:

$$P(\vec{\theta}|D) \propto P(D|\vec{\theta})P(\vec{\theta}) \tag{4.2}$$

In this analysis, the observed data is the binned neutrino event distribution. The likelihood is calculated by comparing Monte Carlo prediction for a given set of model parameters, $\lambda(\vec{\theta})$, with the data, n, in a given bin and calculating the Poisson likelihood. Prior knowledge of parameters is included by modelling them as Gaussian variables with central values $\vec{\mu}$. A covariance matrix, $\mathbf{V}$, is constructed which contains all prior uncertainties and correlations of the model parameters. The posterior distribution can then be expressed as

$$P(\vec{\theta}|D) \propto \prod \mathcal{L}_{\text{Total}} = \underbrace{\prod_{\text{Bins}} \mathcal{L}_{\text{Poisson}}}_{P(D|\vec{\theta})} \times \underbrace{\prod_{\theta} \mathcal{L}_{\text{Gaussian}}}_{P(\vec{\theta})} \tag{4.3}$$

It is often far more convenient to express this in the negative logarithmic space. This reduces the volatility in the likelihood values and computationally addition is cheaper than multiplication. The negative allows us to minimize rather than maximise which is far more common in statistical packages. We can then write

$$-\log \mathcal{L}_{\text{Total}} = \sum_{\text{Bins}} \left[\lambda(\vec{\theta}) - n + n log \frac{n}{\lambda(\vec{\theta})} \right] + \sum_{\theta} \frac{1}{2} \left[(\theta_i - \mu_i) \mathbf{V}_{ij}^{-1} (\theta_j - \mu_j) \right] \tag{4.4}$$

Although Eq 4.4 allows us to calculate the likelihood for a given set of $\vec{\theta}$, it is generally not possible to analytically describe the likelihood in the full parameter space. Instead, it is common practice to sample from the likelihood using Monte Carlo methods.

4.2 Monte Carlo Methods

Monte Carlo methods are a widely used class of computational algorithms, named after the famous Monte Carlo casino[1]. The fundamental concept behind all Monte Carlo methods lies in the approximation of distributions or functions through sampling, thereby enabling the extraction of their true properties. Monte Carlo methods differ from one another based on their chosen sampling algorithm.

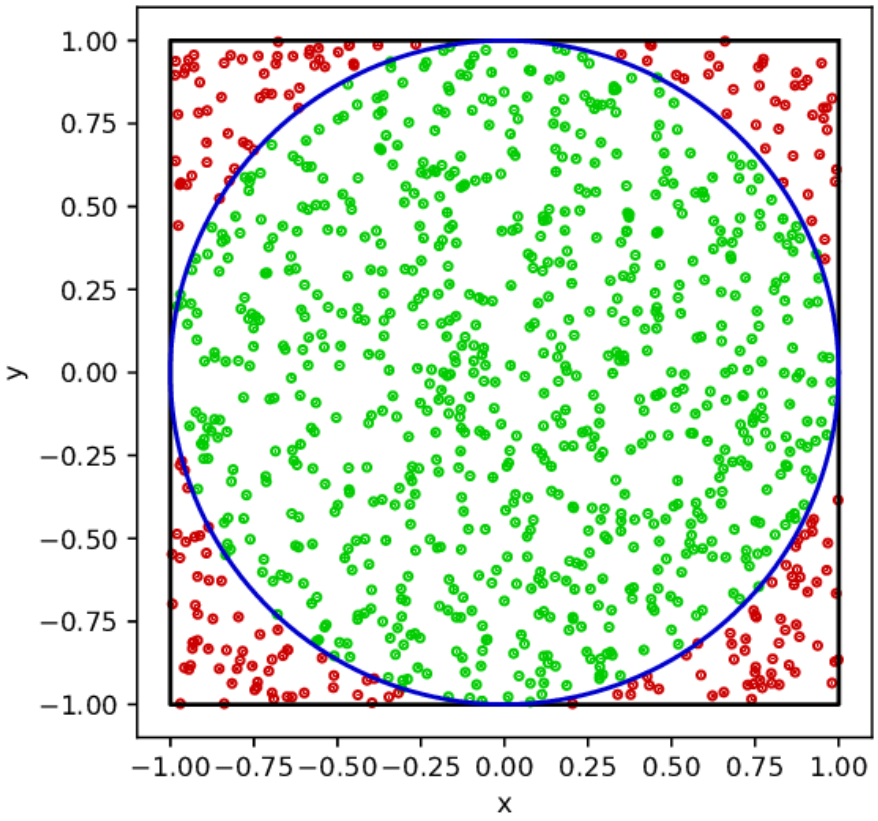

Figure 4.1: A simple demonstration of Monte Carlo used to approximate the area of a circle. Figure taken from [91]

A classic example of the application of Monte Carlo is the approximation of the area of a circle, shown in Figure 4.1. We can randomly scatter points across the square and count the fraction of points which lie inside the circle. The area of the circle is then equal to this fraction multiplied by the area of the square. The more points we scatter, the more accurate

[1]Stanislaw Ulam, the primary developer of Monte Carlo methods, named the technique after the casino of which his uncle was a frequent gambler.

this approximation becomes. For this scenario, this may seem trivial, but now imagine that rather than a circle, we are trying to evaluate the area of shape that has no analytical solution. This same procedure can be used to numerically approximate any integral and demonstrates the power of Monte Carlo methods.

The integration method outlined here randomly distributes points across the square. While this approach may appear efficient in this example, it rapidly becomes impractical as the problem's dimensionality increases. Lots of time might be spent scattering points outside of the regions of interest. Hence, it quickly becomes necessary to choose a more sophisticated sampling method than pure randomness.

4.3 Markov Chain Monte Carlo

Markov Chain Monte Carlo (MCMC) is a specific branch of Monte Carlo algorithms which uses a Markov chain to sample a distribution. MCMC is far more efficient at exploring high-dimensional distributions than the random sampling demonstrated in Figure 4.1.

A Markov chain is a stochastic sequence of events in which the probability of a given event is only dependent on the previous event. If we define $\mathbf{x}$ as a sequence of positions in a Markov chain, the transition to any position x_i, the i^{th} entry in $\mathbf{x}$, depends only on the previous position x_{i-1}. For a multi-dimensional parameter space, each position in the Markov chain is a vector, $\vec{x_i}$, with the same dimensions. Hence, as opposed to random sampling, a Markov chain constructs a "semi-random" walk through the parameter space. By definition, steps in a Markov chain are not independent samples, since they depend fully on the previous step plus random variation.

Convergence of a Markov chain occurs when the chain 'forgets' about its initial distribution and enters a stationary distribution. MCMC methods leverage this property by constructing a Markov chain which has the desired posterior distribution as its stationary distribution. If this is achieved, the chain will contain a set of discrete points whose density is proportional to that of the posterior distribution. Constructing a Markov chain that converges is often not trivial

and the theory describing convergence is extensive [92, 93]. In summary, to obtain convergence, it is necessary and sufficient that a Markov chain satisfies the following conditions:

- **Irreducibility:** From any given state in the parameter space, the probability of reaching all other states is non-zero.
- **Recurrence:** Once the stationary distribution has been found, all subsequent steps must sample from the same stationary distribution
- **Aperiodicity:** The sequence of steps must not be periodic

Provided these conditions are met, a Markov chain will eventually converge on a stationary distribution. It is now necessary to construct a Markov chain whose stationary distribution is equivalent to the desired posterior distribution. Although several algorithms exist, this analysis uses a method known as the Metropolis-Hastings algorithm [94, 95]

4.3.1 The Metropolis-Hastings Algorithm

The Metropolis-Hastings algorithm is a well-established method of constructing a Markov chain. It was first developed by N. Metropolis et al. in 1953 [94] and later generalised by W. K. Hastings [95]. It has become a cornerstone algorithm in the field of MCMC and Bayesian computational methods due to its applicability and simple implementation.

Metropolis-Hastings algorithm is executed at every step in the Markov chain. Remember that $\vec{x}_i$ is defined as a discrete point in the Markov chain, which corresponds to a position (or "step") in the parameter space $\vec{\theta}$. The algorithm is the following (starting at the initial point in the chain):

1. **Initialisation:** The chain starts with a random point $\vec{x}_0$ in the parameter space. The test-statistic, generally and in this analysis is the negative log-likelihood defined in Eq 4.4, is evaluated at $\vec{x}_0$.
2. **Proposal step:** A candidate step $\vec{y}_0$ proposed using a proposal function. Importantly, the proposal function is only dependent on the previous step as required for Markov chains.

Any proposal function can be used, but commonly (and in this analysis) a Gaussian distribution centred around the current step is used

$$\vec{y}_n \;=\; \vec{x}_n \;+\; \vec{c}\,\vec{z}_n \tag{4.5}$$

$$\vec{z}_n \;\sim\; \mathcal{N}(\vec{x}_n, \mathbf{V}) \tag{4.6}$$

The width of the Gaussian distribution is determined by the prior uncertainty and correlations, contained in the covariance matrix $\mathbf{V}$ and a tuneable scaling parameter, $\vec{c}$. Tuning the proposal function can be vital in constructing a converging Markov chain and is discussed in more detail in Section 4.3.2.

3. **Acceptance Probability:** The test-statistic is evaluated at the candidate step c. The acceptance probability of the proposed step is calculated by taking the ratio of the likelihood

$$\alpha \;=\; \min\left[\,1\,,\; \frac{\mathcal{L}_{\text{proposed}}}{\mathcal{L}_{\text{current}}}\right] \;=\; \min\left[\,1\,,\; \frac{\mathcal{L}(\vec{y}_0)}{\mathcal{L}(\vec{x}_0)}\right] \tag{4.7}$$

where the min function returns the smallest of the values.

4. **Accept or Reject:** A random number, r, is generated from a uniform distribution between 0 and 1. This is compared with the acceptance probability, α, and the next step is defined to be

$$\vec{x}_1 = \begin{cases} \vec{y}_0 & \alpha \geq r \\ \vec{x}_0 & \alpha < r \end{cases} \tag{4.8}$$

Due to the min function in Eq 4.7, if the likelihood of the proposed step is greater than the current step, the proposed step is always accepted.

5. **Repeat:** Items 2 - 4 then repeat for each step in the Markov chain.

The acceptance probability in the Metropolis-Hastings algorithm is constructed such that regions in the parameter space are explored in the correct proportion. If point A has half the likelihood than point B, it will only transition half of the time. An added benefit is that, since steps with lower probability are still accepted, the chain avoids getting stuck in local minima.

4.3.2 Proposal function

As proven in [92], any form of proposal function can be used in the Metropolis-Hastings algorithm and the Markov chain will still converge to the same stationary distribution. However, the choice of proposal function has a significant effect on the speed of convergence. It is also desirable to have a Markov chain that explores the full parameter space quickly once it has converged. In this case, the chain is said to 'mix' well. As mentioned in Section 4.3.1, the choice of proposal function in this analysis (and typically used in Random Walk Metropolis-Hastings) is a Gaussian distribution centred on the position of the current step, as stated in Eq 4.5 and Eq 4.6. The task is then to choose values for c, often labelled the "step size", which optimises the Markov chain convergence.

Figure 4.2 shows 1000 steps from three Markov chains using the Metropolis-Hastings algorithm with the same stationary distribution. The sole difference between the chains is the step size in the proposal function. We see that if the step size is too small, as in Figure 4.2a, almost all of the proposed steps are accepted, since $\mathcal{L}(\vec{y}_n) \approx \mathcal{L}(\vec{x}_n)$. However, since the steps are close to one another, convergence may be slow and, once convergence is reached, a large number of steps are required to obtain a good sample of the entire posterior distribution. If the step size is too large, as in Figure 4.2b, the proposed steps are often in regions of low likelihood i.e. $\mathcal{L}(\vec{y}_n) << \mathcal{L}(\vec{x}_n)$. This results in a very low acceptance rate, so the chain rarely moves at all. Similarly to Figure 4.2a, a large number of steps would be required to obtain a reliable sample.

Figure 4.2c shows a resulting Markov chain when the chosen step size is close to optimal. The chain mixes quickly, exploring the full parameter space regularly in a small number of steps.

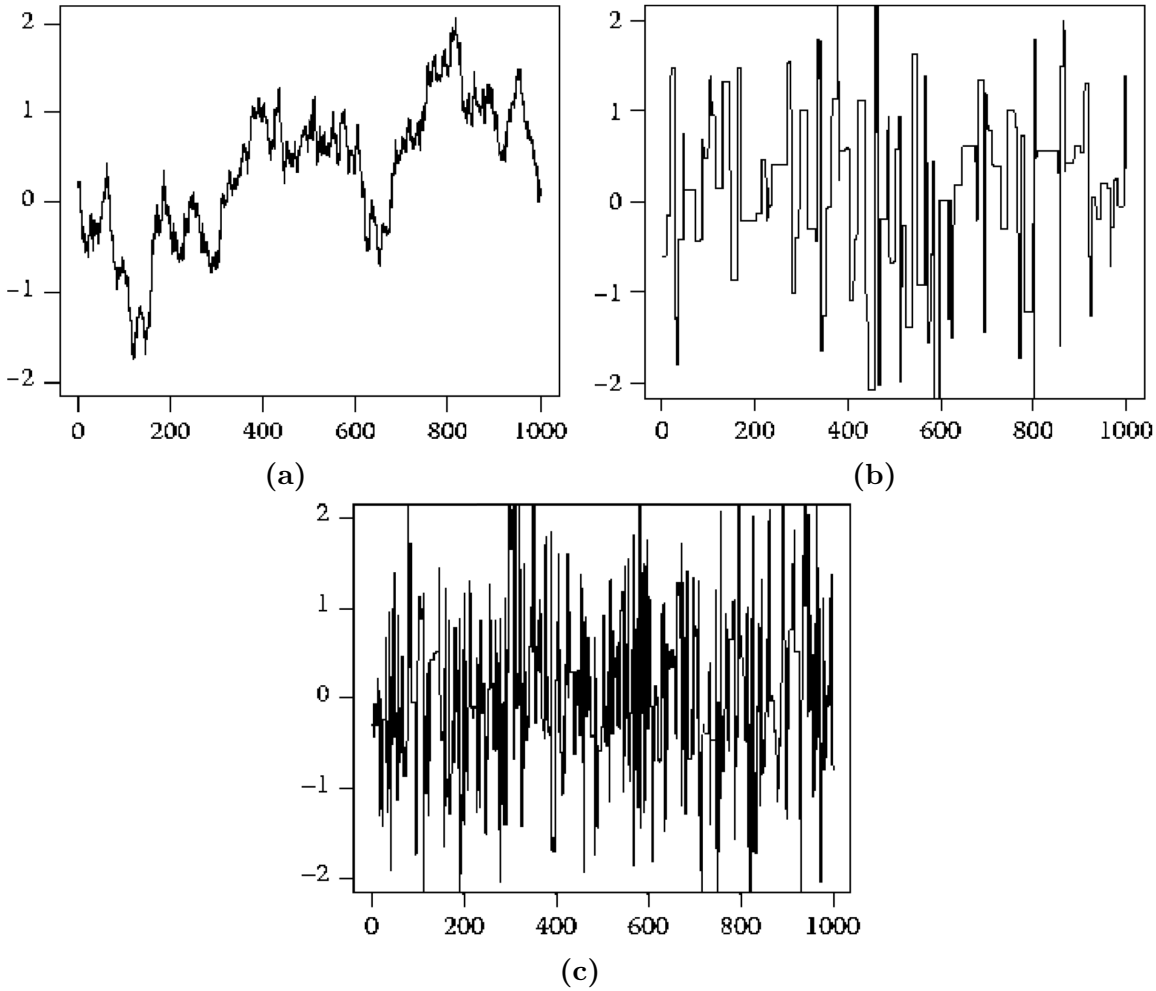

Figure 4.2: MCMC chains with 1000 steps using Metropolis-Hastings algorithm with a Gaussian proposal function where the scaling parameter is (a) too small; (b) too big, (c) optimal. Figures taken from [93]

We can see that the choice of proposal function and step size significantly affects the performance of a Markov chain. Albeit crucial, step size tuning is infamously challenging and can be tedious upon multiple iterations. Early studies on how to determine the optimal scaling were performed by Roberts et al. [96]. A simplistic target distribution, $\pi(\vec{x})$, was considered with the form

$$\pi(\vec{x}) \;=\; \prod_{i=1}^{d} f(x_i) \tag{4.9}$$

where f is a one-dimensional distribution and d is the number of dimensions (i.e. parameters). This is an unrealistic scenario since it suggests that sampling from each of the independent distributions f is equivalent to sampling from the distribution π. Nonetheless, under this assumption and using the Random Walk Metropolis-Hastings, it was shown that as $\mathrm{d} \to \infty$, the acceptance rate for optimal mixing was ≈ 0.234. Of course, one might expect different behaviour for a non-asymptotic scenario. However, numerical studies conducted in [97] found that the infinite-dimensional assumption approximates finite-dimensional cases well. That being said, when the target distribution π takes a more general or complex form, the exact results may no longer hold. Each target distribution will require some unique fine-tuning. However, the general principle that the step size should be neither too small nor too large applies to all MCMC applications.

The final complexity not yet considered is the case where parameters in the Markov chain contain prior correlations. In the context of a neutrino oscillation analysis, systematic parameters that describe the neutrino flux for neutrinos and antineutrinos are likely to be highly correlated. This is because the underlying physics parameters that affect the neutrino flux, i.e. horn current uncertainties, are often very similar for both neutrinos and antineutrinos. These correlations are included in the covariance matrix, $\mathbf{V}$, in Eq 4.4, and hence have an impact on the likelihood. For example, if a proposed parameter step predicts a high neutrino flux and a low antineutrino flux, this would have a low prior probability $P(\vec{\theta})$ and hence a low likelihood. This results in a low acceptance probability. It is therefore necessary to account for these correlations in the proposal function to prevent wasted steps being proposed in regions of low $P(\vec{\theta})$.

In this analysis, Cholesky Decomposition is used to turn an uncorrelated proposal step into

the correlated basis [98]. Cholesky Decomposition acts on a positive definite hermitian matrix and decomposes it into a product of a lower triangle matrix and its conjugate transpose. Given our covariance matrix $\mathbf{V}$, the decomposition is of the form

$$\mathbf{V} = \mathbf{L}\mathbf{L}^* \tag{4.10}$$

where $\mathbf{L}$ is a lower triangular matrix. This allows us to convert an uncorrelated proposal into a correlated proposal by applying $\mathbf{L}$, i.e.

$$\vec{\theta} = \mathbf{L}\vec{\theta}' \tag{4.11}$$

where $\vec{\theta}$ is a correlated step and $\vec{\theta}'$ is an uncorrelated step. Geometrically, Cholesky Decomposition can be visualised as the transformation from some basis axis to the conjugate axis of an ellipse, as shown in Figure 4.3. Randomly distributed points in a circle can be converted to the equivalent point in an ellipse defined by the correlation between the axes. Cholesky Decomposition is commonly used in numerical approximation methods such as Monte Carlo simulation.

4.3.3 Diagnostics

As well as ensuring the acceptance probability is within a reasonable range, other diagnostic tools are important for verifying the Markov chain is in an optimal state. As mentioned previously, the ideal tune of the step size will result in a Markov chain that explores the full parameter space quickly. This property can be checked quantitatively by calculating the correlation between different steps. By definition of a Markov chain, every step $\vec{\theta}_n$ is highly correlated with the previous step $\vec{\theta}_{n-1}$. As a result, $\vec{\theta}_n$ is also highly correlated with $\vec{\theta}_{n-2}$,

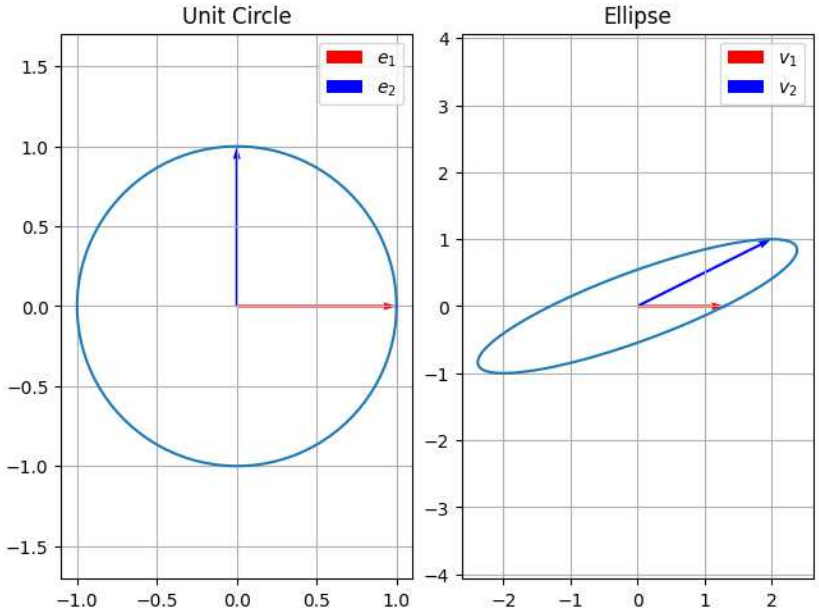

Figure 4.3: Visual representation of Cholesky Decomposition as a linear transformation applied to a unit circle. Figure taken from [99]

since $\vec{\theta}_{n-1}$ is highly correlated with $\vec{\theta}_{n-2}$. The further back in the chain you go, generally, the correlation decreases. Ideally, steps become uncorrelated quickly, as the chain will then contain more independent samples of the target distribution. The quantity describing the correlation between successive steps in a Markov chain is called the auto-correlation. For some parameter P_n at step n, the auto-correlation a_k after k successive steps (also referred to as the 'lag') is calculated by

$$a_k = \frac{\sum_{n=1}^{N-k} (P_n - \hat{P})\,(P_{n+k} - \hat{P})}{\sum_{n=1}^{N-k} (P_n - \hat{P})^2} \tag{4.12}$$

where N is the number of steps in the Markov chain, $\hat{P}$ is the mean of the parameter. Often this is calculated iteratively for increasing values of the lag k, showing how long it takes the parameter to become uncorrelated. This can be seen in Figure 4.4. The speed at which steps become uncorrelated depends significantly on the choice of step size. If the step size is too small, it follows that the number of steps required to reduce the correlation increases. Similarly,

if the step size is too large, the chain often gets stuck in regions of the parameter space. It is also worth noting that the choice of a reasonable auto-correlation is not fixed for all Markov chains. Different target distributions and prior correlations will have different Markov chain behaviours. In this analysis, based on similar neutrino oscillation analyses, an auto-correlation below 0.2 after 15,000 steps is aimed for across all parameters.

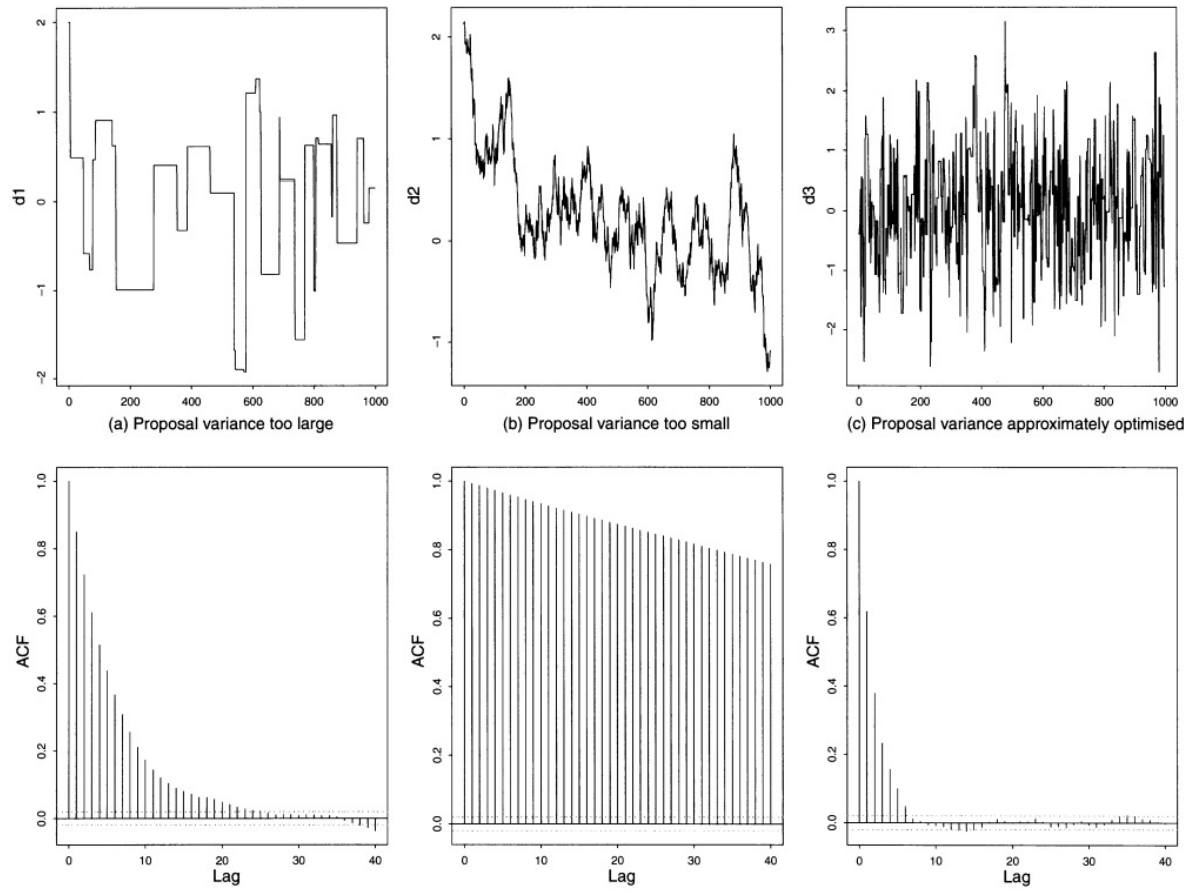

Figure 4.4: Markov chain using Metropolis-Algorithm with (a) too-large step size, (b) too-small step size (middle) and (c) appropriate step size. Trace plots (top) and autocorrelation plots (below) are shown for each case. Figure taken from [97]

In addition to diagnosing the efficiency of a Markov chain, the auto-correlation indicates the number of steps required to make levels of statistical statements. Intuitively, the more samples are made from a distribution, the better the shape of the distribution is understood. However, this relies on the samples being independent. Since each step in a Markov chain is not fully independent, the auto-correlation must be accounted for to calculate the number of independent steps. Simply, if a Markov chain has N steps and a lag of k provides reasonable independence,

the Markov chain contains at least N/k independent steps. Now let's imagine the analysis aims to make a 3σ statement about an oscillation parameter. For a normally distributed variable, a 3σ credible interval should contain 99.73% of the probability distribution, and hence 99.73% of the Markov chain steps. Let's also say we want at least 30 steps outside of our credible region, such that we have a good idea of where the boundary should be. Finally, let's say we consider steps with less than 0.2 auto-correlation to be independent, and that occurs after a lag of 15,000 steps. We can now calculate the total number of steps our Markov chain would need to satisfy these criteria, by doing

$$N = \frac{\mathrm{I}_{\mathrm{tail}} \times \mathrm{k}}{(1 - \delta)} = \frac{30 \times 15000}{1 - 0.9973} \approx 160\,\text{million} \tag{4.13}$$

where $\mathrm{I}_{\mathrm{tail}}$ is the number of independent steps outside the credible region δ. However, the number of steps calculated via Eq 4.13 is only an estimate. Often the stability of the contours produced from the Markov chain is checked to ensure enough steps have been collected to make statistical statements.

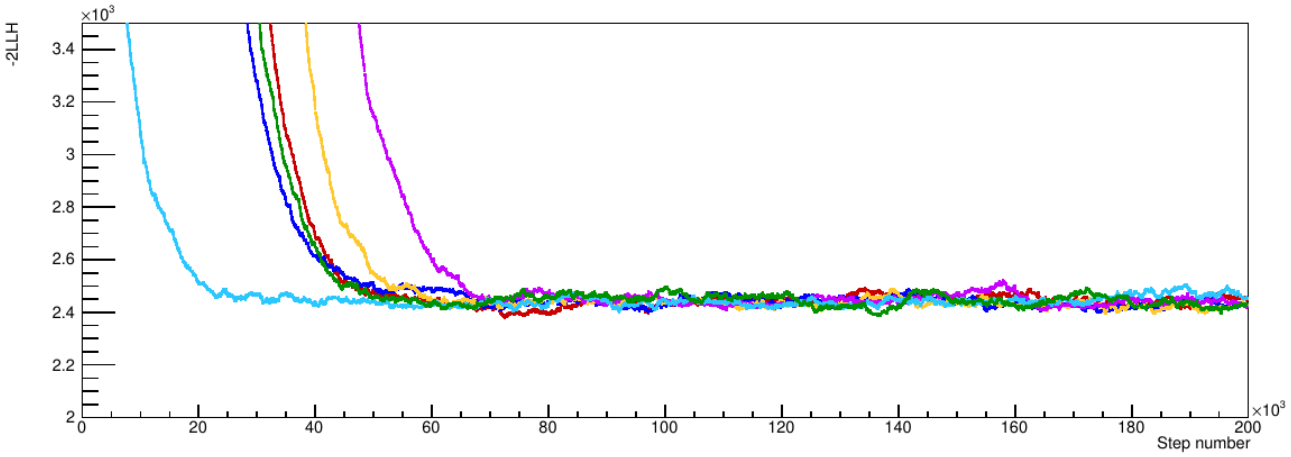

Figure 4.5: The negative loglikelihood against step number for multiple Markov chains sampling the same distribution. The negative loglikelihood begins large since the chain is initiated in a random unlikely region. All of the Markov chains reach the stationary distribution by 100,000 steps and hence 0 - 100,000 is a valid burn-in range. Figure taken from [100]

As mentioned previously, the Markov chain must converge to a stationary distribution that is equivalent to the desired posterior distribution. When a Markov chain is in a stationary state, it has 'forgotten' about its initial state. Let's define $P^n(\vec{x}_n|\vec{x}_0)$ to be the probability distribution of some step $\vec{x}_n$ given some starting point $\vec{x}_0$. A chain is considered to have converged to a stationary distribution when $P^n(\vec{x}_n|\vec{x}_0)$ is independent of the step number n. However, since this does not happen immediately, the steps in the Markov chain before convergence is achieved can bias the sample towards the starting position of the chain. It is therefore necessary to remove these steps before extracting any parameter values (since these steps are not sampling from the desired posterior distribution). These steps are often called the "burn-in". The burn-in region can be identified by plotting the loglikelihood against step number, as shown in Figure 4.5. Once several Markov chains converge to the same stable loglikelihood value, this indicates that the chains are exploring the same stationary distribution. In Figure 4.5, this occurs within the first 100,000 steps for all of the chains. Hence, the first 100,000 steps of each chain would be considered burn-in and should be cut off.

4.4 Inferring from the Posterior

At this stage, we've constructed a Markov chain with a stationary distribution that is equivalent to the target distribution we want to explore. This means we have produced a posterior distribution from which we can extract information about the distribution of any given parameter. From a Bayesian perspective, there is one joint probability distribution for all the parameters in the model and the data, and hence this full posterior distribution is the 'solution'. However, if the posterior distribution has hundreds of dimensions, as in this analysis, it's impossible to draw any conclusions about the parameters of interest. So, we need a way of removing the parameters we aren't interested in.

4.4.1 Marginalisation

In a Bayesian view of statistics, all parameters in a model are fundamentally random variables. However, generally, we are interested in inferring information about a small subset of the parameters. These parameters generally have physical importance to the field being studied. In a long-baseline neutrino oscillation analysis, we want to present best-fit values and uncertainties for the neutrino oscillation parameters that the experiment is sensitive to, i.e. $\sin^2\theta_{23}$, $\sin^2\theta_{13}$, Δm^2_{32}, and δ_{CP}. However, when extracting information about these parameters from the full posterior, the uncertainty introduced by the remaining parameters must be taken into account properly.

The parameters of the full posterior that have no physical importance and hence need to be removed are called 'nuisance parameters'. For example, parameters associated with the energy resolution of the detector are not the critical physics parameters that the community wants to read about in a paper abstract, even though they affect the observations of the experiment. Often, the vast majority of parameters included in the full posterior are nuisance parameters, with only a select few being parameters of interest. In this analysis, the flux, cross-section and detector model parameters are all nuisance parameters, as well as two neutrino oscillation parameters ($\sin^2\theta_{12}$ and Δm^2_{21}) that long-baseline experiments are not sensitive to.

Bayesian statistical methods contain a simple solution for removing nuisance parameters from the full posterior distribution, called marginalisation. Depending on whether the distribution is discrete or continuous, simply one integrates or sums the posterior over the nuisance parameters leaving behind a posterior distribution as a function of only the parameters of interest. Given a full posterior distribution $P(\vec{\theta}|D)$, over a set of parameters $\vec{\theta}$ which are split into parameters of interest $\vec{\theta}_{poi}$ and $\vec{\theta}_n$, the marginal posterior over $\vec{\theta}_{poi}$ can be written as

$$P(\vec{\theta}_{poi}|D) \;=\; \int P(\vec{\theta}_{poi}, \vec{\theta}_n|D)\, \mathrm{d}\vec{\theta}_n \tag{4.14}$$

Marginalisation is an intuitive solution to the problem. To demonstrate this, let's consider the following probability distribution of two random variables, X and Y, which can only take discrete values 1 and 2:

$$P(X,Y) \;=\; \begin{cases} 0.2 & X=1, Y=1 \\ 0.3 & X=2, Y=1 \\ 0.1 & X=1, Y=2 \\ 0.4 & X=2, Y=2 \end{cases} \tag{4.15}$$

To find the marginal probability of X, $P(X)$, it should be obvious that we simply sum over the two possible Y values for each X value:

$$P(X) \;=\; \sum_{y} P(X,y) \;=\; \begin{cases} 0.3 & X=1 \\ 0.7 & X=2 \end{cases} \tag{4.16}$$

Although simple, evaluating Eq 4.14 can be tricky if the number of nuisance parameters is large, since each nuisance parameter adds a dimension to the integral. However, when using MCMC, this integral becomes trivial since the posterior probability is a discrete distribution of steps. Much like the simplified example in Eq 4.16, all that's necessary is to add up the steps in each bin of the dimensions of the parameters of interest ignoring the nuisance parameters. This highlights the power of MCMC to evaluate high-dimensional integrals that cannot be solved analytically.

Although marginalisation is a correct method for removing nuisance parameters, it can introduce features which may not seem to agree with some of the intuitive features (like the maximum probability point) of the full posterior. This often occurs when the nuisance param-

eters are correlated with the parameters of interest, or if they have non-Gaussian distributions. After marginalizing, this can cause a shift in the distribution, moving the region of the highest probability away from where it might be expected.

This effect can become more apparent when compared with a commonly used alternative method called profiling. When profiling a distribution to remove nuisance parameters, the value of the highest probability for the nuisance parameter is selected i.e. $P(\vec{\theta}_{poi}|D) = P(\vec{\theta}_{poi}, \vec{h}_n|D)$ where $\vec{h}_n$ are the highest probability points of the nuisance parameters $\vec{\theta}_n$. The differences in these methods are illustrated in Figure 4.6, which shows a 2D distribution of a signal parameter and a background parameter. The 1D distribution of the signal parameter is found using both marginalisation and profiling.

In Figure 4.6, the profiled distribution seems more intuitive, since the distribution peaks in the same place as the 2D distribution. The marginalised distribution seems to be shifted to the right, such that its peak and the peak of the 2D distribution do not align. When considering the shape of the 2D distribution, the reason for this effect becomes clear. The profiled distribution only considers the variation in the best b value as a function of s. Marginalisation considers both the variation in the best b value as a function of s and the uncertainty in b as a function of s. The width of the distribution of b increases significantly as s increases. Taking this into account, the peak of s shifts upwards. A marginalisation effect can often be confused with bias in the fit. However, if the full posterior distribution has the correct features, then this is not the case. Any effect in the marginalised distribution is a 'true' effect.

4.4.2 Best-fit Points and Credible Intervals

The ultimate goal of this analysis is to present inferences of the parameters of interest. Once the posterior has been marginalized, we need to extract parameter estimates and uncertainties.

The best-fit point is generally defined as the mode of the marginalised distribution, corresponding to the most probable value of the parameter. Since MCMC provides a discrete distribution, this involves finding the bin containing the largest number of steps. This ap-

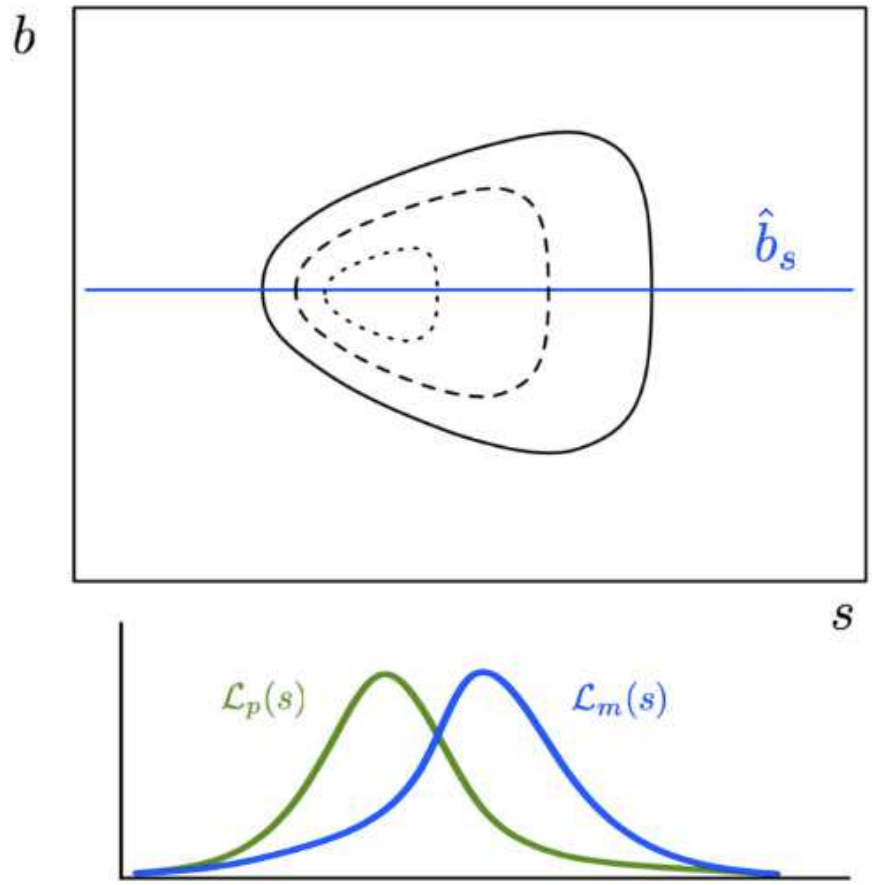

Figure 4.6: A 2D distribution as a function of a signal parameter, s, and a background parameter, b (top). The contours indicate the density of the distribution, showing regions of highest density (dotted) and lowest density (bold). The marginalised distribution as a function of only s is also shown (bottom) using marginalisation (blue) and profiling (green). Figure from [101]

proach is commonly used in MCMC analyses due to its simple implementation and is used in this analysis to present all best-fit points in 2D marginalised distribution of oscillation parameters and 1D marginalised distributions of nuisance parameters.

The one caveat to this approach is that the number of steps in each bin must be sufficient for the statistical uncertainty to be negligible when selecting the bin with the largest number of steps. This can become an issue when locating the mode of marginal distribution of several parameters of interest. In these situations, Markov chains must be run for longer (i.e. contain more steps) to ensure the statistical uncertainty in each bin is indeed negligible. An alternative solution is to increase the width of the bins, however, this reduces the resolution of the parameter

best-fit estimation.

Once the best-fit point has been found, we must also present the corresponding uncertainty. In Bayesian analyses, this is typically done by constructing credible intervals from the marginal distribution. Bayesian credible intervals are defined as parameter ranges that contain a certain fraction of the marginal posterior density. For example, the 99% credible interval is the range of parameter values that contain 99% of the posterior density. This can also be defined with the integral

$$\int_{\theta_1}^{\theta_2} P(\vec{\theta}|D)\, d\vec{\theta} \;=\; \alpha \tag{4.17}$$

where $P(\vec{\theta}|D)$ is the marginal posterior density for the parameters $\vec{\theta}$ and α is the fraction of the posterior density contained within the interval $\theta_1 \to \theta_2$. For a posterior obtained using MCMC, the fraction of the posterior density is equivalent to the fraction of MCMC steps. In this analysis, credible intervals are stated as 1σ, 2σ and 3σ, which correspond to the same fraction of the posterior as standard deviations from the mean in a normal distribution (i.e. 68.3%, 95.5%, 99.7%).

Although credible intervals are standard practice in Bayesian statistics, there are many ways to construct credible intervals such that they contain the correct fraction of the posterior. This analysis uses the highest posterior density (HPD) credible intervals. These are constructed such that all points or bins within the credible interval have a higher posterior density than points outside the interval. To construct an HPD credible interval, initially, the bin with the highest posterior density (i.e. the best-fit point) is located. Bins are then successively added, following a descending order of posterior density, until the desired proportion of the posterior is encompassed within the specified intervals.

For roughly Gaussian parameters, HPD credible intervals appear as one would intuitively expect i.e. the 1σ interval is contained within the 2σ interval symmetrically and so on. However,

for highly non-Gaussian parameters, the credible intervals can be non-contiguous. For example, a distribution with two distinct peaks in different regions may each be contained by non-contiguous 1σ intervals. This irregularity occurs in the distribution of $\sin^2\theta_{23}$ as shown in Figure 4.7.

4.4.3 Prior Reweighting

According to Eq 4.2, the posterior distribution explored by the Markov chain is a product of the likelihood that the model parameter choice agrees with data and prior knowledge of the model parameters. Hence, before running a Markov chain, we must input some prior information about the model parameters. The prior information could be a flat distribution, i.e. any value of a model parameter is equally likely, or contain some shape, making some regions more likely than others.

Often we want to investigate how a different choice of prior would affect the posterior distribution. A handy feature of Markov chains is that they can be reweighted to match a different choice of prior, without needing to run a new Markov chain. For an existing set of steps in a Markov chain, a new choice of prior would change the probability of that step being accepted. Hence, for each step in that chain, a weight is assigned according to the new probability. For example, consider a Markov chain containing N steps constructed with a flat prior in parameter θ. Now we want to look at the posterior distribution given some prior information about θ. In this case, let's consider a Gaussian prior, with a mean of μ_θ and a standard deviation of σ_θ. Any given step in the Markov chain, i, would receive a weight of

$$w_i \;=\; \exp\left(-\frac{(\theta_i - \mu_\theta)^2}{2\sigma_\theta^2}\right) \tag{4.18}$$

where θ_i is the value of the parameter at each step. It should be clear that the values of θ near the mean μ_θ will receive a weight close to one, and values far from the mean will receive

weights close to 0. It is also necessary to divide by the original prior and that the integral of the new prior is equal to 1. However, the one caveat is that the posterior distribution must cover the parameter values which the prior prefers sufficiently. If you impose a prior that weights up (or doesn't weight down in the previous case) regions with very few Markov chain steps, the shape of the distribution will be limited by sampling uncertainty. In this analysis, this technique is used to investigate the impact of applying an external constraint on θ_{13} determined by reactor neutrino experiments.

However, this tool not only allows us to apply external constraints to a Markov chain but can also be used to help make more confident statistical statements with the same number of Markov chain steps. As mentioned in Section 4.3.3, the number of steps required to make statistical statements (e.g. 3σ intervals, 5σ intervals) depends on the number of independent steps which statistically would fall outside the range. For example, for drawing a 3σ interval, as per HPD credible intervals, we would count the highest density bins until 99.73% of the density is contained. The larger the number of Markov chain steps outside of this range, the more confidence we have of where the range should be. As shown in Eq 4.13, roughly 160 million steps are required to obtain 30 *independent* steps outside the 3σ interval. DUNE and other next-generation long-baseline experiments aim to make 5σ statements about neutrino oscillation parameters. According to Eq 4.13, to obtain 30 independent samples outside the 5σ range with the same auto-correlation would require roughly 780 billion Markov chain steps. Even with significant resources, this is incredibly computationally heavy.

The ability to reweight a Markov chain under a different choice of prior allows us to perform a trick called importance sampling. The idea is that you can add an arbitrary prior into the fit that increases the posterior distribution in the regions of low likelihood and decreases the posterior distribution in the regions of high likelihood. This will cause the Markov chain to accumulate more steps in the regions outside the credible intervals. You can then reweight the output posterior with the inverse of the prior that was imposed in the fit, bringing you back to the original posterior. Regardless of the chosen prior, you should arrive at the same posterior

distribution as if a more uniform prior had been used when running the chain. In theory, this will allow statistical statements about parameters to be made more confidently with far fewer Markov chain steps. A proof of concept of this technique is presented in this analysis for δ_{CP}.

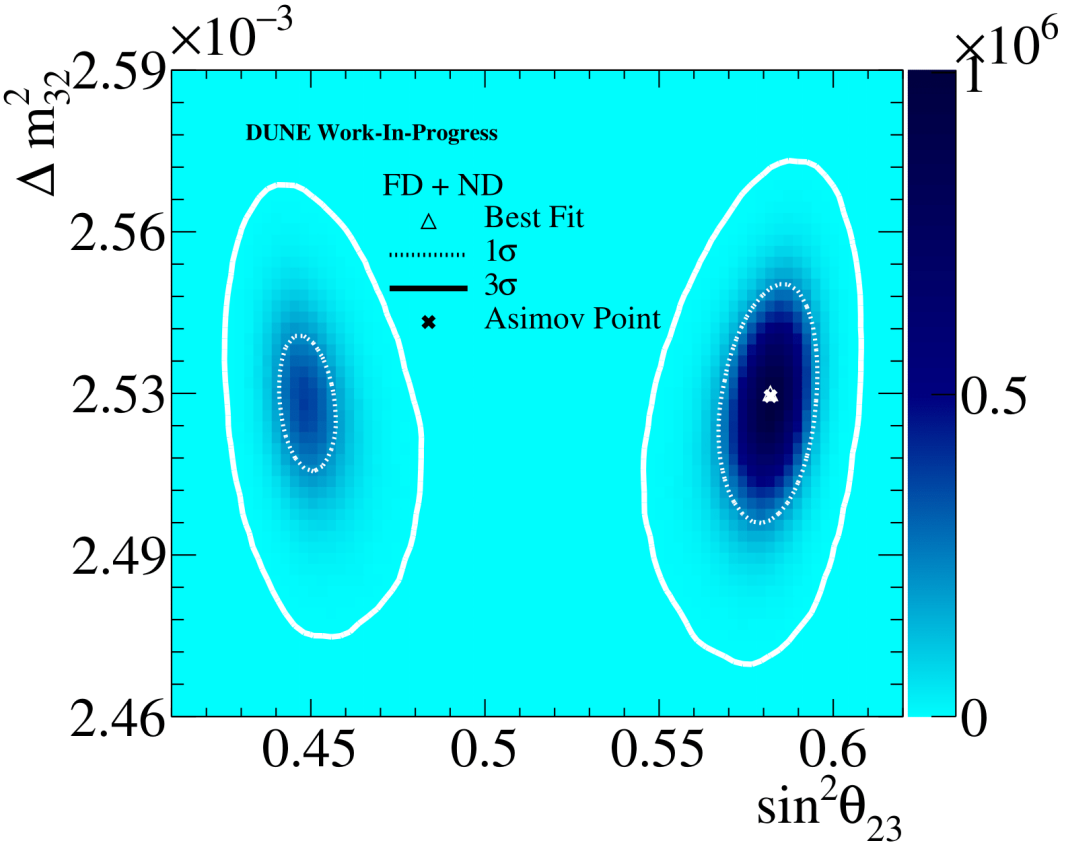

(a) 2D marginal distribution in $\sin^2\theta_{23}$ - Δm^2_{32}

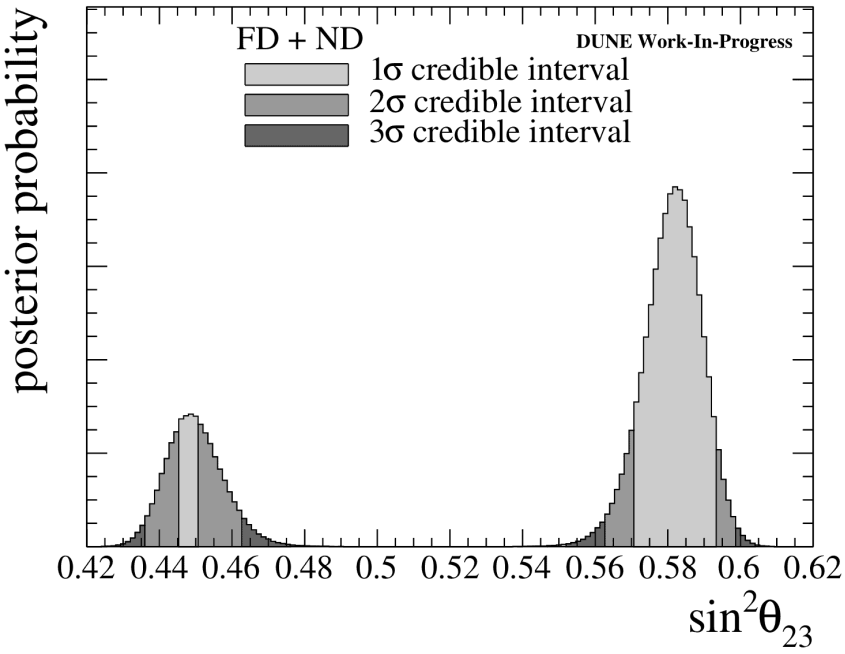

(b) 1D marginal distribution in $\sin^2\theta_{23}$

Figure 4.7: (a) 2D marginal distribution in $\sin^2\theta_{23}$ - Δm^2_{32} and (b) 1D marginal distribution in $\sin^2\theta_{23}$ from a DUNE Asimov fit (copied from Figure 7.7 and Figure 7.8 in Section 7.6). 2D credible intervals, 1D credible intervals, 2D best-fit points and 2D true Asimov point are shown.

Chapter 5

Parallel CPU Computation of Oscillation Probabilities

One of the most computationally demanding operations during a neutrino oscillation analysis is the calculation of the three-flavour oscillation probabilities. In this analysis, using MCMC as described in Chapter 4, oscillation parameters are proposed at each step, for which the oscillation probabilities for each Monte Carlo event must be calculated. The time required to evaluate each step within a Markov chain is crucial, as the total number of steps accumulated directly influences the reliability of the statistical inferences that can be drawn as explained in Section 4.3.3. In previous analyses, the oscillation probability calculation has been GPU-accelerated, reducing the contribution to the step time. However, generally, GPU resources are less readily available. CPU resources are relatively abundant, hence the reliance on GPUs limits the number of resources the analysis can be performed on. This chapter describes the adaptation of a parallelizable oscillation calculator to function for DUNE sensitivity studies. Oscillation probability calculation is discussed in Section 5.1. The oscillation calculator used is then described in Section 5.2 and the speed improvements are presented in Section 5.3.

5.1 Oscillation Probability Calculation

Given the standard three-flavour formulation, neutrinos propagate through space as a superposition of three mass eigenstates. Any given neutrino interaction is governed by its flavour. Since the mass eigenstates propagate at different rates, a given neutrino created in flavour α may be observed in flavour β, with a probability governed by the unitary PMNS transition matrix, shown in Eq 2.13, and as a function of the distance propagated and the energy, shown in Eq 2.29. Furthermore, the presence of matter introduces an extra potential that affects the probability. A more detailed description of neutrino oscillation including the derivation of the oscillation probabilities can be found in Section 2.3.2.

Generally, neutrino oscillation analyses are performed using large MC samples of simulated neutrino events. These samples can be used to approximate the distribution of neutrino events as a function of the incoming neutrino energy. The MC contains simulations of each combination of neutrino flavours i.e. $\nu_\mu \rightarrow \nu_\mu$, $\nu_\mu \rightarrow \nu_e$ etc. For a given choice of neutrino oscillation parameters, the simulated events must be reweighted according to the oscillation probability. For example, if the mixing angles θ_{23}, θ_{12} and θ_{13} are all set to 0 (i.e. there is no mixing between neutrino flavours), the weights for all MC events where $\alpha = \beta$ will be 1 and the weights for all MC events where $\alpha \neq \beta$ will be 0.

Calculating the oscillation probability for a neutrino event is computationally costly due to the necessity of multiple matrix calculations. Since each simulated neutrino event has a different energy, the calculation must be repeated for each event. Generally, $\mathcal{O}(10^6)$ neutrino events are simulated to ensure MC statistical error is negligible in the analysis. There are $\sim$ 3 million FD MC events in this analysis. Calculating oscillation probabilities for all 3 million events using a single core takes $\sim$ 13 seconds. Ignoring all other computations involved in evaluating the total likelihood of a given step in the Markov chain, obtaining the 100 million Markov chain steps would take $\sim$ 40 years. Evidently, this is unfeasible and the calculation must be accelerated in some way.

One method for accelerating the oscillation calculation is using a binned approach. Instead of calculating the oscillation probability for a given event using its true energy, oscillation probabilities are calculated at the bin centres of some pre-defined binning. During the analysis, the bin in which an MC event falls is found and the corresponding oscillation weight is applied. This has the advantage of being much faster since the oscillation calculation is only done N times per Markov chain step, where N is the number of bins, rather than for each MC event. The downside, however, is any shape information within the bin is lost. In some regions of neutrino energy, this might be a safe assumption. However, around the oscillation maxima where the probability changes significantly with energy, the difference can be significant. We are therefore motivated to find a way to keep all the true properties of an MC event in an analysis.

Given that each MC event is independent, the calculation of oscillation weights can be performed in parallel. This opens up the possibility of utilizing GPUs. A typical modern consumer GPU has between 100 - 10,000 cores that are used for graphical calculations. However, using application programming interfaces such as CUDA [102], the architecture can now be exploited for non-graphical calculations. This approach has been taken for recent oscillation analyses, resulting in a significant speed-up in the analysis run-time. However, the downside to this approach is the dependence on GPUs limits the available resources that the analysis can be performed on. To illustrate this, Figure 5.1 shows the number of CPU nodes compared to GPU nodes on the computing clusters that this analysis was able to use. Furthermore, at national computing facilities, the cost of using GPU resources is often higher than the cost of using CPU resources. This means that more hours can be utilised on a CPU node vs a GPU node with the same allocation.

The final option is to accelerate the oscillation calculation on CPUs. As mentioned before, the calculation is parallelizable meaning it can also be multi-threaded on CPUs. The following section describes CUDAProb3, an oscillation calculator with optimal multi-threading capabilities.

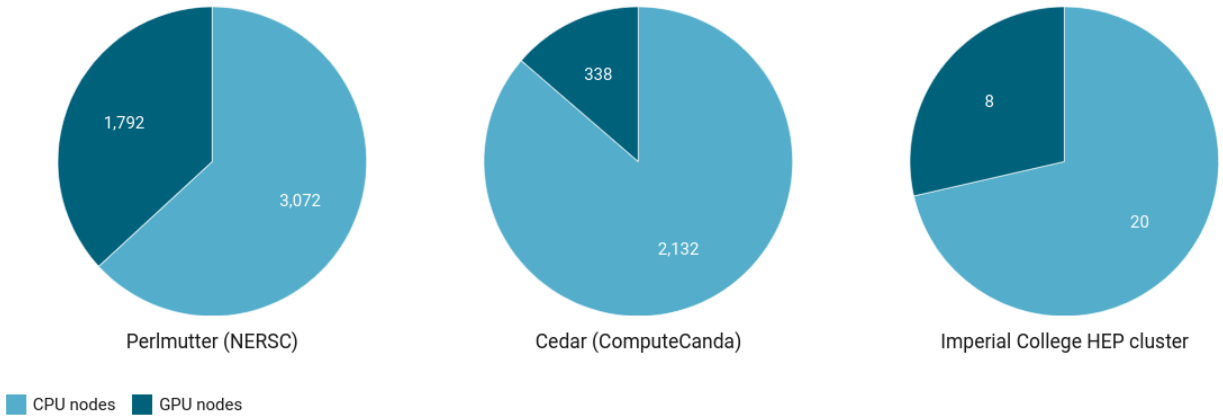

Figure 5.1: Pie charts illustrating the number of CPU and GPU nodes available on computing clusters that this analysis had access to [103–105].

5.2 CUDAProb3 Calculator

CUDAProb3 [106] is a C++ library, adapted from the Prob3++ library [107], initially designed for the calculation of atmospheric neutrino oscillation probabilities. The transition amplitude through a layer of matter with constant density is calculated using an analytic matrix diagonalization method [108]. CUDAProb3 calculates the probabilities for propagation through a sphere with piece-wise constant density. To function for long-baseline experiments, the CPU and GPU classes were split into atmospheric and beam classes, as shown in Figure 5.2.

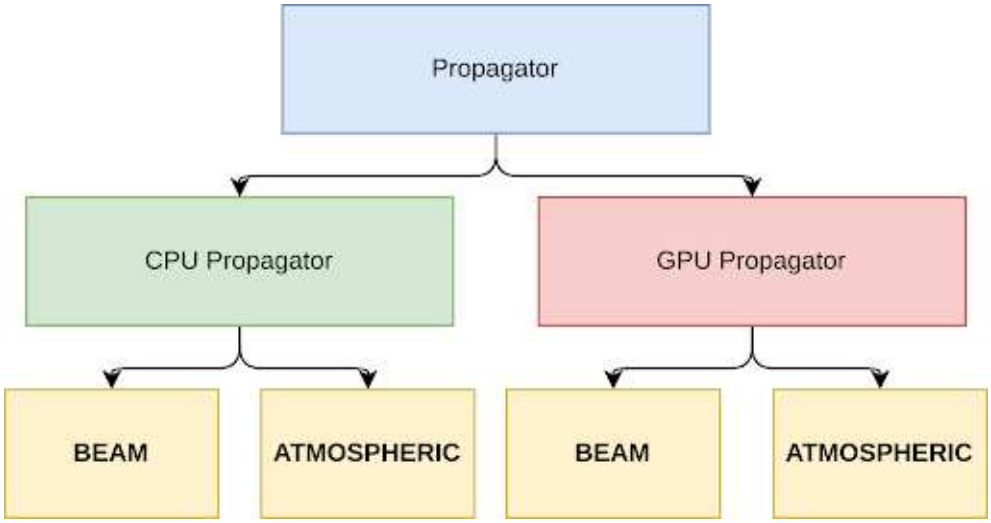

Figure 5.2: Flowchart showing the structure of CUDAProb3 after generalisation to both atmospheric and beamline neutrino oscillation calculation.

5.2.1 Beamline Calculation

For beamline oscillation calculation, the density and neutrino path length are passed to the calculator when initialising the class (i.e. for DUNE, $\rho = 2.848$ g/cm^3 and L = 1284.9 km). The density can be approximated as a constant in long-baseline accelerator experiments since the neutrinos only pass through the uppermost layer of the Earth. Figure 5.3 shows the oscillation probabilities for DUNE calculated by CUDAProb3 for both ν_μ and $\bar{\nu}_\mu$ using the NuFit 4.0 parameter values. We can see that the vast majority of ν_μ and $\bar{\nu}_\mu$ oscillate to ν_τ and $\bar{\nu}_\tau$. The oscillation maxima are indicated in the grey regions and the first two, which DUNE is sensitive to, occur at ~ 2.5 GeV and ~ 0.85 GeV as expected. We can also see a difference in the ν_e and $\bar{\nu}_e$ appearance probability caused by a combination of the matter effects and CP-violation.

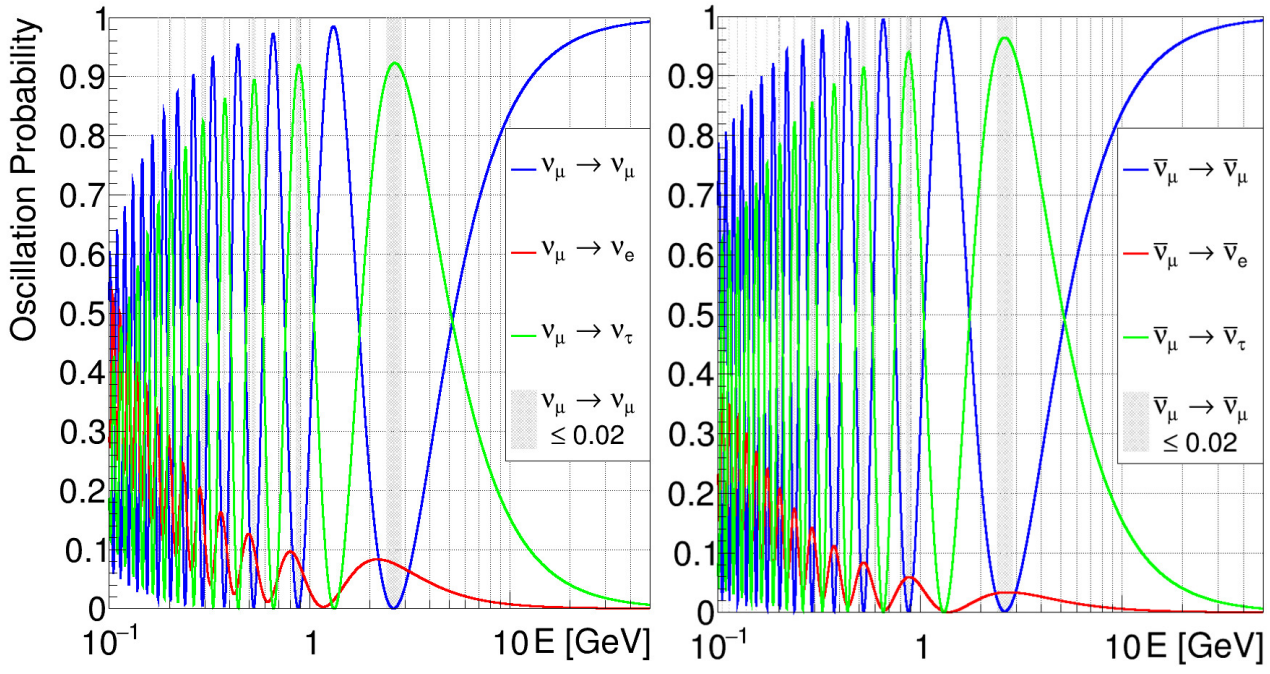

Figure 5.3: Beamline oscillation probabilities calculated by CUDAProb3 of ν_μ (left) and $\bar{\nu}_\mu$ (right) to all three flavours for DUNE baseline and density using NuFit 4.0 oscillation parameter values. The grey regions indicate where the $P_{\nu_\mu \rightarrow \nu_\mu}$ or $P_{\bar{\nu}_\mu \rightarrow \bar{\nu}_\mu}$ is ≤ 0.02 indicating the oscillation maxima.

5.2.2 Atmospheric Calculation

For atmospheric calculation, the neutrinos may pass through multiple Earth layers and travel along a wide range of distances. The path of the neutrino is dependent on the zenith angle, defined as the angle between the line connecting the detector to the centre of the Earth and the momentum of the neutrino. This is illustrated in Figure 5.4. The production height of the neutrinos is chosen and defines the depth of the outermost "atmosphere" layer. The Earth model is also chosen and defined as a set of N layers (not including the atmosphere) each with constant density ρ_i. Depending on the zenith angle, the neutrino may pass through anywhere between 1 - 2N layers (maximum 2N-1 Earth layers plus additional atmosphere layer).

First, the total path length, L, of the neutrino is calculated trigonometrically via the following function

$$\mathrm{L} = \sqrt{(\mathrm{R_E} + \mathrm{h})^2 - \mathrm{R_E^2}\,(1 - \cos\theta^2)} - \mathrm{R_E}\cos\theta \tag{5.1}$$

where $\mathrm{R_E}$ is the radius of Earth, h is the production height (i.e. the height above the Earth's surface where the neutrino was produced) and θ is the zenith angle. We can see that if $\theta = 0°$, in which the neutrino momentum is aligned with the line connecting the detector and the centre of the Earth, then simply L = h. If $\theta = 180°$, in which they are anti-aligned, then $\mathrm{L} = 2\mathrm{R_E} + \mathrm{h}$, referring to the case where the neutrino is produced on the other side of Earth to the detector. The distance traversed through each layer is then found and used to calculate each transition matrix.

As mentioned previously, the oscillation probabilities for atmospheric neutrinos depend on both energy and zenith angle, θ. Figure 5.5 shows oscillation probabilities for atmospheric neutrinos as a function of both energy and cosine zenith angle. We see that for higher energy neutrinos, much larger path lengths are required, i.e. $\theta \sim 180°$, for oscillations to become significant. This also helps to demonstrate why neutrino oscillation in the atmospheric sector

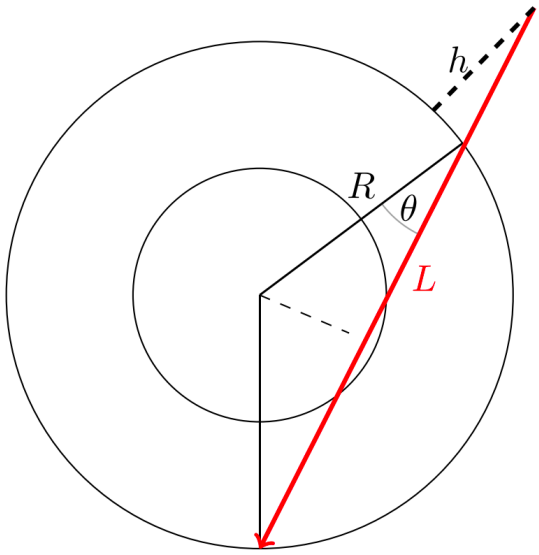

Figure 5.4: Diagram showing the path length, L, of an atmospheric neutrino (red line) through a two-layer solid sphere with radius R. Atmospheric neutrinos are produced at height h and enter the mantel of the earth at a certain energy E and zenith angle . Figure taken from [106].

was observed for up-going neutrinos, which travelled most of the way through Earth, and not down-going neutrinos. Since the production height of atmospheric neutrinos can vary significantly, oscillation probabilities can also be calculated for several heights. These probabilities are then averaged according to some chosen height binning.

5.3 CPU Multi-threaded Performance

The performance of multi-threaded beamline oscillation calculations with CUDAProb3 was tested by comparing against the single-threaded performance and the GPU performance. For all comparisons in this section, the CPU used was an Intel Xeon Gold 6230 20-core 20-thread server-grade processor. The GPU used was a NVIDIA Quadro RTX 4000. Figure 5.6 shows the runtime of CUDAProb3 with single-threaded, multi-threaded and GPU configurations against the number of events being computed. As expected, we see that the GPU massively outperforms even a CPU with 20 threads. Given that GPUs can often round 1000s of computations in parallel this is unsurprising. However, we also find that comparing the CPU configurations, the decrease in runtime is fairly linear. Ideally, when multi-threading, $2 \times$ (number of threads)

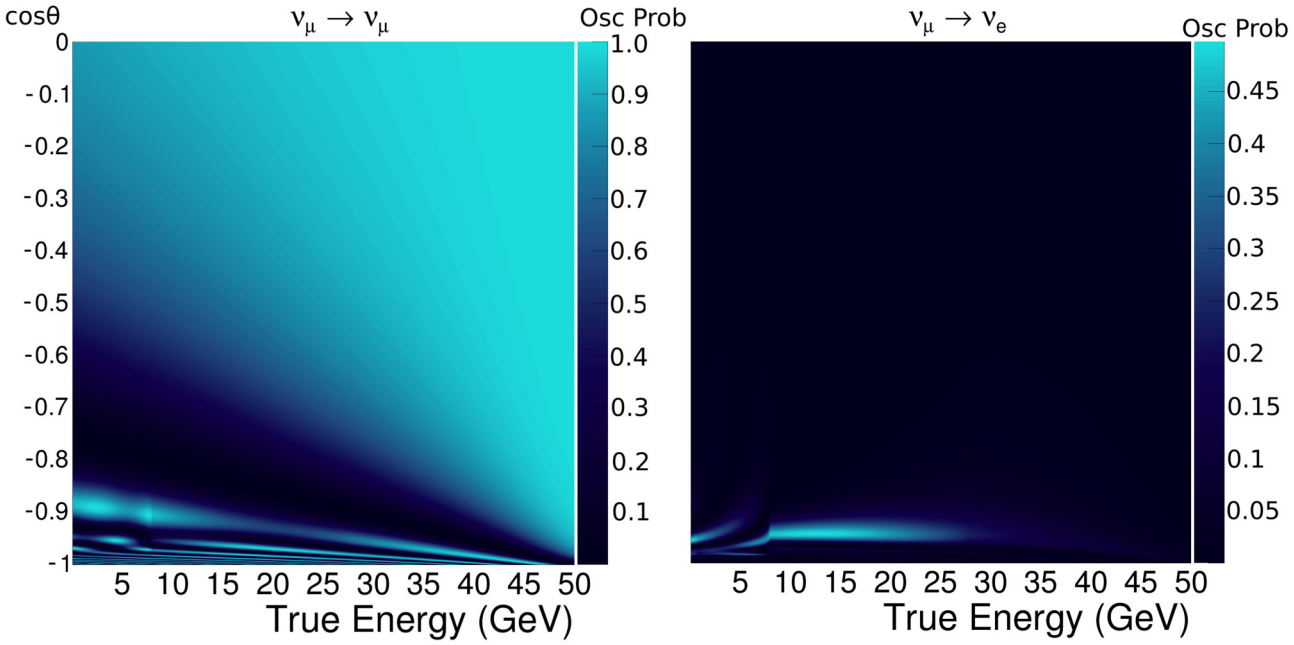

Figure 5.5: Atmospheric oscillation probabilities as a function of cosine zenith angle, $\cos\theta$, and neutrino energy, E, calculated by CUDAPRob3 for $\nu_\mu \rightarrow \nu_\mu$ (left) and $\nu_\mu \rightarrow \nu_e$ (right) using a four-layer density model, a fixed production height of 22.0m and NuFit 4.0 oscillation parameter values.

results in a 0.5× runtime. Often due to imperfections in implementation, doubling the threads results in a slightly less than half runtime effect. However, comparing the performance of the different number of threads, the multi-threading of the calculation is performing efficiently.

The performance of CPU multi-threading in CUDAProb3 was also evaluated using the MCMC step time. The MCMC step time is dominated by the likelihood evaluation and is comprised of two computationally expensive calculations: the oscillation probabilities and the systematic weights. Given a single-threaded CPU configuration, both calculations contribute significantly to the step time. Figure 5.7 shows the MCMC step time distributions using a CPU configuration of CUDAProb with a single-thread, 5 threads and a GPU configuration. We can see that with just 5 threads, the MCMC step-time becomes comparable to that using a GPU for the oscillation calculation. This is because, although the GPU is still an order of magnitude quicker, the systematic weight calculation becomes the largely dominant factor in the step time. Unless both calculations can be moved to a GPU, it's far more beneficial to utilise 5 threads in a CPU than a GPU for the reasons mentioned previously.

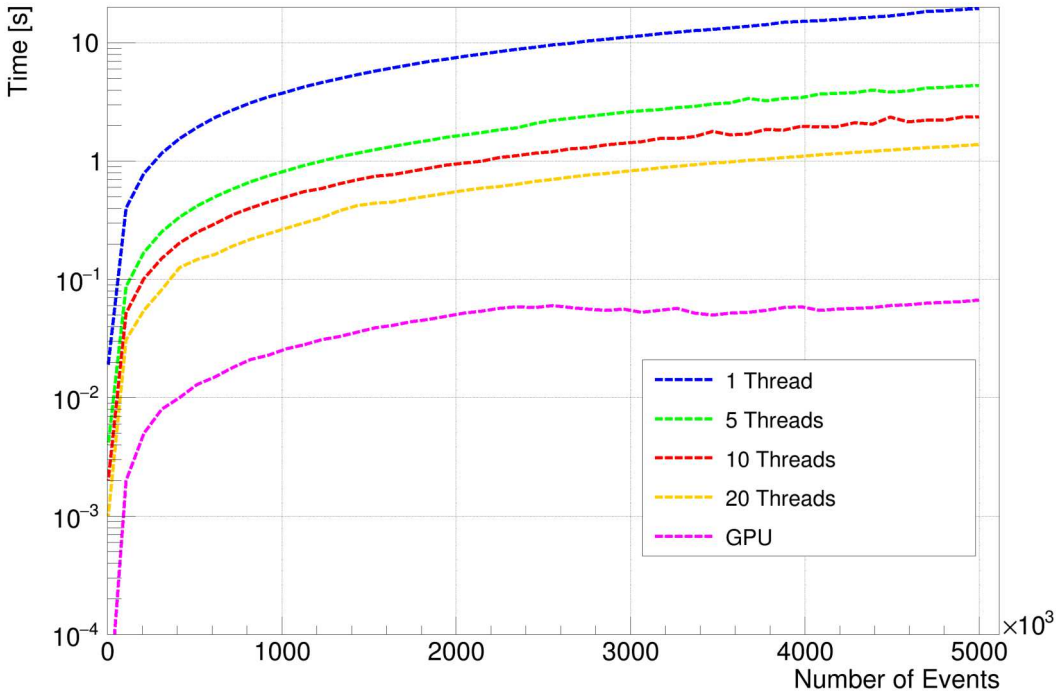

Figure 5.6: Runtime comparison of CUDAProb3 beamline oscillation probabilities for 5×10^3 - 5×10^5 events using a CPU with a single thread, 5 threads, 10 threads, 20 threads, and a GPU.

During this analysis, since no CUDA implementation of the systematic weight calculation was available, a CUDAProb3 multi-threaded CPU configuration was used throughout. A multi-threaded implementation of the systematic weight calculation was also used, allowing both processes to be multi-threaded sequentially.

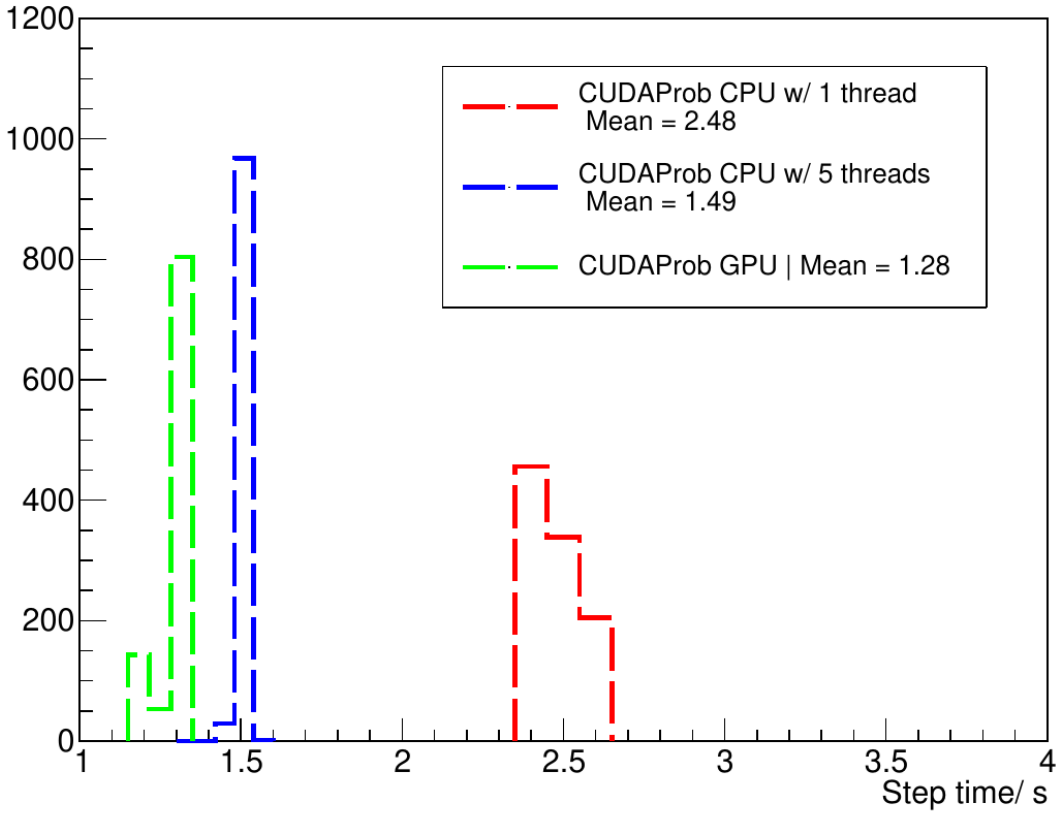

Figure 5.7: MCMC step time distributions of CUDAProb3 running a DUNE beamline FD-only chain using a CPU with a single thread, 5 threads and a GPU (all other calculations in the MCMC are CPU single-threaded).

Chapter 6

Far Detector Oscillation Sensitivity

This chapter uses the techniques described in Chapter 4 to perform the first-ever Bayesian FD-only sensitivity study of DUNE. This is done by creating a "fake" FD dataset that DUNE could observe assuming a set of true oscillation parameters. This study aims to perform an analysis using the same samples and MC inputs as was used in the DUNE Technical Design Report (TDR) [2]. However, there are significant differences in the implementation of the systematic uncertainties between the TDR-era analysis and the analysis presented in this thesis. This analysis takes an approach to systematic uncertainties that makes fewer assumptions, resulting in more accurate predictions of the event spectra given changes to systematic parameters. Next-generation long-baseline neutrino experiments will collect 10x more neutrino events than the current experiments, resulting in a far lower statistical uncertainty. The assumptions made in the oscillation analyses of current long-baseline experiments make little difference to the results due to the large impact of the statistical uncertainty. Since next-generation long-baseline experiments will be limited by systematic uncertainty, these simplifications can have large impacts. It is therefore imperative that systematics uncertainties are treated accurately to produce accurate estimations of the sensitivity of these experiments.

Furthermore, this is the first Bayesian analysis of DUNE. The TDR-era analysis was performed using a frequentist framework. Bayesian analyses allow for the in-depth study of system-

atic parameters with the same fit, since each nuisance parameter is treated as a free variable. As well as this, the posterior obtained from the Markov chain can be used to produce posteriors in other variables which are of interest, without the need for running another fit. Finally, there is an inherent benefit in validating the results of an analysis with a different statistical framework.

As described in Chapter 3, DUNE will collect neutrino events with both the near and far detectors. Although both FD and ND neutrino event samples will be used in the oscillation analysis, this chapter will only describe the FD samples and systematics. Many of the systematic uncertainties and implementations described in this chapter are the same or share similarities with ND samples. This chapter will also show the results of sensitivity studies in which only the FD samples are included, both with and without systematic uncertainties. This allows us to quantify the effect of the systematic uncertainty on the estimation of oscillation parameters. As well as this, these results can be compared with the results of sensitivity studies in which ND samples are also included, providing quantification of the impact of the DUNE ND.

Initially, the event selection for the FD samples is described in Section 6.1. The resulting FD MC event spectra predictions are presented in Section 6.2. The systematic uncertainty models, including flux, cross-section and detector uncertainties, included in the sensitivity study are discussed in Section 6.3. Prior predictive spectra are shown in Section 6.4 and the fitter validations that were done are discussed in Section 6.5. Finally, results of the sensitivity study are presented in Section 6.6, including both a statistics-error only sensitivity and sensitivity including systematic uncertainties.

6.1 Event Selection

At the FD, the aim is to select only CC interactions for both muon and electron neutrinos in both neutrino and antineutrino modes. This results in a total of four event samples: FHC ν_μ, FHC ν_e, RHC $\bar{\nu}_\mu$, RHC $\bar{\nu}_e$. After the MC event simulated hits in a detector are reconstructed,

DUNE uses the Convolutional Visual Network (CVN) event selection algorithm to decipher which sample an event belongs to.

The DUNE CVN is a Convolutional Neural Network (CNN) [109] that classifies neutrino interactions in the DUNE FD via image recognition. The CVN was trained using approximately 3 million MC events, a sample of MC events different from those used to make physics measurement sensitivities. The training input to the DUNE CVN contained three images for each MC event, one for each of the readout views. Each view consists of reconstructed hits on the wire planes, described in Section 3.4.1. An important feature of the DUNE CVN, which differs from traditional CNNs, is the ability for a given layer in the network to access several layers before it. For example, the n^{th} layer in the network can access the $(n-1)^{th}$ layer, as in a traditional CNN, and also the $(n-k)^{th}$ layer where $k \geq 2$. This allows the granular details contained within LArTPCs to propagate further into the network. A detailed description of the DUNE CVN architecture can be found in [110].

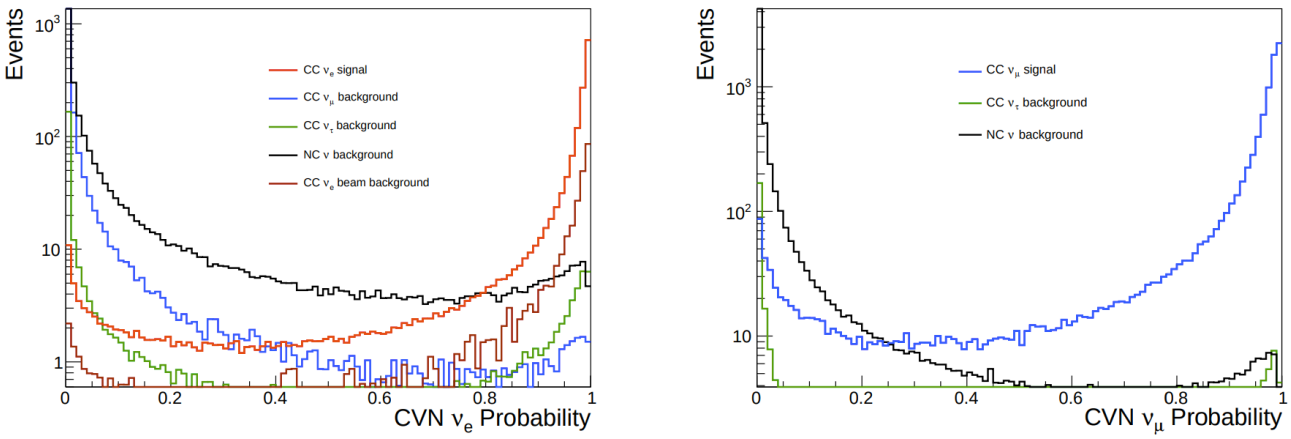

Figure 6.1: Probability distributions of CC ν_μ (left) and CC ν_e (right) from the DUNE CVN for the FHC mode split by the true interaction type. The probability axis is shown in a log scale. Figure from [2]

As mentioned previously, the primary goal of the DUNE CVN is to accurately classify neutrino interactions by flavour (and hence identify only CC interactions). Figure 6.1 shows the score distributions of a given MC event being CC ν_μ and CC ν_e for different true interaction types. For both cases, there is a good separation between signal and background interactions

at higher scores. The selection cut on the score was optimised based on sensitivity to δ_{CP}. The lower the threshold, the more events pass which provides more sensitivity, but the efficiency is worse which harms sensitivity. After optimisation, the following selection cuts were chosen:

- **CC $\boldsymbol{\nu_\mu/\bar{\nu}_\mu}$** : $\mathrm{P_{CVN}}\,(\mathrm{CC}\nu_\mu/\bar{\nu}_\mu) \geq \mathbf{0.5}$
- **CC $\boldsymbol{\nu_e/\bar{\nu}_e}$** : $\mathrm{P_{CVN}}\,(\mathrm{CC}\nu_e/\bar{\nu}_e) \geq \mathbf{0.85}$

Since the flavour classification scores are constructed so that they sum to 1, these selection cuts contain completely independent events (i.e. if $\mathrm{P_{CVN}}\,(\mathrm{CC}\nu_e) \geq 0.85$ then $\mathrm{P_{CVN}}\,(\mathrm{CC}\nu_\mu) < 0.5$ and vice versa).

As well as selection cuts on the flavour classification, fiducial volume cuts are also applied to the MC events. Ideally, these cuts would be applied to the reconstructed interaction vertex. However, due to some limitations in the simulation and reconstruction at the time, these cuts are applied to the true interaction vertex. The fiducial volume cuts were as follows:

- $-3.1\mathrm{m} < \mathrm{v_x} < 3.1\mathrm{m}$
- $-5.5\mathrm{m} < \mathrm{v_y} < 5.5\mathrm{m}$
- $0.5\mathrm{m} < \mathrm{v_z} < 12.44\mathrm{m}$

where $\mathrm{v_x}$, $\mathrm{v_x}$ and $\mathrm{v_z}$ are the true vertex positions in the x, y and z directions. This is measured in metres from the centre of the detector in x and y, and the front of the detector in z.

6.2 FD MC Prediction

Using the generated FD MC and applying the selection cuts described in Section 6.1, the predicted event spectra at the FD can be produced. The MC is generated using a full simulation chain: beam flux → Generates Events for Neutrino Interaction Experiments (GENIE) neutrino interaction generator [111] → Geant-4 based detector models. The beam design, neutrino flux and associated uncertainties are described in Section 3.1. The cross-section model and

uncertainties are described in more detail in Section 6.3.2.

The predicted event rates at the FD correspond to a 1.2 MW neutrino beam which results in 1.1×10^{21} protons on target (POT) per year including both FHC and RHC beam modes. This relies on an average up-time and efficiency of the Fermilab accelerator complex of 56%. The POT is then adjusted according to the planned additions of FD modules and beam power upgrades:

- **Initial beam run:** Two FD modules with a total mass of 20 kt and a 1.2 MW beam
- **After 1 year:** Add one FD module, total mass of 30kt
- **After 3 years:** Add one FD module, total mass of 40kt
- **After 6 years:** Upgrade to 2.4 MW beam
    - **Resulting 336 kt-MW-year POT per beam mode**: 1.3628×10^{23}

All event predictions and fake datasets in this analysis use a 336 kt-MW-year exposure. Predicted event distributions are shown given a choice of oscillation parameters. All predicted oscillated event distributions use the choice of oscillation parameters shown in Table 6.1. Unoscillated event distributions are those where oscillations have not been applied. For the FD samples, there is only one analysis variable: reconstructed neutrino energy.

The predicted unoscillated and oscillated event distributions, broken down by interaction mode as a function of reconstructed neutrino energy for all of the four DUNE FD samples using a 336 kt-MW-year exposure can be seen in Figure 6.2 and Figure 6.3 respectively. The corresponding event rates can be seen in Table 6.2 and Table 6.3. As we'd expect at these energy ranges, the largest contributions for most samples are CC DIS, CC RES, CC QE and CC 2p-2h interaction modes. The exception to this is the unoscillated FHC ν_e and RHC ν_e distributions. These event distributions effectively show us the intrinsic ν_e and $\bar{\nu}_e$ present in the FHC and RHC beam modes respectively, which act as a background to the oscillation signal. Since these intrinsic contributions are low, a significant portion of these are NC backgrounds which can mimic "ν_e-like" events. Since these are caused by the total ν flux, the rate of these

Parameter	Central Value	Uncertainty
$\sin^2\theta_{12}$	0.310	0.013
$\sin^2\theta_{23}$	0.582	0.021
$\sin^2\theta_{13}$	0.0224	0.0007
Δm^2_{21}/ eV^2	7.39×10^{-5}	1.8×10^{-6}
Δm^2_{32}/ eV^2	2.525×10^{-3}	3.4×10^{-5}
δ_{CP}/ rad	-2.498	1.17

Mass Ordering: Normal

Table 6.1: Central value and uncertainty of neutrino oscillation parameters used to make oscillated FD event spectra predictions. These oscillation parameter values correspond to NO.

backgrounds does not change from unoscillated to oscillated distributions.

6.3 Systematic Models and Uncertainties

6.3.1 Flux Uncertainties

The LBNF neutrino beam that will be used for DUNE and the resulting simulated neutrino fluxes are described in Section 3.1. The simulation of the beam uses an accurate description of the beam design [112], including magnetic horns and a target that has been optimised for sensitivity to CP violation.

The uncertainty in the neutrino flux is predominantly caused by uncertainties in the hadron production model and uncertainties in the parameters of the beam design, i.e. horn currents and positions of the horns and targets etc. Given measurements of the hadron production model uncertainty and LBNF design alignment uncertainties, the flux uncertainty at the first oscillation maxima is approximately 8% and 12% at the second oscillation maxima [2].

These uncertainties are highly correlated between true neutrino flavours and energy bins. These underlying flux uncertainties were included in the fit as bin-content normalisations uncertainties in bins of neutrino energies. Normalisation uncertainties are those which scale events according to the value of the systematic parameter. For example, if a normalisation param-

	FHC ν_μ	**FHC ν_e**	**RHC $\bar{\nu}_\mu$**	**RHC $\bar{\nu}_e$**
CC DIS	8692.05	148.79	4060.74	74.30
CC RES	8540.17	82.12	4234.53	48.32
CC QE	6497.34	58.97	2880.76	32.23
CC 2p-2h	1807.83	14.68	1043.02	9.98
CC Coherent	184.31	2.59	160.99	2.51
CC Nu-e Elastic	0.42	0.08	0.33	0.17
NC QE	1.67	0.25	1.16	0.12
NC DIS	180.62	69.39	88.23	32.70
NC RES	37.16	11.05	22.84	5.57
NC Coherent	0.00	1.17	0.00	1.29
NC Nu-e Elastic	0.00	1.76	0.00	1.12
Sample totals	**25941.57**	**390.85**	**12492.62**	**208.32**

Table 6.2: Unoscillated event rate predictions for the 4 samples at the DUNE FD broken down by interaction mode and corresponding to a 336 kt-MW-year exposure.

	FHC ν_μ	**FHC ν_e**	**RHC $\bar{\nu}_\mu$**	**RHC $\bar{\nu}_e$**
CC DIS	2128.56	497.76	1685.43	145.52
CC RES	2400.70	585.59	1249.50	153.61
CC QE	2128.56	449.00	957.00	108.62
CC 2p-2h	574.40	122.71	326.03	34.87
CC Coherent	56.66	13.91	48.89	6.34
CC Nu-e Elastic	0.40	0.55	0.33	0.24
NC QE	1.67	0.25	1.16	0.12
NC DIS	180.62	69.39	88.23	32.70
NC RES	37.16	11.05	22.84	5.57
NC Coherent	0.00	1.17	0.00	1.29
NC Nu-e Elastic	0.00	1.76	0.00	1.12
Sample totals	**8243.92**	**1753.14**	**4379.41**	**490.00**

Table 6.3: Oscillated event rate predictions for the 4 samples at the DUNE FD broken down by interaction mode using the oscillation parameter values in Table 6.1 and corresponding to a 336 kt-MW-year exposure.

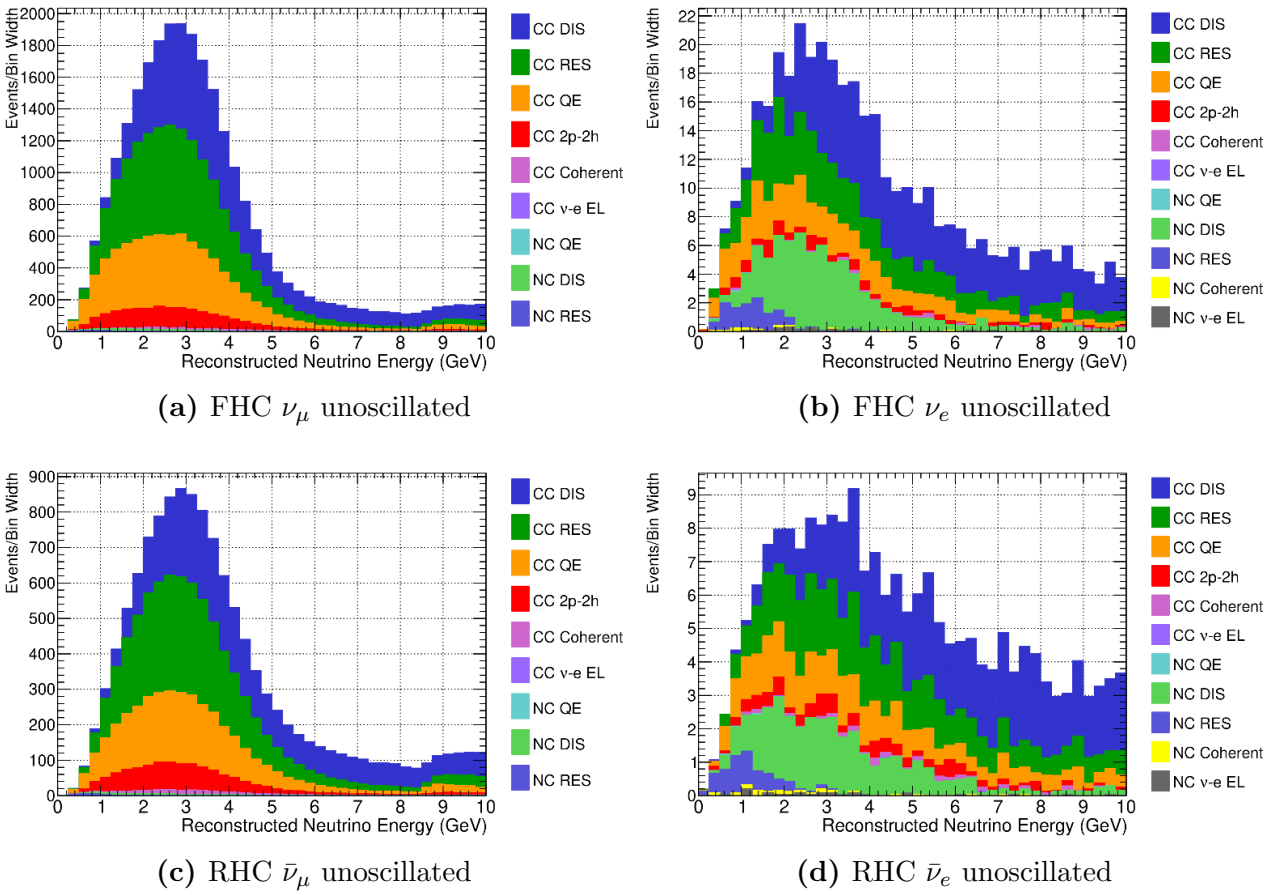

(a) FHC ν_μ unoscillated

(b) FHC ν_e unoscillated

(c) RHC $\bar{\nu}_\mu$ unoscillated

(d) RHC $\bar{\nu}_e$ unoscillated

Figure 6.2: Event spectra predictions as a function of reconstructed neutrino energy of the 4 samples at the DUNE FD: (a) FHC ν_μ, (b) FHC ν_e, (c) RHC $\bar{\nu}_\mu$ and (d) RHC $\bar{\nu}_e$ without oscillations. The predicted event spectra are also broken down by the different interaction modes.

eter shifts up by 20%, each event it affects is scaled by 1.2. The normalisation uncertainties were categorised into beam mode, produced neutrino flavour, energy bin as follows (where the number of energy bins = number of parameters per category): FHC ν_μ (19 parameters), FHC $\bar{\nu}_\mu$ (19 parameters), FHC ν_e (7 parameters), FHC $\bar{\nu}_e$ (7 parameters), RHC ν_μ (19 parameters), RHC $\bar{\nu}_\mu$ (19 parameters), RHC ν_e (7 parameters), RHC $\bar{\nu}_e$ (7 parameters). This results in a total of 104 normalisation parameters for the FD flux and 104 normalisation parameters for the ND. The variation and correlation of event rates in each bin were found by varying the true underlying flux uncertainty parameters. The variation standard deviation and correlations of each bin then define the prior uncertainty and correlations of the bin-to-bin normalisation parameters. The central value of each normalisation parameter is 1. The flux covariance matrix

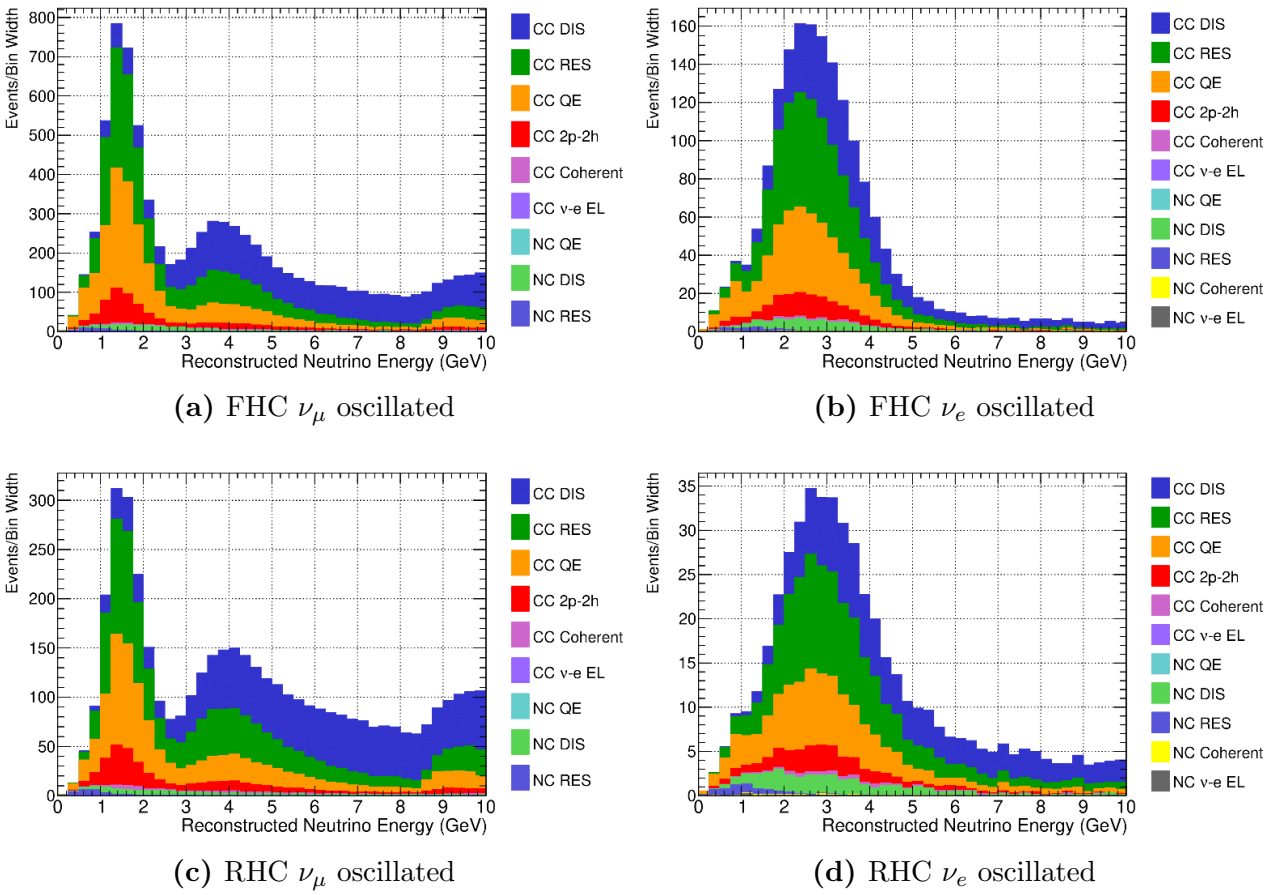

(a) FHC ν_μ oscillated

(b) FHC ν_e oscillated

(c) RHC ν_μ oscillated

(d) RHC ν_e oscillated

Figure 6.3: Event spectra predictions as a function of reconstructed neutrino energy of the 4 samples at the DUNE FD: (a) FHC ν_μ, (b) FHC ν_e, (c) RHC $\bar{\nu}_\mu$ and (d) RHC $\bar{\nu}_e$ with oscillations using the parameter values shown in Table 6.1. The predicted event spectra are also broken down by the different interaction modes.

used in this analysis is shown in Figure 6.4.

6.3.2 Cross Section Uncertainties

The neutrino interaction model used in this analysis is implemented in the GENIE generator. More specifically, the fixed GENIE version v2.12.10 is used. As well as this, some variations to this model were added to parameterise uncertainties currently not implemented in GENIE. Very generally, the model of neutrino-nucleus interactions can be factorised into three components: initial state, nucleon interaction and final state propagation. The initial state describes the momentum and positions of the nucleons inside the nucleus, as well as effects that cause kinematic modifications to the final state, such as binding energy. The final state interactions

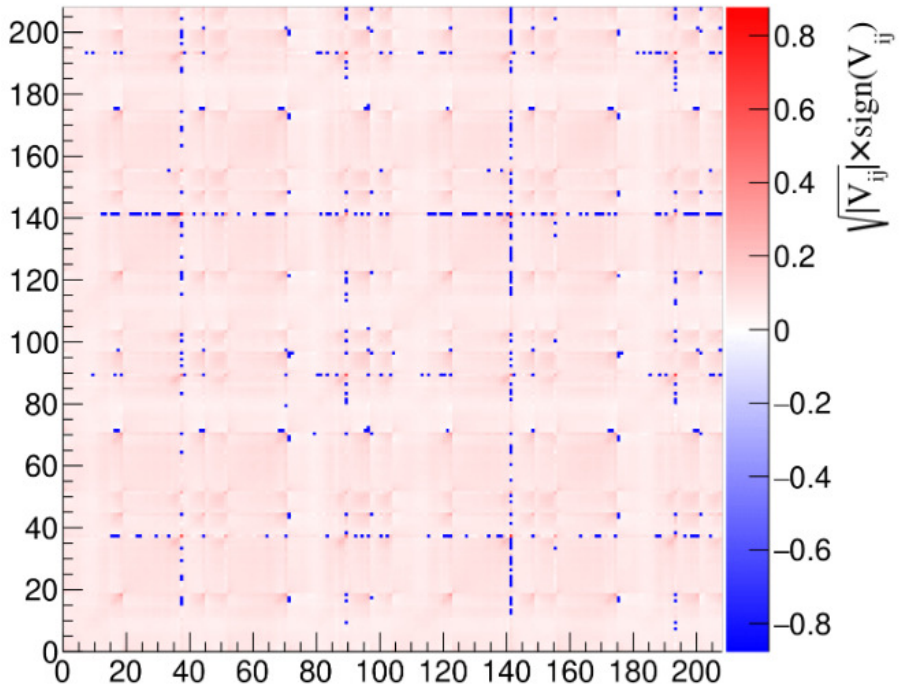

Figure 6.4: Square root of the fractional covariance matrix used for the flux systematic parameters. The x and y axes show the 208 flux parameters corresponding to ND and FD ν_μ and ν_e samples for neutrino and antineutrino modes and certain energy bins. The z-axis shows the square root of the covariance matrix entry while keeping the sign unchanged.

(FSI) refer to the alteration of the momentum and energy of particles produced in the final state or changes to the multiplicity of particles emitted. Within the nucleus, hadrons can be absorbed, scattered or produce further hadrons. These effects can lead to missing energy that is invisible to the detector, which can easily be on the order of hundreds of MeV. Given the importance of neutrino energy in oscillation measurements, there must be enough freedom in the model to prevent a bias in the result.

The interaction uncertainties implemented in this analysis can be categorised into 6 rough groups:

- **QE uncertainties** The uncertainties in QE interactions stem from the axial form factor of CCQE interactions, which is modelled as a dipole with one free parameter, the axial mass, M_A [113]. Uncertainty is also introduced corresponding to the choice of vector form factor. As well as this, low Q^2 interaction cross-sections are modified using the Random Phase Approximation (RPA) which accounts for correlations between nucleons within the

nucleus. The shape of the correction can be parameterised as a function of Q^2 using the Bernstein polynomials [114] and results in 3 Bernstein RPA (BeRPA) parameters (A, B, C) which control the behaviour at increasing Q^2.

- **2*p*2*h* uncertainties** The underlying theoretical model for $2p2h$ interactions in GENIE, the "Valencia" model [115], is not compatible with data from the MINERvA [116] and NOvA experiments[117]. Corrections were applied to the model and uncertainties are included for the energy dependence of these corrections. As well as this, uncertainties stem from the scaling of the $2p2h$ prediction from carbon, in which the mismatch was found, to argon.

- **Resonant uncertainties** The axial and vector form factors also control the interaction model for resonant events. As such, axial and vector mass parameters are also included for CC and NC RES interactions. The angular distribution of pions from the decaying Δ can be modelled as isotropic or follow predictions from the Rein-Sehgal model of pion production [118]. An uncertainty parameter is introduced to switch between both predictions.

- **Non-resonant and DIS uncertainties** More freedom is necessary in the region between single-pion production and DIS which is currently not well-modelled. Additional parameters were added for all combinations of 1,2 or $\geq$ 3 pions in the final state, CC or NC, neutrino or antineutrino and interactions on proton or neutron [117]. Except for merging the neutron and proton parameters for CC neutrino single-pion production, this resulted in a total of 23 dials adding freedom to these parameters. As well as this, 4 parameters are included from the Bodek-Yang DIS model [119].

- **Final state interaction uncertainties** There are several final state uncertainties included which stem from the cascade model implemented in GENIE. These are associated with probabilities of charge exchange, elastic/inelastic interactions, absorption and pion production for both pions and nucleons.

- **Neutrino flavour dependent uncertainties** The neutrino cross-section also includes

terms which are dependent on lepton mass. The form factors for these terms have large uncertainties which results in changes in the cross-section ratio of ν_μ/ν_e as well as $\nu/\bar{\nu}$.

Table 6.4 and Table 6.5 list all the cross-section systematic parameters included in this analysis, showing a description of the uncertainty, their central value, prior uncertainty and valid ranges. In total, 56 cross-section parameters were included. The units of all cross-section parameters are relative to their generated value and prior uncertainty i.e. 0 is the value at which the MC was generated and 1 is 1 standard deviation away. Parameters which have a prior central value that is not 0, shown in Table 6.4 and Table 6.5, have pre-fit values different from the generated value.

The cross-section model is a function of several event kinematics which means different events can be affected differently given a change to the model parameters. Furthermore, model parameter changes may often not result in linear changes and can contain complex shape characteristics. As a result, all cross-section parameters, unlike the flux parameters, are implemented with response functions. For each MC event, GENIE (including all additions to the base model) returns weights given a shift of a single parameter at evenly spaced points i.e. M_A^{QE} at -3σ, -2σ, -1σ, 0, 1σ, 2σ, 3σ. A cubic spline is used to interpolate between these points, creating a continuous response function. To reduce the computational load, events are binned in the analysis variables (for FD this is just reconstructed energy), true neutrino energy, true interaction mode and true initial and final neutrino flavour. The average weight for each parameter shift in each of these bins is evaluated and the spline is created between these bins. When reweighting the Monte Carlo, the spline corresponding to the kinematic bin that the event falls into is evaluated.

6.3.3 FD Detector Uncertainties

Detector effects can impact the quality of reconstructed variables and selection efficiency. Since our FD samples are only binned in reconstructed energy, we care only about how the uncertainty in our detector model affects the reconstructed energy and not other reconstructed quantities.

Group	Parameter	Description	Central Value	1σ Uncertainty	bounds
QE	M_A^{CCQE}	Axial mass for CCQE	0.99 GeV	+0.25/-0.15 GeV	$-\infty \leftrightarrow \infty$
	CCQE VecFF	Choice of CCQE vector form factor	0	N/A	$0 \leftrightarrow 1$
	CCQE KF	Fermi surface momentum for Pauli blocking	0	1	$0 \leftrightarrow 1$
	BeRPA A	RPA model suppression polynomial A	0.59 (*)	$\pm$ 0.12	$-\infty \leftrightarrow \infty$
	BeRPA B	RPA model suppression polynomial B	1.05 (*)	$\pm$ 0.21	$-\infty \leftrightarrow \infty$
	BeRPA D	RPA model suppression polynomial D	1.13 (*)	$\pm$ 0.17	$-\infty \leftrightarrow \infty$
$2p2h$	E_{2p2h} nu A	neutrino $2p2h$ Energy dependence parameter A	0	1	$-\infty \leftrightarrow \infty$
	E_{2p2h} nu B	neutrino $2p2h$ Energy dependence parameter B	0	1	$-\infty \leftrightarrow \infty$
	E_{2p2h} nubar A	antineutrino $2p2h$ Energy dependence parameter A	0	1	$-\infty \leftrightarrow \infty$
	E_{2p2h} nubar B	antineutrino $2p2h$ Energy dependence parameter B	0	1	$-\infty \leftrightarrow \infty$
	ArC$2p2h$ nu	neutrino $2p2h$ Ar/C scaling	0	1	$-\infty \leftrightarrow \infty$
	ArC$2p2h$ nubar	antineutrino neutrino $2p2h$ Ar/C scaling	0	1	$-\infty \leftrightarrow \infty$
RES	M_A^{CCRES}	Axial mass for CC RES	0.94 GeV	$\pm$ 0.05 GeV	$-\infty \leftrightarrow \infty$
	M_V^{CCRES}	Vector mass for CC RES	0.84 GeV	$\pm$ 0.084 GeV	$-\infty \leftrightarrow \infty$
	M_A^{NCRES}	Axial mass for NC RES	1.12 GeV	$\pm$ 0.22 GeV	$-\infty \leftrightarrow \infty$
	M_V^{NCRES}	Vector mass for NC RES	0.84 GeV	$\pm$ 0.084 GeV	$-\infty \leftrightarrow \infty$
	$\theta_\pi^{\Delta\text{decay}}$	θ_π distribution in decaying Δ rest frame	0	1	$0 \leftrightarrow 1$
Non-RES	NR nu n CC 2pi	Non-RES neutrino-neutron CC 2-pion production	0	1	$-\infty \leftrightarrow \infty$
	NR nu n CC 3pi	Non-RES neutrino-neutron CC $\geq$ 3-pion production	0	1	$-\infty \leftrightarrow \infty$
	NR nu p CC 2pi	Non-RES neutrino-proton CC 2-pion production	0	1	$-\infty \leftrightarrow \infty$
	NR nu p CC 3pi	Non-RES neutrino-proton CC $\geq$ 3-pion production	0	1	$-\infty \leftrightarrow \infty$
	NR nu np CC 1pi	Non-RES neutrino-proton/neutron CC 1-pion production	-1.14	0.02	$-\infty \leftrightarrow \infty$
	NR nu n NC 1pi	Non-RES neutrino-neutron NC 1-pion production	0	1	$-\infty \leftrightarrow \infty$
	NR nu n NC 2pi	Non-RES neutrino-neutron NC 2-pion production	0	1	$-\infty \leftrightarrow \infty$
	NR nu n NC 3pi	Non-RES neutrino-neutron NC $\geq$ 3-pion production	0	1	$-\infty \leftrightarrow \infty$
	NR nu p NC 1pi	Non-RES neutrino-proton NC 1-pion production	0	1	$-\infty \leftrightarrow \infty$
	NR nu p NC 2pi	Non-RES neutrino-proton NC 2-pion production	0	1	$-\infty \leftrightarrow \infty$
	NR nu p NC 3pi	Non-RES neutrino-proton NC $\geq$ 3-pion production	0	1	$-\infty \leftrightarrow \infty$
	NR nubar n CC 1pi	Non-RES antineutrino-neutron CC 1-pion production	0	1	$-\infty \leftrightarrow \infty$
	NR nubar n CC 2pi	Non-RES antineutrino-neutron CC 2-pion production	0	1	$-\infty \leftrightarrow \infty$
	NR nubar n CC 3pi	Non-RES antineutrino-neutron CC $\geq$ 3-pion production	0	1	$-\infty \leftrightarrow \infty$
	NR nubar p CC 1pi	Non-RES antineutrino-proton CC 1-pion production	0	1	$-\infty \leftrightarrow \infty$
	NR nubar p CC 2pi	Non-RES antineutrino-proton CC 2-pion production	0	1	$-\infty \leftrightarrow \infty$
	NR nubar p CC 3pi	Non-RES antineutrino-proton CC $\geq$ 3-pion production	0	1	$-\infty \leftrightarrow \infty$
	NR nubar n NC 1pi	Non-RES antineutrino-neutron NC 1-pion production	0	1	$-\infty \leftrightarrow \infty$
	NR nubar n NC 2pi	Non-RES antineutrino-neutron NC 2-pion production	0	1	$-\infty \leftrightarrow \infty$
	NR nubar n NC 3pi	Non-RES antineutrino-neutron NC $\geq$ 3-pion production	0	1	$-\infty \leftrightarrow \infty$
	NR nubar p NC 1pi	Non-RES antineutrino-proton NC 1-pion production	0	1	$-\infty \leftrightarrow \infty$
	NR nubar p NC 2pi	Non-RES antineutrino-proton NC 2-pion production	0	1	$-\infty \leftrightarrow \infty$
	NR nubar p NC 3pi	Non-RES antineutrino-proton NC $\geq$ 3-pion production	0	1	$-\infty \leftrightarrow \infty$

Table 6.4: Neutrino interaction cross-section systematic parameters included in this analysis, within categories of QE, $2p2h$, RES and Non-res, showing the central value, prior uncertainty and bounds of each parameter. * Weights are applied to the MC to reweight for pre-fit tuning in the BeRPA parameters.

Group	Parameter	Description	Central Value	1σ Uncertainty	bounds
DIS	A_{HT}	Higher-twist parameter, A, in Bodek-Yang model	0.538	$\pm$ 0.134	$-\infty \leftrightarrow \infty$
	B_{HT}	Higher-twist parameter, B, in Bodek-Yang model	0.305	$\pm$ 0.076	$-\infty \leftrightarrow \infty$
	C_{V1u}	GRV98 valence quark distribution function correction to Bodek-Yang model	0.291	$\pm$ 0.087	$-\infty \leftrightarrow \infty$
	C_{V2u}	GRV98 valence quark distribution function correction to Bodek-Yang model	0.189	$\pm$ 0.076	$-\infty \leftrightarrow \infty$
FSI	x_{cex}^{N}	Nucleon charge exchange probability	1	$\pm$ 0.50	$-\infty \leftrightarrow \infty$
	x_{el}^{N}	Nucleon elastic reaction probability	1	$\pm$ 0.30	$-\infty \leftrightarrow \infty$
	x_{inel}^{N}	Nucleon inelastic reaction probability	1	$\pm$ 0.40	$-\infty \leftrightarrow \infty$
	x_{abs}^{N}	Nucleon absorption probability	1	$\pm$ 0.20	$-\infty \leftrightarrow \infty$
	x_{π}^{N}	Nucleon π-production probability	1	$\pm$ 0.20	$-\infty \leftrightarrow \infty$
	x_{cex}^{π}	π charge exchange probability	1	$\pm$ 0.50	$-\infty \leftrightarrow \infty$
	x_{el}^{π}	π elastic reaction probability	1	$\pm$ 0.10	$-\infty \leftrightarrow \infty$
	x_{inel}^{π}	π inelastic reaction probability	1	$\pm$ 0.40	$-\infty \leftrightarrow \infty$
	x_{abs}^{π}	π absorption probability	1	$\pm$ 0.30	$-\infty \leftrightarrow \infty$
	x_{π}^{π}	π π-production probability	1	$\pm$ 0.20	$-\infty \leftrightarrow \infty$
Flavour	$\sigma_{\nu_e/\bar{\nu}_e}$	$\nu_e/\bar{\nu}_e$ cross-section ratio	1	1	$0 \leftrightarrow 1$
	σ_{ν_e/ν_μ}	ν_e/ν_μ cross-section ratio	1	1	$0 \leftrightarrow 1$

Table 6.5: Neutrino interaction cross-section systematic parameters included in this analysis, within categories of DIS, FSI and flavour-dependent, showing the central value, prior uncertainty and bounds of each parameter.

We also care about how uncertainty in our event reconstruction might affect which events fall into which samples. The FD detector systematics can be split into three subgroups: energy scale, resolution, and selection efficiency.

Energy Scale

Uncertainty in the energy scale is separated by particle species. Total energy scale as well as four particle species are considered: muons, charged hadrons, neutrons and electromagnetic showers. The total energy scale effectively acts as a calibration uncertainty for calorimetric reconstruction, while the particle species energy scale acts as a particle response uncertainty. Both total and particle-specific uncertainties are parameterised in the same way, in which each systematic has 3 parameters:

Particle group	p_0	p_1	p_2	r
Total (except muons)	2%	1%	2%	
Hadronic (p, $\pi^{\pm}$)	5%	5%	5%	2%
μ	2%	0.5%	2%	2%
n	20%	30%	30%	10%
electromagnetic (e, γ, π^0)	2.5%	2.5%	2.5%	2%

Table 6.6: Prior uncertainty on the three energy scale parameters (p_0, p_1 and p_2) and the resolution (r) for different particle groups.

$$E'_{rec} = E_{rec} \times \left(p_0 + p_1 \sqrt{E_{rec}} + \frac{p_2}{\sqrt{E_{rec}}} \right) \tag{6.1}$$

where E'_{rec} is the shifted reconstructed energy, E_{rec} is the nominal reconstructed energy and p_0, p_1 and p_2 are free systematic parameters. The p_0 parameter controls the energy-independent uncertainty, while the p_1 and p_2 parameters provide some freedom in the shape of the uncertainty with respect to energy. Each of these parameters, for each uncertainty group, has a prior uncertainty that has been derived from previous liquid argon experiments, and previous long-baseline experiments.

Table 6.6 shows the prior uncertainty on each of the energy scale parameters for each particle group. The total energy scale does not include muons since muon energy is reconstructed using range or Coulomb scattering, rather than calorimetrically. Uncertainties associated with neutrons are high due to significant reconstruction issues. The response of neutrons in LAr is still an active area of research, aiming to understand the relationship between neutron capture and scattering to its true kinetic energy.

Resolution

The resolution used in the detector model is also a source of uncertainty. This is also split up by particle group. The effect of resolution on the reconstructed energy is calculated by

$$E'_{rec} = r \times (E_{true} - E_{rec}) \tag{6.2}$$

where r is a free parameter associated with the resolution for the particle group. This has the effect of scaling the smearing of energy. The prior uncertainty caused by the resolution is also shown for each particle group in Table 6.6. Once again, the neutron energy resolution has a much higher prior uncertainty than the other particle groups.

Selection Efficiency

Another uncertainty that needs to be considered is the event selection uncertainty. As mentioned in Section 6.1, the event selection is done via the CVN which assigns each event a score of being ν_μ-like or ν_e-like. The prior uncertainty associated with both scores was chosen to be 0.01. A fiducial volume uncertainty was also included for both ν_μ-like and ν_e-like events.

Except for the fiducial volume uncertainty, the detector uncertainties all change the value of reconstructed variables. The energy scale and resolution uncertainties change the reconstructed energy and the selection efficiency uncertainty changes the CVN score. In this analysis, these have all been implemented as shift systematics. Since the reconstructed variables for each MC event are stored during the Markov chain, they can be varied depending on the value of the free parameter. For example, all 19 resolution and energy scale parameters will shift the reconstructed energy for a given MC event depending on the value of the parameters at that point in the Markov chain. Given the new reconstructed energy after shifting, the event may have crossed an analysis bin boundary. If this is the case, the event is moved into the correct bin. This provides an accurate method of changing the event spectra given the parameterisation of the energy scale and resolution uncertainties. This is different from the TDR-era analysis, which used spline functions to mimic the effect of these shifts. Similarly, for the selection efficiencies, the CVN scores are shifted during the fit. If the CVN probability for a given event crosses the event selection cut, the event moves in or out of the selection, changing the event

spectra. Given the prior uncertainty for the selection efficiency, it is not possible for an event to simultaneously pass both ν_μ and ν_e cuts, preventing double counting.

Given that there is no visible change in the spectra until events move across bins or selection cuts, this often causes discontinuous likelihoods. Common minimisers, such as gradient descent, use the gradient of the likelihood to find global minima and hence run into issues if the likelihood space contains discontinuities. However, since the MCMC algorithm used here does not evaluate the gradient this does not cause an issue.

6.4 Prior Predicitves

The nominal prediction of the event distributions for the samples change depending on the values of the systematic parameters described in section Section 6.3. To get an idea of the effect of the systematic uncertainty on the event distributions in different samples, prior predictive spectra are produced. This involves randomly throwing the values of the systematic parameters according to their prior uncertainties and correlations. The event spectra are then reweighted according to these values and an event rate distribution is created in each bin. The prior predictive spectra for the DUNE FD samples using 10000 throws of the systematic parameters are shown in Figure 6.5.

Although the uncertainty bands are much wider at higher energy values, the fractional width of the bands remains relatively constant. The largest fractional change in the event rate occurs in the two lowest energy bins. The uncertainty on the total number of events can be estimated by calculating the root mean square of the integral of each toy spectra. Table 6.7 shows the percentage uncertainty on the total event rate for each sample and the contributions from each uncertainty group. The statistical uncertainty was also included where, for each bin, a random number was drawn from a Poisson distribution with a mean of the nominal event count in the bin. The distribution of event spectra integrals was done in the same way as the systematic throws. The largest contribution to the prior uncertainty comes from the cross-section for every sample, with the smallest contribution coming from the detector uncertainties.

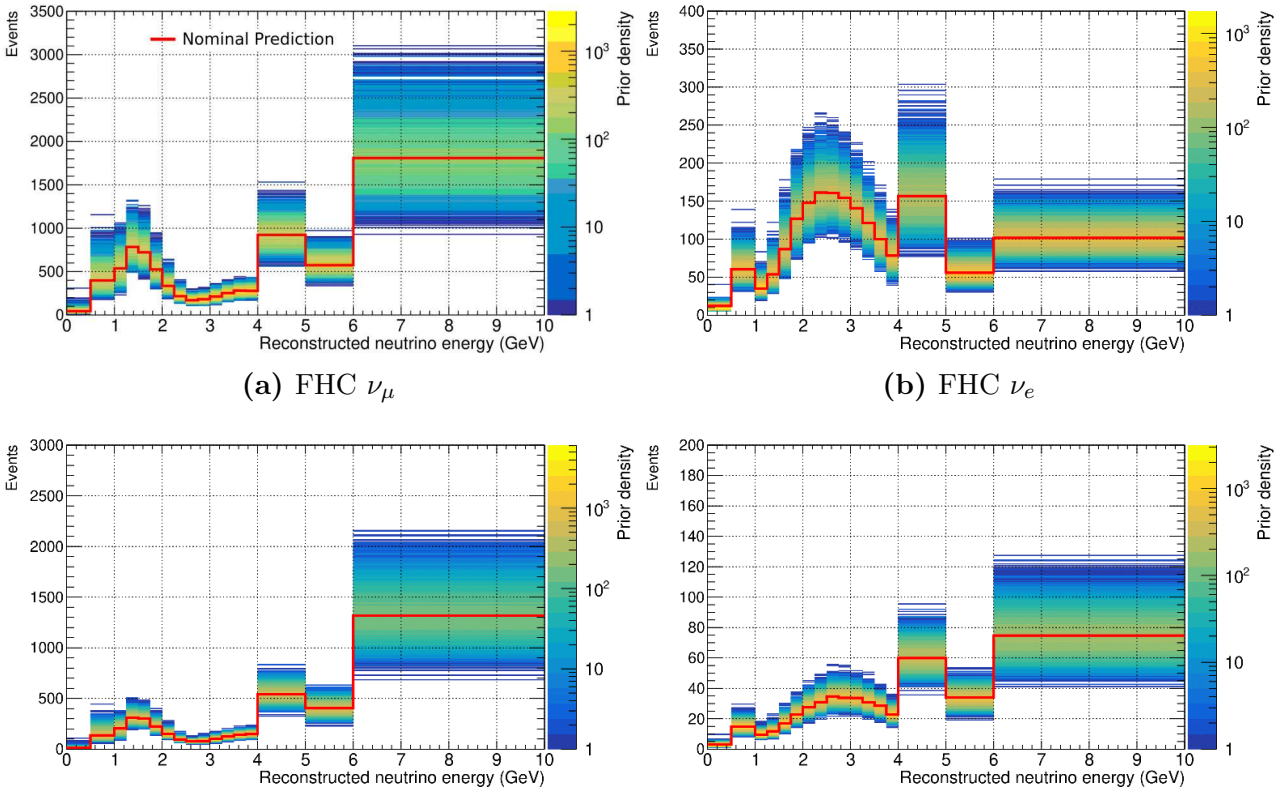

(a) FHC ν_μ **(b)** FHC ν_e

(c) RHC $\bar{\nu}_\mu$ **(d)** RHC $\bar{\nu}_e$

Figure 6.5: Oscillated prior predictive spectra as a function of reconstructed neutrino energy of the 4 samples at the DUNE FD: (a) FHC ν_μ, (b) FHC ν_e, (c) RHC $\bar{\nu}_\mu$ and (d) RHC $\bar{\nu}_e$. The z-axis shows the distribution of event rates in a given bin from 10,000 throws of the full systematic model. The nominal prediction is shown by the red line.

The statistical error is much smaller than the systematic error, even for samples with lower event rates (i.e. FHC ν_e, RHC $\bar{\nu}_e$) It should be emphasised that these are the *prior* systematic uncertainties and will become constrained during the fit.

6.5 Fitter Validation

Once the MC and systematic inputs are constructed, validations confirm that the inputs are correct and that the framework behaves as expected. The two main ways this was done were by comparing outputs with other analysis frameworks and performing scans of the likelihood given changes to the oscillation parameters.

	FHC ν_μ	FHC ν_e	RHC $\bar{\nu}_\mu$	RHC $\bar{\nu}_e$
Flux	7.3%	7.3%	7.5%	7.3%
Cross-section	9.3%	9.2%	7.8%	7.9%
Detector	2.3%	1.7%	2.7%	2.0%
Total Systematic	11.9%	11.7%	11.1%	10.7%
Statistical	1.0%	2.2%	1.5%	4.4%

Table 6.7: Prior percentage uncertainty on the total event rate for each FD sample split by uncertainty group.

6.5.1 Cross-check with previous Fitter

Given that the frameworks have different methods for constructing the likelihood space, we would not expect the final results to be identical. The likelihood constructed in this analysis uses more accurate systematic reweighting, making fewer assumptions. As well as this, here we have made a marginalised likelihood rather than a profiled likelihood, which can also introduce differences. However, the nominal event spectra predictions and spectra given variations to the systematic parameters should be more or less the same, provided an identical systematic implementation. Initially, the nominal event spectra with and without oscillations were checked. After the systematic uncertainties were implemented, the varied spectra for each systematic parameter were also validated and good agreement was found.

6.5.2 Log-Likelihood Scans

As mentioned in Section 4.4.1, marginalising a posterior over nuisance parameters may result in a difference in the HPD for a 1D parameter relative to the HPD in the full posterior. However, to ensure this is not a result of some bias in the fitting method, a log-likelihood scan can be performed. A fake dataset is created using some chosen values of the oscillation parameters and all systematic parameters are set to their prior values. Then one or two of the oscillation parameters are varied and for each point, the MC is reweighted and the LLH is evaluated for the fake dataset. All other parameters remain fixed. The LLH distribution, 1D or 2D, is

then plotted and the point with the bin with minimum LLH is the best-fit point. If the fitter is operating as expected, the best-fit point and the input parameters should agree almost perfectly (very small differences can occur as a result of the binning of the oscillation parameters).

Figure 6.6 and Figure 6.7 show the result of a LLH scan in $\sin^2\theta_{23} - \Delta m^2_{32}$ and $\sin^2\theta_{13} - \delta_{CP}$ respectively. The input value is shown with the red cross and the best-fit bin is shown with the blue triangle. In both cases, the input values and the best-fit point are in identical positions, showing the fitter is calculating the LLH correctly for these oscillation parameters.

6.6 Oscillation Sensitivities

This section will show the sensitivity of DUNE to the oscillation parameters given a fake dataset, specifically in this analysis an Asimov fake dataset is used. An Asimov fake dataset is an MC prediction without the addition of any statistical fluctuation. The MC is reweighted to a given set of oscillation and systematic parameters and the resulting event spectra for each sample are the "fake data". Since no statistical fluctuation is introduced into the dataset, it is expected that the Asimov point, the values of the oscillation parameters used to generate the Asimov data (also referenced as the "true value"), and the best-fit point lie in agreement. An Asimov fit allows us to estimate the sensitivity of an analysis - the contours show us the maximum precision with which we'd expect to be able to measure the oscillation parameters if their true values were the values of the Asimov point.

Statistics-error only results will be presented in Section 6.6.1, in which all systematic uncertainty parameters are fixed at their prior values during the Markov chain. The results including systematic uncertainties will then be presented in Section 6.6.2.

6.6.1 Statistics-error only

This section will show the results of a statistics-error only sensitivity study of DUNE using Asimov FD data given the DUNE 336 kt-MW-year exposure and generated with the oscillation parameter values shown in Table 6.1. In a statistics-error only study, the values of all systematic

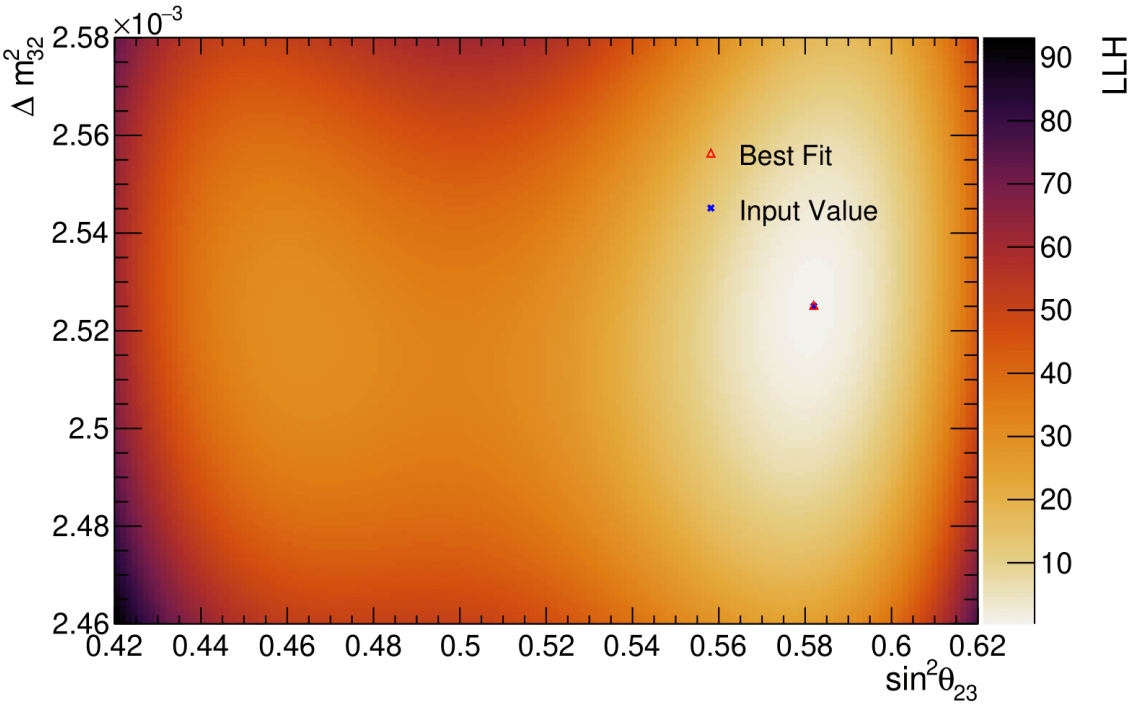

Figure 6.6: 2D LLH scan in $\sin^2\theta_{23} - \Delta m^2_{32}$ where the input data used the values $\sin^2\theta_{23} = 0.582$ and $\Delta m^2_{32} = 2.525 \times 10^{-3}\text{eV}^2$.

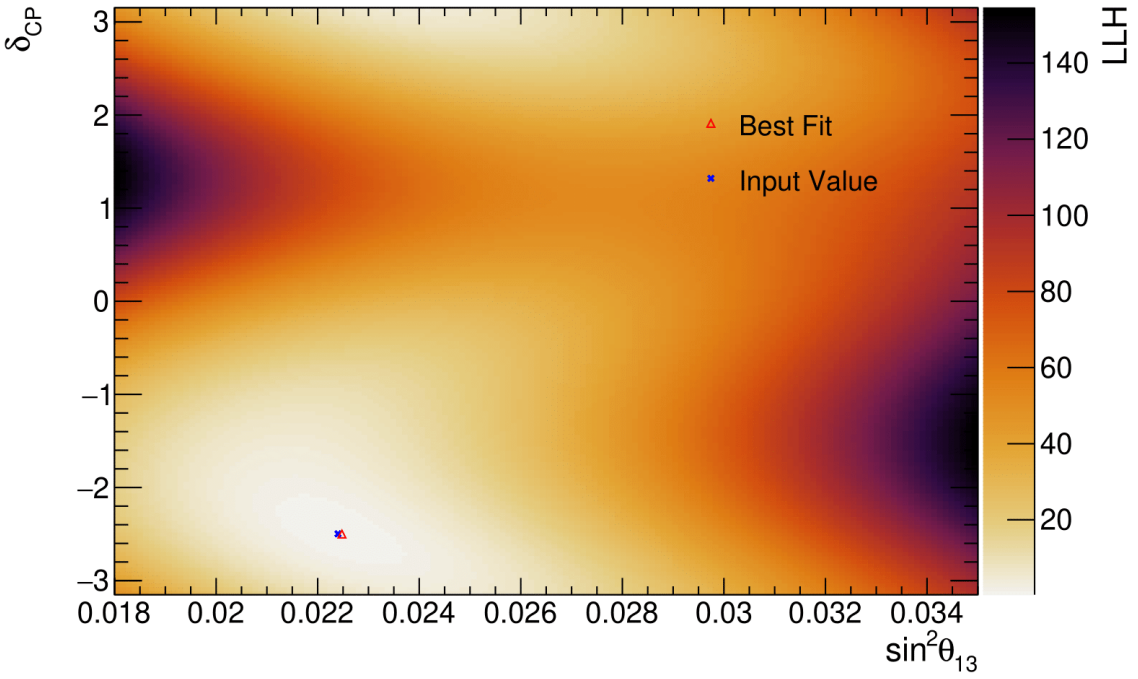

Figure 6.7: 2D LLH scan in $\sin^2\theta_{13} - \delta_{CP}$ where the input data used the values $\sin^2\theta_{13} = 0.0.0224$ and $\delta_{CP} = -2.498$.

parameters (flux, cross-section and detector) are fixed to their prior values. This essentially shows the maximum sensitivity to the oscillation parameters given the Asimov data, since there is no systematic uncertainty considered in the fit. The results presented have been obtained from a Markov chain containing 54 million steps after the burn-in was removed.

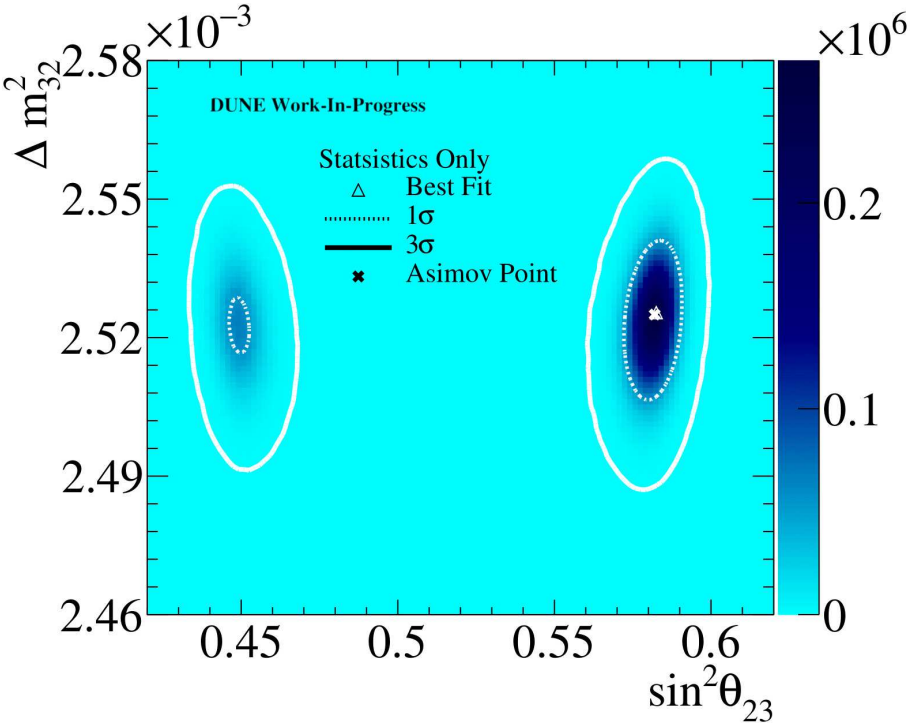

Figure 6.8: 2D posterior in $\sin^2\theta_{23} - \Delta m^2_{32}$ from a statistics-error only sensitivity study. The 1σ and 3σ credible intervals are shown in the dashed and bold lines respectively.

Figure 6.8 shows the posterior distribution in $\sin^2\theta_{23} - \Delta m^2_{32}$. The posterior only shows values of Δm^2_{32} since there is almost no posterior in the IO space. This shows the power of DUNE's sensitivity to the mass ordering. Both octants of θ_{23} are being explored yet the correct octant is preferred. The best-fit point and the Asimov point also lie in perfect agreement, suggesting no marginalisation effects are affecting the position of the HPD when integrating over the other oscillation parameters.

Figure 6.9 shows the 1D posterior distribution for both $\sin^2\theta_{23}$ and Δm^2_{32}. It is more clear from the 1D posterior that the correct θ_{23} octant is significantly preferred, with the wrong octant being excluded at 1σ. The posterior distribution Δm^2_{32} in IO, i.e. $\Delta m^2_{32} < 0$, contains very few Markov chain steps resulting in the unsmooth posterior shape. Of the 54 million steps

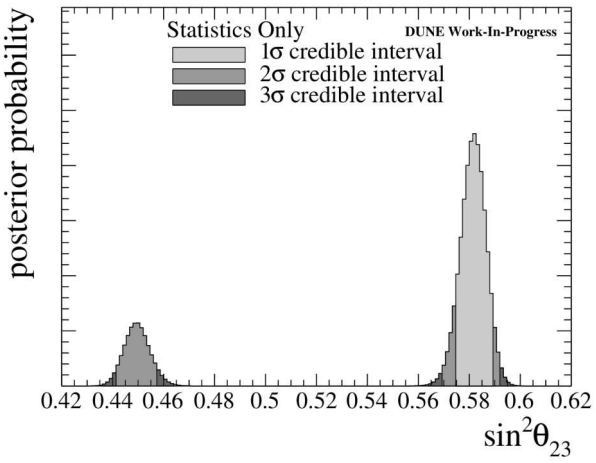

(a) $\sin^2\theta_{23}$

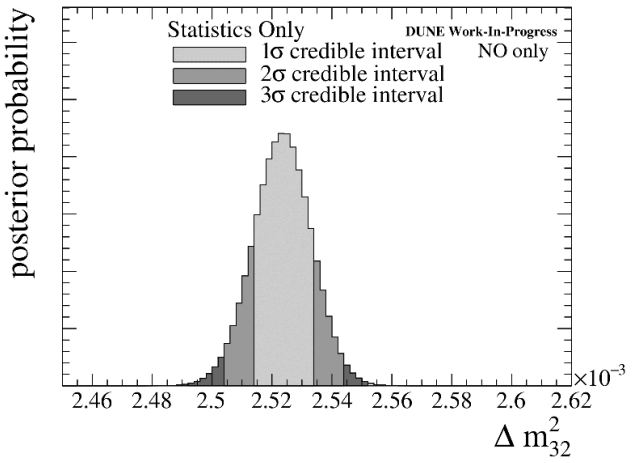

(b) Δm^2_{32} NO

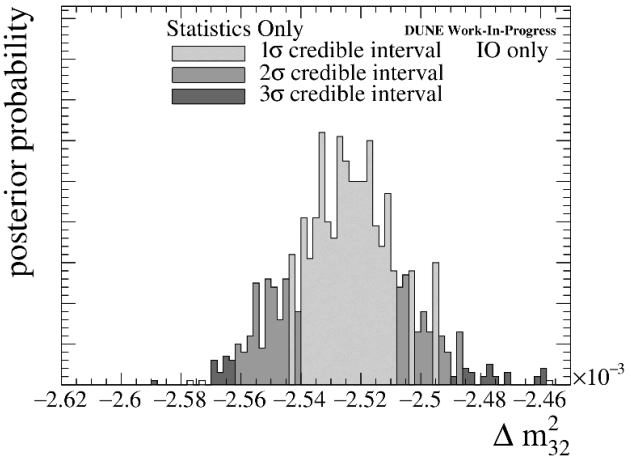

(c) Δm^2_{32} IO

Figure 6.9: 1D posterior in (a) $\sin^2\theta_{23}$, (b) Δm^2_{32} in NO and (c) Δm^2_{32} in IO from a statistics-error only sensitivity study. The 1σ, 2σ and 3σ credible intervals are shown in the light grey, grey and dark grey regions respectively.

obtained post-burn-in, just 1000 steps were accepted in the IO. To get an accurate relative probability of the two hierarchies, more steps are required since IO has not been fully explored yet.

Figure 6.10 shows the posterior distribution in $\sin^2\theta_{13} - \delta_{CP}$. A second region is being explored in $\sin^2\theta_{13}$, however, the majority of the posterior is still centred around the correct value. The best-fit point also lies perfectly on the Asimov point. Figure 6.11 shows the 1D posterior distribution for both $\sin^2\theta_{13}$ and δ_{CP}. The existence of a second peak in $\sin^2\theta_{13}$ becomes

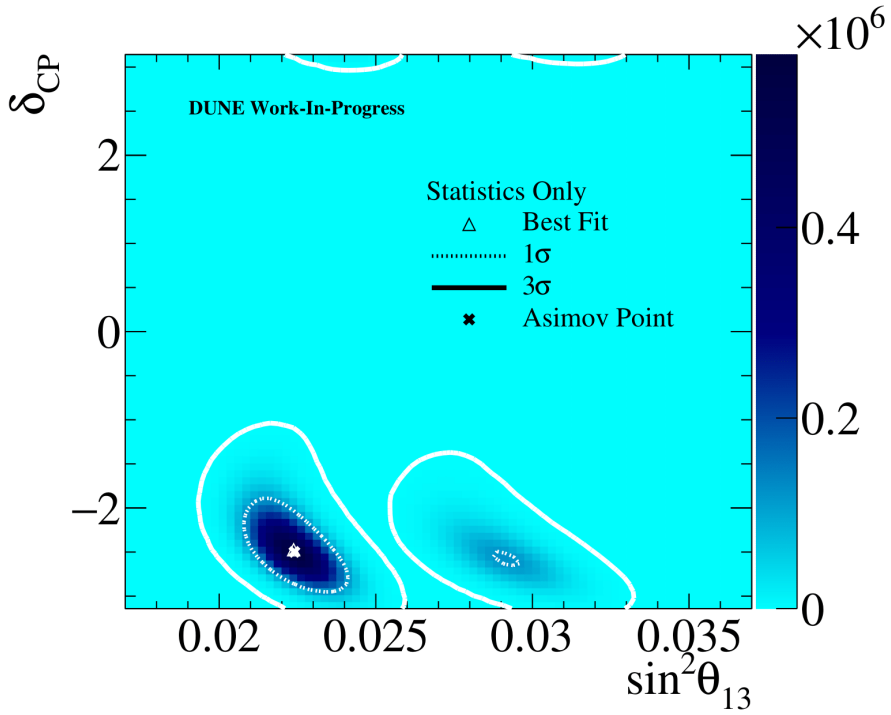

Figure 6.10: 2D posterior in $\sin^2\theta_{13} - \delta_{CP}$ from a statistics-error only sensitivity study. The 1σ and 3σ credible intervals are shown in the dashed and bold lines respectively.

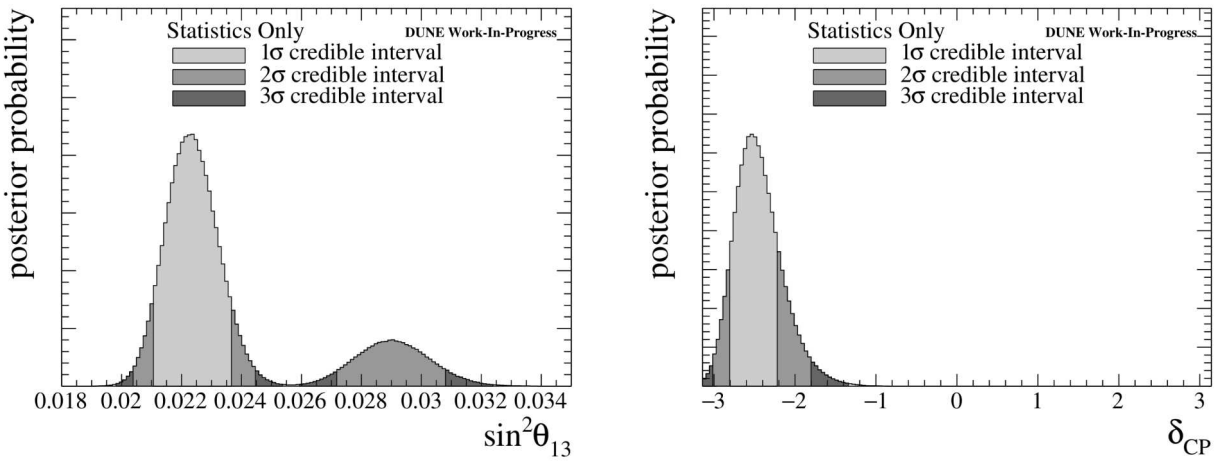

(a) $\sin^2\theta_{13}$ **(b)** δ_{CP}

Figure 6.11: 1D posterior in (a) $\sin^2\theta_{13}$, (b) δ_{CP} from a statistics-error only sensitivity study. The 1σ, 2σ and 3σ credible intervals are shown in the light grey, grey and dark grey regions respectively.

more clear, however, the region is still excluded at 1σ. The 1D δ_{CP} posterior shows that, given this Asimov point, the neutrino events obtained by the 336 kt-MW-year DUNE exposure are sufficient to exclude $\delta_{CP} = 0$ (i.e. CP conservation) given no systematic uncertainty.

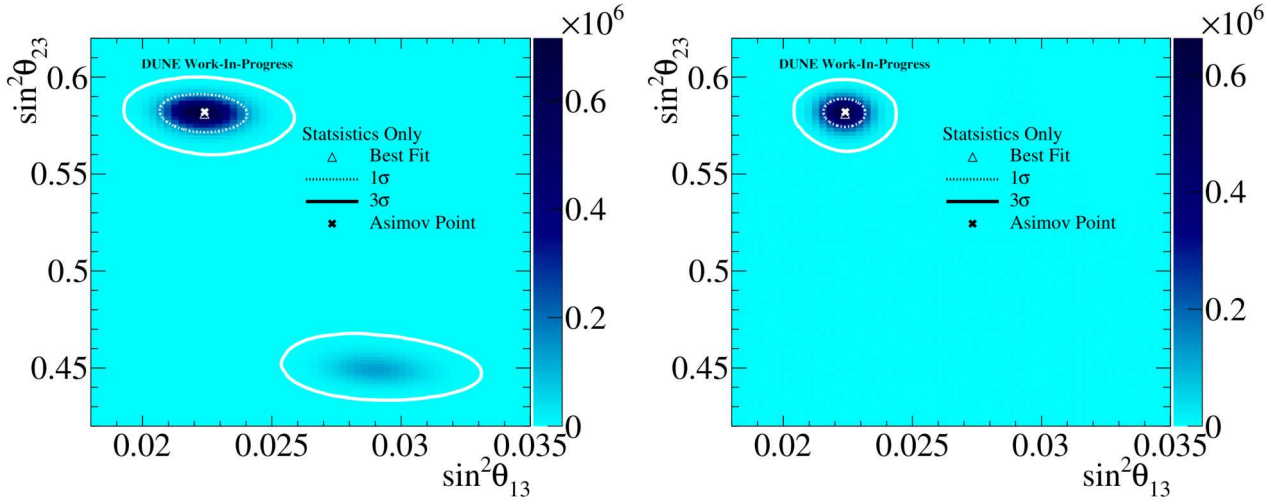

(a) Without Reactor Constraint **(b)** With Reactor Constraint

Figure 6.12: 2D posterior in $\sin^2\theta_{13} - \sin^2\theta_{23}$ (a) without reactor constraint and (b)with reactor constraint from a statistics-error only sensitivity study. The 1σ and 3σ credible intervals are shown in the dashed and bold lines respectively.

Plotting the 2D posterior distribution in $\sin^2\theta_{13}-\sin^2\theta_{23}$, as shown in Figure 6.12, shows the origin of the second peak in θ_{13} results from a correlation with the octant of θ_{23}. As mentioned in Section 4.4.3, a Markov chain can be reweighted based on a change of prior information on any given parameter. Figure 6.12 also shows the effect of applying a Gaussian constraint on θ_{13} (reactor constraint), using the mean and error from Table 6.1, on the $\sin^2\theta_{13} - \sin^2\theta_{23}$ distribution. Due to the correlation between θ_{13} and the octant of θ_{23}, applying the reactor constraint suppresses the posterior in the wrong octant.

6.6.2 With Systematic Uncertainties

The previous section shows the sensitivity of DUNE to the oscillation parameters, given the 336 kt-MW-year exposure and the Asimov point being investigated, if the cross-section, flux and detector models were understood perfectly. Of course, this is not possible, and sensitivity will be lost when the systematic uncertainties that exist in the experiment are taken into

consideration. This section presents the results of sensitivity results where the systematic parameters described in Section 6.3 are allowed to vary during the fit. These results will be referred to as FD-only since they contain the full systematic treatment but contain just FD data. The results presented have been obtained from a Markov chain containing 218 million steps after the burn-in was removed.

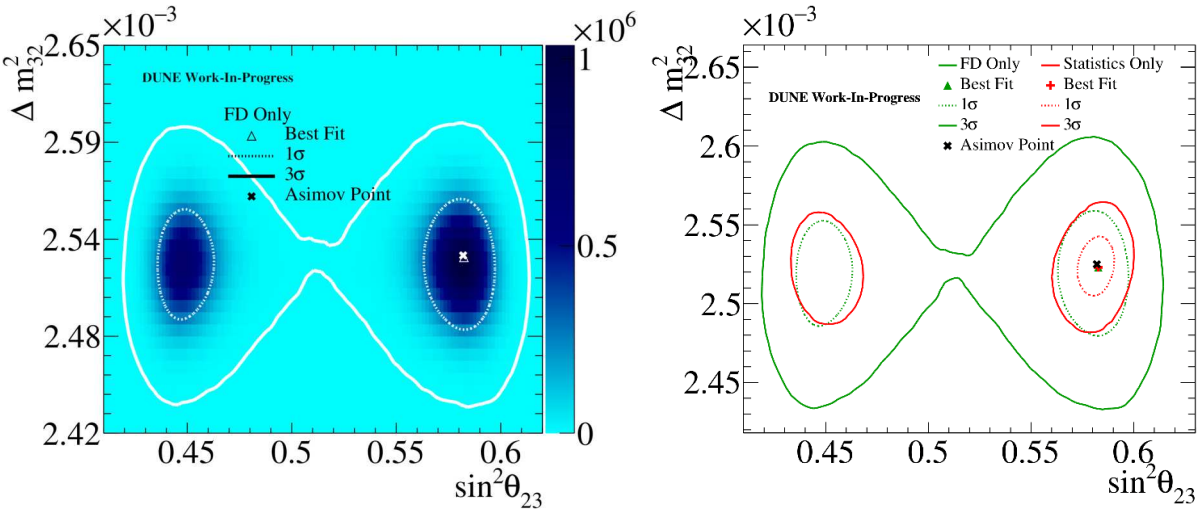

Figure 6.13: 2D posterior in $\sin^2\theta_{23} - \Delta m^2_{32}$ from an FD-only sensitivity study and comparison of the 1σ and 3σ credible intervals with the statistics-error only study. The 1σ and 3σ credible intervals are shown in the dashed and bold lines respectively.

Figure 6.13 shows the posterior distribution in $\sin^2\theta_{23} - \Delta m^2_{32}$ as well as the credible intervals for both FD-Only and statistics-error only sensitivity studies. A significant difference can be seen between the two posteriors, with a much larger posterior existing in the wrong octant when compared with Figure 6.8. The 3σ contour now connects both octants, with maximal $\sin^2\theta_{23}$ ($\sin^2\theta_{23} = 0.5$) no longer being excluded. In both octants, the 1σ interval for FD-only is roughly the same size as the 3σ interval statistics-error only. The effect of the systematic uncertainty can also be seen in the 1D posterior, as shown in Figure 6.14. The height of the peak in the wrong octant is now comparable to that of the correct octant. There is also a non-zero posterior between both octant peaks. The Δm^2_{32} 1D posterior is visibly much wider. However, even when including the systematic uncertainty, the number of steps in the inverted ordering is still negligible.

Figure 6.15 shows the $\sin^2\theta_{13} - \delta_{CP}$ posterior as well as the comparison of FD-only and

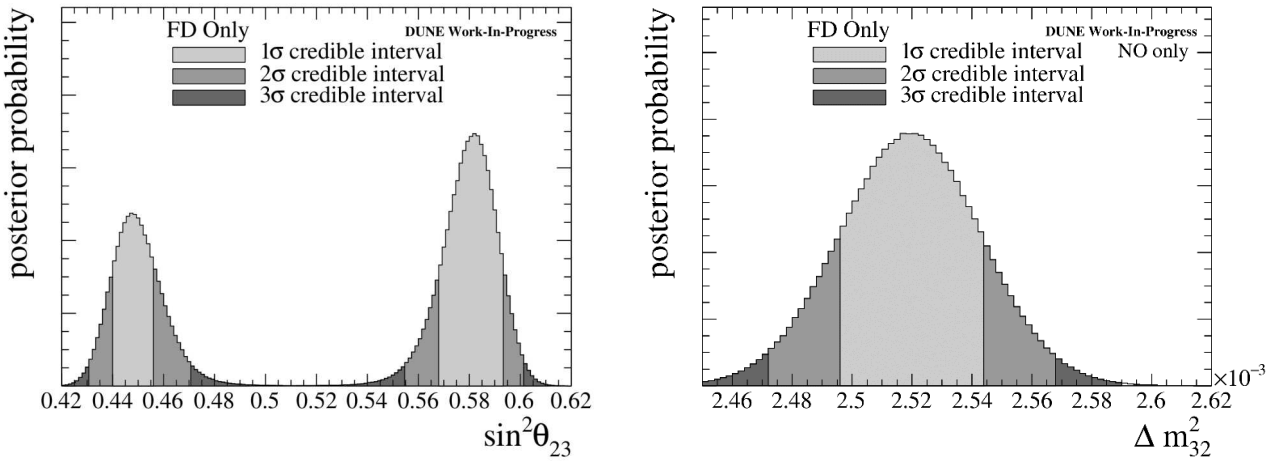

(a) $\sin^2\theta_{23}$ **(b)** Δm^2_{32} NO

Figure 6.14: 1D posterior in (a) $\sin^2\theta_{23}$ and (b) Δm^2_{32} in NO from an FD-only sensitivity study. The 1σ, 2σ and 3σ credible intervals are shown in the light grey, grey and dark grey regions respectively.

statistics-error only credible intervals. The posterior distribution in the second peak in θ_{13} has increased significantly. Comparing the intervals, there is a significant enlargement in the 3σ contour which is now fully connected around both θ_{13} peaks. There is a more significant shift towards 0 in δ_{CP} than across the $\pi/-\pi$ boundary. Despite this, CP-conservation is still excluded at 3σ

Some of the effects of the systematic uncertainty on the sensitivity of the mixing angles can be mitigated by applying the reactor constraint. The 2D posteriors for both $\sin^2\theta_{23}-\Delta m^2_{32}$ and $\sin^2\theta_{13}-\delta_{CP}$ after the reactor constraint is applied is shown in Figure 6.16. In $\sin^2\theta_{13}-\delta_{CP}$, the reactor constraint removes the second posterior region, significantly shrinking the both credible intervals. In $\sin^2\theta_{23}-\Delta m^2_{32}$, the reactor constraint suppresses the lower θ_{23} octant, as a result of the correlation with θ_{13}, excluding it from the 1σ credible interval. There is minimal effect of the reactor constraint on Δm^2_{32} and δ_{CP}.

As mentioned in subsubsection 2.3.2, the Jarlskog invariant, J_{CP}, indicates the magnitude of CP violation as a result of the PMNS mixing angles and δ_{CP}. Although not one of the parameters in the Markov chain, the J_{CP} distribution can be found by evaluating Eq 2.36 at each step in the Markov Chain. Figure 6.17 shows the J_{CP} distribution from the FD-

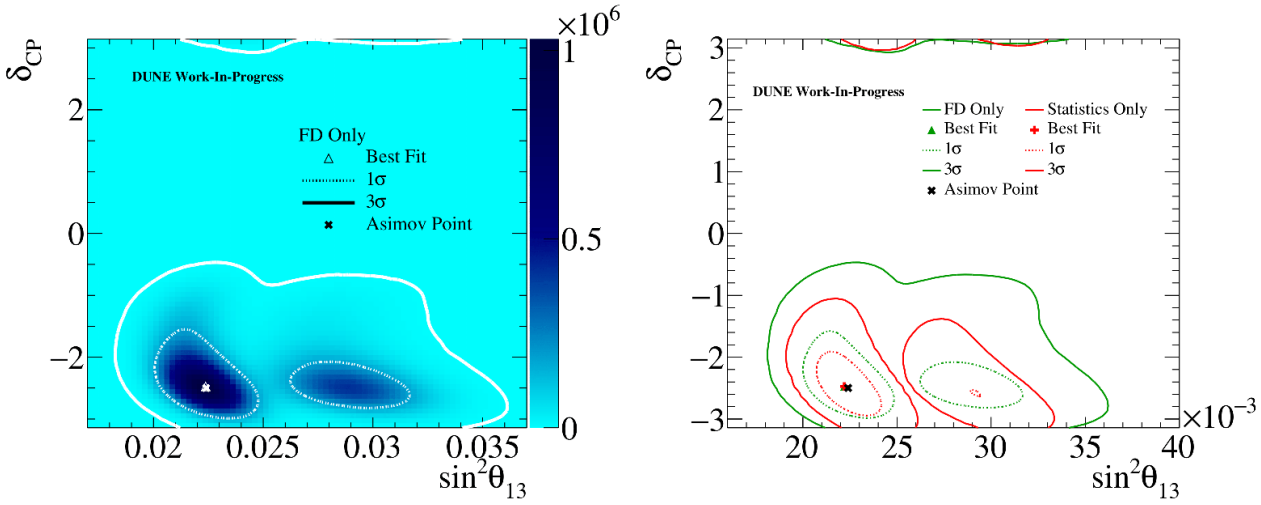

Figure 6.15: 2D posterior in $\sin^2\theta_{13} - \delta_{CP}$ from an FD-only sensitivity study and comparison of the 1σ and 3σ credible intervals with the statistics-error only study. The 1σ and 3σ credible intervals are shown in the dashed and bold lines respectively.

only sensitivity study. The edge of the 3σ contour lies at $J_{CP} = 0$ which corresponds to CP conservation.

Although the aim of this analysis is to make estimates of the oscillation parameters, it can also be helpful to examine how the systematic parameters are affected by the fit. The distributions of the systematic parameters can affect the marginalised distribution of the oscillation parameters.

Figure 6.18 shows the 1σ posterior uncertainties and central values from the FD-only sensitivity study relative to the prior 1σ uncertainties and central values. For the vast majority of systematics, the posterior uncertainty corresponds to the HPD 1σ credible interval, as explained in Section 4.4.2. This allows for asymmetrical error bars to be plotted. For some parameters whose effect is significantly asymmetrical or have ranges restricted such that the 1σ prior uncertainty covers most of the allowed values, the HPD credible intervals are not a suitable uncertainty estimate. In these cases, a Gaussian distribution is fitted to the marginalised posterior and the standard deviation is extracted.

Overall, there is minimal constraint coming from the FD samples for the majority of systematic parameters, as expected given the limited number of neutrino events. The flux parameters

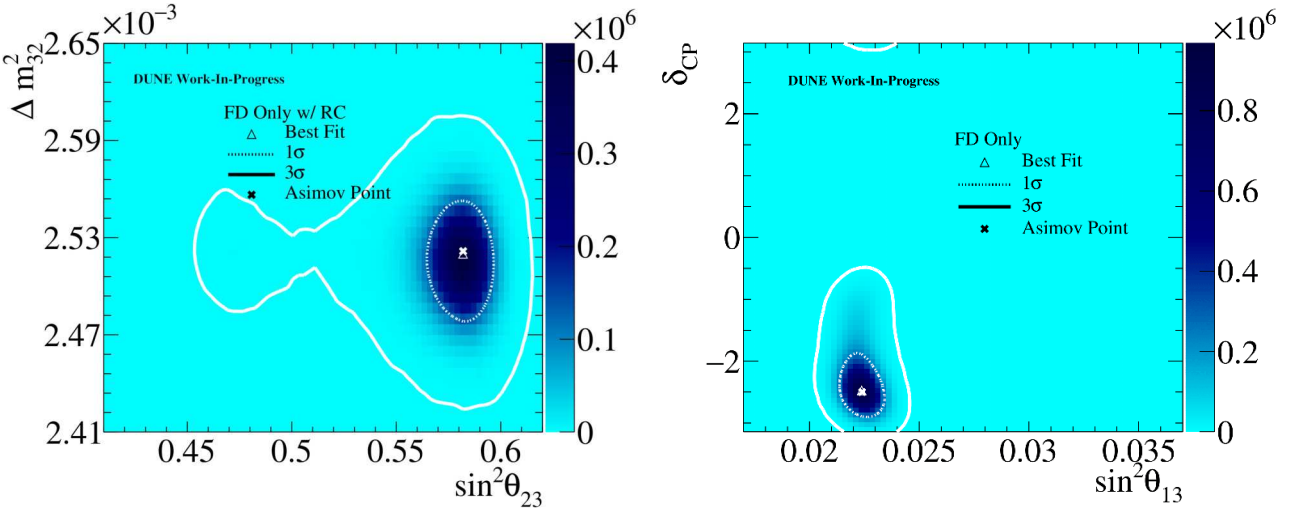

(a) $\sin^2\theta_{23} - \Delta m^2_{32}$ **(b)** $\sin^2\theta_{13} - \delta_{CP}$

Figure 6.16: 2D posterior in (a) $\sin^2\theta_{23} - \Delta m^2_{32}$ and (b) $\sin^2\theta_{13} - \delta_{CP}$ from an FD-only sensitivity study after applying the reactor constraint. The 1σ and 3σ credible intervals are shown in the dashed and bold lines respectively.

(i.e. "b_") have a more consistent level of constraint caused by the correlations between parameters. The constraints on the uncorrelated cross-section parameters vary widely, with the 2p2h energy dependence parameters receiving the strongest constraint. Generally, the FD detector systematics receive very weak constraints, with the exception of the parameter controlling $1/\sqrt{E}$ dependence of the neutron energy scale. There is some variation in the HPD relative to the prior central value, however, the HPD is well within the 1σ range for all parameters. For some parameters, the posterior 1σ lies just outside the prior 1σ. Generally, this is only true in either the positive or negative direction due to the asymmetrical errors, however, for a couple of parameters, the total width of the posterior 1σ is slightly larger than the prior 1σ. This is due to anti-correlations, likely between cross-section and flux parameters, that are not included in the prior description of the parameters. This allows systematics to drift outside of their prior ranges provided other parameters drift in the opposite direction.

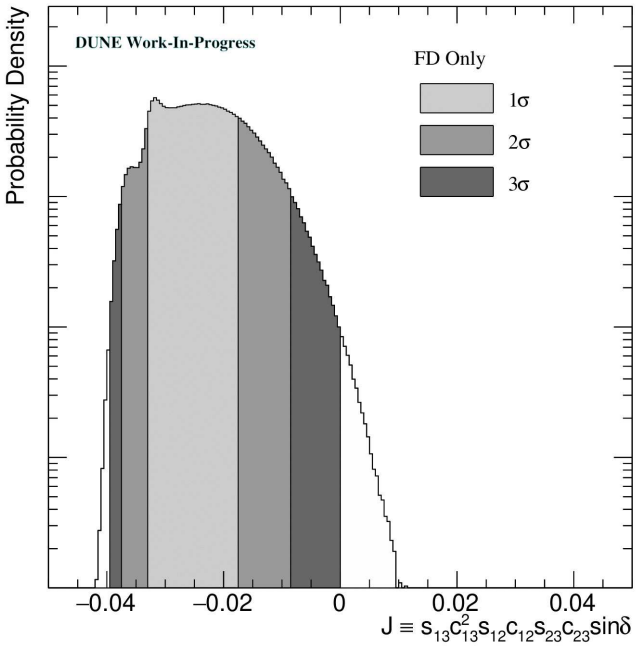

Figure 6.17: 1D posterior in J_{CP} from a FD-only sensitivity study. The 1σ and 3σ credible intervals are shown in the dashed and bold lines respectively.

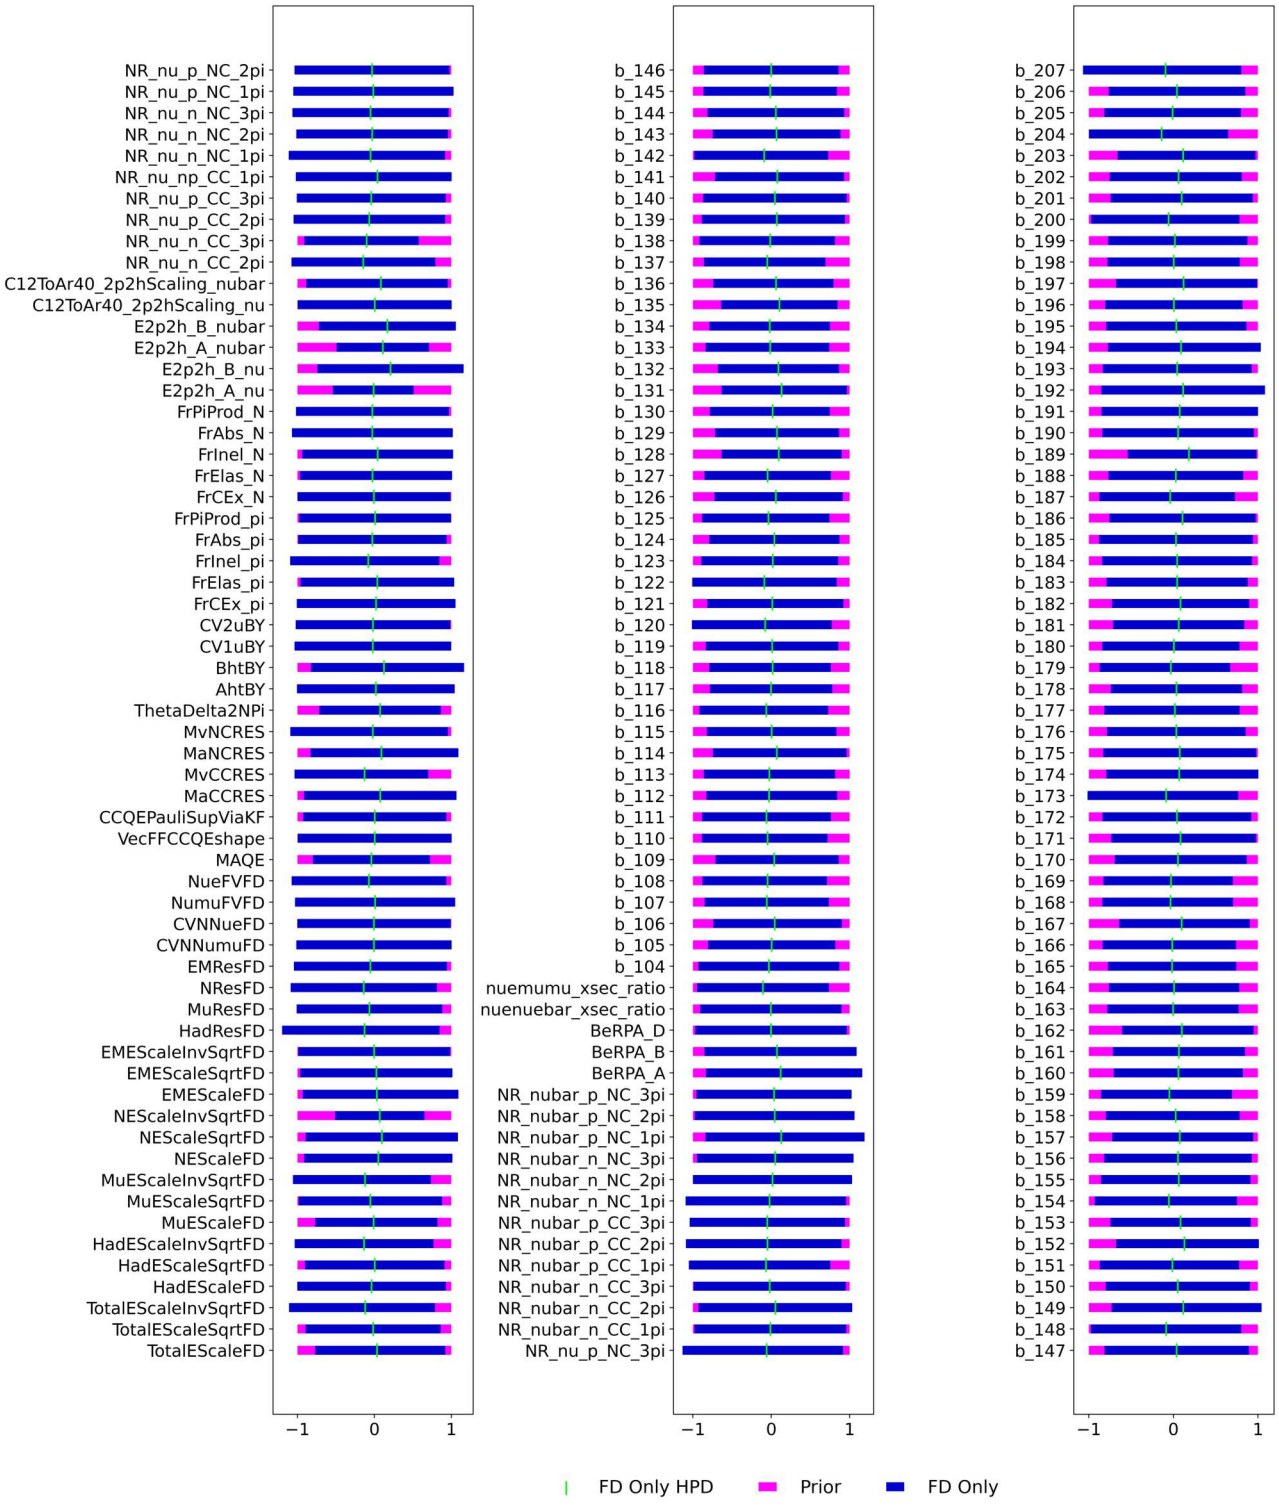

Figure 6.18: The ratio of posterior to prior 1σ uncertainties for all systematic parameters included in the FD-only sensitivity study. The prior uncertainty is shown in magenta and the FD-only posterior uncertainty is shown in blue. The vertical marks show the HPD point relative to the prior central value as a fraction of the prior uncertainty.

Chapter 7

Joint Far and Near Detector Sensitivity

In the previous section, the results from a statistics-error only fit and a fit containing systematic uncertainties were compared, showing a significant reduction in sensitivity when including systematic uncertainty. As mentioned in Figure 3.8, long-baseline experiments use near detectors to constrain the systematic uncertainties, since the event rate at the ND is not affected by oscillations, and is simply a product of the systematics models: flux, cross-section and detector. Comparing the sensitivities with and without the addition of ND samples provides vital justification for the design and construction of DUNE's ND complex. As mentioned in Chapter 3, DUNE aims to make world-leading measurements of the oscillation parameters. This section helps illustrate the necessity of the ND complex in achieving DUNE's physics goals.

As with the FD-Only study, the FD + ND study uses the same MC inputs as was used in the TDR analysis [2], with significant differences in the implementation of systematic uncertainties. As mentioned in Chapter 6, this analysis takes a more comprehensive approach to systematic uncertainties. However, the primary difference is that this is a Bayesian analysis, whereas the TDR analysis was done with a Frequenist approach. As in Chapter 6, this allows us to study the behaviour of nuisance parameters in detail without carrying out further fits. In addition, there is great benefit in cross-validating the sensitivity results using different statistical approaches.

This section presents the results of an oscillation sensitivity study using FD and ND fake

datasets simultaneously. This section also discusses some of the current difficulties in performing a fit to future long-baseline ND data, given unprecedented numbers of neutrino events, and how a realistic uncertainty model is required for reasonable results.

7.1 ND Design

At the time of the development of the simulation inputs used in this analysis, the final design of the near detector was not confirmed. A simpler design was used in this simulation compared with what is described in Figure 3.8. Only samples of neutrino interactions in LAr are explicitly used in this analysis. The ND consists of a LAr detector, a simplified version of the detector described in Figure 3.8 but with similar dimensions. Although a sample of events on GAr is not included, the ND simulation includes a GAr detector downstream of the LAr detector which we use to measure the momentum of particles exiting the LAr. The GAr detectors ECAL is assumed to consist of 30 layers of alternating planes of CH (5mm) and Cu (2mm). The pressure vessel is assumed to be titanium with a thickness of 3 cm. A solenoid magnet was included with an inner radius of 320 cm, with a cut-out section facing the LAr detector to minimise the material between them. SAND is not planned to be coupled with the LAr detector and hence is not included in these simulations. It is worth noting that the official design of the DUNE GAr detector has changed significantly since development of this simulation.

7.2 Event Selection

In this analysis, the aim of the ND is to produce a selection of CCν_μ and CC$\bar{\nu}_\mu$ events that originate in the LAr detector that can constrain the systematic models. Unlike the FD, pattern recognition and full reconstruction software were not yet available for the ND. A parameterised reconstruction is used based on the true energy deposits of particles in the detector.

To produce a sample of CCν_μ and CC$\bar{\nu}_\mu$ events, a μ-like track must be identified. Both muons and charged pions are initially evaluated as potential muon candidates. For the track

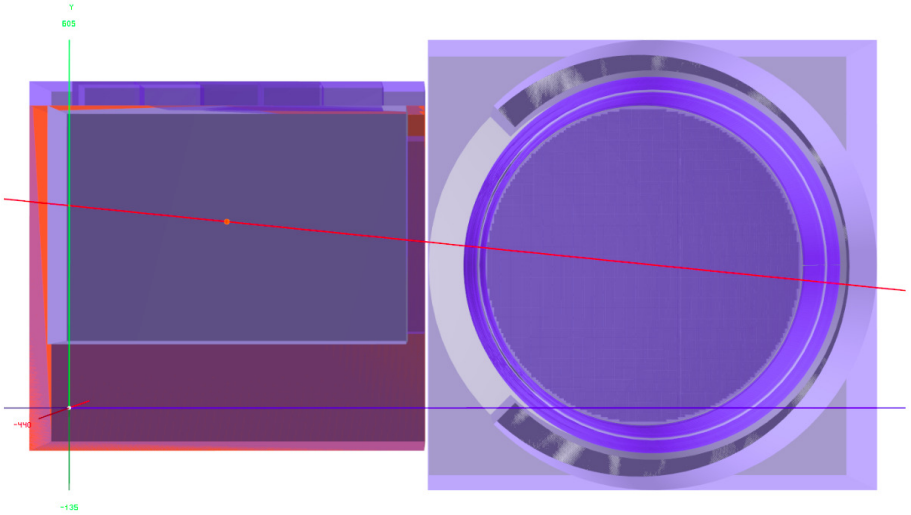

Figure 7.1: Horizontal view of the ND design used in this analysis. Neutrinos are incident from the left, first passing through and interacting in the LAr detector. Exiting muons are captured by the magnetised GAr detector downstream.

to be considered a muon, the track length must be at least 1 m and the mean energy deposit per cm must be less than 3 MeV. The track length cut essentially imposes a muon energy threshold of 200 MeV, but significantly suppresses NC background events which contain non-interacting charged pions. The charge of muons which exit the LAr detector and enter the GAr detector can be determined by the curvature of the track. However, the charge of muons which stop in the LAr detector cannot be determined. In FHC mode, the wrong-sign background is sufficiently small such that all muons are determined to be μ^-. In RHC mode, the background is much larger due to the muon kinematic differences between ν_μ and $\bar{\nu}_\mu$, so the contamination in the sample must be mitigated. The majority of μ^- that stop in LAr will be captured by an Ar nucleus, compared with μ^+ which does not get captured. The Michel electron, an electron from a μ decay, is not observed when the muon is captured in a nucleus. Hence, only events containing a Michel electron are selected in RHC mode, which helps to suppress the wrong-sign background contamination.

Events are also selected based on the location of the interaction in the detector. Events are required to originate in the fiducial volume of the LAr detector, which excludes 50 cm from the sides and upstream face and 150 cm from the downstream face. As well as this, a hadronic

veto region is defined to be 30 cm from each side of the detector. If more than 30 MeV of visible hadronic energy is measured in this veto region the event is excluded since it's likely a significant amount of energy was missed which leads to a poor reconstructed energy estimate.

The selection cuts for both FHC and RHC modes are summarised as follows:

- Event originates within the fiducial volume
- μ-like track identified
- μ track is contained in LAr or enters GAr detector
- μ track that is contained is followed by a Michel electron (**RHC only**)
- Visible Hadronic energy in veto region ¡ 30 MeV

7.3 ND MC Prediction

Similarly to the FD, a prediction of the ND samples can be produced by applying the cuts described in Section 7.2 to the ND MC. The ND MC uses the same flux prediction as FD and the interaction is simulated using the same event generator, GENIE. A description of the flux and cross-section model can be found in Section 6.6.2. The propagation of neutrino interactions through the detectors is simulated using a Geant-4 model.

As mentioned in Chapter 6, the predicted event rates used here correspond to an exposure of 336 kt-MW-yrs. Unlike the FD, which is only binned in reconstructed energy, the ND samples are binned in both reconstructed energy and reconstructed inelasticity. Inelasticity, y, is a measure of the ratio of observed energy which is hadronic, E_{Had}/E_{ν}. The predicted event distributions, broken down by interaction mode, as a function of reconstructed energy and reconstructed inelasticity for both ND samples given a 336 kt-MW-yr exposure are shown in Figure 7.2. Table 7.1 shows the total event rate for both samples and split by the interaction mode. As we'd expect, CC DIS, CC RES, CC QE and CC 2p-2h make up the vast majority of events in all distributions. The distribution in reconstructed energy is very similar to the

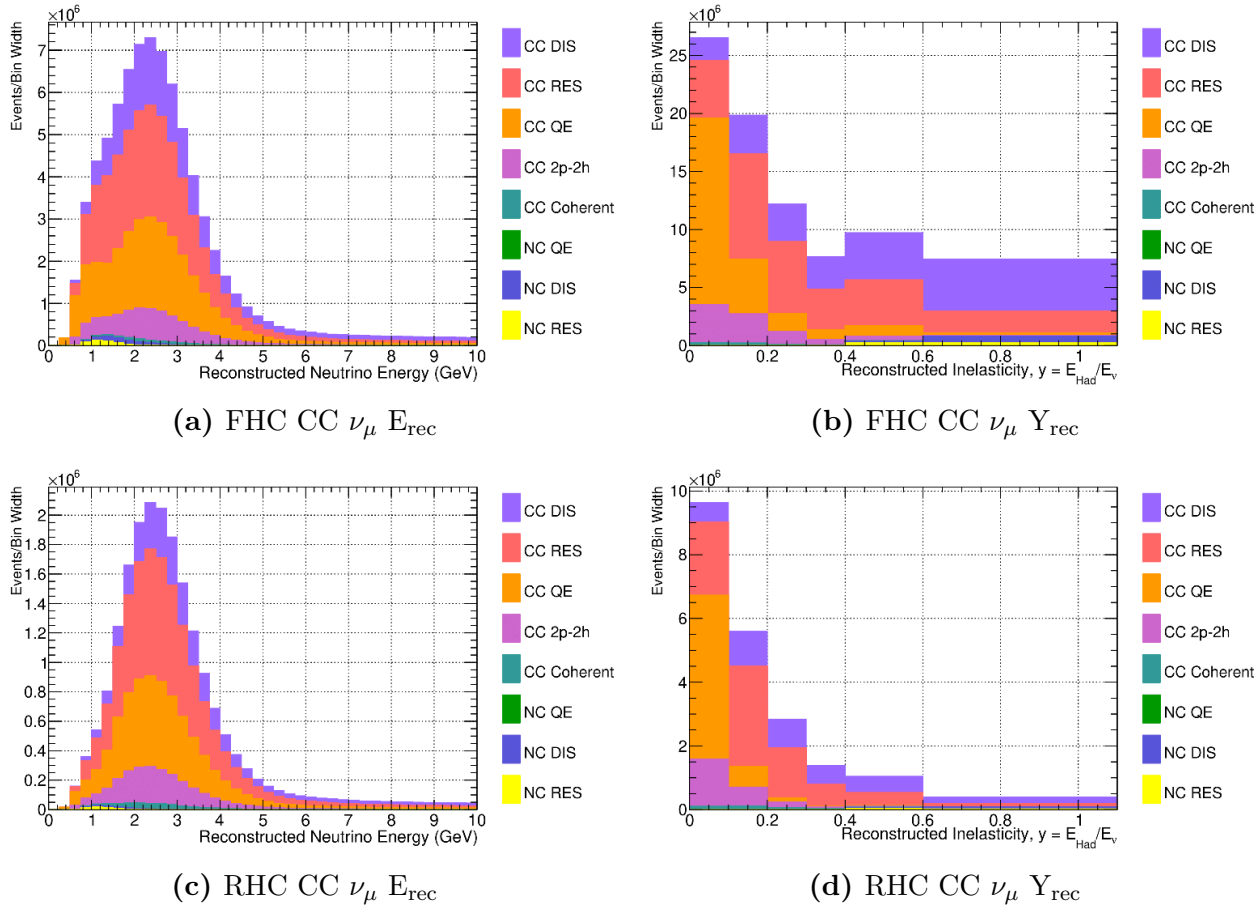

(a) FHC CC ν_μ E_{rec} **(b)** FHC CC ν_μ Y_{rec}

(c) RHC CC ν_μ E_{rec} **(d)** RHC CC ν_μ Y_{rec}

Figure 7.2: Event spectra predictions as a function of reconstructed neutrino energy and reconstructed inelasticity of the FHC CC ν_μ and RHC CC ν_μ samples at the DUNE ND. The predicted event spectra are also broken down by the different interaction modes.

unoscillated distribution for FD. The majority of events have low reconstructed inelasticity, and this is a result of the event selection which identifies muon tracks and minimal hadronic energy in the outer regions of the detector. CC DIS dominates the higher inelasticity events, in which hadronic showers are emitted from the nucleus interaction.

The shape of the distributions between FHC and RHC are considerably different, compared with the FD spectra which are more similar. In reconstructed energy, there is a faster rise in the distribution in FHC compared with RHC. The CC QE and CC 2p-2h distributions in FHC at low energy have a visible 'shoulder' which is not apparent in the RHC spectrum. This is caused by the differences in the muon kinematics between neutrinos and antineutrinos which in turn affects the acceptance. Muons produced in antineutrino interactions are more forward-

	FHC CC ν_μ	**RHC CC $\bar{\nu}_\mu$**
CC DIS	15,173,846	4,807,865
CC RES	24,508,835	10,996,575
CC QE	20,334,032	8,128,637
CC 2p-2h	6,528,339	3,162,026
CC Coherent	564,024	469,903
NC QE	7554	883
NC DIS	626,792	115,189
NC RES	399,838	85.465
Sample totals	**68,143,262**	**27,766,542**

Table 7.1: Event rate predictions for the FHC CC ν_μ and RHC CC $\bar{\nu}_\mu$ samples at the DUNE ND broken down by interaction mode and corresponding to a 336 kt-MW-yr exposure.

going than those produced in neutrino interactions. At low energies, these wide-angle muons produced neutrino interactions are contained, but at higher energies, they exit the LAr detector but are not picked up by the GAr detector. This effect causes a dip in the CC QE and CC 2p2h distributions. This effect is also not apparent in higher Q^2 processes (i.e. CC RES and CC DIS) since the relationship between neutrino energy and muon kinematics is more smeared.

Figure 7.3 shows the predicted event distributions broken down into true signal, wrong-sign background and intrinsic ν_e components. For FHC, ν_μ events are signal and $\bar{\nu}_\mu$ are background and vice versa for RHC. There is a larger contribution of wrong-sign ν_μ in the RHC relative to the FHC, despite the Michel electron tag. The intrinsic ν_e contribution is roughly the same for both FHC and RHC samples, as expected. There is an increase in flux contamination, wrong-sign and intrinsic ν_e, as reconstructed inelasticity increases. The signal event rates also decrease, resulting in a much larger flux contamination contribution. This is even more significant for the intrinsic ν_e since DIS makes up a larger portion of ν_e interactions.

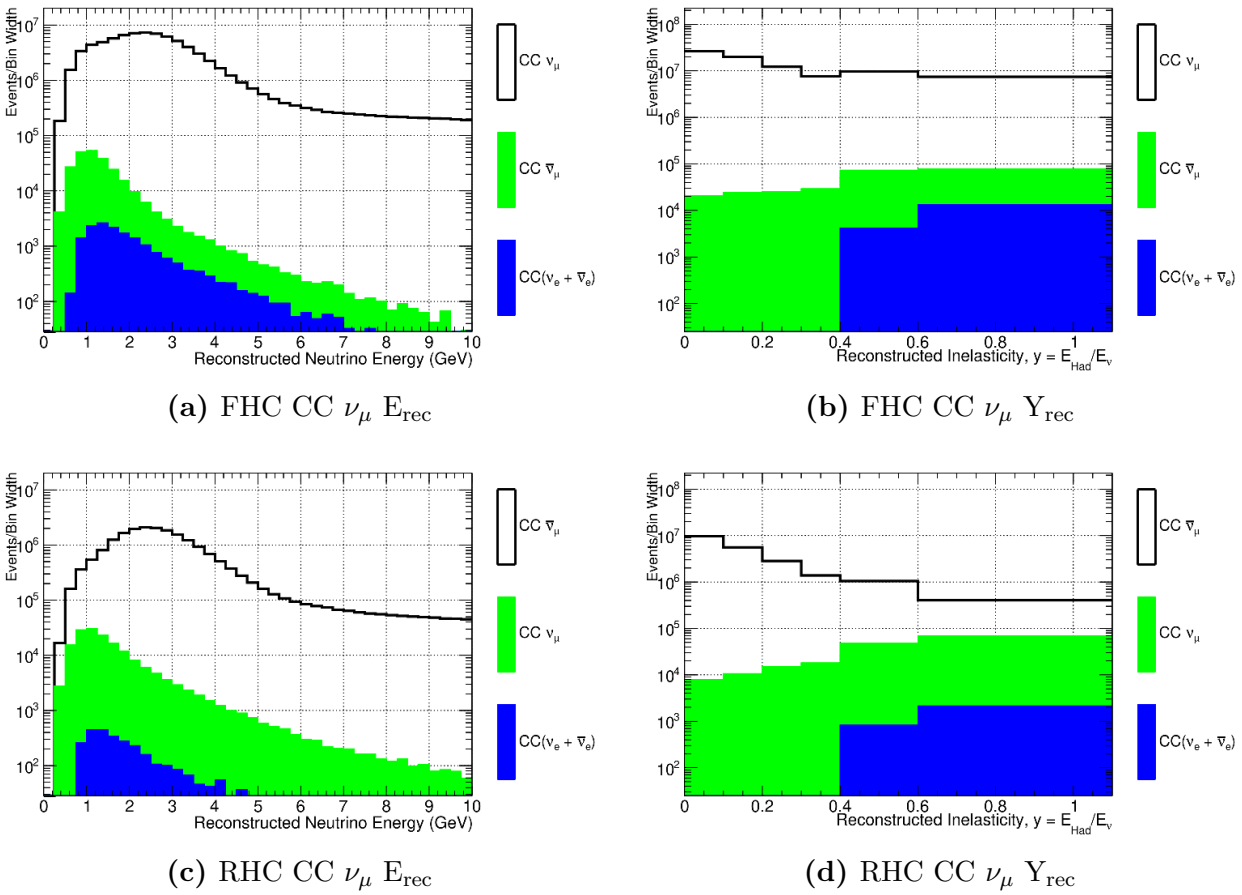

(a) FHC CC ν_μ E_{rec}

(b) FHC CC ν_μ Y_{rec}

(c) RHC CC ν_μ E_{rec}

(d) RHC CC ν_μ Y_{rec}

Figure 7.3: Event spectra predictions as a function of reconstructed neutrino energy and reconstructed inelasticity of the FHC CC ν_μ and RHC CC ν_μ samples at the DUNE ND using log-scale. The predicted event spectra are also broken down into true signal, wrong-sign background and intrinsic electron neutrino components.

7.4 ND Systematic Uncertainties

The systematic model used for the ND samples is largely the same as for the FD samples described in Section 6.6.2. The flux systematics are implemented in the same way, in which bin-content normalisation uncertainties, categorised beam mode, true neutrino flavour and energy bin, are included. Figure 6.4 shows the flux covariance matrix used in this analysis. As one might expect, the same cross-section model was also used for both near and far detector samples.

As for the detector uncertainties, the energy scale uncertainties and resolution uncertainties were implemented in the same method as for FD, with the same prior uncertainties as shown

in Table 6.6. The FD and ND energy scale and resolution parameters were kept uncorrelated. However, since the LAr detector is much smaller than the FD modules, the ND has a very different acceptance. At the FD, it is assumed that the acceptance does not change significantly as a function of the event kinematics. At the ND, there are regions of both the muon and hadron kinematic phase space that have comparatively very poor acceptance. With muons, low acceptance occurs generally at low lepton energies, where events are mis-reconstructed as NC background, or at high transverse momentum, where the muon exits the sides of the LAr detector without entering the GAr detector. Hadronic acceptance decreases with hadronic energy, as the likelihood of events depositing $\geq$ 30 MeV in the veto region increases. The uncertainty in both muon and hadronic acceptance is evaluated by considering the change in acceptance as a function of the relevant event kinematics.

7.5 ND Constraining Power

Table 7.1 shows an extremely large number of events in both samples, $\sim$ 100 million in total, across a 336 kt-MW-yr exposure. These are unprecedented numbers of neutrino events. The T2K experiment has collected $\sim$ 200,000 events in all ND samples across $\sim$ 10 years of running [56]. These large samples will allow significant constraints on the systematic models, improving the sensitivity of the experiment on the oscillation parameters. Previous and current oscillation experiments are limited by the number of neutrino events they can collect at the near and far detectors. This means, although they can constraint their systematic uncertainties with near detector data, the effect of these constraints are limited by the statistical uncertainties at the far detector. For DUNE, the systematic uncertainty will dominate the statistical uncertainty. This means the choice and implementation of systematic uncertainty models will have significant effects on the results of oscillation analyses. Therefore, it is imperative that our systematic models reflect the true uncertainties in the experiment as closely as possible.

One important feature which is necessary to capture in the systematic uncertainty model is degeneracy. In this context, degeneracy refers to situations where different combinations

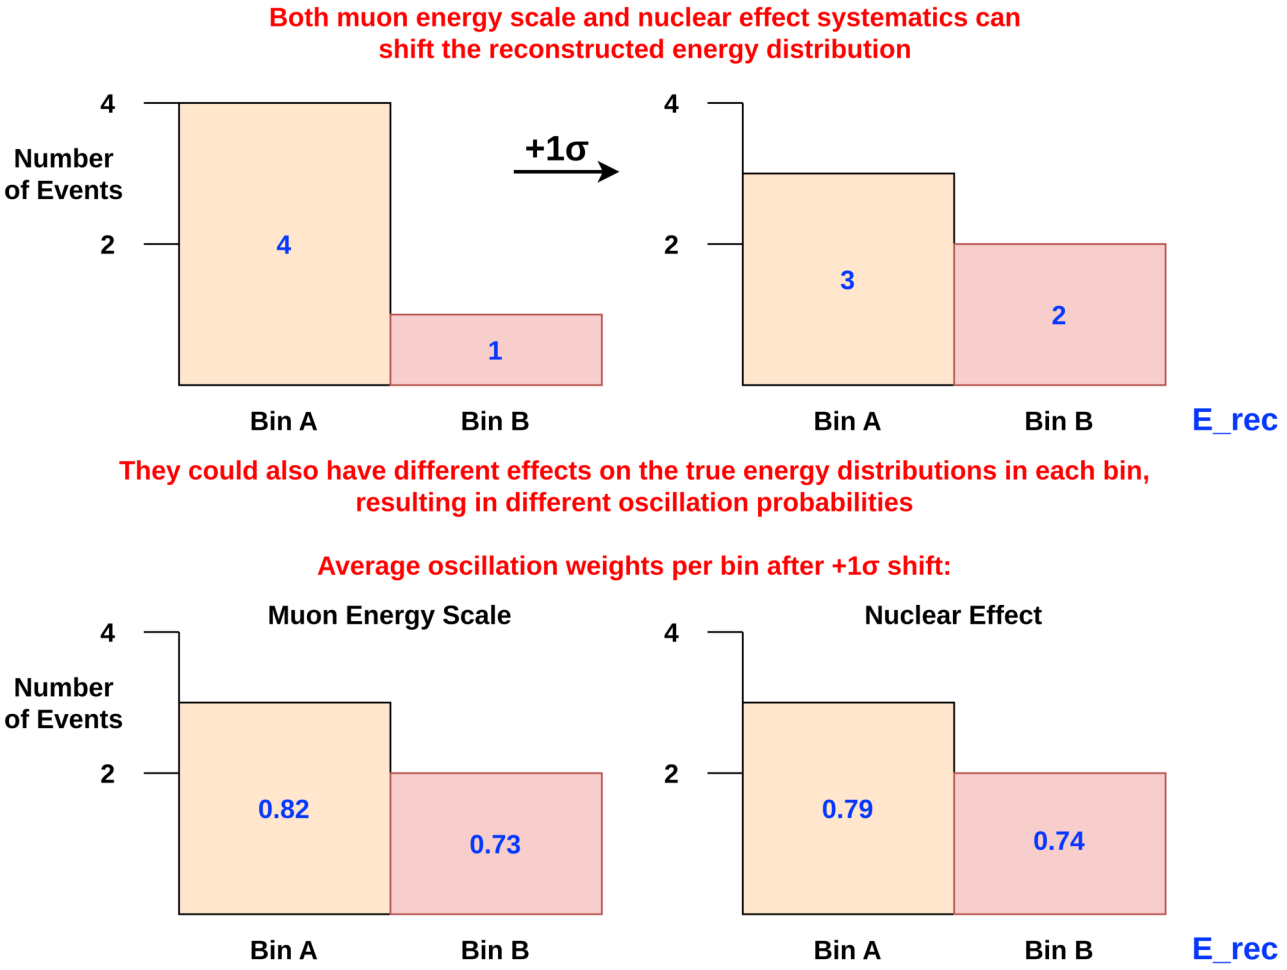

Figure 7.4: Toy example of degeneracy within parameters of the systematic uncertainty model. The top half shows two systematics which have the same effect in the observable distribution, in this case, reconstructed energy. The bottom half shows the effect of each systematic on the true energy distribution, which changes oscillated spectra and can mimic oscillations.

of parameter values produce the same or very similar predictions. The 'true' experiment has systematic uncertainties that have effects on the data which are indistinguishable (or at least close to indistinguishable) from each other. This degeneracy between parameters must be apparent in the systematic models that are included in the analysis.

A toy example of this concept can be seen in Figure 7.4. Let's consider two different systematic uncertainties: muon energy scale and some nuclear effect. Both of these systematics cause a shift in the observable distributions, in this case, reconstructed energy. In this toy example, let's imagine that shifting both these parameters have an identical effect on the reconstructed

energy. For an ND sample, the observables are reconstructed energy and reconstructed inelasticity. Since they have the same effect, there is no observable difference between a $+1\sigma$ shift in muon energy scale or nuclear effect and the same predicted spectra will be obtained. This prevents any one of these systematics from becoming individually constrained, rather they become anti-correlated. If changing these systematics also resembles a change in one of the oscillation parameters, i.e. Δm^2_{32} also shifts the energy distribution, this will result in a weaker constraint on Δm^2_{32}.

Furthermore, each systematic parameter will shift different events. This might change the true energy distribution in each bin, even though the reconstructed energy distribution is the same. This results in different oscillation probabilities for each bin, which also mimics the effect of oscillation parameters. This further weakens the constraint on the oscillation parameters. Also, oscillation sensitivity can be significantly affected if two parameters which both affect the ND samples are degenerate, but only one of them affects the FD. The parameters will become anti-correlated to constrain the ND event rate (essentially cancelling each other out), but since only one of them affects the FD, the FD spectrum will have a much wider distribution, reducing the oscillation parameter constraints. For example, this might occur if a cross-section effect is degenerate with a detector systematic that only affect the ND but not the FD.

Given a large number of events in the DUNE ND samples, and hence almost zero statistical uncertainty, accurate modeling of the degeneracies present is vital, as missing some would lead to significant overconstraint of parameters. The change in each systematic and oscillation parameter can be differentiated from each other, causing each parameter to become constrained. In current neutrino experiments, the large statistical uncertainty creates ambiguity between changes in the systematic parameters and oscillation parameters. In next-generation experiments, such as DUNE, over-simplified systematic uncertainty models will result in unrealistic measurements.

Figure 7.5 shows the posterior distribution of M_A^{QE} after a joint FD + ND sensitivity study. The prior uncertainty on M_A^{QE} is 1. The posterior uncertainty can be determined via

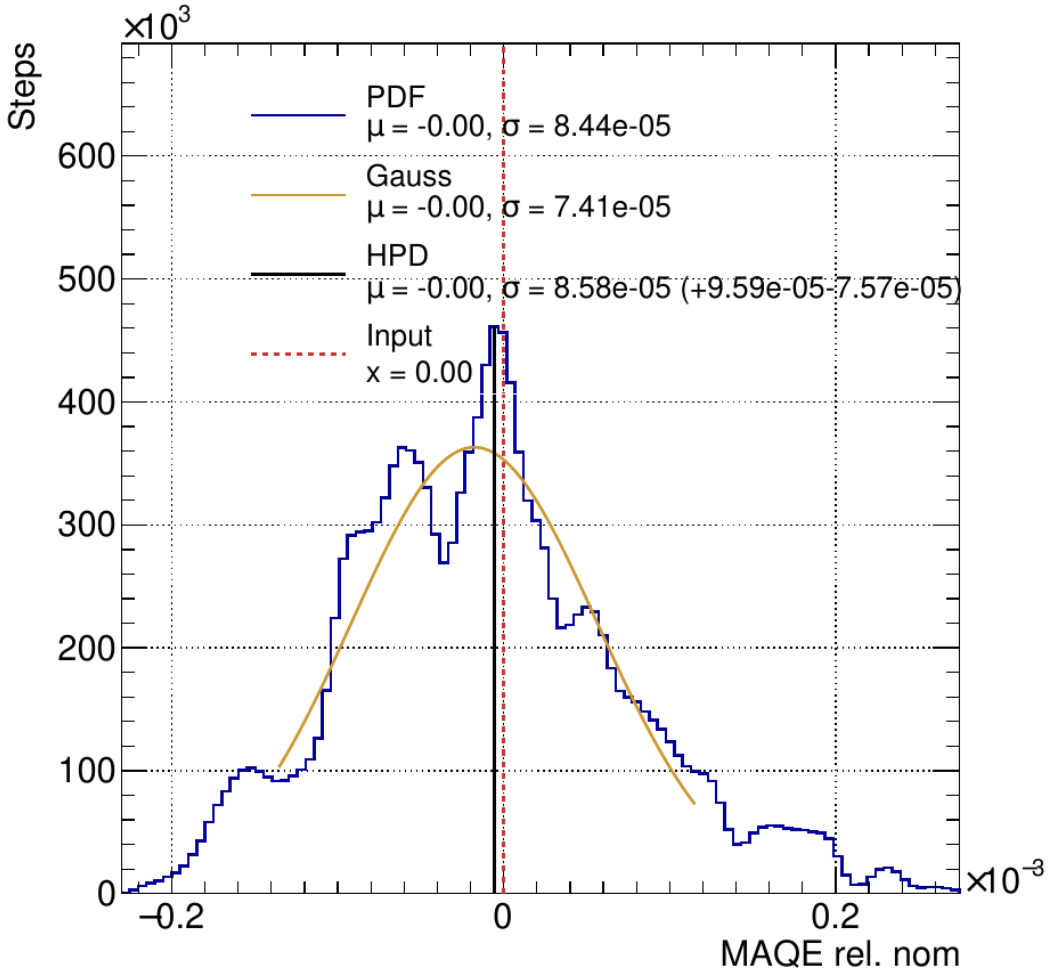

Figure 7.5: Posterior distribution of M_A^{QE} after a joint FD + ND fit. A Gaussian fit is shown along with the HPD and the input value. The prior uncertainty of M_A^{QE} is 1.

several methods, however, using the HPD method we find a posterior 1σ range of 8.6×10^{-5}. This is clearly an unrealistic constraint on this cross-section parameter and is a result of an oversimplified cross-section systematic model. The parameterised reconstruction at the ND also plays a role, since complex reconstruction effects smear the relationship between true and reconstructed energy which reduces the sensitivity of the experiment. The jaggedness in the distribution is caused by poor mixing of the parameters in the Markov chain which occurs when these extreme unrealistic constraints lead to some parameters having very small errors.

To test that degeneracy is missing in the systematic model, a simple study was carried out in which 'fake' degeneracy is introduced. This was done by including a duplicate M_A^{QE} parameter into the fit. The idea is that since both parameters are perfectly degenerate, they

would become anti-correlated, preventing either parameter from becoming heavily constrained. Also, the study is constructed so that only one of these parameters affects the FD samples. Figure 7.6 shows the posterior distribution of M_A^{QE} after a joint FD + ND sensitivity study in which the fake degeneracy was introduced. We see a posterior uncertainty of 6.2×10^{-3} when using the HPD method. This is $\sim 100\times$ larger than posterior uncertainty shown in Figure 7.5. However, this post-fit uncertainty is still unrealistic, as we do not expect to constrain M_A^{QE} to 1% of its prior uncertainty. What is necessary is a complete systematic uncertainty model which accurately reflects the true degeneracies in the experiment.

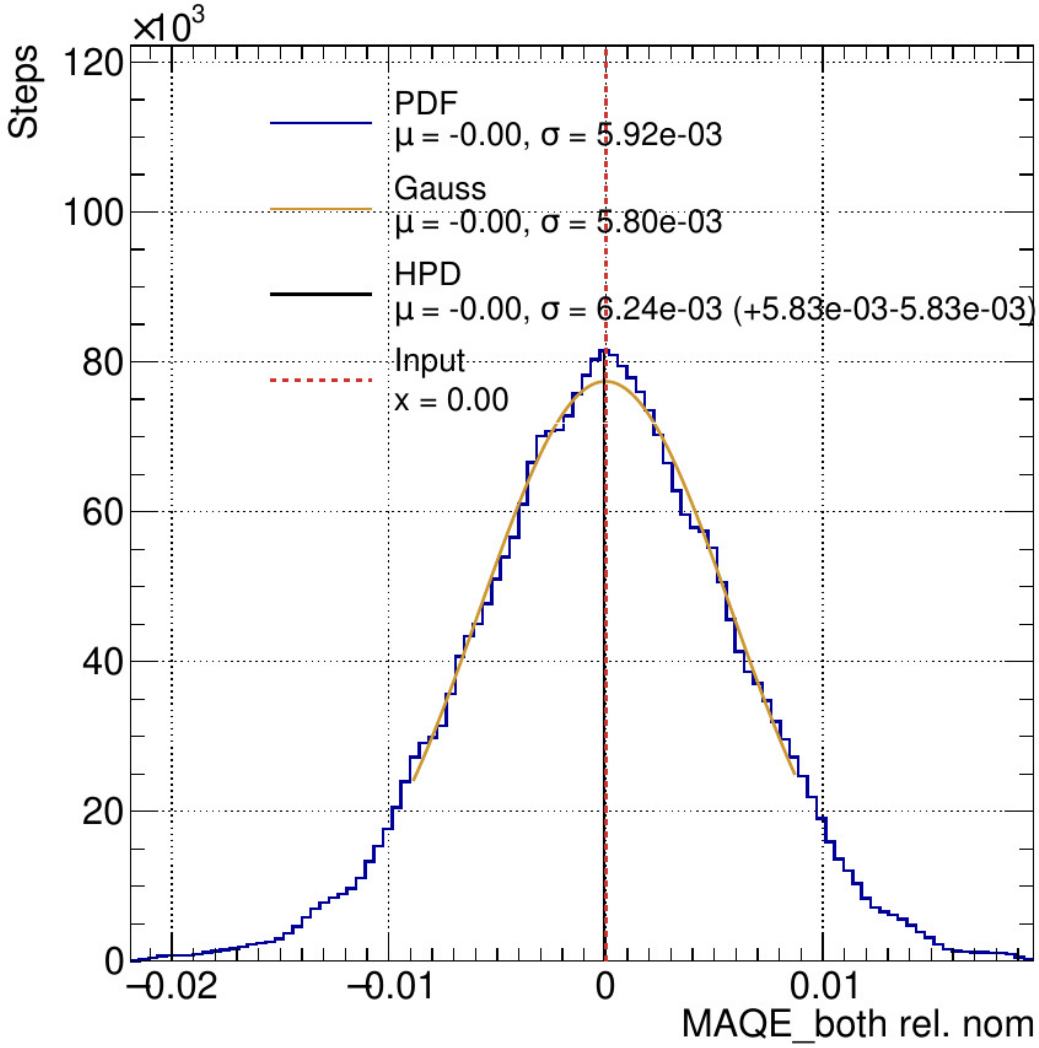

Figure 7.6: Posterior distribution of M_A^{QE} after a joint FD + ND fit including a duplicate parameter to introduce fake degeneracy. A Gaussian fit is shown along with the HPD and the input value. The prior uncertainty of M_A^{QE} is 1.

Using the current uncertainty model, a different approach is necessary to avoid overcon-

straint. In the fits which produced the plots in Figure 7.5 and Figure 7.6, each of the systematic uncertainties are included as free parameters which vary and change the observed spectra, including the ND detector uncertainties. In the final sensitivity studies, the ND detector systematics are not included as free parameters. Instead, a bin-to-bin covariance matrix is constructed from the ND detector systematics, including energy scale, resolution and acceptance uncertainties. This is done using the many-universes technique, in which each parameter is randomly thrown according to their prior uncertainties. The covariance matrix therefore has dimensions equal to the total number of bins in the ND samples. Given 17 reconstructed energy bins and 6 reconstructed inelasticity bins per sample, there are 102 bins in total for each covariance matrix. Hence, the ND detector covariance matrix is a 204 x 204 matrix.

Instead of calculating a Poisson likelihood from the observed event rate and the predicted event rate, a χ^2 is calculated using the ND covariance matrix in the same way as the Gaussian penalty terms for the systematic parameters moving away from prior values. The statistical uncertainty of the nominal prediction is added to the diagonal values of the ND covariance matrix prior to the fit. In reality, the statistical uncertainty of the prediction at each step of the Markov chain should be used. However, this would require re-inverting the matrix at each step which is computationally costly. It was found that the statistical uncertainty of the nominal prediction was a good approximation.

7.6 Oscillation Sensitivities

This section shows the results of a sensitivity study of DUNE using Asimov FD and ND data given a 336 kt-MW-yr exposure and generated with the oscillation parameter values shown in Table 6.1. These results will be referred to as FD + ND. The results presented have been obtained from a Markov chain containing 165 million steps after the burn-in was removed.

Figure 7.7 shows the posterior distribution in $\sin^2\theta_{23} - \Delta m^2_{32}$. Similar to the FD-only studies, only a negligible number of steps, 2500, were accepted in the inverted hierarchy. Hence, the posterior here is shown only under normal hierarchy. There is a good agreement between

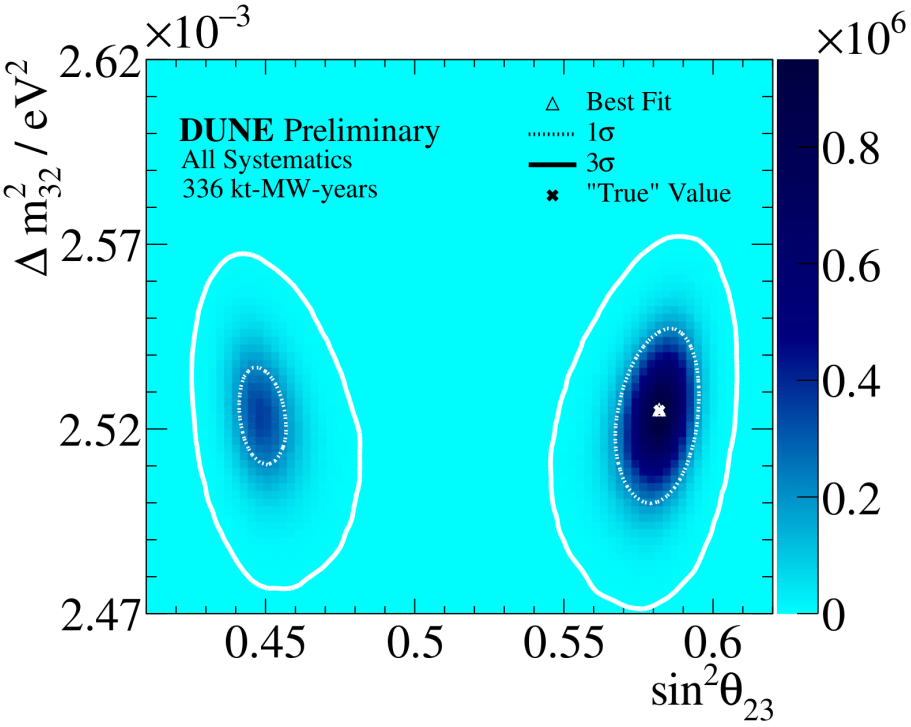

Figure 7.7: 2D posterior in $\sin^2\theta_{23} - \Delta m^2_{32}$ from an FD + ND sensitivity study. The 1σ and 3σ credible intervals are shown in the dashed and bold lines respectively.

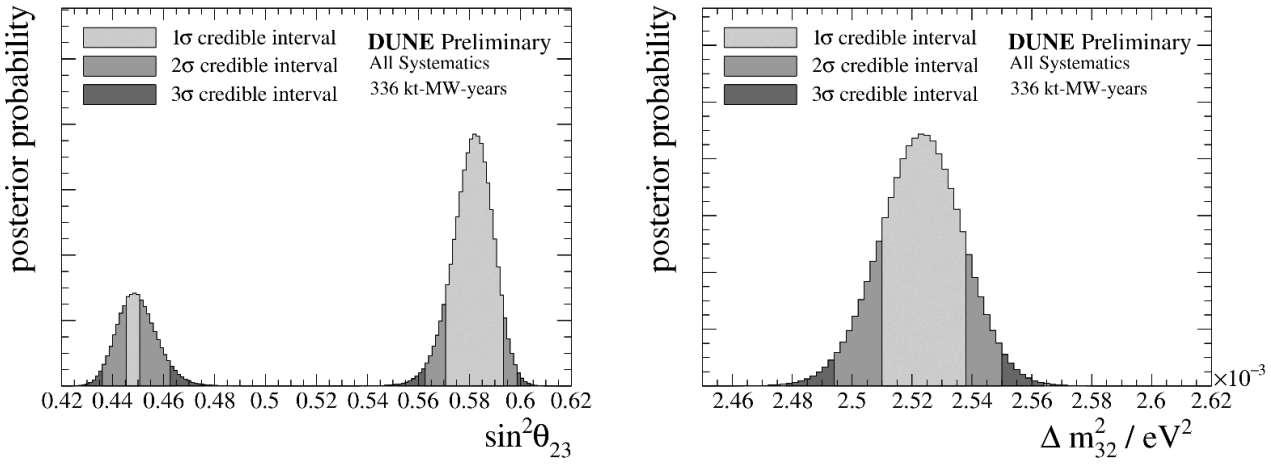

(a) $\sin^2\theta_{23}$ **(b)** Δm^2_{32} NO

Figure 7.8: 1D posterior in (a) $\sin^2\theta_{23}$ and (b) Δm^2_{32} in NO from a FD + ND sensitivity study. The 1σ, 2σ and 3σ credible intervals are shown in the light grey, grey and dark grey regions respectively.

the Asimov and best-fit points. Both θ_{23} octants are being explored, with the 1σ credible interval still containing points in the wrong octant. This can also be seen in Figure 7.8, which shows the 1D distributions in both $\sin^2\theta_{23}$ and Δm^2_{32}. It is clear that the correct octant is still preferred.

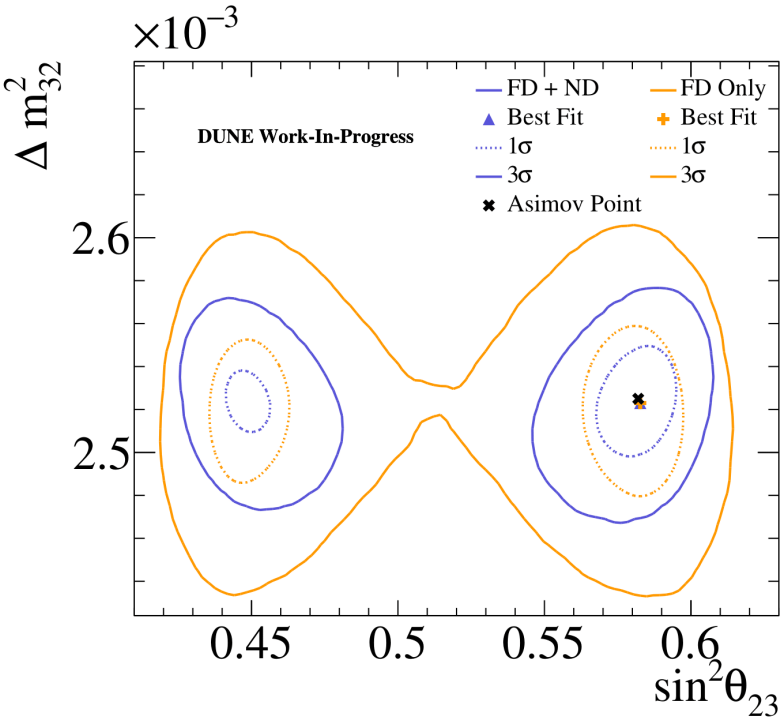

Figure 7.9: Comparison of the 1σ and 3σ credible intervals from the 2D posterior in $\sin^2\theta_{23} - \Delta m^2_{32}$ between FD + ND and FD-only sensitivity studies. The 1σ and 3σ credible intervals are shown in the dashed and bold lines respectively.

Figure 7.9 shows the comparison of FD + ND and FD-only credible intervals for $\sin^2\theta_{23} - \Delta m^2_{32}$, highlighting the effect of the ND samples. There is a significant reduction in the size of both 1σ and 3σ contours. The 3σ contour is no longer fully connected, excluding maximal $\sin^2\theta_{23}$. The effect of the ND samples is more significant on $\sin^2\theta_{23}$ values in between both octants, compared with values outside. There is also a much more significant reduction in the 1σ contour in the wrong octant compared with the reduction in the correct octant. Comparing with Figure 6.8, we see the FD + ND contours more closely resemble the statistics-error only fit than the FD-only fit, emphasising the ND constraining power.

Figure 7.10 shows the posterior distribution in $\sin^2\theta_{13} - \delta_{CP}$. As observed in Section 6.6,

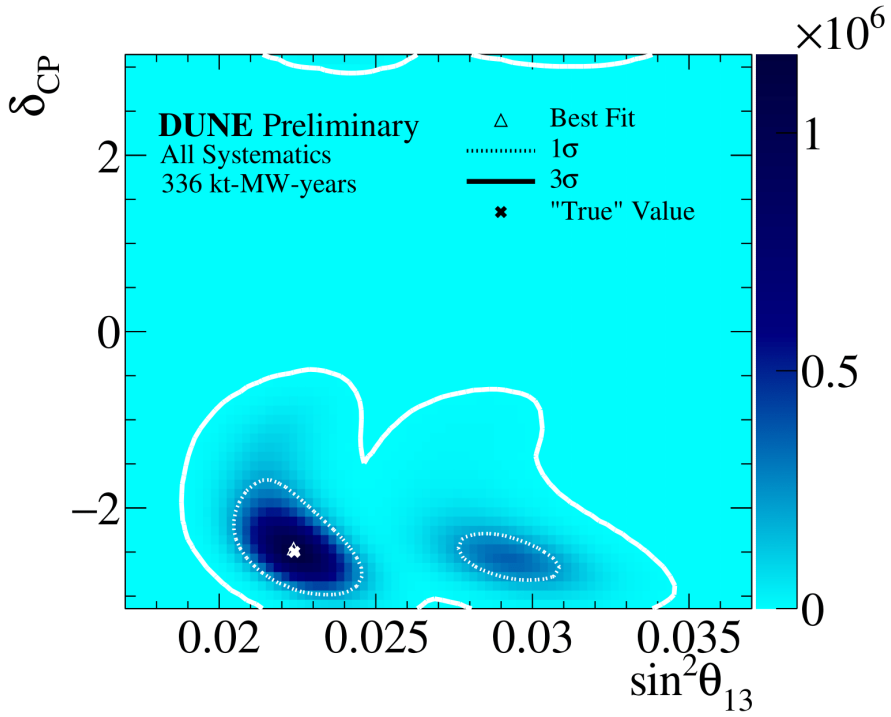

Figure 7.10: 2D posterior in $\sin^2\theta_{13} - \delta_{CP}$ from an FD + ND sensitivity study. The 1σ and 3σ credible intervals are shown in the dashed and bold lines respectively.

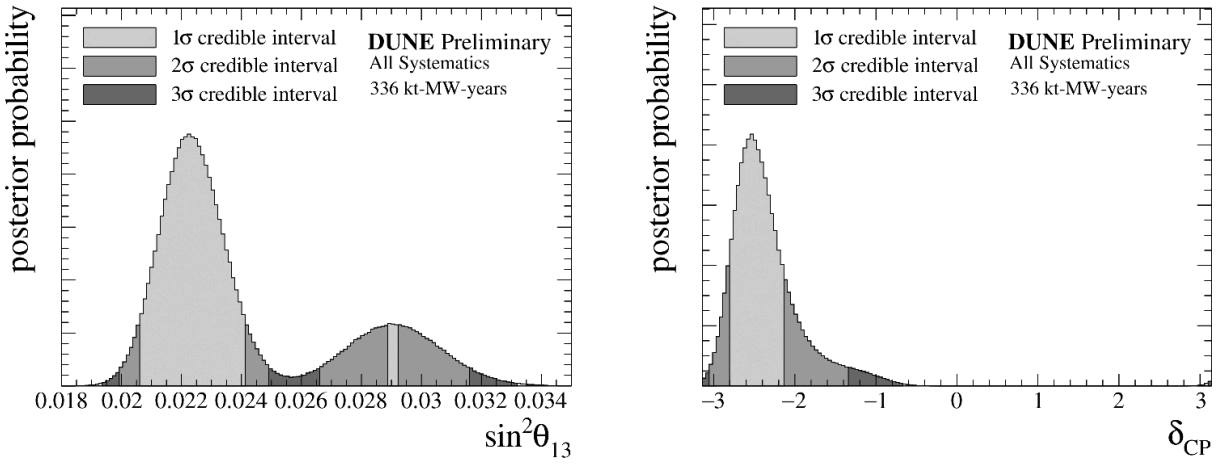

(a) $\sin^2\theta_{13}$ **(b)** δ_{CP}

Figure 7.11: 1D posterior in (a) $\sin^2\theta_{13}$ and (b) δ_{CP} in NO from a FD + ND sensitivity study. The 1σ, 2σ and 3σ credible intervals are shown in the light grey, grey and dark grey regions respectively.

there is a second peak in θ_{13} being explored due to the θ_{23} θ_{13}-octant degeneracy. The 3σ contour remains connected around both θ_{13} peaks. This can also be seen in the 1D distributions shown in Figure 7.11. There is still a significant posterior connecting both θ_{13} peaks, compared with Figure 6.11 where there is a small region with almost no posterior. Both the 2D and 1D distributions show that an exposure of 336 kt-MW-yrs is enough to exclude $\delta_{CP} = 0$ at 3σ, given this true δ_{CP} value. This has been true for all sensitivity studies. Although the contour remains closed, we see some separation beginning to be formed. This is more evident in Figure 7.12, which shows the comparison of FD + ND and FD-only credible intervals for $\sin^2\theta_{13} - \delta_{CP}$. There is a dip in the 3σ interval, concerning δ_{CP}, at roughly $\sin^2\theta_{13} = 0.025$, which gets deeper when the ND samples are included. Otherwise, there is not a significant change in the δ_{CP} sensitivity when the ND samples are included. This is expected since the δ_{CP} measurement is statistically limited until much higher exposures are achieved. There is a more significant constraint from the ND samples in the $\sin^2\theta_{13}$, especially in the larger values.

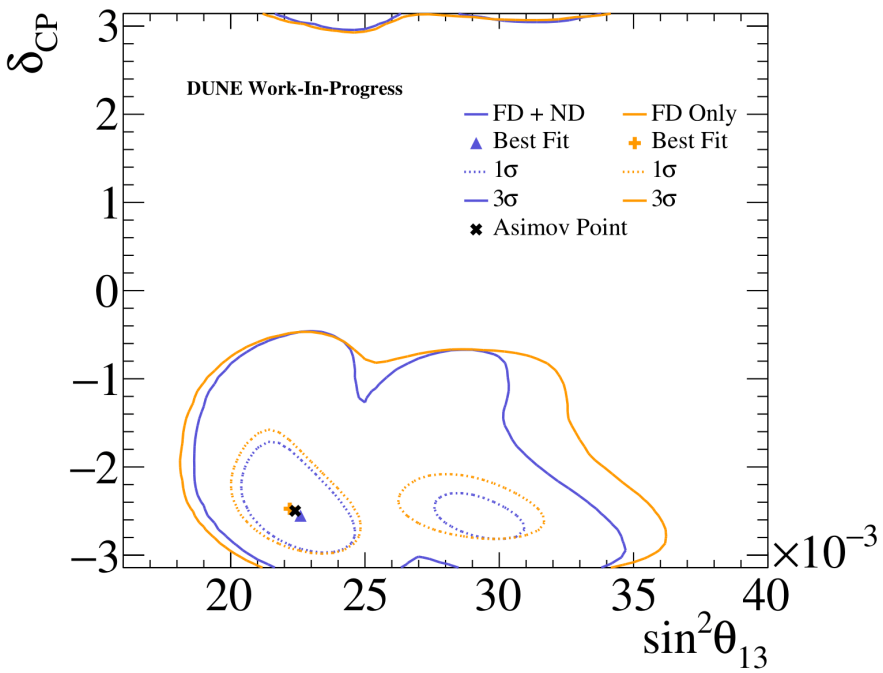

Figure 7.12: Comparison of the 1σ and 3σ credible intervals from the 2D posterior in $\sin^2\theta_{13} - \delta_{CP}$ between FD + ND and FD-only sensitivity studies. The 1σ and 3σ credible intervals are shown in the dashed and bold lines respectively.

Figure 7.13 shows the 2D posterior in both $\sin^2\theta_{23} - \Delta m^2_{32}$ and $\sin^2\theta_{13} - \delta_{CP}$ after the reactor

constraint has been applied. As mentioned in Section 6.6, the reactor constraint consists of a Gaussian constraint on θ_{13}, using the mean and error from Table 6.1. The wrong θ_{23} octant is completely suppressed. Unlike the FD-only study, the 3σ credible interval no longer contains the wrong octant after the reactor constraint is applied. The upper peak in $\sin^2\theta_{13}$ is also completely suppressed, due to the correlation between the upper peak in $\sin^2\theta_{13}$ and the octant in wrong θ_{23}. This demonstrates that at this Asimov value and given an exposure of 336 kt-MW-yrs, the wrong θ_{23} octant is excluded at 3σ after the reactor constraint is applied.

Figure 7.14 shows the comparison of the credible intervals between the FD + ND sensitivity study with and without the reactor constraint. The credible interval in the correct θ_{23} octant is almost identical, with a minor reduction in size after the reactor constraint is applied. In $\sin^2\theta_{13} - \delta_{CP}$, as we'd expect the width of the credible interval is constrained to the Gaussian prior that is applied. There is a minimal effect on the constraint on δ_{CP}.

So far all sensitivities studies have been shown using a flat prior in δ_{CP}. This states that, prior to the observation of data, each value of δ_{CP} is equally likely. However, when considering other quantities, such as $\sin\delta_{CP}$ or $\cos\delta_{CP}$, a flat prior in δ_{CP} results in a non-uniform prior in these other quantities. Hence, one might instead choose a flat prior in some other quantity, which in turn results in a non-uniform prior in δ_{CP}. Here, we consider the effect of using a flat prior in $\sin\delta_{CP}$, since the Jarlskog invariant, shown in Eq 2.36, depends on $\sin\delta_{CP}$. Figure 7.15 shows the 2D posterior in both $\sin^2\theta_{23} - \Delta m^2_{32}$ and $\sin^2\theta_{13} - \delta_{CP}$ when using a flat prior in $\sin\delta_{CP}$. There is minimal change in $\sin^2\theta_{23} - \Delta m^2_{32}$. This can be seen more clearly in Figure 7.16 which shows the comparison of credible intervals between a flat δ_{CP} prior and a flat $\sin\delta_{CP}$ prior. There is almost no difference between the credible intervals in $\sin^2\theta_{23} - \Delta m^2_{32}$. However, in $\sin^2\theta_{13} - \delta_{CP}$ there is a significant change in the shape of the contours. There is a reduction in the size of the 1σ contour in the correct $\sin^2\theta_{13}$ lobe, with regions closer to 0 being excluded. In the wrong $\sin^2\theta_{13}$ lobe, there is a separation of the 3σ contour at values around $\delta_{CP} = -\frac{\pi}{2}$. This is expected since to transform to the flat in $\sin\delta_{CP}$ prior, each step is weighted by $\cos\delta_{CP}$. The effect of this can also be seen in the 3σ contour in the correct $\sin^2\theta_{13}$

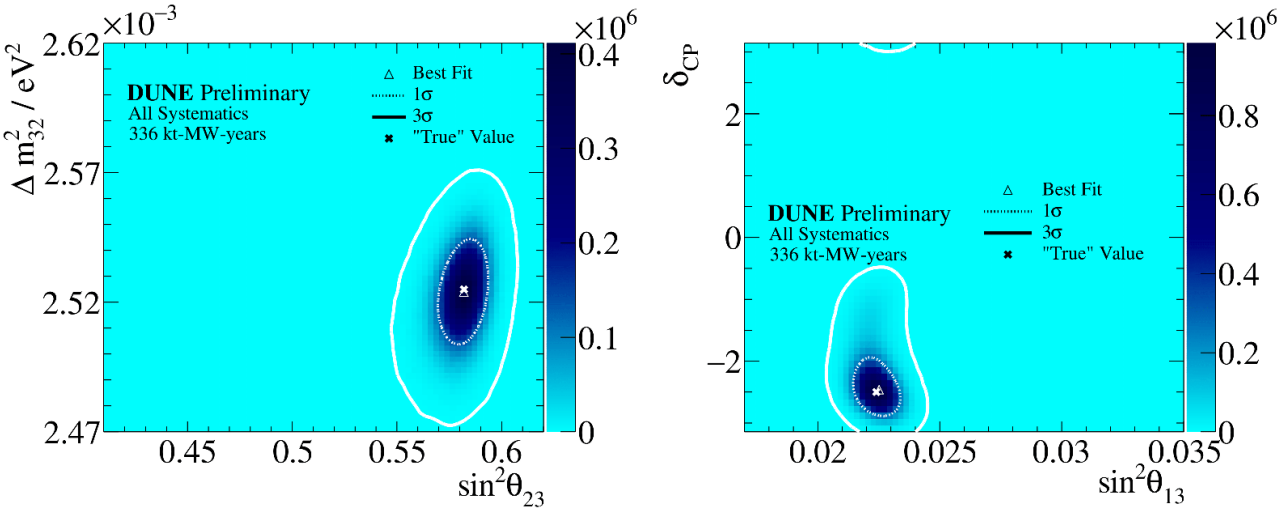

(a) $\sin^2\theta_{23} - \Delta m^2_{32}$ **(b)** $\sin^2\theta_{13} - \delta_{CP}$

Figure 7.13: 2D posterior in (a) $\sin^2\theta_{23} - \Delta m^2_{32}$ and (b) $\sin^2\theta_{13} - \delta_{CP}$ from an FD + ND sensitivity study after applying the reactor constraint. The 1σ and 3σ credible intervals are shown in the dashed and bold lines respectively.

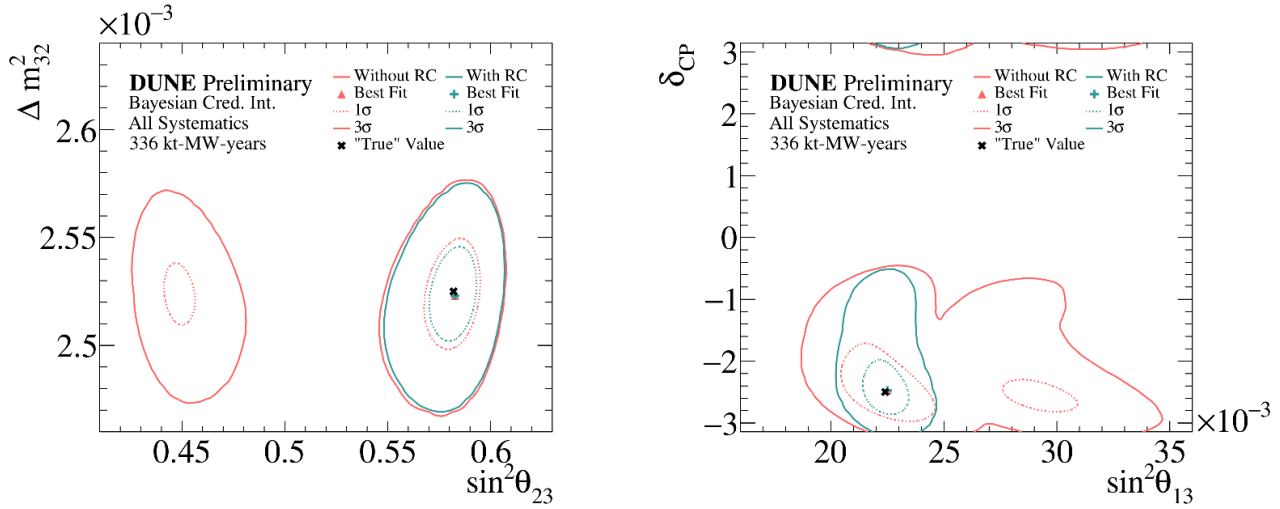

(a) $\sin^2\theta_{23} - \Delta m^2_{32}$ **(b)** $\sin^2\theta_{13} - \delta_{CP}$

Figure 7.14: Comparison of the 1σ and 3σ credible intervals from the 2D posterior in (a) $\sin^2\theta_{23} - \Delta m^2_{32}$ and (b) $\sin^2\theta_{13} - \delta_{CP}$ between the FD + ND sensitivity study without and with the reactor constraint. The 1σ and 3σ credible intervals are shown in the dashed and bold lines respectively.

lobe, in which the contour becomes thinner at $\delta_{CP} = -\frac{\pi}{2}$.

Figure 7.17 shows the J_{CP} posterior distribution from the FD + ND sensitivity study as well as the credible intervals. The distribution looks almost identical to the FD-Only J_{CP} distribution shown in Figure 6.17. The distribution peaks at roughly $J_{CP} = -0.02$ with a smaller second peak occurring at lower values. There is also a second shoulder towards the end of the distribution. The edge of the 3σ interval still rests at $J_{CP} = 0$. Figure 7.18 shows the J_{CP} posterior distribution after applying the reactor constraint and changing to a flat prior in $\sin\delta_{CP}$. After applying the reactor constraint, the width of the distribution decreases slightly. The 3σ credible interval no longer contains $J_{CP} = 0$. This suggests at this exposure and at this Asimov point, CP-conservation can be excluded at 3σ after applying the reactor constraint. We also see that the shoulder at low values disappears, suggesting that it is caused by the posterior in the wrong octant of θ_{23}.

The change in the δ_{CP} prior is especially interesting for J_{CP} since J_{CP} is a function of $\sin\delta_{CP}$, not δ_{CP}. This means that when a flat prior in δ_{CP} is used, as in Figure 7.17, a non-flat prior is introduced into J_{CP}. Switching to a flat prior in $\sin\delta_{CP}$ results in a flat prior in J_{CP}. As can be seen in Figure 7.18b, this change in prior results in a smoother distribution around the peak. The second peak is eliminated, suggesting that the cause of this was the shape in the prior in J_{CP}. The shoulder at the end of the distribution remains since the reactor constraint has not been applied. The distribution also widens, resulting in the 3σ interval extending past $J_{CP} = 0$.

Figure 7.19 shows the 1σ posterior uncertainties and central values from the FD-only sensitivity study and the FD + ND sensitivity study relative to the prior 1σ uncertainties and central values. Similarly to the FD-Only case in Figure 6.18, for the vast majority of systematics the posterior uncertainty corresponds to the HPD 1σ credible interval, whereas others are obtained from a Gaussian fit. As we'd expect, generally there is a much stronger constraint on the systematic parameters upon inclusion of the ND samples. This is a result of the large number of neutrino events at the ND. Although the FD detector systematics have no effect on the ND

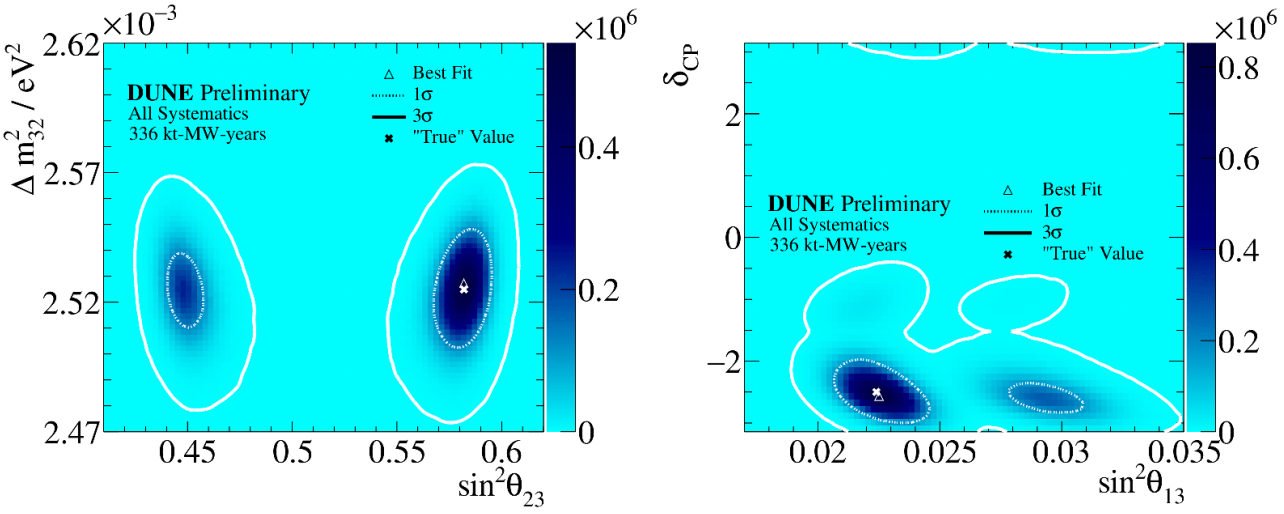

(a) $\sin^2\theta_{23} - \Delta m^2_{32}$ **(b)** $\sin^2\theta_{13} - \delta_{CP}$

Figure 7.15: 2D posterior in (a) $\sin^2\theta_{23} - \Delta m^2_{32}$ and (b) $\sin^2\theta_{13} - \delta_{CP}$ from an FD + ND sensitivity study after applying a flat prior in $\sin\delta_{CP}$. The 1σ and 3σ credible intervals are shown in the dashed and bold lines respectively.

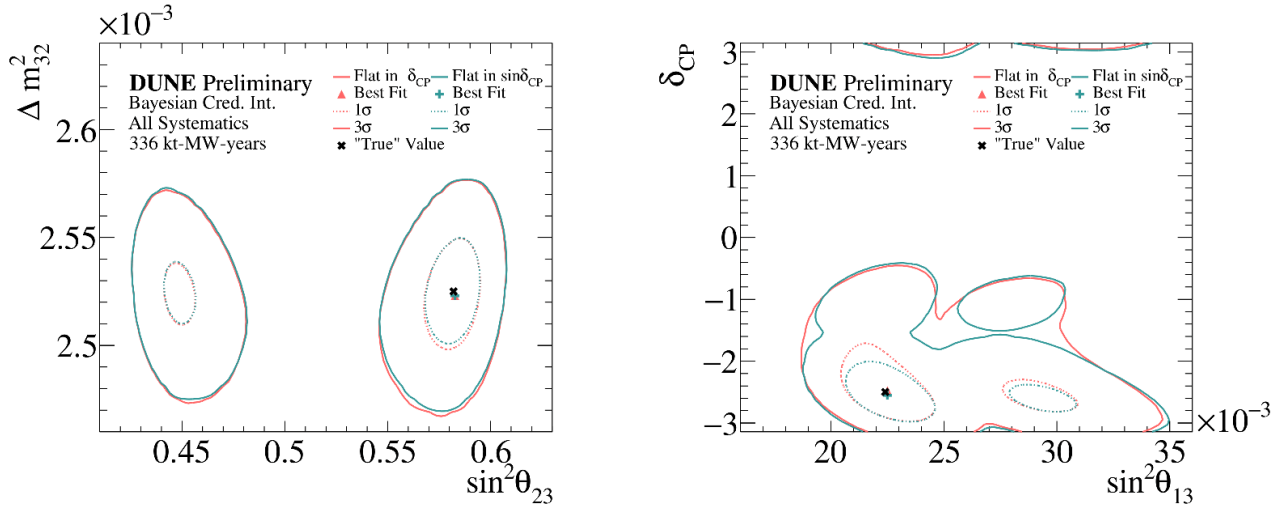

(a) $\sin^2\theta_{23} - \Delta m^2_{32}$ **(b)** $\sin^2\theta_{13} - \delta_{CP}$

Figure 7.16: Comparison of the 1σ and 3σ credible intervals from the 2D posterior in (a) $\sin^2\theta_{23} - \Delta m^2_{32}$ and (b) $\sin^2\theta_{13} - \delta_{CP}$ between the FD + ND sensitivity study with a flat prior in δ_{CP} and $\sin\delta_{CP}$. The 1σ and 3σ credible intervals are shown in the dashed and bold lines respectively.

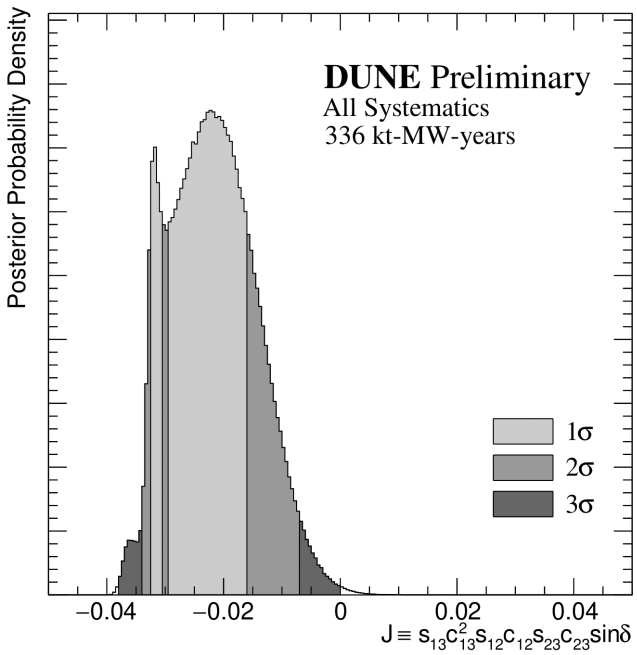

Figure 7.17: 1D posterior in J_{CP} from an FD + ND sensitivity study. The 1σ and 3σ credible intervals are shown in the dashed and bold lines respectively.

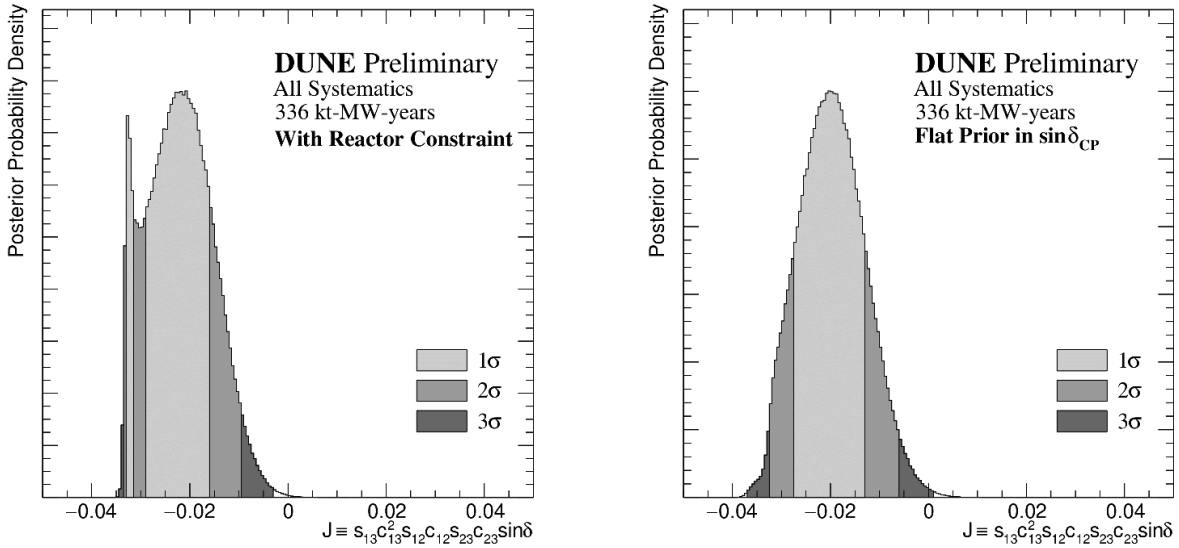

(a) reactor constraint **(b)** flat prior in $\sin\delta_{CP}$

Figure 7.18: 1D posterior in J_{CP} from an FD + ND sensitivity study after (a) applying the reactor constraint and (b) changing to a flat prior in $\sin\delta_{CP}$. The 1σ and 3σ credible intervals are shown in the dashed and bold lines respectively.

MC, there is a shift in the central value between the FD-Only and FD + ND sensitivity studies. This also causes a shift in the 1σ range causing the appeared one-sided constraint coming from the ND samples. The vast majority of cross-section systematics receive a much stronger constraint upon the addition of the ND samples. There are few cross-section systematics that have very weak constraints from both sensitivity studies. Generally, these are systematics that apply to only DIS interactions or only NC interactions. All of the flux systematics receive a much tighter constraint. Here only the FD flux systematics are shown. However, the ND flux systematics follow almost an identical trend as a result of the strong correlations between ND and FD flux systematics which affect the same neutrino flavours and energy ranges.

7.7 Summary

The results of this sensitivity study show the importance of an LAr near detector in making precise measurements of oscillation parameters. Figure 7.9 and Figure 7.12 show a significant reduction in the credible intervals upon the addition of the ND samples, especially for $\sin^2\theta_{23}$, $\sin^2\theta_{13}$ and Δm^2_{32}. The effect is not as significant for δ_{CP} since it is still expected that the δ_{CP} measurement will be statistically limited at this exposure. However, these sensitivities do not use the full exposure of DUNE and do not include all systematic uncertainties, hence the effect of systematic uncertainty on δ_{CP} sensitivity is expected to increase. The effect of the ND samples on the systematic uncertainty parameters is also highlighted, showing a much greater constraint when compared with the FD-only sensitivities.

It should be noted that the ND model and samples included in this analysis do not reflect the current DUNE design. This analysis, which developed on the TDR analysis, is intended to show the benefits of ND-LAr in constraining systematic uncertainties. The ND-GAr and SAND detectors, as well as the off-axis capabilities, in the current DUNE ND design were not included, as well as the uncertainties that these detectors are intended to reduce. Furthermore, the ND-LAr samples included in this analysis did not include full reconstruction. In the future, sensitivity studies should be performed with more realistic ND-LAr samples, samples from all

components of the ND and better systematic models. This is imperative for reliable predictions of DUNE's sensitivity to the oscillation parameters.

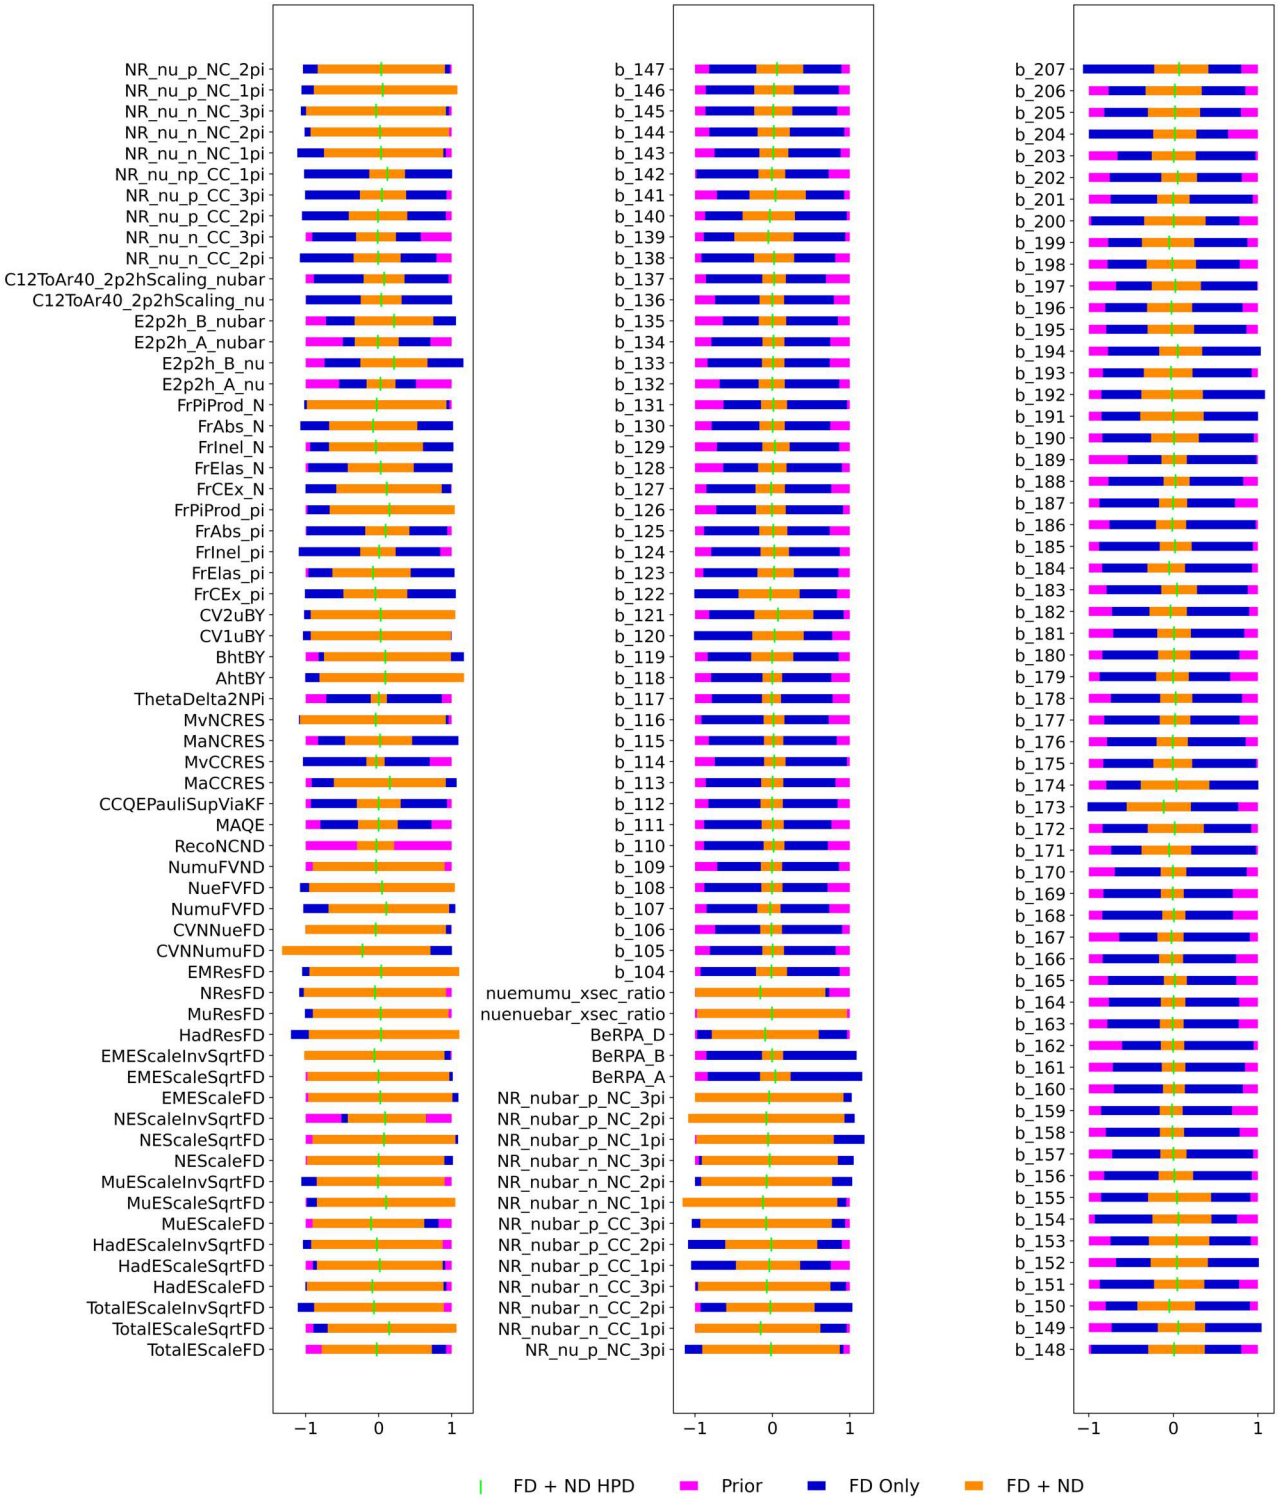

Figure 7.19: The ratio of posterior to prior 1σ uncertainties for systematic parameters FD-only sensitivity study and FD + ND sensitivity study. The prior uncertainty is shown in magenta, the FD-only posterior uncertainty is shown in blue and the FD + ND posterior uncertainty is shown in orange. The vertical marks show the FD + ND HPD point relative to the prior central value as a fraction of the prior uncertainty. The ND flux systematics, i.e. b_0 - b_{103}, have been omitted for convenience since they have similar constraints to the FD flux systematics.

Chapter 8

Conclusions and Outlook

The purpose of this thesis was to perform the first Bayesian neutrino oscillation sensitivity study of DUNE. DUNE is a next-generation long-baseline experiment that aims to make precise measurements of neutrino oscillation parameters, with extra significance on determining whether CP-violation exists in the neutrino sector. We must make predictions of DUNE's sensitivity to the oscillation parameters to ensure the experiment can achieve its physics goals. Given the large number of neutrino events DUNE will collect, the systematic uncertainty will dominate the statistical uncertainty. As a result, the systematic uncertainties must be implemented correctly in analyses to produce realistic sensitivities. The analysis presented in this thesis uses a more comprehensive systematic model compared with previous DUNE sensitivities [2]. Furthermore, taking a Bayesian approach provides a robust validation of frequentist results, as well as allowing a more efficient study of the behaviour of systematic uncertainty parameters.

Chapter 6 presented the results of Asimov sensitivities considering only the predicted FD neutrino event samples. Sensitivities were performed both excluding, i.e. statistics-only in Section 6.6.1 and including the systematic uncertainty parameters in Section 6.6.2. At the Asimov point chosen for these sensitivities, there is a significant degeneracy in the lower and upper octant of $\sin^2\theta_{23}$, as visible in Figure 6.8. This degeneracy results in a second peak

in $\sin^2\theta_{13}$, as shown in Figure 6.10. Figure 6.12 demonstrated that by applying the reactor constraint, the second peak in $\sin^2\theta_{13}$ is suppressed, which in turn suppresses the wrong octant in $\sin^2\theta_{23}$. The statistics-only sensitivity study also showed DUNE's power in resolving the mass ordering, with very minimal posterior distribution found in the wrong hierarchy.

As we'd expect, there is a significant drop in sensitivity to all oscillation parameters upon the inclusion of the systematic uncertainty parameters. There is a more significant effect in $\sin^2\theta_{23}$, $\sin^2\theta_{13}$ and Δm^2_{32} than δ_{CP}, as illustrated by Figure 6.13 and Figure 6.15. This shows, as we'd expect, that at this exposure the measurement of δ_{CP} is more statistically limited than the other oscillation parameters of interest. Applying the reactor constraint suppresses the second peaks in $\sin^2\theta_{23}$ and $\sin^2\theta_{13}$, but has little to no effect on Δm^2_{32} and δ_{CP}. Figure 6.18 demonstrated that the FD samples alone provide a very weak constraint, illustrating the necessity of a near detector.

Chapter 7 built upon the results of Chapter 6, including near detector samples and presenting the results of a joint FD and ND sensitivity study. The constraining power of the near detector is discussed in Section 7.5, highlighting the issue of over-constraints which can occur when the systematic model used for high-statistics is oversimplified. A study was performed in which fake degeneracy between parameters was added into the model and resulted in much wider posterior distributions. This helped confirm that the systematic models currently used in next-generation long-baseline sensitivities are over-simplified which leads to unrealistic constraints.

The result of the joint FD and ND study showed improved sensitivity across all oscillation parameters, as shown in Figure 7.9 and Figure 7.12. There is a more significant improvement in $\sin^2\theta_{23}$, $\sin^2\theta_{13}$ and Δm^2_{32} than δ_{CP}, as we'd expect due to the statistical limitation of the δ_{CP} at this exposure. With regards to CP violation, Figure 7.11 showed that $\delta_{CP} = 0$ can be excluded at 3σ given this Asimov point and a 336 kt-MW-yr exposure. The effect of the reactor constraint is shown in Figure 7.14, highlighting the improvement in the sensitivity of $\sin^2\theta_{23}$ and $\sin^2\theta_{13}$. The effect of changing from a flat prior in δ_{CP} to a flat prior in

$\sin\delta_{CP}$ is shown in Figure 7.16, showing generally little effect on the credible intervals. The distributions of J_{CP} from the FD + ND sensitivity study were produced, both with and without reactor constraint and δ_{CP} prior changes, and are shown in Figure 7.17 and Figure 7.18. This demonstrated that the reactor constraint improves the J_{CP}, as well as suggesting that the $\sin^2\theta_{23}$ octant and $\sin^2\theta_{13}$ correlation and flat δ_{CP} prior are responsible for irregularities in the J_{CP} distribution. When including the reactor constraint, $J_{CP} = 0$ is excluded at 3σ, which corresponds to the exclusion of CP conservation at 3σ. The effect of the ND samples on the systematic uncertainty parameters is shown in Figure 7.19. There is a significant reduction in the posterior uncertainty across cross-section and flux parameters, highlighting the importance of a capable ND in reducing systematic uncertainty.

Although these sensitivities give us some first-order indication of DUNE's capabilities, there are improvements that are necessary before we can produce realistic results. As mentioned in Section 7.2, the ND samples and systematic models in this analysis are designed to show the benefits of ND-LAr. As described in Chapter 3, ND-GAr and SAND will also be part of the DUNE ND complex, and uncertainties that only these detectors can constrain need to be included in the systematic model and studied in detail. Furthermore, it is important that all detectors included in analyses have full simulation and reconstruction. Looking further into the future, we expect more comprehensive cross-section uncertainty models to become available and a better understanding of neutrino flux from hadron production measurements. Better detector uncertainty models for both FD and ND will come from the analysis of prototype data, i.e. Proto-DUNE [86] and the 2x2 Demonstrator [120]. Work is currently ongoing to produce sensitivities including new samples and their associated uncertainties. This will lead us to more realistic predictions of DUNE's sensitivity, and ensure that DUNE will achieve its goal of producing world-leading measurements of oscillation parameters.

Bibliography

[1] J. H. Christenson et al. "Evidence for the 2π Decay of the K_2^0 Meson". In: *Phys. Rev. Lett.* 13 (4 July 1964), pp. 138–140. DOI: 10.1103/PhysRevLett.13.138. URL: https://link.aps.org/doi/10.1103/PhysRevLett.13.138.

[2] B. Abi et al. *Deep Underground Neutrino Experiment (DUNE), Far Detector Technical Design Report, Volume II: DUNE Physics*. 2020. arXiv: 2002.03005 [hep-ex].

[3] J Chadwick. "Intensitätsverteilung im magnetischen Spectrum der β-Strahlen von radium B + C". In: *Verhandl. Dtsc. Phys. Ges.* 16 (1914), pp. 383–391. URL: https://cds.cern.ch/record/262756.

[4] W. Pauli. "Dear radioactive ladies and gentlemen". In: *Phys. Today* 31N9 (1978), p. 27.

[5] E. Fermi. "Versuch einer Theorie der β-Strahlen. I". In: *Zeitschrift für Physik* 88.3 (Mar. 1934), pp. 161–177. ISSN: 0044-3328. DOI: 10.1007/BF01351864. URL: https://doi.org/10.1007/BF01351864.

[6] H. BETHE and R. PEIERLS. "The "Neutrino"". In: *Nature* 133.3362 (Apr. 1934), pp. 532–532. ISSN: 1476-4687. DOI: 10.1038/133532a0. URL: https://doi.org/10.1038/133532a0.

[7] C. L. Cowan et al. "Detection of the free neutrino: A Confirmation". In: *Science* 124 (1956), pp. 103–104. DOI: 10.1126/science.124.3212.103.

[8] F. Reines and C. L. Cowan. "Detection of the Free Neutrino". In: *Phys. Rev.* 92 (3 Nov. 1953), pp. 830–831. DOI: 10.1103/PhysRev.92.830. URL: https://link.aps.org/doi/10.1103/PhysRev.92.830.

[9] F. Reines et al. "Detection of the Free Antineutrino". In: *Phys. Rev.* 117 (1 Jan. 1960), pp. 159–173. DOI: 10.1103/PhysRev.117.159. URL: https://link.aps.org/doi/10.1103/PhysRev.117.159.

[10] Seth H. Neddermeyer and Carl D. Anderson. "Note on the Nature of Cosmic-Ray Particles". In: *Phys. Rev.* 51 (10 May 1937), pp. 884–886. DOI: 10.1103/PhysRev.51.884. URL: https://link.aps.org/doi/10.1103/PhysRev.51.884.

[11] J. C. Street and E. C. Stevenson. "New Evidence for the Existence of a Particle of Mass Intermediate Between the Proton and Electron". In: *Physical Review* 52.9 (Nov. 1937), pp. 1003–1004. DOI: 10.1103/PhysRev.52.1003.

[12] B. Pontecorvo. "Electron and Muon Neutrinos". In: *Zh. Eksp. Teor. Fiz.* 37 (1959), pp. 1751–1757.

[13] G. Danby et al. "Observation of High-Energy Neutrino Reactions and the Existence of Two Kinds of Neutrinos". In: *Phys. Rev. Lett.* 9 (1 July 1962), pp. 36–44. DOI: 10.1103/PhysRevLett.9.36. URL: https://link.aps.org/doi/10.1103/PhysRevLett.9.36.

[14] M.M. Block et al. "Neutrino interactions in the CERN heavy liquid bubble chamber". In: *Physics Letters* 12.3 (1964), pp. 281–285. ISSN: 0031-9163. DOI: https://doi.org/10.1016/0031-9163(64)91104-7. URL: https://www.sciencedirect.com/science/article/pii/0031916364911047.

[15] J.K. Bienlein et al. "Spark chamber study of high-energy neutrino interactions". In: *Physics Letters* 13.1 (1964), pp. 80–86. ISSN: 0031-9163. DOI: https://doi.org/10.1016/0031-9163(64)90316-6. URL: https://www.sciencedirect.com/science/article/pii/0031916364903166.

[16] Martin L. Perl et al. "Evidence for Anomalous Lepton Production in e+ - e- Annihilation". In: *Phys. Rev. Lett.* 35 (1975), pp. 1489–1492. DOI: 10.1103/PhysRevLett.35.1489.

[17] K. Kodama et al. "Observation of tau neutrino interactions". In: *Physics Letters B* 504.3 (Apr. 2001), pp. 218–224. ISSN: 0370-2693. DOI: 10.1016/s0370-2693(01)00307-0. URL: http://dx.doi.org/10.1016/S0370-2693(01)00307-0.

[18] D. DeCamp et al. "Determination of the number of light neutrino species". In: *Physics Letters B* 231.4 (1989), pp. 519–529. ISSN: 0370-2693. DOI: https://doi.org/10.1016/0370-2693(89)90704-1. URL: https://www.sciencedirect.com/science/article/pii/0370269389907041.

[19] S. et al. Schael. "Precision electroweak measurements on the Z resonance". In: *Phys. Rept.* 427 (2006), pp. 257–454. DOI: 10.1016/j.physrep.2005.12.006. arXiv: hep-ex/0509008.

[20] Mark Thomson. *Modern particle physics*. New York: Cambridge University Press, 2013. ISBN: 978-1-107-03426-6. DOI: 10.1017/CBO9781139525367.

[21] C. S. Wu et al. "Experimental Test of Parity Conservation in Beta Decay". In: *Phys. Rev.* 105 (4 Feb. 1957), pp. 1413–1415. DOI: 10.1103/PhysRev.105.1413. URL: https://link.aps.org/doi/10.1103/PhysRev.105.1413.

[22] R. P. Feynman and M. Gell-Mann. "Theory of the Fermi Interaction". In: *Phys. Rev.* 109 (1 Jan. 1958), pp. 193–198. DOI: 10.1103/PhysRev.109.193. URL: https://link.aps.org/doi/10.1103/PhysRev.109.193.

[23] J. H. Christenson et al. "Evidence for the 2π Decay of the K_2^0 Meson". In: *Phys. Rev. Lett.* 13 (4 July 1964), pp. 138–140. DOI: 10.1103/PhysRevLett.13.138. URL: https://link.aps.org/doi/10.1103/PhysRevLett.13.138.

[24] Abdus Salam. "Weak and Electromagnetic Interactions". In: *Conf. Proc. C* 680519 (1968), pp. 367–377. DOI: 10.1142/9789812795915_0034.

[25] Steven Weinberg. "A Model of Leptons". In: *Phys. Rev. Lett.* 19 (1967), pp. 1264–1266. DOI: 10.1103/PhysRevLett.19.1264.

[26] Sam Zeller. "Final State Nucleons for Neutrino-Nucleus". Interactions Workshop at JLAB. 2015. URL: https://www.jlab.org/conferences/funfact/talks/thursday/zeller.pdf.

[27] R. A. Smith and E. J. Moniz. "NEUTRINO REACTIONS ON NUCLEAR TARGETS". In: *Nucl. Phys. B* 43 (1972). [Erratum: Nucl.Phys.B 101, 547 (1975)], p. 605. DOI: 10.1016/0550-3213(75)90612-4.

[28] O. Benhar et al. "Spectral function of finite nuclei and scattering of GeV electrons". In: *Nucl. Phys. A* 579 (1994), pp. 493–517. DOI: 10.1016/0375-9474(94)90920-2.

[29] Raymond Davis, Don S. Harmer, and Kenneth C. Hoffman. "Search for Neutrinos from the Sun". In: *Phys. Rev. Lett.* 20 (21 May 1968), pp. 1205–1209. DOI: 10.1103/PhysRevLett.20.1205. URL: https://link.aps.org/doi/10.1103/PhysRevLett.20.1205.

[30] K. S. Hirata et al. "Observation of ^{8}B solar neutrinos in the Kamiokande-II detector". In: *Phys. Rev. Lett.* 63 (1 July 1989), pp. 16–19. DOI: 10.1103/PhysRevLett.63.16. URL: https://link.aps.org/doi/10.1103/PhysRevLett.63.16.

[31] M. Cribier. "Results of the whole GALLEX experiment". In: *Nucl. Phys. B Proc. Suppl.* 70 (1999). Ed. by A. Bottino, Alessandra Di Credico, and Piero Monacelli, pp. 284–291. DOI: 10.1016/S0920-5632(98)00438-1.

[32] J. N. Abdurashitov et al. "Solar neutrino flux measurements by the Soviet-American Gallium Experiment (SAGE) for half the 22 year solar cycle". In: *J. Exp. Theor. Phys.* 95 (2002), pp. 181–193. DOI: 10.1134/1.1506424. arXiv: astro-ph/0204245.

[33] Q. R. Ahmad et al. "Direct Evidence for Neutrino Flavor Transformation from Neutral-Current Interactions in the Sudbury Neutrino Observatory". In: *Physical Review Letters* 89.1 (2002). ISSN: 1079-7114. DOI: 10.1103/physrevlett.89.011301. URL: http://dx.doi.org/10.1103/PhysRevLett.89.011301.

[34] T. J. Haines et al. "Calculation of Atmospheric Neutrino-Induced Backgrounds in a Nucleon-Decay Search". In: *Phys. Rev. Lett.* 57 (16 Oct. 1986), pp. 1986–1989. DOI: 10.1103/PhysRevLett.57.1986. URL: https://link.aps.org/doi/10.1103/PhysRevLett.57.1986.

[35] Masayuki Nakahata et al. "Atmospheric Neutrino Background and Pion Nuclear Effect for Kamioka Nucleon Decay Experiment". In: *J. Phys. Soc. Jap.* 55 (1986), p. 3786. DOI: 10.1143/JPSJ.55.3786.

[36] Y. Fukuda et al. "Evidence for oscillation of atmospheric neutrinos". In: *Phys. Rev. Lett.* 81 (1998), pp. 1562–1567. DOI: 10.1103/PhysRevLett.81.1562. arXiv: hep-ex/9807003.

[37] K. Eguchi et al. "First results from KamLAND: Evidence for reactor anti-neutrino disappearance". In: *Phys. Rev. Lett.* 90 (2003), p. 021802. DOI: 10.1103/PhysRevLett.90.021802. arXiv: hep-ex/0212021.

[38] Nobel Foundation. *The Nobel Prize in Physics*. 2009. URL: http://nobelprize.org/nobel_prizes/physics/laureates/2009/index.html.

[39] Boris Kayser. "Neutrino physics". In: *eConf* C040802 (2004). Ed. by Joanne Hewett et al., p. L004. arXiv: hep-ph/0506165.

[40] W. Xu et al. "The Majorana Demonstrator: A Search for Neutrinoless Double-beta Decay of ^{76}Ge". In: *Journal of Physics: Conference Series* 606 (May 2015), p. 012004. DOI: 10.1088/1742-6596/606/1/012004. URL: https://doi.org/10.1088/1742-6596/606/1/012004.

[41] M. Agostini et al. "Results on Neutrinoless Double-β Decay of ^{76}Ge from Phase I of the GERDA Experiment". In: *Phys. Rev. Lett.* 111.12 (2013), p. 122503. DOI: 10.1103/PhysRevLett.111.122503. arXiv: 1307.4720 [nucl-ex].

[42] S. Andringa et al. "Current Status and Future Prospects of the SNO+ Experiment". In: *Adv. High Energy Phys.* 2016 (2016), p. 6194250. DOI: 10.1155/2016/6194250. arXiv: 1508.05759 [physics.ins-det].

[43] A. Gando et al. "Limit on Neutrinoless $\beta\beta$ Decay of ^{136}Xe from the First Phase of KamLAND-Zen and Comparison with the Positive Claim in ^{76}Ge". In: *Phys. Rev. Lett.* 110.6 (2013), p. 062502. DOI: 10.1103/PhysRevLett.110.062502. arXiv: 1211.3863 [hep-ex].

[44] R. Arnold et al. "Probing New Physics Models of Neutrinoless Double Beta Decay with SuperNEMO". In: *Eur. Phys. J. C* 70 (2010), pp. 927–943. DOI: 10.1140/epjc/s10052-010-1481-5. arXiv: 1005.1241 [hep-ex].

[45] A. D. Sakharov. "Violation of CP Invariance, C asymmetry, and baryon asymmetry of the universe". In: *Pisma Zh. Eksp. Teor. Fiz.* 5 (1967), pp. 32–35. DOI: 10.1070/PU1991v034n05ABEH002497.

[46] C. Jarlskog. "A Basis Independent Formulation of the Connection Between Quark Mass Matrices, CP Violation and Experiment". In: *Z. Phys. C* 29 (1985), pp. 491–497. DOI: 10.1007/BF01565198.

[47] Olga Mena and Stephen Parke. "Unified graphical summary of neutrino mixing parameters". In: *Phys. Rev. D* 69 (11 June 2004), p. 117301. DOI: 10.1103/PhysRevD.69.117301. URL: https://link.aps.org/doi/10.1103/PhysRevD.69.117301.

[48] Carlo Giunti and Chung W. Kim. *Fundamentals of Neutrino Physics and Astrophysics*. 2007. ISBN: 978-0-19-850871-7.

[49] P. Langacker et al. "Implications of the Mikheev-Smirnov-Wolfenstein (MSW) Mechanism of Amplification of Neutrino Oscillations in Matter". In: *Nucl. Phys. B* 282 (1987), pp. 589–609. DOI: 10.1016/0550-3213(87)90699-7.

[50] Byron Roe. "Matter density versus distance for the neutrino beam from Fermilab to Lead, South Dakota, and comparison of oscillations with variable and constant density". In: *Phys. Rev. D* 95 (11 June 2017), p. 113004. DOI: 10.1103/PhysRevD.95.113004. URL: https://link.aps.org/doi/10.1103/PhysRevD.95.113004.

[51] R. L. Workman et al. "Review of Particle Physics". In: *PTEP* 2022 (2022), p. 083C01. DOI: 10.1093/ptep/ptac097.

[52] Ivan Esteban et al. "The fate of hints: updated global analysis of three-flavor neutrino oscillations". In: *Journal of High Energy Physics* 2020.9 (Sept. 2020). ISSN: 1029-8479. DOI: 10.1007/jhep09(2020)178. URL: http://dx.doi.org/10.1007/JHEP09(2020)178.

[53] K. Abe et al. "The T2K Experiment". In: *Nucl. Instrum. Meth. A* 659 (2011), pp. 106–135. DOI: 10.1016/j.nima.2011.06.067. arXiv: 1106.1238 [physics.ins-det].

[54] M. A. et al Acero. "Improved measurement of neutrino oscillation parameters by the NOvA experiment". In: *Physical Review D* 106.3 (Aug. 2022). ISSN: 2470-0029. DOI:

10.1103/physrevd.106.032004. URL: http://dx.doi.org/10.1103/PhysRevD.106.032004.

[55] P. Adamson et al. "Improved search for muon-neutrino to electron-neutrino oscillations in MINOS". In: *Phys. Rev. Lett.* 107 (2011), p. 181802. DOI: 10.1103/PhysRevLett.107.181802. arXiv: 1108.0015 [hep-ex].

[56] K. Abe et al. "Measurements of neutrino oscillation parameters from the T2K experiment using 3.6×10^{21} protons on target". In: *Eur. Phys. J. C* 83.9 (2023), p. 782. DOI: 10.1140/epjc/s10052-023-11819-x. arXiv: 2303.03222 [hep-ex].

[57] C. Arpesella et al. "Direct Measurement of the Be-7 Solar Neutrino Flux with 192 Days of Borexino Data". In: *Phys. Rev. Lett.* 101 (2008), p. 091302. DOI: 10.1103/PhysRevLett.101.091302. arXiv: 0805.3843 [astro-ph].

[58] J. Hosaka et al. "Solar neutrino measurements in super-Kamiokande-I". In: *Phys. Rev. D* 73 (2006), p. 112001. DOI: 10.1103/PhysRevD.73.112001. arXiv: hep-ex/0508053.

[59] S. Abe et al. "Precision Measurement of Neutrino Oscillation Parameters with KamLAND". In: *Phys. Rev. Lett.* 100 (2008), p. 221803. DOI: 10.1103/PhysRevLett.100.221803. arXiv: 0801.4589 [hep-ex].

[60] F. P. An et al. "Observation of electron-antineutrino disappearance at Daya Bay". In: *Phys. Rev. Lett.* 108 (2012), p. 171803. DOI: 10.1103/PhysRevLett.108.171803. arXiv: 1203.1669 [hep-ex].

[61] Y. Abe et al. "Improved measurements of the neutrino mixing angle θ_{13} with the Double Chooz detector". In: *JHEP* 10 (2014). [Erratum: JHEP 02, 074 (2015)], p. 086. DOI: 10.1007/JHEP02(2015)074. arXiv: 1406.7763 [hep-ex].

[62] J. K. Ahn et al. "Observation of Reactor Electron Antineutrino Disappearance in the RENO Experiment". In: *Phys. Rev. Lett.* 108 (2012), p. 191802. DOI: 10.1103/PhysRevLett.108.191802. arXiv: 1204.0626 [hep-ex].

[63] K. Abe et al. "Hyper-Kamiokande Design Report". In: (May 2018). arXiv: 1805.04163 [physics.ins-det].

[64] B. Abi et al. "Volume I. Introduction to DUNE". In: *Journal of Instrumentation* 15.08 (Aug. 2020), T08008. DOI: 10.1088/1748-0221/15/08/T08008. URL: https://dx.doi.org/10.1088/1748-0221/15/08/T08008.

[65] B. Abi et al. "Volume III. DUNE far detector technical coordination". In: *Journal of Instrumentation* 15.08 (Aug. 2020), T08009. DOI: 10.1088/1748-0221/15/08/T08009. URL: https://dx.doi.org/10.1088/1748-0221/15/08/T08009.

[66] B. Abi et al. "Volume IV. The DUNE far detector single-phase technology". In: *Journal of Instrumentation* 15.08 (Aug. 2020), T08010. DOI: 10.1088/1748-0221/15/08/T08010. URL: https://dx.doi.org/10.1088/1748-0221/15/08/T08010.

[67] Fengpeng An et al. "Neutrino Physics with JUNO". In: *J. Phys. G* 43.3 (2016), p. 030401. DOI: 10.1088/0954-3899/43/3/030401. arXiv: 1507.05613 [physics.ins-det].

[68] DUNE Collaboration et al. *The DUNE Far Detector Vertical Drift Technology, Technical Design Report*. 2023. arXiv: 2312.03130 [hep-ex].

[69] A. Abed Abud et al. *Deep Underground Neutrino Experiment (DUNE) Near Detector Conceptual Design Report*. 2021. arXiv: 2103.13910 [physics.ins-det].

[70] Mary Convery et al. "Fermilab Accelerator Complex: Status and Improvement Plans". In: (Dec. 2018). DOI: 10.2172/1556948.

[71] Jean-Michel Levy. *Kinematics of an off axis neutrino beam*. 2010. arXiv: 1005.0574 [hep-ex].

[72] S. Amerio et al. "Design, construction and tests of the ICARUS T600 detector". In: *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment* 527.3 (2004), pp. 329–410. ISSN: 0168-9002. DOI: https://doi.org/10.1016/j.nima.2004.02.044. URL: https://www.sciencedirect.com/science/article/pii/S0168900204004966.

[73] C Anderson et al. "The ArgoNeuT detector in the NuMI low-energy beam line at Fermilab". In: *Journal of Instrumentation* 7.10 (Oct. 2012), P10019–P10019. ISSN: 1748-0221. DOI: 10.1088/1748-0221/7/10/p10019. URL: http://dx.doi.org/10.1088/1748-0221/7/10/P10019.

[74] R. Acciarri et al. "The Liquid Argon In A Testbeam (LArIAT) experiment". In: *Journal of Instrumentation* 15.04 (Apr. 2020), P04026–P04026. ISSN: 1748-0221. DOI: 10.1088/1748-0221/15/04/p04026. URL: http://dx.doi.org/10.1088/1748-0221/15/04/P04026.

[75] R. Acciarri et al. "Design and construction of the MicroBooNE detector". In: *Journal of Instrumentation* 12.02 (Feb. 2017), P02017–P02017. ISSN: 1748-0221. DOI: 10.1088/1748-0221/12/02/p02017. URL: http://dx.doi.org/10.1088/1748-0221/12/02/P02017.

[76] D. R. Nygren. "The Time Projection Chamber: A New 4 pi Detector for Charged Particles". In: *eConf* C740805 (1974). Ed. by J Kadyk et al., p. 58.

[77] W.J. Willis and V. Radeka. "Liquid-argon ionization chambers as total-absorption detectors". In: *Nuclear Instruments and Methods* 120.2 (1974), pp. 221–236. ISSN: 0029-554X. DOI: https://doi.org/10.1016/0029-554X(74)90039-1. URL: https://www.sciencedirect.com/science/article/pii/0029554X74900391.

[78] C. Rubbia. *The Liquid Argon Time Projection Chamber: A New Concept for Neutrino Detectors*. Tech. rep. CERN, May 1977.

[79] V Chepel and H Araújo. "Liquid noble gas detectors for low energy particle physics". In: *Journal of Instrumentation* 8.04 (Apr. 2013), R04001–R04001. ISSN: 1748-0221. DOI: 10.1088/1748-0221/8/04/r04001. URL: http://dx.doi.org/10.1088/1748-0221/8/04/R04001.

[80] R. Acciarri et al. "Noise Characterization and Filtering in the MicroBooNE Liquid Argon TPC". In: *Journal of Instrumentation* 12.08 (Aug. 2017), P08003–P08003. ISSN: 1748-0221. DOI: 10.1088/1748-0221/12/08/p08003. URL: http://dx.doi.org/10.1088/1748-0221/12/08/P08003.

[81] S. Himi et al. "LIQUID AND SOLID ARGON, AND NITROGEN DOPED LIQUID AND SOLID ARGON SCINTILLATORS". In: *Nucl. Instrum. Meth.* 203 (1982), pp. 153–157. DOI: 10.1016/0167-5087(82)90623-8.

[82] Damian Göldi. "A Novel Liquid Argon Time Projection Chamber Detector: The ArgonCube Concept". PhD thesis. Bern U., 2018. DOI: 10.7892/boris.116931.

[83] J. Alme et al. "The ALICE TPC, a large 3-dimensional tracking device with fast readout for ultra-high multiplicity events". In: *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment* 622.1 (Oct. 2010), pp. 316–367. ISSN: 0168-9002. DOI: 10.1016/j.nima.2010.04.042. URL: http://dx.doi.org/10.1016/j.nima.2010.04.042.

[84] K. et al Abe. "The T2K experiment". In: *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment* 659.1 (Dec. 2011), pp. 106–135. ISSN: 0168-9002. DOI: 10.1016/j.nima.2011.06.067. URL: http://dx.doi.org/10.1016/j.nima.2011.06.067.

[85] M. Adinolfi et al. "The KLOE electromagnetic calorimeter". In: *Nucl. Instrum. Meth. A* 494 (2002). Ed. by G. V. Fedotovich and B. Khazin, pp. 326–331. DOI: 10.1016/S0168-9002(02)01488-2.

[86] The DUNE Collaboration et al. "Design, construction and operation of the ProtoDUNE-SP Liquid Argon TPC". In: *Journal of Instrumentation* 17.01 (Jan. 2022), P01005. DOI: 10.1088/1748-0221/17/01/P01005. URL: https://dx.doi.org/10.1088/1748-0221/17/01/P01005.

[87] DUNE Collaboration et al. *The DUNE Far Detector Interim Design Report, Volume 3: Dual-Phase Module*. 2018. arXiv: 1807.10340 [physics.ins-det].

[88] Hiroshi Nunokawa, Stephen Parke, and José W.F. Valle. "CP violation and neutrino oscillations". In: *Progress in Particle and Nuclear Physics* 60.2 (Apr. 2008), pp. 338–402. ISSN: 0146-6410. DOI: 10.1016/j.ppnp.2007.10.001. URL: http://dx.doi.org/10.1016/j.ppnp.2007.10.001.

[89] R. M. Bionta et al. "Observation of a neutrino burst in coincidence with supernova 1987A in the Large Magellanic Cloud". In: *Phys. Rev. Lett.* 58 (14 Apr. 1987), pp. 1494–1496. DOI: 10.1103/PhysRevLett.58.1494. URL: https://link.aps.org/doi/10.1103/PhysRevLett.58.1494.

[90] K. Hirata et al. "Observation of a Neutrino Burst from the Supernova SN 1987a". In: *Phys. Rev. Lett.* 58 (1987). Ed. by K. C. Wali, pp. 1490–1493. DOI: 10.1103/PhysRevLett.58.1490.

[91] Manuel Rigger and Zhendong Su. "Intramorphic Testing: A New Approach to the Test Oracle Problem". In: *Proceedings of the 2022 ACM SIGPLAN International Symposium on New Ideas, New Paradigms, and Reflections on Programming and Software*. ACM, Nov. 2022. DOI: 10.1145/3563835.3567662. URL: http://dx.doi.org/10.1145/3563835.3567662.

[92] W.R. Gilks, S. Richardson, and D. Spiegelhalter. *Markov Chain Monte Carlo in Practice*. Chapman & Hall/CRC Interdisciplinary Statistics. Taylor & Francis, 1995. ISBN: 9780412055515. URL: https://doi.org/10.1201/b14835.

[93] Steve Brooks et al. *Handbook of Markov Chain Monte Carlo*. CRC press, 2011. URL: https://doi.org/10.1201/b10905.

[94] Nicholas Metropolis et al. "Equation of State Calculations by Fast Computing Machines". In: *The Journal of Chemical Physics* 21.6 (June 1953), pp. 1087–1092. ISSN: 0021-9606. DOI: 10.1063/1.1699114. eprint: https://pubs.aip.org/aip/jcp/article-pdf/21/6/1087/18802390/1087_1_online.pdf. URL: https://doi.org/10.1063/1.1699114.

[95] W. K. Hastings. "Monte Carlo Sampling Methods Using Markov Chains and Their Applications". In: *Biometrika* 57.1 (1970), pp. 97–109. ISSN: 00063444. URL: http://www.jstor.org/stable/2334940 (visited on 02/28/2024).

[96] A. Gelman, W. R. Gilks, and G. O. Roberts. "Weak convergence and optimal scaling of random walk Metropolis algorithms". In: *The Annals of Applied Probability* 7.1 (1997), pp. 110–120. DOI: 10.1214/aoap/1034625254. URL: https://doi.org/10.1214/aoap/1034625254.

[97] Gareth O. Roberts and Jeffrey S. Rosenthal. "Optimal scaling for various Metropolis-Hastings algorithms". In: *Statistical Science* 16.4 (2001), pp. 351–367. DOI: 10.1214/ss/1015346320. URL: https://doi.org/10.1214/ss/1015346320.

[98] "Note Sur Une Méthode de Résolution des équations Normales Provenant de L'Application de la MéThode des Moindres Carrés a un Système D'équations Linéaires en Nombre Inférieur a Celui des Inconnues. — Application de la Méthode a la Résolution D'un Système Defini D'éQuations LinéAires". In: *Bulletin Geodesique* 2.1 (Apr. 1924), pp. 67–77. DOI: 10.1007/BF03031308.

[99] Cosmia Nebula. *Cholesky decomposition*. [Online; accessed 06-March-2024]. 2023. URL: https://en.wikipedia.org/wiki/Cholesky_decomposition#/media/File:Cholesky_decomposition.png.

[100] Edward Thomas Atkin. "Neutrino Oscillations Analysis at the T2K experiment including studies of new uncertainties on interactions involving final state hadrons". PhD thesis. Imperial Coll., London, 2022. DOI: 10.25560/103333.

[101] Tom Loredo. *Marginalization vs. Profiling*. URL: http://www.hep.fsu.edu/~harry/statistics/loredo_profile.pdf (visited on 03/20/2023).

[102] NVIDIA Corporation. *NVIDIA CUDA C++ Programming Guide*. Apr. 2024.

[103] NERSC. *Perlmutter Architecture*. [Online; accessed 03-May-2024]. URL: https://docs.nersc.gov/systems/perlmutter/architecture/.

[104] Digital Research Alliance of Canada. *Cedar*. [Online; accessed 03-May-2024]. URL: https://docs.alliancecan.ca/wiki/Cedar.

[105] Imperial College HEP Group. *Group Batch System*. [Online; accessed 03-May-2024]. URL: https://www.hep.ph.ic.ac.uk/private/computing/htc.shtml.

[106] Felix Kallenborn et al. "Massively parallel computation of atmospheric neutrino oscillations on CUDA-enabled accelerators". In: *Comput. Phys. Commun.* 234 (2019), pp. 235–244. DOI: 10.1016/j.cpc.2018.07.022.

[107] Prob3pplusplus. *rogerwendell/Prob3plusplus*. Version v3.10. Feb. 2025. URL: https://github.com/rogerwendell/Prob3plusplus.

[108] V. Barger et al. "Matter effects on three-neutrino oscillations". In: *Phys. Rev. D* 22 (11 Dec. 1980), pp. 2718–2726. DOI: 10.1103/PhysRevD.22.2718. URL: https://link.aps.org/doi/10.1103/PhysRevD.22.2718.

[109] Yann LeCun, Yoshua Bengio, and Geoffrey Hinton. "Deep learning". In: *Nature* 521.7553 (May 2015), pp. 436–444. ISSN: 1476-4687. DOI: 10.1038/nature14539. URL: https://doi.org/10.1038/nature14539.

[110] B. Abi et al. "Neutrino interaction classification with a convolutional neural network in the DUNE far detector". In: *Phys. Rev. D* 102.9 (2020), p. 092003. DOI: 10.1103/PhysRevD.102.092003. arXiv: 2006.15052 [physics.ins-det].

[111] Costas Andreopoulos et al. *The GENIE Neutrino Monte Carlo Generator: Physics and User Manual*. 2015. arXiv: 1510.05494 [hep-ph]. URL: https://arxiv.org/abs/1510.05494.

[112] P. Adamson et al. LBNF/DUNE Collaboration. ""Long-Baseline Neutrino Facility (LBNF)/DUNE Conceptual Design Report: Annex 3A$_o$*pt*.". LBNF Beam CDR. 2017. URL: https://docs.dunescience.org/cgi-bin/private/RetrieveFile?docid=4559&%20filename=CDR_Optimized_Beam_Oct02.pdf&version=12.

[113] N. J. Baker et al. "Quasielastic neutrino scattering: A measurement of the weak nucleon axial-vector form factor". In: *Phys. Rev. D* 23 (11 June 1981), pp. 2499–2505. DOI: 10.1103/PhysRevD.23.2499. URL: https://link.aps.org/doi/10.1103/PhysRevD.23.2499.

[114] N. J. Baker et al. "Quasielastic neutrino scattering: A measurement of the weak nucleon axial-vector form factor". In: *Phys. Rev. D* 23 (11 June 1981), pp. 2499–2505. DOI: 10.1103/PhysRevD.23.2499. URL: https://link.aps.org/doi/10.1103/PhysRevD.23.2499.

[115] J. Nieves, I. Ruiz Simo, and M. J. Vicente Vacas. "Inclusive charged-current neutrino-nucleus reactions". In: *Physical Review C* 83.4 (Apr. 2011). ISSN: 1089-490X. DOI: 10.1103/physrevc.83.045501. URL: http://dx.doi.org/10.1103/PhysRevC.83.045501.

[116] P. A. Rodrigues et al. "Identification of Nuclear Effects in Neutrino-Carbon Interactions at Low Three-Momentum Transfer". In: *Phys. Rev. Lett.* 116 (7 Feb. 2016), p. 071802.

DOI: 10.1103/PhysRevLett.116.071802. URL: https://link.aps.org/doi/10.1103/PhysRevLett.116.071802.

[117] M. A. et al. Acero. "Adjusting neutrino interaction models and evaluating uncertainties using NOvA near detector data". In: *The European Physical Journal C* 80.12 (Dec. 2020). ISSN: 1434-6052. DOI: 10.1140/epjc/s10052-020-08577-5. URL: http://dx.doi.org/10.1140/epjc/s10052-020-08577-5.

[118] Dieter Rein and Lalit M. Sehgal. "Neutrino Excitation of Baryon Resonances and Single Pion Production". In: *Annals Phys.* 133 (1981), pp. 79–153. DOI: 10.1016/0003-4916(81)90242-6.

[119] Arie Bodek and Un-ki Yang. "Axial and Vector Structure Functions for Electron- and Neutrino- Nucleon Scattering Cross Sections at all Q^2 using Effective Leading order Parton Distribution Functions". In: (Nov. 2010). arXiv: 1011.6592 [hep-ph].

[120] Brooke Russell. "The 2x2 Demonstrator - A demonstrator for the DUNE ND-LAr Near Detector based on the ArgonCube Design". In: *PoS* TAUP2023 (2024), p. 221. DOI: 10.22323/1.441.0221.