



**Improving Measurements of
Neutrino Mixing with long-baseline
Oscillation Experiments**

Margot MacMahon

UNIVERSITY COLLEGE LONDON

High Energy Physics

Submitted to University College London (UCL) in partial
fulfilment of the requirements for the award of the degree of
Doctor of Philosophy.

Supervisor: **Ryan Nichol**

October 2025

Declaration

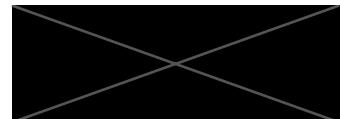
I, Margot MacMahon, confirm that the work presented in this thesis is my own. Where information has been derived from other sources, I confirm that this has been indicated in the thesis.

Chapters 1 and 2 describe the field of neutrino oscillation physics at the time this thesis was written, and the workings of the NOvA experiment, respectively.

The analysis presented in Chapter 3 was the joint effort of the 3-flavour group in the NOvA collaboration. The author performed the work in Sections 3.0.1 and 3.0.7 as part of this effort.

Chapter 4 comprises entirely of original work performed by the author in collaboration with Pedro Machado, Joachim Kopp, and Ivan Martinez-Soler

Chapter 5 includes work performed solely by the author (Sections 5.3 and 5.4) and work performed in collaboration with George Parker, Joachim Kopp and Julia Gehrlein.



Margot MacMahon
London, United Kingdom
3rd October 2025

Abstract

Neutrinos are arguably the least well-understood and most compelling Standard Model particle. The Standard Model does not require they have mass, which they do, albeit on a scale at least six orders of magnitude lighter than other particles. Their mixing properties provide a wealth of information about the lepton sector, in particular the size of leptonic charge-parity violation, which could help explain why we live in a matter dominated universe. In this thesis, we will discuss two experiments, NOvA and DUNE, the former currently running, and the latter set to begin taking data in the coming years. For NOvA, we will briefly describe the workings of the experiment, then focus on optimising cut values and calculating posterior ranges on neutrino event spectra for its latest analysis, which achieved world leading precision for the single experiment measurement of the neutrino mass splitting Δm_{32}^2 . For DUNE we will discuss some phenomenological projects, undertaken in an attempt to maximise the already appreciable sensitivity of DUNE to neutrino oscillation parameters. We will utilise machine learning to demonstrate how the energy resolution could be improved, and quantify the effect of this improvement on its sensitivity. We will also propose a new technique to exploit the capabilities of the DUNE near detector to move off-axis in order to constrain hadron production uncertainties. Again we will show the corresponding improvements this affords on the accuracy and robustness of DUNE’s measurements.

Impact Statement

Although about 100 trillion neutrinos pass through our bodies every second, they are perhaps the least understood particle in the known universe. Yet, the discovery of neutrino oscillations was one of the first indications that the Standard Model of particle physics is incomplete. Studying their properties may also help explain why we live in a world of matter as opposed to antimatter, one of the most pressing questions in modern physics, and help in the search for a grand unified field theory.

The research presented in this thesis was undertaken to deepen our understanding of these fascinating particles, and contribute to answering these fundamental questions. It focuses on the role of current and future long-baseline oscillation experiments, NOvA and DUNE, and explores methods to improve their precision and sensitivity to neutrino oscillation parameters.

Outside its scientific significance, particle physics has also contributed greatly to technologies we rely upon in our day-to-day lives. The World Wide Web was developed at CERN, while novel techniques of information storage and retrieval have been developed for storing huge datasets from particle physics experiments, and particle accelerators are now widely used in medical treatment. Research into neutrinos creates more opportunities for discoveries which benefit both the scientific community, and society as a whole.

Acknowledgements

To my great surprise, this has been the most difficult part of my thesis to write. There are simply too many people to thank for too many things.

Firstly, I would like to thank my supervisor, Ryan Nichol, for always providing excellent advice, (delivered with a trademark pinch of cynicism) throughout my PhD. Thank you also for allowing me to pursue a variety of projects, and helping me with whatever iteration of neutrino physics query I brought to your office door.

My first introduction to neutrino physics came from Jessica Turner, thank you for showing me how fascinating this field is. I am extremely grateful to Pedro Machado and Joachim Kopp, for helping me continue in neutrino phenomenology, and being incredible collaborators throughout multiple projects. I would also like to thank Pedro Machado, Joachim Kopp and Ivan Martinez-Soler for making my first machine learning project so interesting and rewarding. To Joachim Kopp, Julia Gehrlein and George Parker, I appreciated all of our conversations while working on our PRISM project, especially the miscellaneous chats before the official start of meetings.

I am glad to have been a part of the NOvA collaboration throughout my PhD, both for the exciting physics and welcoming people. Thank you in particular to Artur Sztuc, for sharing your knowledge on Bayesian Statistics, all aspects of neutrino physics, and kindly letting me hitch a ride with you at many collaboration meetings.

Thank you to the HEP group at UCL for all the post work pints, which exceeded my expectations as to the social life of a particle physicist, I will miss having you as colleagues. I'm especially grateful to my office mates in C17 for the wide array of mid-work chats and daily puzzles. I am indescribably glad to have found lifelong friends in my fellow students Lucy Bailey and Marcin Jastrebski, who have been fabulous, fun housemates, both in Naperville and London, and made the stressful parts of PhD life infinitely more bearable.

To my Durham friends, thank you for the many holidays, parties unending moral support. To my friends from Northern Ireland, I'm glad to have had a source of sarcasm and wonderful humour both back home, and here in London.

Finally, I am endlessly grateful to Mum, Dad, Izzy and Cece, for your love, support, and valiant attempts to explain to people why neutrinos, which don't appear to do much, are in fact very interesting. I don't think I would have started a PhD, let alone finished it, without you.

Contents

List of Figures	ix
List of Tables	xx
1 Introduction	1
1.1 A short history of neutrinos	1
1.2 Neutrinos in the Standard Model	3
1.2.1 Neutrino Mass	5
1.3 Neutrino Oscillations	6
1.3.1 Neutrino Oscillations in Vacuum	6
1.3.2 Neutrino Oscillations in Matter	9
1.4 Current picture of neutrino oscillations	15
1.4.1 Reactor (θ_{13})	15
1.4.2 Solar ($\theta_{12}, \Delta m_{21}^2$)	16
1.4.3 Atmospheric ($\theta_{23}, \Delta m_{32}^2$)	17
1.4.4 δ_{CP} and Mass Ordering	19
2 NOvA	21
2.1 Physics Goals	21
2.2 NuMI beam	22
2.3 NOvA detectors	26
2.3.1 Near Detector	27
2.3.2 Far Detector	27
2.4 Particle identification	27
2.4.1 Convolutional Visual Network (CVN)	28
2.4.2 Reconstructed Muon Identifier (ReMId)	29

2.4.3	Cosmic Rejection BDTs	30
2.5	Energy Reconstruction	31
2.5.1	Calibration	31
2.5.2	ν_μ Energy Reconstruction	33
2.5.3	ν_e Energy Reconstruction	34
2.6	Simulation	35
3	NOvA 3-Flavour analysis	36
3.0.1	ν_μ PID Cut Optimisation	36
3.0.2	Analysis Binning	41
3.0.3	Decomposition and Extrapolation	44
3.0.4	Systematic Uncertainties	47
3.0.5	Oscillation parameter extraction	49
3.0.6	Bayesian Analysis	51
3.0.7	2024 Bayesian Analysis Results	54
4	Machine Learning for Neutrino Energy Resolution	62
4.1	DUNE experiment	62
4.2	Event topologies and missing energy	63
4.3	Dataset	65
4.4	Neural Network	68
4.5	Results	70
4.6	Impact on Oscillation Analyses	74
4.7	Dependence on neutrino–nucleus cross section modeling	78
5	Reducing systematic uncertainties with DUNE-PRISM	80
5.1	DUNE-PRISM	80
5.2	Generating fluxes	83
5.3	Fit	84
5.4	Results	87
5.5	Impact on Oscillation Measurements	92
6	Summary and Conclusions	97

List of Figures

1.1	Table of all particles in the Standard Model and their spins, electromagnetic charges and approximate masses, from [20].	3
1.2	Feynman diagrams of charged current (left) and neutral current (right) interactions.	5
1.3	Feynman diagrams of a charged current interaction between a ν_e and electron(left) and a neutral current interaction between any neutrino flavour and an electron, proton, or neutron (right). Only electron neutrinos can interact with the Earth's matter via charged current, meaning they experience a stronger matter potential than muon or tau neutrinos.	10
1.4	Neutrino (blue) and anti-neutrino (red) oscillations in matter (top panels), and in vacuum (bottom panels), in the cases of normal (left) and inverted (right) neutrino mass ordering. Three values of δ_{CP} are shown: minimal violation - $\delta_{CP} = 0, \pi$, in dotted and dashed lines, respectively, and maximal violation in solid lines. In vacuum, neutrino and anti-neutrino oscillations look extremely similar regardless of mass ordering, whereas in matter and normal ordering neutrino oscillations are enhanced and anti-neutrino oscillations suppressed, while in inverted ordering the reverse is true.	15
1.5	Solar neutrino fluxes resulting from different processes in the nuclear fusion reaction chain. Only the 8B and hep processes produce electron neutrinos with energies exceeding 2 MeV. Figure from [43].	16

1.6	Electron neutrino survival probabilities as measured by Borexino (points) and as predicted by the solar model (bands) for solar neutrinos produced via different fusion subprocesses. Taken from [43].	18
1.7	Schematic showing the Normal Ordering (NO) and Inverted Ordering (IO) of the neutrino mass eigenstates. The relative composition of the mass eigenstates in terms of flavour eigenstates is also shown, along with the probability of finding the neutrino flavour α in the i^{th} mass state at different values of δ_{CP} . From [57].	20
2.1	Schematic of the NuMI beamline. From [58].	23
2.2	Left: The flux of neutrinos from a pion with energy E_π seen by a detector located at angle θ from the beam axis, over a distance of 810km. Right: Energy of neutrino produced at angle θ relative to the parent pion as a function of E_π . Taken from [59].	25
2.3	Unoscillated ν_μ charged current events calculated 810km from Fermilab at different off axis angles θ . Left and right plots show spectra for two different beam configurations, low energy and medium energy, respectively. Taken from [59].	25
2.4	Schematic showing the orientation of planes in NOvA's detectors and the two possible event views that this allows. Taken from [60].	26
2.5	Simulated 2.15 GeV neutrino interactions showing ν_μ CC (top) ν_e CC (middle) and NC (bottom). Taken from [61].	28
2.6	Illustration of feature extraction on true ν_μ CC event(left). Event hit map is fed to CVN as a heatmap representing deposited charge, and a series of transformations applied to extract 256 features. A feature can be understood as a variable in the underlying image which is key to classifying the interaction. The features are used to construct the feature map (bottom right). The green, blue and purple features seem to be responses to muon tracks, electromagnetic showers and hadronic showers, respectively. The event shown is a ν_μ CC event, so the response in the blue (electron shower) feature is much weaker than the green. Taken from [62].	29

2.7	ν_μ CC CVN scores for appearance and beam ν_e , survived ν_μ and NC events. The network has excellent discrimination between all event types, with vanishing numbers of non ν_μ CC events scoring higher than 0.5. Taken from [63]	30
2.8	Ratio of reconstructed and true energy of a cosmic ray muon track along the length of an FD cell, with and without relative calibration. Taken from [66].	32
2.9	Distribution of attenuation corrected PE/cm as a function of distance to the end of the track, for all hits in selected stopping muon tracks. The black line shows the mean value for the corrected response in each distance bin. The region between 100 and 200 cm is most uniform, so these hits are used for absolute calibration. Taken from [68]. .	33
2.10	Spline functions for converting reconstructed muon track length to true muon energy (left) and reconstructed visible energy to hadronic energy (right) for FHC mode. Taken from [69].	34
3.1	Flowchart of the NOvA oscillation analysis framework. The upper panels (red) represent the ν_μ analysis, and the lower (blue) panels ν_e . Taken from [74].	37
3.2	Example of a NOvA Near Detector event spectrum with data (black points) and MC prediction (purple line). Expected wrong sign events are shown in green, and a 1σ systematic uncertainty error band in purple. Taken from [75].	37
3.3	ν_μ CC CVN scores for appearance and beam ν_e , survived ν_μ and NC events. At scores of 0.8 and higher the only appreciable contribution to the events comes from survived ν_μ . Taken from [63]	38
3.4	2D histograms showing values of figure of merit against number of background events. Here the FOM is $S/\sqrt{S+B}$ over the full or dip (oscillation maximum) energy region. The upper panels show neutrino mode (FHC) events, and lower RHC.	40
3.5	FHC cut stability plots. The colour axis shows the number of cut combinations with a 'good' FOM for a given value of the X and Y axes.	40

3.6 RHC cut stability plots. The colour axis shows the number of cut combinations with a 'good' FOM for a given value of the X and Y axes.	41
3.7 Signal (ν_μ CC and $\bar{\nu}_\mu$ CC), Background(NC and beam ν_e) and Cosmic energy spectra for the PID cuts used in 2020, and those chosen for 2024. The considerable scale differences between FHC and RHC mode samples are due to NOvA having received mostly neutrino beam during its running time.	42
3.8 Spectra of ν_μ CC, background and cosmic events for each fractional hadronic energy quartile, where Quartile 1 has the best resolution (and smallest background and cosmic rates), and Quartile 4 has the worst resolution, and highest background and cosmic rates. The average energy resolution in the FD for each quartile is 6.2%, 8.2%, 10.3% and 12.4% respectively. Taken from [76].	44
3.9 Spectra of ν_e CC, background and cosmic events in low energy, low PID, high PID and peripheral bins. Only the number of events in the peripheral bin are used in the analysis - because these events fail containment cuts, accurate energy reconstruction is not guaranteed and could skew results, hence it is not binned in reconstructed neutrino energy, unlike the other bins. Taken from [76].	45
3.10 Reconstructed neutrino energy distribution of 2024 data events (black crosses), for ν_e and $\bar{\nu}_e$ appearance (top), and ν_μ and $\bar{\nu}_\mu$ disappearance (bottom), without Daya Bay constraints applied. The ν_e samples are divided into bins of low (I) and high (II) PID scores, plus the peripheral sample (III). In the ν_e sample only, Θ indicates the low energy sample which was added for 2024 analysis. The coloured bands correspond to the range of 1σ (darkest), 2σ and 3σ (lightest) confidence intervals from the posterior distribution resulting from our fit to data (including oscillation and systematic parameters, but not statistical error.) . . .	56

3.11 Posterior predictive p -values from real data MCMC samples with the $1D$ reactor constraint applied to θ_{13} . The purple distribution shows the scatterplot of the χ^2 computed between the observed data and predicted datasets pulled from the posterior probability distribution against a similar χ^2 between these predicted datasets and corresponding pseudodata sets, both divided by the number of degrees of freedom in the fit. The dashed and solid contours show the 1σ intervals from the same posterior-predictive distributions calculated only for ν_e (dark blue, solid), $\bar{\nu}_e$ (red, solid), ν_μ (light blue, dashed) and $\bar{\nu}_\mu$ (light red, dashed)	57
3.12 90% Credible Intervals for NOvA and other experiments in the Δm_{32}^2 - $\sin^2 \theta_{23}$ space when considering only Normal (left) or Inverted (right) ordering points. To find the best fit point both orderings are considered, this point is marked with a star, and lies in normal ordering, upper θ_{23} octant. From [88].	58
3.13 68% Confidence Intervals for NOvA and other experiments in the $\sin^2 \theta_{23}$ - δ_{CP} space when considering only Normal (left) or Inverted (right) ordering points. The best fit point lies in NO, upper octant, at $\delta_{CP} = 0.93\pi$. From [88].	59
3.14 Marginalised posterior probability densities for Δm_{32}^2 from a fit to data with different Daya Bay constraints applied-left without constraint, the middle with θ_{13} constraint, and the rightmost with 2D Δm_{32}^2 - θ_{13} constraint. NOvA measurements are complementary to Daya Bay, with the normal ordering preference increasing as constraints are applied. Figure from [87].	59
3.15 Bievent plot showing ν_e against $\bar{\nu}_e$ event counts (black point) compared to the expected value for different values of δ_{CP} in normal (blue) and inverted (red) mass ordering. From [88].	60

3.16 Posterior probability densities for the Jarlskog invariant, marginalized separately over each ordering, and shown for normal ordering (blue, left) and inverted ordering (red, right). On each plot, the top panel shows the posterior with a prior uniform in δ_{CP} and the prior uniform in $\sin \delta_{CP}$ in the upside-down canvas. A Jarlskog invariant of 0 represents CP conservation, and non-zero values represent CP violation. The posterior is extracted from a fit with an external 1D constraint on θ_{13} from the Daya Bay experiment. Figure from [87].	61
4.1 Schematics for four main neutrino-nucleon ($\nu - N$) charged current interaction mechanisms. $2p2h$ stands for two-particle two-hole.	64
4.2 Charged current neutrino-liquid argon cross sections used by the GENIE [73] event generator for ν_μ (top left), $\bar{\nu}_\mu$ (top right). There is no appreciable difference between muon and electron (anti)neutrino cross sections. Antineutrino cross sections tend to be a factor of three lower than the corresponding neutrino cross section due to the parity selecting nature of the weak force.	65
4.3 Example of a Pearson correlation matrix between several observables in charged-current neutrino-argon interactions with at least one proton and one neutron in the final state, for the DUNE neutrino-mode flux and before including detector responses. We include the total kinetic energy of all protons (K_p) and neutrons (K_n) per event, the energies of the neutrino and outgoing lepton (E_ν, E_ℓ), the angle between a given particle and the beam axis $\cos \theta_{\ell,p,n}$, and the angles between particles $\cos \theta_{\ell p, \ell n, p n}$. Note that neutrons are very challenging to reconstruct, so information on the neutron system is typically not available in realistic event records.	66
4.4 Fraction of final state protons with $E < 1500$ MeV per event for NuWro (pink) and GENIE (blue) datasets. GENIE predicts higher numbers of final state low energy protons than NuWro.	69

4.5 DNN loss functions when trained on beam neutrino events (upper panel) and atmospheric neutrino events (lower panel) for training (darker) and validation (brighter) data. These results are for the no neutron reconstruction scenario, but results are similar for the neutron energy and neutron energy and direction reconstruction cases.	69
4.6 Normalized event distribution using DNN energy reconstruction (filled histogram) and the calorimetric method (unfilled histogram). Neutrinos with true energies uniform across six ranges, [500 – 600, 1000 – 1100, 2000 – 2100, 3000 – 3100, 4000 – 4100, 5000 – 5100 MeV], were used. The DNN reconstruction produces a narrower event distribution and reduces the bias in the mean reconstructed energy compared to the calorimetric method.	70
4.7 <i>Top</i> : fractional neutrino energy resolution $\sigma(E_\nu)/E_\nu$ as a function of neutrino energy from DNN-based analyses trained on muon neutrino (left) and antineutrino (right) events with no information on final-state neutrons (solid purple), with limited information on the neutron energy (dashed purple), and with information on the neutron energy and direction (dotted purple). When information on the energy of final-state neutrons is available, the improvement is more than a factor of two at high energies. The energy resolutions anticipated in the DUNE CDR and TDR simulations can be seen in magenta, and the resolution of a simple calorimetric method assuming invisible neutrons in grey. <i>Bottom</i> : reconstruction bias for the DNN compared to the calorimetric method.	71

4.8 *Top:* fractional neutrino energy resolution $\sigma(E_\nu)/E_\nu$ as a function of neutrino energy from DNN-based analyses trained on electron neutrino (left) and antineutrino (right) events with no information on final-state neutrons (solid purple), with limited information on the neutron energy (dashed purple), and with information on the neutron energy and direction (dotted purple). When information on the energy of final-state neutrons is available, the improvement is more than a factor of two at high energies. The energy resolutions anticipated in the DUNE CDR and TDR simulations can be seen in magenta, and the resolution of a simple calorimetric method assuming invisible neutrons in grey. *Bottom:* reconstruction bias for the DNN compared to the calorimetric method. 72

4.9 Fraction of ν_μ CC event energy taken up by neutron as a function of the energy of the incident neutrino. In most ν_μ CC events a proton is produced, but in those events with a neutron, the fraction of energy it takes up is not constant. At energies greater than 4000 MeV there are far fewer events in which the neutron energy accounts for more than 40% of the true neutrino energy, making reconstruction less challenging for the DNN. 74

4.10 Fractional neutrino energy resolution, $\sigma(E_\nu)/E_\nu$ (top), and angular resolution, $\sigma(\theta_\nu)$ (bottom) from DNN-based analyses of atmospheric neutrino events with no information on final-state neutrons (solid cyan), and with information on neutron energies and directions included (dashed/dot-dashed cyan). For comparison, we also show the resolutions achievable with simple calorimetric methods (gray curves). The dotted black curve in the bottom panel is based on only the charged lepton kinematics, as in Cherenkov detectors at low energy. 75

4.11 Impact of neural network-improved energy resolution on the energy spectrum of ν_μ events in DUNE. Here, S and B denote signal and background. 76

4.12 Impact of neural network-improved energy resolution on precision oscillation measurements in DUNE. The solid lines show the DUNE TDR’s predicted sensitivity, and the dashed lines the sensitivity achievable using the improved energy resolution afforded by the DNN. For the sensitivity to CP violation and to non-maximal θ_{23} , the improvement in sensitivity due to the DNN is equivalent to a $\sim 10\%$ increase in exposure.	77
4.13 Impact of neural network-improved energy resolution on precision measurements of Δm_{31}^2 , θ_{23} , δ_{CP} , and θ_{13} in DUNE. The solid lines show the DUNE TDR’s predicted precision, and the dashed lines the precision achievable using the improved energy resolution afforded by the DNN.	77
4.14 <i>Top</i> : fractional neutrino energy resolution, $\sigma(E_\nu)/E_\nu$ from a DNN trained on NuWro events, and applied to GENIE (blue) and NuWro (purple) events. We compare to the performance of the calorimetric method when applied to GENIE events with no neutron information (grey solid line), and with neutron energy included (dashed grey line). <i>Bottom</i> : Energy reconstruction bias for the DNN trained on the “wrong” neutrino–nucleus interaction model, and for the calorimetric method.	79
5.1 Schematic of the DUNE-PRISM setup [116]. The detector will sit on moveable rails, meaning it can take data <i>at any position</i> between 0 m and 36 m off-axis.	80
5.2 Total neutrino flux in the DUNE/LBNF beam at different off-axis positions. The flux moves towards lower neutrino energies and becomes more peaked as the detector moves further off axis.	81
5.3 Total neutrino flux and the constituent flux components split by parent hadron in the on-axis (top) and maximally off-axis (bottom) positions. On-axis the flux is dominated by π^+ decays, while maximally off-axis the K^+ component has a distinct shape and greater relative importance.	82
5.4 Uncertainties on beam focusing (left) and hadron production (right) in the ν_μ flux at DUNE’s Far Detector. From [100]	83

5.5	Correlation matrix for hadron production normalisation parameters extracted from the fit performed.	86
5.6	Projected constraints on the individual components of the DUNE neutrino flux from a fit to PRISM data in FHC mode (orange) and RHC mode (blue), assuming half of the total 6.6e21 POT are collected in the on-axis position, while the rest are equally split over six different off-axis positions (6.3 m, 12.6 m, 18.9 m, 24.3 m, 30.6 m, 36.0 m, corresponding to off-axis angles out to 0.063 radians).	89
5.7	<i>Left:</i> Error bands for on-axis ND event distribution when using hadron production uncertainty shown in Figure 5.4 (navy) and when using uncertainties derived from LENS (magenta). <i>Right:</i> Comparison of fractional hadron production uncertainty from DUNE and LENS as a function of neutrino energy.	90
5.8	Error bands for each neutrino flavour’s on-axis ND event distribution when using the hadron production uncertainty shown in Figure 5.4 for ν_μ , and a flat 5% uncertainty for $\bar{\nu}_\mu$, ν_e and $\bar{\nu}_e$	91
5.9	Impact of the flux uncertainty on the DUNE far detector event spectrum with and without using the LENS procedure to constrain the flux model. All curves were obtained using PRISM to predict the oscillated far detector fluxes based on near detector data. While the near and far detector data were predicted using the nominal flux model, the PRISM coefficients c_j were derived from biased flux models in which the relative importance of individual flux components r_p were allowed to vary either by 5% (red curves) or within the tighter constraints imposed by our LENS fit to on-axis and off-axis near detector data, see Table 5.3. In each case, we show 100 random realizations. The solid black curve shows for comparison the spectrum based on the nominal flux model with no variation in hadron-contribution normalisation. The spread in possible outcomes is substantially reduced when the LENS constraint is included.	93

5.10 Impact of LENS on DUNE’s sensitivity to leptonic CP violation. For each value of the true δ_{CP} , we show the significance at which the CP-conserving hypothesis ($\delta_{CP} \in \{0, \pi\}$) can be excluded. All curves were obtained using PRISM to predict the oscillated far detector fluxes based on near detector data. While the near and far detector data were predicted using the nominal flux model, the PRISM coefficients c_j are derived from biased flux models in which the relative importance of individual flux components r_p is allowed to vary within the constraints set by LENS (blue contours) or within 5% of the nominal value (no LENS constraint, red). We show 100 random realizations of biased flux models for each case. We see that the spread in possible outcomes is substantially reduced when the LENS constraint is included. The solid black curve shows for comparison the spectrum based on the nominal flux model, with no variation in hadron-contribution normalisation, and the dashed black curve shows the same result, but with no systematic uncertainties. 94

5.11 Impact of LENS on the measurement of the oscillation parameters $(\theta_{13}, \delta_{CP})$ (left) and $(\theta_{23}, \Delta m_{31}^2)$ (right). As in Fig. 5.10, red and blue contours correspond to neutrino flux models with and without the LENS constraint. In all cases, the far detector prediction is based on the PRISM procedure, with the superposition coefficients c_j determined from the biased flux models and the data generated using the nominal model. Solid black contours show the anticipated sensitivity using the nominal flux model with all $r_p = 1$, utilising DUNE’s simulated systematic uncertainties as per [105]. Dashed black contours show results from the same procedure, but without systematic uncertainties. For comparison, we also show the current constraints on the oscillation parameters from NuFit 6.0 [130] in beige with an orange outline. 96

List of Tables

1.1	Results for 3ν mixing parameters from NuFit 6.0 [39], and the dominant experiment type used to constrain each parameter.	16
3.1	FHC : Cut combinations considered for the 2024 Analysis and the corresponding figures of merit and expected event counts for each case.	43
3.2	RHC : Cut combinations considered for the 2024 Analysis and the corresponding figures of merit and expected event counts for each case.	43
3.3	Highest Posterior Density (HPD) points for oscillation parameters in both, Normal and Inverted Mass Orderings (Both MO, NO and IO, respectively). Results from [87].	55
3.4	Bayes factor for normal ordering against inverse ordering hypothesis with no, 1D and 2D reactor constraints applied.	58
4.1	The kinetic energy threshold, momentum resolution coefficient α_p (which is used to calculate the resolution as a function of momentum using $\sigma(p) = \alpha_p \sqrt{p/\text{GeV}}$), and angular resolution for different final-state particles. For neutron reconstruction, we list the three different scenarios. The values in this table were developed in consultation with members of the DUNE collaboration for Ref. [96].	67

4.2	Best fit parameters from fitting the no neutrino information ν_μ DNN and calorimetric fractional energy resolution curves shown in Figure 4.7 to (4.5.1). While the α parameter, which dictates how fast the resolution falls with increasing energy, is similar for calorimetric and DNN methods, the β parameter, which controls the overall normalisation of the energy resolution, is smaller for the DNN case. This implies that the DNN improves energy resolution uniformly across energies when compared to the calorimetric method.	73
5.1	Hadron decay processes and the percentage of neutrinos at the DUNE ND (on-axis) produced by each one.	85
5.2	Systematic uncertainties on cross-section normalisation (a) and tilt (γ) used in the LENS fit. The values used for each interaction type were derived by requiring that they covered the difference between the predicted cross sections as given by GENIE and NuWro. The normalisation cross section uncertainty has then been increased to include a 1.5% uncertainty on detector response and efficiency.	86
5.3	Projected fractional uncertainties on the fit parameters r_p (the normalisation of the various flux components) and the nuisance parameters a_k and γ_k (parameterizing the cross section uncertainties) from a fit to 6.6×10^{21} pot of DUNE near detector data. Results are shown for a neutrino-dominated beam (“forward horn current”) and for an anti-neutrino-dominated beam (“reverse horn current”), and for two different running strategies (equal data-taking time at each of the seven detector positions, or 50% of the data-taking time spent in the on-axis positions and the rest equally distributed among six off-axis positions.)	88

Chapter 1

Introduction

1.1 A short history of neutrinos

Neutrinos were first postulated by Pauli in 1930 as a solution to the seeming violation of energy conservation seen in beta minus (β^-) decay [1]. Prior to his suggestion it was thought that β^- decay occurred via a neutron converting to a proton (i.e. a down quark converting to an up quark via the weak force) within the nucleus, and emitting a single electron. However, in 1914, Chadwick observed that the energy spectrum of β^- decay electrons was continuous as opposed to resembling a delta function, which would be expected for a single body decay [2]. Pauli's solution ensured the conservation of energy and momentum in such a decay, but posed another issue. The particle he proposed was electrically neutral, could be massless, and would have penetration length no shorter than a photon. This combination of characteristics would make the detection of this particle very challenging [3]. Indeed, the first detection of neutrinos didn't occur until 1956. Cowan and Reines, using a nuclear reactor as a source of electron anti-neutrinos ($\bar{\nu}_e$) were able to detect inverse beta decay [4]

$$\bar{\nu}_e + p \rightarrow n + e^+ \quad (1.1.1)$$

thus proving the existence of the neutrino. The experiment (dubbed 'Project Poltergeist'), used detectors consisting of cadmium enriched water tanks sandwiched between scintillator layers. When inverse beta decay occurred, the electron-positron annihilation which followed could be detected by measuring the coincident gamma rays emitted, while the final state neutron would be absorbed by the cadmium in the detector, emitting another gamma ray, providing a unique signature for a $\bar{\nu}_e$ interaction. Their observation of inverse beta decay with a cross section of

$6.3 \times 10^{-44} \text{ cm}^2$ agreed with predicted rates, and provided the first direct neutrino measurement [4].

The weak force was proven to be parity violating, or chiral, by Wu, Lee and Yang in 1956 [5, 6]. If parity were conserved by the weak force (which mediates beta decay), electrons emitted by beta decay of spin-polarised nuclei would be equally likely to travel parallel and antiparallel to the nuclear spin. This follows as spin, an axial vector, is invariant under parity transformations, while velocity is not. However, Wu observed an asymmetry in the spatial distribution of emitted electrons relative to their spin. In fact, electrons were predominantly emitted in the opposite direction to the spin vector, indicating that parity is maximally violated by the weak force. Lee and Yang shared the 1957 Nobel Prize for this discovery, while Wu, (the only woman involved), who designed and carried out the experiment, was not recognised.

Building on parity violation in the weak force, in 1957 Goldhaber et al. used Europium decays to determine that neutrinos are left-handed (their spin points opposite to their momentum) [7]. Because no right-handed neutrinos were observed, it was assumed that neutrinos are massless: a massive particle cannot travel at the speed of light, as such, there always exists a reference frame in which, for example, a left-handed particle will appear right-handed. Therefore, the only scenario allowing all observed neutrinos to be left-handed is one in which they are massless. In the same year, contrary to this result, Pontecorvo suggested that neutrino-antineutrino oscillations could occur [8], analogous to neutral kaon mixing [9], which would require neutrinos to have a non-zero mass, however, such oscillations have never been observed. Upon the detection of the muon neutrino by Lederman et al. in 1962 [10] Maki, Nakagawa and Sakata introduced a model describing lepton mixing [11], which expanded on Pontecorvo's work in 1957, but considered mixing between flavour states as opposed to particles and anti-particles. The theory for neutrino flavour oscillation which we recognise today was formalised in 1967 by Gribov and Pontecorvo [12]. Evidence for neutrino flavour oscillations was provided when solar neutrinos (ν_e 's created by fusion in the Sun's core) were first detected in 1968 by Davis and Bahcall at the Homestake experiment at a rate of $\frac{1}{3}$ predicted by the standard solar model [13, 14]. This deficit was dubbed the 'solar neutrino problem', and motivated additional research into the properties of neutrinos.

When the tau (τ) lepton was detected in 1975, a third generation of neutrino, ν_τ , became seemingly likely. While the existence of 3 distinct neutrinos was confirmed by Z-decays at the LEP experiment in 1989 [15],

Standard Model of Elementary Particles

three generations of matter (fermions)			interactions / force carriers (bosons)	
QUARKS	I	II	III	
mass	$\approx 2.16 \text{ MeV}/c^2$	$\approx 1.273 \text{ GeV}/c^2$	$\approx 172.57 \text{ GeV}/c^2$	
charge	$\frac{2}{3}$	$\frac{2}{3}$	$\frac{2}{3}$	
spin	$\frac{1}{2}$	$\frac{1}{2}$	$\frac{1}{2}$	
up	u	c	t	g
down	d	s	b	γ
electron	e	μ	τ	Z
electron neutrino	ν_e	ν_μ	ν_τ	W
muon neutrino				
LEPTONS				
SCALAR BOSONS				
GAUGE BOSONS VECTOR BOSONS				

Figure 1.1: Table of all particles in the Standard Model and their spins, electromagnetic charges and approximate masses, from [20].

the ν_τ itself wasn't detected until 2000, when the DONUT experiment measured nine tau neutrino interactions [16, 17]. By 2015, the OPERA experiment had seen five further events [18], and in 2024 the IceCube experiment reported the first observation of astrophysical ν_τ , adding seven events to the dataset [19]. To date, only 21 tau neutrinos have been detected.

1.2 Neutrinos in the Standard Model

Under the Standard Model (SM) of particle physics, illustrated by Figure 1.1, neutrinos are electrically neutral, massless particles which can only interact via the weak force. The Standard Model is a renormalizable $SU(3) \times SU(2)_L \times U(1)_Y$ quantum field theory in flat 3+1 spacetime, i.e. not including gravitational effects. $SU(3)$ represents the strong force, which is mediated by gluons, and acts on quarks. $SU(2)_L \times U(1)_Y$ is the electroweak gauge symmetry which undergoes spontaneous symmetry breaking. After this breaking, the gauge bosons manifest as the W^\pm

and Z bosons (mediating the weak force, which acts on all fermions), and the photon (mediating the electromagnetic force, which acts only on charged fermions). While the simplest version of the standard model Lagrangian has been successful at describing many areas of particle physics, it falls short at explaining neutrino masses, and by extension, neutrino oscillations [21].

Neutrinos are only acted upon by the $SU(2) \times U(1)$ electroweak subgroup. We can write the electroweak Lagrangian for neutrinos (choosing to include a Dirac mass term, which will be explained in Section 1.2.1), as

$$\begin{aligned} \mathcal{L} \supset & \sum_{\alpha=e,\mu,\tau} \left[\bar{\nu}_\alpha i\partial^\mu \nu_\alpha + \frac{g}{\sqrt{2}} (\bar{\nu}_{\alpha,L} \gamma^\mu l_\alpha W_\mu^+ + h.c.) + \frac{g}{2 \cos \theta_W} \bar{\nu}_{\alpha,L} \gamma^\mu \nu_{\alpha,L} Z_\mu \right] \\ & - \sum_{\alpha=e,\mu,\tau} (m_{\alpha\beta} \bar{\nu}_{\alpha,L} \nu_{\beta,R} + h.c.) \end{aligned} \quad (1.2.1)$$

where g is the weak coupling constant, θ_W is the Weinberg angle, W and Z are the gauge bosons of the weak force [22]. The first term describes the kinetic energy, of the neutrino field. The second and third terms show the couplings between neutrinos and the weak bosons, which allow interactions between neutrinos and charged leptons via coupling to the W^\pm boson (known as the charged current, or CC), or between neutrinos via coupling to the Z boson (known as the neutral current, or NC) as shown in Figure 1.2.

The fields $\nu_{\alpha,L}$ and $\nu_{\alpha,R}$ are obtained by acting on the spinor field ν_α with the chirality projection operators

$$\nu_{\alpha,L} \equiv P_L \nu_\alpha \equiv \frac{1}{2} (1 - \gamma^5) \nu_\alpha, \quad \nu_{\alpha,R} \equiv P_R \nu_\alpha \equiv \frac{1}{2} (1 + \gamma^5) \nu_\alpha \quad (1.2.2)$$

which are constructed using gamma matrices,

$$\gamma^5 = i\gamma^0\gamma^1\gamma^2\gamma^3 \quad (1.2.3)$$

As discussed in Section 1.1 the weak force is chiral, and only interacts with left-handed particles. That is to say that only left-handed fermions are $SU(2)_L$ doublets and therefore couple to W and Z , while right-handed fermions are $SU(2)_L$ singlets and do not participate in weak interactions. We should note that this would mean a right-handed neutrino field, ν_R , has no interactions with any SM force, so is not technically part of the standard model. Nonetheless, we include it here because it is required to construct a gauge invariant neutrino mass term - the final part of

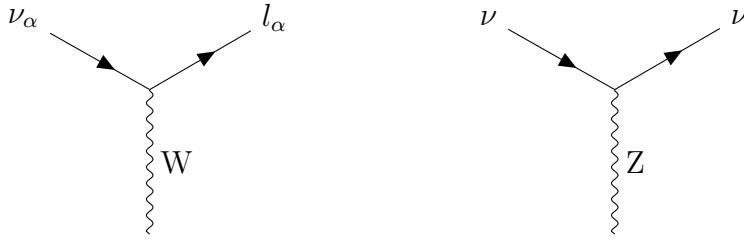


Figure 1.2: Feynman diagrams of charged current (left) and neutral current (right) interactions.

Equation 1.2.1.

1.2.1 Neutrino Mass

The neutrino mass term included in Equation 1.2.1 is analogous to the mass terms used for charged fermions in the Standard Model, known as a Dirac mass term [23]. This term couples fields of opposite chirality together and is the only gauge invariant mass term possible for charged fermions. It arises from the Yukawa interaction with the Higgs field, which generates mass after spontaneous symmetry breaking. While it is plausible that neutrinos acquire mass through this mechanism, it's also possible to construct an alternative mass term which couples a left-handed neutrino and left-handed antineutrino together. This requires the charge conjugation operation \mathcal{C} , which exchanges particles and antiparticles

$$\mathcal{C} : \psi \rightarrow \psi^c \equiv -C\bar{\psi}^T \equiv -i\gamma^2\gamma^0\bar{\psi}^T \equiv -i\gamma^2\psi^*. \quad (1.2.4)$$

By acting on ψ^c with γ^5 we can see that charge conjugation changes the chirality of a field,

$$\gamma^5\psi^c = -i\gamma^5\gamma^2\psi^* = +i\gamma^2\gamma^5\psi^* = -(\gamma^5\psi)^c, \quad (1.2.5)$$

i.e. it interchanges left and right-handed fields. As a result the charge conjugate of a left-handed neutrino field, $(\nu_L)^c$ is right-handed, and we can create an alternative mass term without using the right-handed neutrino field ν_R ,

$$\mathcal{L}_{\text{Majorana}} \supset -\frac{1}{2}m(\bar{\nu}_L)^c\nu_L + h.c. \quad (1.2.6)$$

known as the Majorana mass term [24]. We must note that in the Standard Model, ν_L is part of an $SU(2)_L$ doublet with non-zero weak isospin and hypercharge, meaning Equation 1.2.6 is not invariant under $SU(2)_L \times U(1)_Y$ gauge symmetry. Gauge invariance can be restored us-

ing the dimension-5 Weinberg operator [25], such that the Lagrangian includes

$$\mathcal{L} = \frac{1}{\Lambda} \Phi \bar{L}^c L \Phi + h.c. \quad (1.2.7)$$

where L is the lepton doublet, Φ is the Higgs field with a vacuum expectation value $v = \sqrt{2}\langle\Phi\rangle \approx 246$ GeV and Λ is the energy scale at which some new physics generates the Weinberg operator. After electroweak symmetry breaking (SSB), the Weinberg operator induces a Majorana mass for the active neutrinos of order v^2/Λ . Models which realise neutrino mass generation by the Weinberg operator are known as seesaw mechanisms.

In the type-I seesaw, a heavy right-handed neutrino field N is introduced at energy scales above SSB allowing a Yukawa coupling between N and the left-handed $SU(2)$ doublet, L . Integrating out the heavy neutrino N generates the Weinberg operator, leading to active (left-handed) neutrino masses of $m_\nu \sim v^2/m_N$ (where m_N is on the same scale as Λ). The active neutrino mass is therefore suppressed by the heavy right-handed neutrino, meaning seesaw mechanisms can explain the disparity between the mass of neutrinos and other Standard Model fermions.

It's not clear whether neutrinos are Dirac or Majorana particles, however from the observation of their oscillations we know they must be massive.

1.3 Neutrino Oscillations

We will first derive neutrino oscillations in the 3 flavour picture in a vacuum, before moving on to modify the treatment and take into account propagation through matter.

1.3.1 Neutrino Oscillations in Vacuum

The active neutrino flavour states (i.e. the states which feel the weak interaction) ν_e , ν_μ and ν_τ , are comprised of a superposition of neutrino mass states, which we will label 1, 2 and 3. As discussed in Section 1.1, the relationship between these states can be parameterised by the PMNS matrix [11, 26]

$$\begin{pmatrix} \nu_e \\ \nu_\mu \\ \nu_\tau \end{pmatrix} = \begin{pmatrix} U_{e1} & U_{e2} & U_{e3} \\ U_{\mu 1} & U_{\mu 2} & U_{\mu 3} \\ U_{\tau 1} & U_{\tau 2} & U_{\tau 3} \end{pmatrix} \begin{pmatrix} \nu_1 \\ \nu_2 \\ \nu_3 \end{pmatrix}. \quad (1.3.1)$$

The PMNS matrix is not diagonal, meaning the flavour and mass eigenstates are misaligned, leading to neutrino oscillations. If left-handed neutrinos are the only states, the matrix can be decomposed into 3×3 matrices such that

$$U_{PMNS} = \underbrace{\begin{pmatrix} 1 & 0 & 0 \\ 0 & c_{23} & s_{23} \\ 0 & -s_{23} & c_{23} \end{pmatrix}}_{\text{atmospheric}} \underbrace{\begin{pmatrix} c_{13} & 0 & s_{13}e^{i\delta_{CP}} \\ 0 & 1 & 0 \\ s_{13}e^{-i\delta_{CP}} & 0 & c_{13} \end{pmatrix}}_{\text{reactor}} \underbrace{\begin{pmatrix} c_{12} & s_{12} & 0 \\ -s_{12} & c_{12} & 0 \\ 0 & 0 & 1 \end{pmatrix}}_{\text{solar}} \quad (1.3.2)$$

Where the $s_{ab} \equiv \sin \theta_{ab}$, $c_{ab} \equiv \cos \theta_{ab}$ and δ_{CP} is the Dirac phase. Neutrino flavour states can thus be written

$$|\nu_\alpha\rangle = \sum_{j=1}^3 U_{\alpha j} |\nu_j\rangle \quad (1.3.3)$$

where the mass states $|\nu_j\rangle$ evolve according to the time dependent Schrödinger equation

$$\frac{\partial}{\partial t} |\nu_j\rangle = iH |\nu_j\rangle \quad (1.3.4)$$

where $j = 1, 2, 3$ and the solutions can be approximated as plane waves $|e^{-iE_j t} \nu_j\rangle$, and E is the energy of the neutrino mass state. After propagating over a distance L for time T , the neutrino flavour state becomes

$$|\nu_\alpha(T, L)\rangle = \sum_{j=1}^3 U_{\alpha j}^* e^{-iE_j T} |\nu_j\rangle. \quad (1.3.5)$$

The amplitude for $\nu_\alpha \rightarrow \nu_\beta$, i.e. for an oscillation from ν_α to ν_β is then

$$\langle \nu_\beta | \nu_\alpha(T, L) \rangle = \sum_{j,k=1}^3 U_{\alpha j}^* U_{\beta k} e^{-iE_j T} \langle \nu_k | \nu_j \rangle = \sum_{j=1}^3 U_{\alpha j}^* U_{\beta j} e^{-iE_j T} \quad (1.3.6)$$

where, moving from middle to right, we have used the orthogonality of the mass states, i.e. $\langle \nu_k | \nu_j \rangle = \delta_{jk}$.

In the relativistic limit, $L \sim ct$, (however we use natural units, where $c \equiv 1$) and the energy of the j^{th} neutrino mass state E_j can be rewritten in terms of mass and momentum as

$$E_j = \sqrt{p_j^2 + m_j^2} \simeq p \left(1 + \frac{m_j^2}{2p^2}\right) \simeq E + \frac{m_j^2}{2E}, \quad (1.3.7)$$

where the E is the *total* neutrino energy. The penultimate step is obtained by Taylor expanding, applying the equal momentum approximation (which assumes the mass eigenstates are created with equal momen-

tum) and using the ultrarelativistic limit, such that $p^{-1} \simeq E^{-1}$.

The oscillation probability is then

$$P_{\alpha \rightarrow \beta}(T, L) = |\langle \nu_\beta | \nu_\alpha(T, L) \rangle|^2 = \sum_{j,k=1}^3 U_{\alpha j}^* U_{\beta j} U_{\alpha k} U_{\beta k}^* e^{-i \frac{\Delta m_{jk}^2 L}{2E}}, \quad (1.3.8)$$

with the mass squared splitting $\Delta m_{jk}^2 \equiv m_j^2 - m_k^2$.

Next, we apply the unitarity of the PMNS matrix, particularly that $\sum_j U_{\alpha j} U_{\beta j}^* = \delta_{\alpha \beta}$, and use Euler's Theorem to deconstruct the complex exponential term, then $2 \sin^2(\theta/2) = 1 - \cos \theta$ to write everything in terms of sin, to give

$$\begin{aligned} P_{\alpha \rightarrow \beta} = \delta_{\alpha \beta} - 4 \sum_{j>k} \Re[U_{\alpha j}^* U_{\beta j} U_{\alpha k} U_{\beta k}^*] \sin^2 \frac{\Delta m_{jk}^2 L}{4E} \\ + 2 \sum_{j>k} \Im[U_{\alpha j}^* U_{\beta j} U_{\alpha k} U_{\beta k}^*] \sin \frac{\Delta m_{jk}^2 L}{2E}. \end{aligned} \quad (1.3.9)$$

The oscillation probability therefore depends on the propagation distance (or baseline) L , the neutrino energy, E , the mass squared splittings Δm_{jk}^2 , the mixing angles θ_{jk} and phase δ_{CP} . This result is identical to the one arrived at by more formal (and technically correct) treatments which do not use the equal momentum approximation. One such approach is to treat the eigenstates as quantum mechanical wave packets, and rather than assuming pure momentum eigenstates, to use a superposition of many momentum eigenstates [27, 28].

To obtain the equivalent expression for antineutrinos one must replace $U \rightarrow U^*$, which is equivalent to performing a CP transformation on the neutrino states. If there is CP violation in the lepton sector, this would imply

$$P_{\nu_\alpha \rightarrow \nu_\beta} \neq P_{\bar{\nu}_\alpha \rightarrow \bar{\nu}_\beta}. \quad (1.3.10)$$

The product of charge (C) parity (P) ($\nu_\alpha \xrightarrow{CP} \bar{\nu}_\alpha$) and time (T) ($P_{\nu_\alpha \rightarrow \nu_\beta} \xrightarrow{T} P_{\nu_\beta \rightarrow \nu_\alpha}$) transformations, however, must be conserved in the standard model, meaning for neutrino oscillations we require

$$P_{\nu_\alpha \rightarrow \nu_\beta} = P_{\bar{\nu}_\beta \rightarrow \bar{\nu}_\alpha} \quad (1.3.11)$$

Therefore the observation of CP violation necessarily implies T violation as well.

The difference between neutrino and antineutrino oscillation proba-

bilities,

$$\begin{aligned}
\Delta P_{\alpha \rightarrow \beta}(L) &\simeq \sum_{j,k=1}^3 U_{\alpha j}^* U_{\beta j} U_{\alpha k} U_{\beta k} e^{-i \frac{\Delta m_{jk}^2 L}{2E}} - \sum_{j,k=1}^3 U_{\alpha j} U_{\beta j}^* U_{\alpha k}^* U_{\beta k} e^{-i \frac{\Delta m_{jk}^2 L}{2E}} \\
&= \sum_{j \neq k}^3 (U_{\alpha j}^* U_{\beta j} U_{\alpha k} U_{\beta k}^* - U_{\alpha j} U_{\beta j}^* U_{\alpha k}^* U_{\beta k}) e^{-i \frac{\Delta m_{jk}^2 L}{2E}} \\
&= \sum_{j \neq k}^3 2i \mathbb{I}(U_{\alpha j}^* U_{\beta j} U_{\alpha k} U_{\beta k}^*) e^{-i \frac{\Delta m_{jk}^2 L}{2E}}
\end{aligned} \tag{1.3.12}$$

can be written as $J_{\alpha \beta}^{jk} \equiv \mathbb{I}(U_{\alpha j}^* U_{\beta j} U_{\alpha k} U_{\beta k}^*)$ which is proportional to the *Jarlskog invariant* [29]. For survival probabilities $\alpha = \beta$, $J_{\alpha \beta}^{jk} = 0$, consistent with CPT symmetry. In the three flavour case we can write

$$J = c_{12} s_{12} c_{23} s_{23} c_{13}^2 s_{13} \sin \delta_{CP} \tag{1.3.13}$$

noting that J vanishes if any mixing angle is zero, illustrating that CP-violation is a three flavour effect. This characteristic was first observed in the CKM matrix by Cecilia Jarlskog, but can also be used to quantify CP violation in the lepton sector [30].

Long-baseline oscillation experiments can therefore probe the amount of CP violation in the lepton sector (δ_{CP}), by comparing the oscillation rates of neutrinos and antineutrinos. While this phenomenon has been proven to occur via multiple processes in the quark sector, for example in B_s and D meson decays [31, 32], no such measurement has been made in the lepton sector. CP-violation is a necessary condition for Baryogenesis - the process which produced the matter-antimatter asymmetry in the universe. The current degree of CP-violation observed in the quark sector isn't sufficient to explain the size of this asymmetry, as such, its existence in the lepton sector is both expected and of critical importance to the scientific community.

1.3.2 Neutrino Oscillations in Matter

When neutrinos propagate through matter, their oscillation probabilities are modified by interactions with the particles they encounter. Although incoherent neutrino-nucleon scatterings occur, their cross section is very small and scales linearly with the number of targets N . Therefore, they have a negligible impact on oscillations. In contrast, coherent forward scattering events, wherein the neutrino interacts with the nu-

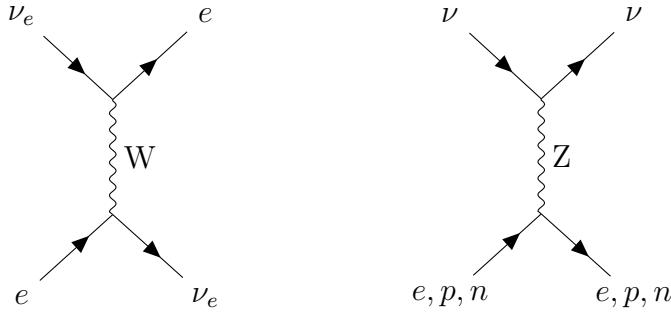


Figure 1.3: Feynman diagrams of a charged current interaction between a ν_e and electron(left) and a neutral current interaction between any neutrino flavour and an electron, proton, or neutron (right). Only electron neutrinos can interact with the Earth's matter via charged current, meaning they experience a stronger matter potential than muon or tau neutrinos.

cleus as a whole, scale with N^2 , as the interference between the phases of constituent nucleons add constructively [33]. This effect introduces a refractive index for the neutrinos, or equivalently an effective matter-induced mass, which alters their flavour evolution. Because ordinary matter consists of protons, neutrons and electrons, all flavours of neutrino can scatter with all matter via NC interactions, and ν_e and $\bar{\nu}_e$ can also scatter with electrons via CC interactions.

To include matter effects in oscillation probabilities, we must construct an effective Hamiltonian describing neutrino-matter coherent scattering. When considering neutrinos with energies $\ll M_W$, it is possible to integrate out the W boson propagator and replace it with $1/M_W^2$, yielding an effective Hamiltonian (for the interaction of ν_e with electrons) of

$$\begin{aligned} H_{\text{eff}} &\supset \frac{g^2}{2M_W^2} [\bar{\nu}_{e,L} \gamma^\mu e_L] [\bar{e}_L \gamma^\mu \nu_{e,L}] \\ &= \frac{G_F}{\sqrt{2}} [\bar{e} \gamma^\mu (1 - \gamma^5) e] [\bar{\nu}_e \gamma^\mu (1 - \gamma^5) \nu_e], \end{aligned} \quad (1.3.14)$$

where the Fermi constant, $G_F \equiv \frac{\sqrt{g^2}}{8M_W^2}$.

The electrons can be treated static relative to neutrinos propagating through matter, meaning we can replace the electron field term in Equation 1.3.14 with its expectation value. This can be calculated by Fourier expanding the electron fields and integrating over their momentum space, assuming they follow a momentum distribution normalised

to the electron number density, N_e . This yields

$$\begin{aligned}\mathcal{H}_{\text{eff}} &\supset \frac{G_F}{\sqrt{2}} N_e [\bar{\nu}_e \gamma^0 (1 - \gamma^5) \nu_e] \\ &= \frac{G_F}{\sqrt{2}} N_e \bar{\nu}_{e,L} \gamma^0 \nu_{e,L}\end{aligned}\quad (1.3.15)$$

where $\frac{G_F}{\sqrt{2}} N_e = V_{CC}$, the potential due to charged current neutrino-electron interactions. For $\bar{\nu}_e$ the sign of V_{CC} is reversed.

The calculation for the neutral current potential for each neutrino flavour is analogous, and gives

$$V_{NC} = -\frac{1}{2} \sqrt{2} G_F N_n \quad (1.3.16)$$

where N_n is the number density of neutrons, and the contributions of protons and electrons cancel because matter is electrically neutral. In fact, because oscillation probabilities are sensitive only to phase differences, and the NC potential contributes equally to all flavours, V_{NC} can be ignored altogether. The Hamiltonian for neutrino oscillation in matter can then be written

$$\mathcal{H} = \mathcal{H}_V + V_{MSW} \quad (1.3.17)$$

where $H_V = U \text{diag}(E_1, E_2, E_3) U^\dagger$ is the vacuum Hamiltonian and V_{MSW} is the Mikhehev Smirnov Wolfenstein potential [34, 35],

$$V_{MSW} = \sqrt{2} G_F \begin{pmatrix} N_e & 0 & 0 \\ 0 & 0 & 0 \\ 0 & 0 & 0 \end{pmatrix}. \quad (1.3.18)$$

The evolution equation of neutrino flavour states in matter is identical to Equation 1.3.4, but uses the modified Hamiltonian in Equation 1.3.17. Considering all three neutrino states,

$$\frac{\partial}{\partial t} \begin{pmatrix} |\nu_e\rangle \\ |\nu_\mu\rangle \\ |\nu_\tau\rangle \end{pmatrix} = i\mathcal{H} \begin{pmatrix} |\nu_e\rangle \\ |\nu_\mu\rangle \\ |\nu_\tau\rangle \end{pmatrix} = \frac{i}{2E} \left[U \begin{pmatrix} m_1^2 & 0 & 0 \\ 0 & m_2^2 & 0 \\ 0 & 0 & m_3^2 \end{pmatrix} U^\dagger + \begin{pmatrix} A & 0 & 0 \\ 0 & 0 & 0 \\ 0 & 0 & 0 \end{pmatrix} \right] \begin{pmatrix} |\nu_e\rangle \\ |\nu_\mu\rangle \\ |\nu_\tau\rangle \end{pmatrix}, \quad (1.3.19)$$

where $A = \sqrt{2} G_F N_e E$, U is the PMNS matrix, and we have used the total flavour state energy E to express the mass state energies E_i , as in Equation 1.3.7, and removed the global factor of E , which affects all flavours identically.

Numerical techniques must be employed to find the modified oscilla-

tion probabilities for an arbitrary matter density profile, however we can find exact solutions for the case where the matter density is constant. This is a very good approximation for long-baseline neutrino oscillation experiments, because the neutrinos only pass through the upper layer of the earth, and don't approach the core. In this case, we diagonalise the Hamiltonian defined in Equation 1.3.17, which yields the effective eigenvalues λ_i and the effective mixing matrix \tilde{U} , such that

$$\tilde{U}^\dagger \mathcal{H} \tilde{U} = \begin{pmatrix} \lambda_1 & 0 & 0 \\ 0 & \lambda_2 & 0 \\ 0 & 0 & \lambda_3 \end{pmatrix}. \quad (1.3.20)$$

The oscillation probability in matter is then

$$\begin{aligned} P_{\alpha \rightarrow \beta} = & \delta_{\alpha\beta} - 4 \sum_{j>k} \mathbb{R}[\tilde{U}_{\alpha j}^* \tilde{U}_{\beta j} \tilde{U}_{\alpha k} \tilde{U}_{\beta k}^*] \sin^2 \frac{(\lambda_j - \lambda_k) L}{4E} \\ & + 2 \sum_{j>k} \mathbb{I}[\tilde{U}_{\alpha j}^* \tilde{U}_{\beta j} \tilde{U}_{\alpha k} \tilde{U}_{\beta k}^*] \sin \frac{(\lambda_j - \lambda_k) L}{2E}. \end{aligned} \quad (1.3.21)$$

Note that the expression for oscillation probability in vacuum, (Equation 1.3.9), can be recovered by replacing $\tilde{U} \rightarrow U$ and $\lambda_j \rightarrow m_j^2$.

Understanding the impact of matter on neutrino oscillations requires solving to the eigenvalues of the modified Hamiltonian, λ . In the three-flavour paradigm this process is non-trivial and the results somewhat difficult to interpret. We will therefore first consider the two-flavour approximation, which is more intuitive.

2 Flavour Approximation

In the 2-flavour picture the mixing matrix simplifies to the standard rotation matrix

$$U_m = \begin{pmatrix} \cos \theta & \sin \theta \\ -\sin \theta & \cos \theta \end{pmatrix} \quad (1.3.22)$$

where θ is the mixing angle between two arbitrary flavour states, ν_α and ν_β . The evolution equation in the flavour basis then becomes

$$\frac{\partial}{\partial t} \begin{pmatrix} |\nu_\alpha\rangle \\ |\nu_\beta\rangle \end{pmatrix} = i\mathcal{H} \begin{pmatrix} |\nu_\alpha\rangle \\ |\nu_\beta\rangle \end{pmatrix} = \frac{i}{2E} \left[U_m \begin{pmatrix} m_1^2 & 0 \\ 0 & m_2^2 \end{pmatrix} U_m^\dagger + \begin{pmatrix} A_\alpha & 0 \\ 0 & A_\beta \end{pmatrix} \right] \begin{pmatrix} |\nu_\alpha\rangle \\ |\nu_\beta\rangle \end{pmatrix}. \quad (1.3.23)$$

The flavour basis Hamiltonian can be written explicitly as

$$\mathcal{H} = \frac{1}{2E} \begin{pmatrix} m_1^2 \cos^2 \theta + A_\alpha & m_1^2 \sin \theta \cos \theta \\ m_1^2 \sin \theta \cos \theta & m_2^2 \sin^2 \theta + A_\beta \end{pmatrix} \quad (1.3.24)$$

and diagonalised to give the effective eigenvalues in matter

$$\begin{aligned} \lambda_1(x) &= \frac{1}{2} \left[m_1^2 + m_2^2 + A_\alpha + A_\beta - \sqrt{(\Delta m^2 \sin 2\theta)^2 + (\Delta m^2 \cos 2\theta - (A_\alpha - A_\beta))^2} \right] \\ \lambda_2(x) &= \frac{1}{2} \left[m_1^2 + m_2^2 + A_\alpha + A_\beta + \sqrt{(\Delta m^2 \sin 2\theta)^2 + (\Delta m^2 \cos 2\theta - (A_\alpha - A_\beta))^2} \right]. \end{aligned} \quad (1.3.25)$$

Using the two-flavour mixing matrix (Equation 1.3.22) and Equation 1.3.8 we can write the two-flavour oscillation probability

$$P_{\alpha \rightarrow \beta} = \sin^2 2\theta_{\text{eff}} \sin^2 \frac{(\lambda_1 - \lambda_2)L}{4E}. \quad (1.3.26)$$

$(\lambda_1 - \lambda_2)$ is the effective mass splitting Δm_{eff}^2 ,

$$\Delta m_{\text{eff}}^2 = \Delta m^2 \sqrt{[2E(V_\alpha - V_\beta)/\Delta m^2 - \cos 2\theta]^2 + \sin^2 2\theta} \quad (1.3.27)$$

and θ_{eff} the effective mixing angle, (which is derived by requiring it diagonalises the Hamiltonian)

$$\tan 2\theta_{\text{eff}} = \frac{\Delta m^2 \sin 2\theta}{\Delta m^2 \cos 2\theta - 2E(V_\alpha - V_\beta)}. \quad (1.3.28)$$

We can therefore see that in matter, the neutrino mixing angle depends on both the neutrino energy and the matter potential (and by extension the matter density). The mixing angle has a resonance at

$$2E(V_\alpha - V_\beta) = \Delta m^2 \cos 2\theta \quad (1.3.29)$$

at which point $\theta_{\text{eff}} \equiv \pi/4$, meaning the mixing between mass and flavour states is maximal, regardless of the size of the vacuum mixing angle θ . This enhancement in mixing in matter is known as the Mikhehyev-Smirnow-Wolfenstein (MSW) resonance condition. Below the resonance $\theta_{\text{eff}} \simeq \theta$, and far above it, $\theta_{\text{eff}} \rightarrow 0$.

Note that reversing the mass ordering by replacing $\Delta m^2 \rightarrow -\Delta m^2$ changes the effective mixing angle. Oscillations in matter therefore provide sensitivity to the sign of the mass squared splitting, or neutrino mass ordering, whereas oscillations in vacuum can probe only the magnitude of this splitting. This is illustrated in Figure 1.4, which shows the oscil-

lation probability $P(\overleftarrow{\nu}_\mu \rightarrow \overleftarrow{\nu}_e)$ in matter and in vacuum for normal and inverted ordering. In vacuum the oscillation probability is unaffected by the mass ordering.

3 Flavour Treatment

An examination of the various methods of Hamiltonian diagonalisation can be found in [36], however we will use the formalism described in [37], which algebraically diagonalises the 3-flavour Hamiltonian in matter, Equation 1.3.19 to find the eigenvalues

$$\lambda_i^2 = -\frac{2}{3}\sqrt{\alpha^2 - 3\beta} \cos\left(\frac{1}{3} \arccos\left(\frac{2\alpha^3 - 9\alpha\beta + 27\gamma}{2(\alpha^2 - 3\beta)^{3/2}}\right)\right) + m_1^2 - \alpha/3 \quad (1.3.30)$$

where the coefficients are defined as

$$\begin{aligned} \alpha &= 2\sqrt{2}EGN_e + \Delta m_{21}^2 + \Delta m_{31}^2 \\ \beta &= \Delta m_{21}^2 \Delta m_{31}^2 + 2\sqrt{2}EGN_e + [\Delta m_{21}^2 (1 - |U_{e2}|^2) + \Delta m_{31}^2 (1 - |U_{e3}|^2)] \\ \gamma &= 2\sqrt{2}EGN_e \Delta m_{21}^2 \Delta m_{31}^2 |U_{e1}|^2. \end{aligned} \quad (1.3.31)$$

Substituting these values into Equation 1.3.21 gives the full oscillation probabilities in matter, however the algebra is impenetrable and the results difficult to interpret. Instead, for long-baseline oscillations such as NOvA and DUNE, a useful approximation is to expand to second order in the mass ordering parameter $\alpha \equiv \Delta m_{21}^2 / \Delta m_{31}^2$, as in [38]. In this regime, the $\nu_\mu \rightarrow \nu_e$ oscillation probability can be written as

$$\begin{aligned} P_{\mu \rightarrow e} &= \alpha^2 \sin^2 2\theta_{12} c_{23}^2 \frac{\sin^2 A\Delta}{A^2} + 4s_{13}^2 s_{23}^2 \frac{\sin^2(A-1)\Delta}{(A-1)^2} \\ &\quad + 2\alpha s_{13} \sin 2\theta_{12} \sin 2\theta_{23} \cos(\Delta + \delta_{CP}) \frac{\sin A\Delta}{A} \frac{\sin(A-1)\Delta}{A-1}, \end{aligned} \quad (1.3.32)$$

where

$$\begin{aligned} \Delta &\equiv \frac{\Delta m_{31}^2 L}{4E} \\ A &\equiv \frac{2EV}{\Delta m_{31}^2}. \end{aligned} \quad (1.3.33)$$

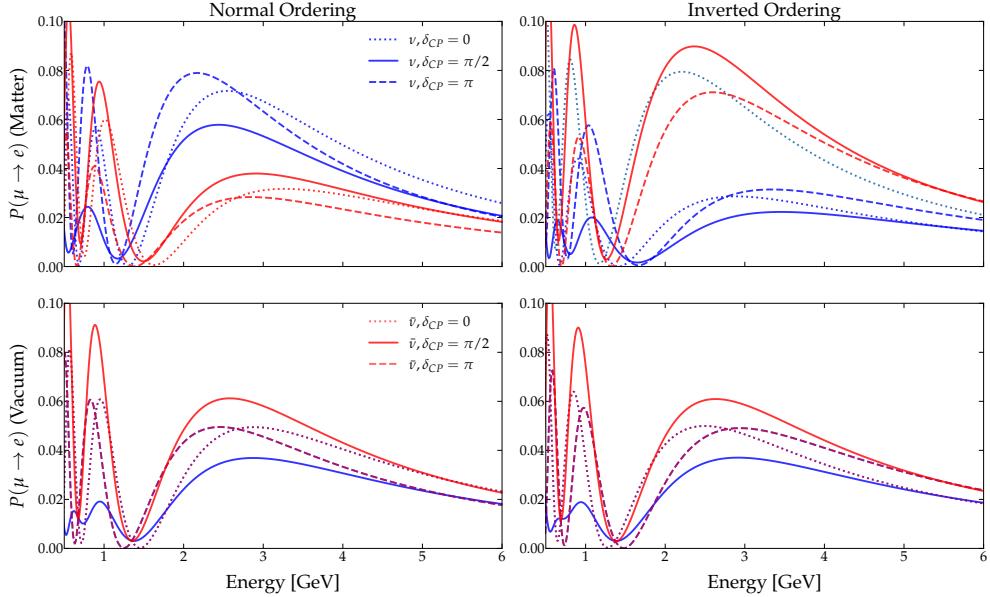


Figure 1.4: Neutrino (blue) and anti-neutrino (red) oscillations in matter (top panels), and in vacuum (bottom panels), in the cases of normal (left) and inverted (right) neutrino mass ordering. Three values of δ_{CP} are shown: minimal violation - $\delta_{CP} = 0, \pi$, in dotted and dashed lines, respectively, and maximal violation in solid lines. In vacuum, neutrino and anti-neutrino oscillations look extremely similar regardless of mass ordering, whereas in matter and normal ordering neutrino oscillations are enhanced and anti-neutrino oscillations suppressed, while in inverted ordering the reverse is true.

1.4 Current picture of neutrino oscillations

Table 1.1 shows the current best fit values for neutrino oscillation parameters dependent on the mass ordering, from [39]. It also details which type of neutrino experiment has the strongest sensitivity to each parameter.

1.4.1 Reactor (θ_{13})

The smallest, and now most accurately measured mixing angle θ_{13} was only verified to be non-zero in 2012. After the first hints of this result from Double-Chooz [40], Daya Bay reported the exclusion of the zero θ_{13} hypothesis at 5.2σ via the observation of reactor electron antineutrino disappearance [41] i.e. the oscillation $\bar{\nu}_e \rightarrow \bar{\nu}_{\mu,\tau}$, which has probability

$$P_{(\bar{\nu}_e \rightarrow \bar{\nu}_e)} \approx 1 - \sin^2 2\theta_{13} \sin \frac{\Delta m_{31}^2 L}{4E_\nu} \quad (1.4.1)$$

Reno in South Korea quickly confirmed this result, observing such oscillations at 4.9σ [42]. Each of these experiments used a source of antineutri-

Table 1.1: Results for 3ν mixing parameters from NuFit 6.0 [39], and the dominant experiment type used to constrain each parameter.

Parameter	Normal Ordering		Inverted Ordering		Dominant source
	BFV $\pm 1\sigma$	3σ range	BFV $\pm 1\sigma$	3σ range	
θ_{12}	$33.68^{+0.73}_{-0.70}$	$31.63 \rightarrow 35.95$	$33.68^{+0.73}_{-0.70}$	$31.63 \rightarrow 35.95$	Solar
θ_{13}	$8.52^{+0.11}_{-0.11}$	$8.18 \rightarrow 8.87$	$8.58^{+0.11}_{-0.11}$	$8.24 \rightarrow 8.91$	Reactor
θ_{23}	$48.5^{+0.7}_{-0.9}$	$41.0 \rightarrow 50.5$	$48.6^{+0.7}_{-0.9}$	$41.4 \rightarrow 50.6$	Accelerator
$\frac{\Delta m_{21}^2}{10^{-5}\text{eV}^2}$	$7.49^{+0.19}_{-0.19}$	$6.92 \rightarrow 8.05$	$7.49^{+0.19}_{-0.19}$	$6.92 \rightarrow 8.05$	Reactor
$\frac{\Delta m_{31}^2}{10^{-3}\text{eV}^2}$	$+2.534^{+0.025}_{-0.023}$	$+2.463 \rightarrow +2.606$	-	-	Reactor
$\frac{\Delta m_{32}^2}{10^{-3}\text{eV}^2}$	-	-	$-2.510^{+0.024}_{-0.025}$	$-2.584 \rightarrow -2.438$	Accelerator
δ_{CP}	212^{+26}_{-41}	$124 \rightarrow 364$	274^{+22}_{-25}	$201 \rightarrow 335$	Accelerator

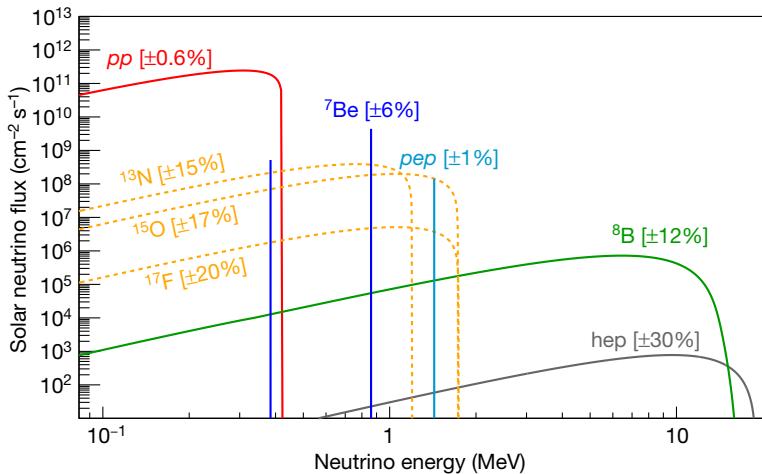


Figure 1.5: Solar neutrino fluxes resulting from different processes in the nuclear fusion reaction chain. Only the 8B and hep processes produce electron neutrinos with energies exceeding 2 MeV. Figure from [43].

nos created in nuclear reactors, and consisted of gadolinium-doped liquid scintillator detectors, with near and far detectors placed $\mathcal{O}(100\text{m})$ and $\mathcal{O}(1\text{km})$ from the source, respectively. In their detectors, $\bar{\nu}_e$'s were detected via the same process as Eq. 1.1.1, with gadolinium increasing the chance of neutron capture as opposed to cadmium, as used by Cowan and Reines [4].

1.4.2 Solar ($\theta_{12}, \Delta m_{21}^2$)

Solar neutrinos are ν_e produced by nuclear fusion reactions in the sun's core, with energies depending on which fusion process has created them, as shown in Figure 1.5. The solar neutrino problem arose when the Homestake experiment (1970) measured the flux of (^7Be and 8B) solar neutrinos to be roughly one third of the predicted value. This discrepancy was resolved almost thirty years later, when it was confirmed that

the deficit was explained by vacuum-like oscillations for lower energy neutrinos, and adiabatic flavour conversion (the MSW effect) for higher energy neutrinos.

8B neutrinos undergo adiabatic flavour conversion because they experience a large matter potential in the sun's core. In the two flavour approximation, which holds well for solar neutrinos, the produced electron neutrino can be written as

$$\nu_e = \cos \theta_{\text{eff}} |\nu_1\rangle + \sin \theta_{\text{eff}} |\nu_2\rangle \quad (1.4.2)$$

where θ_{eff} is, as before, the effective mixing angle in matter. In the sun's core, θ_{eff} is large, so ν_e are produced almost entirely in the $|\nu_2\rangle$ state. As they propagate outward the state remains as $|\nu_2\rangle$, but crosses the MSW resonance region defined in Equation 1.3.29, causing the flavour mixing within $|\nu_2\rangle$ to become maximal. Therefore, at the solar surface, $|\nu_2\rangle$ is a combination of $|\nu_e\rangle$ and $|\nu_\mu\rangle$ [44]. The probability of a ν_e interaction is then approximately

$$P_{ee} \sim |U_{e2}|^2 \sim \sin^2 \theta_{12} \quad (1.4.3)$$

which, from Table 1.1 gives approximately 1/3, agreeing with the deficit observed by Homestake.

Lower energy neutrinos, (e.g. from the pp chain) experience small matter effects, so propagate out of the sun essentially as they would in vacuum. Over the earth-sun baseline, the ν_e survival probability can be written as $1 - \sin^2 2\theta_{12}$, as the $\Delta m_{21}^2 L/4E$ term averages over the long-baseline, yielding $P_{ee} \sim 0.55$.

Borexino [43] has now measured the survival probability of both pp and 8B neutrinos, as shown in Figure 1.6.

Solar neutrino experiments, such as SuperKamiokande [45], SNO [46], and Borexino [47] measure the disappearance and survival rate of solar neutrinos to place constraints on θ_{12} and Δm_{21}^2 .

1.4.3 Atmospheric ($\theta_{23}, \Delta m_{32}^2$)

Atmospheric neutrinos are produced when high-energy cosmic rays [48], primarily protons, collide with nuclei in the earth's atmosphere. These collisions produce cascades of charged mesons; mainly charged pions (π^\pm), and, to a lesser extent due to their higher mass, Kaons (K^\pm) [49]. The charged pions quickly decay into muons and muon neutrinos via $\pi^+ \rightarrow \mu^+ \nu_\mu$ (and charge conjugates), and the resulting muons decay

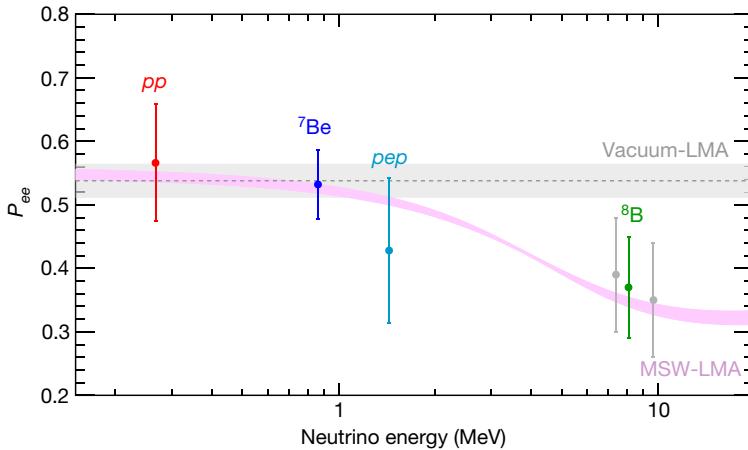


Figure 1.6: Electron neutrino survival probabilities as measured by Borexino (points) and as predicted by the solar model (bands) for solar neutrinos produced via different fusion subprocesses. Taken from [43].

again into electrons and additional muon neutrinos. The more massive kaons decay predominantly into pions and muons, from which more muon (anti)neutrinos and electron (anti) neutrinos are produced.

Although the decay $\pi^+ \rightarrow e^+ \nu_e$ (and its charge conjugate) is kinematically allowed, it is highly spin suppressed due to the electron's low mass. Because the charged pion has spin zero, its decay must emit the final-state lepton and neutrino with opposite spins. The weak interaction couples only to left-handed particles and right-handed antiparticles. The nearly massless neutrino must therefore be produced with left-handed helicity, (in the massless limit helicity and chirality are equivalent). The charged lepton must then be produced with right-handed helicity to conserve spin. For the light electron, helicity and chirality are closely aligned, so producing a right-handed electron is highly suppressed. In contrast, the heavier muon can flip helicity more easily, making the pion-muon decay channel overwhelmingly dominant.

This spin suppression explains why atmospheric neutrinos are produced mainly as $\nu_\mu^{(-)}$, with a smaller fraction of $\nu_e^{(-)}$ from subsequent μ^\pm decay. In absence of oscillations, the expected flavour ratio atmospheric neutrinos the earth's surface would be $\nu_\mu : \nu_e \sim 2 : 1$. As with solar neutrinos, early measurements observed a deficit of muon neutrinos, bolstering the evidence for neutrino oscillation.

Atmospheric neutrino oscillations are governed by the mixing angle θ_{23} and the atmospheric mass squared splitting Δm_{32}^2 . These parameters can be measured by observing atmospheric neutrinos directly, which is the approach taken by experiments such as SuperKamiokande and Ice-

Cube [50, 51]. Long-baseline (LBL) accelerator experiments, such as T2K and NOvA, create beams of muon neutrinos via the decay of pions and kaons, and measure their oscillations over baselines of 295 km and 810 km respectively [52, 53]. LBL experiments access these parameters by measuring the energy spectrum and survival rate (also known as disappearance) of muon neutrinos at the beginning and end of their propagation.

While Δm_{32}^2 and θ_{23} have been measured with reasonable precision, the octant of θ_{23} , i.e. whether it lies above or below $\pi/4$, is unknown. This dictates whether the ν_3 mass state is primarily composed of ν_μ or ν_τ , which can serve to distinguish the neutrino mass ordering in some proposed neutrino mass generation models, such as [54]. This degeneracy also contributes to difficulties in disentangling the effects of mass ordering and CP violation. Resolving the θ_{23} octant is therefore a goal of current and future atmospheric neutrino experiments.

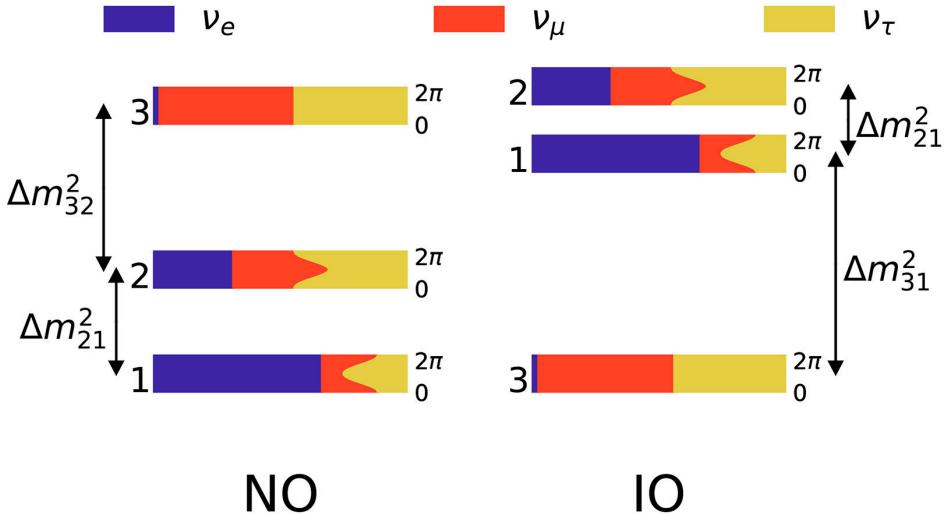
1.4.4 δ_{CP} and Mass Ordering

Long-baseline accelerator experiments can also probe the leptonic CP violating phase, δ_{CP} , by comparing the appearance rates of electron neutrinos and electron antineutrinos in a muon neutrino beam. As discussed in Section 1.3.1 a difference in these probabilities implies CP violation, and is proportional to the Jarlskog invariant. A non-zero value of δ_{CP} would signal that charge-parity is violated in the lepton sector. However, current measurements place only weak constraints on δ_{CP} . This is in part due to limited statistics, but is also due to the observed appearance asymmetry depending not only on δ_{CP} , but also on the neutrino mass ordering. This refers to whether ν_3 is heavier or lighter than ν_1 and ν_2 , as shown in Fig 1.7. As we noted in Sec 1.3.2, matter effects in neutrino oscillations modify the effective mixing angles differently depending on this ordering.

These matter induced modifications make the effects of CP violation and the mass ordering difficult to disentangle. For example as shown in Figure 1.4, even if δ_{CP} is zero, the MSW effect masquerades and CP violation, enhancing $P(\nu_\mu \rightarrow \nu_e)$ and suppressing $P(\bar{\nu}_\mu \rightarrow \bar{\nu}_e)$ in normal ordering, and the reverse in inverted ordering. This degeneracy reduces the sensitivity of current LBL experiments, especially NOvA, due to its baseline and beam energy. This will be discussed in further detail in Chapter 2.

Upcoming long-baseline experiments DUNE and HyperKamiokande

Figure 1.7: Schematic showing the Normal Ordering (NO) and Inverted Ordering (IO) of the neutrino mass eigenstates. The relative composition of the mass eigenstates in terms of flavour eigenstates is also shown, along with the probability of finding the neutrino flavour α in the i^{th} mass state at different values of δ_{CP} . From [57].



are designed with larger detectors than their predecessors, and aim to measure both the mass ordering and CP violation to a high precision [55, 56].

Chapter 2

NOvA

NOvA (NuMI **O**n-Axis ν_e **A**ppearance) is a long-baseline neutrino oscillation experiment based in Fermilab, Illinois, where its near detector sits, and Ash River, Minnesota, which houses the far detector. A beam of mostly muon neutrinos is created by the NuMI beamline at Fermilab and travels 810 km from the near to far detector. Over this distance some ν_μ oscillate to ν_e or ν_τ . By measuring the rate of ν_μ disappearance and ν_e appearance, NOvA can probe neutrino oscillation parameters and the mass ordering. In this chapter we will describe the experiment with a focus on the 3-flavour analysis, specifically that performed in the framework of Bayesian Statistics.

2.1 Physics Goals

NOvA's physics goals are:

- Precisely measure the atmospheric mass splitting Δm_{32}^2 , and in turn determine the neutrino mass ordering.
- Determine the value and octant of θ_{23} .
- Measure δ_{CP} .

The relevant oscillation channels for the NOvA experiment are $\stackrel{(-)}{\nu_\mu} \rightarrow \stackrel{(-)}{\nu_\mu}$ (disappearance) and $\stackrel{(-)}{\nu_\mu} \rightarrow \stackrel{(-)}{\nu_e}$ (appearance). The muon disappearance probability in matter can be taken as the same as in vacuum (muon neutrinos will not coherently scatter off the earth's matter via the charged current, which is the source of matter effect, as discussed in Section 1.3.2). The $\stackrel{(-)}{\nu_\mu} \rightarrow \stackrel{(-)}{\nu_e}$ probability can be simplified from those quoted in Section

1.3.2 to

$$\begin{aligned} R_{\nu_\mu \rightarrow \nu_\mu}^{(-)} &= 1 - \sin^2 2\theta_{23} \sin^2 \Delta_{23} \\ R_{\nu_\mu \rightarrow \nu_e}^{(-)} &= P_{\text{atm}} + P_{\text{solar}} + 2\sqrt{P_{\text{atm}} + P_{\text{solar}}} (\cos \delta_{CP} \cos \Delta_{32} \pm \sin \delta_{CP} \sin \Delta_{32}) \end{aligned} \quad (2.1.1)$$

Where we have defined

$$\begin{aligned} \Delta_{ij} &= \Delta m_{ij}^2 L / 4E \\ \sqrt{P_{\text{atm}}} &\equiv \sin \theta_{23} \sin 2\theta_{13} \sin \left(\frac{\sin(\Delta_{31} - aL)}{\Delta_{31} - aL} \right) \Delta_{31} \\ \sqrt{P_{\text{solar}}} &\equiv \cos \theta_{23} \sin 2\theta_{12} \sin \left(\frac{\sin(aL)}{aL} \right) \Delta_{21} \end{aligned} \quad (2.1.2)$$

and $a \equiv \pm G_F N_e / \sqrt{2}$, the matter potential for neutrinos(positive) and antineutrinos(negative). This simplification takes advantage of θ_{13} being small compared to other mixing angles, meaning $\sin^2 \theta_{13} \sim 0$ and $\cos \theta_{13} \sim 1$, (this approximation is only applied for the disappearance channel). Additionally, terms at order $\sin^2 \theta_{13}$ and higher are dropped in the appearance case.

The disappearance channel provides sensitivity to the atmospheric parameters: the oscillation amplitude (as a function of energy) depends on $\sin^2 2\theta_{23}$, and is used to measure the size of θ_{23} , while the position of the oscillation maximum depends on the size of Δm_{32}^2 . However, it does not distinguish the θ_{23} octant or the sign of the mass splitting.

The appearance channel, however, is sensitive to δ_{CP} , the octant of θ_{23} , and can probe the mass ordering through matter-induced modifications to the oscillation probability. As discussed in Section 1.3.2, the effects of the mass ordering and δ_{CP} are somewhat degenerate, limiting the significance with which NOvA can comment on either phenomenon.

To achieve these goals, NOvA needs the ability to correctly differentiate between different neutrino flavours to measure the disappearance and appearance rates and excellent neutrino energy resolution, to ensure the depth and position of the oscillation maximum is measured accurately.

2.2 NuMI beam

The muon neutrinos studied by NOvA are supplied by the Neutrinos at the Main Injector (NuMI) beam at Fermilab. The beam originates as 400 MeV H^- ions, which are accelerated, stripped down to protons upon entering the booster, where they are accelerated and passed into

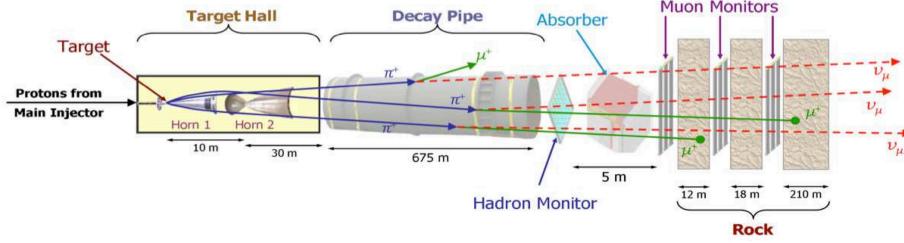


Figure 2.1: Schematic of the NuMI beamline. From [58].

the Main Injector to be boosted to 120 GeV. These protons are extracted in groups of 5×10^{13} , known as spills, and directed towards NOvA's 1.2 m graphite target located 350 m downstream, as shown in Figure 2.1.

Proton-carbon interactions in the target produce a hadron shower comprised mainly of pions and kaons. Two magnetic focusing horns select the charge sign of the hadrons which enter the steel decay pipe, allowing either a neutrino or antineutrino beam to be made further downstream. In the decay pipe the mesons decay predominantly via

$$\begin{aligned} \pi^{+(-)} &\rightarrow \mu^{+(-)} + \bar{\nu}_\mu \\ K^{+(-)} &\rightarrow \mu^{+(-)} + \bar{\nu}_\mu \end{aligned} \quad (2.2.1)$$

A hadron monitor at the end of the decay pipe records the profile of the residual hadrons (protons and undecayed mesons), which are attenuated by an absorber before the beam, now containing primarily muons and neutrinos, passes through the muon flux monitors. These are placed within 240 m of rock, which stops most muons. The beam of neutrinos then passes through the near detector (ND) and travels 810 km to the far detector (FD). In Forward Horn Current (FHC) mode, which selects positive sign hadrons, the beam (the product of the neutrino flux and associated interaction cross section) is 94% ν_μ at the ND, while in Reverse Horn Current (RHC) mode, the beam is 93% $\bar{\nu}_\mu$. Beam contaminants include wrong sign $\bar{\nu}_\mu$ and $\bar{\nu}_e$ which are produced by

$$\begin{aligned} \mu^+ &\rightarrow e^+ + \nu_e + \bar{\nu}_\mu \\ K^+ &\rightarrow \pi^0 + e^+ + \nu_e \end{aligned} \quad (2.2.2)$$

at branching fractions far lower than the ν_μ producing processes in (2.2.1). The beam intensity is quantified by the number of protons delivered to the NOvA target (POT). The number of neutrinos produced is proportional to the POT.

Both NOvA detectors are placed 14.6mrad off the beam axis. This selects a narrower band of energies, peaked around the oscillation maximum for $P(\nu_\mu \rightarrow \nu_e)$ over an 810km baseline. Pions and kaons decay isotropically in their rest frame to produce monoenergetic neutrinos. In the lab frame, however, the mesons and their daughter neutrinos are boosted, meaning the detectors see a broad neutrino energy spectrum, as shown in the left-hand panel of Figure 2.2. The flux (F) and energy (E_ν) of neutrinos produced by a boosted $\pi^+ \rightarrow \mu^+ + \nu_\mu$ decay seen by a detector of area A at distance Z from the decay can be described, for small angles θ , by

$$F = \left(\frac{2\gamma}{1 + \gamma^2\theta^2} \right)^2 \frac{A}{4\pi z^2} \quad (2.2.3)$$

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

where θ is the angle between the pion and neutrino direction, E_π is the pion energy, m_π the pion mass, and $\gamma = \frac{E_\pi}{m_\pi}$ the pion's Lorentz factor. The expression for kaons is identical, however m_π is replaced by m_K . The kaon is 3.5 times more massive than the pion, meaning the energy spectrum for neutrinos from kaon decay is shifted towards higher energies and is broader.

The relationship between pion and neutrino energy for 4 different angles can be seen in the right-hand pane of Figure 2.2. On axis, the neutrino energy has a strong dependency on the pion energy, however moving off axis this dependency becomes weaker, leading to a narrower range of neutrino energies. This can also be seen in Figure 2.3, where the peak of the unoscillated muon neutrino flux becomes sharper as the off axis angle increases.

We note that at 14mrad off axis, the number of ν_μ 's at the oscillation maximum energy is about 5× higher than on axis. Positioning the far detector 14mrad off axis therefore maximises the chance a ν_μ will oscillate by the time it reaches the far detector. It also acts as a natural veto, particularly for background events caused by feed down from high energy NC interactions, as the acceptance window for neutrino energy is narrowed by moving off axis.

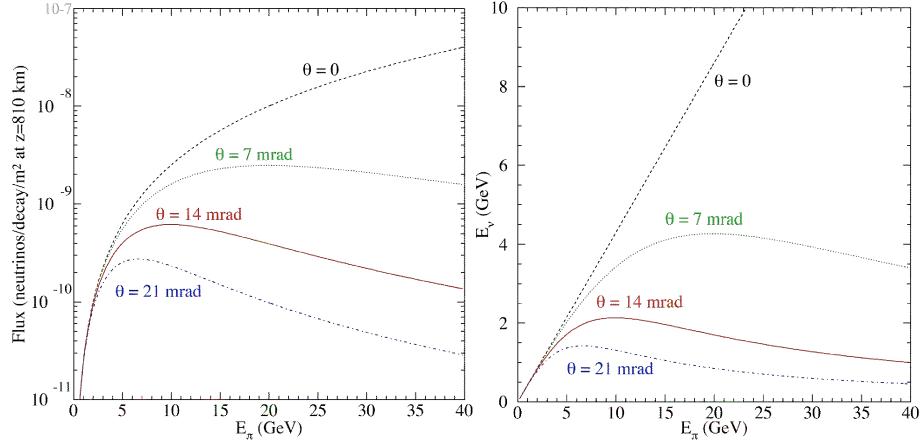


Figure 2.2: Left: The flux of neutrinos from a pion with energy E_π seen by a detector located at angle θ from the beam axis, over a distance of 810km. Right: Energy of neutrino produced at angle θ relative to the parent pion as a function of E_π . Taken from [59].

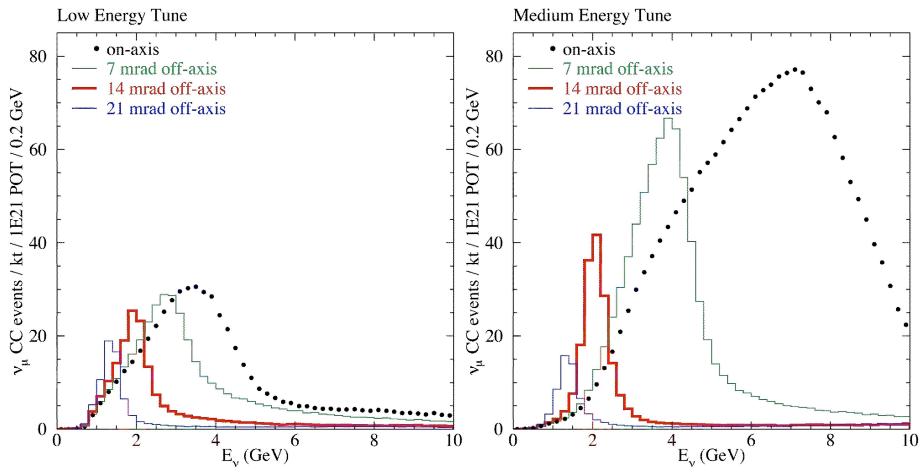


Figure 2.3: Unoscillated ν_μ charged current events calculated 810km from Fermilab at different off axis angles θ . Left and right plots show spectra for two different beam configurations, low energy and medium energy, respectively. Taken from [59].

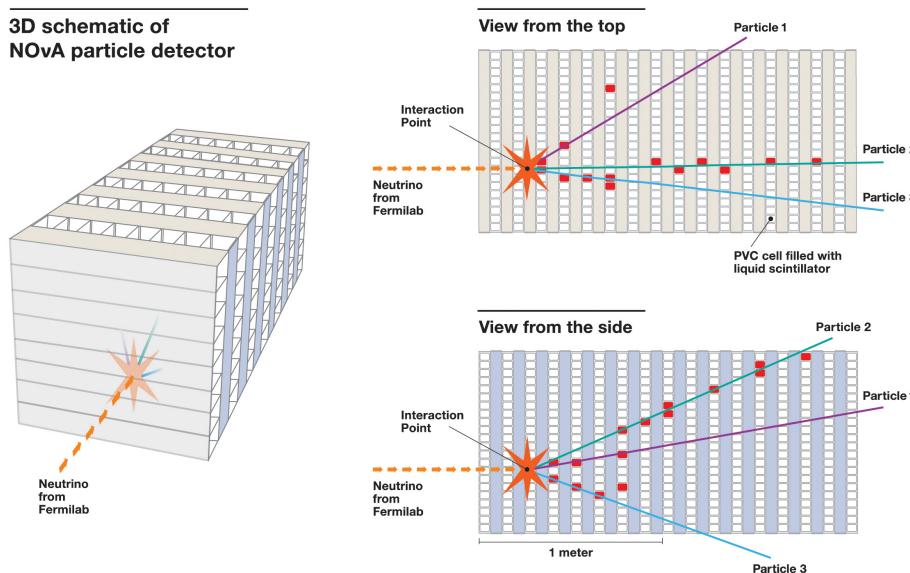


Figure 2.4: Schematic showing the orientation of planes in NOvA’s detectors and the two possible event views that this allows. Taken from [60].

2.3 NOvA detectors

The NOvA experiment consists of two functionally identical detectors, composed of stacked, highly reflective PVC cells filled with liquid scintillator. Charged particles passing through the scintillator produce light, which is collected by wavelength shifting fibres (WLS) and amplified and read out by avalanche photodiodes (APDs). Events occurring within $\pm 15\ \mu\text{s}$ of the NuMI spill time are recorded by NOvA’s Data Acquisition system. Hits (readouts from the detector) which are light-like intervals apart are grouped into slices, which represent event candidates, while isolated hits are labelled as noise. In NOvA, the slice object is often referred to as an event.

The configuration of both the near and far detectors can be seen in Figure 2.4. Each detector consists of planes, (made by sealing 32 scintillator filled plastic cells into modules), which are stacked together. Alternating the orientations of the planes allows 3D tracking within NOvA’s detectors.

NOvA uses the higher statistics at the Near Detector to constrain predictions at the Far Detector.

2.3.1 Near Detector

The near detector sits underground at the bottom of the MINOS shaft at Fermilab. The rock which it sits beneath provides a shield from cosmic rays, which would otherwise form a large background. Its active (fiducial) volume is 222 tons, with dimensions 2.9 m wide, 4.2 m tall, and 14.3 m long. The detector comprises 186 planes, each 96 cells high and 64 cells wide. 10 additional planes alternating with planes of steel are placed at the end of the detector where the beam exits, in order to tag muons exiting the detector's fiducial volume.

2.3.2 Far Detector

The far detector sits 810 km downstream from the ND in Ash River, Minnesota, has a fiducial volume of 15 kton and dimensions 15.6 m wide, 15.6 m tall and 78 m long. Unlike the ND, it sits closer to the Earth's surface, under a small overburden of rocks. The far detector is therefore constantly bombarded by cosmic rays (in particular cosmic muons). As a result, the timing resolution at the far detector must be sufficient to differentiate between beam neutrinos (in the spill window), and neutrinos produced by cosmic ray muons.

2.4 Particle identification

NOvA measures two features of interacting neutrinos, their energy and, for charged current interactions, their flavour. NOvA determines the neutrino flavour by identifying which charged lepton was produced when said neutrino interacted. Examples of how CC and NC interactions appear in NOvA can be seen in Figure 2.5. Muons from ν_μ CC interactions produce long, straight tracks which end with a muon decay. Electrons from ν_e CC interactions behave differently, producing photons which cause electromagnetic showers very close to the interaction vertex. NC interactions preserve the incident neutrino (which is invisible to the detector) but excite the nucleus involved, producing a hadronic shower. As it is not possible to distinguish between ν_μ and ν_e NC interactions, the oscillation analysis takes such events as background to both ν_μ and ν_e CC events. In fact, NC events are easily distinguished from ν_μ CC events unless they contain a charged pion, which can mimic a muon track. However, the hadronic shower of an NC event can be misidentified as an electron when a π^0 is produced, as the two photons from its decay produce electromag-

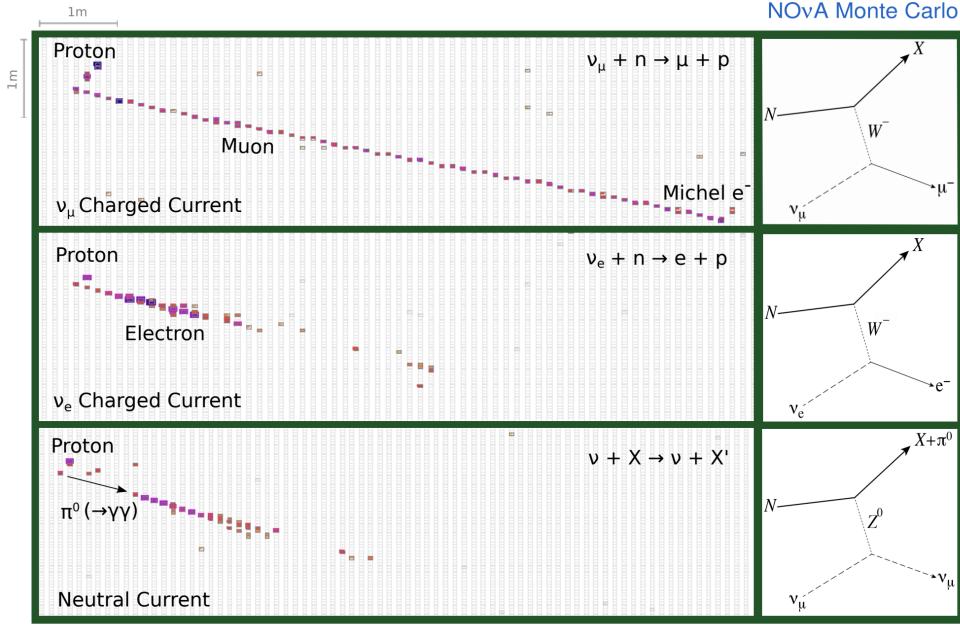


Figure 2.5: Simulated 2.15 GeV neutrino interactions showing ν_μ CC (top) ν_e CC (middle) and NC (bottom). Taken from [61].

netic showers.

The events shown in Figure 2.5 are reconstructed by identifying *slices*: hits (detector readouts) which are clustered together in spacetime. Hits which are light-like time intervals apart (the particles in question will be travelling close to the speed of light) are considered part of the same slice, while isolated hits are labeled as noise. In NOvA, the slice object contains all of the information in an event, so it is often referred to as an event.

In the following section we will briefly review the techniques used by NOvA to discriminate between different event types.

2.4.1 Convolutional Visual Network (CVN)

NOvA utilises a convolutional neural network based tool for particle identification, known as CVN. CNN's are particularly well suited to image recognition tasks due to their ability to parse 2D, or grid-like, datasets, meaning they can be directly applied to hit maps from NOvA. As shown in Figure 2.4, each event has an x and y view. As such, the CVN takes as input two 100×80 pixel hit maps per event. Directly learning features prevents useful information from being discarded. An example of a true ν_μ CC event and the feature maps extracted from it after the first module of the CVN can be seen in Figure 2.6. Feature maps are the results of applying different convolutional filters to the image. These filters activate

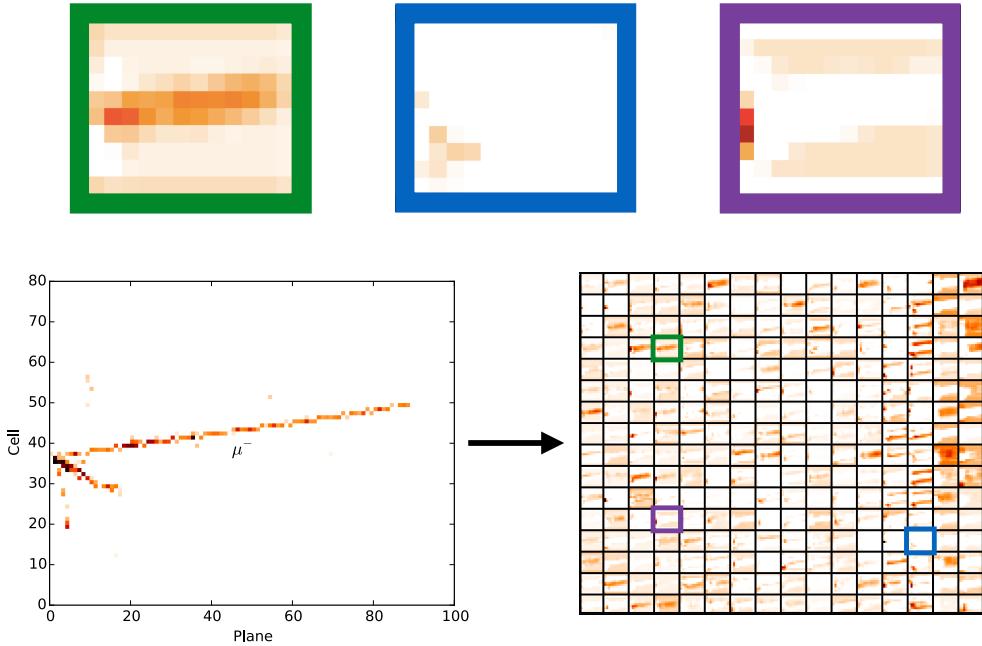


Figure 2.6: Illustration of feature extraction on true ν_μ CC event(left). Event hit map is fed to CVN as a heatmap representing deposited charge, and a series of transformations applied to extract 256 features. A feature can be understood as a variable in the underlying image which is key to classifying the interaction. The features are used to construct the feature map (bottom right). The green, blue and purple features seem to be responses to muon tracks, electromagnetic showers and hadronic showers, respectively. The event shown is a ν_μ CC event, so the response in the blue (electron shower) feature is much weaker than the green. Taken from [62].

certain aspects of the image, e.g. edges or shapes, allowing the network to learn from them. The CVN is trained on simulated beam events and cosmic ray data. The output of the network is a set of scores between zero and one which estimate how likely a given event is to be ν_μ CC, ν_e CC, ν_τ CC, NC, or cosmic activity. The ν_μ CC scores of signal and background events can be seen in Figure 2.7.

2.4.2 Reconstructed Muon Identifier (ReMId)

ReMId is a boosted decision tree (BDT) used to find and score muon like tracks among the different particles in an event. The algorithm uses tracks whose trajectories have been reconstructed using a Kalman-filter algorithm, which iteratively collects hits which only constitute small angular deviations (as would be expected of a muon, which should only be deflected by scattering as opposed to showering), or which are collinear

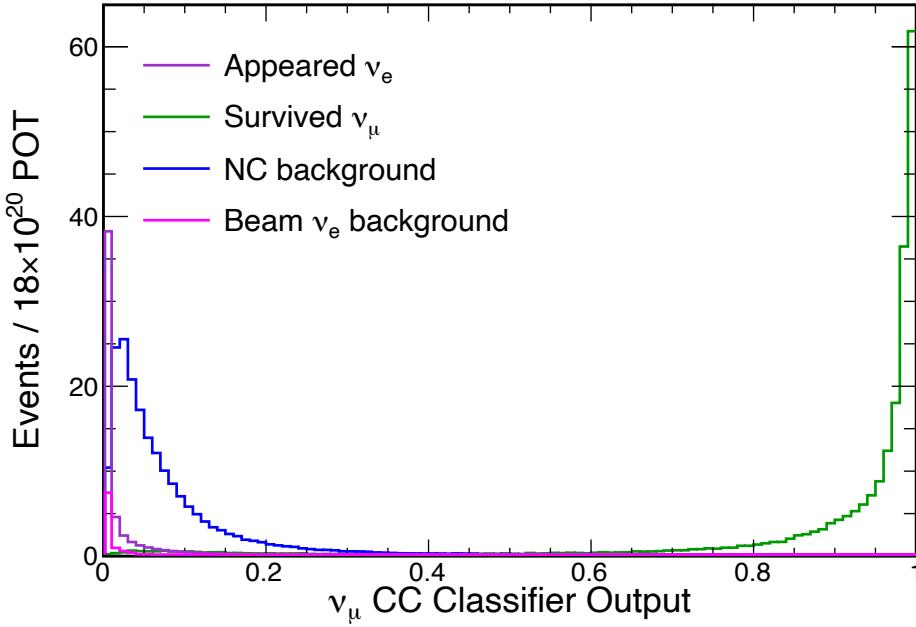


Figure 2.7: ν_μ CC CVN scores for appearance and beam ν_e , survived ν_μ and NC events. The network has excellent discrimination between all event types, with vanishing numbers of non ν_μ CC events scoring higher than 0.5. Taken from [63]

within a few planes. Simulated ν_μ CC and NC events are used for training, and represent signal and background, respectively. Scores are assigned based on

- deposited energy per unit length (dE/dx) likelihood
- scattering likelihood
- track length
- fraction of planes overlapping with hadronic activity.

By calculating the 4D distance between the Kalman tracks in a candidate event and the simulated tracks, the ReMId algorithm assigns a score to each track based on its similarity to a simulated muon track. The highest scoring track is selected as the muon candidate, and its score used for the entire event. An event with no good muon candidates will receive a low score, and can likely be classed as NC background.

2.4.3 Cosmic Rejection BDTs

Dedicated branching decision trees (BDTs) are trained to identify and reject cosmic ray muons which can look similar to ν_μ and ν_e events.

The ν_μ BDT is trained on seven variables; the angle between the lepton and neutrino, the length of the muon track and how vertical said track is, the distance of the track to either end of the detector, distance of the track start/stop to the top or bottom of the detector, the fraction of hits in slice belonging to the track, and the CVN cosmic score. The BDT is trained using real cosmic data (out of beam spill events) as background, and MC simulated ν_μ CC events as signal. All events used in training must have a ReMID score and CVN ν_μ score of at least 0.4. The BDT is therefore trained on events which have been identified as possible muon candidates by previous classifiers, and aims to improve upon them, ideally removing incorrectly labelled cosmic events from the sample. The algorithm outputs a score which quantifies how likely the event is to *not* be a cosmic ray.

Two ν_e BDTs are trained. One for the ‘core’ sample of events, and the other for the ‘peripheral’ sample, where the latter contains events which have failed containment cuts (i.e. do not deposit all of their energy in the active volume of the detector), but have high ν_e CVN scores. The training variables for the core sample are similar to those used for ν_μ . The peripheral BDT is trained on; number of hits, distance from the end of the prong to the top of the detector, minimum distance from the end of the prong to a detector face, normalised X and Y momentum, and vertex position. Similarly to the ν_μ BDT, the networks are trained using MC simulated ν_e CC events as signal, and real cosmic data as background, where only events with reasonably high ν_e CVN scores are included in the training sample.

2.5 Energy Reconstruction

2.5.1 Calibration

As discussed in Section 1.2, neutrinos do not directly ionise the detector medium and cause scintillation. Instead, we infer their energy using the energies of particles they produce, in particular charged leptons (electrons and muons), and hadrons. To ensure accurate energy reconstruction across the volume of the detector and convert the photodetector signals into physically meaningful units, NOvA performs a two part calibration. Both detectors use cosmic-ray muons as standard candles, and beam muons are also used in the ND.

Relative calibration is the first step, and corrects for the effects of at-

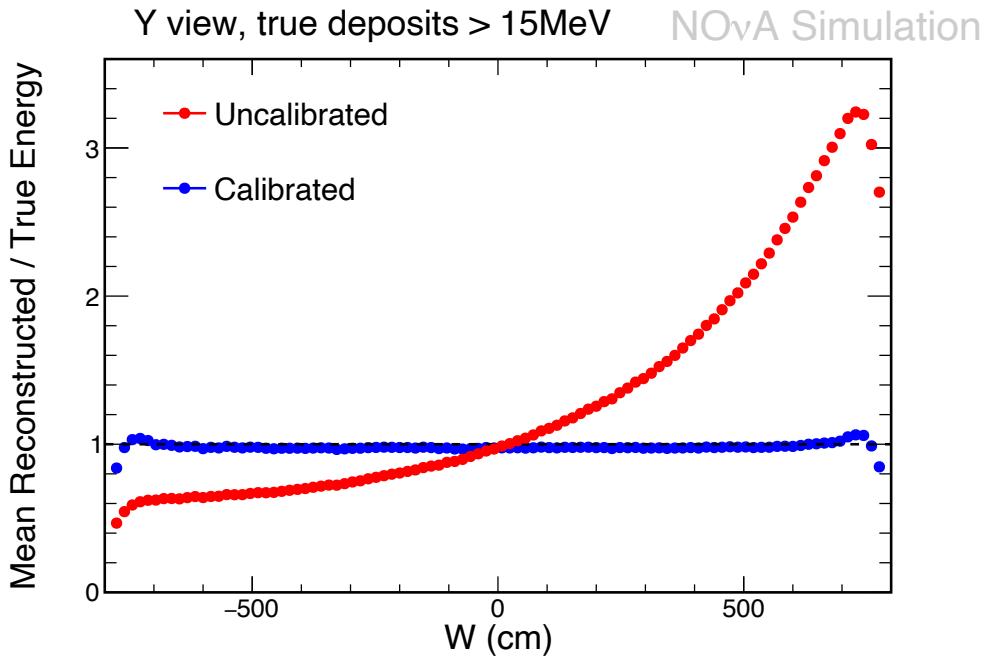


Figure 2.8: Ratio of reconstructed and true energy of a cosmic ray muon track along the length of an FD cell, with and without relative calibration. Taken from [66].

tenuation in wavelength-shifting fibres. Light loses intensity as it travels from its point of production to the readout, meaning without calibration, two identical energy deposits would produce disparate signal strengths depending on where in the cell they originated. (at the end of a detector cell) is attenuated before detection, the same energy deposit will be read out at different values depending on its position in the cell. This is accounted for by using attenuation curves of throughgoing cosmic ray muons to fit the distribution of mean recorded number of photoelectrons per centimetre (PE/cm) as a function of distance from the readout [64]. The fit uses a double exponential distribution, and describes the behaviour of light in both the main body of the cell, and near the non-reflective cell endcaps [65]. An example of the ratio between the true and reconstructed (measured) energy of a simulated cosmic ray muon track, with and without the correction derived from the fit, is shown in Figure 2.8. We see that the relative energy calibration flattens the attenuation curve, meaning the muon energy appears constant across the cell.

The absolute calibration then translates the corrected photoelectron measurement (PE/cm) into GeV/cm. This step uses stopping muons, which deposit all of their energy in the active volume of the detector. The tracks selected for absolute calibration must be away from the cell edges, as the relative calibration in this area is less reliable. The predictions of the Bethe-Bloch equations [67] for the dE/dx (rate of energy loss) a

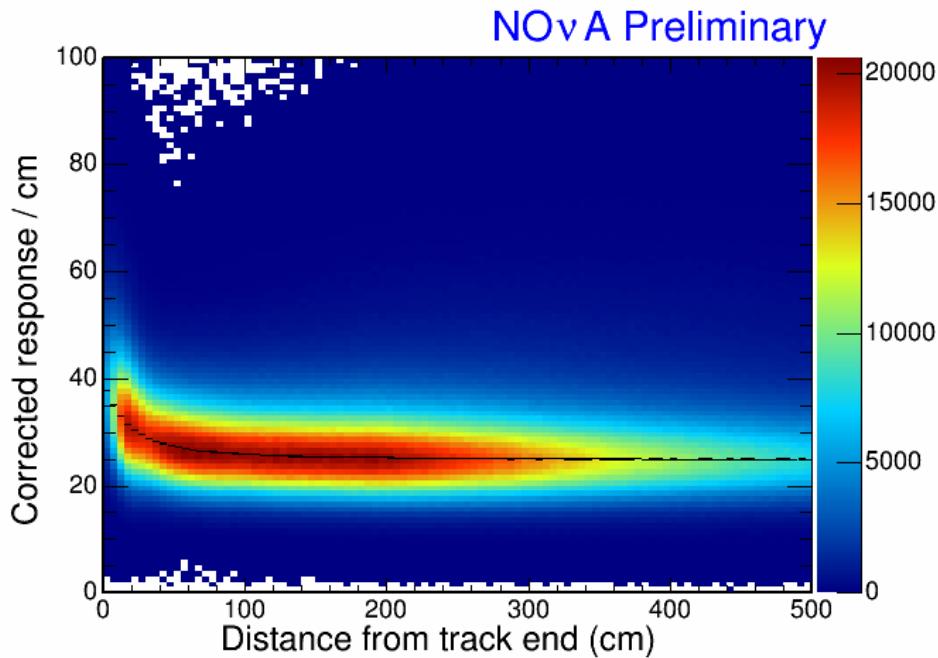


Figure 2.9: Distribution of attenuation corrected PE/cm as a function of distance to the end of the track, for all hits in selected stopping muon tracks. The black line shows the mean value for the corrected response in each distance bin. The region between 100 and 200 cm is most uniform, so these hits are used for absolute calibration. Taken from [68].

particle travelling through a material in units of PE/cm are compared to the corresponding prediction in GeV/cm to determine the conversion factor. Figure 2.9 shows the distribution of attenuation corrected PE/cm for all hits in selected cosmic ray muon tracks.

2.5.2 ν_μ Energy Reconstruction

Charged current muon neutrino interactions are characterised by their long muon track and shower of hadronic activity. The energy of the incoming ν_μ ($\bar{\nu}_\mu$) can therefore be estimated by the sum of muon(anti-muon) energy and hadronic energy.

$$\hat{E}_{\nu_\mu} = \hat{E}_\mu + \hat{E}_{Had} \quad (2.5.1)$$

The muon energy is found using the track length, which has an uncertainty equal to the length of one detector plane. The hadronic energy is estimated by summing the calorimetric energy of all hits not associated with the muon track. Spline (or piecewise) functions then translate the reconstructed muon track length to true muon energy, and the reconstructed visible energy to hadronic energy. These functions are shown in Figure 2.10. The muon spline function lies closer to the diagonal(which

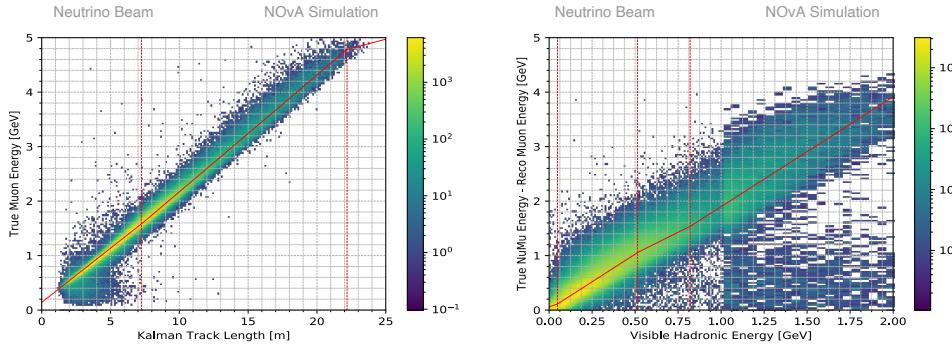


Figure 2.10: Spline functions for converting reconstructed muon track length to true muon energy (left) and reconstructed visible energy to hadronic energy (right) for FHC mode. Taken from [69].

represents a 1:1 conversion) than the hadronic function. This translates to the energy resolution achieved for each component; muon energy resolution is 3%, while hadronic energy resolution is 30%. The resulting total ν_μ energy resolution at the FD ((averaged across the neutrino beam energy distribution)) is 9% in neutrino mode and 8% in antineutrino mode.

2.5.3 ν_e Energy Reconstruction

Electron neutrino reconstruction is more challenging, primarily because electrons shower within the detector, while muons leave clean tracks. Although the ν_e energy is also estimated by the sum of the electron and hadronic energy, the length of the electromagnetic shower produced by the electron is not a reliable energy estimator. The EM shower grows as $\ln E$, and is elongated by stochastic Bremsstrahlung activity towards its end. Calorimetric reconstruction is therefore used for both components, and the total energy calculated using a quadratic fit,

$$\hat{E}_{\nu_e} = \alpha_1 \hat{E}_{EM} + \alpha_2 \hat{E}_{Had} + \alpha_3 \hat{E}_{EM}^2 + \alpha_4 \hat{E}_{Had}^2 \quad (2.5.2)$$

where Prong CVN labels hits as electromagnetic (EM), and the hadronic label is assigned to all non-EM hits. The coefficients α_i are determined by minimising the variance between reconstructed neutrino energy and true neutrino energy in simulated events. This method achieves a ν_e FD energy resolution (averaged across the neutrino beam energy distribution) of $\sim 10\%$ in neutrino mode, and $\sim 9\%$ in antineutrino mode.

2.6 Simulation

To constrain oscillation parameters, NOvA must compare the expected rates of $(\nu_\mu) \nu_e$ (dis)appearance with those observed. Accurate simulation of the beam, neutrino production, neutrino interactions and the interactions of their products within the detector is therefore essential to obtain precise measurements.

The simulation proceeds in several stages. First, a Geant4 [70] based package known as G4NuMI models the production, propagation and decay of hadrons into muons and neutrinos in the NuMI beamline. This includes effects of the geometry and material composition of the target, focusing horns and decay pipe in the NuMI beam. The simulation begins with 120 GeV protons whose profile is Gaussian in the transverse plane, and tracks hadrons until they decay to neutrinos (and, where applicable, continues to track any simultaneously produced muon until its decay). The resulting neutrino flux is then reweighted by the Package to Predict the Flux (PPFX), developed by the MINERvA collaboration [71], which uses external hadroproduction data (e.g. from NA49 [72]) to improve the hadron production prediction and estimate flux uncertainties.

The meson decay information (decay points and momenta) and results from the beam simulation are then passed to a Monte Carlo neutrino event generator. NOvA uses a specially tuned version of GENIE [73] to generate neutrinos and model their interactions with the detector nuclei. The generator outputs the topology of each simulated event, the type of interaction and final state particle kinematics. Cosmic ray interactions, a significant background in the FD, are generated by the Cosmic Ray Shower Library (CRY) generator, and added to the simulation of beam events.

Finally, Geant4 simulates the propagation and energy deposition of all final state particles in the detector. Custom NOvA algorithms then simulate the transport of the scintillation light in the WLS fibres, the collection of these signals by the APDs, and their conversion to electronic signals. This stage accounts for scintillator response, cell reflectivity, fibre absorption and APD noise.

Chapter 3

NOvA 3-Flavour analysis

The results from NOvA’s 2024 analysis were made possible by the work of many collaboration members. In this chapter we give a brief overview of the entire analysis, with more detail on the author’s contributions, followed by a discussion of the results.

A schematic of the NOvA 3-flavour analysis is shown in Figure 3.1. The first step is to classify the event type, which, as discussed in Section 2.4, involves applying the CVN and ReMId algorithms. Event energies are then reconstructed according to their flavour and current, and ν_μ and ν_e CC energy spectra formed. The near detector data (an example of which can be seen in Figure 3.2) is ‘decomposed’ - a procedure which assigns events in the ND dataset to specific neutrino interaction categories present in the MC simulation, for example ν_μ CC or NC. These decomposed samples are then extrapolated using the appropriate oscillation probability, so they are comparable to FD MC simulations. By comparing extrapolated ND data to FD MC simulation, we can account for MC mis-modelling - we assume that any discrepancy arises from inaccuracies in the MC rather than the data. Finally, FD spectra are compared to the corrected simulation, allowing constraints to be placed on oscillation parameters.

3.0.1 ν_μ PID Cut Optimisation

As illustrated in Figure 3.1, the first step in the analysis toolchain is particle identification (PID), where cuts on the CVN, ReMId and Cosmic BDT scores are applied to near and far detector datasets to form samples of ν_e CC, ν_μ CC, NC and cosmic events. These cuts are optimised to maximise the statistical power of NOvA’s measurements. In the most recent analysis, the author was responsible for the ν_μ CC cut optimisation,

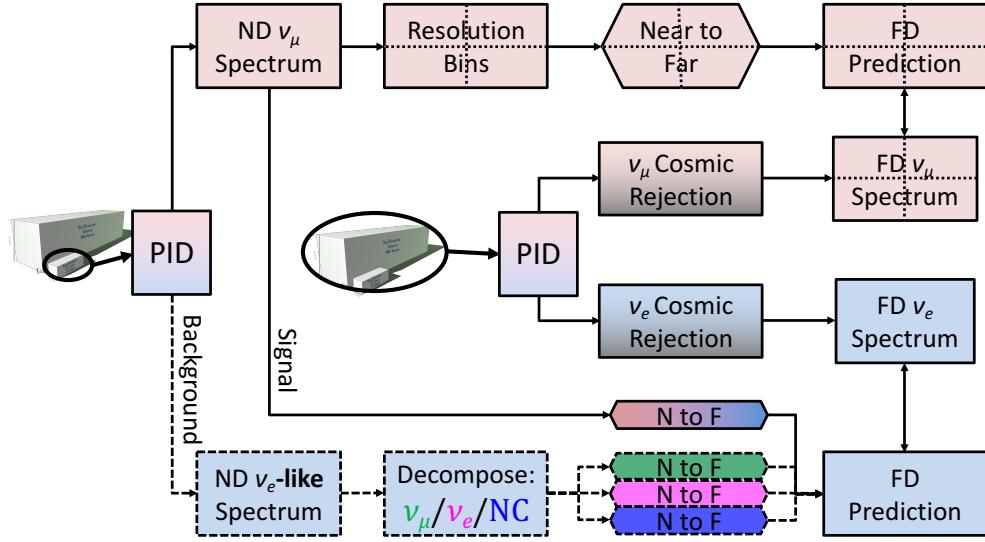


Figure 3.1: Flowchart of the NOvA oscillation analysis framework. The upper panels (red) represent the ν_μ analysis, and the lower (blue) panels ν_e . Taken from [74].

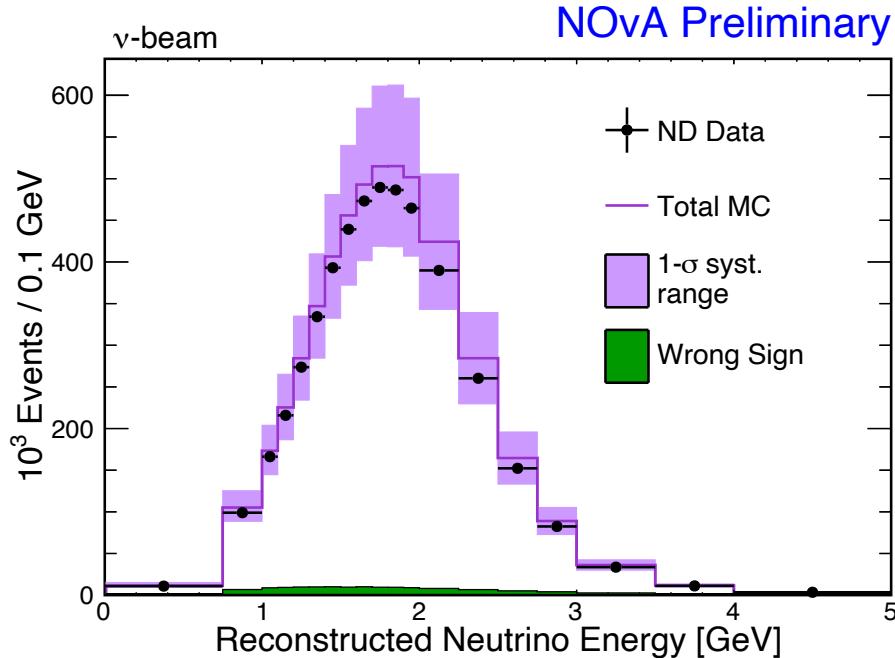


Figure 3.2: Example of a NOvA Near Detector event spectrum with data (black points) and MC prediction (purple line). Expected wrong sign events are shown in green, and a 1σ systematic uncertainty error band in purple. Taken from [75].

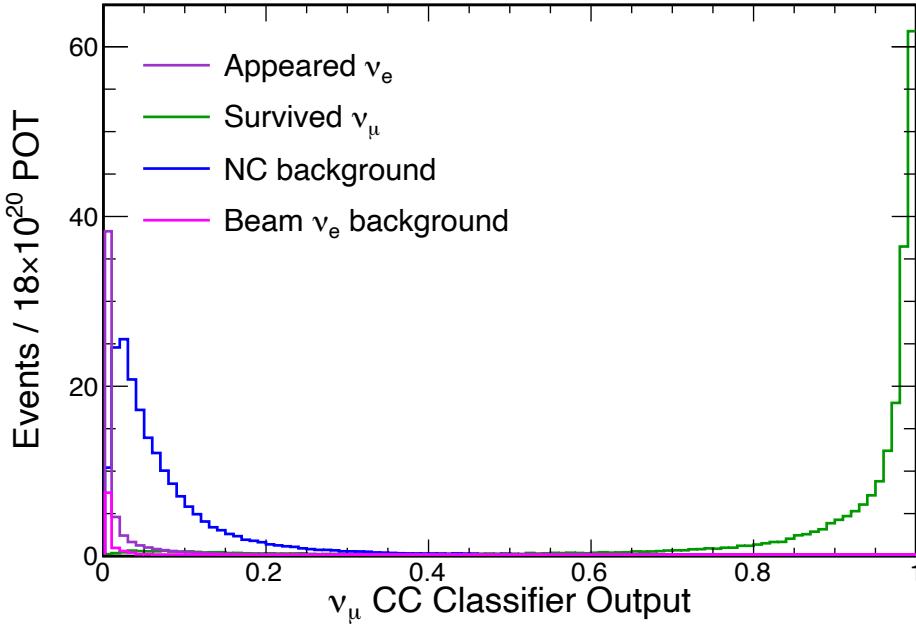


Figure 3.3: ν_μ CC CVN scores for appearance and beam ν_e , survived ν_μ and NC events. At scores of 0.8 and higher the only appreciable contribution to the events comes from survived ν_μ . Taken from [63]

which will be discussed in this section.

To avoid bias, cut optimisation is not performed on real far detector signal. Instead, we use MC signal and background samples, and real out-of-spill data for cosmics. Oscillation weights are applied to MC samples using the previous analysis best fit values for $\sin^2 \theta_{23}$ and Δm_{32}^2 to approximate far detector event spectra.

Cuts on ReMId, CVN and Cosmic rejection scores are optimised simultaneously by scanning over their phase spaces and calculating the number of signal, background and cosmic events at each point. For example, Figure 3.3 shows the ν_μ CC CVN score of survived ν_μ , appeared ν_e , NC background and beam ν_e background events. Most true ν_μ CC signal events lie at values above 0.8, however placing the cut at a higher value could improve signal to background ratio, despite sacrificing some signal events.

A loose preselection, requiring ν_μ CC signal events have $\text{ReMId} > 0.3$, $\text{CVN } \nu_\mu\text{CC} > 0.4$ and $\text{Cosrej} > 0.4$, is applied before scanning over:

- ReMId 0.3 - 0.95
- CVN 0.4 - 0.95
- CosRej 0.4 - 0.65

We consider four figures of merit (FOMs) by which to quantify the

effectiveness of a given cut combination:

- $S/\sqrt{S+B}$ over the entire energy region
- $S/\sqrt{S+B}$ over the energy region of the oscillation maximum, or dip region (between 1 and 2 GeV)
- number of signal events in the dip region (maximum)
- number of cosmic events in the dip region (minimum)

NOvA’s sensitivity to atmospheric parameters depends on the location and depth of the oscillation maximum (also called the dip) in the ν_μ energy spectrum, hence three FOMs focus on this energy range. One FOM calculated over the entire energy region is included to ensure that improvements in the dip region aren’t detrimental to the rest of the spectrum. Previous analyses showed that the signal maximising and cosmic-minimising FOMs were ineffective – the first led to a large background in the dip region while the latter cut too harshly on signal events. Therefore, we focus on $S/\sqrt{S+B}$ metrics and refer to them as the FOM below.

We identify the optimal cut combination as the one yielding the maximum FOM and define any combination within 0.2 of this maximum as ‘good’. This ‘good’ sample is used to assess the stability of the selection. Figure 3.4 shows the FOM distributions of good cut points in the dip and full energy regions in both FHC and RHC beam mode. In FHC mode the number of background events can be increased quite significantly without an outsized impact on the $S/\sqrt{S+B}$. In RHC mode, however, the FOM decreases faster with increasing background because antineutrino events have a lower visible hadronic energy fraction, which lead to clearer anti-muon tracks. Tight CVN and ReMId cuts therefore preserve the majority of the RHC signal, and loosening them only serves to increase background.

Figures 3.5 and 3.6 show projections of the 3D PID phase space in FHC and RHC mode, respectively. These distributions indicate that the FOM is stable over wide regions of phase space. In particular, in both FHC and RHC we see a high concentration of points at low cut values of the Track PID, or ReMId algorithm, but stricter cuts on the CVN score (particularly in RHC mode).

The RHC dip region cuts optimised over all hadronic energy quartiles were chosen for the 2024 analysis. These cuts perform well across the full energy range while maintaining low levels of cosmic background. Tables

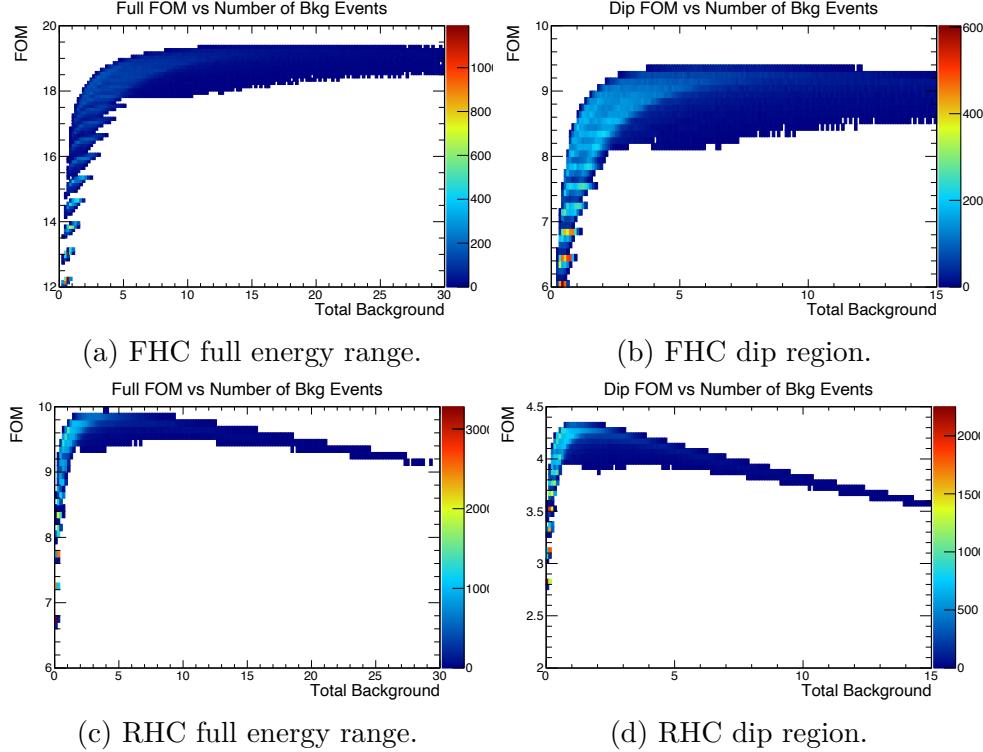


Figure 3.4: 2D histograms showing values of figure of merit against number of background events. Here the FOM is $S/\sqrt{S+B}$ over the full or dip (oscillation maximum) energy region. The upper panels show neutrino mode (FHC) events, and lower RHC.

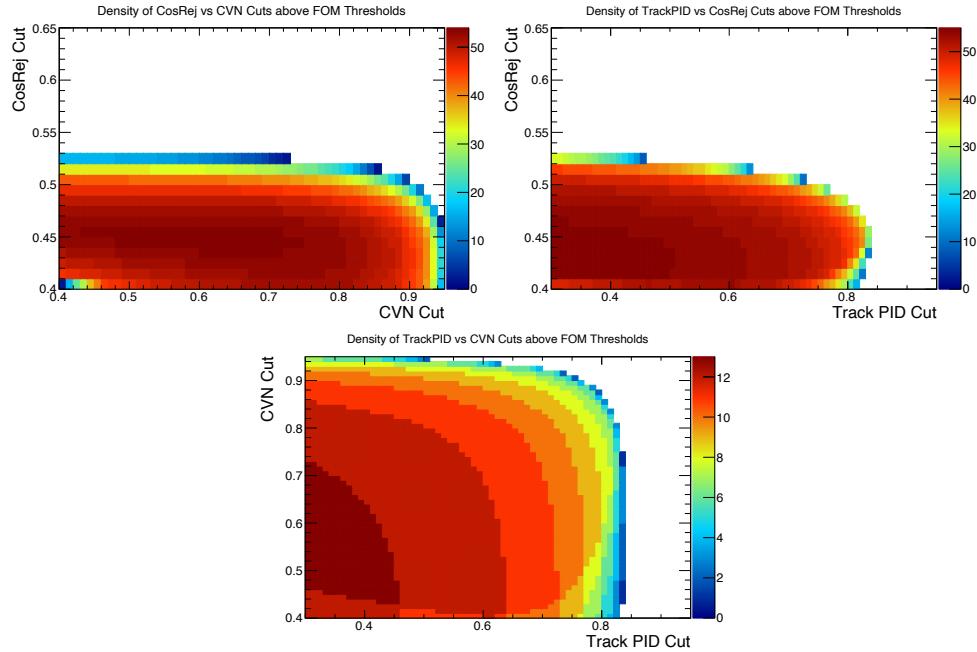


Figure 3.5: FHC cut stability plots. The colour axis shows the number of cut combinations with a 'good' FOM for a given value of the X and Y axes.

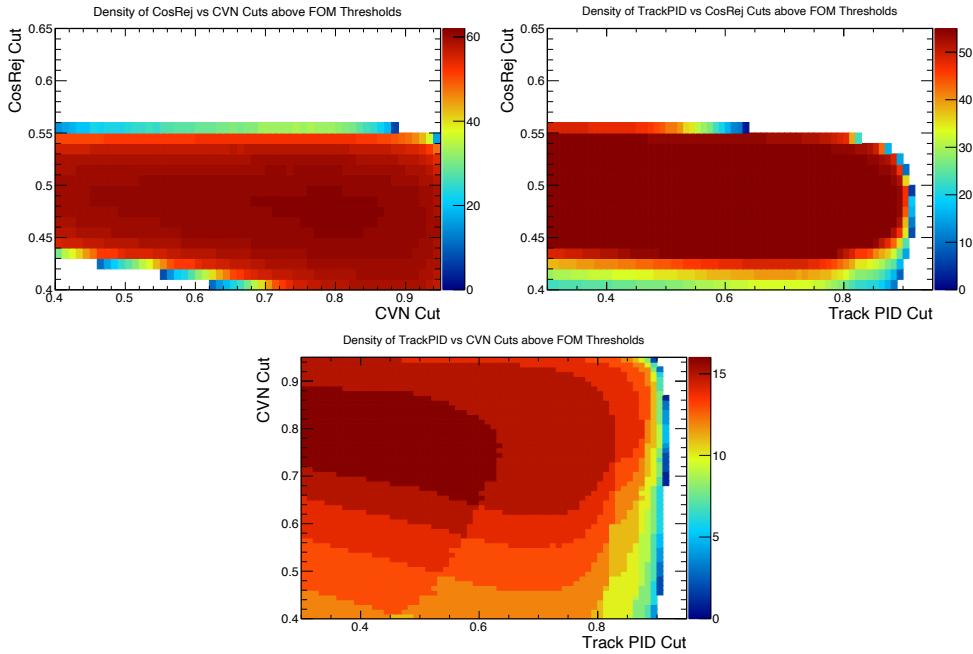


Figure 3.6: RHC cut stability plots. The colour axis shows the number of cut combinations with a 'good' FOM for a given value of the X and Y axes.

3.1 and 3.2 summarise the expected signal and background rates when the FHC dip, RHC dip, and previous analysis cuts are applied. While the FHC dip optimised cuts yield slightly more signal, they admit more cosmic background. Meanwhile, the RHC dip cuts slightly reduce signal, but achieve lower cosmic background.

Figure 3.7 compares the expected signal, background and cosmic event spectra for the 2020 and 2024 selections. While the number of signal events are largely unchanged, the cosmic rates are reduced by the 2024 cuts. The beam background in RHC mode is higher in 2024 than 2020, however its event rate is extremely low for both sets of cuts. A final test to ensure good sensitivity and cut performance in the case of maximal mixing between ν_μ and ν_τ ($\sin^2 \theta_{23} = 0.5$) was performed. The regions of PID space with high concentrations of 'good' FOMs were unchanged, meaning the selected cuts would be appropriate in the case nature has chosen maximal $\mu \leftrightarrow \tau$ mixing.

3.0.2 Analysis Binning

NOvA bins ν_e CC and ν_μ CC samples differently to maximise sensitivity in each analysis. Accurate energy estimation in the ν_μ disappearance channel relies on reconstructing the muon track and hadronic energy. The hadronic component's energy resolution is roughly 10 times worse than

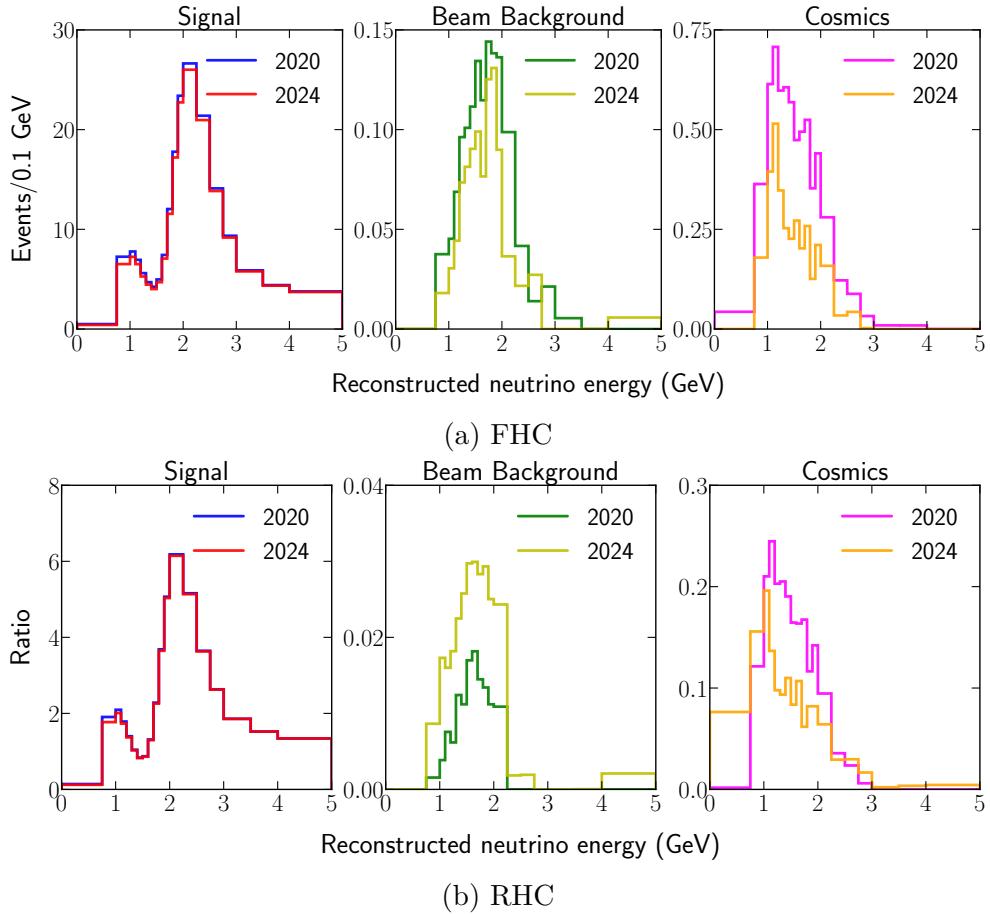


Figure 3.7: Signal (ν_μ CC and $\bar{\nu}_\mu$ CC), Background(NC and beam ν_e) and Cosmic energy spectra for the PID cuts used in 2020, and those chosen for 2024. The considerable scale differences between FHC and RHC mode samples are due to NOvA having received mostly neutrino beam during its running time.

Table 3.1: **FHC**: Cut combinations considered for the 2024 Analysis and the corresponding figures of merit and expected event counts for each case.

Optimisation	Full Energy Range						Dip Energy Range		
	REMid	CVN	CosRej	FOM	N Sig	N Bkg	FOM	N Sig	N Bkg
FHC DIP	0.30	0.71	0.46	19.2	383.8	11.9	9.37	94.1	6.76
RHC DIP	0.30	0.76	0.48	19.1	372.1	7.4	9.31	90.6	4.2
Ana2020	0.30	0.80	0.45	19.23	380	10.6	9.36	93.7	6.45

Table 3.2: **RHC**: Cut combinations considered for the 2024 Analysis and the corresponding figures of merit and expected event counts for each case.

Optimisation	Full Energy Range						Dip Energy Range		
	REMid	CVN	CosRej	FOM	N Sig	N Bkg	FOM	N Sig D	N Bkg
FHC DIP	0.30	0.71	0.46	9.892	101.5	3.806	4.313	20.58	2.19
RHC DIP	0.30	0.76	0.48	9.853	99.27	2.245	4.343	20.08	1.297
Ana2020	0.30	0.80	0.45	9.87	100.7	3.405	4.31	20.49	2.12

the muon’s, (see Section 2.5.2), meaning events with low hadronic activity have superior energy resolution and PID accuracy. Separating events into resolution bins according to their hadronic energy fraction E_{had}/E_ν (in reconstructed energy space) therefore increases the sensitivity of the disappearance analysis. Quartile ranges of fractional hadronic energy in un-oscillated ν_μ CC events define the boundaries used for the resolution bins. As shown in Figure 3.8, quartile 1 (lowest hadronic energy fraction) has the most pronounced dip at the oscillation maximum, while most background events are concentrated in quartile 4.

While energy resolution is paramount in the ν_μ disappearance analysis, the ν_e appearance analysis prioritises discrimination between ν_e CC signal events and background in order to maximise its sensitivity to δ_{CP} and the mass ordering (see Section 2.1). The core ν_e CC sample is therefore separated into two samples based on CVNe scores: a ‘low PID’ sample and a ‘high PID’ sample. The peripheral sample, whose events aren’t guaranteed to be contained within the FD (see Section 2.4) is put into a separate bin with no energy dependence, since their reliable energy reconstruction isn’t guaranteed.

The 2024 analysis includes a new low energy ν_e bin, containing events which were previously excluded due to the high background at low reconstructed neutrino energies. A new BDT has been trained to reject backgrounds in this energy region. Currently, this bin is only included

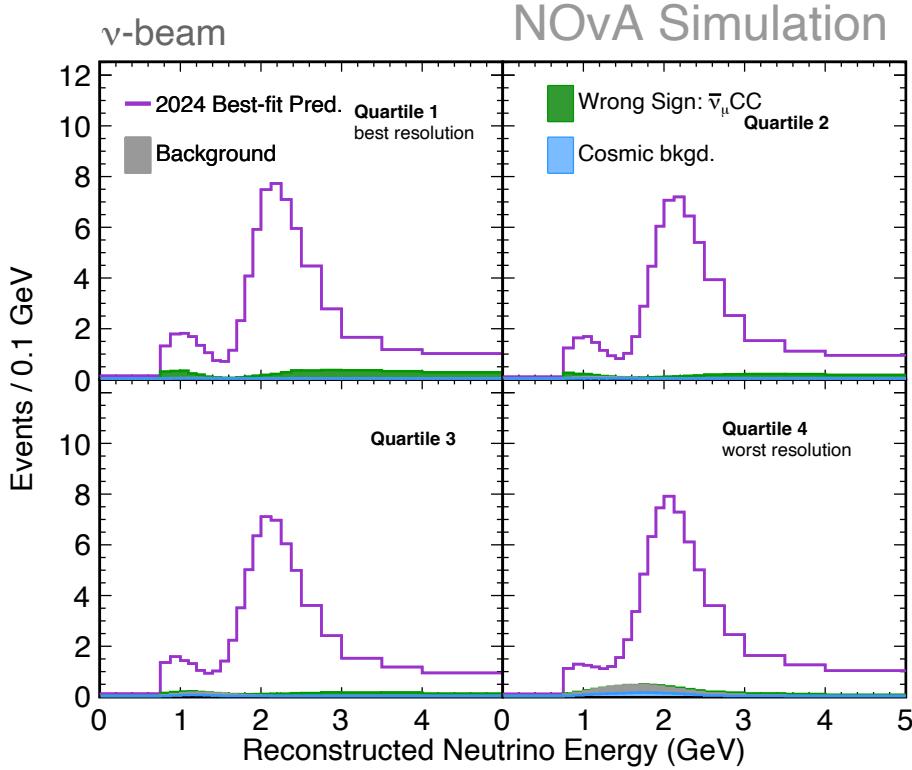


Figure 3.8: Spectra of ν_μ CC, background and cosmic events for each fractional hadronic energy quartile, where Quartile 1 has the best resolution (and smallest background and cosmic rates), and Quartile 4 has the worst resolution, and highest background and cosmic rates. The average energy resolution in the FD for each quartile is 6.2%, 8.2%, 10.3% and 12.4% respectively. Taken from [76].

for FHC data due to limited statistics in RHC mode.

As with the ν_μ CC binning, the PID binning isolates cosmic background in the low PID and peripheral bins, while in the high PID bin, beam neutrino backgrounds dominate, as shown in Figure 3.9.

3.0.3 Decomposition and Extrapolation

The use of near detector data to constrain far detector uncertainties is central to NOvA’s analysis methodology. The ND ν_e sample is used to predict the background to ν_e appearance at the FD, while the ND ν_μ sample is used to predict both the $\nu_\mu \rightarrow \nu_e$ appearance and $\nu_\mu \rightarrow \nu_\mu$ disappearance signal. As shown in Figure 3.1, this is achieved by decomposition (determining the composition of ND samples) followed by extrapolation (projecting ND samples to the FD).

The methods for decomposition and extrapolation differ slightly depending on the sample under consideration. We will outline the processes for each sample below.

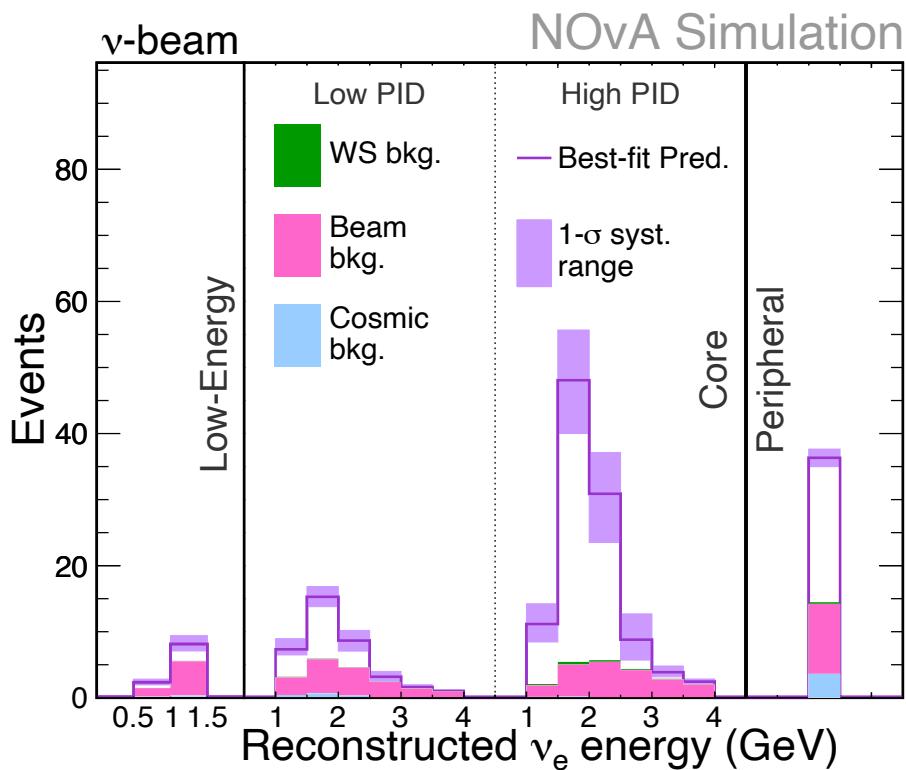


Figure 3.9: Spectra of ν_e CC, background and cosmic events in low energy, low PID, high PID and peripheral bins. Only the number of events in the peripheral bin are used in the analysis - because these events fail containment cuts, accurate energy reconstruction is not guaranteed and could skew results, hence it is not binned in reconstructed neutrino energy, unlike the other bins. Taken from [76].

$\nu_\mu \rightarrow \nu_\mu$ signal

The $\nu_\mu \rightarrow \nu_\mu$ signal at the FD is predicted by the ND ν_μ sample. The NuMI beam is very pure (95% and 97% $(\bar{\nu}_\mu)$, in FHC and RHC mode respectively) at the ND. Therefore, decomposition is simply performed by assigning any difference between data and simulation to the $(\bar{\nu}_\mu)$ component of the sample.

The extrapolation is then performed as follows: first, the background in the ND sample (found using simulation), is subtracted. A smearing matrix (obtained from ND simulation) is applied to the background subtracted sample to unfold it from reconstructed to true energy. Differences in event rates between the near and far detector due to, e.g. detector acceptance or size, are then accounted for by multiplying the true energy ND spectrum by an FD/ND event ratio. The spectrum can now be considered to be an FD sample in true energy. Energy dependent oscillation weights are then applied, followed by another smearing matrix to convert the spectrum back into reconstructed energy space. Backgrounds are then added back in.

$\nu_\mu \rightarrow \nu_\mu$ background

The $\nu_\mu \rightarrow \nu_\mu$ signal has two sources of background, however neither require ND samples to constrain them. The first is background from intrinsic beam electron (anti) neutrinos, which is expected to be small due to high beam purity, so can be taken directly from FD simulation. Cosmic background is estimated using data from the cosmic trigger (which captures events out of time with the beam). The final FD spectra correspond to a specific amount of cosmic livetime, so the cosmic data is scaled to this livetime and the cosmic background added to the total background sample.

$\nu_\mu \rightarrow \nu_e$ signal

The $\nu_\mu \rightarrow \nu_e$ signal sample uses the same extrapolation as the $\nu_\mu \rightarrow \nu_\mu$ signal, however the FD/ND event rate ratio and oscillation weights are altered as required for ν_e events, and $\nu_\mu \rightarrow \nu_e$ oscillations.

$\nu_\mu \rightarrow \nu_e$ background

The ND ν_e sample is used to predict the ν_e appearance background at the far detector. This sample consists of mis-identified NC and ν_μ CC

events, and ν_e CC events from intrinsic beam electron neutrinos. No component is negligible, so simply assigning any discrepancies to one event type is not effective. Instead, the samples are decomposed using different methods depending on the HC mode. FHC decomposition uses additional samples to provide information on event type proportions, then assigns discrepancies accordingly. For example, the ν_μ ND signal sample constrains beam ν_e background by examining the peaks formed by pion and kaon decays in the ν_μ energy spectrum, and using these peaks to estimate the number of electron neutrinos present in the beam from the processes in Equation 2.2.2. Due to lower signal statistics and higher rates of background, proportional decomposition is applied to RHC ND ν_e samples. Here the discrepancies are assigned to all three components, and it is assumed the MC correctly predicts their relative proportions, but not their absolute normalisation in each bin.

Extrapolation of this sample differs because the energy estimator is expected to perform poorly on the ν_e appearance background, in particular on NC events. Converting from reconstructed to true energy space is then more likely to introduce uncertainties. Therefore, the FD/ND event rate ratio is simulated and applied in reconstructed space, and the smearing reco-to-true matrix only used to apply oscillation weights in true space where necessary.

3.0.4 Systematic Uncertainties

60 different systematic uncertainties are considered in the 3-Flavour analysis. Generally, their effects are evaluated by reproducing analysis steps with modified simulated samples. This is achieved by reweighting events, or by simulating new samples with different parameters and performing reconstruction again. The variation in predictions found when considering different values of systematic parameters defines their ± 1 and 2σ values, which are then included in the experiment's fit to data.

Flux

Uncertainties in hadron production and beam transport contribute to the overall neutrino flux uncertainty. Beam transport systematic uncertainties include the horn current, the beam position on the target, the beam spot size and the horn and target positions. Hadron production uncertainties are constrained by external data using the PPFX package [71]. Their effect is assessed by creating multiple universes with varying

proton-target cross sections and theoretical assumptions, to produce a set of alternative flux correction weights. These alternative weights are combined and reduced into a set of uncorrelated weights via Principal Component Analysis (PCA). PCA reduces the dimensionality of highly correlated datasets while preserving most of the underlying information. By computing the covariance matrix of the alternate weights, and calculating its eigenvectors (Principal Components), and eigenvalues (the variance explained by the PC), we can select the largest PCs, which explain much of the dataset’s variance with far fewer parameters. In the case of the hadron production uncertainty the five most dominant PCs, which account for 68% of the total variance, are scaled up to account for the total effect of the systematic and used as uncertainties in the fit.

Cross Section

The GENIE event generator (see Section 2.6) includes many tunable parameters (knobs) that, for example, control neutrino-nucleus interaction models. Instead of regenerating events for each parameter variation, GENIE provides reweight factors which modify existing samples to reflect different knob values. These reweights are used to evaluate the effect each parameter has on the predicted sensitivities. The 2024 analysis includes 78 knobs. Those with small effects on the final sensitivities are condensed to 8 systematic uncertainties using PCA, while the remaining thirty large knobs are each considered a separate systematic uncertainty.

Detector

Uncertainties in detector response and calibration are characterised by regenerating ND and FD samples with the relevant systematics altered.

The dominant detector uncertainty is the overall energy calibration scale ($\pm 5\%$), and is motivated by a data-MC discrepancy in the responses of protons and muons in the ND. This includes an absolute and relative scale uncertainty: in the former we assume the ND and FD have identical calibration, and the latter we assume their calibration is anticorrelated. Spatial variations in detector response necessitates a calibration shape uncertainty, which is found by assuming a linear shape, and characterising the differences in the middle and edge responses using fits to data/simulation ratios.

The 2024 analysis used an improved light model, which describes the amount of light produced in the detector and collected by the fibre. Light

within the detector comes from both scintillation and Cherenkov photons produced as a charged particle passes through its volume. Uncertainties on both light sources are found by varying the relevant parameters when tuning simulated detector light levels to data. NOvA detector response has reduced over its run time due to scintillator and fibre ageing, among other effects. This drift in response is described by a linear decrease in light level as a function of time.

As discussed in Section 2.5.2, a muon’s energy reconstruction relies on its track length. The uncertainty on track length for a muon of known energy is accounted for by a muon energy scale systematic, which includes both correlated and anticorrelated near and far detector uncertainties.

Finally, a new detector response systematic, ‘Geant4 Reweight’ was included in the 2024 analysis. Geant4Reweight is a framework which allows events to be reweighted after production, and is used to account for uncertainties in inelastic hadron scattering in both neutrino nucleus, and resulting secondary interactions.

Other

A new neutron interaction model, known as MENATE [77], replaced the built in Geant4 model used in previous analyses. An additional systematic was thus included to parameterise the differences in sensitivity when each model is used.

The statistical uncertainty on the number of cosmics expected to be seen at the FD (which is used to make predictions) is covered by the cosmic scale systematic, which decreased in 2024 due to the increased dataset size.

Other uncertainties include the total POT, number of events caused by interactions in the rock surrounding the FD, and the matter density of the earth are also included, and are described in further detail in [78].

3.0.5 Oscillation parameter extraction

Oscillation parameters are found by comparing the measured $\nu_\mu^{(-)}$ and $\nu_e^{(-)}$ datasets to extrapolated FD predictions as a function of oscillation parameters and systematic nuisance parameters [78]. We fit the oscillation parameters $\vec{\theta}$ to the observed data by minimising a binned Poisson log-likelihood between the prediction $E(\vec{\theta})$ and the observed data O , in-

cluding Guassian penalty terms for systematic degrees of freedom $\vec{\delta}$:

$$\chi^2 \left(\text{data}, \vec{\theta}, \vec{\delta} \right) = 2 \sum_{i=1}^N \left[E_i \left(\vec{\theta}, \vec{\delta} \right) - O_i + O_i \ln \frac{O_i}{E_i \left(\vec{\theta}, \vec{\delta} \right)} \right] + \sum_{j=1}^M \frac{\delta_j^2}{\sigma_j^2} \quad (3.0.1)$$

Here N is the analysis bins (Section 3.0.2) and M the systematic parameters (Section 3.0.4). $\hat{\theta}$ is the set of oscillation parameters which minimise χ^2 .

In the fit we vary $\sin^2 \theta_{23}$, $\sin^2 2\theta_{13}$, Δm_{32}^2 and δ_{CP} . Solar parameters are fixed at values $\Delta m_{21}^2 = 7.53 \times 10^{-5} \text{eV}^2$ and $\sin^2 \theta_{12} = 0.0307$.

NOvA's 2024 analysis used two different statistical approaches to measure oscillation parameters. Traditionally, NOvA's results were extracted using Frequentist statistics, however a Bayesian analysis was first performed in 2022 (on 2020 data), and was again performed in 2024. Here we will predominantly discuss the Bayesian approach, as per the focus of the author's work.

In brief, Frequentist and Bayesian statistics differ in their interpretation of probability. While a Bayesian approach sees probability as a statement about how likely a particular outcome is, a Frequentist one sees it as a statement of long range likelihood. Consider an experiment with two possible outcomes: a Frequentist would consider the probability of outcome A being 90% to mean that were the experiment performed 100 times over, 90 of the results would be outcome A, while a Bayesian would consider the same probability to mean there is a 90% chance that on performing the experiment, the outcome will be A.

Methods of reporting results also vary between the two approaches. A standard metric used in Frequentist statistics is 'confidence intervals' (or regions in the case of higher order phase spaces) - regions which, if a measurement were made many times with independently collected data, would include the true values with at least a minimum given probability (this is known as coverage). In Bayesian statistics *credible* intervals (or regions) are reported. Perhaps more intuitively, credible intervals represent regions that should contain the true value with a minimum given probability. The Frequentist analysis uses the Feldman-Cousins technique to create confidence intervals, which is explained in [79].

3.0.6 Bayesian Analysis

Bayesian inference is built on Bayes' Theorem [80, 81], which states, for measured values \vec{x} and parameters $\vec{\theta}$,

$$\mathcal{P}(\vec{\theta}|\vec{x}) = \frac{\mathcal{P}(\vec{x}|\vec{\theta}) \mathcal{P}(\vec{\theta})}{\mathcal{P}(\vec{x})} \quad (3.0.2)$$

where $\mathcal{P}(\vec{\theta}|\vec{x})$ is the posterior probability - the probability of the parameters $\vec{\theta}$ being correct given data \vec{x} . $\mathcal{P}(\vec{x}|\vec{\theta})$ is the likelihood, (which for our purposes is the binned Poisson likelihood used to get (3.0.1)) and $\mathcal{P}(\vec{\theta})$, is the prior [82], representing pre-existing knowledge of $\vec{\theta}$. Should there be no previous information about the parameter in question one can choose an 'uninformative' prior which is uniform in said parameter. The combination of the likelihood and the prior is commonly referred to as the model in Bayesian inference. $\mathcal{P}(\vec{x})$ is the probability for observing the data. Given that our goal here is to estimate parameters using existing data, this term is a normalization constant, and can be ignored.

An undoubted advantage of Bayesian inference is the ease with which it incorporates systematic uncertainties. Each systematic uncertainty adds a degree of freedom to the fit which is uninteresting for our analysis - in the frequentist approach this is remedied by profiling; scanning over all possible values of the systematics and finding the combination which maximises the likelihoods of the parameters of interest. In Bayesian inference however, we can account for these parameters and absorb the extra degrees of freedom by marginalizing - integrating the likelihood across them, i.e.

$$\mathcal{P}(\vec{x}|\vec{\theta}) = \int \mathcal{P}(\vec{x}|\vec{\theta}, \vec{\delta}) d\vec{\delta} \quad (3.0.3)$$

where as before, $\vec{\delta}$ are the nuisance parameters.

When comparing hypotheses in Bayesian statistics it is necessary to consider that the probability of a hypothesis (H_0) given data x , $\mathcal{P}(H_0|x)$ is dependent on the prior used in (3.0.2). As a result, for a binary hypothesis test (e.g. whether neutrino mass ordering is normal or inverted) it is often useful to find the Bayes Factor, where, considering two models, H_0 and H_1 with associated priors $\mathcal{P}(H_0)$ and $\mathcal{P}(H_1)$

$$BF = \frac{\mathcal{P}(x|H_1)}{\mathcal{P}(x|H_0)} = \frac{\mathcal{P}(H_1|x)}{\mathcal{P}(H_1)} / \frac{\mathcal{P}(H_0|x)}{\mathcal{P}(H_0)} = \frac{\mathcal{P}(H_1|x) \mathcal{P}(H_0)}{\mathcal{P}(H_0|x) \mathcal{P}(H_1)} \quad (3.0.4)$$

If the priors are identical, this is equal to the ratio of the hypotheses' posterior probabilities, and a large Bayes factor indicates H_1 is preferred to H_0 . For different priors, a large Bayes Factor implies H_1 is preferred to a greater degree than the prior information on H_1 and H_0 would suggest.

Markov Chain Monte Carlo

While Bayes's theorem can be solved analytically to find the posterior probability in problems with few parameters, the dimensionality of the posterior space grows with the number of parameters: it becomes necessary to perform multidimensional integrals to marginalize away nuisance parameters (recall NOvA has ~ 60 systematic uncertainties which must be treated this way). We therefore use Markov Chain Monte Carlo (MCMC) to calculate the posterior probability. MCMC is highly computationally efficient because instead of sampling points equally over the posterior space, it samples according to the shape of the posterior probability. Therefore, areas of high probability are heavily sampled, and computational power isn't wasted on areas of low probability which do not contain the maximum. This is achieved by using the probability at a given point to inform the next sampled point, each sample point is recorded, creating a so-called chain.

NOvA has implemented two different techniques for MCMC sampling, known as ARIA and Stan.

ARIA

ARIA uses the Metropolis-Rosenbluth-Rosenbluth-Teller-Teller (MRRTT) approach, and is named for Arianna Rosenbluth, who first implemented the method computationally. MRRTT uses a proposal-acceptance technique, such that a new point x' is proposed from a symmetric distribution (in NOvA's case a multivariate Gaussian with dimensions matching the number of parameters). The proposed point and previous point are then used to find the acceptance criterion $A = \min \left[\frac{P(x')}{P(x)} \right]$. The proposed point is accepted if a value randomly sampled from a uniform distribution between 0 and 1, u , lies below the acceptance criterion A . If the point is rejected, the previous point is reused. The structure of A means that proposed points which lie in high posterior probability regions are more likely to be accepted. The MRRTT algorithm leads to high levels of autocorrelation (the degree of dependency a given point has on its predecessor) between steps in the chain. The chain must therefore be

‘thinned out’ (steps removed at fixed intervals, effectively increasing the space between steps), before use. It can also be ineffective for posteriors with many dimensions and tight constraints, leading to computational inefficiency.

Stan

Hamiltonian Monte Carlo (HMC) [83], avoids some pitfalls associated with ARIA, and is implemented in NOvA via an external package, Stan [84]. HMC treats the *negative* of the posterior space as a topological object across which a particle is moving. As such, the ‘gravitational’ gradient is always in the direction of a high probability region. By numerically solving Hamilton’s equations for a simulated particle with some initial momentum and whose Hamiltonian is a function of the posterior probability distribution, the algorithm finds the endpoint of the particle’s trajectory (which should move towards regions of high probability), and uses its coordinates to choose the next sample. This proposed point is accepted or rejected as in ARIA. HMC yields lower autocorrelation in the MC chain than MRRT, and, because it makes use of the posterior distribution when proposing new points, requires fewer sample attempts to explore the entire posterior space.

Mass Ordering

Unlike the mixing angles and (within 0 and 2π) δ_{CP} , the mass ordering is a discrete parameter. This presents an issue for MCMC, as it will not naturally ‘jump’ from a positive Δm_{32}^2 to a negative one, meaning only the ordering the chain begins in will be explored. This is remedied in ARIA and Stan by adding a step in the proposal acceptance procedure which gives an acceptance score for the proposed point also moving switching mass ordering. This switch is accepted/rejected in the same way as previously described.

Reactor Constraints

Long-baseline experiments, such as NOvA, have limited sensitivity to $\sin^2 \theta_{13}$ because, as seen in Equation 2.1.1, it only appears in combination with $\sin^2 \theta_{23}$. Applying external constraints from reactor experiment Daya Bay [85] can increase NOvA’s sensitivity to other parameters which it is better placed to measure. Three different treatments of $\sin^2 2\theta_{13}$ are used. In the first, $\sin^2 2\theta_{13}$ is allowed to vary freely. In

the second, the so-called 1D constraint from Daya Bay is used, wherein $\sin^2 2\theta_{13} = 0.0851 \pm 0.0024$. The third utilises ‘2D’ constraints - Daya Bay also reports separate constraints for $\sin^2 2\theta_{13}$ vs Δm_{32}^2 in the normal and inverted mass ordering. These separate constraints allow NOvA to increase its sensitivity to the mass ordering via a technique proposed by Nunokawa et. al in [86], which showed that reactor and long-baseline experiments would measure different values for Δm_{32}^2 when assuming the incorrect mass ordering, but would agree when assuming the correct mass ordering. This effect can be seen in Figure 3.14.

In the Bayesian analysis these constraints are implemented after sampling, which is performed with an uninformative (i.e. uniform) prior on $\sin^2 \theta_{13}$. The resulting posterior distributions are then simply multiplied by the likelihood of the external constraint to find the constrained distributions. Performing the sampling with a prior flat in $\sin^2 \theta_{13}$ also provides a NOvA only measurement of this parameter, which had previously been excluded due to the computational power required to compute this with the Feldman-Cousins method.

3.0.7 2024 Bayesian Analysis Results

NOvA’s 2024 analysis observed 384 ν_μ signal events, of which 11.3 are expected to be background, and 106 $\bar{\nu}_\mu$ events, of which 1.7 are expected to be background. If neutrinos didn’t oscillate, one would expect to see ~ 2100 ν_μ at the FD. In the appearance channel, 169 ν_e events and 32 $\bar{\nu}_e$ signal events were observed, with expected background counts of 54.9 and 12.2, respectively. The data and posterior ranges are shown in Figure 3.10. This dataset was fitted according to the method described in Section 3.0.5. After describing the goodness of fit metric used to check the robustness of our model, we will discuss the preferred regions NOvA has found with this data analysis.

The goodness of fit was evaluated using posterior predictive p-values (PPP), the Bayesian analogue of Frequentist p-values. To calculate the PPP we draw S samples of $\vec{\theta}$ from the posterior probability distribution and use them to create predictions E for the observed data O . We then compute the χ^2 between the observed and predicted values, χ^2_{data} . Another set of predictions is then formed by applying Poisson fluctuations to the predictions E , and calculate the χ^2 between the fluctuated and original predictions, $\chi^2_{\text{pseudodata}}$. χ^2_{data} tests how well the model encapsulates oscillation and systematic parameter variations, while $\chi^2_{\text{pseudodata}}$ treats only statistical uncertainties. By examining the distribution of the

$\chi^2_{\text{data}}, \chi^2_{\text{pseudodata}}$ pairs, we can infer how well our model describes both kinds of uncertainties and oscillation parameters. The PPP is defined as the fraction of $\chi^2_{\text{pseudodata}} > \chi^2_{\text{data}}$. A PPP of 0.5 indicates that the model describes the data perfectly with the exception of statistical variations.

The PPP for the 2024 analysis is 0.48, indicating the data is well described by our model. This is clear in the total posterior predictive distribution in Figure 3.11, which is symmetric around $\chi^2 = 1$ per bin in both axes, implying that the model is capable of describing both statistical fluctuations and systematic parameter variations.

The ν_μ and $\bar{\nu}_\mu$ samples individually have good p -values, and their distributions closely match the combined results. The effect of statistical fluctuations is more visible in the smaller statistics samples. The ν_e , $\bar{\nu}_e$ and low energy ν_e samples all have a wider spread of χ^2 values than their ν_μ equivalents, particularly in $\chi^2_{\text{pseudodata}}$, due to their larger statistical uncertainties. The ν_e sample has a good p -value of 0.58, however, χ^2 between each of the pseudodata spectra and predicted spectra is typically larger than the equivalent between the predictions and the data. The same can be seen, but to a greater extent, in the ν_e low energy sample. This is a result of very low statistics in this sample, and, as can be seen in the top left panel of Figure 3.10, good description of the data points by the prediction in both the nominal and low energy ν_e samples. The opposite effect is seen in the $\bar{\nu}_e$ contour, where the systematic and oscillation parameter fluctuations give rise to larger χ^2 's when compared to data than the statistical fluctuations when compared to predictions. Once again this is evident in Figure 3.10, where we can see, in the top right panel, three datapoints which lie outside the posterior ranges.

Parameter	HPD		
	Both MO	NO	IO
$\Delta m_{32}^2 (\times 10^{-3} \text{eV}^2)$	-	2.431	-2.476
$\sin^2 \theta_{23}$	0.547	0.547	0.474
δ_{CP}	0.89π	0.89π	1.57π

Table 3.3: Highest Posterior Density (HPD) points for oscillation parameters in both, Normal and Inverted Mass Orderings (Both MO, NO and IO, respectively). Results from [87].

The highest posterior density points for oscillation parameters with the 1D reactor constraint applied are shown in Table 3.3. Here, the HPD point for $\sin^2 \theta_{23}$ lies in the upper octant regardless of the mass ordering. Figure 3.12 shows the 90% credible intervals in $\Delta m_{32}^2 - \sin^2 \theta_{23}$ space in both mass ordering cases for both NOvA and other experiments.

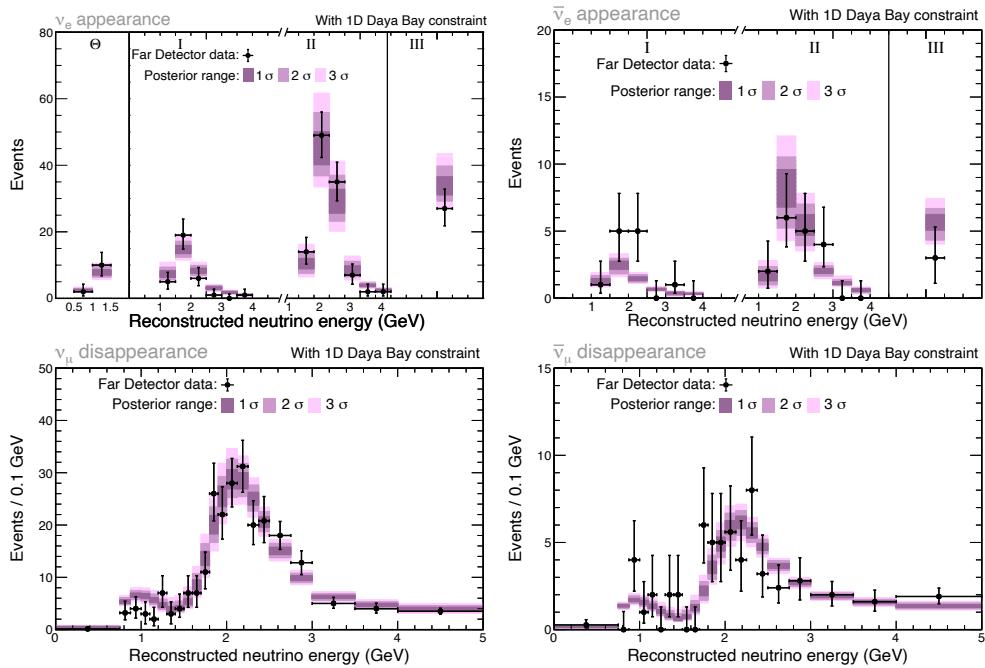


Figure 3.10: Reconstructed neutrino energy distribution of 2024 data events (black crosses), for ν_e and $\bar{\nu}_e$ appearance (top), and ν_μ and $\bar{\nu}_\mu$ disappearance (bottom), without Daya Bay constraints applied. The ν_e samples are divided into bins of low (I) and high (II) PID scores, plus the peripheral sample (III). In the ν_e sample only, Θ indicates the low energy sample which was added for 2024 analysis. The coloured bands correspond to the range of 1σ (darkest), 2σ and 3σ (lightest) confidence intervals from the posterior distribution resulting from our fit to data (including oscillation and systematic parameters, but not statistical error.)

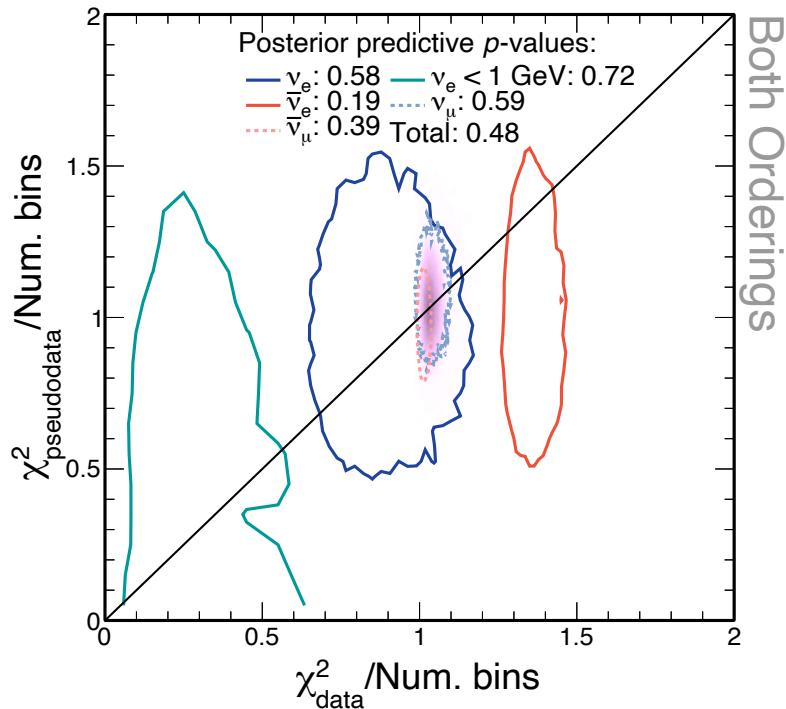


Figure 3.11: Posterior predictive p -values from real data MCMC samples with the $1D$ reactor constraint applied to θ_{13} . The purple distribution shows the scatterplot of the χ^2 computed between the observed data and predicted datasets pulled from the posterior probability distribution against a similar χ^2 between these predicted datasets and corresponding pseudodata sets, both divided by the number of degrees of freedom in the fit. The dashed and solid contours show the 1σ intervals from the same posterior-predictive distributions calculated only for ν_e (dark blue, solid), $\bar{\nu}_e$ (red, solid), ν_μ (light blue, dashed) and $\bar{\nu}_\mu$ (light red, dashed) .

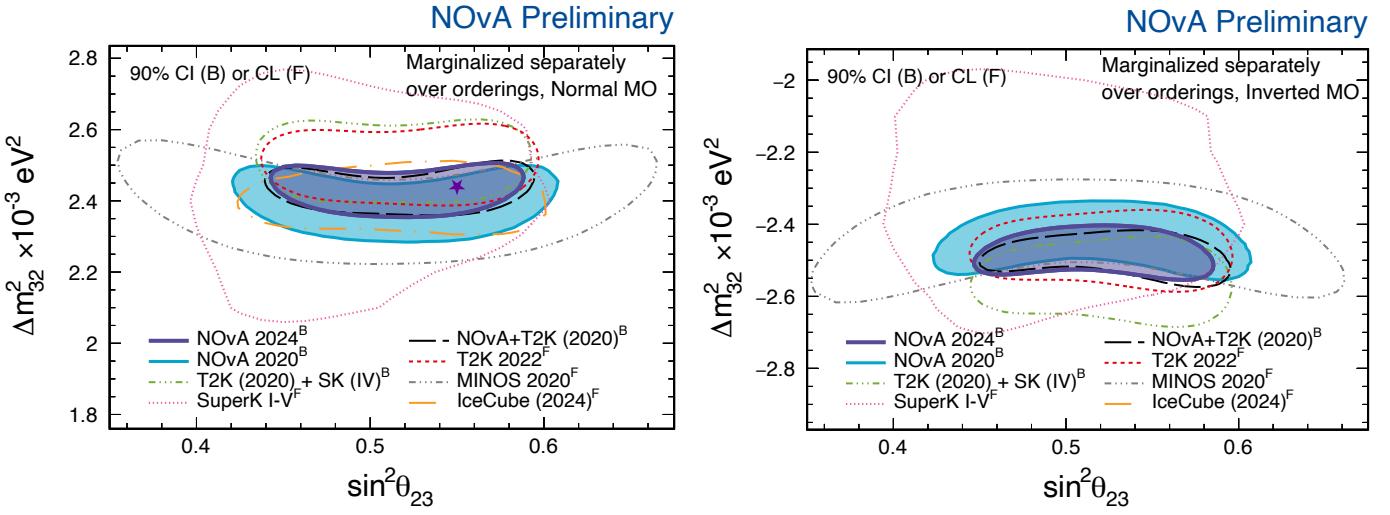


Figure 3.12: 90% Credible Intervals for NOvA and other experiments in the Δm_{32}^2 - $\sin^2 \theta_{23}$ space when considering only Normal (left) or Inverted (right) ordering points. To find the best fit point both orderings are considered, this point is marked with a star, and lies in normal ordering, upper θ_{23} octant. From [88].

NOvA's favoured regions are consistent with those of other experiments, and show improved sensitivity to both parameters in question relative to its previous analysis (shown in pale blue) due to the increase in statistics. NOvA's results are consistent with maximal mixing ($\sin^2 \theta_{23} = 0.5$) at 1σ .

	No constraint		1D constraint		2D constraint	
	BF	Probability	BF	Probability	BF	Probability
NO preference	2.2	69%	3.2	76%	6.8	87%

Table 3.4: Bayes factor for normal ordering against inverse ordering hypothesis with no, 1D and 2D reactor constraints applied.

From Figure 3.12 we can also observe that the HPD point lies in the normal ordering region. Indeed, NOvA prefers normal ordering (albeit weakly, with a Bayes factor of 2.2) without reactor constraints applied. This preference strengthens with the application of 1D and 2D reactor constraints (87% NO preference), as predicted in [86]. The Bayes factors and associated probabilities for mass orderings can be seen in Table 3.4, and the 1D posterior distribution for Δm_{32}^2 with and without reactor constraints applied can be seen in Figure 3.14.

As discussed in Section 2.1, NOvA gleans sensitivity to δ_{CP} by ν_e and $\bar{\nu}_e$ comparing appearance rates, which differ if there is leptonic CP violation. Figure 3.15 shows the observed ν_e and $\bar{\nu}_e$ event counts alongside the values one would expect for various δ_{CP} values and in each mass ordering. Here degeneracy between CP violation and mass ordering on appearance

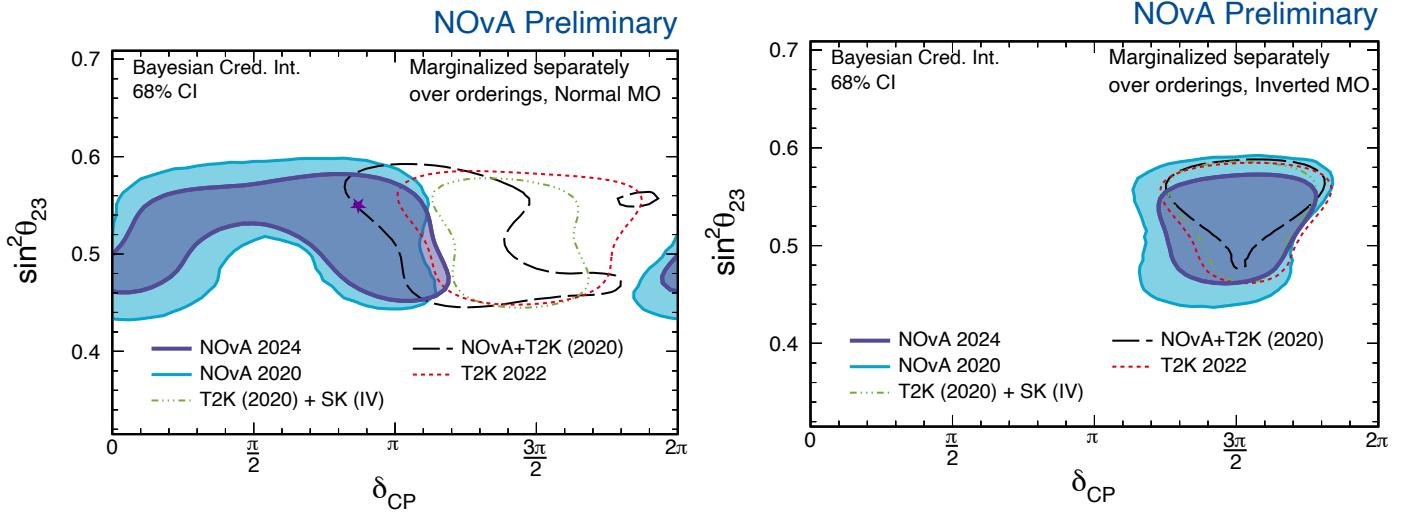


Figure 3.13: 68% Confidence Intervals for NOvA and other experiments in the $\sin^2 \theta_{23}$ - δ_{CP} space when considering only Normal (left) or Inverted (right) ordering points. The best fit point lies in NO, upper octant, at $\delta_{CP} = 0.93\pi$. From [88].

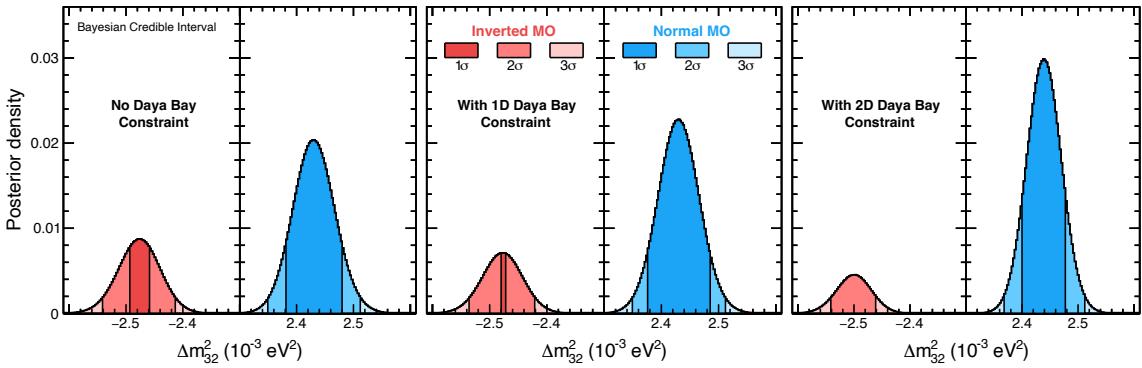


Figure 3.14: Marginalised posterior probability densities for Δm_{32}^2 from a fit to data with different Daya Bay constraints applied- left without constraint, the middle with θ_{13} constraint, and the rightmost with 2D $\Delta m_{32}^2 - \theta_{13}$ constraint. NOvA measurements are complementary to Daya Bay, with the normal ordering preference increasing as constraints are applied. Figure from [87].

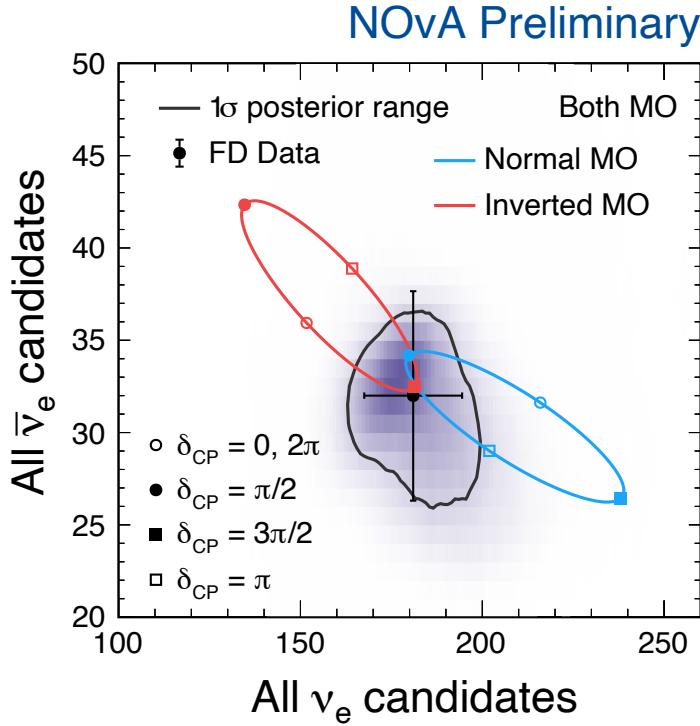


Figure 3.15: Bievent plot showing ν_e against $\bar{\nu}_e$ event counts (black point) compared to the expected value for different values of δ_{CP} in normal (blue) and inverted (red) mass ordering. From [88].

rates is apparent - NOvA's data point lies in a region compatible with maximally violating δ_{CP} and inverted ordering, and less violating δ_{CP} and normal ordering. As seen in Figure 3.13, the best fit point lies near CP conserving values in normal ordering, however the preference is very small. Interestingly, while NOvA's preferred regions are disparate with the equivalents of other experiments in the normal ordering, they agree well in the inverted ordering.

Figure 3.16, shows the posterior densities for the Jarlskog invariant for normal ordering. When applying a prior to $\sin \delta_{CP}$, the posterior density distribution is peaked near $J = 0$, tending slightly towards CP conserving values, in normal ordering, but has a strong preference towards CP violation in inverted ordering. This is confirmed by the Bayes Factors for CP violation over CP conservation extracted from these spectra: 1.1 (51% preference for CP violation) in NO, and 4.3 (81% preference) in IO [87].

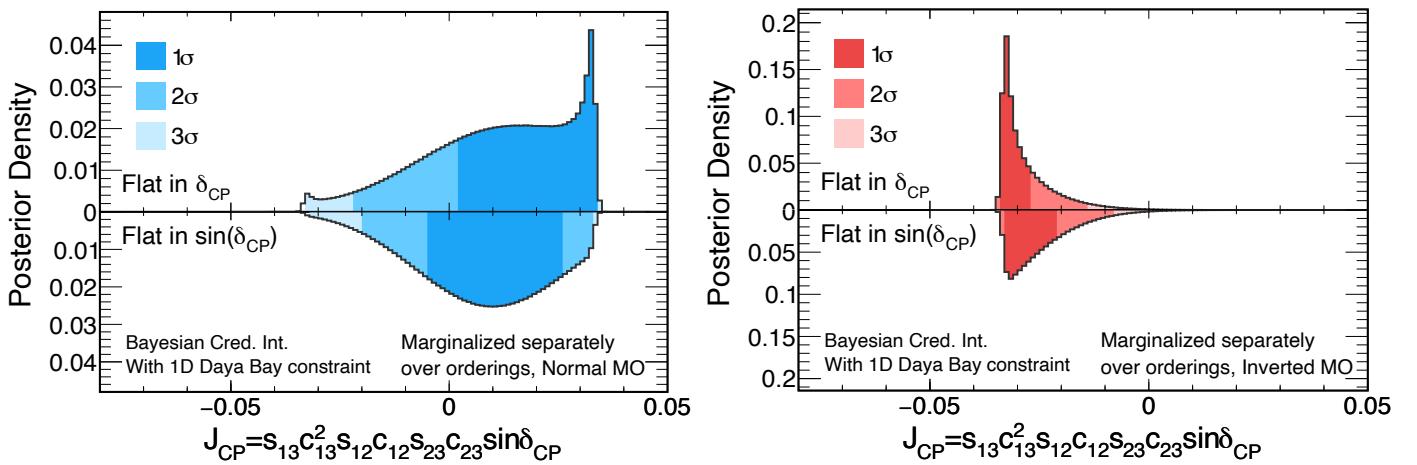


Figure 3.16: Posterior probability densities for the Jarlskog invariant, marginalized separately over each ordering, and shown for normal ordering (blue, left) and inverted ordering (red, right). On each plot, the top panel shows the posterior with a prior uniform in δ_{CP} and the prior uniform in $\sin \delta_{CP}$ in the upside-down canvas. A Jarlskog invariant of 0 represents CP conservation, and non-zero values represent CP violation. The posterior is extracted from a fit with an external 1D constraint on θ_{13} from the Daya Bay experiment. Figure from [87].

Chapter 4

Machine Learning for Neutrino Energy Resolution

Machine learning has been put to great use in neutrino oscillation experiments. In fact, NOvA was the first long-baseline oscillation experiment to use convolutional neural networks (see Section 2.4.1) for event reconstruction/classification. Many other neutrino experiments, such as MicroBooNE, ICECUBE and SBND have also used machine learning to improve their analyses [89–92]. In this chapter we consider an approach wherein, unlike the use discussed in Section 2.4, a neural network is not applied directly to event displays, but is instead used after reconstruction, operating on truth level information to correct for missing energy caused by long-lived, neutral final state particles, such as neutrons. A description of this work can be found in [93].

4.1 DUNE experiment

DUNE is a next-generation long-baseline neutrino oscillation experiment based at Fermilab, Illinois, and the Sanford Underground Research Facility (SAND), South Dakota. Like NOvA, it will use a high-purity neutrino beam produced at Fermilab, where its near detector will be located. However, DUNE’s far detector will be positioned 1297 km downstream, directly on the beam axis. It will employ Liquid Argon Time Projection Chambers (LArTPCs), which offer improved spatial resolution and detailed event imaging. This leads to superior energy resolution and particle identification compared to NOvA. Combined with a more powerful beam, longer baseline, and larger detector mass, these features will enable DUNE to measure quantities unresolved by current experiments such as NOvA and T2K.

In particular, DUNE will break the mass ordering-CP degeneracy seen by NOvA. Its extended baseline increases the matter effect experienced by beam neutrinos, which, along with larger datasets in both neutrino and antineutrino modes will allow DUNE to separate the overlap of normal ordering with $\delta_{CP} = \pi/2$ and inverted ordering with $\delta_{CP} = 3\pi/2$ shown in Figure 3.15. DUNE is consequently expected to determine the neutrino mass ordering at 5σ with only 90 kt-MW-years of exposure [94].

In this chapter and the next we will focus on analysis improvements which could maximise DUNE’s sensitivity, particularly in measuring parameters it will not resolve as easily as the mass ordering, i.e. δ_{CP} and the octant of θ_{23} .

Central to DUNE’s sensitivity are neutrino energy resolution and the reduction of systematic uncertainties. To improve the former, we develop a neural network that uses final state event information to predict the energy of the incident beam neutrino. The DUNE FD will also collect atmospheric neutrino data before the beam is operational. This data will be invaluable, however in this case there is another missing variable - the direction of the incoming neutrino. We therefore also train a network to reconstruct both the energy and direction of the atmospheric neutrinos. The direction is related to the distance the neutrino has travelled – a key ingredient in an oscillation analysis. It is particularly important for sub-GeV neutrinos which can give DUNE sensitivity to CP violation before the beam turns on [95], and which allow for neutrino tomography of the Earth’s interior [96, 97].

Beyond DUNE, our results can also benefit the LArTPCs comprising Fermilab’s Short Baseline Neutrino Program [98].

4.2 Event topologies and missing energy

Neutrino-nucleus charged current interactions occur via four main processes relevant for DUNE, illustrated in Figure 4.1:

- Quasielastic scattering (QE) - neutrino exchanges some energy and momentum with a nucleon via the weak force, leaving the nucleon intact, but with one quark changing flavour.
- Resonance (Res) - neutrino excites a nucleon to a resonant state from which it decays, typically producing a meson.
- Two-particle two-hole (2p2h) - neutrino interacts with two nucleons

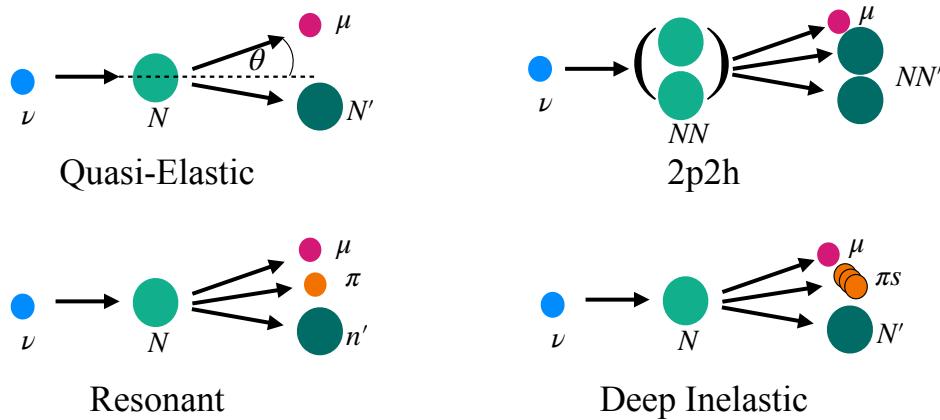


Figure 4.1: Schematics for four main neutrino-nucleon ($\nu - N$) charged current interaction mechanisms. $2p2h$ stands for two-particle two-hole.

in the nucleus causing both to be ejected without an accompanying meson. One mechanism for this is Meson Exchange Currents (MEC), wherein the nucleons interact via meson exchange, but once again, no meson is produced.

- Deep inelastic scattering (DIS) - neutrino interacts with a nucleon's constituent quarks, causing an energetic intranuclear cascade.

For QE interactions the incident neutrino energy can (neglecting detector effects) be reconstructed from simple kinematics. For a ν_μ CC QE interaction,

$$E_\nu = \frac{2M_n E_\mu - M_n^2 + M_\mu^2 - M_p^2}{2(M_n - E_\mu) + \sqrt{E_\mu^2 - M_\mu^2} \cos \theta} \quad (4.2.1)$$

where M_n , M_p and M_μ are the neutron, proton and muon masses, respectively, E_μ is the muon energy, and θ the angle between the muon and the incident neutrino. However, while QE events are the dominant interaction process between 0 – 1 GeV, they are overtaken by Resonant and then DIS at higher energies, as shown in Figure 4.2. These higher energy exchange interactions lead to more complex final states, often containing more neutral particles which the detector cannot reconstruct.

Furthermore, the fraction of the incoming neutrino energy taken up by final state neutral particles is energy dependent, as shown in Figure 4.9, meaning it cannot be corrected with a simple constant offset. Instead, more complex reconstruction techniques are necessary.

After an interaction via one of the processes above, outgoing particles propagate through the nuclear medium and may undergo further interactions with other nucleons before exiting the nucleus. These final state interactions may distort the event signature, meaning the topology seen

by a detector (which, in the case of LArTPCs, will already be missing neutral particles), may differ from the true final state of the interaction.

The combination of neutral long-lived particles, complex event topologies and final state interactions can lead to inaccurate neutrino energy predictions. Nevertheless, it is possible to exploit the complexity of the problem by using the many correlations between the kinematics of final state particles, (as evidenced by the Pearson matrix shown in Figure 4.3), to inform machine learning techniques, which are well suited for dealing with correlations in high-dimensional parameter spaces.

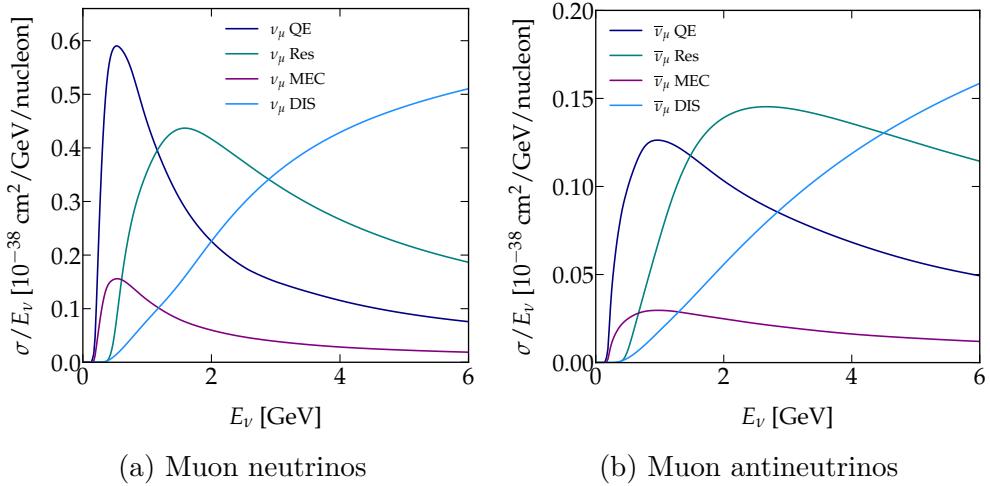


Figure 4.2: Charged current neutrino-liquid argon cross sections used by the GENIE [73] event generator for ν_μ (top left), $\bar{\nu}_\mu$ (top right). There is no appreciable difference between muon and electron (anti)neutrino cross sections. Antineutrino cross sections tend to be a factor of three lower than the corresponding neutrino cross section due to the parity selecting nature of the weak force.

4.3 Dataset

We simulate the DUNE beam by generating 8×10^5 ν_μ -argon charged current events in NuWro 21.09 [99], with neutrino energies distributed according to the DUNE flux in FHC mode [100]. Event generators output ‘true’ particle information., which differs from detector output due to reconstruction efficiencies and detection thresholds. To approximate detector effects, we impose a minimum kinetic energy threshold for charged particles, below which they are excluded from the dataset. We also smear their angles and momenta by the values shown in Table 4.1. To simulate atmospheric neutrinos, we additionally randomize the incoming neutrino direction.

We consider three neutrino reconstruction scenarios. The first, 0_n ,

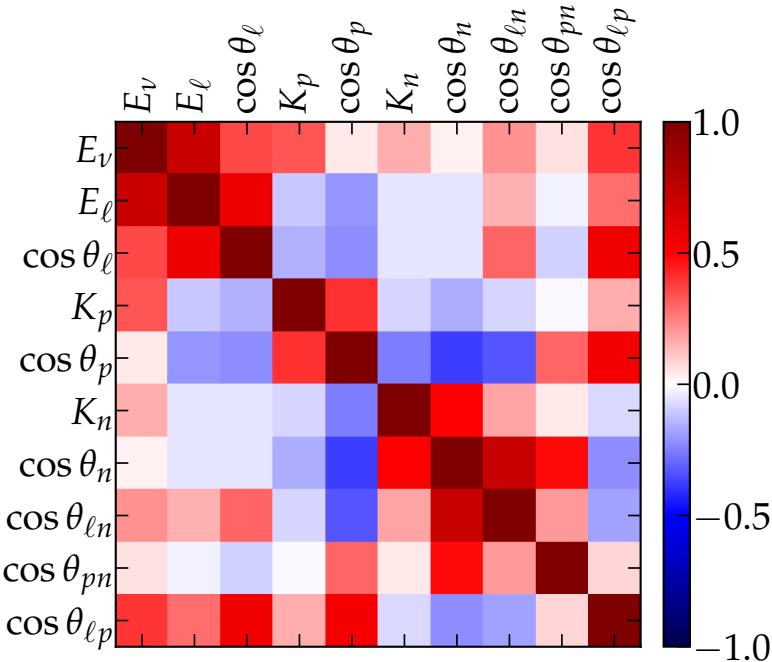


Figure 4.3: Example of a Pearson correlation matrix between several observables in charged-current neutrino–argon interactions with at least one proton and one neutron in the final state, for the DUNE neutrino-mode flux and before including detector responses. We include the total kinetic energy of all protons (K_p) and neutrons (K_n) per event, the energies of the neutrino and outgoing lepton (E_ν , E_ℓ), the angle between a given particle and the beam axis $\cos\theta_{\ell,p,n}$, and the angles between particles $\cos\theta_{\ell p, \ell n, p n}$. Note that neutrons are very challenging to reconstruct, so information on the neutron system is typically not available in realistic event records.

Particle	Threshold	resolution coefficient α_p	$\sigma(\theta)$
μ, e, γ	30 MeV	5%	2°
π, K, proton	30 MeV	10%	10°
neutron “ 0_n ”	invisible	—	—
neutron “ E_n ”	100 MeV	40%	—
neutron “ $E_n + \theta_n$ ”	100 MeV	40%	10°

Table 4.1: The kinetic energy threshold, momentum resolution coefficient α_p (which is used to calculate the resolution as a function of momentum using $\sigma(p) = \alpha_p \sqrt{p/\text{GeV}}$), and angular resolution for different final-state particles. For neutron reconstruction, we list the three different scenarios. The values in this table were developed in consultation with members of the DUNE collaboration for Ref. [96].

assumes no reconstruction capabilities for neutrons. The second, E_n , is motivated by attempts to reconstruct some neutron energy [101] by using the small deposits of energy, or ‘blips’, they leave when propagating through liquid argon. Here we assume reconstruction of the neutron energy with a fractional resolution of $40\%/\sqrt{K_n/\text{GeV}}$. These blips have been detected [102], however neutron reconstruction has not been firmly demonstrated. As such, the third neutron reconstruction scenario, $E_n + \theta_n$ is an optimistic one, where we allow reconstruction of the neutron direction in addition to the energy, with a resolution of 10°. This final configuration serves to quantify the effects of neutron direction information on energy resolution.

Although the detector simulation attempts to account for the differences in simulated and real data, reconstruction of experimental data remains more challenging due to issues such as the misidentification of charged pions and protons. Our work therefore serves as a motivation to what could be achieved by leveraging machine learning to improve neutrino energy resolution.

To provide a benchmark for energy resolution performance in our detector simulation, we will compare our results to a purely calorimetric method, in which the neutrino energy is obtained as

$$E_\nu^{\text{cal}} = E_\ell + \sum_i^{\text{mesons}} E_i + \sum_i^{\text{baryons}} K_i, \quad (4.3.1)$$

where E_i and K_i denote the total and kinetic energy of a particle i , respectively.

4.4 Neural Network

We use a dense neural network (DNN) [103, 104] to predict the neutrino energy for beam events, and the neutrino energy, zenith and azimuth angles for atmospheric events. The DNN input is the summed four-momentum of each particle species produced in the event. For example, the energy of 3-momenta of all protons in a given event are summed and passed to the network as a single 4-vector. We consider sixteen particle types, and each element of the 4-momentum corresponds to an input neuron, so the input layer has 64 notes. The network architecture is $[64 - 16 - 32 - 1]$ for beam neutrinos and $[64 - 16 - 32 - 3]$ atmospheric neutrinos. The reduced dimensionality after the input layer allows for more effective training.

Summing over all particles of a given species improves the network's robustness to inaccuracies in neutrino interaction models used in event generators, or differences in event generator output. As shown in Figure 4.4 GENIE [73] and NuWro [99] predict different numbers of low-energy final state protons. Aggregating four-momenta reduces the networks sensitivity to such discrepancies.

The network is trained using a fractional mean squared error loss function for beam events,

$$L_{\text{energy}} = 10 [1 - E_{\text{Pred}}/E_{\text{True}}]^2 \quad (4.4.1)$$

where E_{True} and E_{Pred} denote the true and predicted neutrino energy.

For atmospheric neutrinos an angular term is added

$$L_{\text{angle}} = 30 \arccos^2 [\hat{v}_{\text{True}(x)} \cdot \hat{v}_{\text{Pred}(x)}] \quad (4.4.2)$$

where $\hat{v}_{\text{True}(x)}$ and $\hat{v}_{\text{Pred}(x)}$ are unit vectors in the true and predicted neutrino momentum directions, respectively. The factor of 30 on L_{angle} weights angular reconstruction more heavily than the less challenging energy reconstruction, and was optimised by training multiple networks with different scaling factors on the atmospheric loss term.

The training curves for the beam and atmospheric neutrino DNN's can be seen in Figure 4.5.

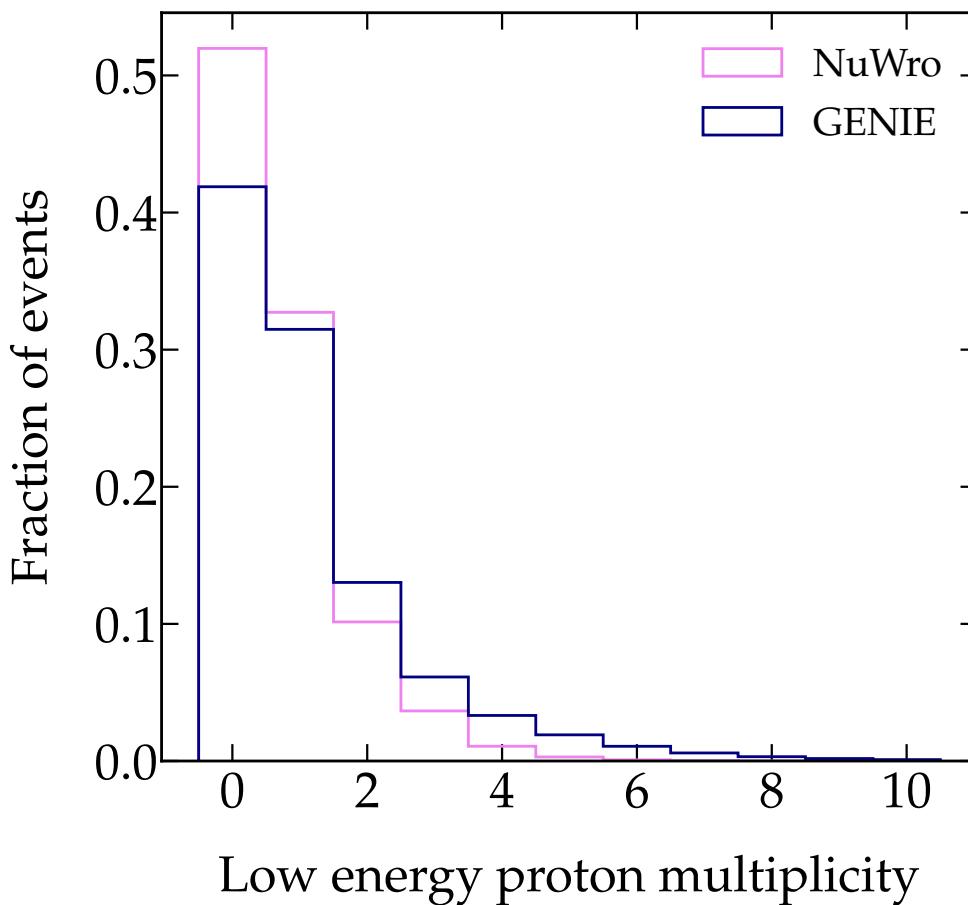


Figure 4.4: Fraction of final state protons with $E < 1500$ MeV per event for NuWro (pink) and GENIE (blue) datasets. GENIE predicts higher numbers of final state low energy protons than NuWro.

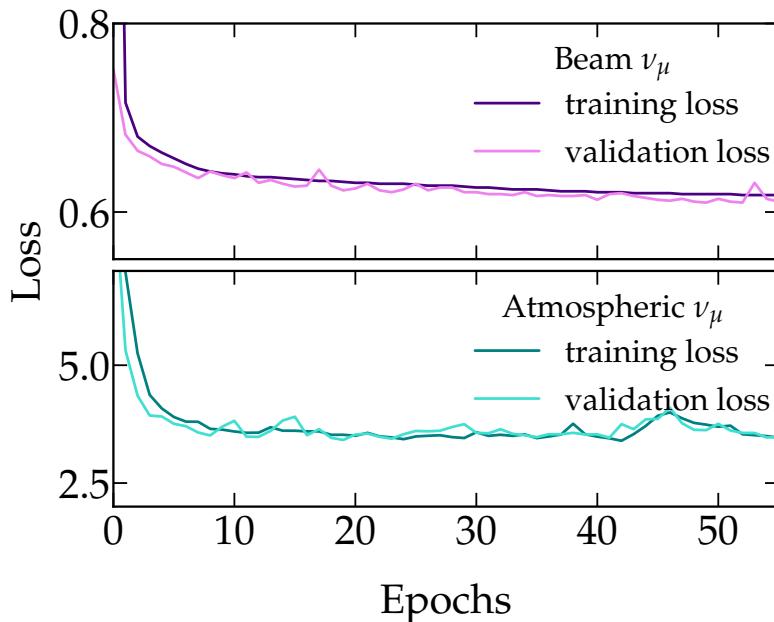


Figure 4.5: DNN loss functions when trained on beam neutrino events (upper panel) and atmospheric neutrino events (lower panel) for training (darker) and validation (brighter) data. These results are for the no neutron reconstruction scenario, but results are similar for the neutron energy and neutron energy and direction reconstruction cases.

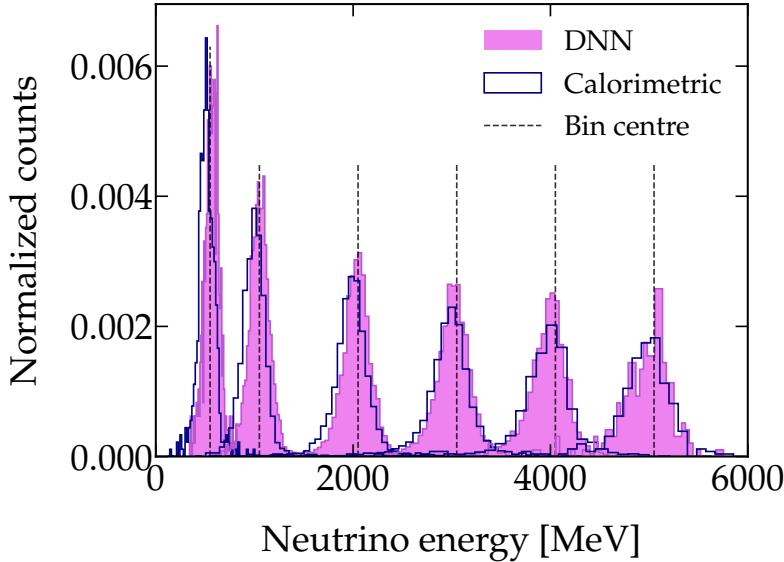


Figure 4.6: Normalized event distribution using DNN energy reconstruction (filled histogram) and the calorimetric method (unfilled histogram). Neutrinos with true energies uniform across six ranges, [500 – 600, 1000 – 1100, 2000 – 2100, 3000 – 3100, 4000 – 4100, 5000 – 5100 MeV], were used. The DNN reconstruction produces a narrower event distribution and reduces the bias in the mean reconstructed energy compared to the calorimetric method.

4.5 Results

Figure 4.6 shows normalised event distributions for beam ν_μ with true energies uniformly distributed across ranges [500 – 600, 1000 – 1100, 2000 – 2100, 3000 – 3100, 4000 – 4100, 5000 – 5100 MeV]. Results are shown for calorimetric reconstruction (unfilled histogram) and the DNN (filled histogram), both with no neutron information. The DNN reduces the bias inherent in the calorimetric method, improving the accuracy of reconstructed neutrino energy.

To further evaluate the model’s performance, we calculate the fractional energy resolution, $\sigma(E_\nu)/E_\nu$, of distributions like those shown in Figure 4.6 across multiple energies. As shown in Figures 4.7 and 4.8 the DNN 0_n outperforms the calorimetric method, and surpasses the resolutions quoted in DUNE’s Technical Design Report (TDR) [105] and Conceptual Design Report (CDR) [106]. This implies that the DNN is able to partially infer the kinematics of invisible particles from correlations between other kinematic variables. Adding limited neutron reconstruction (E_n) leads to further improvement in $\sigma(E_\nu)$, particularly at higher incoming neutrino energies, where complex final states are more likely. Including neutron direction reconstruction yields no further improvement.

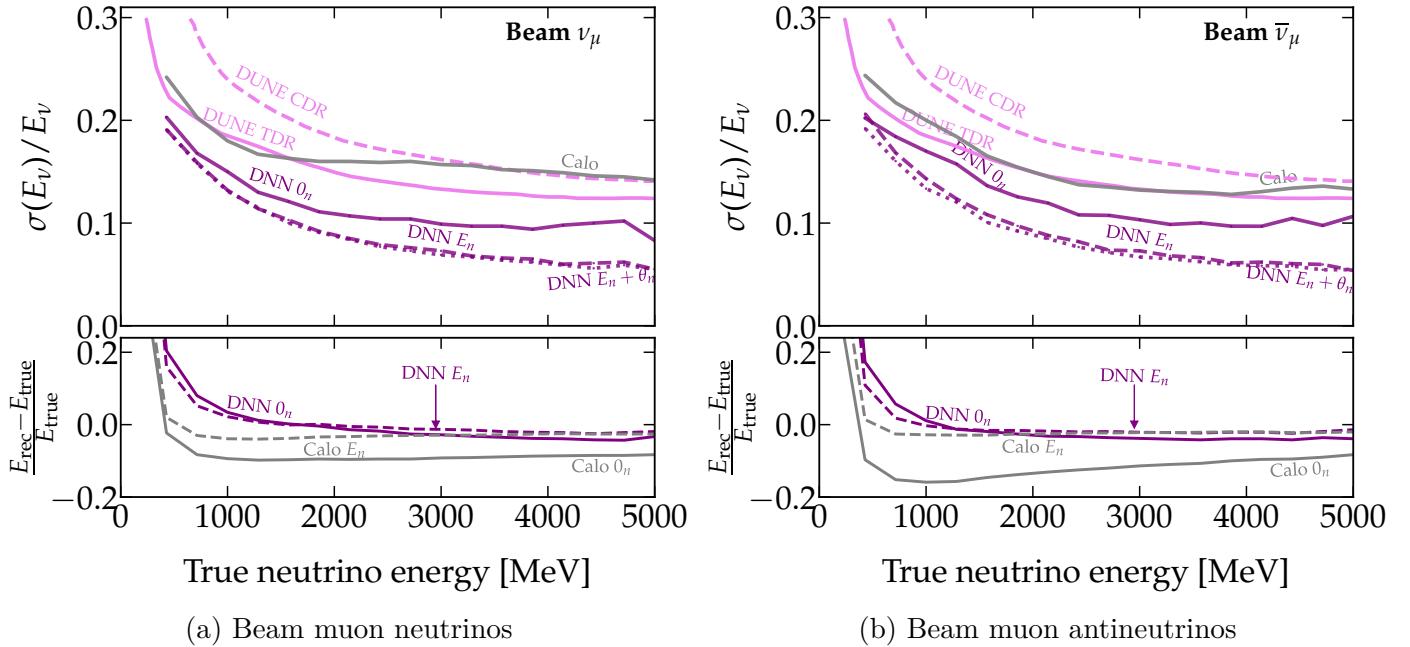
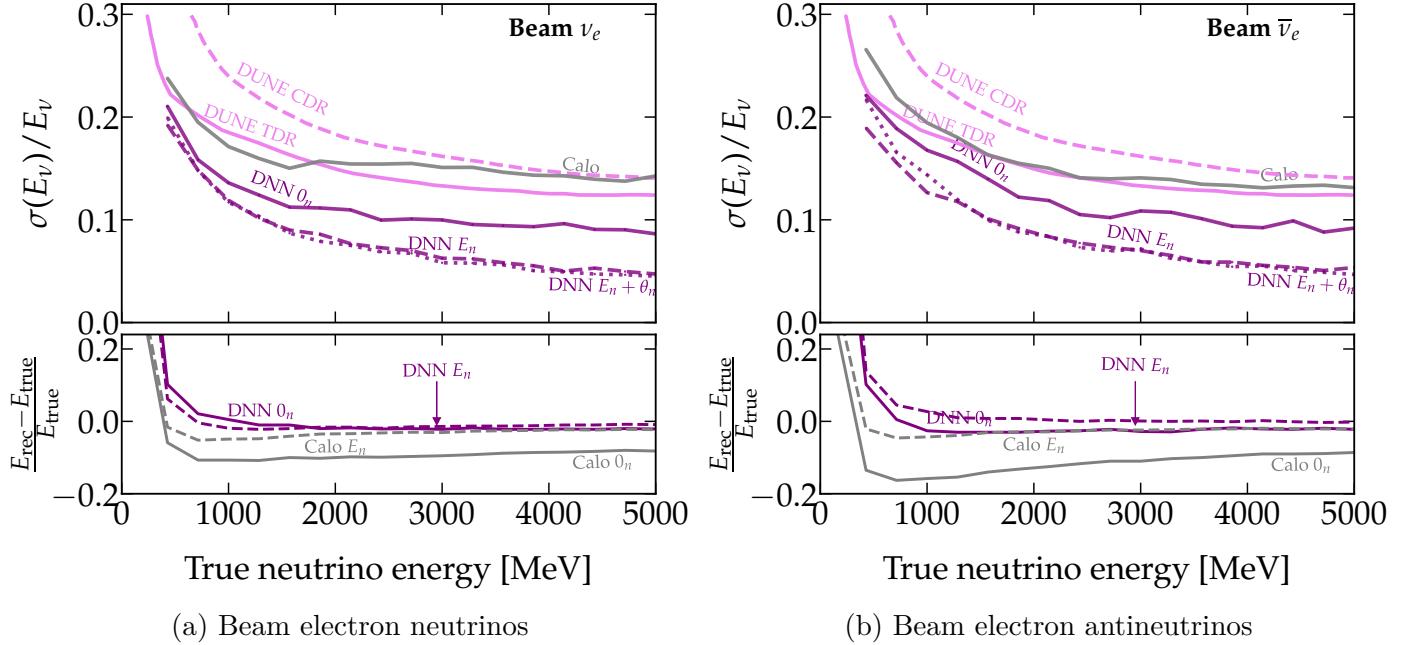


Figure 4.7: *Top:* fractional neutrino energy resolution $\sigma(E_\nu)/E_\nu$ as a function of neutrino energy from DNN-based analyses trained on muon neutrino (left) and antineutrino (right) events with no information on final-state neutrons (solid purple), with limited information on the neutron energy (dashed purple), and with information on the neutron energy and direction (dotted purple). When information on the energy of final-state neutrons is available, the improvement is more than a factor of two at high energies. The energy resolutions anticipated in the DUNE CDR and TDR simulations can be seen in magenta, and the resolution of a simple calorimetric method assuming invisible neutrons in grey. *Bottom:* reconstruction bias for the DNN compared to the calorimetric method.



(a) Beam electron neutrinos

(b) Beam electron antineutrinos

Figure 4.8: *Top:* fractional neutrino energy resolution $\sigma(E_\nu)/E_\nu$ as a function of neutrino energy from DNN-based analyses trained on electron neutrino (left) and antineutrino (right) events with no information on final-state neutrons (solid purple), with limited information on the neutron energy (dashed purple), and with information on the neutron energy and direction (dotted purple). When information on the energy of final-state neutrons is available, the improvement is more than a factor of two at high energies. The energy resolutions anticipated in the DUNE CDR and TDR simulations can be seen in magenta, and the resolution of a simple calorimetric method assuming invisible neutrons in grey. *Bottom:* reconstruction bias for the DNN compared to the calorimetric method.

Reconstruction method	α	β
DNN	0.099	0.046
Calorimetric	0.084	0.102

Table 4.2: Best fit parameters from fitting the no neutrino information ν_μ DNN and calorimetric fractional energy resolution curves shown in Figure 4.7 to (4.5.1). While the α parameter, which dictates how fast the resolution falls with increasing energy, is similar for calorimetric and DNN methods, the β parameter, which controls the overall normalisation of the energy resolution, is smaller for the DNN case. This implies that the DNN improves energy resolution uniformly across energies when compared to the calorimetric method.

Fitting the DNN and calorimetric $\sigma(E_\nu)/E_\nu$ curves to the energy resolution function

$$\sigma(E_\nu)/E_\nu = \alpha/\sqrt{E_\nu} + \beta \quad (4.5.1)$$

gives the parameters shown in Table 4.2. The similar α values indicate that both methods scale similarly with energy, but the DNN improves energy resolution uniformly across neutrino energies, as shown by its smaller β value.

It should be noted that, as shown in Figure 4.9, the fraction of the total neutrino energy taken up by the final state neutron also falls at higher neutrino energies, where the missing energy accounts for less of the total event energy, and reconstruction is less challenging.

Figures 4.7 and 4.8 show energy resolutions for beam $\bar{\nu}_\mu$, ν_e and $\bar{\nu}_e$. The network performance is consistent across flavours, despite the calorimetric method achieving better resolution for antineutrino events at higher energies, and poorer resolution at low energies. This behaviour is expected; $\bar{\nu}$ quasielastic CC scattering typically produces a neutron, whereas ν_{CC} scattering produces a proton. The former (in the 0_n configuration) is invisible to the detector, meaning at $E_\nu < 2.7$ GeV where QE interactions dominate (see Figure 4.2), calorimetric reconstruction is poorer. The DNN shows the same behaviour.

For atmospheric neutrinos, the DNN improves energy resolution, but, as shown in Figure 4.10, does not increase angular resolution. The calorimetric neutrino direction is defined as the vector sum of the three-momenta of all visible outgoing particles. This method's performance proves comparable to the DNN's, except at higher energies. This may be due to lower energy events having a smaller hadronic energy fraction, leaving little room for the angular reconstruction to be improved further.

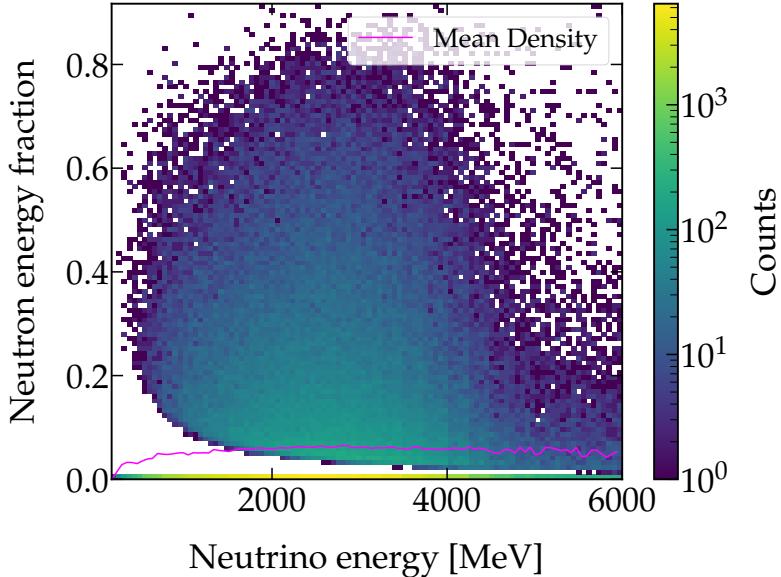


Figure 4.9: Fraction of ν_μ CC event energy taken up by neutron as a function of the energy of the incident neutrino. In most ν_μ CC events a proton is produced, but in those events with a neutron, the fraction of energy it takes up is not constant. At energies greater than 4000 MeV there are far fewer events in which the neutron energy accounts for more than 40% of the true neutrino energy, making reconstruction less challenging for the DNN.

Results at low energies could be improved by using a DNN-based classifier to separate events into several samples based on the accuracy of the energy reconstruction [53, 95].

4.6 Impact on Oscillation Analyses

Here we use the GLoBES framework to quantify the impact of improved energy resolution on DUNE’s sensitivity to leptonic CP violation, the octant of θ_{23} , and the neutrino mass ordering. For CP violation, the sensitivity is found by scanning over true values of δ_{CP} and testing how well CP conservation ($\delta_{CP} \in \{0, \pi\}$) could be excluded by finding the χ^2 at each value. Similarly, for the θ_{23} octant and mass ordering we scan over regions of phase space which cover each scenario (upper or lower octant, inverted or normal ordering), and test exclusion of the alternate scenario at each point in the scan.

Throughout, we compare the energy resolution of the DNN without neutron information (see Equation 4.5.1 and Table 4.2), to results from DUNE TDR configurations, as outlined in [100, 105]. The statistical analysis and fit used code originally developed in [107–109].

Figure 4.11 shows the impact of DNN-achieved energy resolution on

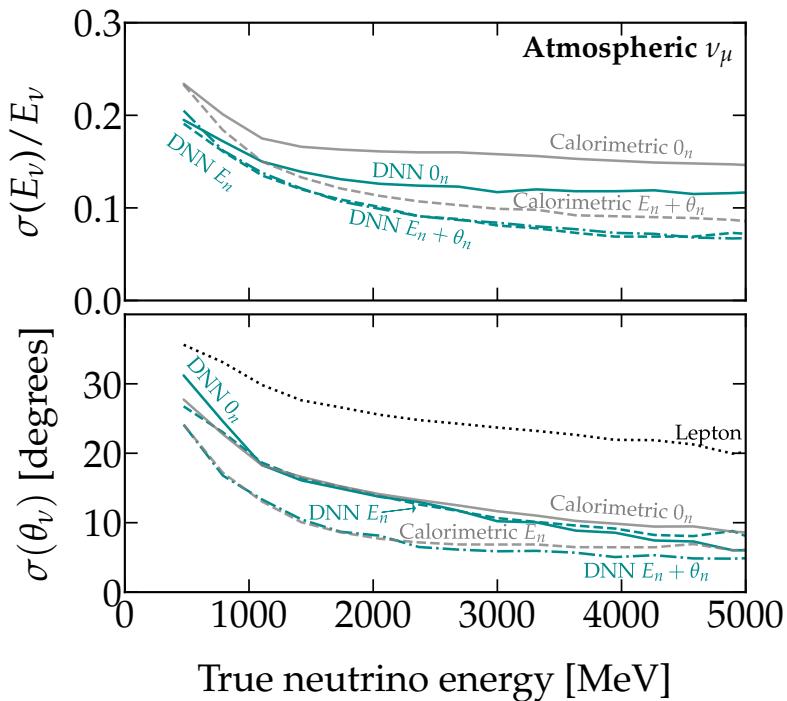


Figure 4.10: Fractional neutrino energy resolution, $\sigma(E_\nu)/E_\nu$ (top), and angular resolution, $\sigma(\theta_\nu)$ (bottom) from DNN-based analyses of atmospheric neutrino events with no information on final-state neutrons (solid cyan), and with information on neutron energies and directions included (dashed/dot-dashed cyan). For comparison, we also show the resolutions achievable with simple calorimetric methods (gray curves). The dotted black curve in the bottom panel is based on only the charged lepton kinematics, as in Cherenkov detectors at low energy.

the ν_μ disappearance and ν_e appearance event spectra at the DUNE FD for the case of maximal CP violation. In the disappearance channel we see oscillations become more pronounced, particularly the oscillation maximum at 2.2 GeV, which translates to more accurate measurements of Δm_{32}^2 and θ_{23} (see Section 2.1). Improvements in the appearance channel are more modest, although the second oscillation maximum, at 1.2 GeV for a baseline of 1300 km, becomes more defined with improved energy resolution. This should boost sensitivity to CP violation, as lower neutrino energies correspond to relatively larger L/E terms, which translate to enhanced CP effects on neutrino oscillation probability (see Equation 2.1.1).

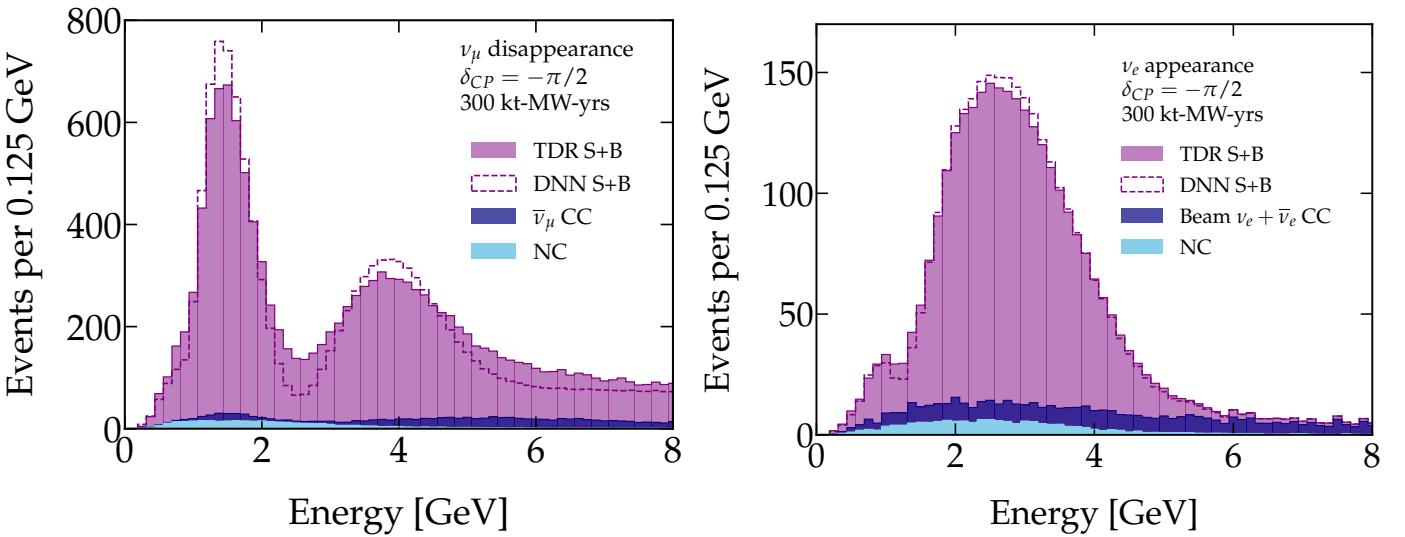


Figure 4.11: Impact of neural network-improved energy resolution on the energy spectrum of ν_μ events in DUNE. Here, S and B denote signal and background.

Looking at the effects on CP violation sensitivity, we can see in Figure 4.12 that the greatest gains occur for $\delta_{CP} = -\pi/2$, i.e. at maximal values. In this scenario the updated energy resolution allows for earlier 5σ sensitivity to CP violation, corresponding to a 10% increase in exposure. Smaller improvements can be seen for the less violating $\delta_{CP} = -\pi/4$. Sensitivity to non-maximal θ_{23} , shown in Figure 4.12 with comparison to $\theta_{23} = 40^\circ$ and 50° , also benefits from increased energy resolution. DUNE will the mass ordering at 5σ with only 100 kt MW years of exposure, meaning this measurement benefits less from improved $\sigma(E_\nu)$ and is not included here [110].

We can also test the effects of improved $\sigma(E_\nu)$ on the relative precision of oscillation parameter measurements. Figure 4.13 shows that the mass splitting Δm_{31}^2 and mixing angles θ_{23} and θ_{13} could be measured more precisely in the case of maximal and minimal CP violation. The precision

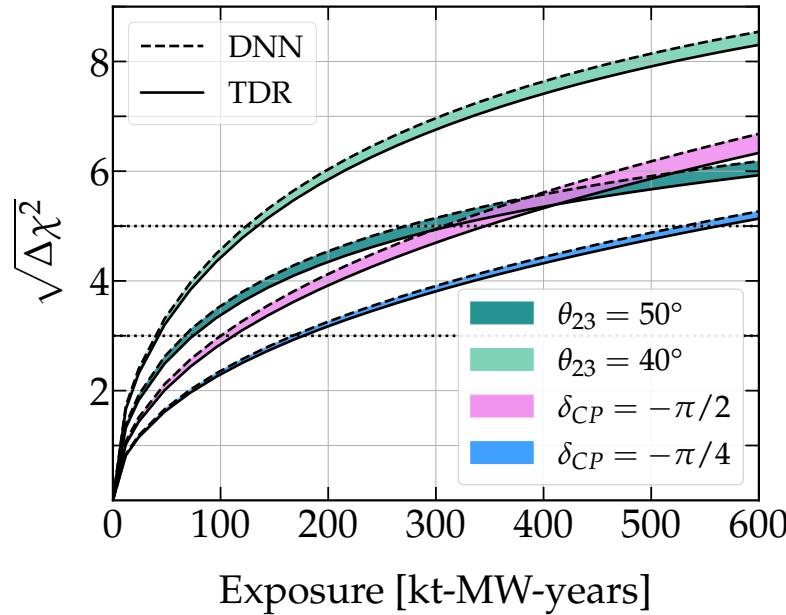


Figure 4.12: Impact of neural network-improved energy resolution on precision oscillation measurements in DUNE. The solid lines show the DUNE TDR's predicted sensitivity, and the dashed lines the sensitivity achievable using the improved energy resolution afforded by the DNN. For the sensitivity to CP violation and to non-maximal θ_{23} , the improvement in sensitivity due to the DNN is equivalent to a $\sim 10\%$ increase in exposure.

of δ_{CP} only sees an increase in the case of maximal violation.

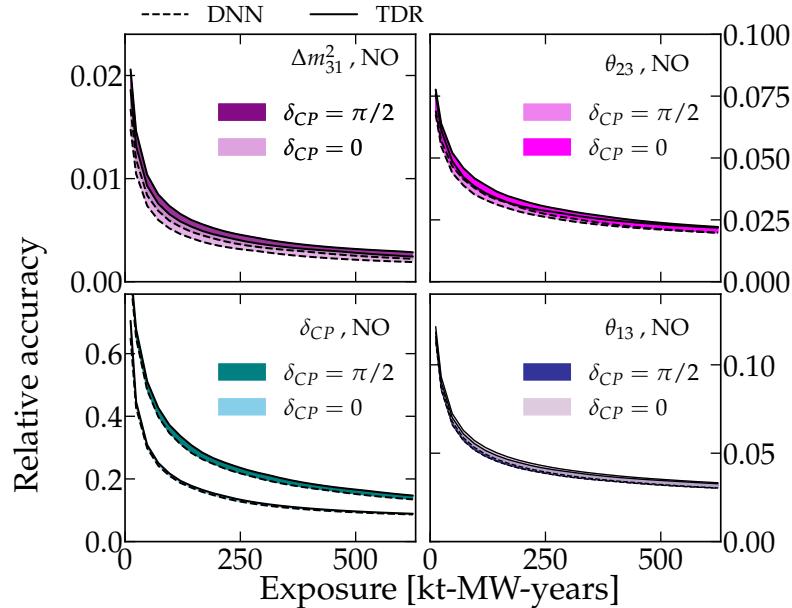


Figure 4.13: Impact of neural network-improved energy resolution on precision measurements of Δm_{31}^2 , θ_{23} , δ_{CP} , and θ_{13} in DUNE. The solid lines show the DUNE TDR's predicted precision, and the dashed lines the precision achievable using the improved energy resolution afforded by the DNN.

4.7 Dependence on neutrino–nucleus cross section modeling

Current models of neutrino–nucleus interactions exhibit significant discrepancies with experimental data [111–115]. As a result, the DNN, which is trained on Monte Carlo events generated by NuWro, could exhibit poorer performance when applied to data if it learns features of mis-modelled interactions.

To obtain a qualitative understanding of the impact neutrino–nucleus interaction uncertainties may have on DNN performance, we apply our network trained on NuWro 21.09 events to mock data generated with GENIE 3.4.0. In the upper panel of Figure 4.14, the navy lines show the performance of the network trained on NuWro events, but tested on Genie events (abbreviated to G-NW DNN). For comparison, the purple lines show the network trained and tested on NuWro only, (NW-NW 0_n DNN). Across all energies, the NW-NW 0_n DNN outperforms the G-NW 0_n DNN. In fact, below 1.5 GeV, the G-NW 0_n DNN even performs more poorly than 0_n calorimetric reconstruction (solid grey line). This is likely due to differences in how the generators calculate intranuclear cascades and low energy final state nucleons. Adding neutron energy information significantly improves performance over the 0_n case, but the G-NW E_n DNN (dashed navy line) still only slightly outperforms the equivalent calorimetric result (dashed grey line).

More importantly, however, is the bias shown in the lower panel of Figure 4.14. Here it is apparent that the G-NW DNN underpredicts neutrino energies at order $\sim 10\%$ across most of the energy spectrum, with larger bias when neutron kinematics are included. This again highlights discrepancies in the treatment of final-state nucleons across event generators. The G-NW DNN 0_n bias is nonetheless smaller than the comparable calorimetric method up to energies of several GeV. Contrastingly, the G-NW DNN E_n exhibits larger bias than E_n calorimetry across all energies.

Overall, this serves as evidence that improved neutrino energy resolution is crucial for maximising the physics potential of the upcoming DUNE experiment, particularly for measuring CP violation and the octant of θ_{23} . Furthermore, we have shown that understanding neutrino–nucleus interactions is critical to ensure the accuracy and precision of next generation experiments.

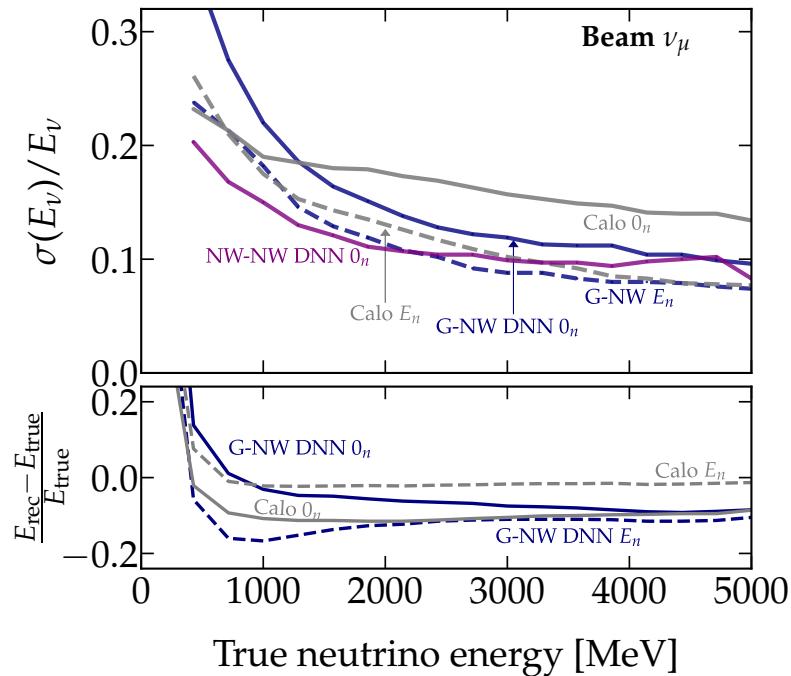


Figure 4.14: *Top*: fractional neutrino energy resolution, $\sigma(E_\nu)/E_\nu$ from a DNN trained on NuWro events, and applied to GENIE (blue) and NuWro (purple) events. We compare to the performance of the calorimetric method when applied to GENIE events with no neutron information (grey solid line), and with neutron energy included (dashed grey line). *Bottom*: Energy reconstruction bias for the DNN trained on the “wrong” neutrino–nucleus interaction model, and for the calorimetric method.

Chapter 5

Reducing systematic uncertainties with DUNE-PRISM

5.1 DUNE-PRISM

Neutrino physics will move from being statistically to systematically limited with the arrival of next generation experiments such as DUNE [105] and HyperKamiokande [55]. These increasingly critical systematic uncertainties will be dictated by our limited understanding of hadron production and neutrino-nucleus interactions, i.e. neutrino fluxes and cross sections.

DUNE’s primary tool for reducing these uncertainties is its suite of near detectors, which, as in the NOvA experiment, can be used to tune simulations to mitigate the effect of mis-modeling. The DUNE ND will have the additional ability to move off-axis from the beamline, as shown in Figure 5.1, enabling measurements of multiple neutrino fluxes from

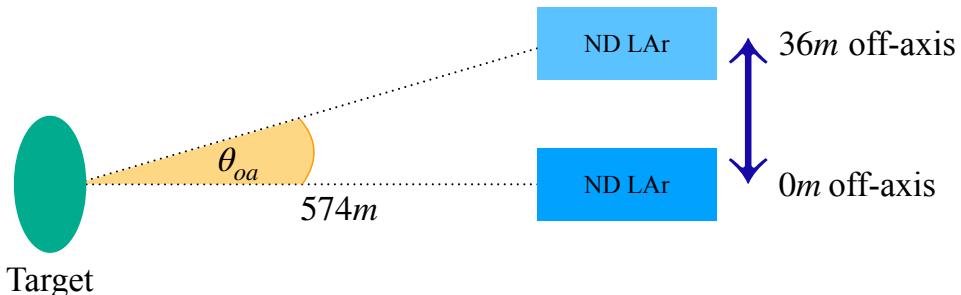


Figure 5.1: Schematic of the DUNE-PRISM setup [116]. The detector will sit on moveable rails, meaning it can take data *at any position* between 0 m and 36 m off-axis.

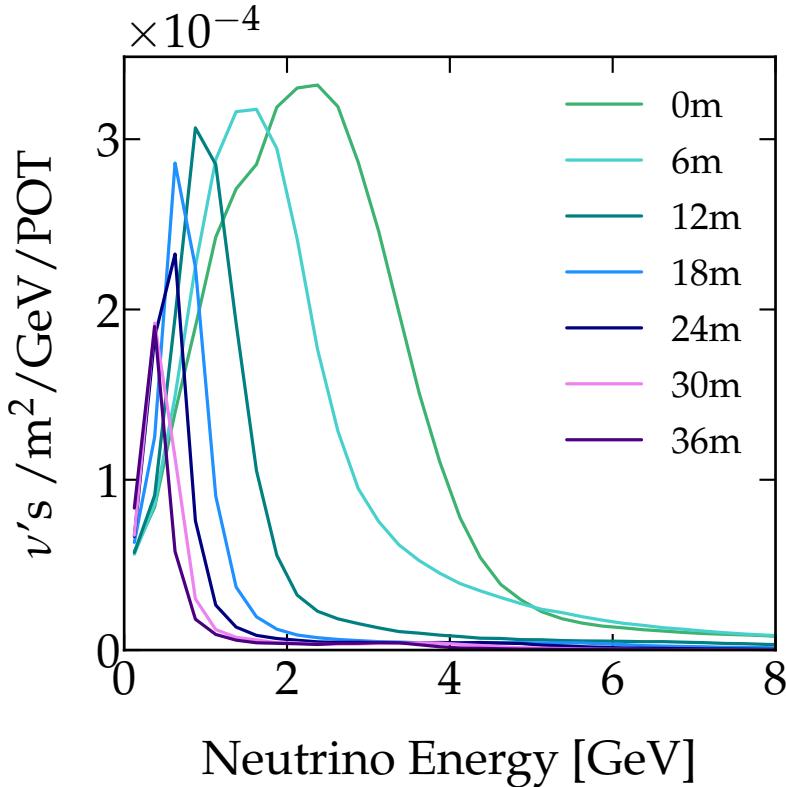


Figure 5.2: Total neutrino flux in the DUNE/LBNF beam at different off-axis positions. The flux moves towards lower neutrino energies and becomes more peaked as the detector moves further off axis.

the same beam [116]. Interestingly, the NuMI beamline initially had a moveable target for a similar purpose. In this chapter, we will discuss how these off-axis measurements can directly constrain neutrino flux uncertainties, the effect these constraints have on DUNE’s sensitivity to oscillation parameters.

While the muon neutrino beam produced by DUNE’s PIP-II accelerator [117] is primarily produced by charged pion decays, (as in the NuMI beam), there is also a non-negligible contribution from kaons and, to a lesser extent, secondary muon decays. The kinematics of two body boosted meson decays at small angles can be described by

$$E_\nu = \frac{(1 - (m_\mu/m_M)^2)E_M}{1 + \gamma^2\theta^2} \quad M = \pi^\pm, K^\pm \quad (5.1.1)$$

where E_M and m_M are the meson’s energy and mass, $\gamma = E_M/m_M$ and θ is the angle between the neutrino and parent meson. From this expression it is clear that neutrinos produced at angles $\theta > 0$ relative to the parent meson have lower energies, as illustrated in Figure 5.2. Because pions are lighter than kaons, they are more strongly boosted in the beam direction, and tend to produce more forward-going neutrinos, leading to the flux being dominated by π^\pm decays on-axis. Kaons, however, pro-

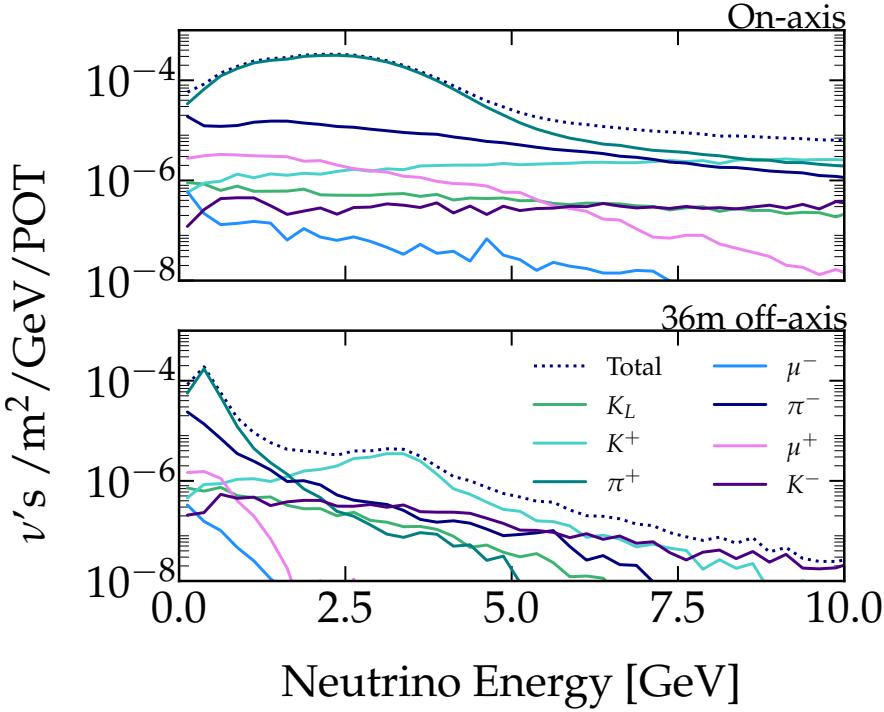


Figure 5.3: Total neutrino flux and the constituent flux components split by parent hadron in the on-axis (top) and maximally off-axis (bottom) positions. On-axis the flux is dominated by π^+ decays, while maximally off-axis the K^+ component has a distinct shape and greater relative importance.

duce neutrinos with a broader energy distribution, and at larger angles relative to the beamline. Consequently, as the detector moves off-axis, the kaon contribution becomes more visible and has increasing relative importance at the flux peak. This effect is illustrated in Figure 5.3.

The PRISM (Precision Reaction Independent Spectrum Measurement) method combines ND flux measurements from multiple off-axis positions to obtain a prediction of the oscillated flux at the FD [118]. For ND and FD fluxes $\phi_{\text{ND}}(E; \theta_{\text{oa}})$ and $\phi_{\text{FD}}(E; \vec{\Theta})$,

$$\phi_{\text{FD}}(E; \vec{\Theta}) = \sum_j c_j(\Theta) \phi_{\text{ND}}(E_i; \theta_{\text{oa},j}) \quad (5.1.2)$$

where $\vec{\Theta} = (\theta_{12}, \theta_{13}, \theta_{23}, \delta_{CP}, \Delta m_{21}^2, \Delta m_{31}^2)$, and i and j the energy bins and off-axis angles, respectively. The coefficients $c_j(\Theta)$ are determined using MC simulated ND and FD fluxes, such that

$$c_j = \sum_i \left[\left(\phi_{\text{ND}}^{MC} \right)^{-1} \right]_{ji} \phi_{\text{FD},i}^{MC} \quad (5.1.3)$$

c_j are applied to measured ND data to obtain an FD prediction with reduced systematic uncertainties compared to ϕ_{FD}^{MC} . FD predictions for

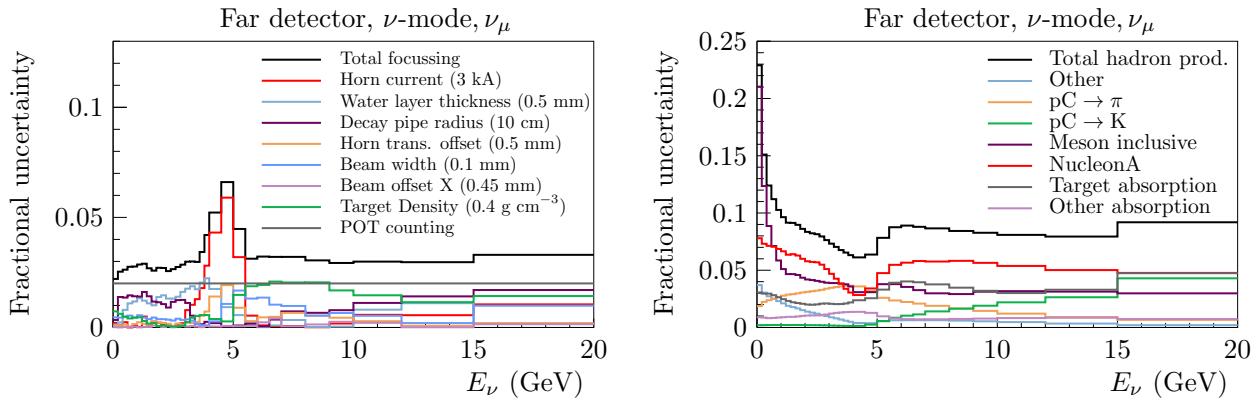


Figure 5.4: Uncertainties on beam focusing (left) and hadron production (right) in the ν_μ flux at DUNE’s Far Detector. From [100]

different sets of oscillation parameters are then compared to the observed FD data using a standard maximum likelihood fit. A new set of coefficients must be calculated for each set of oscillation parameters $\vec{\Theta}$.

This method reduces uncertainties on neutrino cross sections and total neutrino fluxes (because these uncertainties affect the near and far detectors identically), and can improve the mapping between true and reconstructed neutrino energy [118]. However, PRISM is vulnerable to errors in the assumed proportions of parent hadrons in the beam, which could cause differences in the shapes of the assumed and true flux spectra. We propose the **L**ateral **E**xtraction of **N**eutrino **S**pectra (LENS), an ND only joint fit over multiple off-axis angles. By exploiting the additional information on the beam composition afforded by DUNE’s moveable ND, LENS aims to improve PRISM spectrum predictions and reduce hadron production uncertainties which, as seen in Figure 5.4, will be the dominant beam production uncertainty at DUNE.

5.2 Generating fluxes

We generate neutrino fluxes using ROOT Ntuple files provided by the DUNE collaboration [119], which contain outputs of MC simulations of the DUNE beam. This simulation accounts for hadron production in the target, the subsequent absorption or decay of these hadrons, and the full beamline geometry. Only information about hadrons which decay to produce neutrinos is stored in these files. Each simulated neutrino is labelled with its parent particle, said parent particle’s position and momentum at the time of decay, and the decay channel. The relevant decay modes and their contribution to the neutrino flux are listed in Table 5.1. Fluxes are formed using these neutrinos with their weights adjusted

to account for the detector moving off-axis. However, performing the PRISM procedure requires ~ 40 fluxes at varying off-axis positions, so to ensure adequate statistics we decay each parent meson 200 times using the phasespace package [120] which, due to the probabilistic nature of decays, produces many random kinematic configurations of daughter neutrinos.

To assign geometric weights to neutrinos produced by the above decays, we simulate the flux propagating through the ND-LAr (Liquid Argon Near Detector) described in [121]. This detector is a cuboid with dimensions 5 m (along the beam axis) \times 7 m (transverse to the beam axis) \times 3 m, and sits 574 m downstream from the target. Following convention in literature, we halve the width (transverse measurement) of the detector to ensure we consider only neutrinos passing through its centre. We select neutrinos which pass through this simulated ND by calculating their intersection vertices using the Trimesh package [122]. Moving the detector off-axis by a distance d then corresponds to an off-axis angle $\tan \theta = d/574$ rad. For each off-axis position the flux is expressed as a sum of five contributions, corresponding to neutrinos with different parent hadrons,

$$\phi_{\text{ND},\alpha}(E; \theta_{\text{oa}}) = \sum_p r_p \phi_{\text{ND},\alpha}^{(p)}(E; \theta_{\text{oa}}) \quad (5.2.1)$$

where $p = K_L, K^+, K^-, \pi^+, \pi^-$ are the neutrino parent particles and the coefficients r_p are their flux component normalisations we will vary in the fit. We do not treat neutrinos from μ^+ and μ^- independently because 98% of μ^\pm in the beam originate from π^\pm and K^\pm decays. The normalisation of the μ^\pm daughter neutrino flux therefore depends on the pion and kaon normalisation, and can be expressed as

$$\phi_{\text{ND},\alpha}^{\mu^\pm}(E; \theta_{\text{oa}}) = r_{\pi^\pm} \phi_{\text{ND},\alpha}^{\pi^\pm \rightarrow \mu^\pm}(E; \theta_{\text{oa}}) + r_{K^\pm} \phi_{\text{ND},\alpha}^{K^\pm \rightarrow \mu^\pm}(E; \theta_{\text{oa}}) \quad (5.2.2)$$

where $\phi_{\text{ND},\alpha}^{M^\pm \rightarrow \mu^\pm}(E; \theta_{\text{oa}})$ is the flux of neutrinos produced by muons whose parent was meson M .

5.3 Fit

We calculate the neutrino event spectrum by folding the flux $\phi_{\text{ND},\alpha}(E; \theta_{\text{oa}})$ with the associated cross section, $\sigma_\alpha(E)$. As discussed in Section 4.7, neutrino cross sections are poorly understood and prone to mis-modelling. For this reason, although we use cross sections from GENIE for each

Decay process	Beam contribution (%)
$K_L^0 \rightarrow \nu_e + \pi^- + e^+$	0.522
$K_L^0 \rightarrow \bar{\nu}_e + \pi^+ + e^-$	0.540
$K_L^0 \rightarrow \nu_\mu + \pi^- + \mu^+$	0.353
$K_L^0 \rightarrow \bar{\nu}_\mu + \pi^+ + \mu^-$	0.356
$K^+ \rightarrow \nu_\mu + \mu^+$	5.35
$K^+ \rightarrow \nu_e + \pi^0 + e^+$	0.428
$K^+ \rightarrow \nu_\mu + \pi^0 + \mu^+$	0.275
$K^- \rightarrow \bar{\nu}_\mu + \mu^-$	1.78
$K^- \rightarrow \bar{\nu}_e + \pi^0 + e^-$	0.145
$K^- \rightarrow \bar{\nu}_\mu + \pi^0 + \mu^-$	0.0911
$\mu^+ \rightarrow \bar{\nu}_\mu + \nu_e + e^+$	0.843
$\mu^- \rightarrow \nu_\mu + \bar{\nu}_e + e^-$	0.643
$\pi^+ \rightarrow \nu_\mu + \mu^+$	48.3
$\pi^- \rightarrow \bar{\nu}_\mu + \mu^-$	40.3

Table 5.1: Hadron decay processes and the percentage of neutrinos at the DUNE ND (on-axis) produced by each one.

major interaction type (QE, Res, MEC and DIS), we include a nuisance parameter on both the normalisation (a) and tilt (γ) of each cross section component, such that

$$\sigma_{\nu_\alpha} = \sum_{i=\text{QE, Res, DIS, MEC}} a_i \sigma_{\nu_\mu, i} \left(\frac{E}{E_0} \right)^{\gamma_i} \quad \alpha = e, \mu \quad (5.3.1)$$

where E_0 is some base energy around which the cross section distribution can be tilted. This attempts to capture discrepancies between the predicted cross sections as given by GENIE and NuWro. We also include detector response $R_\alpha(E_{\text{reco}}, E)$ and add backgrounds $B_{\alpha,ij}$ as defined in [105], giving the event rate $N_{\alpha,ij}$ in the i^{th} and j^{th} energy and angular bin, respectively,

$$N_{\alpha,ij} = \int dE_\nu \phi_{\text{ND},\alpha,ij} \sigma_\alpha(E_i) R_\alpha(E_{\text{reco},i}, E_\nu) + B_{\alpha,ij} \quad (5.3.2)$$

To study the impact of flux mis-modelling, we simulate off-axis spectra (at 7 different off-axis angles between 0° and 3.59°) based on the nominal fluxes from [119], but randomly vary the parent-normalisation factors r_p within a uniform distribution ranging between 0.9 and 1.1 - corresponding to the total fractional error shown in Figure 5.4. We then attempt to recover the true values of r_p by performing a frequentist maximum-likelihood fit to this dataset, allowing r_p to vary. The systematic uncertainties on the cross section normalisation and tilt, detailed in Table 5.2, which have been increased to include a 1.5% uncertainty on the detector response and efficiency, as per the prediction made in [100].

Parameter	Systematic uncertainty			
	QE	MEC	RES	DIS
a	7%	8%	8%	11%
γ	4%	10%	10%	3%

Table 5.2: Systematic uncertainties on cross-section normalisation (a) and tilt (γ) used in the LENS fit. The values used for each interaction type were derived by requiring that they covered the difference between the predicted cross sections as given by GENIE and NuWro. The normalisation cross section uncertainty has then been increased to include a 1.5% uncertainty on detector response and efficiency.

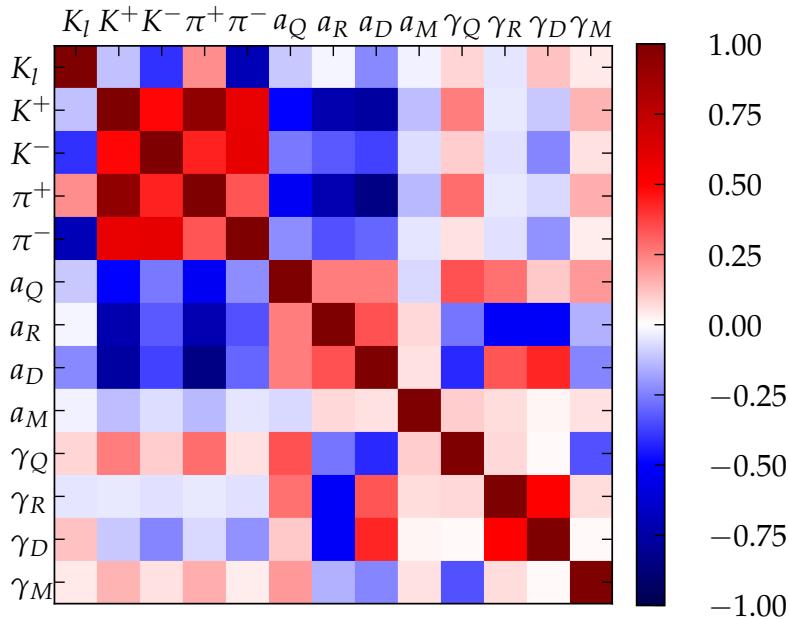


Figure 5.5: Correlation matrix for hadron production normalisation parameters extracted from the fit performed.

The likelihood function is then

$$-2 \log \mathcal{L}(\vec{r}) = \min_{\vec{a}} \left[\sum_{\substack{i=\text{energies} \\ j=\text{angles} \\ \alpha=\text{flavors}}} \left(\frac{N_{\alpha,ij}^{\text{obs}} - N_{\alpha,ij}^{\text{pred}}(\vec{r}; \vec{a})}{N_{\alpha,ij}^{\text{pred}}(\vec{r}; \vec{a})} \right)^2 + \sum_k \frac{a_k^2}{\sigma_k^2} \right], \quad (5.3.3)$$

where $N_{\alpha,ij}^{\text{obs}}$ is the ‘observed’ number of events in the i^{th} energy and j^{th} angular bin, and $N_{\alpha,ij}^{\text{pred}}$ is the corresponding number of predicted events for that bin. The function depends on the coefficients $\vec{r} = (r_p)$ and nuisance parameters \vec{a} , both of which are independent of energy, flavour and angle. From the fit, we extract a covariance matrix for the normalisation parameters, corresponding to parent particle contributions, from which their 1σ error bands, as shown in Figure 5.6, are derived.

5.4 Results

The results of the LENS fit are shown in Table 5.3 and Figure 5.6. We see that by using an ND only fit, PRISM can constrain most flux components to the few percent level, with the dominant π^+ and π^- components further constrained to below 2%. The K_L component is the least constrained, which is expected – K_L is neutral, so cannot be focused into the beamline by the magnetic horns, and decays semileptonically only via three-body processes. As a result it has a different angular dependence to the dominant two-body decay fluxes, and its contribution to the total flux is small, as shown in Figure 5.3. Nonetheless, the fit achieves a significant improvement over an a priori hadroproduction uncertainty of order 5%, which represents the uncertainty on a future flux model based on improved data from hadroproduction experiments and other external measurements, as well as theory improvements. We adopt this optimistic prediction on flux modelling improvements because our predictions on the improvements due to LENS are also optimistic, due to our relatively simple systematics model and our assumption of perfect particle and charge identification when generating the fluxes.

The equal time running strategy considered yields marginally better results than 50% on-axis. The FHC and RHC beam results are broadly similar, however we note that in FHC mode the π^+ normalisation is more strongly constrained than the π^- , while the reverse is true in RHC mode. This follows from the beam composition - in FHC mode π^+ are focused into the beamline while π^- are removed, and vice versa for RHC - meaning this difference is simply due to statistics. More interestingly, the constraints on the different flux components are more similar in RHC mode than in FHC. This is likely due to the neutrino cross section being approximately three times larger than the antineutrino cross section over the relevant energy range. As such, the measured RHC flux has a non-negligible contribution from neutrinos, allowing tighter constraints to be placed on r_{π^+} and r_{K^+} in the RHC fit than on the analogous parameters in the FHC fit.

The 1σ confidence intervals and regions for each normalisation parameter r_p are shown in Figure 5.6. Here we observe significant correlations between charged pion and kaon flux components, implying that a bias in the pion contribution is more easily masked if it is accompanied by a similar bias in the kaon contribution. This occurs because our fit is largely insensitive to the overall flux normalisation (as evidenced by the anticorrelation of the cross-section normalisations and hadron normalisations in

Parameter	ν -dominated beam		$\bar{\nu}$ -dominated beam	
	equal time	50% on-axis	equal time	50% on-axis
r_{K_L}	4.1%	4.5%	3.5%	4.1%
r_{K^+}	1.7%	1.7%	1.8%	1.9%
r_{K^-}	2.7%	3.2%	1.9%	2.0%
r_{π^+}	1.7%	1.6%	1.8%	1.8%
r_{π^-}	2.1%	1.9%	1.6%	1.6%
a_{QE}	2.3%	2.3%	2.4%	2.3%
a_{RES}	2.3%	2.3%	2.2%	2.2%
a_{DIS}	2.1%	2.0%	2.6%	2.6%
a_{MEC}	3.0%	3.0%	3.0%	3.0%
γ_{QE}	1.1%	1.1%	1.2%	1.2%
γ_{RES}	1.1%	1.1%	1.0%	1.0%
γ_{DIS}	0.6%	0.5%	0.7%	0.7%
γ_{MEC}	3.4%	3.4%	3.4%	3.4%

Table 5.3: Projected fractional uncertainties on the fit parameters r_p (the normalisation of the various flux components) and the nuisance parameters a_k and γ_k (parameterizing the cross section uncertainties) from a fit to 6.6×10^{21} pot of DUNE near detector data. Results are shown for a neutrino-dominated beam (“forward horn current”) and for an anti-neutrino-dominated beam (“reverse horn current”), and for two different running strategies (equal data-taking time at each of the seven detector positions, or 50% of the data-taking time spent in the on-axis positions and the rest equally distributed among six off-axis positions.)

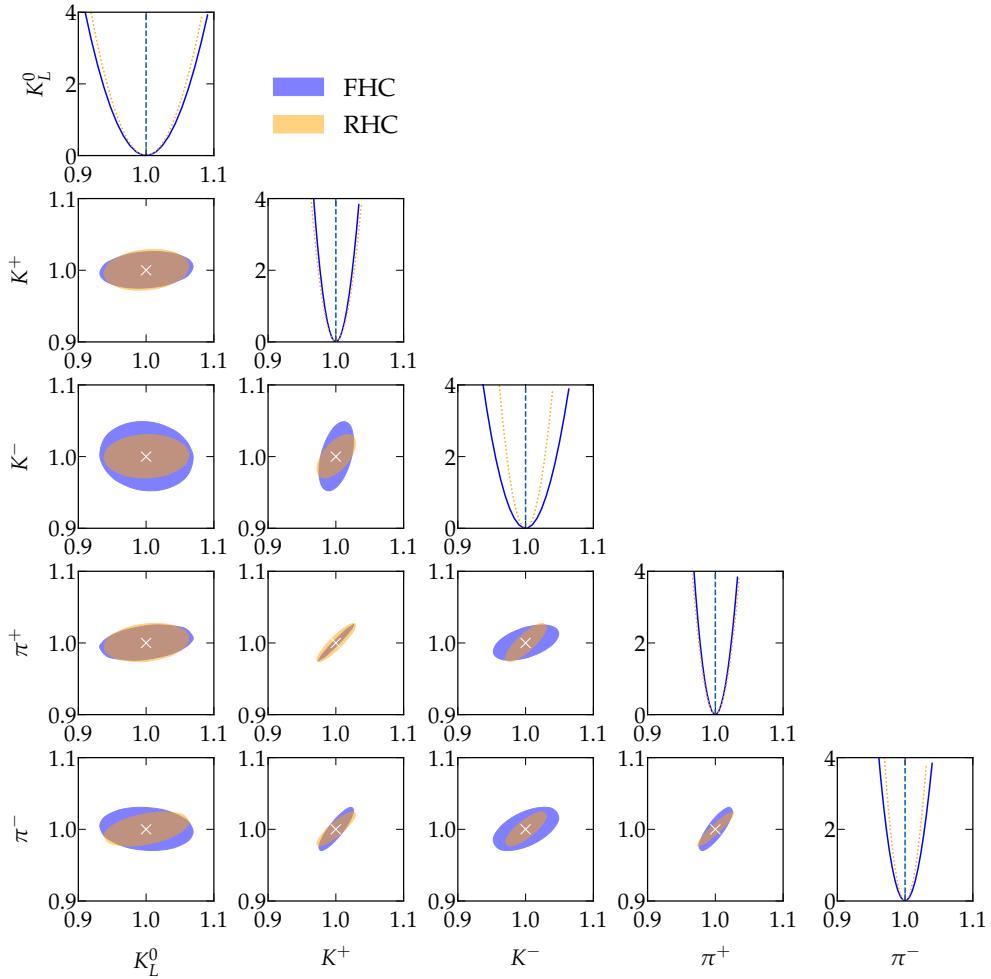


Figure 5.6: Projected constraints on the individual components of the DUNE neutrino flux from a fit to PRISM data in FHC mode (orange) and RHC mode (blue), assuming half of the total 6.6×10^{21} POT are collected in the on-axis position, while the rest are equally split over six different off-axis positions (6.3 m, 12.6 m, 18.9 m, 24.3 m, 30.6 m, 36.0 m, corresponding to off-axis angles out to 0.063 radians).

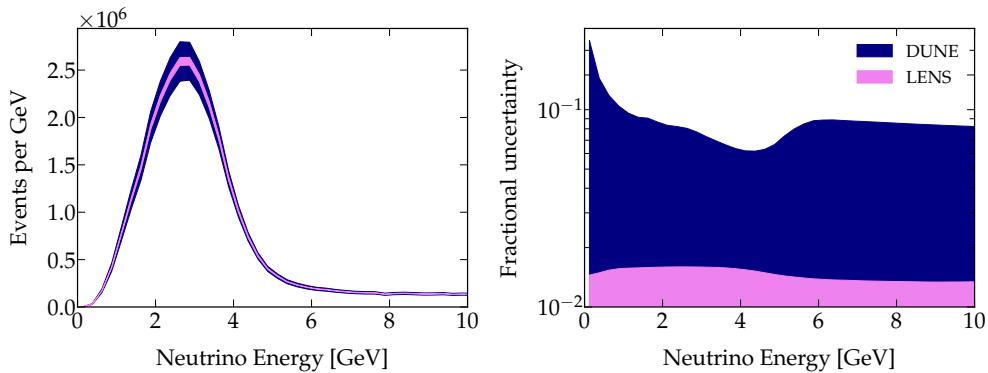


Figure 5.7: *Left*: Error bands for on-axis ND event distribution when using hadron production uncertainty shown in Figure 5.4 (navy) and when using uncertainties derived from LENS (magenta). *Right*: Comparison of fractional hadron production uncertainty from DUNE and LENS as a function of neutrino energy.

Figure 5.5) – this can be absorbed into the nuisance parameters associated with the neutrino cross sections. Changes in the neutrino spectrum however, achieved for example by biasing the pion contribution, but not the kaon, are easily detected by the fit. Small anticorrelations between the negatively charged mesons and K_L^0 are visible in the FHC fit due to the degeneracy in their flux distributions across off-axis positions, as shown in Figure 5.3.

Figure 5.7 compares the total fractional uncertainty on hadron production achieved with the LENS method to that shown in Figure 5.4. Note that we compare the ν_μ only uncertainty from DUNE to the uncertainty on all neutrino types from LENS, because hadron production uncertainties for the other neutrino flavours at DUNE are not publicly available. All uncertainties shown in Figure 5.4 do however apply to the ν_e flux, except the uncertainty on pion production from proton-carbon interactions, $PC \rightarrow \pi$. The LENS uncertainty is notably smaller and more consistent across all energies. The left-hand panel of Figure 5.7 also shows the 1σ uncertainty band on the total on-axis neutrino event distribution at DUNE’s ND for 6 years of exposure, assuming half the runtime is spent on-axis. Figure 5.8 shows the uncertainty band for each neutrino flavour in FHC mode. Unsurprisingly, due to its small contribution to the total FHC flux, the $\bar{\nu}_e$ error band is the largest, however stays within 5%. The results are analogous in RHC mode, where the ν_e flux is the least constrained.

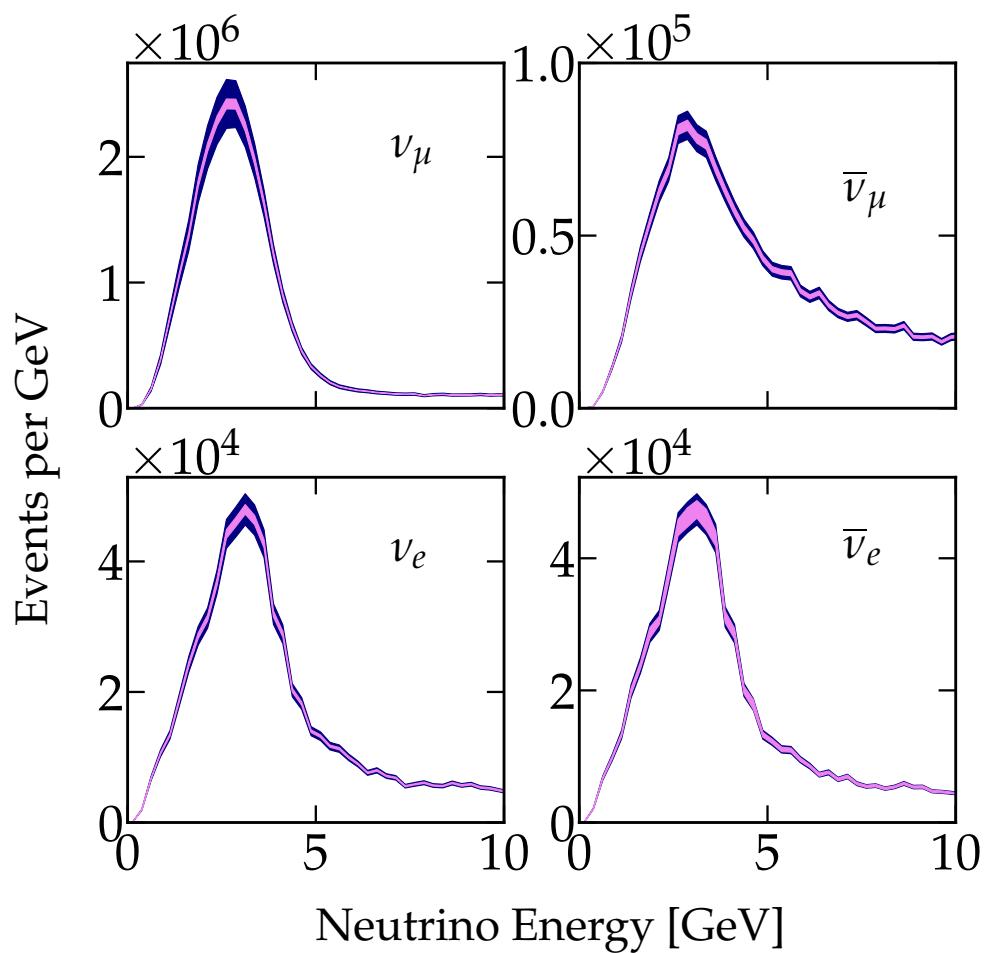


Figure 5.8: Error bands for each neutrino flavour's on-axis ND event distribution when using the hadron production uncertainty shown in Figure 5.4 for ν_μ , and a flat 5% uncertainty for $\bar{\nu}_\mu$, ν_e and $\bar{\nu}_e$.

5.5 Impact on Oscillation Measurements

We now examine how improved flux modelling with LENS affects DUNE’s oscillation sensitivities. Similar benefits could, of course, be achieved by improving neutrino flux modelling with, for example, improved theoretical modelling, monitored and tagged neutrino beams such as EnuBET [123] and NuTag [124], or hadroproduction experiments with replica targets, such as NA61/SHINE [125, 126].

We simulate the PRISM method in GLoBES [107], expanding on the GLoBES implementation of DUNE from [127], and incorporating code developed in [107, 128, 129]. For each set of oscillation parameters under consideration, we compute the PRISM coefficients c_j from simulations in which the meson flux normalisations r_p may vary around their nominal values within a given range. These c_j are then used to form an oscillated FD flux prediction based on simulated ND ‘data’ generated using the nominal flux model QGSP_BERT from [119]. Finally, the FD prediction is compared to simulated FD ‘data’, also generated using the nominal flux model.

Figure 5.9 shows the impact of the LENS procedure on the FD event spectrum predicted by PRISM using nominal near detector data. Each curve represents the results of generating the coefficients c_j from a different flux model, i.e. a flux made with different values of r_p . Blue curves are generated by sampling r_p from a multivariate Gaussian parameterised by the LENS covariance matrix, while the pink curves sample from a multivariate Gaussian allowing uncertainties of up to 5%. The black curve shows the FD spectrum obtained using the nominal model with $r_p = 1$ for all flux components to determine both c_j and the FD prediction, and FD ‘data’.

Figure 5.10 shows DUNE’s projected sensitivity to leptonic CP violation as a function of the true CP phase, δ_{CP}^{true} . For each value of δ_{CP}^{true} , Figure 5.10 we show the significance with which this δ_{CP}^{true} can be distinguished from the case of CP conservation; $\delta_{CP}^{true} \in \{0, \pi\}$. We have marginalised over all other oscillation parameters, and applied an external prior of 5% on the solar parameters Δm_{21}^2 and θ_{21} . As before, blue and pink curves correspond to LENS-constrained and 5% constrained flux models, and the solid black curve the expected result from using the same, unvaried flux, for both FD predictions and ‘data’. From these results it is apparent that constraining the hadroproduction model using LENS prior to using PRISM to predict FD fluxes can reduce the un-

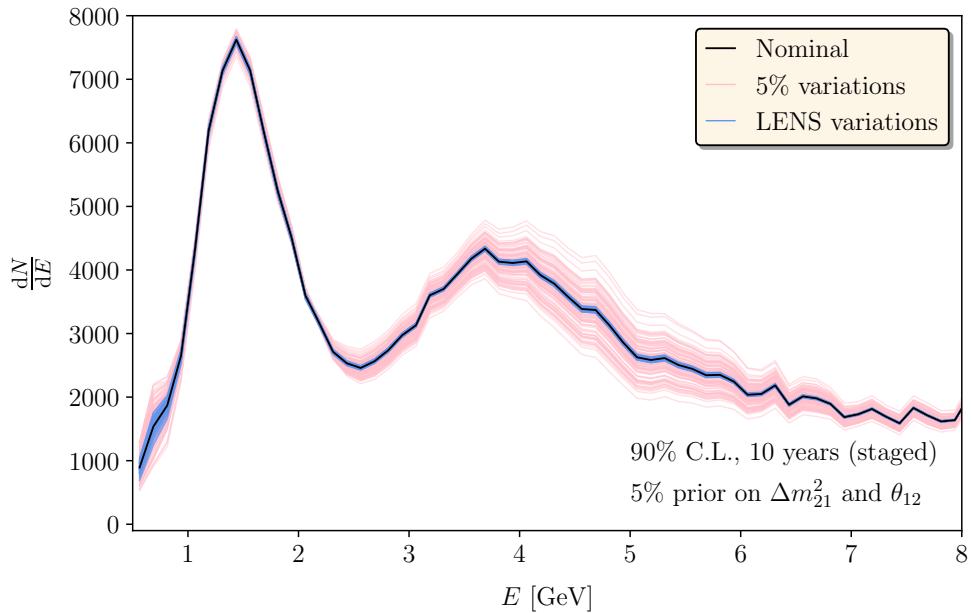


Figure 5.9: Impact of the flux uncertainty on the DUNE far detector event spectrum with and without using the LENS procedure to constrain the flux model. All curves were obtained using PRISM to predict the oscillated far detector fluxes based on near detector data. While the near and far detector data were predicted using the nominal flux model, the PRISM coefficients c_j were derived from biased flux models in which the relative importance of individual flux components r_p were allowed to vary either by 5% (red curves) or within the tighter constraints imposed by our LENS fit to on-axis and off-axis near detector data, see Table 5.3. In each case, we show 100 random realizations. The solid black curve shows for comparison the spectrum based on the nominal flux model with no variation in hadron-contribution normalisation. The spread in possible outcomes is substantially reduced when the LENS constraint is included.

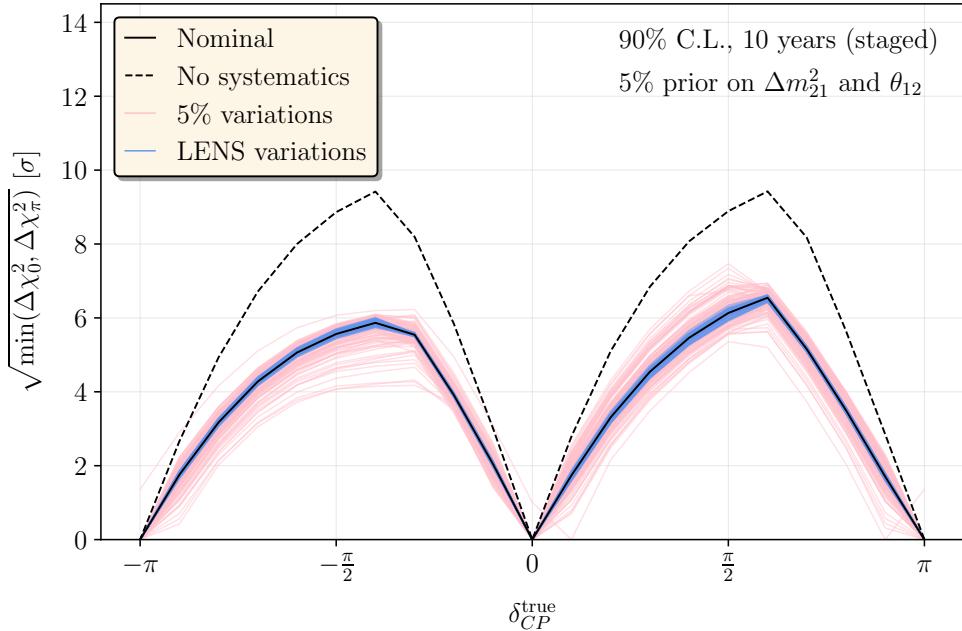


Figure 5.10: Impact of LENS on DUNE's sensitivity to leptonic CP violation. For each value of the true δ_{CP} , we show the significance at which the CP-conserving hypothesis ($\delta_{CP} \in \{0, \pi\}$) can be excluded. All curves were obtained using PRISM to predict the oscillated far detector fluxes based on near detector data. While the near and far detector data were predicted using the nominal flux model, the PRISM coefficients c_j are derived from biased flux models in which the relative importance of individual flux components r_p is allowed to vary within the constraints set by LENS (blue contours) or within 5% of the nominal value (no LENS constraint, red). We show 100 random realizations of biased flux models for each case. We see that the spread in possible outcomes is substantially reduced when the LENS constraint is included. The solid black curve shows for comparison the spectrum based on the nominal flux model, with no variation in hadron-contribution normalisation, and the dashed black curve shows the same result, but with no systematic uncertainties.

certainty on DUNE’s results. Both Figures 5.9 and 5.10 show that in the case of a mismodelled flux, applying the LENS procedure prior to determining c_j means they are based on fluxes which lie closer to the model realised in nature. As shown by the pink curves in Figure 5.10, the sensitivity to CP violation can otherwise be over or underestimated by up to 1σ .

Further evidence of the impact of constraining hadroproduction models is shown in Figure 5.11, which presents DUNE’s anticipated sensitivity contours in $\theta_{13} - \delta_{CP}$, and $\theta_{23} - \Delta m_{31}^2$ planes. Again blue and pink correspond to FD predictions made using LENS and 5% hadroproduction uncertainties, respectively, and stars of the same colours indicate the best fit points for each contour shown. The same marginalisation and priors are used here as in Figure 5.10. In the $\sin^2 2\theta_{13} - \delta_{CP}$ plane, the LENS-constrained hadroproduction uncertainties ensure the best fit points are clustered within the nominal flux contours (black). In the $\Delta m_{31}^2 - \theta_{23}$ plane the LENS variation contours give tighter 90% confidence limits and a smaller spread of best fit points. The results can therefore be considered more reliable and robust.

Of course, improving the uncertainty on the neutrino flux with LENS should yield gains in sensitivity by reducing the systematic uncertainties used in the oscillation fit, as opposed to just making the results more reliable. However, this improvement cannot be shown without disentangling the flux uncertainty from cross-section and detector uncertainties in the systematics model employed in the DUNE simulation [105], which goes beyond the scope of this work. Additionally, as discussed in Section 5.4, the improvements on the flux uncertainty by LENS are optimistic. As such, we instead show the results of a fit with no systematic uncertainties (dashed black line) in Figures 5.10 and 5.11. We expect the sensitivity improvements from LENS to yield results falling somewhere between the no-systematics and current systematics case (solid black line).

This work highlights how long-baseline oscillation experiments can optimally benefit from their near detectors, and implies that highly capable detectors are crucial for their success. We have also shown that, in the systematics dominated regime in which DUNE and HyperKamiokande will operate, even modest improvements in near detector performance can lead to substantial gains in oscillation sensitivities.

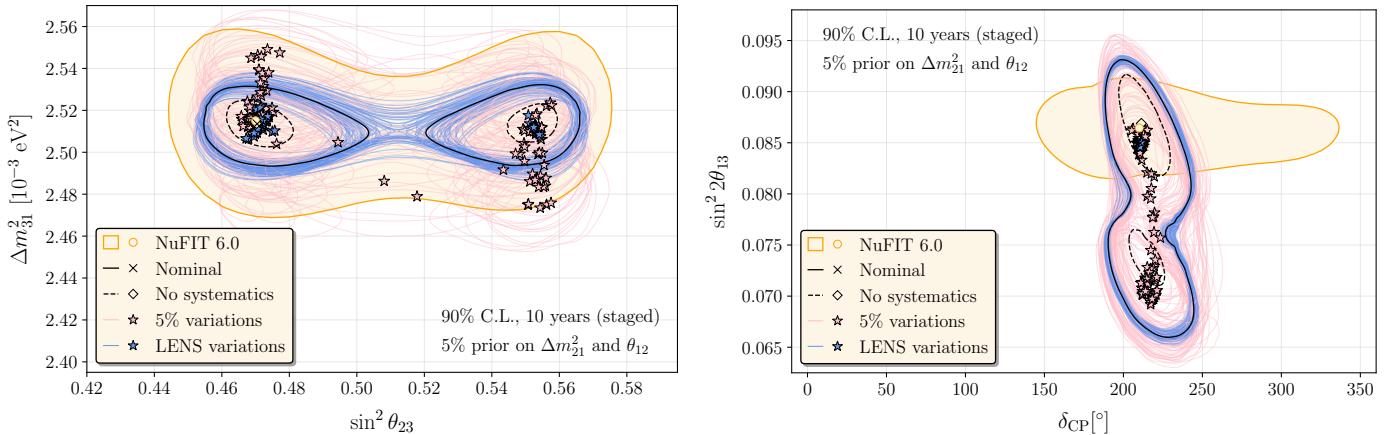


Figure 5.11: Impact of LENS on the measurement of the oscillation parameters $(\theta_{13}, \delta_{CP})$ (left) and $(\theta_{23}, \Delta m_{31}^2)$ (right). As in Fig. 5.10, red and blue contours correspond to neutrino flux models with and without the LENS constraint. In all cases, the far detector prediction is based on the PRISM procedure, with the superposition coefficients c_j determined from the biased flux models and the data generated using the nominal model. Solid black contours show the anticipated sensitivity using the nominal flux model with all $r_p = 1$, utilising DUNE's simulated systematic uncertainties as per [105]. Dashed black contours show results from the same procedure, but without systematic uncertainties. For comparison, we also show the current constraints on the oscillation parameters from NuFit 6.0 [130] in beige with an orange outline.

Chapter 6

Summary and Conclusions

The accurate measurement of neutrino oscillation parameters will answer important outstanding questions in particle physics. The work presented in this thesis has served to maximise the physics potential of both current and next generation experiments from an experimental and phenomenological standpoint.

In Chapter 3, we discussed the most recent NOvA analysis, focussing in particular on the cut optimisation performed by the author to enhance the signal to background ratio of the ν_μ CC dataset, which provides sensitivity to atmospheric mixing parameters. Work was also presented on constructing posterior ranges on neutrino event spectra for the Bayesian analysis, which allow the construction of credible intervals for neutrino oscillation parameters. This analysis achieved the most precise single experiment measurement of the size of the atmospheric mass splitting, $|\Delta m_{32}^2| = 2.431_{-0.039}^{+0.036} \times 10^{-3} \text{ eV}^2$. This corresponds to a 1.7% precision on the absolute mass *difference* between ν_2 and ν_3 , whose individual masses are only loosely constrained. Despite the degeneracy NOvA sees in the δ_{CP} - mass ordering phase space, the latest dataset was also able to disfavour the scenario of minimal charge-parity violation in the inverted mass ordering.

DUNE, a next generation long-baseline oscillation experiment, promises to determine the neutrino mass ordering, and measure the amount of charge-parity violation in the leptonic sector of the Standard Model. While its Liquid Argon Time Projection Chamber detectors promise outstanding reconstruction capabilities, their performance will be hampered by our lack of knowledge on neutrino-nucleus interactions. Due to the many sub-processes which occur when neutrinos interact with nuclei, the final state of a neutrino interaction may confer misleading information about the incident neutrino, making the reconstruction of its true energy

a challenge. This is problematic, as neutrino oscillation experiments rely on only two measurements - the number of neutrinos of a given flavour, and their energy. In Chapter 4, by employing a simple neural network, we showed how the neutrino energy resolution could be improved after reconstruction, and quantified the corresponding gains in DUNE’s sensitivity to charge-parity violation and θ_{23} octant. This work also illustrated the necessity of better understanding neutrino-nucleus interactions at lower energies by showing that a network trained on one neutrino event generator but tested on another underperformed at energies < 2.5 GeV due to differences in the models used by each generator.

The PRISM method will be employed by DUNE in order to circumvent our lack of knowledge about neutrino cross-sections. By using near detector data at multiple off-axis locations to form far detector predictions, PRISM will hugely constrain systematic uncertainties. However, mis-modelling of hadron production in the DUNE PIP-II beam, introducing errors in the relative importance of parent hadrons at different off-axis locations in the predicted flux, could lead to biased oscillation parameter measurements. In Chapter 5 we outlined LENS, a method wherein an all-axis ND only fit can be used alongside PRISM to further constrain hadron production uncertainties. We showed that the LENS method can increase DUNE’s robustness to hadron production mis-modelling, ensuring more reliable oscillation measurements, and may also improve its sensitivity to oscillation parameters.

Hopefully this thesis has convinced the reader both of how exciting the field of neutrino oscillation physics is, and how crucial it is that we keep pushing to further understand these particles and their interactions in order to maximise the potential of experiments in the next decade.

Bibliography

- [1] W. Pauli. “Dear radioactive ladies and gentlemen”. *Phys. Today* 31N9 (1978), p. 27.
- [2] J. Chadwick. “The intensity distribution in the magnetic spectrum of β particles from radium (B + C)”. *Verh. Phys. Gesell.* 16 (1914), pp. 383–391.
- [3] H. Bethe and R. Peierls. “The ‘neutrino’”. *Nature* 133 (1934), p. 532. DOI: [10.1038/133532a0](https://doi.org/10.1038/133532a0).
- [4] F. Reines and C. L. Cowan. “Detection of the Free Neutrino”. *Phys. Rev.* 92 (3 1953), pp. 830–831. DOI: [10.1103/PhysRev.92.830](https://doi.org/10.1103/PhysRev.92.830).
- [5] T. D. Lee and Chen-Ning Yang. “Question of Parity Conservation in Weak Interactions”. *Phys. Rev.* 104 (1956), pp. 254–258. DOI: [10.1103/PhysRev.104.254](https://doi.org/10.1103/PhysRev.104.254).
- [6] C. S. Wu et al. “Experimental Test of Parity Conservation in β Decay”. *Phys. Rev.* 105 (1957), pp. 1413–1414. DOI: [10.1103/PhysRev.105.1413](https://doi.org/10.1103/PhysRev.105.1413).
- [7] M. Goldhaber, L. Grodzins, and A. W. Sunyar. “Helicity of Neutrinos”. *Phys. Rev.* 109 (1958), pp. 1015–1017. DOI: [10.1103/PhysRev.109.1015](https://doi.org/10.1103/PhysRev.109.1015).
- [8] B. Pontecorvo. “Mesonium and anti-mesonium”. *Sov. Phys. JETP* 6 (1957), p. 429.
- [9] Gustavo Castelo Branco, Luís Lavoura, and João Paulo Silva. “The Neutral-Kaon System”. *CP Violation*. Oxford University Press, July 1999. ISBN: 9780198503996. DOI: [10.1093/oso/9780198503996.003.0008](https://doi.org/10.1093/oso/9780198503996.003.0008).
- [10] G. Danby et al. “Observation of High-Energy Neutrino Reactions and the Existence of Two Kinds of Neutrinos”. *Phys. Rev. Lett.* 9 (1 1962), pp. 36–44. DOI: [10.1103/PhysRevLett.9.36](https://doi.org/10.1103/PhysRevLett.9.36).

- [11] Ziro Maki, Masami Nakagawa, and Shoichi Sakata. “Remarks on the unified model of elementary particles”. *Prog. Theor. Phys.* 28 (1962), pp. 870–880. DOI: 10.1143/PTP.28.870.
- [12] V. N. Gribov and B. Pontecorvo. “Neutrino astronomy and lepton charge”. *Phys. Lett. B* 28 (1969), p. 493. DOI: 10.1016/0370-2693(69)90525-5.
- [13] Raymond Davis Jr., Don S. Harmer, and Kenneth C. Hoffman. “Search for neutrinos from the sun”. *Phys. Rev. Lett.* 20 (1968), pp. 1205–1209. DOI: 10.1103/PhysRevLett.20.1205.
- [14] K. Lande et al. “Results from the Homestake solar neutrino observatory”. *Conf. Proc. C* 900802 (1990). Ed. by K. K. Phua and Y. Yamaguchi, pp. 867–675.
- [15] D. Decamp et al. “Determination of the Number of Light Neutrino Species”. *Phys. Lett. B* 231 (1989), pp. 519–529. DOI: 10.1016/0370-2693(89)90704-1.
- [16] K. Kodama et al. “Observation of tau neutrino interactions”. *Physics Letters B* 504.3 (Apr. 2001), pp. 218224. DOI: 10.1016/s0370-2693(01)00307-0.
- [17] K. Kodama et al. “Final tau-neutrino results from the DONuT experiment”. *Phys. Rev. D* 78 (2008), p. 052002. DOI: 10.1103/PhysRevD.78.052002. arXiv: 0711.0728 [hep-ex].
- [18] N. Agafonova et al. “Discovery of τ Neutrino Appearance in the CNGS Neutrino Beam with the OPERA Experiment”. *Physical Review Letters* 115.12 (Sept. 2015). DOI: 10.1103/physrevlett.115.121802.
- [19] R. Abbasi et al. “Observation of Seven Astrophysical Tau Neutrino Candidates with IceCube”. *Physical Review Letters* 132.15 (Apr. 2024). DOI: 10.1103/physrevlett.132.151001.
- [20] *Standard Model of Elementary Particles*. https://en.wikipedia.org/wiki/File:Standard_Model_of_Elementary_Particles.svg. Accessed: 07-12-2025.
- [21] Mary K. Gaillard, Paul D. Grannis, and Frank J. Sciulli. “The standard model of particle physics”. *Reviews of Modern Physics* 71.2 (Mar. 1999), S96S111. DOI: 10.1103/revmodphys.71.s96.
- [22] Steven Weinberg. “A Model of Leptons”. *Phys. Rev. Lett.* 19 (1967), pp. 1264–1266. DOI: 10.1103/PhysRevLett.19.1264.

- [23] Paul A. M. Dirac. “The quantum theory of the electron”. *Proc. Roy. Soc. Lond. A* 117 (1928), pp. 610–624. DOI: [10.1098/rspa.1928.0023](https://doi.org/10.1098/rspa.1928.0023).
- [24] Ettore Majorana. “Teoria simmetrica dell’elettrone e del positrone”. *Nuovo Cim.* 14 (1937), pp. 171–184. DOI: [10.1007/BF02961314](https://doi.org/10.1007/BF02961314).
- [25] Steven Weinberg. “Phenomenological Lagrangians”. *Physica A* 96.1–2 (1979). Ed. by S. Deser, pp. 327–340. DOI: [10.1016/0378-4371\(79\)90223-1](https://doi.org/10.1016/0378-4371(79)90223-1).
- [26] B. Pontecorvo. “Inverse Beta Processes and Nonconservation of Lepton Charge”. *Sov. Phys. JETP* 7 (1958), pp. 172–173.
- [27] O. V. Lychkovskiy. “Neutrino oscillations: Deriving the plane-wave approximation in the wave-packet approach”. *Physics of Atomic Nuclei* 72.9 (Sept. 2009), pp. 15571559. DOI: [10.1134/s1063778809090178](https://doi.org/10.1134/s1063778809090178).
- [28] Carlo Giunti. “Neutrino Wave Packets in Quantum Field Theory”. *Journal of High Energy Physics* 2002.11 (Nov. 2002), pp. 017017. DOI: [10.1088/1126-6708/2002/11/017](https://doi.org/10.1088/1126-6708/2002/11/017).
- [29] C. Jarlskog. “Commutator of the Quark Mass Matrices in the Standard Electroweak Model and a Measure of Maximal CP Nonconservation”. *Phys. Rev. Lett.* 55 (1985), p. 1039. DOI: [10.1103/PhysRevLett.55.1039](https://doi.org/10.1103/PhysRevLett.55.1039).
- [30] Peter B. Denton and Stephen J. Parke. “Simple and precise factorization of the Jarlskog invariant for neutrino oscillations in matter”. *Physical Review D* 100.5 (Sept. 2019). DOI: [10.1103/physrevd.100.053004](https://doi.org/10.1103/physrevd.100.053004).
- [31] R. Aaij et al. “First Observation of CP Violation in the Decays B_S^0 Mesons”. *Physical Review Letters* 110.22 (May 2013). DOI: [10.1103/physrevlett.110.221601](https://doi.org/10.1103/physrevlett.110.221601).
- [32] R. Aaij et al. “Observation of CP Violation in Charm Decays”. *Phys. Rev. Lett.* 122 (21 2019), p. 211803. DOI: [10.1103/PhysRevLett.122.211803](https://doi.org/10.1103/PhysRevLett.122.211803).
- [33] Joachim Kopp. “Solar and Reactor Neutrino Theory”. *International Neutrino Summer School*. 2019.
- [34] S. P. Mikheyev and A. Yu. Smirnov. “Resonance Amplification of Oscillations in Matter and Spectroscopy of Solar Neutrinos”. *Sov. J. Nucl. Phys.* 42 (1985), pp. 913–917.
- [35] L. Wolfenstein. “Neutrino Oscillations and Stellar Collapse”. *Phys. Rev. D* 20 (1979), pp. 2634–2635. DOI: [10.1103/PhysRevD.20.2634](https://doi.org/10.1103/PhysRevD.20.2634).

- [36] Christoph Andreas Ternes. “Neutrino oscillation phenomenology in the standard model and beyond”. PhD thesis. U. Valencia (main), 2020. doi: [10.2172/1767027](https://doi.org/10.2172/1767027).
- [37] Vernon D. Barger et al. “Matter Effects on Three-Neutrino Oscillations”. *Phys. Rev. D* 22 (1980), p. 2718. doi: [10.1103/PhysRevD.22.2718](https://doi.org/10.1103/PhysRevD.22.2718).
- [38] Evgeny K Akhmedov et al. “Series expansions for three-flavor neutrino oscillation probabilities in matter”. *Journal of High Energy Physics* 2004.04 (May 2004), pp. 078078. doi: [10.1088/1126-6708/2004/04/078](https://doi.org/10.1088/1126-6708/2004/04/078).
- [39] Ivan Esteban et al. “NuFit-6.0: updated global analysis of three-flavor neutrino oscillations”. *Journal of High Energy Physics* 2024.12 (Dec. 2024). doi: [10.1007/jhep12\(2024\)216](https://doi.org/10.1007/jhep12(2024)216).
- [40] Y. Abe et al. “Indication of Reactor $\bar{\nu}_e$ Disappearance in the Double Chooz Experiment”. *Physical Review Letters* 108.13 (Mar. 2012). doi: [10.1103/physrevlett.108.131801](https://doi.org/10.1103/physrevlett.108.131801).
- [41] F. P. An et al. “Observation of electron-antineutrino disappearance at Daya Bay”. *Phys. Rev. Lett.* 108 (2012), p. 171803. doi: [10.1103/PhysRevLett.108.171803](https://doi.org/10.1103/PhysRevLett.108.171803). arXiv: [1203.1669 \[hep-ex\]](https://arxiv.org/abs/1203.1669).
- [42] J. K. Ahn et al. “Observation of Reactor Electron Antineutrinos Disappearance in the RENO Experiment”. *Physical Review Letters* 108.19 (May 2012). doi: [10.1103/physrevlett.108.191802](https://doi.org/10.1103/physrevlett.108.191802).
- [43] Mariia Redchuk. “Comprehensive measurement of pp-chain solar neutrinos with Borexino”. *PoS EPS-HEP2019* (2020), p. 400. doi: [10.22323/1.364.0400](https://doi.org/10.22323/1.364.0400).
- [44] B. KAYSER. “Neutrino Mass, Mixing and Oscillation”. *Flavor Physics for the Millennium*. World Scientific, Sept. 2001, pp. 625650. doi: [10.1142/9789812811509_0017](https://doi.org/10.1142/9789812811509_0017).
- [45] K. Abe et al. “Solar neutrino measurements using the full data period of Super-Kamiokande-IV”. *Physical Review D* 109.9 (May 2024). doi: [10.1103/physrevd.109.092001](https://doi.org/10.1103/physrevd.109.092001).
- [46] A. Bellerive et al. “The Sudbury Neutrino Observatory”. *Nuclear Physics B* 908 (July 2016), pp. 3051. doi: [10.1016/j.nuclphysb.2016.04.035](https://doi.org/10.1016/j.nuclphysb.2016.04.035).

- [47] D. Basilico et al. “Recent results from Borexino on solar neutrinos”. *EPJ Web Conf.* 290 (2023), p. 04001. DOI: 10.1051/epjconf/202329004001.
- [48] Victor F. Hess. “Über Beobachtungen der durchdringenden Strahlung bei sieben Freiballonfahrten”. *Phys. Z.* 13 (1912), pp. 1084–1091.
- [49] Takaaki Kajita. “Atmospheric Neutrinos and the Discovery of Neutrino Oscillations”. *Proc. Japan Acad. B* 86 (2010), pp. 303–321. DOI: 10.2183/pjab.86.303.
- [50] T. Wester et al. “Atmospheric neutrino oscillation analysis with neutron tagging and an expanded fiducial volume in Super-Kamiokande I-V” (Nov. 2023). arXiv: 2311.05105 [hep-ex].
- [51] R. Abbasi et al. “Measurement of atmospheric neutrino mixing with improved IceCube DeepCore calibration and data processing”. *Physical Review D* 108.1 (July 2023). DOI: 10.1103/physrevd.108.012014.
- [52] K. Abe et al. “Measurements of neutrino oscillation parameters from the T2K experiment using 3.6×10^{21} protons on target”. *Eur. Phys. J. C* 83.9 (2023), p. 782. DOI: 10.1140/epjc/s10052-023-11819-x. arXiv: 2303.03222 [hep-ex].
- [53] M. A. Acero et al. “Improved measurement of neutrino oscillation parameters by the NOvA experiment”. *Phys. Rev. D* 106.3 (2022), p. 032004. DOI: 10.1103/PhysRevD.106.032004. arXiv: 2108.08219 [hep-ex].
- [54] A. E. Cárcamo Hernández et al. “Predictive Pati-Salam theory of fermion masses and mixing”. *Journal of High Energy Physics* 2017.7 (July 2017). DOI: 10.1007/jhep07(2017)118.
- [55] K. Abe et al. “Hyper-Kamiokande Design Report” (May 2018). arXiv: 1805.04163 [physics.ins-det].
- [56] R. Acciarri et al. “Long-Baseline Neutrino Facility (LBNF) and Deep Underground Neutrino Experiment (DUNE): Conceptual Design Report, Volume 2: The Physics Program for DUNE at LBNF” (Dec. 2015). arXiv: 1512.06148 [physics.ins-det].
- [57] P. F. De Salas et al. “Neutrino Mass Ordering from Oscillations and Beyond: 2018 Status and Future Prospects”. *Front. Astron. Space Sci.* 5 (2018), p. 36. DOI: 10.3389/fspas.2018.00036. arXiv: 1806.11051 [hep-ph].

- [58] P. Adamson et al. “The NuMI neutrino beam”. *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment* 806 (Jan. 2016), pp. 279306. ISSN: 0168-9002. DOI: 10.1016/j.nima.2015.08.063.
- [59] D. S. Ayres et al. “The NOvA Technical Design Report” (Oct. 2007). DOI: 10.2172/935497.
- [60] Erika Catano-Mur. *Recent results from NOvA*. 2022. arXiv: 2206.03542 [hep-ex].
- [61] M. Baird et al. “Event Reconstruction Techniques in NOvA”. *J. Phys. Conf. Ser.* 664.7 (2015), p. 072035. DOI: 10.1088/1742-6596/664/7/072035.
- [62] F Psihas and R Radovic. “Assorted CVN plots for blessing”. *NOvA Internal Document DocDB 15639* (2017).
- [63] A. Aurisano et al. “A Convolutional Neural Network Neutrino Event Classifier”. *JINST* 11.09 (2016), P09001. DOI: 10.1088/1748-0221/11/09/P09001. arXiv: 1604.01444 [hep-ex].
- [64] R Nichol and L Koerner. “Calibration Tech Note 2020 Analysis”. *NOvA Internal Document DocDB 43820* (2020).
- [65] A Radovic and P Singh. “The Attenuation and Threshold Calibration of the NOvA detectors”. *NOvA Internal Document DocDB 13579* (2016).
- [66] P Singh. “Attenuation Calibration of the NOvA Detectors”. *NOvA Internal Document DocDB 15472* (2016).
- [67] H. A. Bethe. “Molière’s Theory of Multiple Scattering”. *Phys. Rev.* 89 (6 1953), pp. 1256–1266. DOI: 10.1103/PhysRev.89.1256.
- [68] A Lister and L Vinton. “Far detector stopping muon energy scale calibration”. *NOvA Internal Document DocDB 13475* (2023).
- [69] W. Mu. “NuMu Energy Spline Plots 2020”. *NOvA Internal Document DocDB 46283* (2020).
- [70] S. Agostinelli et al. “GEANT4 - A Simulation Toolkit”. *Nucl. Instrum. Meth. A* 506 (2003), pp. 250–303. DOI: 10.1016/S0168-9002(03)01368-8.
- [71] Heather Ray. “MINERvA”. *Acta Phys. Polon. B* 40 (2009). Ed. by Arthur Ankowski and Jan Sobczyk, pp. 2647–2652.
- [72] S. Afanasiev et al. “The NA49 large acceptance hadron detector”. *Nucl. Instrum. Meth. A* 430 (1999), pp. 210–244. DOI: 10.1016/S0168-9002(99)00239-9.

- [73] C. Andreopoulos et al. “The GENIE Neutrino Monte Carlo Generator”. *Nucl. Instrum. Meth. A* 614 (2010), pp. 87–104. DOI: 10.1016/j.nima.2009.12.009. arXiv: 0905.2517 [hep-ph].
- [74] A Himmel. “New Oscillation Measurements from NOvA (CERN Seminar)”. *NOvA Internal Document DocDB 26151* (2018).
- [75] A Booth and O Chow. “Blessing Package - 3 Flavour ND Data/MC Comparisons Ana2024”. *NOvA Internal Document DocDB 62843* (2024).
- [76] B Choudhary et al. “Ana2024 FD Data-MC Blessing Package”. *NOvA Internal Document DocDB 62762* (2024).
- [77] Z. Kohley et al. “Modeling interactions of intermediate-energy neutrons in a plastic scintillator array with Geant4”. *Nucl. Instrum. Meth. A* 682 (2012), pp. 59–65. DOI: 10.1016/j.nima.2012.04.060.
- [78] Erika Catano-Mur, Ryan Nichol, and Zoya Vallari. “Three-Flavour Executive Summary 2024”. *NOvA Internal Document DocDB 61655* (2024).
- [79] M. A. Acero et al. *Monte Carlo method for constructing confidence intervals with unconstrained and constrained nuisance parameters in the NOvA experiment*. 2025. DOI: <https://doi.org/10.1088/1748-0221/20/02/T02001>. arXiv: 2207.14353 [hep-ex].
- [80] P. S. Laplace. *A Philosophical Essay on Probabilities*. French. Trans. by F. W. Truscott and F. L. Emory. 6th. New York: Dover Publication, 1951.
- [81] Thomas Bayes. “An essay towards solving a problem in the doctrine of chances.” *Phil. Trans. R. Soc.* 53 (1763), pp. 370–418.
- [82] A. Back et al. “Introduction to Bayesian MCMC fitting”. *NOvA Internal Document DocDB 53021* (2021).
- [83] S. Duane et al. “Hybrid Monte Carlo”. *Phys. Lett. B* 195 (1987), pp. 216–222. DOI: 10.1016/0370-2693(87)91197-X.
- [84] Stan Development Team. *Stan Users Guide, 2.37*. 2025.
- [85] F. P. An et al. “Precision Measurement of Reactor Antineutrino Oscillation at Kilometer-Scale Baselines by Daya Bay”. *Phys. Rev. Lett.* 130.16 (2023), p. 161802. DOI: 10.1103/PhysRevLett.130.161802. arXiv: 2211.14988 [hep-ex].

- [86] Hiroshi Nunokawa, Stephen Parke, and Renata Zukanovich Funchal. “Another possible way to determine the neutrino mass hierarchy”. *Physical Review D* 72.1 (July 2005). DOI: 10.1103/physrevd.72.013009.
- [87] The NOvA Collaboration. *Precision measurement of neutrino oscillation parameters with 10 years of data from the NOvA experiment*. 2025. arXiv: 2509.04361 [hep-ex].
- [88] L Jargowsky B. Kolupaeva and A. Yankelevich. “Blessing Package For Ana2024 3F Bayesian Results Comparisons with Other Experiments”. *NOvA Internal Document DocDB 62924* (2024).
- [89] Fernanda Psihas et al. “A review on machine learning for neutrino experiments”. *International Journal of Modern Physics A* 35.33 (2020), p. 2043005. DOI: 10.1142/S0217751X20430058.
- [90] P. Abratenko et al. “Convolutional neural network for multiple particle identification in the MicroBooNE liquid argon time projection chamber”. *Physical Review D* 103.9 (May 2021). DOI: 10.1103/physrevd.103.092003.
- [91] IceCube Collaboration. *Measurement of atmospheric neutrino oscillation parameters using convolutional neural networks with 9.3 years of data in IceCube DeepCore*. 2024. arXiv: 2405.02163 [hep-ex].
- [92] SBND Collaboration. *Cosmic Background Removal with Deep Neural Networks in SBND*. 2021. arXiv: 2012.01301 [physics.data-an].
- [93] Joachim Kopp et al. *Improving Neutrino Energy Reconstruction with Machine Learning*. 2025. arXiv: 2405.15867 [hep-ph].
- [94] B. Abi et al. “Long-baseline neutrino oscillation physics potential of the DUNE experiment: DUNE Collaboration”. *The European Physical Journal C* 80.10 (Oct. 2020). ISSN: 1434-6052. DOI: 10.1140/epjc/s10052-020-08456-z.
- [95] Kevin James Kelly et al. “Sub-GeV Atmospheric Neutrinos and CP-Violation in DUNE”. *Phys. Rev. Lett.* 123.8 (2019), p. 081801. DOI: 10.1103/PhysRevLett.123.081801. arXiv: 1904.02751 [hep-ph].
- [96] Kevin J. Kelly et al. “DUNE atmospheric neutrinos: Earth tomography”. *JHEP* 05 (2022), p. 187. DOI: 10.1007/JHEP05(2022)187. arXiv: 2110.00003 [hep-ph].

- [97] Peter B. Denton and Rebekah Pestes. “Neutrino oscillations through the Earth’s core”. *Phys. Rev. D* 104.11 (2021), p. 113007. DOI: 10.1103/PhysRevD.104.113007. arXiv: 2110.01148 [hep-ph].
- [98] R. Acciarri et al. “A Proposal for a Three Detector Short-Baseline Neutrino Oscillation Program in the Fermilab Booster Neutrino Beam” (Mar. 2015). arXiv: 1503.01520 [physics.ins-det].
- [99] T. Golan, J. T. Sobczyk, and J. Zmuda. “NuWro: the Wroclaw Monte Carlo Generator of Neutrino Interactions”. *Nucl. Phys. B Proc. Suppl.* 229-232 (2012). Ed. by George S. Tzanakos, pp. 499–499. DOI: 10.1016/j.nuclphysbps.2012.09.136.
- [100] Babak Abi et al. “Deep Underground Neutrino Experiment (DUNE), Far Detector Technical Design Report, Volume II: DUNE Physics” (Feb. 2020). neutrino fluxes available from <https://glaucus.crc.nd.edu/DUNEfluxes/>. arXiv: 2002.03005 [hep-ex].
- [101] Alexander Friedland and Shirley Weishi Li. “Understanding the energy resolution of liquid argon neutrino detectors”. *Phys. Rev. D* 99.3 (2019), p. 036009. DOI: 10.1103/PhysRevD.99.036009. arXiv: 1811.06159 [hep-ph].
- [102] R. Acciarri et al. “Demonstration of MeV-Scale Physics in Liquid Argon Time Projection Chambers Using ArgoNeuT”. *Phys. Rev. D* 99.1 (2019), p. 012002. DOI: 10.1103/PhysRevD.99.012002. arXiv: 1810.06502 [hep-ex].
- [103] Trevor Hastie, Robert Tibshirani, and Jerome Friedman. *The Elements of Statistical Learning*. Springer Series in Statistics. New York, NY, USA: Springer New York Inc., 2001. Chap. 11.
- [104] Christopher M. Bishop. *Neural Networks for Pattern Recognition*. USA: Oxford University Press, Inc., 1995. ISBN: 0198538642.
- [105] B. Abi et al. “Experiment Simulation Configurations Approximating DUNE TDR” (Mar. 2021). arXiv: 2103.04797 [hep-ex].
- [106] T. Alion et al. “Experiment Simulation Configurations Used in DUNE CDR” (June 2016). arXiv: 1606.09550 [physics.ins-det].
- [107] Joachim Kopp et al. “Sterile Neutrino Oscillations: The Global Picture”. *JHEP* 05 (2013), p. 050. DOI: 10.1007/JHEP05(2013)050. arXiv: 1303.3011 [hep-ph].

- [108] Mona Dentler et al. “Updated Global Analysis of Neutrino Oscillations in the Presence of eV-Scale Sterile Neutrinos”. *JHEP* 08 (2018), p. 010. DOI: 10.1007/JHEP08(2018)010. arXiv: 1803.10661 [hep-ph].
- [109] Joachim Kopp. “GitHub repository: <https://github.com/koppj/nu-osc>”. 2024.
- [110] B. Abi et al. *Low exposure long-baseline neutrino oscillation sensitivity of the DUNE experiment*. 2021. arXiv: 2109.01304 [hep-ex].
- [111] S. Dolan, G. D. Megias, and S. Bolognesi. “Implementation of the SuSAv2-meson exchange current 1p1h and 2p2h models in GENIE and analysis of nuclear effects in T2K measurements”. *Phys. Rev. D* 101.3 (2020), p. 033003. DOI: 10.1103/PhysRevD.101.033003. arXiv: 1905.08556 [hep-ex].
- [112] A. Papadopoulou et al. “Inclusive Electron Scattering And The GENIE Neutrino Event Generator”. *Phys. Rev. D* 103 (2021), p. 113003. DOI: 10.1103/PhysRevD.103.113003. arXiv: 2009.07228 [nucl-th].
- [113] A. Filkins et al. “Double-differential inclusive charged-current ν_μ cross sections on hydrocarbon in MINERvA at $\langle E_\nu \rangle \sim 3.5$ GeV”. *Phys. Rev. D* 101.11 (2020), p. 112007. DOI: 10.1103/PhysRevD.101.112007. arXiv: 2002.12496 [hep-ex].
- [114] M. Khachatryan et al. “Electron-beam energy reconstruction for neutrino oscillation measurements”. *Nature* 599.7886 (2021), pp. 565–570. DOI: 10.1038/s41586-021-04046-5.
- [115] Vedran Brdar and Joachim Kopp. “Can standard model and experimental uncertainties resolve the MiniBooNE anomaly?” *Phys. Rev. D* 105.11 (2022), p. 115024. DOI: 10.1103/PhysRevD.105.115024. arXiv: 2109.08157 [hep-ph].
- [116] Ciaran Hasnip. *DUNE-PRISM: Reducing neutrino interaction model dependence with a movable neutrino detector*. 2025. arXiv: 2501.14811 [hep-ex].
- [117] R. Stanek et al. *PIII Project Overview and Status*. 2023. arXiv: 2311.05456 [physics.acc-ph].
- [118] C. Hasnip. “DUNE-PRISM - a new method to measure neutrino oscillations”. PhD thesis. Oxford University, Oxford U., 2023. DOI: 10.5287/zenodo.9082v9m.

- [119] Laura Fields. *DUNE flux repository*. <https://glaucus.crc.nd.edu/-DUNEFluxes/>.
- [120] Albert Puig and Jonas Eschle. “phasespace: n-body phase space generation in Python”. *Journal of Open Source Software* (2019). DOI: [10.21105/joss.01570](https://doi.org/10.21105/joss.01570).
- [121] V. Hewes et al. “Deep Underground Neutrino Experiment (DUNE) Near Detector Conceptual Design Report”. *Instruments* 5.4 (2021), p. 31. DOI: [10.3390/instruments5040031](https://doi.org/10.3390/instruments5040031). arXiv: [2103.13910](https://arxiv.org/abs/2103.13910) [physics.ins-det].
- [122] Michael Dawson-Haggerty. *Trimesh*. <https://github.com/mikedh/trimesh>.
- [123] F. Acerbi et al. “Design and performance of the ENUBET monitored neutrino beam”. *Eur. Phys. J. C* 83.10 (2023), p. 964. DOI: [10.1140/epjc/s10052-023-12116-3](https://doi.org/10.1140/epjc/s10052-023-12116-3). arXiv: [2308.09402](https://arxiv.org/abs/2308.09402) [hep-ex].
- [124] Anna Baratto-Roldán et al. “NuTag: a proof-of-concept study for a long-baseline neutrino beam”. *Eur. Phys. J. C* 84.10 (2024), p. 1024. DOI: [10.1140/epjc/s10052-024-13324-1](https://doi.org/10.1140/epjc/s10052-024-13324-1). arXiv: [2401.17068](https://arxiv.org/abs/2401.17068) [physics.acc-ph].
- [125] N. Abgrall et al. “NA61/SHINE facility at the CERN SPS: beams and detector system”. *JINST* 9 (2014), P06005. DOI: [10.1088/1748-0221/9/06/P06005](https://doi.org/10.1088/1748-0221/9/06/P06005). arXiv: [1401.4699](https://arxiv.org/abs/1401.4699) [physics.ins-det].
- [126] H. Adhikary et al. “Proposal from the NA61/SHINE Collaboration for update of European Strategy for Particle Physics” (July 2025). arXiv: [2507.08602](https://arxiv.org/abs/2507.08602) [nucl-ex].
- [127] Patrick Huber, M. Lindner, and W. Winter. “Simulation of long-baseline neutrino oscillation experiments with GLoBES (General Long Baseline Experiment Simulator)”. *Comput. Phys. Commun.* 167 (2005), p. 195. DOI: [10.1016/j.cpc.2005.01.003](https://doi.org/10.1016/j.cpc.2005.01.003). arXiv: [hep-ph/0407333](https://arxiv.org/abs/hep-ph/0407333).
- [128] Joachim Kopp. “Efficient numerical diagonalization of hermitian 3 x 3 matrices”. *Int. J. Mod. Phys. C* 19 (2008), pp. 523–548. DOI: [10.1142/S0129183108012303](https://doi.org/10.1142/S0129183108012303). arXiv: [physics/0610206](https://arxiv.org/abs/physics/0610206).
- [129] Joachim Kopp, Michele Maltoni, and Thomas Schwetz. “Are There Sterile Neutrinos at the eV Scale?” *Phys. Rev. Lett.* 107 (2011), p. 091801. DOI: [10.1103/PhysRevLett.107.091801](https://doi.org/10.1103/PhysRevLett.107.091801). arXiv: [1103.4570](https://arxiv.org/abs/1103.4570) [hep-ph].

[130] Ivan Esteban et al. “NuFit-6.0: updated global analysis of three-flavor neutrino oscillations”. *JHEP* 12 (2024), p. 216. DOI: 10.1007/JHEP12(2024)216. arXiv: 2410.05380 [hep-ph].