

**Measurements of three-flavor neutrino oscillations
from a PISCES two-detector fit to the NOvA**

Experiment data

A DISSERTATION

Submitted to the Graduate College
of the University of Cincinnati
in partial fulfillment
of the requirements
for the degree of

Doctor of Philosophy

in the Department of Physics
of the College of Arts and Sciences

by

Miriana Rajaoalisoa

M.S., Physics, University of Antananarivo

University of Cincinnati

December, 2025

Committee Chair: Alexandre Sousa

ABSTRACT

NOvA is a long-baseline neutrino oscillation experiment with two functionally identical detectors: a Near Detector (ND) at Fermilab, placed 1 km from the neutrino source, and a Far Detector (FD) located 810 km away from the ND in Minnesota. NOvA's primary physics goals are to measure the neutrino oscillation parameters θ_{23} and Δm_{32}^2 with high precision, determine the neutrino mass hierarchy, and constrain the value of δ_{CP} , primarily via the study of muon neutrino to electron neutrino oscillation. Extracting values for oscillation parameters from fits to data usually relies on treating systematic uncertainties as nuisance parameters, a strategy that suffers from poor scalability as the number of uncertainties becomes larger. This work introduces PISCES (Parameter Inference with Systematic Covariance and Exact Statistics), a novel method that circumvents this scalability problem by encoding systematic uncertainties into a covariance matrix. PISCES utilizes a nested minimization in which optimal systematic pulls are first computed using the covariance matrix in an inner minimization step, then the oscillation parameters are profiled over in the outer minimization. PISCES also uses a Poisson Likelihood term, making it ideal for the inclusion of low-statistic samples in the fits.

PISCES is a flexible framework that also supports complex fits, such as a joint Near and Far detector fit. In the standard NOvA analysis, oscillation parameters are extracted using an extrapolation technique in which the ND data indirectly constrain the FD prediction via a ratio method. PISCES, on the other hand, enables a simultaneous ND+FD fit, allowing the high-statistics ND data to directly constrain systematic uncertainties across all samples.

This thesis presents the full PISCES joint ND+FD fit for the NOvA three-flavor analysis, details its implementation, and evaluates its performance through extensive robustness tests and fake data studies. It also provides a comparison between the PISCES joint ND+FD

results and the standard NOvA extrapolation method using the full NOvA 10-year data set. The results demonstrate that PISCES can successfully fit NOvA data while incorporating the constraints from the ND detectors consistently, using physically motivated systematic uncertainties to account for data/MC discrepancies.

ACKNOWLEDGEMENT

I would like to extend my gratitude to everyone who helped me, whether directly or indirectly, in completing this dissertation.

First and foremost, I am deeply thankful to my advisors, Alexandre Sousa and Adam Aurisano, for their steady support and for sharing their scientific expertise throughout my research years. I am especially thankful for the clarity and perspective they brought to this work. I also appreciate their flexibility and understanding over the years. Their mentorship and the opportunity to work on a wide range of exciting projects, from data analysis to hardware, made my experience in the group genuinely rewarding, and for that, I am extremely grateful.

A special thanks goes to V Hewes. I truly appreciate all of her invaluable help in this analysis. I deeply value that, despite her many responsibilities, she still made time for extra meetings to share her expertise and help guide the development of this work.

I would like to thank the NOvA three-flavor group conveners, Erika Cataño-Mur, Zoya Vallari, Liudmila Kolupaeva, and Artur Sztuc, for their guidance and thoughtful feedback throughout this analysis. Their suggestions, questions, and discussions played an important role in shaping the final version of this work.

I am grateful to Mary Bishai for giving me my first opportunity to work in the field of neutrino physics. One of her talks is what originally convinced me to pursue neutrino physics, and I am glad I made that choice, as I have thoroughly enjoyed the years spent trying to understand neutrinos a little better.

I would also like to express my gratitude to my committee members for their time and effort in reviewing this dissertation.

Last but not least, I am thankful to my family for all of their support. I am especially grateful to my husband, Herilala. Thank you for being by my side since my very first year in physics. Your support has made navigating this field much more doable and much

more enjoyable. I am also thankful to my sister, Sarah, and brother, Nathan, for their encouragement and unwavering moral support. It meant more to me than they know.

TABLE OF CONTENTS

ABSTRACT	i
ACKNOWLEDGEMENT	iv
LIST OF TABLES	ix
LIST OF ILLUSTRATIONS	xi
Introduction	1
1.1 History of Neutrinos	1
1.2 Neutrinos in the Standard Model	7
1.3 Neutrino masses in the minimal SM	8
1.4 SM extension to include neutrino masses	10
1.5 Seesaw mechanism	13
1.6 Physics of Neutrino Oscillations	15
1.7 Neutrino Oscillations in Vacuum	20
1.8 Neutrino Oscillations in Matter	30
1.9 Different types of neutrino interactions	41
1.10 Coherent Pion Production	45
1.11 Types of Neutrino Experiments	46
The NOvA Experiment	50
2.1 The NuMI Beam	51
2.2 NOvA Off-axis Detector Location	58
2.3 The NOvA Detectors	59
2.4 NOvA Simulation	67

2.5	NOvA Reconstruction	73
	The NOvA 3-Flavor Oscillation Analysis	86
3.1	The NOvA Extrapolation method	91
3.2	Recent NOvA Results	99
	Systematic Uncertainties	103
4.1	Flux uncertainties	104
4.2	Detector systematics	107
4.3	Neutrino Cross Sections	112
4.4	Neutron uncertainties	124
4.5	Normalization systematics	126
	The PISCES Method	129
5.1	PISCES	130
5.2	Constructing covariance matrices with PISCES	139
5.3	PISCES 3-Flavor analysis inputs	140
5.4	Comparison of PISCES with the Official 3-flavor analysis framework	152
5.5	PISCES ND+FD fit - Asimov Fake Data Tests	160
	Systematic-pull inference with PISCES:	168
6.1	Problem statement and objectives	168
6.2	Newton-Raphson method	170
6.3	Newton Method limitations and use of Principal Component Analysis (PCA)	175
6.4	PCA Importance	180
6.5	Validating the Newton-Raphson method	181
6.6	Systematic importance	183

Fake Data Studies	190
7.1 MINERvA MEC Fake Data Study	190
7.2 CRPA Fake Data Study	207
7.3 Increased NC ν_e Components Fake Data Study	221
7.4 Difference in cross section central value weight Fake Data Study	231
Results	249
8.1 PISCES ND Fit using NOvA ND data	249
8.2 PISCES ND+FD fit results	253
8.3 PISCES systematic pulls	254
8.4 Systematic Pulls: PISCES vs. Extrapolation	260
8.5 Far Detector Best-Fit Spectra: PISCES vs. Extrapolation	268
8.6 Contours comparison: PISCES vs. Extrapolation	271
Conclusion	281
BIBLIOGRAPHY	285
APPENDIX A :	297

LIST OF TABLES

TABLE 3.1	Critical values of $\Delta\chi^2$ from the χ^2 distribution, corresponding to common confidence levels for one and two parameters of interest.	98
TABLE 3.2	Impact of each uncertainty source on the fitted parameters. Entries are the change in the best fit when each source is varied by its 1σ range, quoted as $+/-$ shifts.	100
TABLE 3.3	Predicted event composition by sample for neutrino and antineutrino beam running. Signal and background totals include systematic uncertainties (superscript/subscript).	102
TABLE 4.1	Effect of selected focusing systematics on the FHC ND neutrino flux. For each variation, the table shows the flux integral between 1 and 3 GeV, the peak position and width of a Gaussian fit in that range, and the percentage change of the integral relative to the nominal prediction. Only variations with $ \Delta I \geq 1.5\%$ are listed [1].	106
TABLE 4.2	Same as Table 4.1, but for the FHC FD neutrino flux.	106
TABLE 4.3	Same as Table 4.1, but for the RHC ND neutrino flux.	106
TABLE 4.4	Same as Table 4.1, but for the RHC FD neutrino flux.	107
TABLE 4.5	Light level systematic uncertainties in NOvA detectors	111
TABLE 5.1	Comparison between the oscillation parameters used to generate the fake data (NOvA 2024 Best-Fit point) and the best-fit values obtained from the PISCES ND+FD fit. In all cases, the fit accurately recovers the injected reference point values.	162
TABLE 5.2	Comparison between the oscillation parameters used to generate the fake data (Asimov B) and the best-fit values obtained from the PISCES ND+FD fit. In all cases, the fit accurately recovers the injected reference point values.	163
TABLE 5.3	Comparison between the oscillation parameters used to generate the fake data (Asimov C) and the best-fit values obtained from the PISCES ND+FD fit. In all cases, the fit accurately recovers the injected reference point values.	164
TABLE 5.4	Comparison between the oscillation parameters used to generate the fake data (Asimov D) and the best-fit values obtained from the PISCES ND+FD fit. In all cases, the fit accurately recovers the injected reference point values.	165
TABLE 5.5	Comparison between the oscillation parameters used to generate the fake data (Asimov E) and the best-fit values obtained from the PISCES ND+FD fit. In all cases, the fit accurately recovers the injected reference point values.	166
TABLE 8.1	Systematic pull values comparison between PISCES ND+FD fit and the Extrapolation method, with the uncertainties ranked by their SHAP importance from PISCES (higher first)	266

TABLE 8.2	PISCES systematics without a clean counterpart in the Extrapolation list.	268
TABLE 8.3	Extrapolation systematics not present (as-is) in the PISCES pull list.	268
TABLE 8.4	Best-fit values for 10 years of NOvA data using the PISCES ND+FD joint fit, compared to the best-fit values obtained with the Extrapolation method. The 1σ intervals in the first and second columns are derived from 1D profile scans (using $\Delta\chi^2 = 1$) and are not Feldman-Cousins corrected. The final column shows the official Extrapolation 1σ intervals reported in [2] using the Feldman-Cousins approach.	279
TABLE A.1	Systematic pulls from the Extrapolation method.	298
TABLE A.2	Subset of pairwise correlation coefficients ρ between the systematic-uncertainty response vectors used in the PISCES ND+FD fit. The pairs shown here are chosen to highlight the main correlated blocks that also appear in the PISCES/Extrapolation pull comparison (Table 8.1): calibration, RES, kaon flux and normalization, FSI, MEC/RPA, and muon energy scales.	302
TABLE A.3	Matched systematics ranked by PISCES FD-only SHAP (higher first). Extrapolation FD-only shifts shown in typewriter .	304

LIST OF ILLUSTRATIONS

FIGURE 1.1	Atmospheric-neutrino anomaly in the Super-Kamiokande detector. A deficit in the upward-going neutrinos was observed.	7
FIGURE 1.2	Comparison between the data and number of expected ν_μ according to the theoretical model without oscillation (in blue) and the model taking oscillation into account (in red), as a function of $\cos\theta$	7
FIGURE 1.3	Standard Model of Elementary Particles	8
FIGURE 1.4	Feynman diagrams of the coherent forward elastic scattering processes that generate the CC potential V_{CC} through W exchange and the NC potential V_{NC} through Z exchange.	30
FIGURE 1.5	Illustration of the two weak interaction processes for neutrinos. Left: Charged Current (CC) interaction mediated by $W^\pm$, producing the corresponding charged lepton and some hadrons. Right: Neutral Current (NC) interaction mediated by Z^0 , where the neutrino exits with the same flavor, and only the hadronic system is observed.	42
FIGURE 1.6	CCQE interaction: $\nu_\ell n \rightarrow \ell^- p$	43
FIGURE 1.7	MEC (2p2h): $\nu_\ell + (NN) \rightarrow \ell^- + N + N'$; interaction with a correlated nucleon pair via meson-exchange currents, producing two outgoing nucleons.	44
FIGURE 1.8	CCRES: $\nu_\ell N \rightarrow \ell^- N' \pi^+$; single-pion production via baryon resonance excitation (Δ^{++}).	44
FIGURE 1.9	CCDIS: $\nu_\ell N \rightarrow \ell^- X$. A charged-current deep inelastic scattering process where a W boson couples the lepton and quark currents; the struck parton hadronizes into the multi-hadron system X	45

FIGURE 1.10	CC coherent pion production on a nucleus A : $\nu_\mu A \rightarrow \mu^- \pi^+ A$. The entire nucleus recoils coherently (small $ t $) and remains intact, producing a very forward π^+	46
FIGURE 2.1	Fermilab Accelerator Complex: H^- ions are accelerated to 400 MeV in the Linac, then injected into the Booster, where they are ramped up to 8 GeV and stripped of their electrons to form a proton beam. The protons are then transferred to the Main Injector, where they are accelerated to 120 GeV before being delivered to the NuMI beamline [3].	52
FIGURE 2.2	Schematic of the NuMI beamline: Protons (~ 120 GeV) from the Fermilab Main Injector strike a graphite target in the Target Hall, producing secondary mesons ($\pi^\pm, K^\pm$). Two pulsed magnetic horns (<i>Horn 1</i> and <i>Horn 2</i>) select the sign and focus the secondaries toward the beam axis, establishing either ν_μ running (focus π^+, K^+) or $\bar{\nu}_\mu$ running (focus π^-, K^-). The focused mesons traverse a 675 m helium-filled decay pipe where they decay in flight, e.g. $\pi^+ \rightarrow \mu^+ \nu_\mu$. Downstream, a hadron absorber stops remnant protons and undecayed hadrons, while muon monitors embedded in a rock sample measure the accompanying muon flux as a proxy for the neutrino beam direction and intensity [4].	54

FIGURE 2.3 **Illustration of magnetic horn focusing in neutrino and antineutrino beam modes.** Protons impinging on the graphite target generate a mix of secondary mesons. These particles are then steered using two pulsed magnetic horns, which produce a toroidal magnetic field that selectively focuses particles of a given charge and defocuses those of the opposite sign. (a) In antineutrino mode, the horn polarity is set to focus negatively charged mesons like π^- , which subsequently decay via $\pi^- \rightarrow \mu^- \bar{\nu}_\mu$, generating a beam dominated by $\bar{\nu}_\mu$. (b) Reversing the horn current shifts the focus to positively charged mesons such as π^+ . These decay through $\pi^+ \rightarrow \mu^+ \nu_\mu$, producing a beam primarily composed of ν_μ (neutrino mode). Particles not aligned with the desired charge are defocused and removed from the beam path. By alternating the direction of the horn current, the NuMI beamline can switch between neutrino and antineutrino modes. 57

FIGURE 2.4 **Off-axis energy spectra in NOvA.** Simulated ν_μ CC event-rate spectra (per 6×10^{20} POT per kton) for different off-axis angles: on-axis (black), 7 mrad (green), 14.6 mrad *NOvA* (red), and 21 mrad (blue). Increasing the off-axis angle narrows the spectrum and shifts it to lower energy due to two-body decay kinematics of $\pi^\pm \rightarrow \mu^\pm \nu_\mu(\bar{\nu}_\mu)$, approximately $E_\nu \simeq 0.43 E_\pi / (1 + \gamma^2 \theta^2)$. NOvA's 14.6 mrad configuration yields a narrow-band beam peaked near 2 GeV, aligning with the first oscillation maximum for the $L \simeq 810$ km baseline [5]. 60

FIGURE 2.5	3D schematic of the NOvA detector and its reconstruction principle. A neutrino from the NuMI beam interacts in the liquid scintillator inside the PVC cells, producing charged particles that generate scintillation light. The signals recorded in successive detector planes are used to reconstruct particle trajectories. Alternating planes of vertically and horizontally oriented cells allow for three-dimensional tracking, as shown in the top and side views, where multiple outgoing particles are identified from a single interaction vertex.	61
FIGURE 2.6	PVC extrusion modules used in the NOvA detectors. Each extrusion consists of 16 individual rectangular cells formed from rigid PVC, with approximate transverse dimensions of $3.9 \times 6.6 \text{ cm}^2$ per cell. The left panel shows the technical drawing with dimensions in millimeters, while the right panels display photographs of the fabricated extrusion. These extrusions are assembled into planes, with cells filled with liquid scintillator and read out by wavelength-shifting fibers coupled to APDs.	62
FIGURE 2.7	NOvA detector cell: a 15 m long tube filled with liquid scintillator. A wavelength-shifting fiber is located on one end of the readout, which absorbs and re-emits scintillation light to enhance the signal. At the far end of the cell is an avalanche photodiode to collect the measured light output.	64
FIGURE 2.8	Overview of the NOvA Data Acquisition (DAQ) system with the ART-DAQ Data-Driven Trigger (DDT) framework [6].	66

FIGURE 2.9	Example event displays of different neutrino interaction topologies in NOvA. Top: a ν_μ charged-current (CC) interaction producing a long muon track and a short proton recoil. Middle: a ν_e CC interaction producing an electromagnetic shower from the outgoing electron, together with a proton. Bottom: a neutral-current (NC) interaction producing a pion and hadronic activity.	74
FIGURE 2.10	Illustration of the Hough transform. Left: hits in detector (x, y) space corresponding to a track (blue), with the true underlying line shown (blue dashed). Right: the same hits mapped into (ρ, θ) Hough space. Points lying on the same straight line in (x, y) intersect at a single peak in (ρ, θ) space, which identifies the parameters of the track.	78
FIGURE 2.11	Confusion matrices for Event CVN classification in NOvA simulation. Left: efficiency (fraction of true events correctly classified). Right: purity (fraction of predicted events that are truly that class). The CVN achieves high efficiency and purity across ν_μ CC, ν_e CC, NC, and cosmic categories.	82
FIGURE 2.12	Schematic layout of the NOvA Convolutional Visual Network (CVN) architecture, inspired by MobileNetV2. Bottleneck blocks, depthwise convolutions, and squeeze-excite modules allow efficient feature extraction from event or prong images.	84
FIGURE 2.13	Confusion matrices for Prong CVN classification in NOvA simulation. Left: efficiency. Right: purity. Prong CVN classifies individual prongs as electron, muon, proton, pion, or photon with high accuracy, complementing the Event CVN in particle identification.	85

FIGURE 3.1	Bi-event plot for NOvA, showing expected ν_e vs. $\bar{\nu}_e$ appearance event counts at the Far Detector for normal hierarchy (NH, blue) and inverted hierarchy (IH, red). Each ellipse corresponds to variation of δ_{CP} over $[0, 2\pi]$. The displacement between the NH and IH ellipses comes from the matter effect, while the orientation of each ellipse is determined by δ_{CP}	92
FIGURE 3.2	NOvA disappearance channel at $L = 810$ km (approximate vacuum expression), extended to 5 GeV. The first minimum occurs near $E \simeq 1.65$ GeV, with the dip depth set by $\sin^2 2\theta_{23} c_{13}^4$ and the position by $ \Delta m_{32}^2 $	93
FIGURE 3.3	NOvA appearance channel at $L = 810$ km (leading term only). The matter effect enters via $(1 - \hat{A})^{-2}$	93
FIGURE 3.4	NOvA's extrapolation technique - The ND measured data reconstructed energy spectrum (1) is first converted into a true energy spectrum (3) via the migration matrix (2). The ratio of the ND data and the simulated Monte Carlo in true energy bins (4) is calculated and used as a weight for the simulated events in the FD (6), along with the probability oscillation weight (5). Finally, the reconstructed-to-true migration matrix (7) is applied to the FD true energy spectrum to obtain the FD prediction (8).	94
FIGURE 3.5	Allowed regions in the $(\sin^2 \theta_{23}, \Delta m_{32}^2)$ plane. The solid band shows the current NOvA result (90% C.L.), with previous NOvA results and external constraints overlaid for comparison. Contours are marginalized over mass ordering as indicated (Normal MO shown here). The “B” and “F” superscripts indicate Bayesian or Frequentist analyses, respectively.	101

FIGURE 5.1	Evaluation of χ^2 at each point in the oscillation parameter space - starting from the nominal predictions and observed data, the optimal set of systematic scales that would shift the nominal prediction to match the data is determined. Once the optimal shifted prediction is found, a fit for the oscillation parameters is performed. The determination of systematic scales is carried out simultaneously across all samples in the analysis, including both ND and FD data.	138
FIGURE 5.2	Simulated ν_μ CC distributions of E_{had}/E_ν versus reconstructed E_ν for FHC (left) and RHC (right) running. The contour lines indicate the quartile boundaries used to categorize events by hadronic energy fraction, which correlates strongly with energy resolution.	141
FIGURE 5.3	PISCES ν_e samples distribution - The plots represent ν_e in the ND, separated into Low PID and High PID, for both FHC and RHC. The ν_e samples are binned as a function of reconstructed energy and number of Michel electrons to enhance information content.	144
FIGURE 5.5	Example of a covariance matrix incorporating combined ND and FD samples, used in a joint PISCES ND+FD fit where the ND and FD samples are fit simultaneously.	152
FIGURE 5.6	Stats-only comparison of PISCES and official 3-Flavor fits using FD-only information considering an Asimov fake data generated at the NOvA 2024 best-fit point. The results obtained from both frameworks are virtually identical.	153

FIGURE 5.7	Comparison of PISCES and official 3-Flavor fits using FD-only information considering an Asimov fake data generated at the NOvA 2024 best-fit point. In this case, the fit includes systematic uncertainties that are identical between the two frameworks. The contours obtained from both methods are also identical.	154
FIGURE 5.8	Comparison of sensitivity contours using the full set of systematic uncertainties in both the PISCES FD-only fit and the 3-Flavor non-extrapolated prediction fit. The results show strong agreement between the two approaches, with slight differences attributed to the distinctions in the systematic lists used in each analysis.	155
FIGURE 5.9	Comparison of contours considering the full systematic uncertainties for the official 3-Flavor analysis with and without ND information. The ND information is incorporated via the extrapolation method, with the corresponding contours produced using extrapolated predictions from ND fake data. The addition of ND information leads to tighter contours due to the constraints on systematic uncertainties provided by the extrapolation.	157
FIGURE 5.10	Comparison of FD-only and simultaneous ND+FD fits in the PISCES framework. Including ND samples directly in the fit improves the constraint on systematic uncertainties, resulting in tighter contours. The shifts are more pronounced than those observed with the extrapolation method, demonstrating the benefit of a joint two-detector fit.	158

FIGURE 5.11	Comparison of a stats-only fit using PISCES and the official 3-Flavor framework with the ND information included. In this case, all ND and FD samples in PISCES are fit simultaneously, while for the official 3-Flavor framework, the extrapolated predictions derived using the ND fake data were used in the fit. The results obtained from both methods are consistent in this case.	159
FIGURE 5.12	Comparison between the PISCES simultaneous ND+FD fit and the official 3-Flavor fit using extrapolated predictions. Tighter contours are obtained in the PISCES joint fit, indicating greater information gain from the direct inclusion of ND samples.	160
FIGURE 5.13	1σ , 2σ and 3σ contours comparison considering an Asimov fake data, generated at the 2024 Best-Fit point, for the $\sin^2 \theta_{23}$ vs Δm_{32}^2 and δ_{CP} vs $\sin^2 \theta_{23}$ parameter space. The first rows consider the normal ordering for Δm_{32}^2 , the second rows are for the inverted ordering.	163
FIGURE 5.14	1σ , 2σ and 3σ contours comparison considering an Asimov fake data, generated at the Asimov B reference point, for the $\sin^2 \theta_{23}$ vs Δm_{32}^2 and δ_{CP} vs $\sin^2 \theta_{23}$ parameter space. The first rows consider the normal ordering for Δm_{32}^2 , the second rows are for the inverted ordering. . . .	164
FIGURE 5.15	1σ , 2σ and 3σ contours comparison considering an Asimov fake data, generated at the Asimov C reference point, for the $\sin^2 \theta_{23}$ vs Δm_{32}^2 and δ_{CP} vs $\sin^2 \theta_{23}$ parameter space. The first rows consider the normal ordering for Δm_{32}^2 , the second rows are for the inverted ordering. . . .	165
FIGURE 5.16	1σ , 2σ and 3σ contours comparison considering an Asimov fake data, generated at the Asimov D reference point, for the $\sin^2 \theta_{23}$ vs Δm_{32}^2 and δ_{CP} vs $\sin^2 \theta_{23}$ parameter space. The first rows consider the normal ordering for Δm_{32}^2 , the second rows are for the inverted ordering. . . .	166

FIGURE 5.17	1 σ , 2 σ and 3 σ contours comparison considering an Asimov fake data, generated at the Asimov E reference point, for the $\sin^2 \theta_{23}$ vs Δm_{32}^2 and δ_{CP} vs $\sin^2 \theta_{23}$ parameter space. The first rows consider the normal ordering for Δm_{32}^2 , the second rows are for the inverted ordering. . . .	167
FIGURE 6.1	Example of the geometric view of Newton-Raphson updates for a non-linear equation $f(x) = 0$, assuming $f(x) = x^3 - 6x + 4$. Starting from an initial guess $x^{(0)}$ on the curve (solid black), the dashed line shows the tangent at $x^{(0)}$; its x -intercept (colored $\times$) is the next iteration $x^{(1)}$. Repeating from $x^{(1)}$ (and again from $x^{(2)}$) illustrates how each step moves to the x -intercept of the current tangent, typically converging rapidly to the root x^*	171
FIGURE 6.2	Preliminary validation test: comparison of injected versus recovered pulls for single-shift toy experiments. Results are shown for MEC Comp 1, MEC Comp 2, PPFX Comp 1, and PPFX Comp 2. The diagonal line indicates perfect agreement, and the clustering of points along this line indicates that the Newton-Raphson method accurately recovers the injected pulls in the one systematics at a time test.	182
FIGURE 6.3	Validation of the Newton-Raphson with PCA method using 200 pseudoexperiments. The method successfully recovers the injected pulls for leading systematics (top), while subdominant systematics remain less constrained but have a negligible impact on the overall result (bottom).	184

FIGURE 6.4	SHAP importance ranking for the fake data test in which the Neutral Current (NC) background in the ν_e samples is artificially doubled. The plot shows the top 20 systematic uncertainties driving the fit, together with their recovered pulls from the Newton-Raphson method. NC-related systematics (e.g. MaNCRES, NC-DIS terms) respond strongly, as expected when the NC background is inflated.	187
FIGURE 6.5	SHAP importance ranking for fake data generated at Asimov E and fit using the Asimov A prediction. The top 20 systematic uncertainties with the highest SHAP importance are shown. Appearance-sensitive systematics (e.g. ν_e radiative corrections, axial CCQE knobs, second-class current terms) acquire significant pulls, reflecting the shifted $\nu_e/\bar{\nu}_e$ rates and shapes in the Asimov E fake data.	189
FIGURE 7.2	1σ , 2σ and 3σ contours comparison considering the MINERvA MEC weighted fake data, generated at the 2024 Best-Fit point (dashed lines), and the nominal Asimov (solid lines), for the $\sin^2\theta_{23}$ vs Δm_{32}^2 and δ_{CP} vs $\sin^2\theta_{23}$ parameter space. The first rows are considering the normal ordering for Δm_{32}^2 , the second rows are for the inverted ordering. The shifts in contours and Best-fit points are minimal, showing that no significant bias was found and that the fitter is robust against the considered change.	199
FIGURE 7.3	Difference in $\Delta\chi^2$ surfaces between Asimov and MINERvA MEC fake data, generated at the 2024 Best-Fit point. The variation across all the different analysis bins is minimal in all 4 parameter space combinations.	200
FIGURE 7.5	Systematic pulls using the MINERvA MEC model fake data. The fitter responds primarily in MEC and QE-like systematic directions, as expected.	206

FIGURE 7.7	1 σ , 2 σ and 3 σ contours comparison considering the CRPA weighted fake data, generated at the 2024 Best-Fit point (dashed lines), and the nominal Asimov (solid lines), for the $\sin^2 \theta_{23}$ vs Δm_{32}^2 and δ_{CP} vs $\sin^2 \theta_{23}$ parameter space. The first rows consider the normal ordering for Δm_{32}^2 , the second rows are for the inverted ordering. The shifts in contours and Best-fit points are minimal, showing that no significant bias was found and that the fitter is robust against the considered change.	215
FIGURE 7.8	Difference in $\Delta\chi^2$ surfaces between Asimov and CRPA fake data, generated at the 2024 Best-Fit point. The variation across all the different analysis bins is minimal in all 4 parameter space combinations.	216
FIGURE 7.10	Systematic pulls for the CRPA fake data. The response is dominated by strong pulls in RPA-related directions as expected.	222
FIGURE 7.12	1 σ , 2 σ and 3 σ contours comparison considering the NC ν_e background doubled fake data, generated at the 2024 Best-Fit point (dashed lines), and the nominal Asimov (solid lines), for the $\sin^2 \theta_{23}$ vs Δm_{32}^2 and δ_{CP} vs $\sin^2 \theta_{23}$ parameter space. The first rows are considering the normal ordering for Δm_{32}^2 , the second rows are for the inverted ordering. The shifts in contours and Best-fit points are minimal, showing that no significant bias was found and that the fitter is robust against the considered change.	226
FIGURE 7.13	Difference in $\Delta\chi^2$ surfaces between Asimov and NC ν_e background doubled fake data, generated at the 2024 Best-Fit point. The variation across all the different analysis bins is minimal in all 4 parameter space combinations.	227

FIGURE 7.15	Recovered pulls for the NC ν_e background doubled fake data. The panels show the pull values across systematics, ordered by PCA importance. NC-related DIS parameters show significant responses, consistent with the altered background composition, while detector, flux, and most cross section systematics remain weakly pulled.	232
FIGURE 7.17	1σ , 2σ and 3σ contours comparison considering predictions with (dashed lines)/without (solid line) cross section CV weight, fitted for a data without cross section CV weight, generated at the 2024 Best-Fit point. The contours are for the $\sin^2 \theta_{23}$ vs Δm_{32}^2 and δ_{CP} vs $\sin^2 \theta_{23}$ parameter space. The first rows are considering the normal ordering for Δm_{32}^2 , the second rows are for the inverted ordering. The shifts in contours and Best-fit points are minimal, showing that no significant bias was found even in the case of considerable discrepancy in central values between predictions and data.	242
FIGURE 7.18	Difference in $\Delta\chi^2$ surfaces between using predictions with/without cross section CV weight for fitting for a fake data without cross section CV weight. The variations across all considered surfaces are minimal. . . .	243
FIGURE 7.20	Recovered pulls for the fake data generated without applying the NOvA central value (CV) weight, fit with predictions including this weight. The fitter responds by shifting resonance (MaCCRES), MEC, and QE-like parameters, consistent with the expected mismodeling introduced by omitting the MEC rescaling.	248

FIGURE 8.1	Best-fit spectra (blue histograms) obtained by performing a PISCES fit to the real NOvA ND data (black points), using ND samples. PISCES successfully reproduces the data across all ν_μ and ν_e ND samples. There is excellent agreement for the ν_μ and low-energy ν_e samples, and reasonable agreement for the ν_e Low PID and High PID samples.	250
FIGURE 8.2	Best-fit spectra (blue histograms) obtained by performing a PISCES fit to the real NOvA ND data (black points), using ND samples. PISCES successfully reproduces the data across all ν_μ and ν_e ND samples. There is excellent agreement for the ν_μ and low-energy ν_e samples, and reasonable agreement for the ν_e Low PID and High PID samples.	251
FIGURE 8.3	Recovered pulls from the ND-only fit to the NOvA real ND data. The first panel shows the systematic pulls for the most important uncertainties. The second panel represents the subdominant systematic uncertainties.	253
FIGURE 8.4	Systematic pulls for the PISCES ND+FD Fit.	256
FIGURE 8.5	(a): ν_μ Q1 samples - Comparison of NOvA Data (black) and the best-fit spectra using PISCES ND+FD fit with post-fit systematic weights (blue histograms)	271
FIGURE 8.6	(b): ν_μ Q2 samples - Comparison of NOvA Data (black) and the best-fit spectra using PISCES ND+FD fit with post-fit systematic weights (blue histograms).	272
FIGURE 8.7	(c): ν_μ Q3 samples - Comparison of NOvA Data (black) and the best-fit spectra using PISCES ND+FD fit with post-fit systematic weights (blue histograms).	272

FIGURE 8.8	(d): ν_μ Q4 samples - Comparison of NOvA Data (black) and the best-fit spectra using PISCES ND+FD fit with post-fit systematic weights (blue histograms).	273
FIGURE 8.9	Figure 8.9(e): ν_e samples - Comparison of NOvA Data (black) and the best-fit spectra using PISCES ND+FD fit with post-fit systematic weights (blue histograms).	274
FIGURE 8.10	1σ , 2σ and 3σ contours comparison using the Extrapolation method (dashed lines), and the PISCES ND+FD fit (solid lines) considering the NOvA Real Data, for the $\sin^2\theta_{23}$ vs Δm_{32}^2 and δ_{CP} vs $\sin^2\theta_{23}$ parameter space. The first rows consider the normal ordering for Δm_{32}^2 , the second rows are for the inverted ordering.	280
FIGURE A.1	1σ , 2σ and 3σ contours comparison considering the MINERvA MEC weighted (dashed lines), and the nominal Asimov (solid lines) fake data, for the $\sin^2\theta_{23}$ vs Δm_{32}^2 and δ_{CP} vs $\sin^2\theta_{23}$ parameter space, for fake data generated at the Asimov B point. Asimov B is the point closest to T2K's 2020 Best-Fit point. The first rows consider the normal hierarchy for Δm_{32}^2 , the second rows are for the inverted hierarchy. The shifts in contours and best fit points are minimal, showing that no significant bias was found and that the fitter is robust against the considered change.	306

FIGURE A.2 1σ , 2σ and 3σ contours comparison considering the MINERvA MEC weighted (dashed lines), and the nominal Asimov (solid lines) fake data, for the $\sin^2\theta_{23}$ vs Δm_{32}^2 and δ_{CP} vs $\sin^2\theta_{23}$ parameter space, for fake data generated at the Asimov C point. Asimov C probes for inverted hierarchy and maximum CP. The first rows consider the normal hierarchy for Δm_{32}^2 , the second rows are for the inverted hierarchy. The shifts in contours and best fit points are minimal, showing that no significant bias was found and that the fitter is robust against the considered change. 307

FIGURE A.3 1σ , 2σ and 3σ contours comparison considering the MINERvA MEC weighted (dashed lines), and the nominal Asimov (solid lines) fake data, for the $\sin^2\theta_{23}$ vs Δm_{32}^2 and δ_{CP} vs $\sin^2\theta_{23}$ parameter space, for fake data generated at the Asimov D point. Asimov D probes for normal hierarchy, lower octant, and maximum CP. The first rows consider the normal hierarchy for Δm_{32}^2 , the second rows are for the inverted hierarchy. The shifts in contours and best fit points are minimal, showing that no significant bias was found and that the fitter is robust against the considered change. 308

FIGURE A.4 1σ , 2σ and 3σ contours comparison considering the MINERvA MEC weighted (dashed lines), and the nominal Asimov (solid lines) fake data, for the $\sin^2\theta_{23}$ vs Δm_{32}^2 and δ_{CP} vs $\sin^2\theta_{23}$ parameter space, for fake data generated at the Asimov E point. Asimov E was chosen as it is located far from NOvA's preferred oscillation point. Asimov E probes for inverted hierarchy, lower octant, and $\delta_{CP} = \pi/2$. The first rows consider the normal hierarchy for Δm_{32}^2 , the second rows are for the inverted hierarchy. The shifts in contours and best fit points are minimal, showing that no significant bias was found and that the fitter is robust against the considered change. 309

FIGURE A.5 1σ , 2σ and 3σ contours comparison considering the CRPA weighted (dashed lines), and the nominal Asimov (solid lines) fake data, for the $\sin^2\theta_{23}$ vs Δm_{32}^2 and δ_{CP} vs $\sin^2\theta_{23}$ parameter space, for fake data generated at the Asimov B point. Asimov B is the point closest to T2K's 2020 Best-Fit point. The first rows consider the normal hierarchy for Δm_{32}^2 , the second rows are for the inverted hierarchy. The shifts in contours and best fit points are minimal, showing that no significant bias was found and that the fitter is robust against the considered change. 310

FIGURE A.6 1σ , 2σ and 3σ contours comparison considering the CRPA weighted (dashed lines), and the nominal Asimov (solid lines) fake data, for the $\sin^2\theta_{23}$ vs Δm_{32}^2 and δ_{CP} vs $\sin^2\theta_{23}$ parameter space, for fake data generated at the Asimov C point. Asimov C probes for inverted hierarchy and maximum CP. The first rows consider the normal hierarchy for Δm_{32}^2 , the second rows are for the inverted hierarchy. The shifts in contours and best fit points are minimal, showing that no significant bias was found and that the fitter is robust against the considered change. 311

FIGURE A.7 1σ , 2σ and 3σ contours comparison considering the CRPA weighted (dashed lines), and the nominal Asimov (solid lines) fake data, for the $\sin^2\theta_{23}$ vs Δm_{32}^2 and δ_{CP} vs $\sin^2\theta_{23}$ parameter space, for fake data generated at the Asimov D point. Asimov D probes for normal hierarchy, lower octant, and maximum CP. The first rows consider the normal hierarchy for Δm_{32}^2 , the second rows are for the inverted hierarchy. The shifts in contours and best fit points are minimal, showing that no significant bias was found and that the fitter is robust against the considered change. 312

FIGURE A.8 1σ , 2σ and 3σ contours comparison considering the CRPA weighted (dashed lines), and the nominal Asimov (solid lines) fake data, for the $\sin^2\theta_{23}$ vs Δm_{32}^2 and δ_{CP} vs $\sin^2\theta_{23}$ parameter space, for fake data generated at the Asimov E point. Asimov E was chosen as it is located far from NOvA's preferred oscillation point. Asimov E probes for inverted hierarchy, lower octant, and $\delta_{CP} = \pi/2$. The first rows consider the normal hierarchy for Δm_{32}^2 , the second rows are for the inverted hierarchy. The shifts in contours and best fit points are minimal, showing that no significant bias was found and that the fitter is robust against the considered change. 313

FIGURE A.9 1σ , 2σ and 3σ contours comparison considering the NC ν_e background doubled (dashed lines), and the nominal Asimov (solid lines) fake data, for the $\sin^2\theta_{23}$ vs Δm_{32}^2 and δ_{CP} vs $\sin^2\theta_{23}$ parameter space, for fake data generated at the Asimov B point. Asimov B is the point closest to T2K's 2020 Best-Fit point. The first rows consider the normal hierarchy for Δm_{32}^2 , the second rows are for the inverted hierarchy. The shifts in contours and best fit points are minimal, showing that no significant bias was found and that the fitter is robust against the considered change. 315

FIGURE A.10 1σ , 2σ and 3σ contours comparison considering the NC ν_e background weighted (dashed lines), and the nominal Asimov (solid lines) fake data, for the $\sin^2\theta_{23}$ vs Δm_{32}^2 and δ_{CP} vs $\sin^2\theta_{23}$ parameter space, for fake data generated at the Asimov C point. Asimov C probes for inverted hierarchy and maximum CP. The first rows consider the normal hierarchy for Δm_{32}^2 , the second rows are for the inverted hierarchy. The shifts in contours and best fit points are minimal, showing that no significant bias was found and that the fitter is robust against the considered change. 316

FIGURE A.11 1σ , 2σ and 3σ contours comparison considering the NC ν_e background doubled (dashed lines), and the nominal Asimov (solid lines) fake data, for the $\sin^2\theta_{23}$ vs Δm_{32}^2 and δ_{CP} vs $\sin^2\theta_{23}$ parameter space, for fake data generated at the Asimov D point. Asimov D probes for normal hierarchy, lower octant, and maximum CP. The first rows consider the normal hierarchy for Δm_{32}^2 , the second rows are for the inverted hierarchy. The shifts in contours and best fit points are minimal, showing that no significant bias was found and that the fitter is robust against the considered change. 317

FIGURE A.12 1σ , 2σ and 3σ contours comparison considering the NC ν_e background doubled (dashed lines), and the nominal Asimov (solid lines) fake data, for the $\sin^2\theta_{23}$ vs Δm_{32}^2 and δ_{CP} vs $\sin^2\theta_{23}$ parameter space, for fake data generated at the Asimov E point. Asimov E was chosen as it is located far from NOvA's preferred oscillation point. Asimov E probes for inverted hierarchy, lower octant, and $\delta_{CP} = \pi/2$. The first rows consider the normal hierarchy for Δm_{32}^2 , the second rows are for the inverted hierarchy. The shifts in contours and best fit points are minimal, showing that no significant bias was found and that the fitter is robust against the considered change. 318

FIGURE A.13 1σ , 2σ and 3σ contours comparison considering predictions with (dashed lines)/without (solid lines) cross-section CV weight fitting for a fake data without the weight. The comparisons are for the $\sin^2\theta_{23}$ vs Δm_{32}^2 and δ_{CP} vs $\sin^2\theta_{23}$ parameter space, where the fake data are generated at the Asimov B point. Asimov B is the point closest to T2K's 2020 Best-Fit point. The first rows consider the normal hierarchy for Δm_{32}^2 , the second rows are for the inverted hierarchy. The shifts in contours and best fit points are minimal, showing that no significant bias was found and that the fitter is robust against the considered change. . . . 319

FIGURE A.14 1σ , 2σ and 3σ contours comparison considering predictions with (dashed lines)/without (solid lines) cross-section CV weight fitting for a fake data without the weight. The comparisons are for the $\sin^2\theta_{23}$ vs Δm_{32}^2 and δ_{CP} vs $\sin^2\theta_{23}$ parameter space, where the fake data are generated at the Asimov C point. Asimov C probes for inverted hierarchy and maximum CP. The first rows consider the normal hierarchy for Δm_{32}^2 , the second rows are for the inverted hierarchy. The shifts in contours and best fit points are minimal, showing that no significant bias was found and that the fitter is robust against the considered change. . . . 320

FIGURE A.15 1σ , 2σ and 3σ contours comparison considering predictions with (dashed lines)/without (solid lines) cross-section CV weight fitting for a fake data without the weight. The comparisons are for the $\sin^2\theta_{23}$ vs Δm_{32}^2 and δ_{CP} vs $\sin^2\theta_{23}$ parameter space, where the fake data are generated at the Asimov D point. Asimov D probes for normal hierarchy, lower octant, and maximum CP. The first rows consider the normal hierarchy for Δm_{32}^2 , the second rows are for the inverted hierarchy. The shifts in contours and best fit points are minimal, showing that no significant bias was found and that the fitter is robust against the considered change. 321

FIGURE A.16 1σ , 2σ and 3σ contours comparison considering predictions with (dashed lines)/without (solid lines) cross-section CV weight fitting for a fake data without the weight. The comparisons are for the $\sin^2\theta_{23}$ vs Δm_{32}^2 and δ_{CP} vs $\sin^2\theta_{23}$ parameter space, where the fake data are generated at the Asimov E point. Asimov E was chosen as it is located far from NOvA's preferred oscillation point. Asimov E probes for inverted hierarchy, lower octant, and $\delta_{CP} = \pi/2$. The first rows consider the normal hierarchy for Δm_{32}^2 , the second rows are for the inverted hierarchy. The shifts in contours and best fit points are reasonable, showing that no significant bias was found and that the fitter is robust against the considered change. 322

Chapter 1. Introduction

1.1 History of Neutrinos

In 1930, the apparent existence of a decay mode that violated the law of energy conservation led Wolfgang Pauli to propose the existence of a light, neutral particle, now known as the neutrino. This apparent non-conservation of energy was observed in β decay, a process in which a β particle (a β^- electron or a β^+ positron) is emitted. Analysis of the decay products showed that a portion of the initial energy and momentum appeared to be missing. Pauli resolved this paradox by postulating the existence of an undetected particle carrying away the missing energy. At the time, this proposal perplexed the scientific community, since, apart from the missing energy, no other evidence of such a particle could be observed.

In 1956, an experiment conducted at the Savannah River Plant by Clyde Cowan Jr. and Frederick Reines provided the first direct evidence for the existence of neutrinos. The detection was based on observing electron antineutrinos produced in inverse β decay within a nuclear reactor. The detector used was a large water tank, which was doped with cadmium chloride to allow neutron detection. The antineutrino interacts with a proton in the hydrogen of water through inverse beta decay:

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

The resulting positron quickly annihilates with an electron, producing two 511 keV gamma rays emitted in opposite directions, which generate a first flash of light in the detector. A fraction of a millisecond later, the neutron from the interaction is captured by a cadmium

nucleus (mainly ^{113}Cd), forming an excited $^{114}\text{Cd}^*$ state, which subsequently de-excites by emitting gamma rays. The appearance of the first burst of gamma rays, followed by the second one emitted from the neutron capture, provides a distinctive signature of antineutrino interactions that could be recorded via delayed coincidence. The experiment recorded roughly three antineutrino events per hour [7, 8]. This provided the first direct experimental evidence for the existence of neutrinos. For this discovery, Frederick Reines was awarded half of the 1995 Nobel Prize in Physics. Although his collaborator Clyde Cowan had passed away in 1974, Reines emphasized that the achievement was a joint effort between the two [9].

In 1962, the existence of a second type of neutrino, the muon neutrino, was experimentally demonstrated. The neutrinos discovered by the Reines and Cowan experiment were associated with neutrino interactions, including electrons or positrons. Therefore, when the muon was discovered, the existence of the corresponding neutrino was suggested. To test this idea, Leon Lederman, Melvin Schwartz, and Jack Steinberger conducted an experiment using the Alternating Gradient Synchrotron (AGS) at Brookhaven National Laboratory. A high-energy proton beam from the AGS impinged on a target, producing π^+ mesons, which decayed into muons and muon neutrinos through

$$\pi^+ \rightarrow \mu^+ + \nu_\mu.$$

The resulting particle beam passed through a 13.5 m thick (about 5000 ton) steel shield that absorbed all charged particles. Only the muon neutrinos pass through the wall and travel towards a "spark chamber", which is a detector filled with neon. In this chamber, the impact of the ν_μ on some aluminum produces scintillating traces of muons that can be detected and photographed. This conclusively proved the existence of the muon neutrino [10].

In 1975, following the discovery of the τ lepton by Martin L. Perl and collaborators, the existence of a third type of neutrino, the tau neutrino (ν_τ), was postulated [11]. To detect

this particle, the Direct Observation of the NU Tau (DONUT) experiment was established at Fermilab [12]. In DONUT, 800 GeV protons from the Tevatron fixed-target line were directed onto a tungsten target, producing charmed mesons D_s that decay into a tau lepton and a tau neutrino:

$$D_s \rightarrow \tau^+ + \nu_\tau.$$

The τ lepton then decays, producing a secondary antineutrino $\bar{\nu}_\tau$. The ν_τ beam from the meson decay was directed toward a hybrid detector composed of alternating sheets of nuclear emulsion and electronic detectors.

When a ν_τ interacts inside the emulsion, it produces a τ lepton which decays and produces a characteristic signature in the emulsion. The τ actually travels for a short distance (a fraction of a millimeter) before decaying, which creates a short, straight track that suddenly changes direction at the decay point. This "kink" recorded in the emulsion is a clear signature of a τ appearance event.

The electronic detectors, which were composed of scintillator triggers and drift chambers, were used to identify which regions of the emulsion might contain a neutrino event and were therefore worth scanning in detail. Those selected emulsion sheets were then analyzed with automated scanning systems in Nagoya. Out of roughly 6.6 million neutrino events recorded, they selected around 1000 emulsion plates that were thoroughly scanned [12]. In July 2000, the DONUT collaboration announced the first direct observation of tau neutrinos, based on four candidate events, and by 2008 the experiment reported nine observed ν_τ interactions with an estimated background of 1.5 events [13].

In the late 1960s, the Homestake experiment, led by Raymond Davis Jr. and John Bahcall, set out to measure the flux of electron neutrinos (ν_e) produced by nuclear fusion reactions in the Sun [14]. Located deep in the Homestake gold mine in South Dakota, the detector consisted of a large tank filled with 615 tons of perchloroethylene (C_2Cl_4) [15]. Solar neutrinos,

interacting with chlorine nuclei, occasionally produce argon atoms through the reaction

$$\nu_e + {}^{37}\text{Cl} \rightarrow e^- + {}^{37}\text{Ar}.$$

The team inferred the solar neutrino flux by counting the number of argon atoms produced each month. They found that the measured flux was only about one-third of the prediction from solar models, which suggested that a large fraction of the expected electron neutrinos was missing. Several other solar-neutrino experiments confirmed this deficit later on, including experiments such as Kamiokande in Japan [16], GALLEX in Italy [17], and SAGE in Russia [18]. This was called the *solar neutrino problem*, and was an indication that our understanding of either the solar models or the neutrino properties was incomplete. This was also the first hint that neutrinos change flavor as they propagate.

In 1998, the first direct evidence of neutrino oscillation (neutrino flavor change) was made by the Super-Kamiokande experiment in Japan. This had major consequences, as it implied that neutrinos, which were previously assumed to be massless, must in fact have mass. The Super-Kamiokande experiment was originally designed to search for proton decay. Super-Kamiokande used a detector which consisted of a cylindrical tank containing 50,000 tons of ultra-pure water, with over 11,000 photomultiplier tubes (PMTs). When charged particles pass through the water, they emit Cherenkov light that forms visible rings on the detector walls. The detector was able to clearly identify ν_μ and ν_e interactions, as they left distinctive features in the detector, where muon rings appear clear and well-defined, whereas electron rings look more diffuse. The Super-Kamiokande detector was placed about 1 km underground to shield it from cosmic rays.

Super-Kamiokande observed atmospheric neutrinos, which are created when high-energy cosmic rays, mostly protons, collide with the atoms in the Earth's atmosphere. These interactions produce particles such as pions and kaons, which then decay into muons and

neutrinos. The experiment measured the number of ν_μ and ν_e as a function of the neutrino's zenith angle $\cos \theta$, which is defined as the angle between the neutrino direction and the vertical axis of the detector, and observed that the number of upward-going muons ($\cos \theta = -1$) was significantly lower than predicted. However, the number of downward-going muons ($\cos \theta = +1$) agreed with expectations. The main difference between those two is that the upward-going muons travel through nearly the entire Earth before reaching the detector, whereas the downward-going muons travel for only a few kilometers [19]. This deficit was later interpreted as the first evidence of neutrino oscillations. The upward-going neutrinos travelled a sufficiently long distance for the ν_μ to oscillate into another flavor, whereas the path for the downward-going neutrino was too short for significant oscillation to occur, explaining why their flux matched the prediction [20].

In 2001, the Sudbury Neutrino Observatory (SNO) experiment in Canada, which measured the 8B solar spectrum, provided key insight into the *solar neutrino problem* [21]. In the Sun, all neutrinos are produced as electron neutrinos through beta-decay processes in the fusion chain. Using heavy water (D_2O), SNO could distinguish neutrino interactions that produced charged leptons (charged-current interactions) from those where the neutrinos interacted through the neutral weak force, leaving the neutrino flavor unchanged but freeing a neutron from the nucleus (neutral-current interactions). Its results demonstrated that the total active neutrino flux inferred from neutral-current interactions was consistent with the solar-model predictions; however, the charged-current rate showed that only about one-third of the flux arrived as electron neutrinos [21]. This deficit suggested that the remaining flux had transformed into other flavors.

This flavor change actually comes from another interesting property of neutrinos: they do not propagate the same way in matter as they do in vacuum. Inside the Sun's dense core, electron neutrinos interact with the surrounding electrons through coherent forward scattering. This produces an additional effective potential that changes how the neutrino evolves as it

travels outward through the Sun. This phenomenon is known as the Mikheyev–Smirnov–Wolfenstein (MSW) effect [22].

For the high-energy ${}^8\text{B}$ neutrinos that SNO studied, what is actually happening is that the electron neutrino undergoes an adiabatic flavor transformation inside the Sun’s gradually decreasing electron density, as explained in Section 1.8.2. Due to this MSW effect, by the time the neutrino reaches the surface of the Sun, the neutrino is essentially in a pure mass eigenstate. Such a state does not undergo further oscillation in vacuum, so for SNO, the flavor change occurred entirely inside the Sun. This is quite different from the atmospheric neutrino case, where flavor change comes from oscillations over long distances. In the solar case, the observed flavor conversion is fully driven by the MSW effect inside the Sun.

Over the last two decades, oscillation data from solar, atmospheric, reactor, and accelerator experiments have constrained many of the three-flavor parameters with high precision. However, several fundamental properties of neutrinos remain unknown, including the ordering of their masses (often also referred to as mass hierarchy), the existence and magnitude of CP violation in the neutrino sector, the origin of neutrino mass, and the possible existence of sterile neutrinos.




Figure 1.1: Atmospheric-neutrino anomaly in the Super-Kamiokande detector. A deficit in the upward-going neutrinos was observed.




Figure 1.2: Comparison between the data and number of expected ν_μ according to the theoretical model without oscillation (in blue) and the model taking oscillation into account (in red), as a function of $\cos\theta$

1.2 Neutrinos in the Standard Model

In the Standard Model (SM) of particle physics, neutrinos belong to the lepton sector. They appear in each of the three generations of fundamental fermions. A neutrino and its associated lepton can be grouped into an electroweak doublet:

$$L_\alpha = \begin{pmatrix} \nu_{\alpha L} \\ \ell_{\alpha L} \end{pmatrix}, \quad \alpha = e, \mu, \tau.$$

Like all fermions, neutrinos are spin- $\frac{1}{2}$ particles obeying Fermi–Dirac statistics. They carry no color charge and are therefore singlets under the strong interaction symmetry group $SU(3)_C$. They participate only in weak interactions, mediated by the $W^\pm$ and Z^0 bosons, and possess a lepton number $L = +1$. Neutrinos in the minimal SM are predicted to be

exactly massless, which is in contradiction with the evidence from neutrino oscillations.




Figure 1.3: The Standard Model, showing the three generations of quarks and leptons, including neutrinos, and the bosons mediating the fundamental forces.

1.3 Neutrino masses in the minimal SM

In the Standard Model (SM), the charged weak interaction is chiral: it couples to left-handed (spin in the opposite direction of the momentum) fermions but not to right-handed (spin aligned with the momentum) ones. Experiments such as β -decay, muon decay, and collider data all show that the charged weak interaction is maximally chiral (left-handed only) to very high accuracy [23, 24, 25]. Thus, the left-handed leptons of each flavor $\alpha = e, \mu, \tau$

are placed in an electroweak doublet to show that they interact weakly via $W^\pm$:

$$L_\alpha = \begin{pmatrix} \nu_{\alpha L} \\ \ell_{\alpha L} \end{pmatrix} \sim (2, Y = -\frac{1}{2}),$$

whereas the right-handed charged lepton is an $SU(2)_L$ singlet, so that we won't have $W^\pm$ coupling for right-handed leptons:

$$e_{\alpha R} \sim (1, Y = -1).$$

Here $(2, Y)$ indicates the $SU(2)_L$ representation and hypercharge Y (with $Q = T_3 + Y$).

In the Standard Model, the mass term for a Dirac fermion (a particle distinct from its antiparticle) is generated through the Dirac mass term. A Dirac mass couples left and right chiralities via

$$\mathcal{L} \supset -m \bar{\Psi} \Psi = -m (\bar{\Psi}_L \Psi_R + \bar{\Psi}_R \Psi_L),$$

where $\Psi_{L,R} = P_{L,R} \Psi$ and $P_{L,R} = \frac{1}{2}(1 \mp \gamma^5)$. Since the neutrino field has only a left-handed component and no right-handed field to couple to, a Dirac mass term is excluded in the minimal SM.

Another option for mass generation is the Majorana term. In the minimal SM, a Majorana mass would be the following

$$\mathcal{L} \supset -\frac{1}{2} m_L \nu_L^T C \nu_L + \text{h.c.} \quad (1.1)$$

where C is the charge-conjugation matrix and the transpose contracts spinor indices (Lorentz scalar).

However, all terms in the SM Lagrangian must be:

- a Lorentz scalar: every term must be invariant under Lorentz transformations,

- a gauge singlet under $SU(3)_C \times SU(2)_L \times U(1)_Y$, which ensures gauge invariance,
- have mass dimension ≤ 4 , which is required for renormalizability in 4D spacetime.

In the Majorana mass term of Eq. 1.1, the expression is a Lorentz scalar of dimension 3, so it satisfies the first and third conditions above. The issue arises with electroweak gauge invariance. The term is invariant under $SU(2)_L$ when written as $L^T C i\sigma^2 L$, where the antisymmetric tensor $i\sigma^2$ contracts the doublet indices. However, under a $U(1)_Y$ hypercharge transformation with local parameter $\alpha(x)$,

$$L(x) \rightarrow e^{iY\alpha(x)} L(x), \quad Y(L) = -\frac{1}{2},$$

each lepton doublet acquires a phase, so the product transforms as

$$LL \rightarrow e^{i(Y+Y)\alpha} LL = e^{-i\alpha} LL.$$

Hence, the total hypercharge of this operator is $Y = -1$. This violates the hypercharge condition as a gauge-invariant operator $U(1)_Y$ must be neutral ($Y = 0$) [26]. For this reason, a Majorana mass term for left-handed neutrinos is not possible in the minimal Standard Model [26]. An extension of the minimal Standard Model is necessary to explain neutrino masses.

1.4 SM extension to include neutrino masses

Neutrino masses can be generated by extending the SM. Right-handed neutrinos can be introduced to allow the construction of Dirac mass terms. Alternatively, neutrinos can acquire Majorana masses through mechanisms such as the Weinberg operator or the seesaw mechanism.

1.4.1 Dirac Masses

In the Dirac case, the minimal SM can be extended by adding a sterile singlet N_R , which:

- is a singlet under the electroweak group: $N_R \sim (1, Y = 0)$, it is a gauge singlet under $SU(2)_L \times U(1)_Y$ since it has no weak isospin and zero hypercharge;
- is sterile: it carries no Standard Model gauge charges, so it doesn't couple to $W^\pm$ and Z , explaining the fact that right-handed neutrinos are not observed in experiments;
- is a right-handed field, and thus can be used to create the Dirac mass term.

The neutrino Yukawa coupling term in the Lagrangian would be:

$$\mathcal{L}_Y \supset -y_\nu^{\alpha i} \bar{L}_\alpha \tilde{H} N_{iR} + \text{h.c.}, \quad \tilde{H} \equiv i\sigma^2 H^*, \quad (1.2)$$

with L being the electroweak doublet $L \sim (2, -\frac{1}{2})$, H being the Higgs field $H \sim (2, +\frac{1}{2})$, $N_R \sim (1, 0)$ being the right-handed neutrino.

Under $SU(2)$, $\bar{L} \tilde{H}$ is an $SU(2)$ singlet. Concerning hypercharge conservation, under $U(1)_Y$: $Y(\bar{L}) = +\frac{1}{2}$, $Y(\tilde{H}) = -\frac{1}{2}$, and $Y(N_R) = 0$, so the total hypercharge for Eq. 1.2 is 0. Hence, this new Dirac operator is gauge invariant under $SU(2)_L \times U(1)_Y$. Furthermore, the term has a mass dimension 4, making it renormalizable, and hence a valid term in the Standard Model Lagrangian [27, 25].

In the Standard Model, particle masses arise via electroweak symmetry breaking: a complex Higgs doublet $H \sim (2, +\frac{1}{2})$ with potential

$$V(H) = \mu^2 H^\dagger H + \lambda (H^\dagger H)^2$$

acquires a nonzero vacuum expectation value

$$\langle H \rangle = \begin{pmatrix} 0 \\ v/\sqrt{2} \end{pmatrix}, \quad v \simeq 246 \text{ GeV},$$

spontaneously breaking $SU(2)_L \times U(1)_Y \rightarrow U(1)_{\text{em}}$ and generating

$$M_W = \frac{g v}{2}, \quad M_Z = \frac{v}{2} \sqrt{g^2 + g'^2}, \quad m_f = \frac{y_f v}{\sqrt{2}}$$

for the W , Z , and a generic fermion f (with Yukawa coupling y_f).

If we include sterile right-handed neutrinos $N_{iR} \sim (1, 0)$, the gauge invariant Dirac Yukawa term (1.2) will lead to a Dirac mass matrix after electroweak symmetry breaking

$$(m_D)_{\alpha i} = \frac{y_\nu^{\alpha i} v}{\sqrt{2}},$$

where $\alpha = e, \mu, \tau$ labels the charged-lepton flavor indices, and $i = 1, 2, 3$ labels the right-handed neutrino fields.

Obtaining the measured scale of neutrino masses within this pure Dirac framework requires an extremely small Yukawa coupling. Given that $v \simeq 246 \text{ GeV}$, the Yukawa coupling must be of the order of $y_\nu \sim 10^{-13}$, compared with $y_e \simeq 3 \times 10^{-6}$ for the electron. In other words, the neutrino Yukawa coupling would need to be about 10^7 times smaller than that of the electron, while the scale differences between the electron and the heavier leptons μ and τ of 10^2 and 10^3 , respectively, are much smaller. The scale of the neutrino Yukawa coupling is not fully understood, hence the introduction of other theories [25].

1.4.2 Majorana masses

As mentioned in Section 1.3, Majorana neutrino masses under the minimal SM are not possible, as they would violate $U(1)_Y$ gauge invariance. However, the model can be extended

by introducing two Higgs fields into the Majorana operator, with each doublet being $H \sim (2, Y = +\frac{1}{2})$, such that each lepton doublet is paired with one H in the unique dimension-5 Weinberg operator:

$$\mathcal{O}_5^{\alpha\beta} = \frac{c_{\alpha\beta}}{\Lambda} (L_\alpha^T C i\sigma^2 H) (H^T i\sigma^2 L_\beta) + \text{h.c.} \quad (1.3)$$

This new Majorana term forms an $SU(2)_L$ singlet since each term in the product uses the invariant antisymmetric tensor $i\sigma^2$. For the hypercharge, each LH pair has $Y = -\frac{1}{2} + \frac{1}{2} = 0$, so the full operator is neutral. Therefore, it is an $SU(2)_L \times U(1)_Y$ singlet. The expression is a Lorentz scalar as the spinor indices are contracted via C . However, considering the operator's mass dimension, with $[L] = 3/2$ and $[H] = 1$, each LH combination has dimension $5/2$. Two such factors give dimension 5. This term is beyond the renormalizable Standard Model as operators in the minimal SM must have dimension ≤ 4 to be renormalizable. It therefore signals new physics at a high scale Λ , where after electroweak symmetry breaking, the induced Majorana mass term is

$$(m_\nu)_{\alpha\beta} = \frac{v^2}{2\Lambda} c_{\alpha\beta}, \quad (1.4)$$

which defines a symmetric mass matrix for the active neutrinos [28]. The observed neutrino masses suggest a scale $\Lambda \sim 10^{14-15}$ GeV (for $c_{\alpha\beta} \sim 1$), pointing to heavy new states near the grand unification scale.

1.5 Seesaw mechanism

An extension of the Dirac theory would require a tiny Yukawa coupling, whereas a Majorana term would require a huge mass scale in the order of $\Lambda \sim 10^{14-15}$ GeV. A theory that adds both extensions at once is the ‘‘Seesaw Mechanism’’. In the type-1 seesaw mechanism, there is a Dirac Yukawa term that couples the active left-handed neutrino to a sterile right-

handed neutrino, where the sterile state has a Majorana mass. In the seesaw mechanism, the interplay between those two would explain the observed neutrino masses.

In the type-1 seesaw, the corresponding operator incorporates both the Yukawa interaction and the Majorana mass term,

$$\mathcal{L} \supset - y_\nu^{\alpha i} \overline{L}_\alpha \tilde{H} N_{iR} - \frac{1}{2} \overline{N_{iR}^c} (M_R)_{ij} N_{jR} + \text{h.c.} \quad (1.5)$$

for the sterile right-handed neutrinos N_R . The expression is gauge invariant. After electroweak symmetry breaking, the neutral-fermion mass matrix becomes

$$\mathcal{M} = \begin{pmatrix} 0 & m_D \\ m_D^T & M_R \end{pmatrix}, \quad (m_D)_{\alpha i} = \frac{y_\nu^{\alpha i} v}{\sqrt{2}}, \quad (1.6)$$

with M_R the Majorana mass matrix for the sterile states. In this case, if the Majorana scale is much *larger* than the Dirac scale ($M_R \gg m_D$), then the eigenvalues are approximately in the order of the observed neutrino mass:

$$m_\nu \simeq -m_D M_R^{-1} m_D^T, \quad \theta \simeq m_D M_R^{-1} \ll 1, \quad (1.7)$$

where θ is the corresponding active-sterile mixing angle. In this theory, as one side goes up, such as a larger M_R , the other side goes down (the light mass m_ν gets smaller), like a seesaw, hence the name.

This explains the small neutrino masses without requiring an extremely small Yukawa coupling. For instance, if $y_\nu \sim 0.1$ – 1 (as for other fermions), then m_D lies around 10–100 GeV, and a neutrino mass $m_\nu \simeq 0.05$ eV implies

$$M_R \sim \frac{m_D^2}{m_\nu} \sim \frac{(100 \text{ GeV})^2}{0.05 \text{ eV}} \sim 2 \times 10^{14} \text{ GeV}.$$

This mass scale is close to the typical grand-unification scale, which adds to the theoretical appeal of the seesaw mechanism [29, 30, 31, 32, 33, 25].

1.6 Physics of Neutrino Oscillations

In quantum mechanics, neutrinos are produced and detected in definite flavor states (ν_e , ν_μ , or ν_τ). However, as they propagate between their production and detection points, they exist in a superposition of these flavor states [33]. This is because the flavor states are not the energy eigenstates of the Hamiltonian; instead, the physical states that evolve in time are the mass eigenstates (ν_1 , ν_2 , and ν_3):

- a neutrino produced with a given flavor is not in a defined state, but is in fact a mixture of the three possible mass eigenstates in a certain proportion,
- when the neutrino moves relative to its point of production, the three mass eigenstates do not move at the same speed; therefore, at some point, the relative proportion of the various mass states of the neutrino will be different from that at the origin,
- when the neutrino interacts again, since it is now in a mixed state of flavor, it may later be detected as its initial flavor or as other flavors. This is the oscillation phenomenon.

This oscillation can be viewed as a rotation between the flavor and mass bases, where the following rotation matrix connects the mass eigenstates and the flavor eigenstates:

$$\underbrace{\begin{pmatrix} \nu_e \\ \nu_\mu \\ \nu_\tau \end{pmatrix}}_{\text{Flavor eigenstate}} = \begin{pmatrix} U_{e1} & U_{e2} & U_{e3} \\ U_{\mu1} & U_{\mu2} & U_{\mu3} \\ U_{\tau1} & U_{\tau2} & U_{\tau3} \end{pmatrix} \underbrace{\begin{pmatrix} \nu_1 \\ \nu_2 \\ \nu_3 \end{pmatrix}}_{\text{Mass eigenstate}} \quad (1.8)$$

where U is the PMNS (Pontecorvo-Maki-Nakagawa-Sakata) matrix [34].

1.6.1 Phase Redundancies and Physical Parameters in the PMNS Matrix (Dirac Case)

A general 3×3 unitary matrix can be described by 9 real parameters: 3 mixing angles and 6 complex phases. In the Dirac neutrino case, field rephasings can remove several of these phases through field redefinitions [33] [35].

The PMNS matrix U appears in the charged-current weak interaction:

$$\mathcal{L}_{CC} = \frac{g}{\sqrt{2}} \bar{\ell}_{\alpha L} \gamma^\mu U_{\alpha i} \nu_{iL} W_\mu^- + \text{h.c.} \quad (1.9)$$

Here, $\alpha = e, \mu, \tau$ runs over the charged-lepton flavors, and $i = 1, 2, 3$ labels the neutrino mass eigenstates.

Under field rephasings of the charged-lepton and neutrino fields,

$$\ell_{\alpha L} \rightarrow e^{i\phi_\alpha} \ell_{\alpha L}, \quad \nu_{iL} \rightarrow e^{i\psi_i} \nu_{iL}, \quad (1.10)$$

The PMNS matrix transforms as:

$$U_{\alpha i} \rightarrow e^{i\phi_\alpha} U_{\alpha i} e^{-i\psi_i}. \quad (1.11)$$

The complex phases of U depend on the choices of phase conventions for the charged-lepton and neutrino mass eigenstate fields.

In the Dirac framework, we introduce the right-handed neutrinos ν_{iR} (see Section 1.4.1), leading to the Dirac mass term:

$$\mathcal{L}_m^{\text{Dirac}} = -m \bar{\nu} \nu = -m(\bar{\nu}_L \nu_R + \bar{\nu}_R \nu_L), \quad (1.12)$$

which remains invariant under the rephasing

$$\nu_L \rightarrow e^{i\psi} \nu_L, \quad \nu_R \rightarrow e^{i\psi} \nu_R, \quad (1.13)$$

since the phase factors cancel: $\bar{\nu}_L \nu_R \rightarrow e^{-i\psi} e^{i\psi} \bar{\nu}_L \nu_R = \bar{\nu}_L \nu_R$.

The kinetic term $\bar{\nu} i \not{\partial} \nu$ is also invariant under this transformation. Consequently, each Dirac neutrino eigenstate retains an independent global $U(1)$ rephasing symmetry.

Similarly, the charged-lepton mass term

$$\mathcal{L}_m^\ell = -m_\ell \bar{\ell}_L \ell_R + \text{h.c.} \quad (1.14)$$

is invariant under $\ell_\alpha \rightarrow e^{i\phi_\alpha} \ell_\alpha$, so each charged-lepton flavor also has an independent $U(1)$ rephasing symmetry.

As there are three charged leptons and three neutrinos, there should be six independent rephasing freedoms. However, the overall phase transformation,

$$\ell_{\alpha L} \rightarrow e^{i\chi} \ell_{\alpha L}, \quad \nu_{iL} \rightarrow e^{i\chi} \nu_{iL}, \quad (1.15)$$

leaves the matrix U completely unchanged:

$$U_{\alpha i} \rightarrow e^{i\chi} U_{\alpha i} e^{-i\chi} = U_{\alpha i}.$$

Thus, only $6 - 1 = 5$ of these phase freedoms are independent. These five rephasings can be used to eliminate five of the six complex phases in a general unitary 3×3 matrix.

The remaining phase is physically meaningful as it cannot be removed by rephasing. It corresponds to the Dirac CP-violating phase δ_{CP} .

It follows that, for Dirac neutrinos, 5 phases can be absorbed. This leaves a total of 4

physical parameters in the PMNS matrix, 3 mixing angles and 1 physical CP-violating Dirac phase.

The PMNS matrix can then be expressed as a product of three matrices :

$$U = \begin{pmatrix} 1 & 0 & 0 \\ 0 & \cos \theta_{23} & \sin \theta_{23} \\ 0 & -\sin \theta_{23} & \cos \theta_{23} \end{pmatrix} \times \begin{pmatrix} \cos \theta_{13} & 0 & \sin \theta_{13} e^{-i\delta_{CP}} \\ 0 & 1 & 0 \\ -\sin \theta_{13} e^{i\delta_{CP}} & 0 & \cos \theta_{13} \end{pmatrix} \\ \times \begin{pmatrix} \cos \theta_{12} & \sin \theta_{12} & 0 \\ -\sin \theta_{12} & \cos \theta_{12} & 0 \\ 0 & 0 & 1 \end{pmatrix}$$

Atmospheric Cross-mixing Solar

where δ_{CP} is the CP-violating phase.

1.6.2 Phase Counting in the PMNS Matrix: Majorana Case

If neutrinos are Majorana particles, their mass can be written as

$$\mathcal{L}_m^{\text{Majorana}} = -\frac{1}{2} \nu_L^T C^{-1} M \nu_L + \text{h.c.}, \quad (1.16)$$

Here, M denotes the complex symmetric matrix and C the charge-conjugation matrix. It violates lepton number by two units and treats neutrinos as their own antiparticles. It is not invariant under a general rephasing of the neutrino fields:

$$\nu_{iL} \rightarrow e^{i\psi_i} \nu_{iL}, \quad (1.17)$$

since the mass term transforms as

$$\nu_L^T C^{-1} M \nu_L \rightarrow e^{2i\psi_i} \nu_L^T C^{-1} M \nu_L,$$

which modifies the Lagrangian unless $e^{2i\psi_i} = 1$, i.e., $\psi_i = 0$ or π . Consequently, unlike in the Dirac case, the neutrino fields in the Majorana case do not have a continuous $U(1)$ rephasing symmetry.

The charged-lepton mass term remains unchanged:

$$\mathcal{L}_m^\ell = -m_\ell \bar{\ell}_L \ell_R + \text{h.c.}, \quad (1.18)$$

so the three charged-lepton fields $\ell_{\alpha L}$ may still be rephased:

$$\ell_{\alpha L} \rightarrow e^{i\phi_\alpha} \ell_{\alpha L},$$

This affects the rows of the PMNS matrix.

Since the neutrino fields cannot be rephased, only the charged-lepton rephasings can be used to absorb phases from the PMNS matrix:

$$U_{\alpha i} \rightarrow e^{i\phi_\alpha} U_{\alpha i}. \quad (1.19)$$

Out of the six complex phases in a general unitary 3×3 matrix, only three can be removed. This leaves

$$6 - 3 = 3$$

physical phases in the PMNS matrix.

These include:

- one Dirac CP-violating phase δ_{CP} , which affects oscillation probabilities, and
- two additional Majorana phases α_{21} and α_{31} , which do not affect oscillations but enter lepton-number-violating processes such as neutrinoless double-beta decay.

For Majorana neutrinos, only three charged-lepton rephasings can remove phases. In this

case, the PMNS matrix has three mixing angles and three physical CP-violating phases.

1.7 Neutrino Oscillations in Vacuum

Neutrino oscillations result from the fact that the flavor eigenstates $\{\nu_\alpha\}$ in weak interactions are not the same as the mass eigenstates $\{\nu_k\}$. The mass eigenstates diagonalize the Hamiltonian which governs the neutrino propagation through spacetime. In charged-current interactions, a neutrino of flavor $\alpha \in \{e, \mu, \tau\}$ can be represented as a coherent superposition of mass eigenstates [33]:

$$|\nu_\alpha\rangle = \sum_k U_{\alpha k}^* |\nu_k\rangle, \quad (1.20)$$

where U is the PMNS matrix and $|\nu_k\rangle$ is an eigenstate of the vacuum Hamiltonian H_0 , i.e.

$$H_0 |\nu_k\rangle = E_k |\nu_k\rangle, \quad E_k = \sqrt{p^2 + m_k^2}. \quad (1.21)$$

In the Schrödinger picture, neutrino time evolution obeys:

$$i \frac{d}{dt} |\nu_\alpha(t)\rangle = H_0 |\nu_\alpha(t)\rangle, \quad |\nu_\alpha(0)\rangle = |\nu_\alpha\rangle. \quad (1.22)$$

Since the Hamiltonian is diagonal in the mass basis, each $|\nu_k\rangle$ propagates freely:

$$|\nu_k(t)\rangle = e^{-iE_k t} |\nu_k\rangle. \quad (1.23)$$

1.7.1 Transition Amplitude and Probability

To calculate the probability that a neutrino produced in the flavor state ν_α is later detected as ν_β , we consider the time evolution of the neutrino state $|\nu_\alpha(t)\rangle$:

$$|\nu_\alpha(t)\rangle = \sum_k U_{\alpha k}^* e^{-iE_k t} |\nu_k\rangle. \quad (1.24)$$

Projecting onto the flavor state $|\nu_\beta\rangle$ yields the transition amplitude:

$$A_{\nu_\alpha \rightarrow \nu_\beta}(t) = \langle \nu_\beta | \nu_\alpha(t) \rangle = \sum_k U_{\alpha k}^* U_{\beta k} e^{-iE_k t}. \quad (1.25)$$

Squaring the amplitude gives the transition probability:

$$P_{\nu_\alpha \rightarrow \nu_\beta}(t) = |A_{\nu_\alpha \rightarrow \nu_\beta}(t)|^2 = \sum_{k,j} (U_{\alpha k}^* U_{\beta k} e^{-iE_k t}) (U_{\alpha j}^* U_{\beta j} e^{-iE_j t})^* \quad (1.26)$$

We obtain the expression for the oscillation probability:

$$P_{\nu_\alpha \rightarrow \nu_\beta}(t) = \sum_{k,j} U_{\alpha k}^* U_{\beta k} U_{\alpha j} U_{\beta j}^* e^{-i(E_k - E_j)t}. \quad (1.27)$$

1.7.2 Ultra-Relativistic Limit

In most experiments, neutrinos are highly relativistic particles ($E_k \gg m_k$). In this regime, we use the following relativistic energy relation for each mass eigenstate [33, 36]:

$$E_k = \sqrt{p^2 + m_k^2} \simeq p + \frac{m_k^2}{2p}, \quad (1.28)$$

Here, the binomial expansion $\sqrt{1+x} \simeq 1 + x/2$ valid for small $x \equiv m_k^2/p^2 \ll 1$ has been used.

Since all mass eigenstates are produced with nearly the same momentum, we define $E \equiv p$ as a common reference energy, and E_k becomes:

$$E_k \simeq E + \frac{m_k^2}{2E}, \quad (1.29)$$

Hence, the phase difference between two mass eigenstates ν_k and ν_j becomes:

$$E_k - E_j \simeq \frac{m_k^2 - m_j^2}{2E} = \frac{\Delta m_{kj}^2}{2E}. \quad (1.30)$$

Furthermore, in natural units where $c = 1$, the neutrino propagation time t is approximately equal to the distance L traveled from source to detector, for an ultra-relativistic neutrino:

$$t \approx L. \quad (1.31)$$

Substituting both approximations into the transition amplitude leads to the expression of the oscillation probability [33]:

$$P_{\nu_\alpha \rightarrow \nu_\beta}(L, E) = \sum_{k,j} U_{\alpha k}^* U_{\beta k} U_{\alpha j} U_{\beta j}^* \exp\left(-i \frac{\Delta m_{kj}^2 L}{2E}\right). \quad (1.32)$$

We can rearrange the oscillation probability in Eq. 1.32, such that the real and imaginary terms are isolated, as well as the terms sensitive to CP violation. We do this by rewriting the double sum by pairing terms with indices (k, j) and (j, k) .

The dimensionless oscillation phase is defined as:

$$\Delta_{kj} := \frac{\Delta m_{kj}^2 L}{4E}. \quad (1.33)$$

The oscillation probability then becomes:

$$P_{\nu_\alpha \rightarrow \nu_\beta} = \sum_{k,j} U_{\alpha k}^* U_{\beta k} U_{\alpha j} U_{\beta j}^* \exp(-2i\Delta_{kj}). \quad (1.34)$$

In the case where $k \neq j$, we complex conjugate the terms with (k, j) , to obtain the (j, k) terms:

$$U_{\alpha j}^* U_{\beta j} U_{\alpha k} U_{\beta k}^* e^{+2i\Delta_{kj}} = (U_{\alpha k}^* U_{\beta k} U_{\alpha j} U_{\beta j}^* e^{-2i\Delta_{kj}})^*.$$

Adding this with the terms with $k = j$:

$$\begin{aligned} U_{\alpha k}^* U_{\beta k} U_{\alpha j} U_{\beta j}^* e^{-2i\Delta_{kj}} + \text{c.c.} &= U_{\alpha k}^* U_{\beta k} U_{\alpha j} U_{\beta j}^* e^{-2i\Delta_{kj}} \\ &+ U_{\alpha j}^* U_{\beta j} U_{\alpha k} U_{\beta k}^* e^{+2i\Delta_{kj}} \\ &= 2 \operatorname{Re} [U_{\alpha k}^* U_{\beta k} U_{\alpha j} U_{\beta j}^* e^{-2i\Delta_{kj}}]. \end{aligned} \quad (1.35)$$

The exponential can be rewritten as:

$$e^{-2i\Delta_{kj}} = \cos(2\Delta_{kj}) - i \sin(2\Delta_{kj}),$$

where the expression can be separated into real and imaginary parts:

$$\operatorname{Re} [A e^{-2i\Delta_{kj}}] = \operatorname{Re}(A) \cos(2\Delta_{kj}) + \operatorname{Im}(A) \sin(2\Delta_{kj}),$$

where $A := U_{\alpha k}^* U_{\beta k} U_{\alpha j} U_{\beta j}^*$.

So each symmetrized term gives:

$$2 \operatorname{Re} [A e^{-2i\Delta_{kj}}] = 2 \operatorname{Re}(A) \cos(2\Delta_{kj}) + 2 \operatorname{Im}(A) \sin(2\Delta_{kj}).$$

Using $\cos(2\Delta_{kj}) = 1 - 2\sin^2(\Delta_{kj})$, the real part becomes:

$$2 \operatorname{Re}[A e^{-2i\Delta_{kj}}] = 2 \operatorname{Re}(A) \cos(2\Delta_{kj}) + 2 \operatorname{Im}(A) \sin(2\Delta_{kj}) \quad (1.36)$$

$$= 2 \operatorname{Re}(A) - 4 \operatorname{Re}(A) \sin^2(\Delta_{kj}) + 2 \operatorname{Im}(A) \sin(2\Delta_{kj}). \quad (1.37)$$

The constant terms $2 \operatorname{Re}(A)$ from all pairs (k, j) combine with the diagonal contributions ($k = j$) of the double sum:

$$\sum_k |U_{\alpha k}|^2 |U_{\beta k}|^2 + 2 \sum_{k>j} \operatorname{Re}[U_{\alpha k}^* U_{\beta k} U_{\alpha j} U_{\beta j}^*] = \sum_{k,j} \operatorname{Re}[U_{\alpha k}^* U_{\beta k} U_{\alpha j} U_{\beta j}^*] \quad (1.38)$$

$$= \operatorname{Re} \left[\left(\sum_k U_{\alpha k}^* U_{\beta k} \right) \left(\sum_j U_{\alpha j} U_{\beta j}^* \right) \right] \quad (1.39)$$

$$= \operatorname{Re}[\delta_{\alpha\beta} \delta_{\alpha\beta}] = \delta_{\alpha\beta}. \quad (1.40)$$

Thus, the constant part of the expansion reproduces the normalization $\delta_{\alpha\beta}$ required by probability conservation.

Summing over all $k > j$ (to avoid double-counting symmetric terms), the full probability becomes:

$$\begin{aligned} P_{\nu_\alpha \rightarrow \nu_\beta} &= \delta_{\alpha\beta} - 4 \sum_{k>j} \operatorname{Re}[U_{\alpha k}^* U_{\beta k} U_{\alpha j} U_{\beta j}^*] \sin^2(\Delta_{kj}) \\ &\quad + 2 \sum_{k>j} \operatorname{Im}[U_{\alpha k}^* U_{\beta k} U_{\alpha j} U_{\beta j}^*] \sin(2\Delta_{kj}). \end{aligned} \quad (1.41)$$

For antineutrinos, the PMNS matrix is replaced by its complex conjugate $U \rightarrow U^*$, the

oscillation probability is then:

$$\begin{aligned}
P_{\bar{\nu}_\alpha \rightarrow \bar{\nu}_\beta} = & \delta_{\alpha\beta} - 4 \sum_{k>j} \text{Re}[U_{\alpha k}^* U_{\beta k} U_{\alpha j} U_{\beta j}^*] \sin^2 \Delta_{kj} \\
& - 2 \sum_{k>j} \text{Im}[U_{\alpha k}^* U_{\beta k} U_{\alpha j} U_{\beta j}^*] \sin(2\Delta_{kj}).
\end{aligned} \tag{1.42}$$

We reverse the sign of the imaginary term to account for the complex conjugation of the PMNS matrix under CP. This step exchanges neutrinos with their corresponding antineutrinos.

1.7.3 CP and T Symmetries in Neutrino Oscillations

In the three-flavor neutrino oscillations framework, discrete symmetries such as charge-parity (CP), time reversal (T), and the combined CPT are useful in the interpretation of oscillation probabilities.

CP symmetry turns each neutrino into its corresponding antineutrino and, at the same time, flips spatial coordinates. In flavor transitions:

$$P_{\nu_\alpha \rightarrow \nu_\beta}(L) \longrightarrow P_{\bar{\nu}_\alpha \rightarrow \bar{\nu}_\beta}(L),$$

which corresponds to complex conjugating all elements of the PMNS matrix: $U \rightarrow U^*$. The T symmetry, on the other hand, reverses the direction of time and exchanges initial and final states:

$$P_{\nu_\alpha \rightarrow \nu_\beta}(L) \longrightarrow P_{\nu_\beta \rightarrow \nu_\alpha}(L).$$

In the general vacuum oscillation probability,

$$\begin{aligned}
P_{\nu_\alpha \rightarrow \nu_\beta} = & \delta_{\alpha\beta} - 4 \sum_{k>j} \text{Re}[U_{\alpha k}^* U_{\beta k} U_{\alpha j} U_{\beta j}^*] \sin^2 \Delta_{kj} \\
& + 2 \sum_{k>j} \text{Im}[U_{\alpha k}^* U_{\beta k} U_{\alpha j} U_{\beta j}^*] \sin(2\Delta_{kj}),
\end{aligned} \tag{1.43}$$

The second term involving the real parts is invariant under both CP and T transformations. It remains unchanged when complex conjugating the matrix elements or swapping initial/-final flavors, and is thus called CP-even and T-even [33]. The third term, however, depends on the imaginary part of rephasing-invariant combinations of PMNS matrix elements:

$$\text{Im}[U_{\alpha k}^* U_{\beta k} U_{\alpha j} U_{\beta j}^*].$$

It is sensitive to the Dirac CP phase δ , which enters the PMNS matrix through complex phases. If $\delta = 0$ or π , all PMNS elements can be made real via rephasing, and this imaginary part vanishes. Thus, it is the only component of the oscillation probability sensitive to CP violation:

$$\delta = 0, \pi \quad \Rightarrow \quad \text{Im}[U_{\alpha k}^* U_{\beta k} U_{\alpha j} U_{\beta j}^*] = 0.$$

By construction, this third term changes sign under both CP and T operations. It is CP-odd and T-odd, and disappears if the CP symmetry is unbroken. Any observed difference between:

$$P_{\nu_\mu \rightarrow \nu_e} \quad \text{and} \quad P_{\bar{\nu}_\mu \rightarrow \bar{\nu}_e}$$

(or their T-reversed counterparts) would provide direct evidence for CP violation in the lepton sector [33]. By the CPT theorem, if CP is violated, this implies that T is violated as

well. This is reflected in the symmetry relation:

$$P_{\nu_\alpha \rightarrow \nu_\beta} \neq P_{\bar{\nu}_\alpha \rightarrow \bar{\nu}_\beta} \quad \Rightarrow \quad P_{\nu_\alpha \rightarrow \nu_\beta} \neq P_{\nu_\beta \rightarrow \nu_\alpha}.$$

These asymmetries can be evaluated experimentally, for example, by comparing appearance probabilities of $\nu_\mu \rightarrow \nu_e$ and $\bar{\nu}_\mu \rightarrow \bar{\nu}_e$.

1.7.4 Majorana Phases and Oscillation Probabilities

For Majorana neutrinos, the PMNS matrix includes two additional complex phases [33] [35]:

$$U_{\text{Majorana}} = U_{\text{PMNS}} \cdot \text{diag}(1, e^{i\alpha_1/2}, e^{i\alpha_2/2}).$$

This affects only the elements in the second and third columns of the mixing matrix:

$$U_{\alpha 1}^{\text{Maj}} = U_{\alpha 1}, \tag{1.44}$$

$$U_{\alpha 2}^{\text{Maj}} = U_{\alpha 2} e^{i\alpha_1/2} \tag{1.45}$$

$$U_{\alpha 3}^{\text{Maj}} = U_{\alpha 3} e^{i\alpha_2/2} \tag{1.46}$$

However, the oscillation probability involves the product $U_{\beta k} U_{\alpha k}^*$, which becomes:

$$U_{\beta k} U_{\alpha k}^* \rightarrow (U_{\beta k} e^{i\alpha_k/2})(U_{\alpha k}^* e^{-i\alpha_k/2}) = U_{\beta k} U_{\alpha k}^*.$$

Thus, the Majorana phases cancel in the oscillation amplitude and have no effect on the oscillation probability:

$$P_{\nu_\alpha \rightarrow \nu_\beta} = \left| \sum_k U_{\beta k} U_{\alpha k}^* e^{-iE_k t} \right|^2,$$

unlike the Dirac CP phase δ , which appears in rephasing-invariant combinations such as:

$$\text{Im} \left[U_{\alpha k}^* U_{\beta k} U_{\alpha j} U_{\beta j}^* \right],$$

and survives in the CP-odd terms of the oscillation probability, leading to observable CP-violating effects.

1.7.5 Two-Flavor Oscillation Approximation

Even though neutrino oscillations are fundamentally a three-flavor phenomenon, many experimental situations can be described to good accuracy with a two-flavor approximation. This can be particularly useful, especially when one mass-squared difference is dominating and certain mixing angles are small.

Derivation from the Schrödinger Equation

Considering two flavor states, $|\nu_\alpha\rangle$ and $|\nu_\beta\rangle$, that are related to two mass eigenstates $|\nu_1\rangle$ and $|\nu_2\rangle$ via a single mixing angle θ :

$$\begin{pmatrix} |\nu_\alpha\rangle \\ |\nu_\beta\rangle \end{pmatrix} = \begin{pmatrix} \cos \theta & \sin \theta \\ -\sin \theta & \cos \theta \end{pmatrix} \begin{pmatrix} |\nu_1\rangle \\ |\nu_2\rangle \end{pmatrix}.$$

In the mass basis, the Hamiltonian is diagonal, and the time evolution of each eigenstate is:

$$|\nu_k(t)\rangle = e^{-iE_k t} |\nu_k\rangle, \quad k = 1, 2.$$

Starting from an initial flavor state $|\nu_\alpha\rangle$ at $t = 0$, the time-evolved state is:

$$|\nu_\alpha(t)\rangle = \cos \theta e^{-iE_1 t} |\nu_1\rangle + \sin \theta e^{-iE_2 t} |\nu_2\rangle.$$

Projecting onto $|\nu_\beta\rangle$ gives the amplitude:

$$A_{\nu_\alpha \rightarrow \nu_\beta}(t) = \langle \nu_\beta | \nu_\alpha(t) \rangle = \sin \theta \cos \theta (e^{-iE_2 t} - e^{-iE_1 t}).$$

Squaring the amplitude gives the probability:

$$P_{\nu_\alpha \rightarrow \nu_\beta}(t) = |A_{\nu_\alpha \rightarrow \nu_\beta}(t)|^2 \tag{1.47}$$

$$= \sin^2(2\theta) \sin^2\left(\frac{E_2 - E_1}{2}t\right). \tag{1.48}$$

For ultra-relativistic neutrinos, where $E_k \simeq E + m_k^2/(2E)$:

$$E_2 - E_1 \simeq \frac{\Delta m^2}{2E}, \quad \Delta m^2 := m_2^2 - m_1^2,$$

Replacing $t \approx L$ in natural units yields the standard two-flavor probability [35]:

$$P_{\nu_\alpha \rightarrow \nu_\beta}(L, E) = \sin^2(2\theta) \sin^2\left(\frac{\Delta m^2 L}{4E}\right). \tag{1.49}$$

Physical Interpretation and Limitations

The two-flavor formula shows some important features of neutrino oscillations:

- The amplitude of the oscillation is controlled by the mixing angle θ .
- The oscillation phase depends on the mass-squared difference Δm^2 and the ratio L/E .

This approximation is valid only when a single mass-squared difference is dominating, for instance:

- **For solar neutrinos:** Dominated by Δm_{21}^2 and θ_{12} . The larger splitting Δm_{31}^2 averages out over the long solar baseline.

- **For short-baseline reactor experiments (e.g., Daya Bay):** Dominated by Δm_{31}^2 and θ_{13} , while the smaller Δm_{21}^2 effect is negligible.

1.8 Neutrino Oscillations in Matter

When neutrinos propagate through matter, they can interact with the particles in the medium, which can modify their oscillation probabilities. Neutrinos can undergo coherent forward elastic scattering with electrons, protons, and neutrons. When passing through matter, neutrinos exchange weak bosons with the medium. Even though these interactions do not change the neutrino's momentum, each one of them contributes a tiny, well-defined phase to the neutrino wave function. These small phase shifts add up to produce an effective potential that modifies the propagation of each flavor state [22] [33].

Two types of weak interactions contribute to this potential:

- the **charged-current (CC)** forward scattering process, mediated by the exchange of a W boson; and
- the **neutral-current (NC)** forward scattering process, mediated by the exchange of a Z boson (Figure 1.4).




Figure 1.4: Feynman diagrams of the coherent forward elastic scattering processes that generate the CC potential V_{CC} through W exchange and the NC potential V_{NC} through Z exchange.

Electron neutrinos experience both charged-current and neutral-current forward scattering. On the other hand, as ordinary matter contains no muons or taus, muon and tau neutrinos only undergo neutral-current scattering. The effective potential generated by charged-current scattering of ν_e is given by [33]:

$$V_{CC} = \sqrt{2} G_F N_e \quad (1.50)$$

where G_F is the Fermi coupling constant and N_e is the local number density of electrons in the medium.

Similarly, the neutral-current potential from coherent interactions with a background fermion f (electron, proton, or neutron) is the following where N_f is [33]:

$$V_{NC}^f = \sqrt{2} G_F N_f g_V^f \quad (1.51)$$

where N_f is the number density of the fermion f , and g_V^f is its weak vector coupling to the Z boson.

The total matter potential combines the charged-current (CC) and neutral-current (NC) contributions [33]. The effective potential acts diagonally in the flavor basis as the coherent forward scattering does not change the neutrino flavor:

$$V_{\alpha\beta} = \delta_{\alpha\beta} V_\alpha, \quad V_\alpha = V_{NC} + \delta_{\alpha e} V_{CC}. \quad (1.52)$$

Each flavor experiences a different effective potential:

$$V_e = V_{CC} + V_{NC}, \quad V_\mu = V_{NC}, \quad V_\tau = V_{NC}. \quad (1.53)$$

Explicitly, the matter potential matrix in the flavor basis takes the form:

$$V = \begin{pmatrix} V_e & 0 & 0 \\ 0 & V_\mu & 0 \\ 0 & 0 & V_\tau \end{pmatrix} = \begin{pmatrix} V_{CC} + V_{NC} & 0 & 0 \\ 0 & V_{NC} & 0 \\ 0 & 0 & V_{NC} \end{pmatrix} = V_{NC} I + \begin{pmatrix} V_{CC} & 0 & 0 \\ 0 & 0 & 0 \\ 0 & 0 & 0 \end{pmatrix}. \quad (1.54)$$

The neutral-current contribution is the same for all active flavors so it does not affect the transition probabilities. Only the charged-current potential acting on the ν_e modifies the neutrino evolution, which leads to an effective matter potential that modifies the oscillation probability in matter, which is known as the Mikheyev–Smirnov–Wolfenstein (MSW) effect [22].

1.8.1 Probability of Oscillation

From Section 1.7, a neutrino of flavor $\alpha \in \{e, \mu, \tau\}$ and momentum $\vec{p}$ can be written as a superposition of mass eigenstates:

$$|\nu_\alpha\rangle = \sum_k U_{\alpha k}^* |\nu_k\rangle, \quad (1.55)$$

where each $|\nu_k\rangle$ is an eigenstate of the vacuum Hamiltonian H_0 ,

$$H_0 |\nu_k\rangle = E_k |\nu_k\rangle, \quad E_k = \sqrt{\vec{p}^2 + m_k^2}. \quad (1.56)$$

Coherent forward scattering in matter adds an effective potential term to the Hamiltonian:

$$H = H_0 + H_I, \quad H_I |\nu_\alpha\rangle = V_\alpha |\nu_\alpha\rangle, \quad (1.57)$$

where V_α represents the matter-induced potential for each flavor.

In the Schrödinger framework, the flavor state evolves according to

$$i \frac{d}{dt} |\nu_\alpha(t)\rangle = H |\nu_\alpha(t)\rangle, \quad |\nu_\alpha(0)\rangle = |\nu_\alpha\rangle. \quad (1.58)$$

Projecting this time-evolved state onto another flavor basis state defines the transition amplitudes:

$$\psi_{\alpha\beta}(t) = \langle \nu_\beta | \nu_\alpha(t) \rangle, \quad \psi_{\alpha\beta}(0) = \delta_{\alpha\beta}. \quad (1.59)$$

With the corresponding transition probability

$$P_{\nu_\alpha \rightarrow \nu_\beta}(t) = |\psi_{\alpha\beta}(t)|^2. \quad (1.60)$$

Differentiating with respect to time and inserting the completeness relation $\sum_\eta |\nu_\eta\rangle \langle \nu_\eta| = I$, we obtain

$$i \frac{d}{dt} \psi_{\alpha\beta}(t) = \sum_\eta \langle \nu_\beta | H | \nu_\eta \rangle \psi_{\alpha\eta}(t). \quad (1.61)$$

Since $H = H_0 + H_I$, the matrix elements can be expressed as

$$\langle \nu_\beta | H_0 | \nu_\eta \rangle = \sum_k U_{\beta k} E_k U_{\eta k}^*, \quad \langle \nu_\beta | H_I | \nu_\eta \rangle = \delta_{\beta\eta} V_\beta, \quad (1.62)$$

where E_k are the mass eigenstate energies and V_β the matter potential for flavor ν_β . Substituting these gives the evolution equation (corresponding to Eq. (9.47)):

$$i \frac{d}{dt} \psi_{\alpha\beta}(t) = \sum_\eta \left(\sum_k U_{\beta k} E_k U_{\eta k}^* + \delta_{\beta\eta} V_\beta \right) \psi_{\alpha\eta}(t). \quad (1.63)$$

The Hamiltonian remains hermitian after the addition of the interaction term H_I . Since both H_0 and H_I are Hermitian, the total Hamiltonian $H = H_0 + H_I$ is also Hermitian.

Therefore, the evolution is still unitary:

$$U(t) = e^{-iHt}, \quad U^\dagger(t)U(t) = I,$$

which leads to the conservation of the total probability:

$$\sum_{\beta} P_{\nu_{\alpha} \rightarrow \nu_{\beta}}(t) = 1 \quad (1.64)$$

For ultrarelativistic neutrinos, the following approximations can be used:

$$E_k \simeq E + \frac{m_k^2}{2E}, \quad p \simeq E, \quad t \simeq x, \quad (1.65)$$

allowing the evolution to be written in terms of the propagation distance x .

Substituting Eq. (1.52) into the evolution equation, we obtain:

$$i \frac{d}{dt} \psi_{\alpha\beta}(t) = \sum_{\eta} \left(\sum_k U_{\beta k} E_k U_{\eta k}^* \right) \psi_{\alpha\eta}(t) + V_{\beta} \psi_{\alpha\beta}(t), \quad (1.66)$$

where $V_{\beta} = V_{NC} + \delta_{\beta e} V_{CC}$. The potential only contributes to the term with $\eta = \beta$ as it is diagonal in the flavor basis. This removes the second summation. The neutral-current part V_{NC} is identical for all active flavors. It can be factored out as a global phase. Only the charged-current contribution V_{CC} affects the flavor evolution [33].

Using the ultrarelativistic approximation $E_k \simeq p + \frac{m_k^2}{2E}$ and the unitarity relation $\sum_k U_{\beta k} U_{\eta k}^* = \delta_{\beta\eta}$, gives

$$\sum_k U_{\beta k} E_k U_{\eta k}^* = \left(p + \frac{m_1^2}{2E} \right) \delta_{\beta\eta} + \sum_k U_{\beta k} \frac{\Delta m_{k1}^2}{2E} U_{\eta k}^*, \quad (1.67)$$

where $\Delta m_{k1}^2 = m_k^2 - m_1^2$.

We substitute $V_{\beta} = V_{NC} + \delta_{\beta e} V_{CC}$, the evolution equation becomes

$$i \frac{d}{dt} \psi_{\alpha\beta}(t) = \left(p + \frac{m_1^2}{2E} + V_{NC} \right) \psi_{\alpha\beta}(t) + \sum_{\eta} \left(\sum_k U_{\beta k} \frac{\Delta m_{k1}^2}{2E} U_{\eta k}^* + \delta_{\beta e} \delta_{\eta e} V_{CC} \right) \psi_{\alpha\eta}(t). \quad (1.68)$$

Finally, using $t \simeq x$ for ultrarelativistic neutrinos yields

$$i \frac{d}{dx} \psi_{\alpha\beta}(x) = \left(p + \frac{m_1^2}{2E} + V_{NC} \right) \psi_{\alpha\beta}(x) + \sum_{\eta} \left(\sum_k U_{\beta k} \frac{\Delta m_{k1}^2}{2E} U_{\eta k}^* + \delta_{\beta e} \delta_{\eta e} V_{CC} \right) \psi_{\alpha\eta}(x). \quad (1.69)$$

The overall phase term $(p + m_1^2/2E + V_{NC})$ can be factored out as it does not affect the oscillation probabilities. The effective flavor evolution equation is then:

$$i \frac{d}{dx} \psi_{\alpha\beta}(x) = \sum_{\eta} \left(\sum_k U_{\beta k} \frac{\Delta m_{k1}^2}{2E} U_{\eta k}^* + \delta_{\beta e} \delta_{\eta e} V_{CC} \right) \psi_{\alpha\eta}(x). \quad (1.70)$$

This can be expressed in compact matrix form:

$$i \frac{d}{dx} \Psi_{\alpha} = H_F \Psi_{\alpha}, \quad (1.71)$$

with the effective flavor-basis Hamiltonian:

$$H_F = \frac{1}{2E} (U M^2 U^{\dagger} + A). \quad (1.72)$$

For the three-flavor case, these matrices are:

$$\Psi_{\alpha} = \begin{pmatrix} \psi_{\alpha e} \\ \psi_{\alpha \mu} \\ \psi_{\alpha \tau} \end{pmatrix}, \quad M^2 = \begin{pmatrix} 0 & 0 & 0 \\ 0 & \Delta m_{21}^2 & 0 \\ 0 & 0 & \Delta m_{31}^2 \end{pmatrix}, \quad A = \begin{pmatrix} A_{CC} & 0 & 0 \\ 0 & 0 & 0 \\ 0 & 0 & 0 \end{pmatrix}. \quad (1.73)$$

where the matter potential term is defined as [33]:

$$A_{CC} = 2EV_{CC} = 2\sqrt{2}EG_F N_e. \quad (1.74)$$

1.8.2 The MSW Effect in the two-flavor approximation

To illustrate the physical implications of the matter potential, its effect can be analyzed in a simplified two-flavor system involving ν_e and ν_μ . The corresponding flavor-basis Hamiltonian is given by:

$$H_F = \frac{1}{2E} (U M^2 U^\dagger + A), \quad (1.75)$$

where, in the two-flavor (ν_e, ν_μ) system,

$$U = \begin{pmatrix} \cos \theta & \sin \theta \\ -\sin \theta & \cos \theta \end{pmatrix}, \quad M^2 = \text{diag}(m_1^2, m_2^2), \quad A = \begin{pmatrix} A_{CC} & 0 \\ 0 & 0 \end{pmatrix}. \quad (1.76)$$

The vacuum expression is then:

$$U M^2 U^\dagger = \begin{pmatrix} m_1^2 \cos^2 \theta + m_2^2 \sin^2 \theta & (m_2^2 - m_1^2) \sin \theta \cos \theta \\ (m_2^2 - m_1^2) \sin \theta \cos \theta & m_1^2 \sin^2 \theta + m_2^2 \cos^2 \theta \end{pmatrix}. \quad (1.77)$$

The trace term $\frac{1}{2}(m_1^2 + m_2^2)I$ can be subtracted from the Hamiltonian without changing the physics, as it just introduces a common phase factor in the evolution operator. Therefore, it has no consequence on the oscillation probabilities. This leads to the simplified expression:

$$U M^2 U^\dagger - \frac{m_1^2 + m_2^2}{2} I = \frac{\Delta m^2}{2} \begin{pmatrix} -\cos 2\theta & \sin 2\theta \\ \sin 2\theta & \cos 2\theta \end{pmatrix}. \quad (1.78)$$

where $\Delta m^2 \equiv m_2^2 - m_1^2$.

Including the matter potential and factoring out $1/(4E)$, we obtain:

$$H_F = \frac{1}{2E} \left[\frac{\Delta m^2}{2} \begin{pmatrix} -\cos 2\theta & \sin 2\theta \\ \sin 2\theta & \cos 2\theta \end{pmatrix} + \begin{pmatrix} A_{CC} & 0 \\ 0 & 0 \end{pmatrix} \right] \quad (1.79)$$

$$= \frac{1}{4E} \begin{pmatrix} -\Delta m^2 \cos 2\theta + 2A_{CC} & \Delta m^2 \sin 2\theta \\ \Delta m^2 \sin 2\theta & \Delta m^2 \cos 2\theta \end{pmatrix}. \quad (1.80)$$

Here also, the trace term can be removed, leading to the modified Hamiltonian:

$$H'_F = H_F - \frac{1}{4E} \begin{pmatrix} A_{CC} & 0 \\ 0 & A_{CC} \end{pmatrix}. \quad (1.81)$$

The resulting effective Hamiltonian matrix becomes

$$H'_F = \frac{1}{4E} \begin{pmatrix} -\Delta m^2 \cos 2\theta + A_{CC} & \Delta m^2 \sin 2\theta \\ \Delta m^2 \sin 2\theta & \Delta m^2 \cos 2\theta - A_{CC} \end{pmatrix}, \quad (1.82)$$

which leads to the evolution equation:

$$i \frac{d}{dx} \begin{pmatrix} \psi_{ee} \\ \psi_{e\mu} \end{pmatrix} = \frac{1}{4E} \begin{pmatrix} -\Delta m^2 \cos 2\theta + A_{CC} & \Delta m^2 \sin 2\theta \\ \Delta m^2 \sin 2\theta & \Delta m^2 \cos 2\theta - A_{CC} \end{pmatrix} \begin{pmatrix} \psi_{ee} \\ \psi_{e\mu} \end{pmatrix}. \quad (1.83)$$

To solve for the evolution equation, the Hamiltonian must be diagonal in the considered basis. H_F is not diagonal in the flavor basis. However, we can perform a rotation that transforms H_F into a diagonal form. As H_F is a real symmetric 2×2 matrix, we define an orthogonal rotation of angle θ_M , such that the off-diagonal elements vanish in the rotated

basis. This rotation is represented by the matrix

$$U_M = \begin{pmatrix} \cos \theta_M & \sin \theta_M \\ -\sin \theta_M & \cos \theta_M \end{pmatrix}. \quad (1.84)$$

Requiring that the off-diagonal elements are zero in the rotated basis leads to :

$$\tan 2\theta_M = \frac{2H_{12}}{H_{22} - H_{11}} = \frac{\Delta m^2 \sin 2\theta}{\Delta m^2 \cos 2\theta - A_{CC}}. \quad (1.85)$$

Diagonalizing the two-flavor Hamiltonian yields two effective eigenvalues,

$$\lambda_{\pm} = \frac{1}{4E} \left[\pm \sqrt{(\Delta m^2 \cos 2\theta - A_{CC})^2 + (\Delta m^2 \sin 2\theta)^2} \right], \quad (1.86)$$

corresponding to the propagation eigenstates in matter.

Their difference defines the effective mass-squared splitting

$$\Delta m_M^2 = 2E(\lambda_+ - \lambda_-) = \sqrt{(\Delta m^2 \cos 2\theta - A_{CC})^2 + (\Delta m^2 \sin 2\theta)^2}. \quad (1.87)$$

This quantity controls the oscillation frequency in matter, similarly to Δm^2 in vacuum.

With the new change of basis, the oscillation probability retains the same functional form as in vacuum, with the substitutions $\theta \rightarrow \theta_M$ and $\Delta m^2 \rightarrow \Delta m_M^2$:

$$P(\nu_\alpha \rightarrow \nu_\beta) = \sin^2(2\theta_M) \sin^2\left(\frac{\Delta m_M^2 L}{4E}\right). \quad (1.88)$$

where both the effective mixing angle θ_M and the mass-squared difference Δm_M^2 depend on the local electron density via the potential A_{CC} .

A **resonance condition** occurs when $\theta_M = \pi/4$ [37], meaning that:

$$\tan 2\theta_M = \frac{\Delta m^2 \sin 2\theta}{\Delta m^2 \cos 2\theta - A_{CC}} \rightarrow \inf \Rightarrow \Delta m^2 \cos 2\theta - A_{CC} \rightarrow 0$$

At resonance, the mixing is maximal. This corresponds to a resonance value of the charged-current matter potential A_{CC}^R [37]:

$$A_{CC}^R = \Delta m^2 \cos 2\theta \quad (1.89)$$

In this regime, the charged-current potential A_{CC}^R cancels out the $\Delta m^2 \cos 2\theta$ term, which causes the effective mass splitting to reach its minimum value:

$$\Delta m_M^2|_R = \Delta m^2 \sin 2\theta, \quad (1.90)$$

corresponding to a lower oscillation frequency but with the largest possible amplitude.

The resonance condition corresponds to a specific electron density N_e^R [37]:

$$A_{CC}^R = \Delta m^2 \cos 2\theta \quad \Rightarrow \quad N_e^R = \frac{\Delta m^2 \cos 2\theta}{2\sqrt{2} E G_F}, \quad (1.91)$$

meaning that when a neutrino propagates through a medium with electron density $N_e = N_e^R$ (for a given E , Δm^2 , and θ), it can undergo maximal flavor conversion.

Adiabatic MSW conversion for solar neutrinos

In the Sun, the electron density $N_e(r)$ changes rapidly with the radius. $N_e(r)$ is very large at the core, where the solar neutrinos are produced, and then drops continuously towards the surface. The matter potential $A_{CC}(r)$ carries this radial dependence on the electron density

$$A_{CC}(r) = 2\sqrt{2} E G_F N_e(r).$$

High-energy ^8B solar neutrinos are produced deep in the core, where the density is well above the MSW resonance, meaning that $A_{CC} \gg \Delta m^2 \cos 2\theta$. The denominator of Equation (1.85) is dominated by $-A_{CC}$. As a result, $\tan 2\theta_M \rightarrow 0^-$ and $\cos 2\theta_M \rightarrow -1$, which

corresponds to

$$\theta_M \longrightarrow \frac{\pi}{2}. \quad (1.92)$$

In this limit, considering the two-flavor approximation (ν_e, ν_x) , where ν_x is a mixture of ν_μ and ν_τ , we have

$$\begin{pmatrix} \nu_e \\ \nu_x \end{pmatrix} = \underbrace{\begin{pmatrix} \cos \theta_M & \sin \theta_M \\ -\sin \theta_M & \cos \theta_M \end{pmatrix}}_{U_M} \begin{pmatrix} \nu_1^m \\ \nu_2^m \end{pmatrix} \xrightarrow{\theta_M \rightarrow \pi/2} \begin{pmatrix} 0 & 1 \\ -1 & 0 \end{pmatrix} \begin{pmatrix} \nu_1^m \\ \nu_2^m \end{pmatrix} = \begin{pmatrix} \nu_2^m \\ -\nu_1^m \end{pmatrix}. \quad (1.93)$$

This shows that, in this regime, the electron neutrino becomes almost identical to one of the matter eigenstates ν_2^m . If the density changes slowly enough inside the Sun (adiabatic limit), by the time the neutrino reaches the surface of the Sun (where $A_{CC} \rightarrow 0$), the matter eigenstate ν_2^m has continuously transitioned to the vacuum eigenstate ν_2 .

Once in vacuum, a pure mass eigenstate does not further oscillate. Therefore, the neutrino stays in this pure ν_2 state as it travels from the Sun's surface to the Earth. The probability that the neutrino reinteracts as ν_e is then [22, 37]:

$$P_{ee}^\odot \simeq |U_{e2}|^2 = \sin^2 \theta_{12}. \quad (1.94)$$

This predicts that only about one-third of the original ν_e flux reaches Earth as ν_e , with the remaining being ν_μ and ν_τ , which explains the solar neutrino problem in SNO [38]. The solar neutrino problem is resolved not by long-baseline neutrino oscillations from the Sun to the Earth, but by the MSW flavor conversion that occurs deep inside the Sun.

Neutrinos, antineutrinos, and the NOvA matter effect

For long-baseline accelerator experiments such as NOvA, the Earth's crust density can be approximated as constant with a value of $\rho \simeq 2.7 \text{ g/cm}^3$. At the beam energy of interest

in NOvA, this leads to a significant matter effect.

Although there are no positrons in the Earth's crust, antineutrinos still interact with the electrons in matter via the same forward scattering process, but with opposite helicity and opposite coupling sign. Therefore, the contribution of the matter potential in the Hamiltonian carries the opposite sign:

$$V_{CC}(\bar{\nu}_e) = -\sqrt{2} G_F N_e. \quad (1.95)$$

This sign flip means that if neutrinos experience a matter enhancement, antineutrinos will experience a matter suppression. This can produce an asymmetry between neutrinos and antineutrinos. For NOvA, the effect of this asymmetry depends on the sign of Δm_{31}^2 (the mass ordering) and on the value of the CP-violating phase δ_{CP} . The matter effects and the intrinsic CP violation can produce similar patterns in the oscillation probabilities. Therefore, their contributions are often degenerate in measurements. The matter effect for the NOvA experiment is further discussed in Section 3.

1.9 Different types of neutrino interactions

Neutrinos interact with matter via the weak force through the exchange of a charged $W^\pm$ boson in **Charged-Current (CC)** interactions or a neutral Z^0 gauge boson, in **Neutral-Current (NC)** interactions.

In a CC interaction, the neutrino exchanges a charged $W^\pm$ boson with a nucleon (or lepton), producing the corresponding charged lepton together with a hadronic system:

$$\nu_\ell + n \rightarrow \ell^- + p + X,$$

where $\ell = e, \mu, \tau$ and X represents possible hadronic jets. With CC interactions, the flavor of the produced charged lepton allows the identification of the incoming neutrino's flavor,

making CC processes important for neutrino experiments.

The NC interactions involve the exchange of a neutral Z^0 boson:

$$\nu_\ell + N \rightarrow \nu_\ell + X.$$

Here, the result of the interaction is an outgoing neutrino that retains its initial flavor and a hadronic system X .




Figure 1.5: Illustration of the two weak interaction processes for neutrinos. **Left:** Charged Current (CC) interaction mediated by $W^\pm$, producing the corresponding charged lepton and some hadrons. **Right:** Neutral Current (NC) interaction mediated by Z^0 , where the neutrino exits with the same flavor, and only the hadronic system is observed.

For the NOvA experiment, where the typical neutrino energies are in the few GeV range, the most common neutrino interactions are quasielastic (QE) scattering, meson-exchange currents (MEC), resonance (RES) production, and deep inelastic scattering (DIS).

1.9.1 Charged-Current Quasielastic (CCQE) Interaction

In a charged-current quasielastic (CCQE) interaction, the neutrino exchanges a charged W boson with a bound nucleon, producing the corresponding charged lepton with the nucleon being ejected from the nucleus:

$$\nu_l + n \rightarrow l^- + p.$$

The CCQE interaction is the dominant charged-current interaction channel at neutrino energies in the range around 1 GeV [39]. It produces a specific topology with a single outgoing charged lepton and a proton.

Assuming quasielastic scattering, the neutrino energy can be reconstructed approximately from the measured lepton kinematics. However, there are nuclear effects such as Fermi motion, which cause deviations from the simple two-body reconstruction [40].




Figure 1.6: CCQE interaction: $\nu_l n \rightarrow l^- p$.

1.9.2 Meson-Exchange Current (MEC / 2p2h)

The Meson-Exchange Current (MEC), or two-particle-two-hole (2p2h), describes a process in which the neutrino interacts with a correlated pair of nucleons inside the nucleus (pp or np). In the nucleus, nucleons are bound together by the strong nuclear force, which can be modelled as being mediated by mesons (virtual pions or ρ mesons) being constantly exchanged between nucleons [40].

When a neutrino interacts with the nucleus via the weak force, it is possible that the weak boson couples to the meson that is exchanged between two nearby nucleons. This is referred to as the "meson-exchange current" interaction. In this case, the weak current acts on the correlated two-nucleon pair. The pair can then be ejected, leaving two "holes" in the nuclear state. Hence the name "two-particle-two-hole" (2p2h) [40].




Figure 1.7: MEC (2p2h): $\nu_\ell + (NN) \rightarrow \ell^- + N + N'$; interaction with a correlated nucleon pair via meson-exchange currents, producing two outgoing nucleons.

1.9.3 Charged-Current Resonant (CCRES)

In charged-current resonant (CCRES) interactions, a neutrino exchanges a W boson with a nucleon (proton or neutron), exciting it into a baryon resonance, usually a $\Delta(1232)$. This resonance is an excited state of the nucleon, which quickly decays into a lower-energy nucleon (N) and a pion [39, 41]:

$$\nu_\ell + p \rightarrow \ell^- + \Delta^{++}, \quad \Delta^{++} \rightarrow p + \pi^+.$$




Figure 1.8: CCRES: $\nu_\ell N \rightarrow \ell^- N' \pi^+$; single-pion production via baryon resonance excitation (Δ^{++}).

1.9.4 Charged-Current Deep Inelastic Scattering (CCDIS)

In charged-current deep inelastic scattering (CCDIS), a neutrino ν_ℓ scatters from a parton (quark) inside a nucleon N via W^+ exchange, producing a charged lepton ℓ^- . The struck

quark recoils and hadronizes into a multi-hadron final state X [42]:

$$\nu_\ell + N \rightarrow \ell^- + X.$$

In CCDIS interactions, the momentum transfer is large, and the target does not remain intact but is broken up into a multi-hadron system X [42].




Figure 1.9: CCDIS: $\nu_\ell N \rightarrow \ell^- X$. A charged-current deep inelastic scattering process where a W boson couples the lepton and quark currents; the struck parton hadronizes into the multi-hadron system X .

1.10 Coherent Pion Production

In this process, the neutrino interacts with the entire nucleus A as one collective target. The process is characterized by a small four-momentum transfer $|t|$ to the nucleus. This means that the pion produced in the interaction carries almost all the forward momentum of the incoming neutrino. Consequently, the pion comes out nearly in the same direction as the neutrino beam (forward pion) [42, 41].

For charged-current interactions, the reaction is:

$$\nu_\mu + A \rightarrow \mu^- + \pi^+ + A.$$

For the neutral-current case, we have:

$$\nu_\mu + A \rightarrow \nu_\mu + \pi^0 + A.$$

The scattering amplitudes from all nucleons in the nucleus add up coherently. This is why the process is called "coherent". This is the case when the momentum is small enough ($|t|R^2 \ll 1$, where R is the nuclear radius).




Figure 1.10: CC coherent pion production on a nucleus A : $\nu_\mu A \rightarrow \mu^- \pi^+ A$. The entire nucleus recoils coherently (small $|t|$) and remains intact, producing a very forward π^+ .

1.11 Types of Neutrino Experiments

Neutrino experiments can be classified according to the source of neutrinos, the energy/baseline regime, and the corresponding physics goals. The objectives of these experiments include oscillation measurements, neutrino mass studies, interaction cross sections, and search for rare processes.

1.11.1 Accelerator-Based Oscillation Experiments

For accelerator neutrino experiments, neutrinos are created by directing an intense beam of protons at a stationary target. This proton-target collision generates hadrons, mostly π and kaons K . The hadrons are subsequently guided using magnetic devices called "horns" to select their charge sign. As these particles travel, they decay into neutrinos, producing a beam made up mainly of muon neutrinos or antineutrinos. These experiments are often defined by the distance the neutrinos travel (the baseline, L) and their energy E_ν .

- **Short-baseline (SBL):** These experiments have a baseline length of $L \sim 10\text{--}1000$ m and neutrino energy of $E_\nu \sim 0.1\text{--}1$ GeV. SBL experiments are important for the analysis of neutrino-nucleus cross sections, for detector R&D, or for probing potential

sterile neutrinos. Some examples of short-baseline neutrino experiments are LSND [43], MiniBooNE [44], SBND [45], MicroBooNE [46], and ICARUS [47].

- **Long-baseline (LBL):** LBL are experiments with a baseline length of $L \sim 100\text{--}1300$ km and neutrino energy of $E_\nu \sim 0.5\text{--}5$ GeV. LBL experiments are sensitive to δ_{CP} , the neutrino mass ordering, and the θ_{23} octant. They measure those parameters through the analysis of $\nu_\mu \rightarrow \nu_e$ appearance and ν_μ disappearance channels. LBL programs typically use near detectors to measure the unoscillated flux and constrain cross-section systematics, and far detectors to observe the oscillated spectra. For these experiments, matter effects are important. Examples of LBL experiments include OPERA [48], K2K [49], MINOS/MINOS+ [50], NOvA [51], DUNE [52], and T2K/Hyper-K [53, 54].

1.11.2 Reactor Antineutrino Experiments

Thermal nuclear reactors provide an excellent source of $\bar{\nu}_e$, produced in the β decay of fission fragments. There are three different types of reactor neutrino experiments:

- **Short/near detectors** ($L \sim 5\text{--}20$ m): They are designed to test anomalies in the predicted flux and to search for sterile neutrinos at the eV mass scale. Examples of such experiments include PROSPECT, STEREO, and DANSS [55, 56, 57].
- **Km-scale experiments** ($L \sim 300\text{--}2000$ m): These experiments have enabled the precision measurement of the mixing angle θ_{13} through comparison of data from near and far detectors. The detection is done through inverse beta decay, $\bar{\nu}_e + p \rightarrow e^+ + n$. Some examples of Km-scale experiments are Daya Bay [58], RENO [59], and Double Chooz [60].
- **Medium-baseline** ($L \sim 50\text{--}60$ km): At these baselines, the oscillations driven by Δm_{21}^2 and Δm_{31}^2 begin to overlap. This creates a subtle interference pattern in the observed energy spectrum. These experiments can disentangle these effects, due to their

high statistics and excellent energy resolution. They allow for a precise determination of Δm_{21}^2 and offer sensitivity to the neutrino mass ordering. The leading example for medium-baseline experiments is JUNO [61].

- **Long-baseline reactor experiment ($L \sim 180$ km):** One example of long-baseline reactor neutrino experiment is KamLAND. KamLAND is located in the Kamioka mine in Japan, and was designed to detect reactor antineutrinos via the inverse beta decay. The antineutrinos observed by KamLAND originated from dozens of commercial power plants distributed across Japan, with an average baseline of about 180 km. Such baselines make KamLAND sensitive to oscillations governed by Δm_{21}^2 . KamLAND measured Δm_{21}^2 with high precision. The observations in SNO due to the MSW effect could have been explained by several theories, including Small Mixing Angle (SMA, which, as its name suggests, assumes a very small mixing angle) or the Large Mixing Angle (LMA) solution. This debate was only settled after KamLAND's measurement, which showed that only the LMA solution is compatible with SNO's observations.

1.11.3 Solar Neutrino Experiments

The Sun generates electron neutrinos through nuclear fusion, with their energies ranging from less than 1 MeV to up to around 20 MeV. These neutrinos come from several key reactions happening in the solar core:

- **proton-proton chain:** In this interaction, two protons fuse to form deuterium, which releases a low-energy ν_e . The majority of the total solar neutrino flux comes from these low-energy ν_e . Even if they are abundant, these neutrinos are difficult to detect directly due to their low-energy.
- **${}^7\text{Be}$ line:** The electron capture on ${}^7\text{Be}$ produces monoenergetic neutrinos at two energies, 0.38 and 0.86 MeV. They are important for analyzing MSW matter effects.

- **pep reaction:** The proton–electron–proton (pep) process produces monoenergetic neutrinos at 1.44 MeV. While pep is much less frequent than pp neutrinos, they are used as an important calibration point for solar neutrino experiments.
- **^8B decay:** The neutrinos from beta decay of ^8B have energies of up to ~ 15 MeV. They dominate the high-energy tail of the solar spectrum. Although ^8B decays are quite rare, they were the first solar neutrinos observed in real-time detectors such as Super-K and SNO.
- **CNO cycle:** This is a secondary fusion cycle in the Sun that produces neutrinos in the 1–2 MeV range from ^{13}N , ^{15}O , and ^{17}F decays. While CNO cycles contribute only about 1% of the Sun’s energy, they are important for testing solar metallicity models. Evidence for CNO neutrinos has recently been reported by Borexino.

Solar neutrino experiments fall into two broad categories:

- **Radiochemical:** These experiments detect neutrinos indirectly by measuring how many target nuclei are transformed over long periods of time. The pioneering Homestake experiment (using a chlorine target) [15], along with the gallium-based SAGE and GALLEX detectors [62, 63], are examples of experiments using this method.
- **Real-time:** Real-time experiments detect each neutrino interaction as it happens. This allows for the measurement of not only the neutrino energy, but also its arrival direction in some cases. The Super-K and Hyper-K [64] are some examples of real-time solar neutrino experiments, along with Borexino [65], and SNO [66].

Solar neutrino experiments have provided some of the most accurate measurements of the mixing angle θ_{12} and the solar mass-squared difference Δm_{21}^2 . They have also helped confirm the presence of the Mikheyev–Smirnov–Wolfenstein (MSW) matter effect in the Sun.

Chapter 2. The NOvA Experiment

The NOvA Experiment (NuMI Off-Axis ν_e Appearance Experiment) is a long-baseline neutrino oscillation experiment, which has as primary physics goals the precise measurement of the neutrino oscillation parameters, especially the mixing angle θ_{23} , and the determination of the neutrino mass hierarchy via the study of $\nu_\mu \rightarrow \nu_e$ oscillation. The experiment utilizes two liquid scintillator detectors: the NOvA Near Detector (ND), situated 1 km away from the production target and 105 m underground at Fermilab, and the NOvA Far Detector (FD), located 810 km away on the surface at Ash River, Minnesota. The two detectors are functionally identical tracking calorimeters consisting of highly-reflective cells made of PVC (polyvinyl chloride) infused with 15% TiO₂, filled with liquid scintillator, which consists of mineral oil blended with 5% pseudocumene, which serves as the primary scintillating agent, and smaller fractions of PPO and Bis-MSB431, which serve as wavelength shifters [67]. The Near Detector has a fiducial mass of 310 tons, whereas the Far Detector has a mass of 14 ktons. The neutrinos observed by the NOvA experiment are provided by the Neutrinos at the Main Injector (NuMI) beam facility at Fermilab. NuMI uses a 120 GeV proton beam from its Main Injector for neutrino production. The high-energy proton beam impinges on a fixed graphite target to produce pions and kaons, which will decay into muons and neutrinos. The produced neutrino beam is composed of 97.5 % ν_μ , 1.8% $\bar{\nu}_\mu$ and 0.7% combined ν_e and $\bar{\nu}_e$. To optimize observation of $\nu_\mu \rightarrow \nu_e$ transitions, the NOvA detectors are placed 14.6 mrad off-axis of the NuMI beam, in such a way that the observed neutrino energy spectrum would peak at 2 GeV. This narrow band neutrino beam of energies between 1 and 3 GeV closely matches the first 3-flavor oscillation maximum for the NOvA baseline, expected at 1.6 GeV [5].

2.1 The NuMI Beam

The neutrino beam in NOvA is provided by the NuMI (Neutrinos at the Main Injector) neutrino source at Fermilab [4]. NuMI supplies neutrinos for several experiments, including NOvA. Below, we discuss how the NuMI beam is created using the Fermilab Accelerator Complex.

2.1.1 Extracting protons from the Fermilab Accelerator Complex

The protons used to generate the NuMI beam are extracted from the Fermilab Accelerator Complex, illustrated in Fig. 2.1. First, H^- ions are created through plasma discharges on hydrogen gas and accelerated to 400 MeV in the Linac (Linear Accelerator) before being injected into the Booster, a synchrotron accelerator that further increases their energy using radio-frequency (RF) electric fields [68]. During this process, magnetic fields are gradually ramped up, or "synchronized" with the particles' rising momentum, to maintain a stable circular orbit. The Booster accelerates the ions to 8 GeV, after which they pass through a carbon foil that strips away their electrons, producing a high-energy proton beam. In the booster, alternating RF cavities, placed out of phase with each other, can accelerate or decelerate parts of the proton beam. This process causes the protons to cluster together, forming a "bunch". The Booster, producing 84 bunches per cycle, generates one "batch" containing 4.3×10^{12} protons. Each batch is then supplied to the Recycler Ring in $1.6 \mu\text{s}$ intervals [68]. The Ring combines 6 pairs of batches, which are sent to the Main Injector to be accelerated to 120 GeV. A total of 13.6×10^{13} protons are delivered to the target at a beam power of up to 900 kW in short bursts ("spills") lasting $10 \mu\text{s}$, occurring as frequently as 1.07 seconds. These protons are then extracted and directed to the NuMI beamline [68].

Fermilab Accelerator Complex




Figure 2.1: Fermilab Accelerator Complex: H^- ions are accelerated to 400 MeV in the Linac, then injected into the Booster, where they are ramped up to 8 GeV and stripped of their electrons to form a proton beam. The protons are then transferred to the Main Injector, where they are accelerated to 120 GeV before being delivered to the NuMI beamline [3].

2.1.2 The NuMI Beamline energy

From the two-flavor approximation probability, neutrino oscillations depend on the baseline length L , the neutrino energy E , and the mass splitting Δm^2 . For the impact of neutrino oscillations to be apparent, $\Delta m^2 L$ should be equal to or larger than E (in the mentioned equation units). The proton accelerator in the NuMI beamline produces 120 GeV protons and is well-suited for producing neutrino beams in the 1-10 GeV energy range. Considering that long-baseline neutrino experiments operate over distances of several hundred kilometers, a neutrino beam in that energy range would be optimum for probing oscillations with Δm^2 in the order of 10^{-3} eV^2 . Since the neutrino flux decreases proportionally to the square of the distance from the decay point, the initial neutrino flux from NuMI must be powerful enough to generate a sufficient event rate in the Far Detector. To ensure this, the NuMI beamline was designed to produce a high-intensity proton beam with a power of up to 900 kW [68].

2.1.3 The NuMI Beam Components

To produce neutrinos, high-energy protons are extracted from the Fermilab Main Injector using a kicker magnet and travel in a straight line until they impinge upon a graphite target of approximately 6.4 mm in width, 15 mm in height, and 1 m in length. This interaction generates secondary charged hadrons, mostly pions and kaons, which are focused by two magnetic horns and then decay within a 675 m long decay pipe. Positively charged mesons (π^+ , K^+) decay predominantly into μ^+ and ν_μ , while negatively charged mesons (π^- , K^-) decay into μ^- and $\bar{\nu}_\mu$. It is possible to operate the beamline in either neutrino or antineutrino mode by selecting the desired meson charge with the horns. In neutrino mode, the hadrons

mostly decay into muons and ν_μ via

$$\pi^+ \rightarrow \mu^+ + \nu_\mu \quad (2.1)$$

$$K^+ \rightarrow \mu^+ + \nu_\mu \quad (2.2)$$

with less than 1% contamination of ν_e from rare kaon and muon decay channels [69].

To measure the leftover hadrons, a hadron monitor is placed right after the decay pipe. A 5 m absorber would then capture the residual hadrons. The last part of the NuMI Beamline apparatus consists of a 240 m dolomite rock section, from which alcoves were excavated to host three muon monitors. The rock serves as a muon shield, and the muon monitors enable measurement of the residual muon flux.




Figure 2.2: Schematic of the NuMI beamline: Protons (~ 120 GeV) from the Fermilab Main Injector strike a graphite target in the Target Hall, producing secondary mesons ($\pi^\pm$, $K^\pm$). Two pulsed magnetic horns (*Horn 1* and *Horn 2*) select the sign and focus the secondaries toward the beam axis, establishing either ν_μ running (focus π^+ , K^+) or $\bar{\nu}_\mu$ running (focus π^- , K^-). The focused mesons traverse a 675 m helium-filled decay pipe where they decay in flight, e.g. $\pi^+ \rightarrow \mu^+ \nu_\mu$. Downstream, a hadron absorber stops remnant protons and undecayed hadrons, while muon monitors embedded in a rock sample measure the accompanying muon flux as a proxy for the neutrino beam direction and intensity [4].

2.1.4 The NuMI Target

Graphite was chosen for the NuMI target because, as a low atomic number (low-Z) material, its interaction with the proton beam produces a more forward-peaked distribution of hadrons, which enhances neutrino production efficiency. Specifically, the NuMI target is made of an ultra-pure, high-density, superfine ZXF-5Q graphite produced by POCO Graphite Inc. This type of graphite has high thermal conductivity and excellent shock resistance, making it ideal for withstanding the intense heat loads generated by the proton beams. The target has a density of 1.78 g/cm^3 and 47 graphite fins, each measuring 20 mm in length, 15 mm in height, and 6.5 mm in width [4]. The total target length is 95.38 cm, with each fin placed 0.3 mm apart. The graphite target's properties have been thoroughly analyzed, including its thermal tolerance. It was determined that the material can withstand temperatures up to 344°C , which is significantly higher than the heat load imposed by the proton beam.

2.1.5 The NuMI focusing horns

The role of the magnetic horns is to select and focus the charged hadrons produced by the proton-target collision. With the use of pulsed magnetic fields, the horns enable the enhancement of the neutrino flux in a specific direction and energy range. Each horn is composed of two coaxial inner and outer conductors shaped to create a toroidal magnetic field when pulsed with current. A pulse of 200 kA current at around 5Hz is applied to each horn, creating a magnetic field of up to 3 T between the conductors in the azimuthal direction. The resulting Lorentz Force focuses the positively charged particles when current flows downstream the beam line. This mode is called "Forward Horn Current" (FHC), and is used to produce a neutrino beam. When the current flows in the reverse direction, negatively charged particles get selected, creating an antineutrino beam. This is the "Reverse Horn Current" (RHC) mode.

There are two focusing horns in the NuMI beamline, labeled as Horn 1 and Horn 2. They are both made from two aluminum parabolic sections joined together, with the narrow first section focusing low-angle particles, and the second flaring section focusing wide-angle particles. Horn 1 is positioned close to the target, with Horn 2 located 10 meters downstream. This configuration has been chosen to maximize neutrino yields in the energy range of interest of the MINOS experiment, which is around 2-5 GeV, representing the first oscillation maximum for the MINOS baseline, where sensitivity to Δm^2 and $\sin^2(2\theta)$ is highest. For the NOvA experiment, the desired energy range is slightly different. NOvA observes a neutrino beam that is 0.8 degrees off-axis (Section 2.2) in the 1-3 GeV range. This corresponds to an on-axis beam peaked at around 6-7 GeV. This was achieved by repositioning the target 1 m upstream of Horn 1 to focus higher-momentum pions [69].

The hadrons selected by the horns will then travel through the "decay pipe" to produce neutrinos or antineutrinos. The decay pipe is a cylindrical tunnel with a length of 675 meters and a diameter of 2 meters. The length was chosen such that most mesons will decay before being absorbed by the downstream apparatus. As an example, a 10GeV π^+ has a decay length of around 560 meters, and about 70% of the pions will decay within 675 meters. The decay pipe is filled with helium gas, as its lower density reduces the interactions of hadrons and muons compared to air. Furthermore, helium also limits the production of secondary neutrons and photons. The decay pipe is encased in a concrete shielding to contain the intense secondary radiation generated when undecayed hadrons interact with the pipe walls [4].

2.1.6 Downstream apparatus

Downstream of the decay pipe are the absorber, the hadron and muon monitors, and finally the rock shielding (Fig. 2.2). The absorber, constructed from steel, aluminum, and concrete, stops surviving protons and undecayed hadrons from the decay pipe to prevent




Figure 2.3: Illustration of magnetic horn focusing in neutrino and antineutrino beam modes. Protons impinging on the graphite target generate a mix of secondary mesons. These particles are then steered using two pulsed magnetic horns, which produce a toroidal magnetic field that selectively focuses particles of a given charge and defocuses those of the opposite sign. (a) In antineutrino mode, the horn polarity is set to focus negatively charged mesons like π^- , which subsequently decay via $\pi^- \rightarrow \mu^- \bar{\nu}_\mu$, generating a beam dominated by $\bar{\nu}_\mu$. (b) Reversing the horn current shifts the focus to positively charged mesons such as π^+ . These decay through $\pi^+ \rightarrow \mu^+ \nu_\mu$, producing a beam primarily composed of ν_μ (neutrino mode). Particles not aligned with the desired charge are defocused and removed from the beam path. By alternating the direction of the horn current, the NuMI beamline can switch between neutrino and antineutrino modes.

them from contributing backgrounds in the detectors. Positioned immediately upstream of the absorber is a hadron monitor, which measures the flux of hadrons reaching this point and is used to monitor horn focusing efficiency. A series of muon monitors is placed downstream of the absorber and embedded in rock alcoves to detect muons from meson decays, providing an indirect measurement of the neutrino flux and beam alignment. The rock serves as additional shielding to stop residual muons, ensuring that only neutrinos reach the NOvA Near Detector.

2.2 NOvA Off-axis Detector Location

As mentioned earlier, NOvA's baseline is 810 km with the first oscillation maximum peaked at around 2 GeV. If the NOvA Far Detector were placed on the beam axis, the obtained neutrino energies would form a broad spectrum peaking at an energy of around 6-7 GeV (Fig. 2.4), with a high-energy tail. However, placing the Far Detector slightly off-axis would produce a narrow-band neutrino energy spectrum centered at approximately 2 GeV with minimal high-energy background. Indeed, when a meson decays in flight, the energy of the resulting neutrino depends on the angle between the neutrino and the meson direction.

For the pion case, for instance, in the pion rest frame (denoted by a *), the decay is two-body. The neutrino has a fixed energy given by:

$$E_\nu^* = \frac{m_\pi^2 - m_\mu^2}{2m_\pi}$$

If we plug in the masses:

$$m_\pi = 139.6 \text{ MeV}, \quad m_\mu = 105.7 \text{ MeV}$$

$$E_\nu^* \approx \frac{(139.6)^2 - (105.7)^2}{2 \times 139.6} \text{ MeV} \approx 29.8 \text{ MeV}$$

We boost in the lab frame with a Lorentz boost $\gamma = \frac{E_\pi}{m_\pi}$, and $\beta = \sqrt{1 - \frac{1}{\gamma^2}}$. The neutrino

energy in the lab frame depends on the angle θ^* (between the neutrino and pion direction in the pion rest frame):

$$E_\nu = \gamma E_\nu^* (1 + \beta \cos \theta^*)$$

A change of variables (using relativistic transformations) leads to the following approximate formula for small θ :

$$E_\nu(\theta) \approx \frac{E_\nu^*}{1 + \gamma^2 \theta^2}$$

We multiply numerator and denominator by γ :

$$E_\nu(\theta) \approx \frac{\gamma E_\nu^*}{1 + \gamma^2 \theta^2}$$

Then, substitute $\gamma = \frac{E_\pi}{m_\pi}$ and $E_\nu^* \approx 0.43 m_\pi$:

$$E_\nu(\theta) \approx \frac{0.43 E_\pi}{1 + \left(\frac{E_\pi}{m_\pi}\right)^2 \theta^2}$$

It can be deduced from the formula that the neutrino receives its **maximum energy** when $\theta = 0$ (on-axis): $E_\nu \approx 0.43 E_\pi$ and as θ increases (off-axis), E_ν decreases and the spectrum becomes narrower. Therefore, if we choose a fixed θ , specifically 14 mrad for NOvA, we obtain a narrow-band neutrino energy spectrum centered at the desired oscillation maximum [5].

2.3 The NOvA Detectors

The NOvA Near and Far detectors were designed to be functionally identical. This was done such that systematic uncertainties with the same effects between the detectors cancel out or are reduced. Both detectors are made of polyvinyl chloride (PVC) extrusions filled with liquid scintillator. When a neutrino from the NuMI beam interacts within the liquid scintillator of the PVC cells, it produces charged particles that deposit energy along




Figure 2.4: Off-axis energy spectra in NOvA. Simulated ν_μ CC event-rate spectra (per 6×10^{20} POT per kton) for different off-axis angles: on-axis (black), 7 mrad (green), 14.6 mrad *NOvA* (red), and 21 mrad (blue). Increasing the off-axis angle narrows the spectrum and shifts it to lower energy due to two-body decay kinematics of $\pi^\pm \rightarrow \mu^\pm \nu_\mu (\bar{\nu}_\mu)$, approximately $E_\nu \simeq 0.43 E_\pi / (1 + \gamma^2 \theta^2)$. NOvA’s 14.6 mrad configuration yields a narrow-band beam peaked near 2 GeV, aligning with the first oscillation maximum for the $L \simeq 810$ km baseline [5].

their paths. This energy deposition generates scintillation light. Each Far detector extrusion measures about 15.5 m in length (3.9 m for the Near detector) and is divided into 16 elongated rectangular cells with internal dimensions of approximately $3.9 \times 6.6 \text{ cm}^2$. The extrusions are arranged side by side to form planes, with each successive plane rotated orthogonally to the one before it. This alternating configuration gives two-dimensional projections of the particle tracks. Combining views from these planes allows for a full three-dimensional reconstruction of the particle trajectories. A wavelength-shifting fiber threaded through each detector cell collects scintillation light generated by the charged particles passing through the detector. The light signal is then read out by avalanche photodiodes (APDs) [5]. Figure 2.5 illustrates how neutrino interactions are reconstructed in the NOvA detector.

3D schematic of NOvA particle detector




Figure 2.5: 3D schematic of the NOvA detector and its reconstruction principle. A neutrino from the NuMI beam interacts in the liquid scintillator inside the PVC cells, producing charged particles that generate scintillation light. The signals recorded in successive detector planes are used to reconstruct particle trajectories. Alternating planes of vertically and horizontally oriented cells allow for three-dimensional tracking, as shown in the top and side views, where multiple outgoing particles are identified from a single interaction vertex.

2.3.1 PVC cells and Liquid Scintillator

NOvA uses a mineral oil-based liquid scintillator doped with 4.1% pseudocumene (1,2,4-trimethylbenzene) together with the primary fluor PPO (2,5-diphenyloxazole) and the secondary fluor bis-MSB (1,4-di(methylstyryl)benzene) [5]. Energy deposited in the pseudocumene excites ultraviolet emission, which is absorbed and re-emitted by the PPO at ~ 350 - 380 nm and subsequently shifted by the bis-MSB to ~ 420 nm, matching the absorption spectrum of the wavelength-shifting fibers. The scintillator has a light yield of $\sim 10^4$ photons/MeV, providing nanosecond-scale timing resolution and excellent light output. It is contained within extruded PVC cells whose walls are infused with 15% titanium dioxide

to increase reflectivity. Each cell has a rectangular cross-section of $3.9 \text{ cm} \times 6.0 \text{ cm}$, with lengths up to 15.5 m in the Far Detector (shorter in the Near Detector). Each detector plane is made of 32-cell blocks. They are formed by joining 16-cell extrusions side-by-side.



Figure 2.6: PVC extrusion modules used in the NOvA detectors. Each extrusion consists of 16 individual rectangular cells formed from rigid PVC, with approximate transverse dimensions of $3.9 \times 6.6 \text{ cm}^2$ per cell. The left panel shows the technical drawing with dimensions in millimeters, while the right panels display photographs of the fabricated extrusion. These extrusions are assembled into planes, with cells filled with liquid scintillator and read out by wavelength-shifting fibers coupled to APDs.

2.3.2 Wavelength-shifting fibers

As particles traverse the detector, they produce scintillation light, which is collected by the wavelength-shifting fibers. They are made of 0.8 mm diameter Kuraray Y11 fibers that are looped through each PVC cell. The fibers enhance the light signal by absorbing the original scintillation light and re-emitting it at a longer wavelength. Specifically, the scintillation photons emitted by charged particles in the detector have a wavelength of approximately 400 nm (blue light), which are shifted to about 500 nm (green light) by the wavelength-shifting fibers. The ends of the fiber are connected to an avalanche photodiode (APD) [5].

2.3.3 Avalanche Photodiodes (APDs)

The photodetector used in NOvA is an APD, which is a silicon semiconductor-based detector designed to sense very low levels of light. When photons enter the silicon, they are absorbed and create electron-hole pairs. The APD is operated under a high reverse voltage, which creates a strong electric field that accelerates these charge carriers. As the charges speed up, they can produce additional electron-hole pairs, creating a chain reaction or "avalanche." This multiplies the original signal by about a factor of 100, and enables the APDs to detect the very faint flashes of light from the NOvA scintillator cells. The APD has a quantum efficiency of about 85% for the green light (520-550 nm) emerging from the wavelength-shifting fibers. This means that roughly 85 out of every 100 incident photons in this wavelength range are absorbed in the silicon and produce a photoelectron [5]. This high efficiency maximizes light collection, ensuring that even the small number of photons transmitted through the long scintillator cells can be converted into an electrical signal. In NOvA, each APD is thermoelectrically cooled to about -15°C to suppress dark current and reduce noise [5].

2.3.4 Data Acquisition System

The NOvA Data Acquisition (DAQ) system converts the scintillation light collected by the APDs in over 360,000 detector channels into synchronized digital data streams that can be used for physics analysis. It combines electronics that amplify and digitize the signals. There are also modules that collect and organize the data at high speed, a GPS-based clock system that keeps all channels synchronized, and a flexible triggering system that decides which data to record. In the NOvA DAQ system, the signal collected by the APDs is first amplified and shaped by a custom application-specific integrated circuit (ASIC) mounted on the Front-End Boards (FEBs). Each channel is then digitized at a rate of 2 MHz [6].

To suppress electronic noise and baseline drift, the readout employs a Dual Correlated




Figure 2.7: NOvA detector cell: a 15 m long tube filled with liquid scintillator. A wavelength-shifting fiber is located on one end of the readout, which absorbs and re-emits scintillation light to enhance the signal. At the far end of the cell is an avalanche photodiode to collect the measured light output.

Sampling (DCS) technique. In the first sample, the pedestal (the electronics baseline plus noise) is recorded, and in the second sample, the pulse plus baseline is measured. Subtracting the two samples cancels out slow drifts and low-frequency noise, leaving a clean signal. However, there are residual fluctuations of a few ADC counts after DCS, so a zero suppression is applied by the FEBs. In this case, only signals above a channel-dependent threshold are transmitted to the Data Acquisition (DAQ) system, while sub-threshold hits are discarded locally [6].

Data Concentrator Module

The digitized signals from the different FEBs are transmitted to a Data Concentrator Module (DCM). Each DCM is a custom-built computer board based on a PowerPC 8347 processor and a Xilinx Field Programmable Gate Array (FPGA). The FPGA handles fast data ordering; it groups hits into 50 μs microslices. Once a slice is complete, the PowerPC

processor is interrupted, and the block is moved into system memory. These $50\ \mu\text{s}$ blocks are then combined into larger 5 ms millislices, which are better suited for network transfer. The DCM sends the 5 ms slices over the network using jumbo-frame Ethernet packets, with a sustained throughput of 24 MB/s per DCM [6].

Timing System

A sophisticated timing system is used to ensure that the readout regions defined by the DCMs are synchronized. To make sense of data from hundreds of thousands of channels, all FEBs and DCMs must share a standard time base provided by the timing system master clock. NOvA uses a GPS-locked master timing unit to provide a stable 16/32/64 MHz clock and sync signals to the entire readout system. Every hit is time-stamped in this global NOvA time. The timing system also supplies the beam spill timing information from the Fermilab accelerator complex. These timing systems enable the DAQ to synchronize and read out 368,640 far detector readout cells spread over 11520 different FEBs. This synchronization facilitates the correlation of the data taken by the detector and the beam spills produced by the accelerator [6].

Triggers

For neutrino oscillation measurements, the primary triggers are the NuMI beam spill signal and a calibration pulser. The beam trigger requires that a beam spill occur, then it selects a time window centered around the time of the beam spill. The window includes both the $10\ \mu\text{s}$ spill and the sidebands on each side of the spill. This enables the estimation of background rates. At the far detector, a $30\ \mu\text{s}$ window around each spill is extracted from the buffers and written to disk. Calibration triggers taken during off-spill periods are also included to monitor detector stability. The DAQ can also support software-based Data-Driven Triggers (DDT) [6], which is useful as the detector is 100% live. The DDT framework is a software-based trigger that analyzes buffered data in real time using parallel processors.

These special triggers continuously monitor the live data stream for unusual topologies (e.g., bursts of low-energy hits, long muon tracks, or upward-going tracks) and request readout of those periods. This enables the monitoring of rare-event physics such as supernova neutrinos, magnetic monopoles, or WIMP-induced signals.




Figure 2.8: Overview of the NOvA Data Acquisition (DAQ) system with the ARTDAQ Data-Driven Trigger (DDT) framework [6].

2.3.5 The NOvA Near Detector

The NOvA Near Detector's primary role is to measure the unoscillated neutrino flux. It also provides constraints on the neutrino energy spectrum and flavor composition, and plays a crucial role in constraining systematic uncertainties. The NOvA ND is located approximately 1 km downstream of the NuMI target station and is placed about 105 m underground to shield it from cosmic-ray backgrounds. It is positioned 14.6 mrad (0.8°) off the beam axis (like the NOvA Far Detector). The active volume of the ND is $4 \text{ m} \times 4 \text{ m} \times 15 \text{ m}$, with a total mass of 310 tons, of which 193 tons are active. The detector consists of 192 planes. At the downstream end of the NOvA ND, there is a dedicated muon catcher composed of alternating

steel absorber plates and scintillator planes, which improves the energy measurement of muons from ν_μ charged-current interactions.

2.3.6 The NOvA Far Detector

The NOvA Far Detector is located in Ash River, Minnesota, 810 km away from the neutrino source, with an active volume of $15.7 \times 15.6 \times 60.0$ m. The Far Detector (FD) is composed of 896 planes, each built from 12 extrusion modules, giving a total of 384 cells per plane, with one cell measuring 15.5 m. The FD has a total mass of 14 kton. The FD contains 344,064 readout channels implemented using 10,752 front-end boards (FEBs) with 32 channels each [70]. The Far Detector is placed on the surface, not underground. However, there is a 3 m water-equivalent of rock overburden to reduce cosmic rays. The latter are further mitigated by analyzing the timing, tracking, and directionality.

2.4 NOvA Simulation

2.4.1 NuMI beamline simulation

To simulate the NuMI beamline, NOvA uses the g4numi Beamline Monte Carlo, which is a Geant4-based simulation. It allows us to simulate the 120 GeV protons from the main injector, their interactions within the target, and the particle distribution after focusing the hadrons. G4numi also models the transport of those particles through the decay pipe and the hadron absorber.

The raw flux obtained by g4numi is then corrected with another package called "Package to Predict the Flux" (PPFX) [69]. This package uses weights from NA49/NA61 and thick-target measurements.

2.4.2 Neutrino Event Generator

To simulate neutrino interactions in the NOvA detectors' scintillator and their corresponding final-state particles, NOvA uses the GENIE Monte Carlo generator [71] [72]. GENIE allows the selection of the model to be used for different neutrino interaction channels. In NOvA, for Charged-Current Quasi-Elastic (CCQE) interactions (Section 1.9.1), the default model used is the Relativistic Fermi Gas (RFG) model [73]. In RFG, the nucleons are modeled as a degenerate gas of free particles that obey Fermi-Dirac statistics. They occupy all momentum states up to the Fermi momentum p_F , forming a "Fermi sphere," with no nucleon allowed to have momentum above p_F in the ground state. The nuclear binding effects are accounted for by applying a constant average binding energy E_B . In a CCQE interaction, the struck nucleon is selected from the Fermi sea, with an initial momentum sampled uniformly up to p_F . Then, Pauli blocking is applied, in which an interaction where the nucleon falls back into an already occupied state below p_F is forbidden.

Two-particle-two-hole (2p2h), or MEC interactions (Section 1.9.2), are modeled using the Valencia model developed by Nieves, Ruiz Simo, and Vicente Vacas [40]. It is a model derived from many-body nuclear theory. The Valencia model calculates the two-body nuclear currents from the meson exchange between nucleons. These contributions are used to compute nuclear response functions. They are then integrated to obtain inclusive cross sections. Electron-scattering data have been analyzed to validate the model.

To simulate Resonance (Section 1.9.3) and Coherent pion production (Section 1.10), the Berger-Sehgal model is employed. This framework is an extension of the default Rein-Sehgal (1981) model for single-pion production in GENIE, which includes lepton mass effects. Those effects are especially important for low-energy ν_μ and ν_τ interactions [74]. For coherent pion production, the Berger-Sehgal model includes both lepton mass corrections and updated experimental inputs, which improves agreement with pion production data in the few-GeV

region. [74].

Deep-inelastic scattering (DIS) interactions (Section 1.9.4) are simulated using the Bodek-Yang model. At neutrino energies of a few GeV, purely perturbative QCD and parton distribution functions (PDFs) do not provide an accurate description of DIS interactions. Bodek and Yang introduced phenomenological corrections that are based on the data from SLAC, BCDMS, NMC, and JLab. This was done such that the model actually matches the measured cross-sections in the few-GeV range. The Bodek-Yang model modifies the leading-order PDFs by scaling variables and by introducing K -factors that include target mass corrections, higher-twist effects, and also take into account the low- W case in which resonances overlap with the DIS onset [75, 76].

Furthermore, as a neutrino interacts with the nucleon, the outgoing hadrons have to travel through the nuclear medium before exiting. Therefore, they can re-interact. They can scatter elastically or inelastically on other nucleons, be absorbed, undergo charge-exchange, lose energy, or change direction. These processes are collectively referred to as Final State Interactions (FSI). FSI significantly alters the visible energy and particle content in the detector; therefore, accurately modeling those interactions is crucial. For NOvA, the hN model in GENIE is used to handle FSI. The hN model uses a semi-classical intranuclear cascade approach. In the model, the particles are tracked classically (not using full quantum many-body theory), and the probabilities for re-interactions are derived from experimental hadron-nucleon cross sections. Each produced hadron is then propagated step by step through the nucleus, where at each step, the model checks whether re-interaction occurred or not, based on its mean free path and interaction probability. This "cascade" continues until the hadron either escapes the nucleus or is absorbed [71].

2.4.3 NOvA ND data tuning

Neutrino interactions with nuclei are challenging to model and to measure, particularly for low-energy neutrinos due to Final State Interactions (FSIs) as discussed in the previous section. This is why neutrino interaction cross sections usually represent one of the dominant sources of systematic uncertainties in oscillation analyses. The complexity of this modeling often means that the nominal Monte Carlo (MC) prediction can differ from measured data, as is the case with NOvA. Each interaction modeled in GENIE, however, possesses distinct degrees of freedom and associated uncertainties, which can be independently tuned to enhance agreement between data and simulation. In NOvA, the GENIE models are first optimized to maximize consistency with the Near Detector data. External data constraints are also applied to the simulation. Nevertheless, following these adjustments, residual differences remain between simulation and data, largely attributed to the uncertainty in meson exchange current (MEC) events. To rectify this discrepancy, the MEC components are rescaled such that the resulting simulation matches the NOvA ND data [77].

The following summarizes the set of changes in the central value prediction to enhance data and MC agreement:

- From external constraints:
 - **Axial form factor value:** the default value used in GENIE for the axial form factor, M_A , is 0.99. A reevaluation of the ANL and BNL bubble chamber data indicates a larger value with an error-weighted mean of $1.04 \text{ GeV}/c^2$. Consequently, this revised value was the axial mass used in the analysis.
 - **Applying nuclear model weights from the MINERvA experiment to QE interactions:** more sophisticated QE models substantially reduce QE interactions at low Q^2 and slightly increase them at high Q^2 compared to the nominal

model. Therefore, the QE nuclear model weights developed by the MINERvA experiment are applied.

- **Reduction of non-resonant single pion production by 57%:** a reassessment of bubble chamber data and the recent MINERvA results showed that the single pion production rate is significantly lower than what was previously assumed. The two analyses indicate a 57% difference. Therefore, the corresponding rate in the simulation was reduced accordingly.
- **NOvA data-driven constraints:** current MEC models in GENIE severely underestimate MEC event rates, especially at low hadronic energy. This is attributed to the considerable uncertainty in the MEC spectral shape. Since no specific model can account for this difference, the MEC interactions are rescaled by applying separate weights obtained directly from the NOvA ND data, such that the sum of MEC components agrees with the ND observations [77].

2.4.4 Detector simulation

For the detector simulation, the final state particles obtained from GENIE are transported through the liquid-scintillator detectors using the Geant4 detector transport [78] [79]. Geant4 tracks all final-state particles and their subsequent interactions within the detector using the detailed NOvA geometry. Geant4 simulates the trajectory of the particles and their energy deposition, along with the scintillation light produced as the particles travel through the detector.

Optical modeling

Concerning the optical modeling of scintillation light, Geant4 provides an extremely detailed optical photon tracking that propagates every single photon through the detector. However, this approach is too slow and computationally inefficient for large-scale simula-

tions such as those required by NOvA. Since all cells share the same geometry and materials (identical scintillator, PVC reflectivity, and fiber properties), the simulation can be optimized by performing a single detailed ray-tracing study instead of re-simulating photon transport for every cell. This ray-tracing study determines how light transport depends on the scintillator emission spectrum, the reflectivity of the PVC walls, the attenuation length of the wavelength-shifting fiber, and the fiber-APD coupling efficiency. The results are converted into parametrized lookup tables, which are then used during the NOvA Geant4 simulation in place of full photon tracking.

NOvA uses custom simulation models to reproduce the detector readout chain and generate digitized signals. These models include the APD gain, which internally amplifies the optical signal. Then there's the shaping applied by the front-end ASIC, which matches the readout bandwidth. Finally, realistic electronic noise is added to match real data. The signal processing steps described in Section 2.3.4 are also applied. These allow us to obtain a simulation that matches the characteristics of the real detector data.

Geant4 Physics lists

For the detector simulation, NOvA employs the QGSP_BERT_HP_EMZ physics list in Geant4. This list combines several models to describe the different particle interactions in NOvA. For high-energy hadronic interactions, for instance, the Quark-Gluon String Precompound (QGSP) model is used. It uses a quark-gluon string model to describe the initial interaction between a high-energy hadron and a nucleus. It then applies a precompound model to simulate how the excited nucleus releases energy by emitting secondary particles. QGSP is needed for high-energy protons, pions, and kaons from neutrino interactions or cosmic rays that produce showers in the detector.

The Bertini Cascade (BERT) models intranuclear cascades for intermediate-energy particles (hundreds of MeV up to a few GeV). BERT racks the hadron as it enters the nucleus and

simulates successive collisions with nucleons. This creates a cascade of secondary particles until further hard collisions cannot occur anymore due to the available kinetic energy. BERT is particularly useful for describing hadronic re-interactions inside the detector [80, 81].

For neutrino interactions below 20 MeV, the High Precision (HP) neutrons model provides the corresponding cross-sections. HP uses nuclear data libraries, allowing it to accurately model neutron scattering, absorption, and capture at low energies [82].

Electromagnetic interactions are handled by the Electromagnetic "Option Z" (EMZ or EM Opt4) package. It provides the most accurate electromagnetic interaction model in Geant4, specifically at low energies [80]. EMZ includes effects like ionization fluctuations, bremsstrahlung, Compton scattering, and photon conversion.

Cosmic and rock backgrounds

The NOvA Far Detector is located on the surface and is exposed to significant cosmic rays. These are backgrounds to the neutrino beam, and they also need to be simulated. To model cosmic muons and secondary particles, NOvA uses the CRY generator. When needed, NOvA also uses CORSIKA for more detailed air-shower simulations.

2.5 NOvA Reconstruction

For a proper evaluation of the oscillation parameters, it is essential to distinguish between the different interaction topologies observed in the detector. Each neutrino interaction leaves a characteristic signature in the detector. In particular:

- ν_μ CC interactions produce a long, minimum-ionizing muon track. It typically crosses many planes of the detector. The muon track is often accompanied by short hadronic tracks.
- ν_e CC interactions produce an electromagnetic shower from the outgoing electron. The shower spreads as the electron undergoes bremsstrahlung, and the resulting photons

convert via pair production.

- For neutral-current interactions, there is no outgoing charged lepton, only hadronic activity. In the detector, this usually shows up as a short, diffuse cluster of hits.

Differentiating these topologies is important because different interaction types can end up looking very similar in the detector. This classification relies on precise reconstruction of the neutrino interactions [5].




Figure 2.9: Example event displays of different neutrino interaction topologies in NOvA. Top: a ν_μ charged-current (CC) interaction producing a long muon track and a short proton recoil. Middle: a ν_e CC interaction producing an electromagnetic shower from the outgoing electron, together with a proton. Bottom: a neutral-current (NC) interaction producing a pion and hadronic activity.

The NOvA reconstruction chain begins with the hits recorded in the individual detector

cells. The hits are first grouped into clusters, and during clustering, the clusters are classified as track-like (particle trajectories) or shower-like. Then the interaction vertex is located. Particle identities and kinematics are then determined using particle identification (PID) and energy reconstruction. Finally, the event is classified into a specific interaction topology.

2.5.1 Raw data to calibrated hits

The goal of this step is to convert the electronic signals into the corresponding energy depositions (hits). Signals from energy deposits go through a few processing steps before we can reconstruct particle kinematics. First, the readout is digitized, as described in Section 2.3.4. We then apply a deconvolution to remove the effect of the electronics response and recover the deposited charge from the digitized waveform.

After deconvolution, we proceed with the following calibration corrections:

- **Pedestal and gain calibration:** this corrects for variations in the baseline and amplification between channels. This guarantees that the same energy depositions give consistent signals across the detector.
- **Light attenuation correction:** as light propagates along the wavelength-shifting fiber, its amplitude depends on the distance of the deposition from the readout end. Due to this, a position-based correction is performed.
- **Birks' law correction:** when charged particles lose energy in the scintillator, the excitation density can become so high that scintillator quenching occurs. This reduces the light yield relative to the deposited energy. The Birks' law correction takes this effect into account.

2.5.2 Clustering

The next step is clustering, where we decide which hits most likely came from the same physical interaction. In NOvA, hits are first grouped into "slices", which are hits occurring

close together in time. A slice might contain a single beam neutrino interaction, a cosmic-ray muon, or sometimes a mix of both if multiple particles pass through the detector at the same time. To sort this out, NOvA uses the Time Density Slicer (TDSlicer). It looks through the readout window for localized concentrations of hits. Each slice should correspond to one complete neutrino event or one through-going cosmic muon.

Clustering proceeds in the following stages:

- **Density calculation:** For each hit, the algorithm counts how many neighboring hits fall within a set spatial radius and time window, assigning a local density score.
- **Centroid selection:** Hits that are both locally dense and relatively isolated are chosen as centroids, which act as seeds for clusters.
- **2D clustering:** Hits are grouped around centroids by proximity, separately in the x - z - t and y - z - t views.
- **3D merging:** The two sets of 2D clusters are matched and combined to build complete 3D slices. Each slice is expected to contain one physical interaction.

2.5.3 Vertex reconstruction

The goal of vertex reconstruction is to determine the position of the neutrino interaction in the detector. NOvA uses a 3D vertexing method where track intersections and shower back-extrapolations are combined to determine the origin of the interaction.

Seed vertexing

When a neutrino interacts in the NOvA detector, there are multiple tracks and showers that emerge from a common vertex. To identify tracks, NOvA applies a modified Hough transform:

- Each hit in detector coordinates (x, y) is mapped into the Hough parameter space (ρ, θ) , where $\rho = x \cos \theta + y \sin \theta$. This produces a sinusoidal curve for each hit.
- The curves for hits that are lying in the same straight line in real space intersect in the Hough space. This forms a peak at (ρ_0, θ_0) , as can be seen in Fig 2.10.
- Each peak (ρ_0, θ_0) corresponds to a candidate line, meaning that it is a possible particle trajectory. The associated line in (x, y) space has the equation

$$y = -\frac{\cos \theta}{\sin \theta} x + \frac{\rho}{\sin \theta}.$$

- After mapping hits into (ρ, θ) , the algorithm does not just look for exact curve intersections. Instead, each hit contributes a Gaussian vote in parameter space. When multiple hits line up in real space, their Gaussians overlap. This creates a strong peak at (ρ_0, θ_0) . This peak is then taken as the best candidate line direction. The Gaussian vote is expressed as:

$$\text{vote}(\rho, \theta) = \exp\left[-\frac{(\rho - \rho_0)^2}{2\sigma_\rho^2}\right] \exp\left[-\frac{(\theta - \theta_0)^2}{2\sigma_\theta^2}\right].$$

with σ_ρ being the spread in the distance-to-origin coordinate and controls how broadly the hit contributes along the ρ axis, and σ_θ is the spread in the angle coordinate and controls how broadly the hit contributes along the θ axis.

This probabilistic weighting makes the algorithm more robust to noise and detector granularity, since nearby hits reinforce each other instead of relying only on exact curve intersections. This procedure provides an initial estimate of the “radiating arms” from the vertex [83, 84].




Figure 2.10: Illustration of the Hough transform. Left: hits in detector (x, y) space corresponding to a track (blue), with the true underlying line shown (blue dashed). Right: the same hits mapped into (ρ, θ) Hough space. Points lying on the same straight line in (x, y) intersect at a single peak in (ρ, θ) space, which identifies the parameters of the track.

Elastic Arms for vertex finding

The Hough transform provides line candidates radiating from a common interaction point. The **Elastic Arms** algorithm then treats these line segments as "arms" deformable around a vertex. It minimizes a cost function that balances hit-arm alignment, penalizes unused hits, and searches for vertex consistency. The minimum yields the most likely vertex and primary track directions [5].

The Elastic Arms algorithm minimizes a global cost function of the form

$$E = \sum_i d^2(\mathbf{h}_i, \text{nearest arm}) + \lambda N_{\text{unassigned}} + \mu P_{\text{vertex}}, \quad (2.3)$$

where:

- $d^2(\mathbf{h}_i, \text{nearest arm})$ is the squared distance of hit $\mathbf{h}_i$ to its closest candidate arm, summed over all hits. This term enforces alignment of hits along reconstructed track-like features.

- $N_{\text{unassigned}}$ counts the number of hits that cannot be reasonably associated with any arm. The parameter λ penalizes unused hits, encouraging the fit to account for as much of the observed detector activity as possible.
- P_{vertex} is a penalty term that enforces vertex consistency (e.g. discouraging arms that fail to converge at a common origin or are poorly aligned near the vertex). The coefficient μ controls the strength of this constraint.

Minimizing E yields the best-fit vertex position and set of primary arm directions radiating from the interaction point [84].

2.5.4 Prong Reconstruction and 3D Matching

After the vertex is found, the next step involves evaluating the cluster of hits that would correspond to a specific outgoing particle, called a "prong." Instead of using hard clustering, where each hit is assigned to exactly one cluster, NOvA uses "Fuzzy-k" clustering. In this case, each hit can belong partially to more than one cluster. The hits are given a probability of a "membership" score. This is particularly useful in cases such as when two showers overlap, or when hits are noisy. An example of this would be an electron shower overlapping with a hadronic shower. In this case, some hits are shared by both prongs with partial weight. The "Fuzzy-k" algorithm uses the angular distribution of energy depositions around the reconstructed vertex in each two-dimensional view of the detector. We identify prong centers along directions where significant energy is deposited. Hits are then assigned to these prongs with a membership weight determined by their angular and spatial distance from the candidate prong axis. In fuzzy clustering (fuzzy k -means / fuzzy c -means [85]), each hit i receives a membership weight $u_{ij} \in [0, 1]$ for every candidate prong j , with $\sum_{j=1}^K u_{ij} = 1$. The standard membership rule is :

$$u_{ij} = \frac{1}{\sum_{k=1}^K \left(\frac{d_{ij}}{d_{ik}} \right)^{\frac{2}{m-1}}}, \quad (2.4)$$

where K is the number of prongs, $m > 1$ is the fuzziness parameter (typically $m \simeq 2$), and d_{ij} is the distance from hit i to prong j .

Hits that are near a prong axis (small d_{ij}) receive larger u_{ij} , i.e., hits close to the axis contribute strongly to that prong, while hits further away contribute weakly. Ambiguous hits can share fractional membership across multiple prongs.

Once two-dimensional prongs are formed in each view ($xz-t$ and $yz-t$), they are matched to form 3D prongs. This matching is done using the Kuiper metric, which compares the longitudinal energy profiles of prongs between different views. The Kuiper test responds to both shape differences and cumulative deviations, which helps align tracks and showers reconstructed in the two projections. This approach achieves high efficiency for both track-like muons and diffuse shower-like electrons or hadrons [86, 84].

2.5.5 Identifying tracks

After prong reconstruction, several dedicated tracking algorithms are used to refine the reconstruction of long, straight tracks. First, there is the cosmic tracker, which identifies tracks from cosmic-ray muons. It looks at hits across consecutive planes of the detector within a sliding time window. At each step, it checks whether the hits being added are consistent with the considered direction. Since cosmic muons typically travel in nearly straight trajectories, this check is enough to confirm that a candidate track is a cosmic track.

Then there is the Kalman Filter Tracker, which aims to reconstruct tracks from charged particles such as muons, which tend to be long minimum-ionizing tracks. Starting from a small set of seed hits, the Kalman Filter algorithm extrapolates the candidate track into

the next plane and searches for compatible hits. At each step, the track parameters, such as position, direction, and an estimate of the momentum, are updated. The update uses the prediction from the previous step and the newly observed hit. The algorithm also keeps track of the parameter uncertainties, represented by a covariance matrix. This matrix evolves as the track grows, which shows how the confidence in the track state changes as more information is incorporated. The Kalman filter is able to account for effects such as multiple scattering and energy loss, yielding a robust reconstruction of muons and other track-like particles [83].

2.5.6 Deep Learning Classifiers

To distinguish between different interaction types, NOvA uses deep convolutional networks known as **Convolutional Visual Networks (CVN)** [87, 88, 89]. These networks treat the detector data as images, where each pixel corresponds to an individual cell and the pixel intensity represents the deposited energy.

Event CVN

The Event CVN is designed to classify entire neutrino events. It takes as input the two detector views (xz and yz) of a slice, each represented as a two-dimensional pixel map. These views are first convolved separately and later merged in a deep network architecture inspired by MobileNetV2 (Fig. 2.12). MobileNetV2 is a lightweight convolutional neural network originally developed for efficient image classification. In a standard convolution, each filter spans the full set of input channels (here the xz and yz pixel maps), which usually results in a relatively high computational cost. MobileNetV2 reduces this cost by using depthwise separable convolutions. They factorize the operation into two stages, where a depthwise convolution first applies partial filters independently to each detector view, then, 1×1 pointwise convolution combines the information across views. This reduces the number of model parameters compared to a full convolution, while retaining similar information.

In addition, MobileNetV2 uses inverted residual bottleneck blocks. In these layers, the input features are first expanded into a higher-dimensional space, then filtered with depth-wise convolutions, and finally projected back into a compact representation. To preserve important features in the compression step, a residual connection links the input and output of the block. This design enables the network to capture more detailed features while keeping efficiency. In this way, MobileNetV2 achieves high accuracy with far fewer parameters than a conventional CNN [90].

The Event CVN outputs four normalized scores corresponding to ν_μ CC, ν_e CC, NC, and cosmic events classification. The Event CVN has been trained on millions of simulated events, and it achieves classification efficiencies above 90% in validation samples. Its performance is summarized in the confusion matrices shown in Fig. 2.11, where both efficiency (left) and purity (right) are high across all event categories.




Figure 2.11: Confusion matrices for Event CVN classification in NOvA simulation. Left: efficiency (fraction of true events correctly classified). Right: purity (fraction of predicted events that are truly that class). The CVN achieves high efficiency and purity across ν_μ CC, ν_e CC, NC, and cosmic categories.

Prong CVN

While the Event CVN provides classification at the event level, the NOvA Prong CVN handles individual particle identification. It classifies prongs, i.e. groups of hits associated with outgoing particles from the vertex, as e , μ , π , γ , or p . When training the Prong CVN, two different image perspectives are used:

- **Cropped prong-centered images:** around each reconstructed prong, a smaller region of interest is first extracted from the detector view around each reconstructed prong. This zoomed-in image highlights the local features of the prong, which allows for a clearer distinction between the different particles. For example, a muon typically produces a narrow, straight profile compared to the diffuse spread of an electromagnetic shower from an electron. This cropping reduces noise from the rest of the event and lets the network focus on the candidate particle.
- **Whole-event context:** in parallel, we also feed the network the full event image. This gives it the global context, such as where the prong sits relative to the vertex and how the rest of the event is laid out. That context matters because some particle signatures can look very similar up close. For example, photon and electron showers often resemble each other locally, but only an electron shower should trace back to the interaction vertex.

The classifier is able to produce more accurate and robust outputs by combining these two inputs. This is particularly useful for cases like overlapping showers or secondary photons. The Prong CVN architecture is again based on MobileNetV2, using bottleneck blocks and depthwise separable convolutions (Fig. 2.12).

About 1.9 million simulated events were used for training. Each event is converted into pixel maps, where the pixel intensity corresponds to the energy deposited in detector cells.




Figure 2.12: Schematic layout of the NOvA Convolutional Visual Network (CVN) architecture, inspired by MobileNetV2. Bottleneck blocks, depthwise convolutions, and squeeze-excite modules allow efficient feature extraction from event or prong images.




Figure 2.13: Confusion matrices for Prong CVN classification in NOvA simulation. Left: efficiency. Right: purity. Prong CVN classifies individual prongs as electron, muon, proton, pion, or photon with high accuracy, complementing the Event CVN in particle identification.

One issue is that the network can become too sensitive to overall brightness, which we do not want because the absolute energy scale in real data is not perfectly fixed. To reduce this dependence, we randomly rescale the pixel intensities within a chosen range during training. This encourages the network to focus on topology and shape rather than absolute brightness. Figure 2.13 shows the resulting particle-identification performance: the separation is strong, with muons and electrons reaching efficiencies and purities above 90

Chapter 3. The NOvA 3-Flavor Oscillation Analysis

This chapter discusses the NOvA three-flavor oscillation analysis. It introduces the three-flavor framework in the context of NOvA and outlines the parameters to which the experiment is most sensitive. It then describes the framework currently used by NOvA to extract oscillation parameters and highlights its main limitations and possible improvements. Finally, it presents the most recent NOvA measurements of the oscillation parameters.

The MSW effect for the NOvA Experiment

From Section 1.8 Eq. 1.70, the equation of evolution including matter effect is:

$$i \frac{d}{dx} \psi_{\alpha\beta}(x) = \sum_{\eta} \left(\sum_k U_{\beta k} \frac{\Delta m_{k1}^2}{2E} U_{\eta k}^* + \delta_{\beta e} \delta_{\eta e} V_{CC} \right) \psi_{\alpha\eta}(x), \quad (3.1)$$

The full three-flavor oscillation probabilities can be obtained by solving for the evolution operator. In practice, there are tools such as the OscProb framework [91], which is used by NOvA, to numerically integrate the three-flavor equation in matter without approximations. However, approximate analytic expressions can be derived by expanding in the small parameters $\alpha \equiv \frac{\Delta m_{21}^2}{\Delta m_{31}^2}$ and $s_{13} \equiv \sin \theta_{13}$. These approximations are useful as they provide insight into how the matter potential, the CP-violating phase, and mass ordering modify the $\nu_{\mu} \rightarrow \nu_e$ and $\nu_{\mu} \rightarrow \nu_{\mu}$ probabilities.

3.0.1 Muon-to-Electron Appearance Probability $P_{\mu e}$

For neutrino propagation through matter of approximately constant density (such as the Earth's crust), the evolution equation can be solved perturbatively in α and s_{13} . In doing so, the parameter

$$\hat{A} \equiv \frac{2E V_{CC}}{\Delta m_{31}^2} \quad (3.2)$$

is kept to all orders [92, 93].

We further define

$$\Delta \equiv \frac{\Delta m_{31}^2 L}{4E},$$

and use the following notations $s_{ij} \equiv \sin \theta_{ij}$, $c_{ij} \equiv \cos \theta_{ij}$.

As a reminder, the Hamiltonian in flavor space can be written as:

$$H = \frac{1}{2E} U \begin{pmatrix} 0 & 0 & 0 \\ 0 & \Delta m_{21}^2 & 0 \\ 0 & 0 & \Delta m_{31}^2 \end{pmatrix} U^\dagger + \begin{pmatrix} V_{CC} & 0 & 0 \\ 0 & 0 & 0 \\ 0 & 0 & 0 \end{pmatrix}, \quad (3.3)$$

where the first term describes neutrino oscillations in vacuum, and the second term, the additional potential that electron neutrinos experience due to their interactions with matter.

To obtain approximate analytic probabilities, we switch to the *propagation basis*, where the largest terms of the Hamiltonian become diagonal. Since Δm_{31}^2 and Δm_{32}^2 are of the order of 10^{-3} eV^2 , while Δm_{21}^2 is of order 10^{-5} eV^2 , the terms proportional to Δm_{31}^2 and Δm_{32}^2 dominate. Therefore, the terms related to Δm_{21}^2 can be treated as perturbations.

The PMNS matrix can be written as:

$$U = U_{23}U_{13}U_{12} \quad (3.4)$$

To get the propagation basis, we apply the rotation:

$$\tilde{H} = U_{13}^\dagger U_{23}^\dagger H U_{23} U_{13} \quad (3.5)$$

With this transformation, the large atmospheric and reactor mixings no longer appear in the off-diagonal terms. The new Hamiltonian H' is nearly diagonal, with terms proportional to Δm_{31}^2 as leading terms. The remaining off-diagonal terms are proportional to θ_{12} and Δm_{21}^2 , which are treated as small perturbations. The Hamiltonian then can be separated into a dominant part and a perturbation:

$$\tilde{H} = \tilde{H}_0 + \tilde{H}_1. \quad (3.6)$$

with $\tilde{H}_0$ depending only on Δm_{31}^2 and the matter potential V_{CC} . $\tilde{H}_1$ contains the solar parameters θ_{12} and Δm_{21}^2 , and the CP phase δ_{CP} .

$\Delta m_{21}^2 \ll \Delta m_{31}^2$ means that we can define $\alpha \equiv \frac{\Delta m_{21}^2}{\Delta m_{31}^2} \ll 1$, which can be used to expand the evolution operator perturbatively as a function of α . Moreover, since $\sin \theta_{13}$ is relatively small, a combined expansion in α and $\sin \theta_{13}$ can be done. This expansion makes it possible to compute the oscillation amplitude $A_{\mu e}$ and, therefore, the probability $P_{\mu e}^{\text{mat}}$ analytically up to second order.

The resulting approximate probability in constant-density matter is [92]:

$$P_{\mu e}^{\text{mat}}(\nu) \simeq s_{23}^2 \sin^2 2\theta_{13} \frac{\sin^2[(1 - \hat{A})\Delta]}{(1 - \hat{A})^2} \quad (3.7)$$

$$+ 2\alpha s_{13} \sin 2\theta_{12} \sin 2\theta_{23} \cos(\Delta + \delta_{CP}) \frac{\sin(\hat{A}\Delta)}{\hat{A}} \frac{\sin[(1 - \hat{A})\Delta]}{(1 - \hat{A})} \quad (3.8)$$

$$+ \alpha^2 c_{23}^2 \sin^2 2\theta_{12} \frac{\sin^2(\hat{A}\Delta)}{\hat{A}^2}, \quad (3.9)$$

with Δ and $\hat{A}$ as defined above.

For antineutrinos, one replaces $\hat{A} \rightarrow -\hat{A}$ and $\delta_{CP} \rightarrow -\delta_{CP}$.

$P_{\mu e}$ can be subdivided into 3 different terms :

- **Leading term** (3.7) $\propto s_{13}^2 s_{23}^2$: this term drives the overall appearance rate. It contains the θ_{23} octant information through s_{23}^2 . It is modified by matter effects via the $(1 - \hat{A})$ term. Because $\hat{A} \equiv \frac{2EV_{CC}}{\Delta m_{31}^2}$, the sign of Δm_{31}^2 directly changes the direction of the matter effect. For neutrinos in the normal ordering ($\Delta m_{31}^2 > 0$), $\hat{A} > 0$ increases ν_e appearance, while in the inverted ordering ($\Delta m_{31}^2 < 0$), $\hat{A} < 0$ suppresses it. For antineutrinos, the opposite is true. This allows experiments like NOvA to probe mass ordering as the enhancement/suppression effect differs between normal and inverted orderings.
- **Interference term** (3.8): this term carries the δ_{CP} dependence through $\cos(\Delta + \delta_{CP})$. Since it changes sign under $\nu \leftrightarrow \bar{\nu}$ ($\hat{A} \rightarrow -\hat{A}$, $\delta_{CP} \rightarrow -\delta_{CP}$), the effect of this term is different between neutrino and antineutrino. Depending on the value of δ_{CP} , it can either increase ν_e appearance or diminish it. These differences are the source of CP sensitivity. However, since certain combinations of ordering and δ_{CP} can produce similar asymmetries, this introduces a degeneracy with the mass ordering. Breaking this degeneracy requires combining both neutrino and antineutrino appearance data, hence the advantage of a joint fit for neutrinos and antineutrinos.
- **Solar term** (3.9): of order $\mathcal{O}(\alpha^2)$. Although subdominant, this term is important to increase precision in the fits as it carries the dependence on the solar parameters θ_{12} and Δm_{21}^2 .

3.0.2 Sensitivity to Mass Ordering, Octant, and CP Violation in NOvA

The interplay between CP violation and matter effect can be illustrated using the bi-event plot in Fig. 3.1. The plot shows the expected number of ν_e appearance events (horizontal axis) and $\bar{\nu}_e$ appearance events (vertical axis) in the NOvA Far Detector. Each point on the plot corresponds to a particular choice of oscillation parameters.

For fixed values of the mixing angles and mass splittings, changing the CP-violating phase δ_{CP} moves the prediction along an ellipse. The blue ellipses correspond to the normal mass ordering, while the red ones correspond to the inverted ordering. The separation between the two sets of ellipses comes from the matter effect, where neutrinos and antineutrinos interact differently with the electrons inside the Earth. This difference depends on the sign of the parameter $\hat{A}$ in the oscillation probability.

The ν_μ disappearance probability in the three-flavor framework with matter effects can be written in a perturbative expansion [93] as

$$P_{\mu\mu} = 1 - \sin^2 2\theta_{23} \sin^2\left(\frac{\Delta m_{32}^2 L}{4E}\right) + \mathcal{O}(s_{13}^2, \alpha, \hat{A} s_{13}^2, \dots), \quad (3.10)$$

where the leading term depends on $\sin^2 2\theta_{23}$ and is symmetric under the transformation $\theta_{23} \rightarrow \frac{\pi}{2} - \theta_{23}$. Thus, ν_μ disappearance data alone cannot determine the value of θ_{23} , specifically whether it is in the lower octant (less than or greater than 45°), or in the upper octant.

The ν_e appearance probability, on the other hand, depends on $\sin^2 \theta_{23}$. Therefore, it can be used to determine the octant value of θ_{23} . From Eq. 3.7, the leading term is proportional to:

$$P_{\mu e}^{\text{lead}} \propto s_{13}^2 s_{23}^2 \frac{\sin^2[(1 - \hat{A})\Delta]}{(1 - \hat{A})^2},$$

so the overall rate of ν_e appearance (and $\bar{\nu}_e$) changes depending on whether θ_{23} is in the lower or upper octant. Therefore, combining ν_μ disappearance and ν_e appearance samples (for both neutrinos and antineutrinos) allows experiments such as NOvA to simultaneously constrain the θ_{23} octant, the CP-phase δ_{CP} , and the mass ordering.

To better understand how each oscillation parameter influences the observed event rates, it is useful to look at the probabilities near the first oscillation maximum ($\Delta \simeq \pi/2$), where NOvA has its highest sensitivity. In this region, the dominant terms can be written as:

$$P_{\mu\mu} \approx 1 - \sin^2 2\theta_{23} c_{13}^4, \quad P_{\mu e} \approx 4 s_{13}^2 s_{23}^2 \frac{1}{(1 - \hat{A})^2} + (\text{CP/matter interference}).$$

Here, the position and depth of the ν_μ disappearance dip determine $|\Delta m_{32}^2|$ and $\sin^2 2\theta_{23}$, while the overall ν_e and $\bar{\nu}_e$ appearance rates, provide information about δ_{CP} , the θ_{23} octant, and the mass ordering. A joint ν_μ disappearance and ν_e appearance fit, considering both neutrino and antineutrino runs, enables NOvA to separate the effects of δ_{CP} , the θ_{23} octant, and the mass ordering.

3.1 The NOvA Extrapolation method

A defining feature of NOvA's analysis is the use of the Near Detector (ND) to constrain unoscillated flux and interaction uncertainties. The procedure is:

1. Measure the ν_μ spectrum in the ND.
2. Compare ND data to ND Monte Carlo (MC), forming data/MC ratios that capture discrepancies in flux and cross-section modeling (Equation 3.4).
3. Apply these corrections to the FD Monte Carlo prediction before oscillations are applied.

To account for oscillation effects, oscillation weights are applied to the unoscillated FD




Figure 3.1: Bi-event plot for NOvA, showing expected ν_e vs. $\bar{\nu}_e$ appearance event counts at the Far Detector for normal hierarchy (NH, blue) and inverted hierarchy (IH, red). Each ellipse corresponds to variation of δ_{CP} over $[0, 2\pi]$. The displacement between the NH and IH ellipses comes from the matter effect, while the orientation of each ellipse is determined by δ_{CP} .




Figure 3.2: NOvA disappearance channel at $L = 810$ km (approximate vacuum expression), extended to 5 GeV. The first minimum occurs near $E \simeq 1.65$ GeV, with the dip depth set by $\sin^2 2\theta_{23} c_{13}^4$ and the position by $|\Delta m_{32}^2|$.




Figure 3.3: NOvA appearance channel at $L = 810$ km (leading term only). The matter effect enters via $(1 - \hat{A})^{-2}$.

spectrum to get the final expected FD prediction (Fig. 3.4). As the ND and FD are functionally identical, this ratio method allows the reduction or cancellation of systematic uncertainties that share the same effect between the two detectors, especially systematic uncertainties related to flux and cross-sections.

$$FD_{pred_{unoscillated}} = \frac{FD_{MC_{sim}}}{ND_{MC_{sim}}} \times ND_{data} \quad (3.11)$$




Figure 3.4: NOvA’s extrapolation technique - The ND measured data reconstructed energy spectrum (1) is first converted into a true energy spectrum (3) via the migration matrix (2). The ratio of the ND data and the simulated Monte Carlo in true energy bins (4) is calculated and used as a weight for the simulated events in the FD (6), along with the probability oscillation weight (5). Finally, the reconstructed-to-true migration matrix (7) is applied to the FD true energy spectrum to obtain the FD prediction (8).

Although the extrapolation technique is effective in reducing the systematic uncertainties exhibiting similar effects in the ND and FD, it still has a few drawbacks. The main one is that, since the ND information enters the method via the ratio, the extrapolation approach does not take full advantage of the high-statistics samples in the NOvA ND, which can be crucial in further constraining the systematic uncertainties. As a result, there have been several efforts to develop two-detector fits for the NOvA 3-Flavor analysis. One such method is the PISCES joint ND+FD fit, which is the focus of this thesis. In the PISCES framework, the

ND samples can be fit simultaneously with the FD samples, fully leveraging information from the high-statistics ND samples. Another downside of the current 3-Flavor framework is that it uses the ROOT MINUIT fitter [94], in which the systematic uncertainties are incorporated as nuisance parameters. This sets a limit on the number of systematic uncertainties that can be incorporated into the analysis, as in practice, MINUIT cannot properly handle nuisance parameters past a certain number.

Currently, this is addressed in the official 3-Flavor analysis by combining systematic uncertainties and excluding those with negligible effects; however, PISCES provides an alternative solution to overcome this limitation. In the PISCES framework, the systematic uncertainties are encoded in a covariance matrix; therefore, it can accommodate as many uncertainties as necessary. Another benefit brought forth by PISCES is that it treats statistical uncertainties as a Poisson distribution, offering an accurate treatment of the low-statistics samples in NOvA inaccessible through the more common Gaussian approach, especially for the low event-rate ν_e samples in the FD. All of these factors strongly motivated the exploration of the PISCES two-detector fit for the NOvA 3-Flavor analysis.

3.1.1 NOvA Fitting Framework

To extract the oscillation parameters, the neutrino samples observed in the Far Detector data are analyzed. For that, a data fit using the modified Far Detector model predictions using the extrapolated method (section 3.1) is performed. The Far Detector spectra are binned event counts, with relatively limited statistics, especially for the ν_e RHC samples. The samples i follow a Poisson distribution with mean prediction $N_i^{\text{pred}}(\vec{\theta}, \vec{\eta})$ for a given oscillation hypothesis $\vec{\theta}$ and nuisance (systematic) parameters $\vec{\eta}$.

The probability to observe N_i^{obs} events, following a Poisson distribution with mean $N_i^{\text{pred}}(\vec{\theta}, \vec{\eta})$, is [95] :

$$P\left(N_i^{\text{obs}} \mid N_i^{\text{pred}}(\vec{\theta}, \vec{\eta})\right) = \frac{\left[N_i^{\text{pred}}(\vec{\theta}, \vec{\eta})\right]^{N_i^{\text{obs}}}}{N_i^{\text{obs}}!} e^{-N_i^{\text{pred}}(\vec{\theta}, \vec{\eta})}. \quad (3.12)$$

Assuming independent bins, the likelihood of the observed data given the model is the product

$$\mathcal{L}(\vec{\theta}, \vec{\eta}) = \prod_{i=1}^{N_{\text{bins}}} \frac{[N_i^{\text{pred}}(\vec{\theta}, \vec{\eta})]^{N_i^{\text{obs}}}}{N_i^{\text{obs}}!} e^{-N_i^{\text{pred}}(\vec{\theta}, \vec{\eta})}. \quad (3.13)$$

Taking the logarithm of the likelihood converts the product in Eq. (3.13) into a sum, improving numerical stability and simplifying the algebra. Multiplying by (-2) defines a test statistic that, in the large-sample limit, asymptotically follows a χ^2 distribution. Thus, we obtain:

$$-2 \ln \mathcal{L}(\vec{\theta}, \vec{\eta}) = -2 \sum_{i=1}^{N_{\text{bins}}} \left[N_i^{\text{obs}} \ln N_i^{\text{pred}}(\vec{\theta}, \vec{\eta}) - N_i^{\text{pred}}(\vec{\theta}, \vec{\eta}) - \ln(N_i^{\text{obs}}!) \right]. \quad (3.14)$$

Dropping constant terms independent of $(\vec{\theta}, \vec{\eta})$ and rearranging by adding and subtracting $N_i^{\text{obs}} \ln N_i^{\text{obs}}$, one obtains the standard binned Poisson χ^2 [96]:

$$\chi^2(\vec{\theta}, \vec{\eta}) = 2 \sum_{i=1}^{N_{\text{bins}}} \left[N_i^{\text{pred}}(\vec{\theta}, \vec{\eta}) - N_i^{\text{obs}} + N_i^{\text{obs}} \ln \frac{N_i^{\text{obs}}}{N_i^{\text{pred}}(\vec{\theta}, \vec{\eta})} \right]. \quad (3.15)$$

Equation (3.15) is exact for Poisson-distributed counts and reduces, in the Gaussian (large-sample) limit, to the Pearson form

$$\chi^2 \simeq \sum_{i=1}^{N_{\text{bins}}} \frac{(N_i^{\text{obs}} - N_i^{\text{pred}})^2}{N_i^{\text{pred}}}. \quad (3.16)$$

Systematic uncertainties are incorporated into the fit as nuisance parameters, noted η_j . They represent effects such as flux normalization, cross-section modeling, or detector calibration, which can modify the predicted spectra considerably. The nuisance parameters are usually determined from external measurements or calibration studies. They take the form of a Gaussian distribution centered at zero (no shift) with width σ corresponding to the estimated 1σ uncertainty.

This corresponds to multiplying the likelihood by $\prod_j \exp[-\eta_j^2/(2\sigma_j^2)]$, or equivalently adding quadratic penalty terms to the test statistic:

$$\chi^2(\vec{\theta}, \vec{\eta}) = 2 \sum_{i=1}^{N_{\text{bins}}} \left[N_i^{\text{pred}}(\vec{\theta}, \vec{\eta}) - N_i^{\text{obs}} + N_i^{\text{obs}} \ln \frac{N_i^{\text{obs}}}{N_i^{\text{pred}}(\vec{\theta}, \vec{\eta})} \right] + \sum_j \left(\frac{\eta_j}{\sigma_j} \right)^2. \quad (3.17)$$

Here, the oscillation parameters $\vec{\theta}$ are the quantities of interest, while the nuisance parameters $\vec{\eta}$ are profiled: for each oscillation hypothesis, χ^2 is minimized with respect to $\vec{\eta}$, so that their uncertainties and correlations with the physics parameters are propagated consistently into the final confidence intervals. Equation (3.17) is the likelihood function minimized in the NOvA 3-Flavor oscillation analyses [51, 97, 98].

Confidence intervals and coverage. Confidence intervals in oscillation analyses represent the regions of parameter space (for one or more parameters) that are statistically compatible with the observed data at a given confidence level. In the likelihood framework, these intervals are constructed using the likelihood-ratio statistic

$$\Delta\chi^2(\vec{\theta}) \equiv \chi^2(\vec{\theta}) - \chi_{\text{min}}^2, \quad (3.18)$$

where χ_{min}^2 is the global best fit. $\Delta\chi^2$ is small in regions where the model parameters explain the data almost as well as the best-fit point.

A confidence interval is the set of parameter values whose $\Delta\chi^2$ is below a critical value. The cutoff is taken from the chi-square distribution via Wilks' theorem and depends on the number of parameters of interest. For one parameter, the 68.3% level corresponds to $\Delta\chi^2 = 1.0$; for two parameters it is $\Delta\chi^2 = 2.30$. In the Gaussian limit, the 68.3% level is the 1σ interval: if the measurement were repeated many times, about 68.3% of the estimates would fall inside that interval. Likewise, the 2σ and 3σ intervals correspond to 95.4% and 99.73% confidence levels, respectively, containing the true value in 95.4% or 99.73% of repeated

experiments. Table 3.1 summarizes the most commonly used critical values.

Table 3.1: Critical values of $\Delta\chi^2$ from the χ^2 distribution, corresponding to common confidence levels for one and two parameters of interest.

Degrees of freedom	68.3% (1σ)	90%	95%
1	1.00	2.71	3.84
2	2.30	4.61	5.99

3.1.2 Feldman-Cousins Corrections

As discussed above, confidence intervals are commonly constructed using the likelihood-ratio statistic $\Delta\chi^2(\vec{\theta}) = \chi^2(\vec{\theta}) - \chi^2_{\min}$, and interpreted according to Wilks' theorem. In the large-sample limit, $-2\ln\lambda$ is distributed as χ^2 with degrees of freedom equal to the number of parameters of interest. However, in neutrino oscillation experiments such as NOvA, two important complications can invalidate the Wilks' approximation:

- **Low statistics:** In appearance channels we often observe only $\mathcal{O}(100)$ events or fewer. In this regime, Poisson fluctuations are non-negligible, and the test statistic may deviate from the asymptotic χ^2 form.
- **Bounded parameters:** Some oscillation parameters are bounded, for example $\sin^2\theta_{23} \in [0, 1]$ and nonnegative event rates. When the true value lies near a boundary, the test-statistic distribution can differ from the usual χ^2 expectation.

To account for these effects, NOvA uses the *Feldman-Cousins (FC) procedure* [99, 100], which constructs confidence intervals from Monte Carlo pseudo-experiment ensembles instead of relying on the asymptotic assumption. The method proceeds as follows:

1. For each parameter point $\vec{\theta}$, generate many pseudo-experiments from the prediction $N_i^{\text{pred}(\vec{\theta}, \vec{\eta})}$, with Poisson fluctuations and with systematic variations included.

2. For each pseudo-experiment, compute the test statistic $\Delta\chi^2(\vec{\theta})$ relative to the global best fit, and collect the resulting distribution under the hypothesis $\vec{\theta}$.
3. Choose the critical value $\Delta\chi_{\text{crit}}^2(\vec{\theta})$ such that a fraction α (e.g. 90%) of the pseudo-experiments satisfy $\Delta\chi^2 \leq \Delta\chi_{\text{crit}}^2$.
4. At confidence level α , we take the confidence region to be the set of parameter points $\vec{\theta}$ where the observed $\Delta\chi^2(\vec{\theta})$ is at or below the critical value $\Delta\chi_{\text{crit}}^2(\vec{\theta})$.

This construction has correct frequentist coverage, meaning that even with low statistics or bounded parameters, the region contains the true value with the stated probability. The FC construction is computationally intensive because it requires generating ensembles of pseudo-experiments at many grid points in parameter space. For this reason, NOvA typically applies Feldman-Cousins corrections to the most relevant parameters, such as $\sin^2 \theta_{23}$ and Δm_{32}^2 , and verifies that coverage is maintained in joint fits including δ_{CP} [51, 97, 98].

3.2 Recent NOvA Results

NOvA has been collecting data since February 6, 2014. The latest official results from NOvA use neutrino data collected from the start of data taking until March 13, 2024, representing about 10 years of data. This corresponds to a neutrino beam exposure of 26.61×10^{20} protons-on-target (POT), and an antineutrino beam exposure of 12.50×10^{20} POT. NOvA recorded a total of 181 ν_e and 384 ν_μ candidates in FHC mode, and 32 $\bar{\nu}_e$ and 106 $\bar{\nu}_\mu$ in RHC mode [2]. Table 3.3 shows the observed neutrinos and their origin, and the corresponding expected MC prediction. After fitting the data, the following best-fit values were extracted [2]:

$$\Delta m_{32}^2 = 2.441_{-0.031}^{+0.035} \times 10^{-3} \text{ eV}^2$$

in the Normal MO, and

$$\Delta m_{32}^2 = -2.481_{-0.034}^{+0.034} \times 10^{-3} \text{ eV}^2$$

in the Inverted MO. The Inverted MO is rejected at a 1.4σ confidence level. The best-fit point for δ_{CP} is $0.87^{+0.30}_{-0.90}\pi$ and $\sin^2 \theta_{23} = 0.55^{+0.02}_{-0.06}$. These results represent the most precise single-experiment measurement of Δm_{32}^2 with an uncertainty of 1.5 %. Figure 3.5 represents the 90 % credible regions in the Δm_{32}^2 vs $\sin^2 \theta_{23}$ parameter space for NOvA, considering Normal mass ordering, compared to different experiments [2].

Table 3.2: Impact of each uncertainty source on the fitted parameters. Entries are the change in the best fit when each source is varied by its 1σ range, quoted as $+/-$ shifts.

Source	$\sin^2 \theta_{23}$	δ_{CP}	$ \Delta m_{32}^2 $
Beam flux	+0.07/ $-$ 0.14	+0.18/ $-$ 1.15	+0.02/ $-$ 0.05
Calibration	+0.58/ $-$ 2.91	+1.95/ $-$ 19.5	+0.57/ $-$ 0.74
Detector model	+0.07/ $-$ 0.71	+0.43/ $-$ 3.79	+0.07/ $-$ 0.11
Lepton Reco.	+0.49/ $-$ 0.82	+0.70/ $-$ 3.33	+0.37/ $-$ 0.57
ND–FD Uncor.	+0.44/ $-$ 0.45	+0.91/ $-$ 4.60	+0.10/ $-$ 0.14
Cross Sections	+0.56/ $-$ 0.93	+2.30/ $-$ 11.2	+0.27/ $-$ 0.45
Neutron model	+0.49/ $-$ 0.09	+0.49/ $-$ 1.49	+0.10/ $-$ 0.16
Systematic Unc.	+1.20/ $-$ 3.27	+3.33/ $-$ 24.1	+0.70/ $-$ 0.94
Statistical Unc.	+4.00/ $-$ 15.5	+7.59/ $-$ 85.1	+1.31/ $-$ 1.80




Figure 3.5: Allowed regions in the $(\sin^2 \theta_{23}, \Delta m_{32}^2)$ plane. The solid band shows the current NOvA result (90% C.L.), with previous NOvA results and external constraints overlaid for comparison. Contours are marginalized over mass ordering as indicated (Normal MO shown here). The “B” and “F” superscripts indicate Bayesian or Frequentist analyses, respectively.

Table 3.3: Predicted event composition by sample for neutrino and antineutrino beam running. Signal and background totals include systematic uncertainties (superscript/subscript).

Sample	Neutrino Beam			Antineutrino Beam	
	ν_μ	ν_e	Low-energy	$\bar{\nu}_\mu$	$\bar{\nu}_e$
$\nu_\mu \rightarrow \nu_\mu$	372.3	4.3	0.3	24.4	0.2
$\bar{\nu}_\mu \rightarrow \nu_\mu$	24.5	0.1	0.0	71.5	0.2
$\nu_\mu \rightarrow \nu_e$	0.4	125.3	3.4	0.0	2.1
$\bar{\nu}_\mu \rightarrow \bar{\nu}_e$	0.0	1.8	0.1	0.0	18.9
Beam $\nu_e + \bar{\nu}_e$	0.1	26.1	0.8	0.0	6.5
NC	5.5	16.8	5.3	0.8	2.0
Cosmic	4.4	5.5	0.5	0.7	1.1
Others	1.5	0.8	0.1	0.2	0.1
Signal	$397.6^{+26.8}_{-24.2}$	$125.3^{+5.1}_{-5.1}$	$3.4^{+1.0}_{-0.8}$	$96.0^{+5.2}_{-5.2}$	$18.9^{+0.9}_{-0.9}$
Background	$11.0^{+1.3}_{-1.2}$	$55.4^{+2.6}_{-1.7}$	$7.1^{+1.3}_{-1.1}$	$1.7^{+0.2}_{-0.2}$	$12.2^{+0.5}_{-0.7}$
Predicted	408.6	180.7	10.5	97.7	31.1
Observed	384	169	12	106	32

Chapter 4. Systematic Uncertainties

Given NOvA’s accumulated large data samples, the statistical uncertainties in the FD are comparable to systematic uncertainties for the three-flavor oscillation analysis, so it is essential to correctly identify and quantify the systematic sources of error in the analysis. These uncertainties arise from imperfect knowledge of the neutrino flux, interaction cross sections, and detector response, as well as from external calibration and modeling limitations. This chapter gives an overview of the different systematic uncertainties considered in NOvA and the motivation behind them.

In this analysis, a total of about 160 individual systematic uncertainties are considered. They are grouped into six major categories according to their physical origin and implementation:

- **Flux uncertainties**, describing uncertainties in the production and focusing of hadrons in the NuMI beamline.
- **Detector systematics**, encompassing calibration, light modeling, Cherenkov light contributions, and aging effects in the Near and Far Detectors.
- **Neutrino cross-section uncertainties**, covering the modeling of CCQE, RES, MEC, and DIS interactions, as well as radiative and second-class current effects.
- **Final-State Interaction (FSI) uncertainties**, describing the intranuclear propagation and reinteraction of hadrons, primarily pions, within the GENIE hN cascade model.

- **Neutron response uncertainties**, associated with discrepancies in neutron transport modeling between the GEANT and MENATE simulations.
- **Normalization uncertainties**, which represent global rate-scale effects such as proton-on-target counting, detector mass, and data/MC efficiency differences.

Each category is implemented through reweightable "systematic knobs" in the simulation framework, which are propagated through the fit. The resulting shifts are incorporated into the covariance matrices for the PISCES case.

4.1 Flux uncertainties

As discussed in chapter 2, the neutrinos detected by NOvA are obtained from the NuMI beamline, where protons are accelerated before hitting a graphite target to produce secondary hadrons, mainly pions and kaons. Magnetic horns are used to focus these charged mesons, directing particles of one sign into the beam while defocusing mesons with the opposite charge. The focused pions and kaons then decay to produce ν_μ and $\bar{\nu}_\mu$ in the decay pipe.

The pion and kaon production cross sections on the target material are defined up to a certain precision. However, multiple interactions can occur within the target and in downstream materials, making the determination of the neutrino flux more challenging. In addition, pion and kaon production is governed by non-perturbative QCD, so purely theoretical predictions are limited in accuracy. For this reason, NOvA uses a detailed beamline simulation together with phenomenological models of hadronic cascades tuned to external hadron production data. Differences between data and model give rise to systematic uncertainties on the predicted neutrino flux.

In NOvA, the flux central value and its associated uncertainties are obtained from the Package to Predict the Flux (PPFX), a tool originally developed in the MINERvA Experiment [69]. PPFX is a flux prediction framework that not only exploits all available

hadron-production data but also incorporates measurements of particle production from thin targets to reduce hadron-production uncertainties. PPFX is used to generate a large number of simulations with random fluctuations on given parameters.

On top of the hadron production uncertainties, additional uncertainties arise from the beam transport and beam focusing configuration. In the NOvA analysis, the following beam transport parameters were varied around their nominal values [1] :

- the horn currents were fluctuated around ± 2 kA with the nominal horn current of ± 200 kA (for FHC and RHC)
- the positions of horn 1 and 2 are shifted around ± 3 mm in the X and Y directions separately
- the beam position on the target is shifted by ± 1 mm in X and Y separately
- the beam spot size is shifted by ± 0.2 mm in both the X and Y axis from the nominal size of 1.3 mm.
- the thickness of the water layer on the horns is shifted by ± 1 mm
- the target position is shifted in the Z axis by ± 7 mm

These shifts correspond to the estimated 1σ uncertainties from external constraints on the beamline parameters obtained through alignment surveys and instrumentation analyses.

Tables 4.1 - 4.4 provide estimates for the most impactful beam focusing systematics. For each horn or beam-transport variation, the tables list the neutrino flux integral between 1 and 3 GeV, together with the peak position and width of a Gaussian fit in this energy range. The last column gives the percentage difference between the shifted and nominal integrals [1].

Table 4.1: Effect of selected focusing systematics on the FHC ND neutrino flux. For each variation, the table shows the flux integral between 1 and 3 GeV, the peak position and width of a Gaussian fit in that range, and the percentage change of the integral relative to the nominal prediction. Only variations with $|\Delta I| \geq 1.5\%$ are listed [1].

Systematic	Integral 1-3 GeV [10^{-14}]	Peak [GeV]	Width [GeV]	ΔI [%]
Nominal	4.985	1.9799	0.5274	0.00
+3mmHorn1X	5.188	1.9854	0.5311	4.08
-3mmHorn1X	4.789	1.9746	0.5249	-3.93
+3mmHorn1Y	4.910	1.9783	0.5269	-1.50
+3mmHorn2X	5.117	1.9907	0.5318	2.66
-3mmHorn2X	4.815	1.9663	0.5232	-3.42
+1mmHornWater	4.899	1.9805	0.5295	-1.72
-1mmHornWater	5.077	1.9805	0.5265	1.85

Table 4.2: Same as Table 4.1, but for the FHC FD neutrino flux.

Systematic	Integral 1-3 GeV [10^{-20}]	Peak [GeV]	Width [GeV]	ΔI [%]
Nominal	4.665	1.9956	0.3586	0.00
+3mmHorn1X	4.827	1.9994	0.3598	3.49
-3mmHorn1X	4.510	1.9919	0.3590	-3.32
+3mmHorn1Y	4.590	1.9940	0.3588	-1.61
+3mmHorn2X	4.778	2.0037	0.3611	2.42
-3mmHorn2X	4.522	1.9855	0.3562	-3.06
+1mmHornWater	4.568	1.9960	0.3615	-2.08
-1mmHornWater	4.753	1.9959	0.3560	1.88
Magnetic field in decay pipe	4.802	2.0087	0.3612	2.94

Table 4.3: Same as Table 4.1, but for the RHC ND neutrino flux.

Systematic	Integral 1-3 GeV [10^{-14}]	Peak [GeV]	Width [GeV]	ΔI [%]
Nominal	2.591	2.0086	0.5453	0.00
+3mmHorn1X	2.703	2.0140	0.5494	4.32
-3mmHorn1X	2.470	2.0036	0.5420	-4.68
+3mmHorn1Y	2.544	2.0063	0.5444	-1.82
+3mmHorn2X	2.657	2.0217	0.5499	2.54
-3mmHorn2X	2.525	1.9950	0.5411	-2.53
+1mmHornWater	2.552	2.0091	0.5489	-1.50

Table 4.4: Same as Table 4.1, but for the RHC FD neutrino flux.

Systematic	Integral 1-3 GeV [10^{-20}]	Peak [GeV]	Width [GeV]	ΔI [%]
Nominal	2.580	2.0114	0.3807	0.00
+3mmHorn1X	2.686	2.0151	0.3823	4.13
-3mmHorn1X	2.471	2.0081	0.3807	-4.20
+3mmHorn1Y	2.529	2.0097	0.3808	-1.97
+3mmHorn2X	2.633	2.0209	0.3833	2.08
-3mmHorn2X	2.526	2.0013	0.3785	-2.07
Magnetic field in decay pipe	2.528	1.9992	0.3794	-2.00

To propagate these uncertainties into the oscillation analysis, NOvA combines the hadron-production and focusing variations into an ensemble of flux "universes". In the flux systematics study used for this work, 100 hadron-production universes are generated in PPFX. In addition, for each of these, 20 universes are produced with random variations of the beam transport parameters within their external constraints, yielding a total of 2000 universes that sample the full flux uncertainty. For each universe, the predicted Near and Far Detector spectra are computed, and the spread of these spectra is used to build a flux covariance matrix [101].

A Principal Component Analysis (PCA) is then performed on this covariance matrix, producing a set of orthogonal eigenvectors labeled "flux principal components". The amplitudes of these principal components are used as nuisance parameters in the oscillation fit, instead of the original individual hadron production and beamline parameters. In the PISCES analysis, the first 50 flux principal components are retained.

4.2 Detector systematics

4.2.1 Calibration energy scale

Calibration uncertainties are one of the most impactful systematic uncertainties for NOvA. These uncertainties arise from the imperfect modeling of the detector energy response, i.e., the conversion from detector signals to deposited particle energy. As particles travel through

the NOvA detectors, the deposited energy is converted into scintillation light, which the APD subsequently records. The analog signal recorded by the APD is then converted into a digital signal by the ADC. Calibration uncertainties quantify how accurately we can convert the recorded detector signal (for example, the ADC counts) into the corresponding particle energy measured in MeV. In NOvA, specific particle samples in the detectors are used to provide a probe of the overall calorimetric energy scale of the NOvA detectors. Specifically, the mass peak of the π^0 , the Michel electrons, and the energy loss per unit length dE/dx of the beam muons and beam protons are used as standard candle control samples :

- **π^0 mass peak:** In neutrino interactions, neutral pions usually decay via $\pi^0 \rightarrow \gamma\gamma$, where each photon is reconstructed as an electromagnetic (EM) shower. From the reconstructed photon energies E_1 , E_2 and their opening angle θ_{12} , the invariant mass $m_{\gamma\gamma}^2$ is obtained as

$$m_{\gamma\gamma}^2 = 2E_1E_2(1 - \cos\theta_{12}).$$

In an ideal detector, the reconstructed $m_{\gamma\gamma}$ should peak at the true π^0 mass (≈ 134.98 MeV). Therefore, evaluating the measured $m_{\gamma\gamma}$ peak from π^0 samples helps determine potential bias in the EM energy scale of the detector. In NOvA, the π^0 sample data is compared to the MC simulation, taking into account the detector responses, and the modeling mismatch in the EM energy scale is quantified as

$$\Delta = \frac{m_{\pi^0}^{\text{data}} - m_{\pi^0}^{\text{MC}}}{m_{\pi^0}^{\text{MC}}}.$$

- **Michel electron spectrum:** Many muons stop in the detector and decay at rest via $\mu^\pm \rightarrow e^\pm + \nu + \bar{\nu}$. The energy spectrum of the electron from this decay (the Michel electron) is well defined from theory and previous precision measurements. Specifically, the Michel electron energy spectrum is continuous with a kinematic endpoint (maxi-

mum energy) at 52.8 MeV [102]. Therefore, measurements of the reconstructed Michel electron energies can also be used to evaluate the detector calibration, especially for low-energy electromagnetic showers.

- **Beam muon dE/dx :** Muons leave long tracks in the detector. The mean energy loss per unit length dE/dx for the muon, as a function of momentum, is given by the Bethe-Bloch formula and is very well constrained. For these long beam muon tracks, the reconstructed dE/dx is studied as a function of several variables and compared between data and simulation. For instance, an analysis of dE/dx as a function of position in the detector can reveal regions with potential calibration offsets.
- **Beam proton dE/dx :** Unlike muons, protons lose significantly more energy as they travel through the detector. The energy deposited dE/dx by the proton is moderate at first, but then increases sharply as the proton slows down before stopping. This sharp increase, known as the Bragg peak, gives the proton dE/dx a characteristic rising shape near the end of the track. This distinctive feature allows for a test of how well hadronic energy deposition and scintillator response are modeled in the detector, by comparing the reconstructed beam proton dE/dx distributions in data and simulation.

Uncertainties in the absolute energy scale are estimated from differences between measurements and simulations from those different samples. Analyses showed that the largest uncertainty stemmed from the ND beam proton samples, at the level of about 5% [103] [104]. This 5% value is taken as the calibration systematic uncertainty and is implemented in the simulation as a global energy-scale variation. Two dedicated calibration samples are produced, one with the reconstructed energy of all hits increased by +5 % ("Calib Up") and another decreased by -5% ("Calib Down"). These samples are then propagated through the full event reconstruction and selection chain to generate the calibration systematics used as systematic knobs in the oscillation fit.

4.2.2 Calibration shape

Analysis of the detector response showed that the energy response varies across the detector cells, with the response at the ends differing from that in the middle. Light signals generated far from the APD end are attenuated, as some photons are absorbed or scattered before reaching the readout. Although correction procedures were performed to account for these effects, small residual differences remain between the reconstructed and true energies. The calibration shape systematic uncertainty was introduced to cover these residual position-dependent biases. This systematic uncertainty is treated as fully correlated between the two NOvA detectors.

4.2.3 Detector aging

The number of hits recorded by the NOvA detectors appears to be decreasing gradually. This is primarily due to scintillator degradation and fiber attenuation. Over time, the scintillator's light yield declines progressively, and the efficiency of the wavelength-shifting fibers is reduced. These effects can be modeled as a linear decrease in the light model as a function of time. NOvA monitors detector aging by analyzing cosmic-ray muon samples, and the observed shift in the data suggested a decrease in the light model with a gradient of 4.5% per year [103]. This drift is implemented as a time-dependent correction to the light model, and then propagated as a systematic knob in the fits. This systematic uncertainty is also treated as fully correlated between the detectors.

4.2.4 Light level

The light-level systematic accounts for differences between two successive light-model calibrations of the NOvA detectors. The original light model was found to be suboptimal, as there were discrepancies between the data and MC, especially in the light-collection response. To address this, a tuning procedure was performed, including modifications to the spectrum

binning, limiting fit to high-statistics bins, and removing degenerate parameters from the fit [103]. After these changes, the updated light model parameters provided an improved data/MC agreement. The light-level uncertainties are defined as shifts in the view-scaling factors that determine the detectors’ overall light yield. The differences between the nominal and improved tune were used to produce a dedicated light level systematic sample that captures the possible bias in the light collection for each detector view. Table 4.5 summarizes the magnitude of the light-level shifts applied in the Near and Far Detectors.

Table 4.5: Summary of light-level systematic variations applied to the Near and Far Detectors [105].

Detector	View	Light Up (%)	Light Down (%)
Near Detector	X, Y views	+10	−10
Far Detector	X view	+16	−16
Far Detector	Y view	−6	+6

4.2.5 Cherenkov

The Cherenkov systematic uncertainty was introduced to account for discrepancies between data and simulation in proton dE/dx . It was determined that the simulated energy deposition for protons was about 5% lower than in data, compared to only 1.5% for muons. This suggested that the Cherenkov light model contribution in NOvA was underestimated. To correct for this, the scintillator yield Y_s and the Cherenkov scaling parameter ε_c were tuned, while ensuring that the muon response remained unchanged. The optimal configuration from the analysis showed that reducing Y_s by 4.4% and increasing ε_c by 83% brought the proton dE/dx data/MC discrepancy from around 5% to below 1% [105]. The difference between the nominal and tuned light models defines the Cherenkov light systematic uncertainty. As the same optical calibration applies to both detectors, this uncertainty is implemented as a fully correlated shift between the Near and Far Detectors.

4.3 Neutrino Cross Sections

Neutrino cross-section uncertainties represent some of the largest systematics in the NOvA 3-Flavor analysis. These are associated with neutrino-nucleus interactions in the detector, as several nuclear effects across a broad kinematical range need to be considered in the analysis.

As seen in Section 2.4.2, NOvA uses GENIE to simulate neutrino interactions, a choice driven in part by the fact that GENIE offers the freedom to choose the interaction models. Generally, preferences are given to the most theory-driven models. Nonetheless, it should be noted that comparing the MC to the NOvA data showcases that some tuning is still required to improve data-MC agreement (Section 2.4.3) [77].

The overwhelming majority of neutrino charged-current interactions that occur in the NOvA detectors are Charged Current Quasi-Elastic (QE) Scattering (Section 1.9.1), Deep Inelastic Scattering (DIS) (Section 1.9.4), Resonant (RES) neutrino interaction (Section 1.9.3), and two-particles two-holes scattering (Section 1.9.2).

4.3.1 Charged Current Quasi-Elastic (CCQE) Scattering

CCQE interactions are one of the most abundant interactions in NOvA and constitute key signals for oscillation searches. In CCQE interactions, the target nucleon has an initial momentum that can influence the observed scattering. This nucleon behavior can be modeled as independent particles moving in the mean field of the nucleus, with the initial-state momenta having a Fermi distribution. In the Fermi Gas model, the nucleons are filling some available states up to some initial state Fermi momentum k_f . This is why the Relativistic Fermi Gas (RFG) model is frequently used to describe the nucleus. Another factor to consider is the possibility that the outgoing nucleon also interacts with the target nucleus, which can be treated as a simple binding energy.

Within the free nucleon formalism, the CCQE interaction cross-section is the following :

$$\frac{d\sigma}{dQ_{QE}^2} = \frac{M^2 G_F^2 \cos^2 \theta_C}{8\pi E_\nu^2} \left\{ A(Q^2) \pm B(Q^2) \frac{s-u}{M^2} + C(Q^2) \frac{(s-u)^2}{M^4} \right\} \quad (4.1)$$

where Q^2 is the 4-momentum transfer expressed by $Q^2 = 2E_\nu(E_l - p_l \cos \theta_l) - m_l^2$; $G_F = 1.166 \times 10^{-5} \text{ GeV}^{-2}$ is the Fermi constant; $M = 938.92 \text{ MeV}$ is the average nucleon mass; $\cos \theta_C = 0.9742$ is the Cabbibo angle; s and u are the Mandelstam variables.

The variables $A(Q^2)$, $B(Q^2)$, and $C(Q^2)$ are:

$$A(Q^2) = \frac{m_l^2 + Q^2}{M^2} \left\{ \left(1 + \frac{Q^2}{4M^2}\right) F_A^2 - \left(1 - \frac{Q^2}{4M^2}\right) F_1^2 + \frac{Q^2}{4M^2} \left(1 - \frac{Q^2}{4M^2}\right) (\xi F_2)^2 + \frac{Q^2}{M^2} \text{Re}(F_1^* \xi F_2) \right. \\ \left. - \frac{Q^2}{M^2} \left(1 + \frac{Q^2}{4M^2}\right) (F_A^3)^2 - \frac{m_l^2}{4M^2} [|F_1 + \xi F_2|^2 + |F_A + 2F_P|^2 - 4 \left(1 + \frac{Q^2}{4M^2}\right) (F_V^3)^2 + F_P^2] \right\} \quad (4.2)$$

$$B(Q^2) = \frac{Q^2}{M^2} \text{Re}[F *_A (F_1 + \xi F_2)] - \frac{m_l^2}{M^2} \text{Re} \left[(F_1 - \tau \xi F_2) F_V^{3*} - \left(F *_A - \frac{Q^2}{2M^2} F_P \right) F_A^3 \right] \quad (4.3)$$

$$C(Q^2) = \frac{1}{4} \left\{ F_A^2 + F_1^2 + \tau (\xi F_2)^2 + \frac{Q^2}{2M^2} (F_A^3)^2 \right\} \quad (4.4)$$

where F_i are the different form factors.

The values of the form factors F_1 and F_2 have been determined from electron-nucleon scattering. The axial form factor F_A , however, is not known but can be parametrized as a dipole with the following expression :

$$F_A(Q^2) = \frac{F_A(0)}{\left(1 - \frac{q^2}{M_A^2}\right)^2} \quad (4.5)$$

where M_A is the axial mass and $F_A(0)$, the axial form factor at $Q^2 = 0$ which has been constrained from neutron beta decay.

Z-Expansion Uncertainties

Deuterium target experiments have allowed us to obtain a world average value of $1.014 \pm 0.014 \text{ GeV}$ for M_A . However, the uncertainties on the axial mass M_A can influence the

CCQE cross section considerably. In NOvA, the neutrino physics model used to simulate interactions uses the Valencia model for CCQE. In this approach, the form factor is written as a polynomial in a reparametrized variable $z(Q^2)$ [106]:

$$z(Q^2) = \frac{\sqrt{t_{\text{cut}} + Q^2} - \sqrt{t_{\text{cut}}}}{\sqrt{t_{\text{cut}} + Q^2} + \sqrt{t_{\text{cut}}}}, \quad (4.6)$$

$$F_A(Q^2) = \sum_{k=0}^{k_{\text{max}}} a_k z(Q^2)^k, \quad (4.7)$$

where t_{cut} is the pion-production threshold.

The advantage of the Z-expansion is that it is a model-independent description of the axial form factor and allows an easier quantification of systematic errors. Furthermore, the coefficients in the Z-expansion fall off as the order increases, which reduces the coefficients required to represent the form factor while also ensuring good convergence. Besides, the Z-expansion's flexibility provides a way of introducing more parameters as more data are available.

In this analysis, the systematic uncertainties associated with the Z-expansion are the ZExp_AxialFFSyst2020_EV[1-4], with 1σ bands ranging from 10 – 20% from the nominal prediction [107].

A normalization uncertainty is also applied to the number of CCQE events with the ZNormCCQE, for which the $+1\sigma$ value corresponds to +20% and the -1σ to -15% [107].

4.3.2 Random Phase Approximation Uncertainties

In the Relativistic Fermi Gas (RFG) model, nucleons inside the nucleus occupy momentum states only up to a sharp Fermi momentum. This description, however, is inconsistent with modern electron- and neutrino-scattering data, which show that roughly 20% of nucleons have momenta above the Fermi momentum due to correlations not captured in the simple RFG picture.

In NOvA, the nuclear model used in the analysis is the Local Fermi Gas (LFG) model, which improves upon the RFG by incorporating position-dependent densities and by allowing additional nuclear correlations. Long-range correlations, in particular, are included through the Random Phase Approximation (RPA).

The Random Phase Approximation (RPA) accounts for long-range nuclear correlations that are not captured by a simple Fermi-gas picture. Instead of treating each nucleon as responding independently to the neutrino interaction, RPA sums coherent particle-hole excitations in the nuclear medium, which modifies the effective nuclear response. This leads to a characteristic suppression of CCQE events at low momentum transfer, with a modest increase at the intermediate Q^2 , consistent with scattering data.

The uncertainties associated with the RPA modeling are implemented in this analysis via the parameters `RPACCQEEEnhSyst2020` and `RPACCQESuppSyst2020`. The 1σ variation of these systematic uncertainties represents around a 10% shift in the CCQE cross section in the relevant q_0 and Q^2 ranges.

4.3.3 Resonance interactions uncertainties

MaCCRES and MvCCRES

In CCRES interactions, the neutrino excites a baryon resonance (usually a $\Delta(1232)$) that then decays to $N\pi$. There is an axial form factor for the $N \rightarrow \Delta$ transition [41] [108]:

$$C_5^A(Q^2) = \frac{C_5^A(0)}{\left(1 + \frac{Q^2}{M_{A, \text{RES}}^2}\right)^2} \quad (4.8)$$

This dipole form factor describes the Q^2 dependence of the axial transition amplitude for the $N \rightarrow \Delta$ resonance process. $M_{A, \text{RES}}$ is the axial mass for resonant production. In GENIE, M_A^{RES} has a nominal value around 1.12 GeV, a value obtained from global fits to bubble-chamber data [109]. The assigned fractional uncertainty on the resonance axial mass

is 20%. This uncertainty is implemented via the **MaCCRES** systematic.

Similarly, there is a vector form factor M_v^{RES} in resonance production, which affects how the vector-current part of the transition changes with Q^2 [108]:

$$C_i^V(Q^2) \propto \frac{1}{\left(1 + \frac{Q^2}{M_{V,RES}^2}\right)^2} \quad (4.9)$$

The typical value for M_V^{RES} is around 0.84 GeV [109]. There is an uncertainty of order 10% on this parameter. This leads to cross-section variations that can reach 20% for high q^2 [110]. In the analysis, this uncertainty is represented by the **MvCCRES** systematic.

LowQ2RESSupp2020

The MINOS and MINERvA neutrino experiments observed lower CCRES events than expected from GENIE, especially at low momentum transfer Q^2 [111]. The Rein-Seghal resonance model used in GENIE is tuned to bubble-chamber data on free nucleons, but when applied to nuclear targets, it over-predicts the cross section at low Q^2 . The suppression of CCQE events observed in MINOS and MINERvA is attributed to nuclear effects, such as long-range correlations, which modify the Q^2 distribution for neutrino-nucleus interaction. Because of this, an ad-hoc, Q^2 -dependent low- Q^2 resonant suppression factor in GENIE was introduced to improve agreement with the MINOS/MINERvA CCRES data at low Q^2 . The associated uncertainty is called **LowQ2RESSupp**. It has the form:

$$1 - Ae^{-bQ^2} \quad (4.10)$$

with $A = 0.391$ and $b = 12.9$ for ν and $A = 0.429$ and $b = 20.9$ for $\bar{\nu}$. This uncertainty suppresses RES interactions by a factor of 10% at $Q^2 = 0.1 \text{ GeV}^2$ [110].

4.3.4 Deep-Inelastic Scattering interactions uncertainties

DIS normalizations

The non-resonant $N\pi$ contribution in the DIS region in GENIE is poorly constrained. Therefore, GENIE recommends assigning a large 50% uncertainty to these channels [109]. In NOvA, however, the systematics are implemented in a differential way instead of a single overall 50% uncertainty. Separate normalization systematics for $N\pi$ production, binned in pion multiplicity ($0, 1, 2, \geq 3$ pions), are introduced for ν and $\bar{\nu}$, CC and NC interactions, and neutron and proton targets. This results in $4 \times 2 \times 2 \times 2 = 32$ individual systematic uncertainties. These systematic uncertainties are dependent on the invariant mass W . For the region $W < 3$ GeV, where the modeling is least reliable, a 50% uncertainty is applied. Between $W = 3$ and 5 GeV, the uncertainties are reduced linearly from 50% to 5%, and for $W > 5$ GeV a 5% uncertainty is used. The most impactful DIS normalization systematics are $\text{DIS}\bar{\nu}\text{nCC3pi}$, $\text{DIS}\bar{\nu}\text{pCC1pi}$, $\text{DIS}\bar{\nu}\text{pCC3pi}$, $\text{DIS}\nu\text{nCC1pi}$, $\text{DIS}\nu\text{nCC2pi}$, $\text{DIS}\nu\text{nCC3pi}$, $\text{DIS}\nu\text{pCC0pi}$, $\text{DIS}\nu\text{pCC2pi}$, $\text{DIS}\nu\text{pCC3pi}$, where the naming represents the struck nucleon (n or p), the interaction type (CC or NC), and the number of pions in the final state.

DIS formation zone

In DIS interactions, GENIE produces outgoing partons that then hadronize into pions and nucleons. In a nuclear environment, these hadrons are not formed at the interaction point. Instead, the partons may propagate some distance before turning into hadrons. This distance is referred to as "formation zone". In GENIE, the formation zone is modeled by sampling an exponential distribution with a characteristic length scale [109]. The length of the formation zone determines how much intranuclear reinteraction the DIS hadrons undergo. If the formation zone is long, partons remain in a pre-hadronic state over a longer distance and often form hadrons outside the nucleus, so they experience fewer intranuclear interactions. If it is short, hadrons form earlier and are more likely to reinteract in the

nuclear medium. The `FormZone2020` uncertainty accounts for this by scaling the pion and nucleon formation zone. A 50% uncertainty is assigned to the nominal formation zone [110].

DIS hadronization model

The DIS hadronization uncertainties were included to modulate the final state hadron production in the DIS regime and in the Shallow Inelastic Scattering (SIS) regime. In the analysis, they are implemented via the `kDISNuHadroQ1Syst` and `kDISNuBarHadroQ0Syst` systematics. They affect ν DIS interactions that strike a neutron (or a proton for the $\bar{\nu}$ case) and have a final state with multiplicity 2, where the final state charge $Q = +1$ ($Q = 0$ for $\bar{\nu}$). These correspond to channels such as $(p + \pi^0)$ versus $(n + \pi^+)$ for neutrinos, and $(n + \pi^0)$ versus $(p + \pi^-)$ for antineutrinos [110]. The hadronization model systematics were introduced because the nominal DIS hadronization model in GENIE does not perfectly reproduce the neutrino-nucleus data on nuclear targets. Specifically, the relative rates of these two-body hadronic final states were found to be mismodeled in the ND data. To address this, the DIS hadronization systematics reweight those specific DIS/SIS channels while preserving the overall DIS cross-section.

These systematic uncertainties were originally developed and constrained for specific NOvA ND topological samples [110]. However, their effect is more general, as they modify the hadronic composition, visible hadronic energy, and the reconstructed prongs of DIS events. Therefore, they were also included in this analysis.

4.3.5 MEC Uncertainties

In the PISCES 3-Flavor analysis, the model-based uncertainties for MEC, initially developed for the NOvA Sterile analysis, were adapted for the 3-Flavor inputs. In this case, several MEC models were used to generate the systematic uncertainties, including the Valencia MEC (the default model in NOvA), the SUSY MEC, and the Dytman MEC. The motivation behind this is that the MEC systematics usually used in NOvA were tuned using

the NOvA ND data. In the PISCES two-detector fit, a preference is given to model-based systematics over ND data-tuned ones, hence the use of the MEC model spread uncertainties.

MEC shape uncertainties were produced by forming 1,000 different combinations of the predicted distributions from the aforementioned three models in the q_0 – q_3 space, each using a randomly selected set of weights. The method for combining these distributions is as follows:

- Three random numbers between 1 and 1000 were generated, then each number was normalized by dividing it by the sum of the three.
- Each of the three model distributions was scaled by one of the normalized random numbers.
- The scaled distributions were added together to create a unique model combination.
- The resulting combined distribution was divided by the default Valencia prediction to produce a weight that could be used to rescale the simulation.
- This procedure was repeated 1,000 times to generate a range of weights.

The resulting weights were then used to perform a Principal Component Analysis (PCA) to derive the corresponding systematic uncertainties. In the PCA, the weights were applied to generate spectra for the relevant samples, separated into hadronic energy quantiles and oscillation channels. These were then used to construct a fractional covariance matrix via:

$$S_{ij} = \frac{1}{n_u} \sum_u \left(\frac{\mu_{ui} + \epsilon}{m_i + \epsilon} \right) \left(\frac{\mu_{uj} + \epsilon}{m_j + \epsilon} \right) \quad (4.11)$$

where i, j denote the analysis bins, m is the Valencia central-value prediction, μ_{ui} is the shifted prediction for MEC weight universe u and bin i , and $\epsilon = 0.01$ is a regularization constant introduced to stabilize the ratio.

An eigenvalue decomposition of the covariance matrix was then performed to extract the principal components, yielding eigenvalues that characterize the scale of variation and

eigenvectors that describe the shape of the shifts in each bin. Only the first two components were found to be significant; therefore, only these were retained to construct the MEC systematic uncertainties [112].

4.3.6 Radiative corrections and second-class currents uncertainties

Radiative corrections

In the Standard Model, charged-current interactions are flavor-universal. However, the QED radiative corrections are not identical for ν_e and ν_μ due to the difference between the final-state lepton masses [113]. The radiative corrections for electrons and muons differ by a few percent at the NOvA energy scale. This leads to percent-level shifts in the ν_e/ν_μ and $\bar{\nu}_e/\bar{\nu}_\mu$ CCQE cross-section ratios [114]. Most event generators, including GENIE, do not implement full radiative corrections, so long-baseline experiments treat them as a theory uncertainty on the ν_e vs ν_μ cross sections. The T2K experiment studied these effects and adopted a model in which the residual theoretical uncertainty on the ν_e/ν_μ and $\bar{\nu}_e/\bar{\nu}_\mu$ cross-section ratios at around 1 GeV is taken to be about 2% [115] [113] [114].

NOvA follows this T2k prescription, and parametrizes it via two radiative correction uncertainties:

- **radcorrnu**: a 2% normalization uncertainty on the ν_e/ν_μ CC cross-section ratio.
- **radcorrnu****bar**: a 2% normalization uncertainty on the $\bar{\nu}_e/\bar{\nu}_\mu$ CC cross-section ratio.

They are treated as uncorrelated between neutrinos and antineutrinos. These systematic uncertainties act as a global multiplicative factor rescaling the predicted ν_e and $\bar{\nu}_e$ CC event rates relative to the ν_μ and $\bar{\nu}_\mu$ samples. They affect primarily the FD $\nu_e/\bar{\nu}_e$ appearance samples, exactly where mis-modelled ν_e cross-sections would bias δ_{CP} .

Second class currents

In CCQE scattering :

$$\nu_l + n \rightarrow l^- + p,$$

The hadronic current (the neutron to proton transition), is described by six Q^2 -dependent form factors: three vector (F_1^V , F_2^V , F_3^V), and three axial (F_A , F_P , F_3^A). In the limit of isospin symmetry and G-parity conservation (combination of charge conjugation and isospin rotation), the second-class form factors F_3^V and F_3^A must vanish. Any non-zero value of these form factors would therefore indicate G -parity violation and would add extra, lepton-mass-suppressed contributions to the CCQE cross section [113].

Experimental constraints from beta decay and related measurements strongly limit second-class currents. However, they do not rule out contributions at the few-percent level to the difference between ν_e and ν_μ cross sections in the energy range of interest of T2K and NOvA. Analyses have shown that allowing non-zero second-class form factors can change the ν_e/ν_μ and $\bar{\nu}_e/\bar{\nu}_\mu$ CCQE cross-sections ratios by order of 1-2% [113].

GENIE sets $F_3^V = F_3^A = 0$ by default. Therefore, the possible residual effect of second-class currents is treated as another theoretical uncertainty on ν_e vs ν_μ cross sections.

Here also, following T2K's model, NOvA implements the second-class currents uncertainty `2ndclasscurr` as a 2% normalization uncertainty that is fully anticorrelated between ν_e and $\bar{\nu}_e$. In this case, a $+1\sigma$ shift of `2ndclasscurr` scales the ν_e CC cross section up by 2% and the $\bar{\nu}_e$ CC cross section down by 2%, whereas a -1σ shift does the opposite. This reflects the theoretical expectation that second-class currents could affect neutrino and antineutrino cross sections in opposite ways, and can bias δ_{CP} measurement as it depends on comparing the ν_e and $\bar{\nu}_e$ appearance.

4.3.7 Final-state interaction uncertainties

When a neutrino scatters on a nucleon bound inside a nucleus, the outgoing hadrons (especially pions) undergo further interactions such as absorption, charge exchange or quasi-elastic scattering. These re-interactions are referred to as final-state interactions (FSI). FSIs can change the visible hadronic energy, as particles can be absorbed. They can also modify pion multiplicity and charge, and therefore the neutrino energy and the event topologies. A realistic FSI model, together with a quantified uncertainty, is therefore an essential part of the NOvA cross-section and oscillation framework.

For simulating FSI, NOvA uses the GENIE **hN** semi-classical intranuclear cascade model. In **hN**, hadrons are propagated step by step through the nucleus, and at each step GENIE evaluates interaction probabilities using theory-driven cross sections that are constrained by external π -nucleus scattering data [109]. Comparisons of the nominal **hN** model to $\pi^\pm$ -nucleus data show noticeable discrepancies, particularly in the absorption and total cross sections [116]. For this reason, NOvA performs a dedicated tuning of the **hN** parameters and builds reweightable uncertainties around this tuned central value [110].

The tuned NOvA **hN** model is expressed in terms of three pion-FSI "fate fractions", which represent the probabilities for absorption (f_{ABS}), single charge exchange (f_{CX}), and quasi-elastic scattering (f_{QE}), with an overall mean free path (MFP) that sets the total cross section. The tuning gives a set of best-fit values for $(f_{\text{ABS}}, f_{\text{CX}}, f_{\text{QE}}, \text{MFP})$ and a covariance matrix describing the correlated uncertainties on the three fate fractions. Following the strategy developed by T2K for a similar intranuclear cascade model [116], NOvA diagonalizes the 3×3 fate-fraction covariance matrix to obtain three eigenvalues and eigenvectors. Each eigenvector corresponds to a correlated variation of $(f_{\text{ABS}}, f_{\text{CX}}, f_{\text{QE}})$ that leaves the total reactive cross section approximately the same while modifying the relative probabilities of absorption, charge exchange, and quasi-elastic scattering [110].

NOvA encodes these as three reweightable fate-fraction eigenvector systematics, `hNFSI_FateFracEV1_2020`, `hNFSI_FateFracEV2_2020`, and `hNFSI_FateFracEV3_2020`. A $\pm 1\sigma$ shift of one of these knobs corresponds to adding or subtracting the appropriate eigenvalue-times-eigenvector combination to the tuned fate fractions. This defines three orthogonal uncertainty directions in the space of pion-FSI modes: for example, one direction may primarily increase absorption while decreasing charge exchange and quasi-elastic scattering, while another trades charge exchange against quasi-elastic. Since these variations conserve the total reactive cross section, they mainly change the distribution of pion topologies (CC0 π vs. CC1 π vs. multi- π) and the detailed hadronic final state, instead of simply scaling the overall event rate [110].

A fourth FSI systematic, `hNFSI_MFP_2020`, controls the pion mean free path. In `hN`, the MFP is inversely proportional to the total pion-nucleus cross section. Varying `hNFSI_MFP_2020` coherently increases or decreases the overall probability that a pion will interact inside the nucleus. In the NOvA implementation, a $\pm 1\sigma$ shift changes the tuned MFP by $\pm 33.3\%$, and the variation is taken to scale linearly up to $\pm 3\sigma$. A shorter MFP (negative shift) makes pions interact more often inside the nucleus, increasing absorption and other FSI channels and reducing the number and energy of outgoing pions. A longer MFP (positive shift) has the opposite effect, it allows more pions to escape with less modification to the primary interaction [110].

These FSI systematics induce $\mathcal{O}(5\text{-}10)\%$ variations in pion kinematic distributions in NOvA’s simulated samples [110]. This alters the visible hadronic energy, the multiplicity and charge of reconstructed prongs, and the migration of events between ND and FD topological categories.

In the PISCES fit, we follow the standard NOvA analysis and use the same tuned `hN` model and reweightable parameters. For the 3-flavor oscillation fit, one of the fate-fraction eigenvector parameters (the dominant mode in terms of impact on the ND samples), together

with the mean-free-path parameter `hNFSI_MFP_2020`, are included in the analysis. The full set of FSI knobs is usually used for dedicated studies.

4.4 Neutron uncertainties

Accurately modeling neutron production is crucial in neutrino oscillation experiments, as incorrect treatment of neutrons can lead to a significant bias in energy reconstruction. As mentioned earlier, neutrino interactions in NOvA are simulated using GENIE. The outgoing particles from the simulated final-state interactions are then passed to GEANT, which models particle transit through the detector by accounting for detector materials and geometry. GEANT provides the energy deposited by the particles as they travel through the detector, along with their eventual re-interactions. For neutrons specifically, since their interactions in the detectors tend to be more complex, a model called "MCNP" has been implemented in the NOvA GEANT simulation chain to model neutron transport.

Low-energy neutrons are a considerable background in the analysis. A study of Michel Electron samples in NOvA has shown a significant over-simulation of neutrons at low energy in the ND ν_e preselected sample. For neutrons below 20 MeV, the neutron transport model used in GEANT is the High Precision (HP) model. It is a model based on the numerous data available for low-energy neutron scattering. The HP process supports diverse neutron interactions such as elastic or inelastic scattering, nucleon capture, or fission reaction. In modeling neutrons, HP takes into account the cross-section and final state tables from the evaluated nuclear data files. In the previous GEANT version, the HP model simulated low-energy neutrons reasonably well; however, in the more recent version, an over-simulation issue appeared. Further analysis of the problem showed that a tweak in the HP model caused a sizable change in the low-energy neutron backgrounds [117].

Similar issues have been seen for a low-energy neutron-enhanced quasi-elastic sample, which has a significant impact on the 3-flavor analysis. The neutron-enhanced samples were

produced with the following criteria:

- Longest prongs, which are deemed to be likely muons, are removed.
- Prongs located within 20 cm of the neutrino interaction vertex are removed.
- Prongs that have more than 6 hits are removed.

With those selections, a neutron-enhanced sample that is 66% pure and 71% efficient is obtained for the RHC samples [117]. When comparing the simulated prongs to the data, an over-simulation happens, especially for the bins that contain photon daughters. Further studies have shown that it is during the intranuclear cascade, pre-compound, evaporation (IPE) stage of the interaction modeling in GEANT that we begin to see the discrepancies. In the IPE process, hadrons interact with the nucleus in the following way. First, a particle entering a nucleus will trigger an intranuclear cascade as it interacts with the nuclear medium. This process continues until no more outgoing particles have enough energy to escape the nuclear medium. After that cut-off, the nucleus enters a pre-equilibrium state in which remaining nuclear fragments are in an unevenly distributed excited state. During this stage, excitons are either created or destroyed until an optimum number is reached for equilibrium. It is in that stage that protons and neutrons are typically created. They can be emitted from the nucleus either as independent particles or as a recombined compound nucleus. Once equilibrium is reached, the simulation is passed to the evaporation model, which determines the probability to emit a given particle from an excited nucleus considering the excitation levels, the kinetic energy of the emitted particle and the cross-section of the case of particle absorption.

MENATE is an alternative neutron transport model that can be used instead of the IPE model in GEANT. What makes MENATE particular is that it simulates neutron transport in the MeV-scale energy in much greater detail than GEANT. Specifically, MENATE uses

a collection of nuclear physics data to tune a separate cross-section for all the possibilities that a neutron can inelastically scatter off the nucleus. Since NOvA’s liquid scintillator detectors are made of hydrocarbons, the main concerns in NOvA are neutron interactions with a carbon nucleus. With MENATE, the model is based on a specific neutron on carbon data-driven cross section, whereas GEANT relies on cascade simulation to determine neutron inelastic interactions.

With the neutron prongs over-simulation issue, there are up to 40% more prongs in the simulation using the GEANT IPE model compared to data for the lower energy neutrons. There is also a difference in shape in the MC to data ratio. When switching to MENATE, the discrepancy is reduced to a flat 20% difference [117]. In the analysis, the full simulation is still done with GEANT, with the neutron transport being handled by MCNP. What was done, however, is that a MENATE neutron sample has been generated and was used to create systematic uncertainties associated with neutrons.

4.5 Normalization systematics

The overall event rate predicted for a given NOvA sample scales as

$$N \propto \Phi \times \sigma \times \epsilon \times M \times \text{POT},$$

where Φ is the flux, σ the interaction cross section, ϵ the selection efficiency, M the detector mass, and POT the number of protons that hit the target. The normalization systematics are the uncertainties that scale the overall event rate, rather than changing the event distribution shape. In NOvA, the normalization uncertainties come from four main contributions: POT counting, data/MC selection efficiency (overlay), the far detector mass, and air bubbles in the detector, which were not included in the MC model.

4.5.1 POT normalization

The POT normalization systematic is the uncertainty on the POT number used to normalize the simulation to the data. The number of protons delivered to the NuMI target is monitored spill by spill, but the instrumentation and calibration of the beam monitoring system are not perfect. This leads to an overall $\sim 0.55\%$ uncertainty on the POT normalization [118]. Since the same beam delivers neutrinos to both detectors and to all running modes, this POT uncertainty is taken as fully correlated across ND and FD, FHC and RHC, and across all flavor channels.

4.5.2 Selection efficiency (overlay) normalization

Another source of normalization uncertainty comes from possible differences in selection efficiency between data and MC simulation. To evaluate the selection efficiency, NOvA uses an "overlay" technique, where single simulated MC neutrinos are overlaid into a subset of data and simulated ND spills. This allows an estimate of how many true neutrino interactions would be selected in data and MC. This way, the selection procedure in MC gives an efficiency directly comparable to data, and their difference is then treated as a normalization uncertainty. The overlay study yields separate uncertainties for ν_μ and ν_e samples, and for forward horn current (FHC) and reverse horn current (RHC) running ($\sim 0.21\%/0.48\%$ for ν_μ FHC/RHC, and $\sim 0.41\%/0.53\%$ for ν_e FHC/RHC) [119]. These are combined into a single representative value of

$$\textit{Overlay normalization uncertainty} \simeq 0.5\%$$

for each FHC/RHC, ν_μ/ν_e combination.

4.5.3 Far detector mass and air-bubble normalization

These normalization uncertainties are specific to the far detector (FD). The first is the FD mass uncertainty. The total active mass is inferred from the detector geometry, material densities, PVC and steel components mass. However, there are uncertainties associated with each of those. These lead to a $\sim 0.19\%$ uncertainty on the FD mass [120]. This is treated as an FD normalization uncertainty that is uncorrelated with the ND. The second FD-specific normalization arises from unmodelled air bubbles in the detector scintillator. Due to small deformities during fabrication of the long FD PVC extrusions, some horizontal cells in the FD contain small regions of unfilled scintillator occupied by air bubbles, which reduce the local active mass and light yield. The standard detector simulation does not include such effects. By comparing data and simulation in control samples, NOvA assigns a $\sim 0.4\%$ normalization uncertainty to account for the effect of these unmodelled air bubbles on the overall FD response [120]. This term is also treated as fully uncorrelated with the ND.

Chapter 5. The PISCES Method

As discussed in Section 3.1, to extract oscillation parameters, NOvA currently uses the "Extrapolation method", a beam fit mechanism in which the measured energy spectrum at the Near Detector (ND) is used to predict the energy spectrum that will be seen at the NOvA Far Detector (FD). This modified FD "extrapolated prediction" is then used during the fit. Though the ratio method allows systematic uncertainties that share the same effects in the detector to be reduced, this method still performs an FD-only likelihood fit, which means that much of the high-statistics ND information is not used optimally in the parameter estimation.

The PISCES method, on the other hand, allows for a joint ND+FD fit, directly incorporating the ND measurement in the likelihood. This significantly strengthens constraints on flux, cross-section, and detector-response systematics, especially those that affect the ND and FD differently.

Another key difference between the two methods is how they treat systematic uncertainties. The Extrapolation method introduces systematic uncertainties as explicit nuisance parameters in a ROOT Minuit-based fit. In practice, the dimensionality of the fit space should remain moderate to ensure stable minimization and reasonable computing time. As the number of parameters becomes large, the fit becomes increasingly difficult due to correlations and non-linearities. For this reason, systematic uncertainties are combined in the Extrapolation method, and only those with the most important impact are included explicitly in the fit. In contrast, PISCES encodes all systematic uncertainties in a full covariance matrix. This allows the framework to handle the full 160 systematics in the analysis simultaneously. As a result, PISCES naturally incorporates more information from the detector

and beam modeling. Furthermore, the covariance matrix spans all the ND and FD samples and analysis bins, and its off-diagonal terms encode how different bins and samples change coherently in response to the same systematics. As a result, correlations between ND and FD, as well as between different channels, are handled automatically and consistently with PISCES.

This chapter describes the PISCES framework in detail, including the PISCES χ^2 function, the construction of the covariance matrix, and the samples used in PISCES. It also compares Asimov data tests between PISCES and the Extrapolation method.

5.1 PISCES

PISCES is a hybrid test statistic combining a Poisson likelihood for statistical uncertainties and a multivariate Gaussian treatment of systematic uncertainties. The Poisson term provides a correct handling of statistical fluctuations, especially for the low-statistics samples such as the ν_e appearance channel in the FD. The Gaussian term incorporates the systematic uncertainties through a covariance matrix, allowing for complex fits that account for correlations between multiple detectors, such as a joint Near and Far detector fit.

In the PISCES framework, systematic uncertainties are not introduced as explicit nuisance parameters. Instead, PISCES uses a nested minimization, where for a fixed point in oscillation-parameter space, PISCES determines a set of optimal systematic shifts that modify the nominal prediction to a shifted prediction that best matches the observed data. These optimal systematic scales are obtained by minimizing the hybrid χ^2_{PISCES} , which combines the Poisson statistical term with the Gaussian systematic penalty encoded in the covariance matrix. An outer minimization over the oscillation parameters is then performed using ROOT's MINUIT fitter, with PISCES providing both the objective function χ^2 and shifted predictions with optimum pulls at each point in parameter space.

5.1.1 PISCES χ^2

For a set of N binned measurements, the Gaussian form of the likelihood, incorporating a covariance matrix V , is defined as

$$\mathcal{L} = \frac{1}{\sqrt{(2\pi)^N |V|}} \exp \left(-\frac{1}{2} \sum_{i,j} (x_i - \mu_i) V_{ij}^{-1} (x_j - \mu_j) \right), \quad (5.1)$$

where x_i denotes the observed event counts in the analysis bin i , and μ_i the corresponding predicted counts. The total covariance matrix V_{ij} encodes how fluctuations in bins i and j are correlated.

This multivariate Gaussian likelihood can be rewritten as

$$-2 \ln \mathcal{L} = \sum_{i,j} (x_i - \mu_i) V_{ij}^{-1} (x_j - \mu_j) + \log |V| + N \log(2\pi). \quad (5.2)$$

The constant term $N \log(2\pi)$ is typically omitted in practical applications. The determinant term $\log |V|$ can also be removed if the covariance matrix V is constant across parameter space. However, this is generally not the case. In PISCES, specific transformations are applied to construct a covariance matrix that can be treated as effectively constant across oscillation parameter space, allowing the $\log |V|$ term to be safely dropped from the test statistic:

- **Componentwise Matrix:** If the covariance matrix were built directly on the total reconstructed neutrino energy spectrum, where all oscillation channels are already summed up in each bin, then the matrix would change according to the considered oscillation parameters. This is due to different systematic uncertainties, such as the flux, cross-section, or detector systematics, not acting on the total spectrum in a uniform way. These systematics depend on underlying physics properties such as the neutrino production mechanism in the beam, the interaction mode in the detector, and

the final-state topology of the event (*e.g.*, a long muon track in $\nu_\mu \rightarrow \nu_\mu$ CC events, an electromagnetic shower in $\nu_\mu \rightarrow \nu_e$ CC events, or a mostly-hadronic signature for NC events).

When all of these different components are mixed in a single total bin, their relative contributions shift with the oscillation parameters. As a result, the effective response of that total bin to a systematic uncertainty also changes. This makes the covariance matrix explicitly dependent on the oscillation parameters, and therefore not constant across parameter space.

To avoid this, PISCES subdivides the total spectrum into the different underlying physics components that the systematic uncertainties actually act on. Instead of constructing the covariance matrix on the summed prediction, PISCES builds the matrix on the separate oscillation channels, each of which have a well-defined flux origin, interaction type, and detector topology. Once the samples are separated by channel, each systematic uncertainty shifts that channel in a consistent manner, essentially unaffected by the choice of oscillation parameters. In practice, each ND sample in PISCES is subdivided into 5 different ND oscillation channels, and each FD sample is subdivided into 13 oscillation channels defined in Section 5.3.4. In this case, the nominal and systematically shifted predictions (μ and m) are compared within each channel, and the covariance is built in this component space. In this analysis, the indices used are:

- Greek letters (α, β) for beam components (oscillation channels),
 - Roman letters (i, j) for energy bins.
- **Fractional Matrix:** Even after dividing the prediction into separate oscillation channels, some systematics can still introduce oscillation-dependent effects. This happens because many systematics are implemented as relative (percentage) changes to the

event rates. In this case, their impact scales with the predicted counts in absolute event-count space. If the covariance matrix were expressed in absolute event count units, its entries would scale with the predicted event rates and would therefore change whenever the overall normalization or the oscillation probabilities change. This would make the covariance matrix depend explicitly on the oscillation parameters.

To avoid this, PISCES works with fractional deviations, so that the covariance depends on the relative response of each channel to a given systematic. This keeps the matrix much more stable across the different oscillation parameters and across the very different statistical regimes of the ND and FD samples. Specifically, after splitting the samples into channels, each systematic covariance element is normalized by the corresponding nominal predictions, transforming Σ into a fractional covariance:

$$F_{\alpha i \beta j} = \frac{\Sigma_{\alpha i \beta j}}{\mu_{\alpha i} \mu_{\beta j}}.$$

This matrix F is the version used in the final PISCES χ^2 .

Decomposition into Statistical and Systematic Components

The Gaussian likelihood expression can be decomposed into contributions from statistical and systematic uncertainties:

$$\chi_{\text{Gaussian}}^2 = \sum_{i,j} [(x_i - m_i) U_{ij}^{-1} (x_j - m_j) + (\mu_i - m_i) \Sigma_{ij}^{-1} (\mu_j - m_j)], \quad (5.3)$$

where:

- U is the diagonal matrix of uncorrelated statistical uncertainties,
- Σ is the covariance matrix representing systematic uncertainties,
- x is the observed data,

- m is the best-fit prediction,
- μ is the nominal prediction,
- i, j are the energy bins.

In PISCES, the statistical Gaussian term is replaced by a Poisson likelihood, leading to the hybrid Gaussian/Poisson test statistic:

$$\chi_{\text{PISCES}}^2 = 2 \sum_i \left[m_i - x_i + x_i \log \left(\frac{x_i}{m_i} \right) \right] + \sum_{i,j} (\mu_i - m_i) \Sigma_{ij}^{-1} (\mu_j - m_j). \quad (5.4)$$

By taking Equation 5.1.1, decomposing each sample into the different oscillation channels, and applying the fractional matrix transformation,

$$F_{\alpha i \beta j} = \frac{\Sigma_{\alpha i \beta j}}{\mu_{\alpha i} \mu_{\beta j}},$$

the χ^2 expression becomes

$$\chi^2 = 2 \sum_i \left[m_i - x_i + x_i \log \left(\frac{x_i}{m_i} \right) \right] + \sum_{\alpha \beta} \sum_{ij} \left(\frac{m_{\alpha i}}{\mu_{\alpha i}} - 1 \right) F_{\alpha i \beta j}^{-1} \left(\frac{m_{\beta j}}{\mu_{\beta j}} - 1 \right), \quad (5.5)$$

where $F_{\alpha i \beta j}$ is the covariance matrix with each beam component binned independently, and with each matrix element divided by the nominal contents of the corresponding bins to form a fractional matrix.

We can now simplify these expressions by performing a parameter transformation to replace the systematically shifted prediction m as a function of the multiplicative systematic shift s , and the nominal prediction μ , where

$$m_{\alpha i} = s_{\alpha i} \mu_{\alpha i}. \quad (5.6)$$

The final χ^2 in the PISCES technique is then expressed as

$$\chi^2 = 2 \sum_{i=1}^N \left[\left(\sum_{\alpha=1}^M \mu_{\alpha i} s_{\alpha i} \right) - x_i + x_i \log \left(\frac{x_i}{\sum_{\alpha=1}^M \mu_{\alpha i} s_{\alpha i}} \right) \right] + \sum_{ij=1}^N \sum_{\alpha\beta=1}^M (s_{\alpha i} - 1) F_{\alpha i \beta j}^{-1} (s_{\beta j} - 1), \quad (5.7)$$

with $\mu_{\alpha i}$ and $s_{\alpha i}$ being the nominal expectation and fractional systematic shift for spectrum component α and spectrum bin i , respectively, and F the full covariance matrix. [121]

5.1.2 Fitting for oscillation parameters with PISCES

As mentioned previously, PISCES uses a nested minimization. The fit of the oscillation parameters is performed by ROOT's MINUIT. The PISCES test statistic depends explicitly on the oscillation parameters via the nominal predictions $\mu_{\alpha i}(\boldsymbol{\theta})$, where $\boldsymbol{\theta}$ denotes the set of oscillation parameters (for instance, $\sin^2 \theta_{23}$, Δm_{32}^2 , δ_{CP} , etc.). For any fixed choice of $\boldsymbol{\theta}$, PISCES evaluates a test statistic that depends both on the oscillation parameters and on the systematic scales $s_{\alpha i}$,

$$\chi^2(\boldsymbol{\theta}, s).$$

At this stage, $\boldsymbol{\theta}$ is fixed by MINUIT, while PISCES determines the values of $s_{\alpha i}$ that minimize this χ^2 .

For every point in oscillation-parameter space that MINUIT tries, PISCES performs an *inner* minimization where it solves for the set of systematic scales $s_{\alpha i}$ that shift the nominal predictions so they best match the data. The resulting optimal pulls and the corresponding shifted prediction $m_{\alpha i}$ are then used to evaluate the PISCES χ^2 at that oscillation point. MINUIT uses this value as its objective function to decide where to move next in the oscillation parameter space. In this way, PISCES provides MINUIT with a profiled χ^2 in which the systematic uncertainties have already been optimally accommodated at every step of the oscillation fit.

As a reminder, the shifted prediction is written as

$$m_{\alpha i} = s_{\alpha i} \mu_{\alpha i}, \quad (5.8)$$

where $\mu_{\alpha i}$ is the nominal prediction for oscillation channel α and energy bin i , and $s_{\alpha i}$ is the fractional systematic shift.

Using this definition, the PISCES χ^2 expression is:

$$\chi_{PISCES}^2 = 2 \sum_i \left[m_i - x_i + x_i \log \left(\frac{x_i}{m_i} \right) \right] + \sum_{\alpha\beta} \sum_{ij} \left(\frac{m_{\alpha i}}{\mu_{\alpha i}} - 1 \right) F_{\alpha i \beta j}^{-1} \left(\frac{m_{\beta j}}{\mu_{\beta j}} - 1 \right) \quad (5.9)$$

with F being the fractional covariance matrix defined previously.

Inner minimization for optimal systematic scales

At a fixed point in parameter space, PISCES first solves for the set of systematic shifts $s_{\alpha i}$ that rescale the prediction to best match the data. These optimal pulls that minimize the χ^2 are found iteratively using the Levenberg-Marquardt (LM) algorithm [122, 123, 124], which is an efficient minimization technique that blends gradient descent and Newton's method. The LM algorithm interpolates between a gradient-descent step and a Newton step depending on a damping parameter λ . When λ is large, the update behaves like a gradient descent, and as λ gets reduced to 0, the algorithm transitions to a Newton-like step.

The algorithm solves for the roots of the derivative of χ^2 with respect to each systematic pull by solving a system of linear equations defined by the gradient and Hessian matrix of χ^2 with respect to the systematic shifts.

The gradient vector of the χ^2 with respect to the systematic shifts in each beam component α and energy bin i is :

$$\frac{\partial \chi^2}{\partial s_{\gamma k}} = 2 \left[\mu_{\gamma k} - \frac{\mu_{\gamma k} x_k}{\sum_{\alpha} \mu_{\alpha k} s_{\alpha k}} + \sum_{\alpha i} (s_{\alpha i} - 1) F_{\alpha i, \gamma k}^{-1} \right], \quad (5.10)$$

The Hessian matrix is :

$$\frac{\partial^2 \chi^2}{\partial s_{\gamma k} \partial s_{\delta l}} = 2(1 + \delta_{\gamma\delta} \delta_{kl} \lambda) \left[\frac{\mu_{\gamma k} x_k}{(\sum_{\alpha} \mu_{\alpha k} s_{\alpha k})^2} \delta_{kl} + F_{\gamma k, \delta l}^{-1} \right]. \quad (5.11)$$

where the λ factor in the Hessian Matrix facilitates stable convergence of the LM algorithm by determining the appropriate minimization method at that step. λ is initially set to a certain value and is reduced iteratively (by a factor of 1.5 per step) until it becomes negligible, and the algorithm transitions to a pure Newton minimization.

It should be noted that during this inner minimization, PISCES does not solve for the individual systematic scales associated with each uncertainty. Instead, the Levenberg-Marquardt minimization determines the set of effective fractional scales $s_{\alpha i}$ that multiply the nominal prediction $\mu_{\alpha i}$, to produce the shifted prediction $m_{\alpha i} = s_{\alpha i} \mu_{\alpha i}$. These scale factors $s_{\alpha i}$ represent the combined effect of all the systematic uncertainties as constrained by the covariance matrix. The individual systematic pulls can be inferred through a separate process described in Chapter 6.

Outer minimization over the oscillation parameters

Once the optimal values for the systematic scales $s_{\alpha i}$ have been obtained at a fixed point in oscillation parameter space, PISCES evaluates the corresponding test statistic $\chi_{\text{prof}}^2(\boldsymbol{\theta})$ using the shifted predictions $m_{\alpha i} = s_{\alpha i} \mu_{\alpha i}$. This profiled $\chi_{\text{prof}}^2(\boldsymbol{\theta})$ value is then returned to MINUIT, which minimizes it with respect to the oscillation parameters.

MINUIT updates the oscillation parameters $\boldsymbol{\theta}$. Then the nominal predictions $\mu_{\alpha i}$ are re-computed, and PISCES repeats the inner minimization over the systematic scales $s_{\alpha i}$ for the new parameter space. This nested procedure, with the inner minimization over $s_{\alpha i}$, and outer minimization over $\boldsymbol{\theta}$, is repeated until the change in χ^2 between successive MINUIT iterations is smaller than 10^{-10} , at which point the fit is considered converged.




Figure 5.1: Evaluation of χ^2 at each point in the oscillation parameter space - starting from the nominal predictions and observed data, the optimal set of systematic scales that would shift the nominal prediction to match the data is determined. Once the optimal shifted prediction is found, a fit for the oscillation parameters is performed. The determination of systematic scales is carried out simultaneously across all samples in the analysis, including both ND and FD data.

5.1.3 Notes on fitting oscillation parameters in PISCES

In fitting oscillation parameters, a key feature of **PISCES** is that it does not require relying on manual seeding of variables. In traditional fitting frameworks, parameters often need to be seeded with initial values. For complex fits, multiple seeds may be required to avoid local minima, which can be inconvenient and error-prone, especially since appropriate seeds may vary depending on the input.

For example, in the official Three-Flavor analysis, the parameter δ_{CP} is typically seeded with values like 0 , $\frac{\pi}{2}$, π , and $\frac{3\pi}{2}$.

In contrast, **PISCES** employs a quadrant-based approach. When fitting δ_{CP} , the space is divided into four constrained variables, each exploring a different region:

- **q1:** constrained to $[0, \frac{\pi}{2}]$
- **q2:** constrained to $[\frac{\pi}{2}, \pi]$

- **q3**: constrained to $[\pi, \frac{3\pi}{2}]$
- **q4**: constrained to $[\frac{3\pi}{2}, 2\pi]$

Each quadrant is explored independently, ensuring the fitter searches the full parameter space and does not become trapped in local minima, improving the likelihood of finding the global minimum.

A similar strategy is also applied to Δm_{32}^2 , where both the **normal** and **inverted mass orderings** are considered, and to $\sin^2 \theta_{23}$, where both the **upper** and **lower octants** are explored.

As a result, a global fit involving δ_{CP} , Δm_{32}^2 , and $\sin^2 \theta_{23}$ explores all 16 combinations of δ_{CP} quadrants, mass orderings, and octant scenarios. This exhaustive exploration leads to a more robust and reliable fitting process, eliminating the need for guesswork in choosing seeds that might or might not lead to the true global minimum.

5.2 Constructing covariance matrices with PISCES

To construct a covariance matrix in PISCES, a Monte Carlo-based sampling method is used. The procedure begins by evaluating the *nominal prediction*, $Pred_i^{\text{nom}}$, in each analysis bin i , based on central values of the oscillation and systematic parameters.

A total of 10^5 pseudo-experiments, or *universes*, are generated to explore the effect of these systematic uncertainties. In each universe, all systematic parameters are varied simultaneously by drawing random shifts. In the NOvA framework, these are drawn according to their defined prior distributions, which are typically Gaussian. For one-sided systematics, only positive shifts are applied. This approach captures the collective influence of all parameter uncertainties on the predicted spectra.

For every shifted universe u , the modified prediction $Pred_i^{(u)}$ is obtained using interpolation from precomputed prediction objects in the NOvA framework. The *fractional deviation*

from the nominal prediction is then calculated for each bin:

$$f_i^{(u)} = \frac{Pred_i^{(u)} - Pred_i^{\text{nom}}}{Pred_i^{\text{nom}} + \epsilon} \quad (5.12)$$

where ϵ is a small regularization constant used to prevent numerical instability in bins with low nominal counts.

The covariance matrix V_{ij} is assembled by averaging the outer products of the fractional deviations across all universes:

$$V_{ij} = \frac{1}{N} \sum_{u=1}^N f_i^{(u)} f_j^{(u)} \quad (5.13)$$

This matrix captures both the variance in each bin (on the diagonal) and the correlations between bins (off-diagonal elements).

5.3 PISCES 3-Flavor analysis inputs

5.3.1 ν_μ Samples

In the NOvA detectors, ν_μ CC interactions are characterized by a long muon track accompanied by some hadronic deposition. The energy of the incoming neutrino is reconstructed from the muon deposited energy E_μ and the visible energy from the hadronic system E_{had} . E_μ is reconstructed using the procedure detailed in Section 2.5. The reconstructed E_ν is then subtracted from the total calorimetric energy deposition to get the remaining hadronic system energy E_{had} .

Muon energies are typically reconstructed with excellent resolution (around 3%). However, the hadronic energy resolution is worse, and it degrades with the amount of hadronic activity in the detector. For this reason, NOvA splits the ν_μ CC samples in four categories, called "quartiles", based on the fraction of hadronic energy relative to the total reconstructed energy E_{had}/E_ν . The 4 quartiles are represented in Fig. 5.2. The ν_μ CC events categorized

in the first quartiles, for which the muon carries most of the energy, have the best overall energy resolution. Quartiles with larger hadronic fractions, on the other hand, have poorer resolution. In this analysis, the 4 different quartile samples are referred to as the ν_μ Q1-Q4 samples.




Figure 5.2: Simulated ν_μ CC distributions of E_{had}/E_ν versus reconstructed E_ν for FHC (left) and RHC (right) running. The contour lines indicate the quartile boundaries used to categorize events by hadronic energy fraction, which correlates strongly with energy resolution.

Section 3 shows that the ν_μ survival probability $P_{\nu_\mu \rightarrow \nu_\mu}$ depends mainly on the oscillation parameters θ_{23} and Δm_{32}^2 at leading order, with θ_{23} setting the amplitude and Δm_{32}^2 setting the oscillation frequency. In the FD, the oscillation maximum occurs between the 1 to 2 GeV region in the ν_μ energy spectra. To better capture the changes within that region, ν_μ CC samples are binned more finely in that oscillation dip. In total, the ν_μ samples are subdivided into 19 reconstructed-energy bins, with 10 of them located between 1 to 2 GeV.

5.3.2 PISCES ν_e samples

ν_e FD samples

In the FD, the ν_e events are split into two categories named "core" and "peripheral". The core region contains most of the well-contained ν_e events, whereas the peripheral region regroups ν_e events that failed containment cuts, but are still relevant.

Concerning the core samples, they are further subdivided into two other categories according to their Particle Identification (PID) score, which is determined by the process in Section 2.5. The first category is "High PID" events, in which the events present strong ν_e -like characteristics. The second category is "Low PID", in which the events still pass the ν_e selection but with lower classifier confidence.

ν_e ND samples

One particularity of the PISCES 3-Flavor inputs is the binning scheme used for the ν_e samples in the ND. Specifically, the ν_e ND samples are binned as a function of reconstructed neutrino energy and the number of Michel electrons observed in the event. The motivation behind this is to enhance the information contained in the samples.

In the standard NOvA 3-Flavor extrapolation method, the ν_e ND samples undergo a dedicated decomposition procedure. This is a multi-step process that uses ND control samples, such as the contained and uncontained ν_μ CC events, to constrain the parent pion and kaon fluxes that generate the beam ν_e . These constraints are propagated to the intrinsic beam ν_e components. After that, the distribution of Michel electrons in the ND ν_e candidates is used as a handle to determine the balance between true ν_e CC events, ν_μ CC events misidentified as ν_e , and the NC backgrounds.

A Michel electron is a low-energy electron from the decay of a stopped muon. As CC events with muons or charged pions often produce Michel electrons, whereas NC events and events with dominating electromagnetic-shower like ν_e do not, the Michel-electron multiplic-

ity gives information on the underlying interaction type. For this reason, in the extrapolation framework, the Michel distribution provides direct information about the CC/NC composition of the ND ν_e background, which can be used to rescale the ν_e samples accordingly.

PISCES does not apply this full decomposition process. Instead of computing explicit scaling factors for each background component, PISCES incorporates the same information provided by the Michel electrons by binning the ND ν_e events directly in the Michel-electron multiplicity, in addition to reconstructed energy. This approach allows the fit to have the same useful information from the decomposition process: bins with zero Michel electrons tend to be EM-shower dominated or NC events, while bins with one or more Michel electrons are more likely to contain ν_μ CC backgrounds. It should be noted that, similar to the FD, those 2D ν_e samples are also further subdivided into Low PID and High PID categories.

5.3.3 PISCES total Samples

As the ν_μ samples are subdivided into the four quartiles, and accounting for the FHC and RHC samples, in the ND and FD, there are a total of 16 ν_μ samples in PISCES. For the ν_e samples, they are subdivided into the Low PID and High PID in the ND, and subdivided into the Low PID, High PID, and peripheral samples in the FD. This means that there are 4 ν_e samples in the ND, and 6 ν_e samples in the FD when taking into account the FHC and RHC samples. Furthermore, the ν_e low-energy samples were also included. These add 2 extra samples in the analysis from the low energy ν_e (one sample in the ND, and another one in the FD), and two extra samples from the uncontained ν_μ in the ND (one FHC sample and one RHC sample). These make a total of 30 samples. It should be noted that each of those samples is subdivided into the corresponding oscillation channels. This means that for ND samples, they are subdivided into the 5 oscillation channels in the ND, and for FD samples, they are subdivided into the 13 oscillation channels in the FD.




Figure 5.3: PISCES ν_e samples distribution - The plots represent ν_e in the ND, separated into Low PID and High PID, for both FHC and RHC. The ν_e samples are binned as a function of reconstructed energy and number of Michel electrons to enhance information content.

5.3.4 Oscillation Channels in PISCES

In the PISCES framework, neutrino samples from both the ND and FD are subdivided into multiple oscillation channels. Each channel corresponds to a specific neutrino flavor transition, such as $\nu_\mu \rightarrow \nu_\mu$ or $\nu_\mu \rightarrow \nu_e$, etc. This subdivision allows PISCES to separately model each oscillation channel and propagate correlations among them in the fit.

ND Channels

At the ND baseline ($L \simeq 1$ km), oscillation effects are negligible; however, PISCES still computes the same oscillation-channel structure for consistency between detectors. This feature has been inherited from the PISCES Sterile Neutrino analysis, in which short baseline oscillations are non-negligible. Five oscillation channels are defined for each of the ND samples, corresponding to possible flavor transitions explaining the origin of the considered ν :

- $\nu_e \rightarrow \nu_e$ (**electron-neutrino survival**) This channel represents the intrinsic ν_e component in the beam that remains as ν_e in FHC mode.
- $\bar{\nu}_e \rightarrow \bar{\nu}_e$ (**electron-antineutrino survival**) This channel represents the intrinsic $\bar{\nu}_e$ component in the beam that remains as $\bar{\nu}_e$ in RHC mode.
- $\nu_\mu \rightarrow \nu_\mu$ (**muon-neutrino survival**) This is the dominant contribution to the ND sample in FHC mode.
- $\bar{\nu}_\mu \rightarrow \bar{\nu}_\mu$ (**muon-neutrino survival**) This is the dominant contribution to the ND sample in RHC mode.
- **Neutral-current** Flavor-independent NC interactions, which contribute as background to the ND samples.

FD Channels

At the FD ($L \simeq 810$ km), oscillations are fully developed. Thirteen distinct oscillation channels are defined to represent all flavor and charge combinations for both neutrinos and antineutrinos. These channels represent the possible flavor transitions that occur as neutrinos propagate from the source to the FD.

- $\nu_e \rightarrow \nu_e$: Electron-neutrino survival in neutrino mode (FHC).
- $\bar{\nu}_e \rightarrow \bar{\nu}_e$: Electron-antineutrino survival in antineutrino mode (RHC).
- $\nu_\mu \rightarrow \nu_\mu$: Muon-neutrino survival in FHC.
- $\bar{\nu}_\mu \rightarrow \bar{\nu}_\mu$: Muon-antineutrino survival in RHC.
- $\nu_e \rightarrow \nu_\mu$: Electron-to-muon appearance in FHC mode.
- $\bar{\nu}_e \rightarrow \bar{\nu}_\mu$: Electron-antineutrino to muon-antineutrino transition in RHC.
- $\nu_e \rightarrow \nu_\tau$: Electron-neutrino to tau-neutrino appearance in FHC. A rare channel, included for completeness and consistency with three-flavor mixing.
- $\bar{\nu}_e \rightarrow \bar{\nu}_\tau$: Electron-antineutrino to tau-antineutrino appearance in RHC.
- $\nu_\mu \rightarrow \nu_e$: Muon-to-electron appearance in FHC mode.
- $\bar{\nu}_\mu \rightarrow \bar{\nu}_e$: Muon-antineutrino to electron-antineutrino appearance in RHC.
- $\nu_\mu \rightarrow \nu_\tau$: Muon-to-tau appearance in FHC. Although τ leptons are not reconstructed, this channel impacts event rates.
- $\bar{\nu}_\mu \rightarrow \bar{\nu}_\tau$: Muon-antineutrino to tau-antineutrino appearance in RHC.
- **Neutral-current**: Flavor-independent NC sample.

Figure 5.4 represents the different nominal distributions for each sample in PISCES and the corresponding 1σ systematic ranges, taking into account all the systematic uncertainties in the analysis, where the different oscillation channels contributions have been summed up.




Figure 5.4(a): ν_μ Q1 samples - PISCES nominal samples distribution and the corresponding 1σ error bands considering all systematic uncertainties in PISCES.




Figure 5.4(b): ν_μ Q2 samples - PISCES nominal samples distribution and the corresponding 1σ error bands considering all systematic uncertainties in PISCES.




Figure 5.4(c): ν_μ Q3 samples - PISCES nominal samples distribution and the corresponding 1σ error bands considering all systematic uncertainties in PISCES.




Figure 5.4(d): ν_μ Q4 samples - PISCES nominal samples distribution and the corresponding 1σ error bands considering all systematic uncertainties in PISCES.




Figure 5.4(e): ν_e FD samples - PISCES nominal samples distribution and the corresponding 1σ error bands considering all systematic uncertainties in PISCES.




Figure 5.4(f): ν_e low-energy samples - PISCES nominal samples distribution and the corresponding 1σ error bands considering all systematic uncertainties in PISCES.




Figure 5.4(g): ν_e ND samples - PISCES nominal samples distribution and the corresponding 1σ error bands considering all systematic uncertainties in PISCES.

5.3.5 PISCES Joint ND+FD Fit

In a joint ND+FD fit, both ND and FD samples are used simultaneously during the fit, and the corresponding full covariance matrix is employed to constrain systematics coherently across both detectors and all analysis samples. Figure 5.5 shows an example of a covariance matrix that incorporates both ND and FD information. The matrix includes contributions from both the ND and the FD, with each quadrant representing different combinations of ND and FD sample correlations. The diagonal elements represent the variance, quantifying the uncertainty in each individual analysis bin. The off-diagonal elements represent the covariance, which captures correlations between different bins, including those between the two detectors.




Figure 5.5: Example of a covariance matrix incorporating combined ND and FD samples, used in a joint PISCES ND+FD fit where the ND and FD samples are fit simultaneously.

5.4 Comparison of PISCES with the Official 3-flavor analysis framework

The following plots showcase the comparison of sensitivity contours obtained from the PISCES framework with those produced by the official 3-Flavor analysis framework. The comparisons were done using an Asimov data set, which is a fake data sample equal to the nominal prediction in each bin, without any statistical or systematic fluctuations.

The first set of comparisons in Section 5.4.1 shows results using only FD information. For PISCES, this means that only the FD samples were included in the fit. In the official 3-Flavor analysis, this corresponds to using non-extrapolated FD predictions. Section 5.4.2 presents comparison plots for cases where ND information is included in both frameworks. In the official 3-Flavor analysis, incorporating ND information means that extrapolated predictions are used during the fit. In the PISCES framework, this translates to a joint ND+FD fit, where both ND and FD samples are used simultaneously.




Figure 5.6: Stats-only comparison of PISCES and official 3-Flavor fits using FD-only information considering an Asimov fake data generated at the NOvA 2024 best-fit point. The results obtained from both frameworks are virtually identical.

5.4.1 Using FD-only information

FD-only / stats-only comparison

Figure 5.6 shows the stats-only comparison, with a fit considering only FD information. The fake data used in both cases is the Asimov data generated at the NOvA 2024 Best-Fit points. The comparison is done in the $\sin^2 \theta_{23}$ vs Δm_{32}^2 , and in the δ_{CP} vs $\sin^2 \theta_{23}$ space, considering both normal ordering (top rows) and inverted ordering (bottom rows). In this case, the contours obtained from PISCES and the official 3-Flavor analysis are virtually identical.




Figure 5.7: Comparison of PISCES and official 3-Flavor fits using FD-only information considering an Asimov fake data generated at the NOvA 2024 best-fit point. In this case, the fit includes systematic uncertainties that are identical between the two frameworks. The contours obtained from both methods are also identical.

FD-only / with identical systematic uncertainties comparison

Still focusing solely on FD information, but now including systematic uncertainties, Fig. 5.7 compares the contours obtained from PISCES and the official 3-Flavor analysis, using a set of systematic uncertainties that are identical in both frameworks. The systematics list includes detector responses, flux and cross-section uncertainties, and excludes systematics that are specific to the Extrapolation method. The results show that, for the same set of systematics, the contours produced by both methods are virtually identical.




Figure 5.8: Comparison of sensitivity contours using the full set of systematic uncertainties in both the PISCES FD-only fit and the 3-Flavor non-extrapolated prediction fit. The results show strong agreement between the two approaches, with slight differences attributed to the distinctions in the systematic lists used in each analysis.

FD-only / with the full list of systematics comparison

Figure 5.8 shows the comparison of the case where the full systematic list is considered for the PISCES fit and the official 3-Flavor fit. The contours obtained from the PISCES FD-ONLY fit and the 3-Flavor non-extrapolated prediction fit are very similar when considering FD-ONLY information and using the full systematic lists in both analyses. The slight differences in the contours arise from the differences in the systematic lists.

5.4.2 Including ND information

Including ND information - Official 3-Flavor framework

Figure 5.9 illustrates the impact of including ND information compared to using FD-only samples in the official 3-Flavor analysis, considering the full list of systematics. In this context, incorporating ND information means using extrapolated FD predictions during the fit. The extrapolation was performed using ND fake data, and the resulting contours with ND information (shown in Fig. 5.9) were based on these extrapolated predictions. The comparison illustrates the information gain achieved by incorporating ND information through extrapolation, as systematic uncertainties with similar effects on both detectors are constrained.

Including ND information - PISCES Joint ND+FD Fit

For the PISCES framework, Fig. 5.10 compares an FD-only fit with a simultaneous ND+FD fit. Both contours were produced using the full set of systematic uncertainties in PISCES. The comparison shows that including ND samples helps to more effectively constrain systematic uncertainties, resulting in tighter contours. The difference between the FD-only and joint ND+FD fits is more pronounced than in the extrapolation case, highlighting the advantage of incorporating ND samples directly into the fit.

5.4.3 Comparison of PISCES Joint ND+FD Fit and the Extrapolation method

PISCES Joint ND+FD Fit and Extrapolation method - stats-only comparison

For the following comparison plots between the PISCES and official 3-Flavor frameworks, ND information is incorporated, and only the stats-only case is considered. As mentioned previously, in the official 3-Flavor analysis, ND information enters via the extrapolation




Figure 5.9: Comparison of contours considering the full systematic uncertainties for the official 3-Flavor analysis with and without ND information. The ND information is incorporated via the extrapolation method, with the corresponding contours produced using extrapolated predictions from ND fake data. The addition of ND information leads to tighter contours due to the constraints on systematic uncertainties provided by the extrapolation.




Figure 5.10: Comparison of FD-only and simultaneous ND+FD fits in the PISCES framework. Including ND samples directly in the fit improves the constraint on systematic uncertainties, resulting in tighter contours. The shifts are more pronounced than those observed with the extrapolation method, demonstrating the benefit of a joint two-detector fit.




Figure 5.11: Comparison of a stats-only fit using PISCES and the official 3-Flavor framework with the ND information included. In this case, all ND and FD samples in PISCES are fit simultaneously, while for the official 3-Flavor framework, the extrapolated predictions derived using the ND fake data were used in the fit. The results obtained from both methods are consistent in this case.

method. In the PISCES framework, this corresponds to using both Near and FD samples simultaneously in the fit. Figure 5.11 shows the comparison between the two methods with ND information included, demonstrating that in the stats-only case, both methods yield very similar results.

PISCES Joint ND+FD Fit and Extrapolation method - FULL SYSTS lists comparison

Figure 5.12 shows a comparison between the PISCES framework and the official 3-Flavor analysis when both ND and FD information are considered, and the full list of systematic un-




Figure 5.12: Comparison between the PISCES simultaneous ND+FD fit and the official 3-Flavor fit using extrapolated predictions. Tighter contours are obtained in the PISCES joint fit, indicating greater information gain from the direct inclusion of ND samples.

certainties is included in both methods. As shown in the results, the PISCES joint ND+FD fit yields tighter contours than the extrapolation method. These comparisons demonstrate that incorporating ND information leads to improved constraints in both approaches; however, greater information gain is achieved through the direct use of ND samples in the simultaneous two-detector PISCES fit.

5.5 PISCES ND+FD fit - Asimov Fake Data Tests

To test the robustness of the technique, several fake data studies were conducted. The first set of tests involved evaluating the fitter using multiple fake datasets generated at different

Asimov oscillation points, to verify whether the fitter successfully recovers the injected values.

In this analysis, five distinct oscillation points were selected to test the fitter:

- Asimov A - This is the NOvA 2024 Best-Fit Point, with a preference for normal ordering for Δm_{32}^2 , upper octant for $\sin^2 \theta_{23}$, and a δ_{CP} value of 0.88π :

$$\Delta m_{32}^2 = 2.44 \times 10^{-3} eV^2 \quad (5.14)$$

$$\sin^2 \theta_{23} = 0.55 \quad (5.15)$$

$$\delta_{CP} = 0.88\pi \quad (5.16)$$

- Asimov B - This is a point close to T2K's 2020 Best-Fit, with preference for normal ordering and upper octant with the following values :

$$\Delta m_{32}^2 = 2.51 \times 10^{-3} eV^2 \quad (5.17)$$

$$\sin^2 \theta_{23} = 0.53 \quad (5.18)$$

$$\delta_{CP} = 1.5\pi \quad (5.19)$$

- Asimov C - This is a point used to probe cases with preference for inverted ordering and maximum CP violation :

$$\Delta m_{32}^2 = -2.45 \times 10^{-3} eV^2 \quad (5.20)$$

$$\sin^2 \theta_{23} = 0.50 \quad (5.21)$$

$$\delta_{CP} = 1.5\pi \quad (5.22)$$

- Asimov D - This is a point to probe lower octant preference with a CP-conserving

Table 5.1: Comparison between the oscillation parameters used to generate the fake data (NOvA 2024 Best-Fit point) and the best-fit values obtained from the PISCES ND+FD fit. In all cases, the fit accurately recovers the injected reference point values.

	Asimov Point	Best Fit Value
$\sin^2 \theta_{23}$	0.55	0.549974
Δm_{32}^2	$2.44 \times 10^{-3} \text{ eV}^2$	$2.44 \times 10^{-3} \text{ eV}^2$
δ_{CP}	0.88π	0.870176π

value for δ_{CP} :

$$\Delta m_{32}^2 = 2.40 \times 10^{-3} \text{ eV}^2 \quad (5.23)$$

$$\sin^2 \theta_{23} = 0.46 \quad (5.24)$$

$$\delta_{CP} = 1\pi \quad (5.25)$$

- Asimov E - This is a point located very far from NOvA's preferred oscillation point, with a preference for inverted ordering, $\sin^2 \theta_{23}$ in the lower octant, and $\delta_{CP} = \pi/2$:

$$\Delta m_{32}^2 = -2.44 \times 10^{-3} \text{ eV}^2 \quad (5.26)$$

$$\sin^2 \theta_{23} = 0.46 \quad (5.27)$$

$$\delta_{CP} = 0.5\pi \quad (5.28)$$

Figures 5.13–5.17 show the contours obtained in the $\sin^2 \theta_{23}$ vs. Δm_{32}^2 and δ_{CP} vs. $\sin^2 \theta_{23}$ parameter spaces, considering both normal and inverted ordering, for fake data generated at different reference oscillation points. Tables 5.1–5.5 compare the oscillation parameter values used to generate the fake data with the best-fit values obtained from the fit. In all cases, the PISCES ND+FD fit accurately recovers the values of the injected reference points.




Figure 5.13: 1σ , 2σ and 3σ contours comparison considering an Asimov fake data, generated at the 2024 Best-Fit point, for the $\sin^2 \theta_{23}$ vs Δm_{32}^2 and δ_{CP} vs $\sin^2 \theta_{23}$ parameter space. The first rows consider the normal ordering for Δm_{32}^2 , the second rows are for the inverted ordering.

Table 5.2: Comparison between the oscillation parameters used to generate the fake data (Asimov B) and the best-fit values obtained from the PISCES ND+FD fit. In all cases, the fit accurately recovers the injected reference point values.

	Asimov Point	Best Fit Value
$\sin^2 \theta_{23}$	0.53	0.529996
Δm_{32}^2	$2.51 \times 10^{-3} \text{ eV}^2$	$2.51 \times 10^{-3} \text{ eV}^2$
δ_{CP}	1.5π	1.5π




Figure 5.14: 1σ, 2σ and 3σ contours comparison considering an Asimov fake data, generated at the Asimov B reference point, for the $\sin^2 \theta_{23}$ vs Δm_{32}^2 and δ_{CP} vs $\sin^2 \theta_{23}$ parameter space. The first rows consider the normal ordering for Δm_{32}^2 , the second rows are for the inverted ordering.

Table 5.3: Comparison between the oscillation parameters used to generate the fake data (Asimov C) and the best-fit values obtained from the PISCES ND+FD fit. In all cases, the fit accurately recovers the injected reference point values.

	Asimov Point	Best Fit Value
$\sin^2 \theta_{23}$	0.50	0.500008
Δm_{32}^2	$-2.45 \times 10^{-3} \text{ eV}^2$	$-2.45 \times 10^{-3} \text{ eV}^2$
δ_{CP}	1.5π	1.50003π




Figure 5.15: 1σ, 2σ and 3σ contours comparison considering an Asimov fake data, generated at the Asimov C reference point, for the $\sin^2 \theta_{23}$ vs Δm_{32}^2 and δ_{CP} vs $\sin^2 \theta_{23}$ parameter space. The first rows consider the normal ordering for Δm_{32}^2 , the second rows are for the inverted ordering.

Table 5.4: Comparison between the oscillation parameters used to generate the fake data (Asimov D) and the best-fit values obtained from the PISCES ND+FD fit. In all cases, the fit accurately recovers the injected reference point values.

	Asimov Point	Best Fit Value
$\sin^2 \theta_{23}$	0.46	0.460016
Δm_{32}^2	$2.40 \times 10^{-3} \text{ eV}^2$	$2.3999 \times 10^{-3} \text{ eV}^2$
δ_{CP}	1.0π	0.9999π




Figure 5.16: 1σ , 2σ and 3σ contours comparison considering an Asimov fake data, generated at the Asimov D reference point, for the $\sin^2 \theta_{23}$ vs Δm_{32}^2 and δ_{CP} vs $\sin^2 \theta_{23}$ parameter space. The first rows consider the normal ordering for Δm_{32}^2 , the second rows are for the inverted ordering.

Table 5.5: Comparison between the oscillation parameters used to generate the fake data (Asimov E) and the best-fit values obtained from the PISCES ND+FD fit. In all cases, the fit accurately recovers the injected reference point values.

	Asimov Point	Best Fit Value
$\sin^2 \theta_{23}$	0.46	0.460016
Δm_{32}^2	$-2.44 \times 10^{-3} \text{ eV}^2$	$-2.44 \times 10^{-3} \text{ eV}^2$
δ_{CP}	0.5π	0.5π




Figure 5.17: 1σ, 2σ and 3σ contours comparison considering an Asimov fake data, generated at the Asimov E reference point, for the $\sin^2 \theta_{23}$ vs Δm_{32}^2 and δ_{CP} vs $\sin^2 \theta_{23}$ parameter space. The first rows consider the normal ordering for Δm_{32}^2 , the second rows are for the inverted ordering.

Chapter 6. Systematic–pull inference with PISCES:

6.1 Problem statement and objectives

As seen in Chapter 5, when fitting oscillation parameters, one of the key steps in the PISCES inner minimization involves solving for the optimal systematic scale values, which are the scaling factors that shift the nominal predictions to best match the considered data. This solution is obtained using the full covariance matrix, which incorporates all systematic uncertainties in the analysis. As such, PISCES provides the overall scaling vector, which combines the effects of all systematic uncertainties into a single effective scale factor per analysis bin. This means that PISCES does not directly provide the individual systematic pull values, *i.e.*, it does not give direct information on which underlying systematic uncertainty is responsible for a given shift. Such information, however, is valuable for interpreting the fit results and for comparing the PISCES outputs with other methods, such as the standard NOvA Extrapolation method.

To evaluate the contribution from each individual systematic in PISCES, additional calculations are required. This chapter details the process by which we infer individual systematic pulls for a given data fit with PISCES. The objective of this analysis is to calculate those per-systematic pulls, noted by $\mathbf{a}$, which correspond to the obtained PISCES systematic scaling factor, using the Newton-Raphson method [125]. This method has been tested and validated with systematically fluctuated fake data toy experiments.

In determining the pull value for each systematic uncertainty, we seek to infer a vector

of optimal systematic pulls $\mathbf{a} \in \mathbb{R}^{n_{\text{syst}}}$ corresponding to the analyzed data, given the overall systematic scale obtained from PISCES. In this analysis, the optimal systematic scaling factor is denoted as the vector β , with each element addressed by a linearized index i spanning all analysis bins (both oscillation channels and analysis bins).

The shifted prediction obtained from PISCES can thus be expressed as :

$$m_i = \beta_i \mu_i \quad (6.1)$$

where m is the shifted prediction, β the scaling factor obtained after the inner PISCES systematics fit, and μ the nominal prediction.

The shifted prediction, however, can be rewritten empirically as a function of the individual systematic pulls as:

$$m_i = \Pi_j (1 + \alpha_j V_{ij}) \mu_i \quad (6.2)$$

with i the analysis bin index, and j the systematic uncertainty index. m is the shifted prediction, α_j , the pull value corresponding to systematic uncertainty j , and V the matrix encoding the $+1\sigma$ change of systematic j . V is the fractional response matrix, whose column $\mathbf{v}_j$ encodes the fractional bin shifts for a $+1\sigma$ change of systematic j .

Equating 6.1 and 6.2 gives:

$$\Pi_j (1 + \alpha_j V_{ij}) = \beta_i \quad (6.3)$$

Where we can take the logarithm form:

$$\sum_j \log(1 + \alpha_j V_{ij}) = \log(\beta_i) \quad (6.4)$$

Solving for α_j in Equation 6.4 is equivalent to solving

$$F_i(\alpha) = \log(\beta_i) - \sum_j \log(1 + \alpha_j V_{ij}) = 0 \quad (6.5)$$

To get the individual systematic pulls α_j , the Newton-Raphson method can be applied to obtain the solution to the nonlinear system in Equation 6.5.

6.2 Newton-Raphson method

The Newton-Raphson method is an algorithm that can solve for the roots of a nonlinear system [125]:

$$\mathbf{F}(\boldsymbol{\alpha}) = \mathbf{0}, \quad \mathbf{F} : \mathbb{R}^n \rightarrow \mathbb{R}^m,$$

by iteratively linearizing the system $\mathbf{F}$ around a current guess, and solving the resulting linear system to take corrective steps.

In the Scalar case ($n = m = 1$)

Given a scalar function $f(x) : \mathbb{R} \rightarrow \mathbb{R}$, we want to find a root x^* such that $f(x^*) = 0$. In the Newton-Raphson method, starting from an initial guess $x^{(0)}$, a sequence of $x^{(k)}$ is generated such that they will eventually converge to the root x^* .

In the algorithm, at each iteration k , f is approximated by its tangent line near x^k using the first-order Taylor expansion:

$$f(x^{(k)} + \Delta) \approx f(x^{(k)}) + f'(x^{(k)})\Delta \quad (6.6)$$

The size Δ for each step is chosen such that the linearized approximation is 0, i.e.:

$$f(x^{(k)}) + f'(x^{(k)})\Delta \approx 0 \quad (6.7)$$




Figure 6.1: Example of the geometric view of Newton-Raphson updates for a nonlinear equation $f(x) = 0$, assuming $f(x) = x^3 - 6x + 4$. Starting from an initial guess $x^{(0)}$ on the curve (solid black), the dashed line shows the tangent at $x^{(0)}$; its x -intercept (colored $\times$) is the next iteration $x^{(1)}$. Repeating from $x^{(1)}$ (and again from $x^{(2)}$) illustrates how each step moves to the x -intercept of the current tangent, typically converging rapidly to the root x^* .

Solving for Δ gives:

$$\Delta = -\frac{f(x^{(k)})}{f'(x^{(k)})} \quad (6.8)$$

We add that increment value Δ to get the next iteration $x^{(k+1)}$:

$$x^{(k+1)} = x^{(k)} - \frac{f(x^{(k)})}{f'(x^{(k)})} \quad (6.9)$$

With this method, each step moves to the x -intercept of the tangent at $x^{(k)}$. The algorithm basically follows the slope down to where the tangent hits zero. As iterations continue, the steps become smaller, and the estimate converges to the root, as illustrated in Fig. 6.1.

Multivariate case (n parameters, m equations)

In the multivariate case where $\mathbf{F} : \mathbb{R}^n \rightarrow \mathbb{R}^m$, we want to find a vector $\boldsymbol{\alpha}$ such that $\mathbf{F}(\boldsymbol{\alpha}) = \mathbf{0}$.

Instead of taking the Taylor expansion as in the scalar case, we linearize $\mathbf{F}$ around the current iterate $\boldsymbol{\alpha}^{(k)}$ using its Jacobian :

$$J(\boldsymbol{\alpha}) = \partial \mathbf{F} / \partial \boldsymbol{\alpha} \in \mathbb{R}^{m \times n}$$

A first-order expansion gives

$$\mathbf{F}(\boldsymbol{\alpha}^{(k)} + \Delta \boldsymbol{\alpha}) \approx \mathbf{F}(\boldsymbol{\alpha}^{(k)}) + J(\boldsymbol{\alpha}^{(k)}) \Delta \boldsymbol{\alpha}. \quad (6.10)$$

Setting this linearized form equal to zero

$$\mathbf{F}(\boldsymbol{\alpha}^{(k)}) + J(\boldsymbol{\alpha}^{(k)}) \Delta \boldsymbol{\alpha} \approx 0 \quad (6.11)$$

leads to the following Newton system

$$J(\boldsymbol{\alpha}^{(k)}) \Delta \boldsymbol{\alpha} = -\mathbf{F}(\boldsymbol{\alpha}^{(k)}), \quad \boldsymbol{\alpha}^{(k+1)} = \boldsymbol{\alpha}^{(k)} + \Delta \boldsymbol{\alpha}.$$

which is the standard Newton update for nonlinear equations [125].

In practice, the update is scaled as:

$$\boldsymbol{\alpha}^{(k+1)} = \boldsymbol{\alpha}^{(k)} + s \Delta \boldsymbol{\alpha} \quad (6.12)$$

where $s \in (0, 1]$ is chosen through a simple backtracking line search.

This step-scaling has two purposes:

- **Ensuring convergence:** The line search makes sure that the residual norm $\|\mathbf{F}\|_2$ decreases. This is done so that Newton's method does not diverge far away from the solution.
- **Maintain physical predictions and valid $\log$ term:** We also require that all multiplicative factors in the prediction remain positive ($1 + \alpha_j v_{ij} > 0$), so that the $\log$ term in $\mathbf{F}$ is well-defined, and that the predicted rates stay physical.

This iteration is repeated until the update becomes sufficiently small

$$\|\Delta\boldsymbol{\alpha}\|_{\infty} < \varepsilon$$

If $m = n$ and J is nonsingular, the Newton step can be solved directly.

However, in our problem

$$n_{bins} \gg n_{syst},$$

as the number of analysis bins is in the order of $\mathcal{O}(4000)$ and the number of systematic uncertainties is $n_{syst} = 160$.

Therefore, the system is overconstrained. In this case, we use the least-squares (Gauss-Newton method) update [125] :

$$(J^{\top} J) \Delta\boldsymbol{\alpha} = -J^{\top} \mathbf{F}.$$

As a reminder, solving for the systematic pulls corresponds to solving for $\mathbf{F}(\boldsymbol{\alpha}) = 0$, where the residual function is

$$\mathbf{F}(\boldsymbol{\alpha}) = \log \boldsymbol{\beta} - \sum_{j=1}^{n_{syst}} \log(1 + \alpha_j \mathbf{V}_j) \quad (6.13)$$

where $\mathbf{V}_j$ is the column of fractional bin shifts for systematic j .

The Jacobian columns are

$$J_{:j}(\boldsymbol{\alpha}) = -\mathbf{V}_j \oslash (1 + \alpha_j \mathbf{V}_j).$$

where the division is element-wise.

For the Gauss–Newton algorithm [125], the step update minimizes the quadratic model

$$\frac{1}{2} \|\mathbf{F} + J \Delta \boldsymbol{\alpha}\|_2^2 \tag{6.14}$$

which leads to the normal equations

$$(J^\top J) \Delta \boldsymbol{\alpha} = -J^\top \mathbf{F}.$$

For a stable minimization, we use the Levenberg–Marquardt modification [122, 123, 124]:

$$(J^\top J + \lambda I) \Delta \boldsymbol{\alpha} = -J^\top \mathbf{F}, \quad \lambda \geq 0,$$

The LM approach continuously interpolates between steepest descent and a Gauss–Newton method, as seen in Section 5.1.2. For large λ , the update behaves as a steepest descent, and for small λ , it behaves as a Gauss–Newton. This stabilizes the iterations when the Gauss–Newton step is not well-behaved or when the initial guess is far from the solution.

After forming the symmetric positive (semi)-definite system

$$(J^\top J + \lambda I) \Delta \boldsymbol{\alpha} = -J^\top \mathbf{F},$$

we solve it using a direct factorization method. In this case, the $LDL^\top$ factorization was used (a unit lower-triangular matrix L times a diagonal matrix D times its transpose). $LDL^\top$ is well suited to symmetric positive-(semi)definite matrices and avoids unnecessary square

roots. It is also more numerically stable when the system is nearly singular [126, 127].

The resulting candidate step $\Delta\boldsymbol{\alpha}$ is then scaled using the line search described above, while making sure that

$$1 + \alpha_j V_{ij} > 0 \quad \text{for all } i, j.$$

The iterations stop when the update becomes sufficiently small,

$$\|\Delta\boldsymbol{\alpha}\|_\infty < \varepsilon,$$

or, equivalently, when the gradient of the least-squares objective is small,

$$\|J^\top \mathbf{F}\|_\infty \text{ is below tolerance.}$$

At that point, further iterations would not meaningfully change the inferred pulls.

Once convergence is reached, the algorithm returns a final, converged set of pull values $\hat{\boldsymbol{\alpha}}$. These are simply the last values of the parameters after the updates have become negligible. This final vector indicates how each individual systematic uncertainty contributes to the overall PISCES scaling $\boldsymbol{\beta}$, allowing us to identify which uncertainties are effectively driving the observed shifts in the prediction. In practice, this makes the PISCES output directly comparable to the standard per-systematic pulls obtained in the Extrapolation method.

6.3 Newton Method limitations and use of Principal Component Analysis (PCA)

The Newton-Raphson and Gauss-Newton methods are well behaved when the Jacobian J is well conditioned and not nearly rank deficient. In practice, this means that J should follow the conditions below:

- **Full column rank:** $\text{rank}(J) = p$. Equivalently [128, 129]:

- the p columns are linearly independent (so $m \geq p$);
- all p singular values are positive, i.e. $\sigma_{\min}(J) > 0$;
- $J^\top J$ is positive definite and invertible;
- the only vector mapped to zero is the zero vector:

$$\text{Null}(J) = \{0\}.$$

- **Smallest singular value is away from zero:** If $\sigma_{\min}$ is close to zero, then some direction in parameter space is almost unconstrained, which makes the Newton step unstable, therefore [130]

$$\sigma_{\min}(J) > 0.$$

- **Moderate condition number:** $\kappa_2(J) = \frac{\sigma_{\max}(J)}{\sigma_{\min}(J)}$ must not be extremely large, otherwise the Newton update becomes dominated by numerical noise [131, 128].

When these conditions are met, the orthogonal projector onto the column space,

$$P = JJ^+ = J(J^\top J)^{-1}J^\top,$$

is well behaved [129, 128], and the Gauss-Newton update

$$(J^\top J) \delta\alpha = -J^\top r$$

is numerically stable and well posed [132, 133].

In the systematic pull analysis, the Jacobian J gets its structure from the fractional-response matrix V . Therefore, any near-linear dependencies among the columns of V appear

directly in J .

This means that, if two systematics shift the spectrum in almost the same way, then two columns of V , and therefore J , are nearly linearly dependent. When this happens:

- $J^\top J$ becomes ill-conditioned or nearly singular,
- many different degenerate combinations of pulls fit the data equally well,
- the Newton updates can become erratic, unstable, or diverge,
- direct solutions for α may return meaningless values or NaNs.

In the 3-Flavor analysis, several of the systematic uncertainties present considerable linear dependence. An example of this is the different normalization uncertainties, which shift the event rates up or down in a very similar pattern across bins. Because of this, their corresponding fractional response vectors in V are nearly linearly dependent. Other uncertainties, such as certain calibration or cross-section shape parameters, also produce partially overlapping responses, creating moderate linear dependence. Overall, these patterns are a major source of ill-conditioning in $J^\top J$.

6.3.1 Using PCA to remove degeneracies

To avoid the issues mentioned above, we developed a solution that involves moving to a subspace where there is no linear dependency. A standard way to achieve this is to perform a Principal Component Analysis (PCA) on the fractional responses V .

If two or more systematics produce very similar shifts of the spectrum, then their "directions" in pull-space are not orthogonal. In practice, this means that a change in one pull can leak into the others, making it impossible for the Newton algorithm to decide which systematic uncertainty caused the effect. When performing a Principal Component Analysis, however, PCA rotates the systematic parameter space into a new set of axes, the principal components, where:

- the directions are orthogonal
- the degeneracies are removed
- each principal component is a linear combination of the original systematics.

In this PCA basis, J becomes numerically stable, and the Newton-Raphson method can appropriately solve for the pulls. Once the pulls are determined in the PCA space, we then map the pulls back to the original systematic uncertainties space, using the PCA rotation matrix information.

An additional benefit is that the PCA eigenvalues quantify how important each principal direction is. Components with extremely small eigenvalues correspond to directions in which the contribution of the systematic uncertainties is negligible. Removing those negligible components prevents singularities and makes the solving algorithm more robust, since those directions correspond to nearly-zero or noisy response vectors.

6.3.2 Solving for the pulls in the PCA space

To remove the linear dependencies of the fractional response matrix V , we proceed to the following eigendecomposition, allowing us to work in the principal subspace of $V^\top V$:

$$V^\top V = U \Lambda U^\top, \quad \Lambda = \text{diag}(\lambda_1 \geq \lambda_2 \geq \dots \geq 0), \quad (6.15)$$

The eigenvectors in U define the principal directions, and the eigenvalues λ_i determine how much variance lies in each direction. We keep the first k principal components, the ones with the largest eigenvalues, and denote this basis by U_k .

Reparametrizing the Problem

Instead of solving directly for the full pull vector $\mathbf{a} \in \mathbb{R}^p$, we represent it in reduced coordinates:

$$\mathbf{a} = U_k \mathbf{z},$$

where $\mathbf{z} \in \mathbb{R}^k$ are the new PCA-based parameters. This reparametrization simply rotates into a basis where degeneracies are removed, allowing Newton's method to be well behaved.

We define the residual and Jacobian in the reduced coordinates as:

$$r_k(\mathbf{z}) = r(U_k \mathbf{z}), \quad J_k(\mathbf{z}) = J(U_k \mathbf{z}) U_k \in \mathbb{R}^{n \times k}.$$

Solving the Reduced Least-Squares Problem

We solve the reduced problem

$$\min_{\mathbf{z}} \frac{1}{2} \|r_k(\mathbf{z})\|_2^2 + \frac{\alpha}{2} \|\mathbf{z}\|_2^2,$$

using Gauss-Newton or Levenberg-Marquardt iterations:

$$(J_k^\top J_k + \alpha I) \Delta \mathbf{z} = -J_k^\top r_k, \quad \mathbf{z} \leftarrow \mathbf{z} + \Delta \mathbf{z}.$$

Since we truncated the near-zero directions, these iterations are stable and converge without much numerical noise.

We stop when the step size becomes sufficiently small, $\|\Delta \mathbf{z}\| < \varepsilon$, and denote the converged reduced systematic pulls obtained in the PCA space by $\hat{\mathbf{z}}$.

Mapping Back to the Original Systematic Space

Once the reduced-space solution is found, we map it back to the original systematic space via:

$$\hat{\mathbf{a}} = U_k \hat{\mathbf{z}}.$$

This simply expresses the solution in the original systematic coordinates. It is also the orthogonal projection of the full-space solution onto the principal subspace $\text{span}(U_k)$, meaning that $\hat{\mathbf{a}}$ is the best k -dimensional approximation of the systematic pulls consistent with the

stable directions of the problem [133, 132, 134].

6.4 PCA Importance

To understand which systematic uncertainties contribute most strongly within the retained PCA subspace, we compute a simple “PCA importance” metric.

Starting again from the eigendecomposition

$$V^\top V = U \Lambda U^\top,$$

each principal direction corresponds to a column u_i of U . The entry U_{ji} is the *loading* of systematic j on principal component i . The sign of U_{ji} is arbitrary, but the magnitude $|U_{ji}|$ indicates how strongly parameter j participates in PC i .

To scale each loading by the importance of its principal component, we define:

$$I_j = \sum_{i=1}^k \lambda_i U_{ji}^2, \quad (6.16)$$

which measures how much of the constrained variance (in the k -dimensional principal subspace) aligns with systematic j .

We then normalize the importances:

$$p_j = 100 \cdot \frac{I_j}{\sum_\ell I_\ell} \quad (\%). \quad (6.17)$$

Sorting the systematics by decreasing p_j yields a natural ranking: systematics with large p_j dominate the constrained space, while those with tiny p_j live mostly in unconstrained or degenerate directions.

This provides a useful and interpretable diagnostic of which systematic parameters drive the variations within the PCA space [135, 136, 137].

6.5 Validating the Newton-Raphson method

6.5.1 Pseudoexperiment test with a single systematic shift

A first validation was performed by generating fake data where only one systematic shift was applied at a time, repeated over 40 random values. The PISCES fit provided the corresponding scaling vector β , and the Newton–Raphson method was then used to recover the injected pulls. Figure 6.2 displays injected versus recovered pulls for the PPFX and MEC systematics. In the plots, points lie close to the diagonal, showing that the method is able to recover the appropriate injected shifts.

6.5.2 Pseudoexperiment test with multiple systematic shifts

In this test, 200 pseudoexperiments were generated. Each pseudoexperiment included a set of known systematic shifts. For each pseudoexperiment t , true pulls were randomly injected as $\mathbf{a}^{\text{true},(t)}$ for up to 3σ . The PCA-Newton-Raphson solver was then applied to obtain the recovered pulls $\hat{\mathbf{a}}^{(t)}$. Where the solution was found first in the PCA space and then mapped back to the original systematic basis.

For each systematic j , the following quantities are defined:

$$\Delta a_j^{(t)} = \hat{a}_j^{(t)} - a_j^{\text{true},(t)}, \quad \mu_j = \frac{1}{N} \sum_t \Delta a_j^{(t)}, \quad \sigma_j = \sqrt{\frac{1}{N} \sum_t (\Delta a_j^{(t)})^2 - \mu_j^2}. \quad (6.18)$$

Here, $\Delta a_j^{(t)}$ represents the difference between the recovered and injected pulls. Ideally, the mean μ_j should remain close to zero, while the RMS σ_j is expected to be small for systematics with large PCA importance. Larger RMS values may appear for systematics of low importance, as these correspond to systematic uncertainties that are not fully captured within the retained top- k PCA components.

Figure 6.3 presents the validation results for the ND samples with 200 pseudoexperiments.




Figure 6.2: Preliminary validation test: comparison of injected versus recovered pulls for single-shift toy experiments. Results are shown for MEC Comp 1, MEC Comp 2, PPFX Comp 1, and PPFX Comp 2. The diagonal line indicates perfect agreement, and the clustering of points along this line indicates that the Newton–Raphson method accurately recovers the injected pulls in the one systematics at a time test.


The upper panel shows the recovered-true pulls for the most important systematics, with individual and cumulative variance contributions indicated above and below each entry. The lower panel shows the corresponding results for less important systematics, whose contributions are at the sub-percent level and therefore less constrained. High-importance systematics show small bias ($|\mu_j|$) and modest σ_j . As the importance p_j decreases, σ_j grows and μ_j fluctuates, which is expected as these low importance systematics are not fully represented in the retained PCA components; however, the larger fluctuations in these components do not affect the interpretation of the pulls, since their contributions are negligible.

6.6 Systematic importance


6.6.1 SHAP importance

The PCA importance described in the previous section ranks systematic uncertainties by how much overall variance they contribute to the covariance matrix. This quantity is the same for a given matrix, regardless of which data we are fitting. Another very useful piece of information about the fit, however, would be to identify which systematic uncertainties are driving the fit for a specific dataset. This is crucial to our understanding of the fit behavior, as a systematic uncertainty with a small pull can shift the prediction more than a less impactful uncertainty with a larger pull, depending on the corresponding V_{ij} responses. To access this information, we calculate a new quantity, called "SHAP importance", for each systematic after computing the pull values from the Newton-Raphson+PCA method. The SHAP importance evaluates both the size of the pull and how impactful that pull is in shifting the prediction (via V_{ij}) [138].

SHAP stands for SHapley Additive exPlanations. It originates from Shapley values in cooperative game theory, which fairly split a model's prediction among the features. SHAP is commonly used in machine learning to determine the local attribution of a feature by



(a) Recovered – True pulls (Δa_j) for the most important systematics. Percentages above each bar show the relative importance of each systematic, while cumulative contributions are displayed below. The pulls remain centered around zero as expected.



(b) Recovered – True pulls for less important systematics. Their contributions are below the percent level, leading to larger RMS fluctuations and looser constraints. The cumulative importance remains close to 100%.

Figure 6.3: Validation of the Newton–Raphson with PCA method using 200 pseudoexperiments. The method successfully recovers the injected pulls for leading systematics (top), while subdominant systematics remain less constrained but have a negligible impact on the overall result (bottom).

calculating how much each feature pushes the prediction. The global importance is then the average attribution across many examples, used to rank which features matter most overall.

SHAP provides a way to break a model's prediction into contributions from each input feature. For a purely linear model,

$$f(x) = b + \sum_{j=1}^d w_j x_j,$$

SHAP values take a simple closed form. Using a baseline input $x^{(0)}$, the contribution of feature j is [138]:

$$\phi_j = w_j (x_j - x_j^{(0)}), \quad \phi_0 = f(x^{(0)}).$$

Each SHAP value is just the model weight w_j multiplied by how far the feature is from its baseline. Because the model is linear and has no interactions, these values exactly match the Shapley values from cooperative game theory.

This can be applied to the PISCES systematic-pulls analysis. The predicted event rate in bin i can be written as

$$f_i(a) \approx f_i^{(0)} + \sum_j V_{ij} (a_{\text{pulls}})_j,$$

where $f_i^{(0)}$ is the nominal prediction, V_{ij} is the 1σ response of bin i to systematic j , and $(a_{\text{pulls}})_j$ is the fitted systematic pull (obtained from the Newton-Raphson+PCA method). Taking the baseline to be the no-pull point ($a^{(0)} = 0$), the SHAP attribution for bin i and systematic j becomes:

$$\phi_{ij} = V_{ij} (a_{\text{pulls}})_j.$$

The global importance of a given systematic uncertainty is obtained with:

$$I_j = \frac{1}{N_{\text{bins}}} \sum_i |V_{ij} (a_{\text{pulls}})_j|, \quad (6.19)$$

which is the sum of the contribution of systematic j across the different bins i .

The SHAP importance is useful for the systematic-pulls analysis, as ϕ_{ij} shows how much systematic uncertainty j is moving bin i , and I_j ranks which systematic uncertainties matter the most in the considered data fit. This makes SHAP an effective tool for identifying which systematics influence the fit result. Examples of how the SHAP importance highlights these shifts are shown in the next section and in Chapter 7.

6.6.2 Examples of SHAP importance rankings

Fake data test with the Neutral Current background doubled

One way to illustrate the need for SHAP importance ranking is to test it on fake data samples. In this example, the fake data is created by doubling the nominal Neutral Current background interactions in the ν_e samples. After performing a fit and calculating the pulls with the Newton-Raphson method, the SHAP importance for each systematic uncertainty is computed to rank them. The pulls and importance from this study are shown in Fig. 6.4. The figure shows the first 20 systematic uncertainties that are driving the fit. The analysis indicates that when the NC background is artificially inflated, all NC-related systematic knobs are pulled strongly, as anticipated. For instance, **MaNCRES**, which is related to the axial mass for NC resonant interactions, is pulled by 2.6σ . Other NC-related systematic uncertainties are significantly pulled as well, such as the NC-DIS (Deep Inelastic Scattering) responses. The different **DISvNC** systematics receive strong positive pulls, which is the expected behavior when enhancing the NC background. Some Final State Interaction uncertainties, such as **FormZone** (formation zone) and **hNFSI_MFP**, also respond, helping resolve remaining tensions as the NC backgrounds are FSI-sensitive. Systematic uncertainties that are not sensitive to NC-background changes show small pulls and low SHAP importance. This demonstrates that the SHAP importance effectively identifies physics-motivated systematics that drive the fit.




Figure 6.4: SHAP importance ranking for the fake data test in which the Neutral Current (NC) background in the ν_e samples is artificially doubled. The plot shows the top 20 systematic uncertainties driving the fit, together with their recovered pulls from the Newton-Raphson method. NC-related systematics (e.g. MaNCRES, NC-DIS terms) respond strongly, as expected when the NC background is inflated.

Test with fake data generated at a different oscillation parameter

Another test of the SHAP importance is to create fake data generated at an Asimov point far from the preferred oscillation parameters in NOvA. As a reminder, the preferred oscillation point in NOvA is Asimov A, the 2024 Best-Fit Point, with a preference for the normal hierarchy for Δm_{32}^2 , the upper octant for $\sin^2 \theta_{23}$, and $\delta_{CP} \approx 0.88\pi$:

$$\Delta m_{32}^2 = 2.44 \times 10^{-3} \text{ eV}^2, \quad (6.20)$$

$$\sin^2 \theta_{23} = 0.55, \quad (6.21)$$

$$\delta_{CP} = 0.88\pi. \quad (6.22)$$

A point far from this is Asimov E, which favors the inverted hierarchy, a lower-octant

value of $\sin^2 \theta_{23}$, and $\delta_{CP} = \frac{\pi}{2}$:

$$\Delta m_{32}^2 = -2.44 \times 10^{-3} \text{ eV}^2, \quad (6.23)$$

$$\sin^2 \theta_{23} = 0.46, \quad (6.24)$$

$$\delta_{CP} = 0.5\pi. \quad (6.25)$$

Creating fake data in Asimov E and performing a fit with the prediction fixed at Asimov A changes the $\nu_e/\bar{\nu}_e$ rates and their energy dependence. In such a fit, the systematic uncertainties that respond should be those sensitive to ν_e appearance. Figure 6.5 shows the 20 systematic uncertainties with the highest SHAP importance and their corresponding pulls. Appearance-specific systematics respond appropriately. For example, the ν_e radiative correction shifts by -0.92σ , helping match the ν_e vs. $\bar{\nu}_e$ rate and shape asymmetry in the Asimov E fake data. The axial CCQE knobs **EV1_4**, which encode ν vs. $\bar{\nu}$ differences through axial-vector interference, are also used to resolve tensions. The **2ndclasscurr** systematic, which adjusts second-class current terms affecting $\nu/\bar{\nu}$ differently, is pulled by -1.3σ . This is the kind of systematic that correctly accounts for the $\nu_e/\bar{\nu}_e$ asymmetry when fitting fake data far from the preferred parameter region. This again illustrates how SHAP importance identifies which systematic uncertainties drive the fit and whether any unexpected, non-physical pulls would dominate. Additional fake data studies and their SHAP analyses are presented in Chapter 7.



Figure 6.5: SHAP importance ranking for fake data generated at Asimov E and fit using the Asimov A prediction. The top 20 systematic uncertainties with the highest SHAP importance are shown. Appearance-sensitive systematics (e.g. ν_e radiative corrections, axial CCQE knobs, second-class current terms) acquire significant pulls, reflecting the shifted $\nu_e/\bar{\nu}_e$ rates and shapes in the Asimov E fake data.

Chapter 7. Fake Data Studies

To evaluate the robustness of the PISCES method, we performed several tests using dedicated fake-data studies. These included tests where specific models in the interaction have been modified, and tests where predictions with different central values are used in the fit. This was done to assess effects from potential biases in the fitter. For each different Fake Data analysis, the shifts in contours and Best-fit points are compared to those from nominal Asimov fits. A robust fitter leads to minimal shifts in contours and Best-fit values, whereas a considerable shift would indicate bias in the fitter. This chapter presents results from the various fake data tests done with PISCES. The following cases were considered in the fake data tests, where each fake data set has been evaluated in the 5 different Asimov reference points :

1. MINERvA MEC Fake Data Study;
2. CRPA Fake Data Study;
3. Increased NC background for ν_e samples Fake Data Study;
4. Using predictions with different cross section central value weights.

7.1 MINERvA MEC Fake Data Study

This Fake Data Study demonstrates a case where specific interaction models in the simulation are modified. In this analysis, fake data is generated using a model for the Meson Exchange Current (MEC) that differs from the nominal one used in NOvA. The default MEC model for NOvA is the Valencia model, whereas in this study the MINERvA MEC model

was used for the fake data. The motivation for exploring different MEC Models stems from the fact that the largest change to the NOvA central value during cross section tuning arises from the MEC interactions. Therefore, it is essential to assess how the results change when the model used to describe MEC interactions is altered. To investigate this, the default Valencia MEC used in the NOvA analysis has been compared to the case where the MINERvA MEC model is employed. The MINERvA MEC tune represents a distinct parametrization of an adjustment to the Valencia model, featuring a 2D Gaussian enhancement in (q_3, q_0) , based on data from the MINERvA experiment. [139]

In the NOvA framework, this comparison is implemented using a weight associated with the MINERvA MEC tune: the `kMINERVA_Wgt_MEC` weight. A comparison of the effect of the weight on the different samples can be found in Fig. 7.1. For the fit, the predictions and other inputs used remain consistent with those used for the nominal Asimov analysis. It should be noted that the MEC uncertainty applied in the analysis is the MEC Model spread uncertainty. The predictions used are the predictions without the ND Data-tuned cross section CV weight.

7.1.1 MINERvA MEC fake data distribution compared to nominal Asimov data

To evaluate the impact of the MINERvA MEC weight, Fig. 7.1 shows a comparison between the nominal Asimov with fake data generated using the MEC weight across the ND and FD samples included in PISCES. The largest discrepancies appear in the ν_μ samples, particularly in ν_μ Q1 FHC, ν_μ Q2 RHC, and ν_μ Q3 RHC, where variations can reach up to 15% for the FHC samples and 20% for RHC. Other ν_μ samples exhibit more moderate shifts. As for the ν_e samples, the most noticeable differences occur in the High PID categories, with deviations around 10%, while for ν_e Low PID, the change is under 5%.




Figure 7.1(a): ν_μ Q1 samples - Comparison between the nominal Asimov (blue) and fake data generated with the MINERvA MEC weight (orange) for all the ND and FD samples used in PISCES. The largest deviations are observed in ν_μ Q1 FHC, Q2 RHC, and Q3 RHC samples (up to 15–20%), with moderate effects elsewhere.




Figure 7.1(b): ν_μ Q2 samples - Comparison between the nominal Asimov (blue) and fake data generated with the MINERvA MEC weight (orange) for all the ND and FD samples used in PISCES. The largest deviations are observed in ν_μ Q1 FHC, Q2 RHC, and Q3 RHC samples (up to 15–20%), with moderate effects elsewhere.




Figure 7.1(c): ν_μ Q3 samples - Comparison between the nominal Asimov (blue) and fake data generated with the MINERvA MEC weight (orange) for all the ND and FD samples used in PISCES. The largest deviations are observed in ν_μ Q1 FHC, Q2 RHC, and Q3 RHC samples (up to 15–20%), with moderate effects elsewhere.




Figure 7.1(d): ν_μ Q4 samples - Comparison between the nominal Asimov (blue) and fake data generated with the MINERvA MEC weight (orange) for all the ND and FD samples used in PISCES. The largest deviations are observed in ν_μ Q1 FHC, Q2 RHC, and Q3 RHC samples (up to 15–20%), with moderate effects elsewhere.




Figure 7.1(e): ν_e FD samples - Asimov vs MINERvA MEC.




Figure 7.1(f): ν_e ND samples - Comparison between the nominal Asimov (blue) and fake data generated with the MINERvA MEC weight (orange) for all the ND and FD samples used in PISCES. For ν_e samples, High PID categories show deviations around 10%, and Low PID under 8%.




Figure 7.1(g): ν_e low-energy samples - Comparison between the nominal Asimov (blue) and fake data generated with the MINERvA MEC weight (orange) for all the ND and FD samples used in PISCES. Effects on the low-energy samples are under 5%.

7.1.2 Comparison of MINERvA MEC with the Asimov contours for fake data generated at the NOvA 2024 Best-Fit point

To probe for potential biases in the fitting framework, sensitivity contours obtained with the MINERvA MEC fake data, generated at the NOvA 2024 Best-Fit point, were compared to the nominal Asimov. The sets of plots in Fig. 7.2 display the 1σ , 2σ , 3σ contours for the $\sin^2 \theta_{23}$ vs Δm_{32}^2 and δ_{CP} vs $\sin^2 \theta_{23}$ parameter space, considering normal ordering and inverted ordering. The shifts in contours are minimal in all 4 combinations. Concerning the best-fit values of the profiled parameters with the MINERvA MEC fake data, the extracted values for δ_{CP} and Δm_{32}^2 remained the same as the Asimov case, while the deviation in $\sin^2 \theta_{23}$ is less than 0.4%. The contours and best-fit values remain stable, indicating a robust fit.

Furthermore, the changes in $\sqrt{\Delta\chi^2}$ have been examined across the different parameter space, where the sensitivity difference ΔS is defined as:

$$\Delta S = \sqrt{\Delta\chi_{Asimov}^2} - \sqrt{\Delta\chi_{FakeData}^2} \quad (7.1)$$

Figure 7.3 illustrates this difference in the $\Delta\chi^2$ surfaces for the two analyzed cases. Across the different points in parameter space, the difference ΔS is small, indicating that the MINERvA MEC model variation has minimal effect on the fit sensitivity.

For a more detailed analysis of the best-fit values, Fig. 7.4 showcases the comparison of the fake data spectra and the best-fit spectra with the systematic weights obtained after the fit. Across all instances, the best-fit spectra match the data well, showing that the PISCES ND+FD fit is able to provide a good fit to the considered fake data.




Figure 7.2: 1σ , 2σ and 3σ contours comparison considering the MINERvA MEC weighted fake data, generated at the 2024 Best-Fit point (dashed lines), and the nominal Asimov (solid lines), for the $\sin^2\theta_{23}$ vs Δm_{32}^2 and δ_{CP} vs $\sin^2\theta_{23}$ parameter space. The first rows are considering the normal ordering for Δm_{32}^2 , the second rows are for the inverted ordering. The shifts in contours and Best-fit points are minimal, showing that no significant bias was found and that the fitter is robust against the considered change.




Figure 7.3: Difference in $\Delta\chi^2$ surfaces between Asimov and MINERvA MEC fake data, generated at the 2024 Best-Fit point. The variation across all the different analysis bins is minimal in all 4 parameter space combinations.




Figure 7.4(a): ν_μ Q1 samples - Comparison of the MINERvA MEC fake data generated at the 2024 NOvA best-fit point (black) and the best-fit spectra using PISCES ND+FD fit with post-fit systematic weights (blue histograms). The best-fit spectra match the data well across all ND and FD samples.




Figure 7.4(b): ν_μ Q2 samples - Comparison of the MINERvA MEC fake data generated at the 2024 NOvA best-fit point (black) and the best-fit spectra using PISCES ND+FD fit with post-fit systematic weights (blue histograms). The best-fit spectra match the data well across all ND and FD samples.




Figure 7.4(c): ν_μ Q3 samples - Comparison of the MINERvA MEC fake data generated at the 2024 NOvA best-fit point (black) and the best-fit spectra using PISCES ND+FD fit with post-fit systematic weights (blue histograms). The best-fit spectra match the data well across all ND and FD samples.




Figure 7.4(d): ν_μ Q4 samples - PISCES ND+FD fit best-fit spectra.




Figure 7.4(e): ν_e FD samples - Comparison of the MINERvA MEC fake data generated at the 2024 NOvA best-fit point (black) and the best-fit spectra using PISCES ND+FD fit with post-fit systematic weights (blue histograms). The best-fit spectra match the data well across all samples.




Figure 7.4(f): ν_e ND samples - Comparison of the MINERvA MEC fake data generated at the 2024 NOvA best-fit point (black) and the best-fit spectra using PISCES ND+FD fit with post-fit systematic weights (blue histograms). The best-fit spectra match the data well across all samples.




Figure 7.4(g): ν_e low energy samples - Comparison of the MINERvA MEC fake data generated at the 2024 NOvA best-fit point (black) and the best-fit spectra using PISCES ND+FD fit with post-fit systematic weights (blue histograms). The best-fit spectra match the data well across all samples.

7.1.3 Systematic pulls for the MINERvA MEC Fake Data

The pull distribution obtained from the MINERvA MEC fake data shows several characteristic responses that are consistent with the nature of the fake dataset. Detector systematics exhibit negligible pulls, confirming that the deviation is not driven by detector modeling. Instead, the fit responds predominantly in MEC/QE-like directions, as expected. In particular:

- The **MEC component 2** direction is pulled by nearly -3σ , even though it carries a very small importance fraction ($\sim 0.02\%$). This indicates that the fit is attempting to absorb the shape difference between the Valencia MEC nominal model and the MINERvA MEC fake data.
- Related shape systematics such as **RPA Shape Suppression (2020)** and **RPA Shape Enhancement (2020)** also show positive pulls of $+2.1\sigma$ and $+2.0\sigma$, respectively, consistent with their role in modifying QE-like kinematic regions affected by MEC contributions.

Other systematics exhibit smaller responses:

- **MvCCRES** shows a modest positive pull of $+0.6\sigma$ (1.34% importance).
- **RES Δ Scale** shifts negatively by -1.3σ (0.53% importance).

These results show that the Newton-Raphson method is able to correctly identify which are the systematic uncertainties driving the fit. They also show that PISCES is able to assign the data/MC differences to the appropriate systematic uncertainties, which are the MEC uncertainties in this case.




Figure 7.5: Systematic pulls using the MINERvA MEC model fake data. The fitter responds primarily in MEC and QE-like systematic directions, as expected.

7.2 CRPA Fake Data Study

This Fake Data Study explores the effect of the new CCQE model implemented in GENIE called CRPA (Continuum Random Phase Approximation) [140]. Implementing neutrino interaction models in event generators requires a fast method to calculate the differential cross section as a function of outgoing particle kinematics. For CCQE interactions, calculating a five-dimensional cross section as a function of outgoing lepton and nucleon kinematics is required. For the model to properly represent the underlying microscopic theory, it is indispensable that the fully exclusive cross section must be calculable (*i.e.*, the cross section needs to be available as a function of the kinematic variables describing the entire final state). Most models available, however, provide only an inclusive cross section, in which the results depend solely on the outgoing lepton kinematics. In this case, MC generators use a "factorization approach", in which lepton kinematics are calculated directly from the microscopic model, and the hadronic system is subsequently added using approximations. This method can generate an accurate reflection of the microscopic models for the lepton kinematics, but is limited in its description of the outgoing hadron kinematics, as it provides only a broad estimation. The CRPA model was developed to mitigate this issue. It employs a multipole decomposition [140] to calculate the energy and angle of the outgoing nucleon, thus providing a more reliable description of the hadronic system and offering additional information beyond inclusive observables.

CRPA predicts a significantly different evolution of the cross section when considering neutrinos or anti-neutrinos and when varying the neutrino energies. The discrepancies are most pronounced in interactions involving low energy and low momentum transfer. The impact of nucleon interactions within the CRPA framework differs significantly from the RPA corrections within GENIE's standard model. Hence, it is crucial to test the current framework against the new CRPA model to ensure that the systematic errors properly account

for any differences and that no significant bias is found, and the fitter is robust against this CRPA model modification. [141]

To analyze CRPA effects within NOvA, the authors of the CRPA model generated 1 million CCQE GENIE events using the NOvA flux on a Carbon target, considering CRPA and the NOvA 2020 GENIE tune. From these events, weights to modify CCQE interactions in the NOvA MC were computed, and then used to generate the fake data for this analysis. In the NOvA samples, specifically for ν_μ CCQE events, the CRPA model results in an increase in CCQE total cross section as a function of neutrino energy, on the order of 8%. This increase originates primarily from the low Q^2 region, shifting the peaks to lower q_0 and q_3 . The CRPA model has shown the strongest effects for the ν_μ Q1 samples for both FHC and RHC, as illustrated in Fig. 7.6, which shows the difference between the nominal prediction for NOvA and the expected distribution when switching to the CRPA model.

7.2.1 CRPA fake data distribution compared to nominal Asimov data

Figure 7.6 illustrates the effect of the CRPA weight for each sample distribution compared to the nominal Asimov with fake data. ν_μ Q1 FHC and RHC are the samples with the largest discrepancies, with variations reaching up to 35% for certain bins. This effect is seen both in the ND and the FD. CRPA has minimal influence on the other samples with overall variations of less than 10%.

7.2.2 Comparison of CRPA with the Asimov contours for fake data generated at the NOvA 2024 Best-Fit point

In probing for potential bias in the fitter in the case where we consider the CRPA model for CCQE events, sensitivity contours produced with the CRPA weighted fake data were compared to the nominal Asimov. Figure 7.7 illustrates the 2σ , 3σ contours comparison




Figure 7.6(a): ν_μ Q1 samples - Comparison between the nominal Asimov (blue) and fake data generated with the CRPA weight (orange) for all the ND and FD samples used in PISCES. The largest deviations are observed in ν_μ Q1 FHC and RHC samples, for both the ND and FD (up to 35%), with moderate effects elsewhere.




Figure 7.6(b): ν_μ Q2 samples - Comparison between the nominal Asimov (blue) and fake data generated with the CRPA weight (orange) for all the ND and FD samples used in PISCES. The effect of CRPA is moderate on these samples.




Figure 7.6(c): ν_μ Q3 samples - Comparison between the nominal Asimov (blue) and fake data generated with the CRPA weight (orange) for all the ND and FD samples used in PISCES. The effect of CRPA is moderate on these samples.




Figure 7.6(d): ν_μ Q4 samples - Comparison between the nominal Asimov (blue) and fake data generated with the CRPA weight (orange) for all the ND and FD samples used in PISCES. The effect of CRPA is moderate on these samples.




Figure 7.6(e): ν_e samples - Nominal Asimov vs CRPA fake data.




Figure 7.6(f): ν_e low-energy samples - Comparison between the nominal Asimov (blue) and fake data generated with the CRPA weight (orange) for all the ND and FD samples used in PISCES. The effect of CRPA is moderate on these samples.

for the $\sin^2 \theta_{23}$ vs Δm_{32}^2 and δ_{CP} vs $\sin^2 \theta_{23}$ parameter space, considering normal ordering and inverted ordering. The shifts in contours are minimal across all four parameter-space combinations, as are the shifts in best-fit values, where the variations in δ_{CP} and Δm_{32}^2 are negligible, and the difference in $\sin^2 \theta_{23}$ is less than 0.6%. In this case, no significant bias was found, and the fitter is robust against these modifications to the RPA modeling for CCQE events.

Changes in $\sqrt{\Delta\chi^2}$ have also been evaluated across the different parameter space, where the sensitivity difference ΔS is defined in Eq. 7.1. Figure 7.8 depicts this difference in the $\Delta\chi^2$ surfaces for the two studied cases. Overall, the difference ΔS is minimal, indicating that the change in the RPA model has a negligible effect on the fit sensitivity.

A detailed analysis of the best-fit spectra is shown in Fig. 7.9. The comparison of the fake data spectra and the best-fit spectra with the systematic weights obtained after the fit indicates that, across all instances, the best-fit spectra match the data well. This shows that the PISCES ND+FD fit can provide a good fit to the CRPA modified fake data.




Figure 7.7: 1σ , 2σ and 3σ contours comparison considering the CRPA weighted fake data, generated at the 2024 Best-Fit point (dashed lines), and the nominal Asimov (solid lines), for the $\sin^2 \theta_{23}$ vs Δm_{32}^2 and δ_{CP} vs $\sin^2 \theta_{23}$ parameter space. The first rows consider the normal ordering for Δm_{32}^2 , the second rows are for the inverted ordering. The shifts in contours and Best-fit points are minimal, showing that no significant bias was found and that the fitter is robust against the considered change.




Figure 7.8: Difference in $\Delta\chi^2$ surfaces between Asimov and CRPA fake data, generated at the 2024 Best-Fit point. The variation across all the different analysis bins is minimal in all 4 parameter space combinations.




Figure 7.9(a): ν_μ Q1 samples - Comparison of the CRPA fake data generated at the 2024 NOvA best-fit point (black) and the best-fit spectra using PISCES ND+FD fit with post-fit systematic weights (blue histograms). The best-fit spectra match the data well across all ND and FD samples.




Figure 7.9(b): ν_μ Q2 samples - PISCES best-fit spectra for CRPA fake data.




Figure 7.9(c): ν_μ Q3 samples - Comparison of the CRPA fake data generated at the 2024 NOvA best-fit point (black) and the best-fit spectra using PISCES ND+FD fit with post-fit systematic weights (blue histograms). The best-fit spectra match the data well across all ND and FD samples.




Figure 7.9(d): ν_μ Q4 samples - PISCES best-fit spectra for CRPA fake data.




Figure 7.9(e): ν_e FD samples - PISCES best-fit spectra for CRPA fake data.




Figure 7.9(f): ν_e low energy samples - Comparison of the CRPA fake data generated at the 2024 NOvA best-fit point (black) and the best-fit spectra using PISCES ND+FD fit with post-fit systematic weights (blue histograms). The best-fit spectra match the data well across all ND and FD samples.

7.2.3 Systematic pulls for the CRPA Fake Data

This Fake Data Study explores the impact of using the CRPA (Continuum Random Phase Approximation) model for CCQE interactions in place of the nominal description. The CRPA model provides a more differential treatment of the hadronic system compared to the factorized approaches commonly used in generators, thereby allowing a more realistic description of both lepton and nucleon kinematics. This analysis helps assess which systematic uncertainties are responding when the underlying CCQE model is modified.

The pull distributions show that detector systematics are barely pulled on, consistent with expectations since the fake data modification targets the interaction model rather than detector response. Resonance-related parameters such as `MaCCRES`, `MvCCRES`, and `RESDeltaScale` show only mild shifts, indicating that the CRPA modifications are not absorbed through resonance degrees of freedom. In contrast, the RPA shape knobs respond strongly: `RPAShapesupp2020` exhibits a significant positive pull of $+3.49\sigma$ (PCA importance 0.042%), while `RPAShapeenh2020` shifts by $+0.79\sigma$ (PCA importance 0.033%). These large responses are consistent with the fact that CRPA explicitly modifies the nuclear response functions in a way that overlaps with the phenomenological role of the RPA knobs in the

nominal tuning. Additional smaller pulls are observed for low- Q^2 resonance suppression (LowQ2RESSupp, $+1.26\sigma$) and the Z -exp axial form factor components (e.g. $EV3 = -0.56\sigma$, $EV2/EV4 \approx +0.65\sigma$).

Overall, the pull pattern demonstrates that the PCA–Newton–Raphson method attributes the differences induced by the CRPA model primarily to RPA-related directions, with secondary responses in low- Q^2 suppression and axial form factors. These results correspond to what is expected, since a modification in CRPA modifies long-range correlation effects that are otherwise parametrized with RPA and other related shape systematics.

7.3 Increased NC ν_e Components Fake Data Study

In this analysis, the robustness of the fitter is being tested by altering the proportion of the NC background components in the ν_e samples. Specifically, the NC background is doubled. The objective of injecting such an extreme change is to probe the scenario in which the NC backgrounds are not properly estimated, and to study their impact, especially for the ν_e samples.

7.3.1 NC ν_e Background Doubled Fake Data distribution compared to nominal Asimov data

Figure 7.11 shows a comparison between the nominal Asimov with the case where the NC background components are doubled for ν_e samples. The largest effect can be seen for Low PID ν_e samples, for both FHC and RHC. Doubling the NC ν_e background modifies the fake data spectra by up to 50% for Low PID bins. On the High PID samples, the changes are around 30%. In the analysis, only NC backgrounds in the ν_e samples were doubled; the NC backgrounds in ν_μ samples were kept at their nominal values.




Figure 7.10: Systematic pulls for the CRPA fake data. The response is dominated by strong pulls in RPA-related directions as expected.




Figure 7.11(a): ν_e ND samples - Comparison of nominal Asimov and modified fake data with doubled NC backgrounds in ν_e samples, highlighting up to 50% spectral changes for ν_e Low PID bins for both FHC and RHC. The variations are in the order of 30% for High PID samples. The NC backgrounds in ν_μ samples were kept at their nominal values.




Figure 7.11(c): ν_e low-energy samples - Comparison of nominal Asimov and modified fake data with doubled NC backgrounds in ν_e low-energy samples. The variations can reach up to 60%.




Figure 7.11(b): ν_e FD samples - Asimov vs NC ν_e doubled fake data.

7.3.2 Comparison of NC ν_e backgrounds doubled with the Asimov contours for fake data generated at the NOvA 2024 Best-Fit point

To evaluate how the contours change when the NC background in the ν_e samples are doubled, Fig. 7.12 shows the 1σ , 2σ , 3σ contours comparison for the new fake data with the Asimov case, for the $\sin^2 \theta_{23}$ vs Δm_{32}^2 and δ_{CP} vs $\sin^2 \theta_{23}$ parameter space, considering normal ordering and inverted ordering. Here also, the shifts in contours are minimal in all 4 combinations. For the best-fit values of the profiled parameters with the NC ν_e background doubled, the deviation in $\sin^2 \theta_{23}$ is less than 1.5%, whereas other parameters remained around the same values. These results show no evidence of appreciable bias is observed, demonstrating the fitter's robustness.

Concerning the changes in $\sqrt{\Delta\chi^2}$, Fig. 7.13 illustrates the difference in the $\Delta\chi^2$ surfaces for the two analyzed fake data. Across the different points in parameter space, the difference ΔS is small, indicating that doubling the NC ν_e background has minimal effect on the fit sensitivity.

A more detailed analysis of the best-fit values can be seen in Fig. 7.14, which shows the comparison of the fake data spectra and the best-fit spectra with the systematic weights obtained after the fit. Across all samples, the best-fit spectra match the data well, showing that the PISCES ND+FD fit can provide a good fit to the considered fake data.

7.3.3 Systematic pulls for the NC ν_e background doubled Fake Data

The pulls recovered from the Newton-Raphson method provide insight into how the fit responds when the NC background in the ν_e samples is doubled. As expected, detector systematics are barely pulled on. The dominant responses come from parameters related to




Figure 7.12: 1σ , 2σ and 3σ contours comparison considering the NC ν_e background doubled fake data, generated at the 2024 Best-Fit point (dashed lines), and the nominal Asimov (solid lines), for the $\sin^2 \theta_{23}$ vs Δm_{32}^2 and δ_{CP} vs $\sin^2 \theta_{23}$ parameter space. The first rows are considering the normal ordering for Δm_{32}^2 , the second rows are for the inverted ordering. The shifts in contours and Best-fit points are minimal, showing that no significant bias was found and that the fitter is robust against the considered change.




Figure 7.13: Difference in $\Delta\chi^2$ surfaces between Asimov and NC ν_e background doubled fake data, generated at the 2024 Best-Fit point. The variation across all the different analysis bins is minimal in all 4 parameter space combinations.




Figure 7.14(a): ν_μ Q1 samples - Comparison of the NC ν_e background doubled fake data generated at the 2024 NOvA best-fit point (black) and the best-fit spectra using PISCES ND+FD fit with post-fit systematic weights (blue histograms). The best-fit spectra match the data well across all ND and FD samples.




Figure 7.14(b): ν_μ Q2 samples - PISCES ND+FD fit best-fit spectra.




Figure 7.14(c): ν_μ Q3 samples - Comparison of the NC ν_e background doubled fake data generated at the 2024 NOvA best-fit point (black) and the best-fit spectra using PISCES ND+FD fit with post-fit systematic weights (blue histograms). The best-fit spectra match the data well across all ND and FD samples.




Figure 7.14(d): ν_μ Q4 samples - PISCES ND+FD fit best-fit spectra.




Figure 7.14(f): ν_e FD samples - PISCES ND+FD fit best-fit spectra.




Figure 7.14(e): ν_e ND samples - Comparison of the NC ν_e background doubled fake data generated at the 2024 NOvA best-fit point (black) and the best-fit spectra using PISCES ND+FD fit with post-fit systematic weights (blue histograms). The best-fit spectra match the data well across all samples.




Figure 7.14(g): ν_e low energy samples - Comparison of the NC ν_e background doubled fake data generated at the 2024 NOvA best-fit point (black) and the best-fit spectra using PISCES ND+FD fit with post-fit systematic weights (blue histograms). The best-fit spectra match the data well across all samples.

the ν_e background modeling. Some of those are the NC DIS knobs, which exhibit noticeable shifts. This reflects the fact that doubling the NC ν_e background alters the shape and normalization of ν_e events. For example, pulls are observed in the neutral current DIS shape parameters, consistent with the increased weight of NC background in the fake data. Additionally, modest shifts in resonance parameters such as **MaNCRES** are also recovered ($+2.65\sigma$, 0.54% importance), indicating that the fit compensates partially through resonance-like variations. Overall, the results demonstrate that the fitter is sensitive to mismodeling of NC backgrounds in ν_e samples: the NC-related knobs carry the main response, while other directions (detector, flux, and subdominant cross section systematics) remain only weakly pulled.

7.4 Difference in cross section central value weight Fake Data Study

Neutrino interactions with nuclei are challenging to model and to measure, especially for low-energy neutrinos. The final state interactions are particularly difficult to evaluate. This is why neutrino interaction cross sections usually represent one of the dominant sources of systematic uncertainties in neutrino analysis. This challenging modeling process often

means that the nominal MC prediction can be different from the measured data. However, each interaction modeled in GENIE possesses distinct degrees of freedom and associated uncertainties, which can be independently tuned to enhance the agreement between data and Monte Carlo simulations. In NOvA, the different GENIE models used are first optimized to maximize consistency with the data. Furthermore, external data constraints are also applied to the simulation. Nevertheless, following these adjustments, a residual difference remains between the simulation and the data. This has been attributed to the large uncertainty in the Meson Exchange Currents. To rectify this discrepancy, those MEC components have been rescaled such that the resulting simulation matches the NOvA ND collected data. [77]

As a reminder, the main updates applied to the cross section central value prediction—introduced to improve the agreement between the simulation and NOvA data—are summarized below:

- **External constraints:**

- *Axial form factor value:* GENIE uses a default axial mass of $M_A = 0.99 \text{ GeV}/c^2$. A recent reanalysis of ANL and BNL data indicates a slightly larger value, with an error-weighted mean of $M_A = 1.04 \text{ GeV}/c^2$. The axial mass in the NOvA tune was updated accordingly.
- *Nuclear-model weights from MINERvA (QE):* More sophisticated nuclear models predict a nucleon momentum distribution that differs from the local Fermi Gas model in GENIE. These models reduce QE strength at low Q^2 and mildly increase it at high Q^2 . To account for this, MINERvA QE nuclear-model weights were applied.
- *Reduced non-resonant single-pion production:* A re-evaluation of bubble-chamber data, supported by recent MINERvA measurements, shows that GENIE overestimates non-resonant single-pion production. The data suggest a $\sim 57\%$ reduction,

so the corresponding rate in the simulation was scaled down.

- **NOvA data-driven constraints:** Existing MEC models in GENIE systematically underpredict MEC event rates, especially at low hadronic energy, mainly due to uncertainties in the MEC spectral shape. Since no available model fully describes the observed data, NOvA applies additional MEC weights derived directly from ND data. These “CV MEC weights” rescale the various MEC components so that the simulated MEC contribution matches the measured ND spectra [77].

Specifically, these modifications for the MEC events are described by the NOvA cross section central value weights. Including or removing the cross section central value weight leads to considerable differences in the prediction’s central values. In this analysis, fits using predictions where we add the ND data tuned cross section central value weight were compared with the untuned case. The comparison between the two predictions was performed by doing a fit to the same fake data sample, which was generated without the central value weight.

7.4.1 Comparison prediction with and without the cross section CV weight

Figure 7.16 illustrates the difference between including or forgoing the cross section CV weight in the predictions. The differences between the two types of predictions are significant, with variations reaching more than 50% for several bins. The discrepancies are more pronounced for ν_μ samples. For the ν_e case, High PID samples are more affected, with differences in the order of 20%. The cross section CV weight impacts nearly all samples, indicating a significant difference between the two predictions.




Figure 7.16(a): ν_μ Q1 samples - Comparison between prediction with the cross section CV weight (blue) and without (orange) for all the ND and FD samples used in PISCES. The largest deviations are observed in ν_μ Q1 samples (up to 50%).




Figure 7.16(b): ν_μ Q2 samples - Comparison between prediction with the cross section CV weight (blue) and without (orange) for all the ND and FD samples used in PISCES. Deviations observed in ν_μ Q2 samples are up to 35%.




Figure 7.16(c): ν_μ Q3 samples - Comparison between prediction with the cross section CV weight (blue) and without (orange) for all the ND and FD samples used in PISCES. The largest deviations observed in ν_μ Q3 samples are around 20%.




Figure 7.16(d): ν_μ Q4 samples - Comparison between prediction with the cross section CV weight (blue) and without (orange) for all the ND and FD samples used in PISCES. Maximum deviations of 20% for ν_μ Q4 samples.




Figure 7.16(e): ν_e FD samples - With vs without CV weight comparison.




Figure 7.16(f): ν_e ND samples - Comparison between prediction with the cross section CV weight (blue) and without (orange) for all the ND and FD samples used in PISCES. The effect on the Low PID samples is modest, while the deviations in the High PID samples can go up to 20%.




Figure 7.16(g): ν_e low energy samples - Comparison between prediction with the cross section CV weight (blue) and without (orange) for all the ND and FD samples used in PISCES. The effect on the low-energy samples is minimal.

7.4.2 Comparison of sensitivity obtained using predictions with cross section CV weight and without the CV weight

The sets of plots in Fig. 7.17 represent the 1σ , 2σ , 3σ comparison contours using predictions with/without cross section CV weight, in the $\sin^2 \theta_{23}$ vs Δm_{32}^2 and δ_{CP} vs $\sin^2 \theta_{23}$ parameter space, considering normal ordering and inverted ordering. Despite the considerable differences between the predictions, the shifts in contours were minimal in all 4 combinations. Here, both predictions yielded closely matching best-fit values, indicating that no significant bias was observed in this case, further proving the fitter’s robustness.

Figure 7.18 illustrates this difference in the $\Delta\chi^2$ surfaces for the sensitivity analysis using the two different predictions. The difference ΔS is small in all 4 surfaces, indicating that whether we include the cross section CV weight or not, the results obtained with the PISCES ND+FD fit are consistent.

Figure 7.19 showcases the comparison of the fake data generated without the cross section CV weight with the best-fit spectra obtained using prediction with the cross section CV weight. Across all samples, the best-fit spectra match the data well, showing that the PISCES ND+FD fit can provide a good fit even in the case of a considerable difference between the prediction central value and the data.

7.4.3 Systematic pulls in the CV-weight comparison study

Figure 7.20 represents the recovered pulls for the CV weight fake data. The results show that detector systematics are barely shifted. However, the strongest pulls appear in the resonance and MEC systematic uncertainties. In this case, M_a^{CCRES} is pulled by $+1.0\sigma$ (contributing about 3.4% of the PCA variance), and MEC component 1 is shifted by about $+1.1\sigma$ (roughly 1.9% share). These are consistent with the expectation when not including the ND-tuned MEC weight. Indeed, removing the weight means leaving the mismodeling in




Figure 7.17: 1σ , 2σ and 3σ contours comparison considering predictions with (dashed lines)/without (solid line) cross section CV weight, fitted for a data without cross section CV weight, generated at the 2024 Best-Fit point. The contours are for the $\sin^2 \theta_{23}$ vs Δm_{32}^2 and δ_{CP} vs $\sin^2 \theta_{23}$ parameter space. The first rows are considering the normal ordering for Δm_{32}^2 , the second rows are for the inverted ordering. The shifts in contours and Best-fit points are minimal, showing that no significant bias was found even in the case of considerable discrepancy in central values between predictions and data.




Figure 7.18: Difference in $\Delta\chi^2$ surfaces between using predictions with/without cross section CV weight for fitting for a fake data without cross section CV weight. The variations across all considered surfaces are minimal.




Figure 7.19(a): ν_μ Q1 samples - Comparison of the fake data generated without a cross section CV weight (black) and the best-fit spectra obtained using PISCES ND+FD fit with the predictions with the cross section CV weight, with post-fit systematic weights included (blue histograms). The best-fit spectra match the data well across all samples.




Figure 7.19(b): ν_μ Q2 samples - PISCES ND+FD fit best-fit spectra.




Figure 7.19(c): ν_μ Q3 samples - Comparison of the fake data generated without a cross section CV weight (black) and the best-fit spectra obtained using PISCES ND+FD fit with the predictions with the cross section CV weight, with post-fit systematic weights included (blue histograms). The best-fit spectra match the data well across all samples.




Figure 7.19(d): ν_μ Q4 samples - PISCES ND+FD fit best-fit spectra.




Figure 7.19(e): ν_e FD samples - PISCES ND+FD fit best-fit spectra.




Figure 7.19(f): ν_e ND samples - Comparison of the fake data generated without a cross section CV weight (black) and the best-fit spectra obtained using PISCES ND+FD fit with the predictions with the cross section CV weight, with post-fit systematic weights included (blue histograms). The best-fit spectra match the data well across all samples.




Figure 7.19(g): ν_e ND samples - Comparison of the fake data generated without a cross section CV weight (black) and the best-fit spectra obtained using PISCES ND+FD fit with the predictions with the cross section CV weight, with post-fit systematic weights included (blue histograms). The best-fit spectra match the data well across all samples.

MEC and resonance regions. Here, the fitter is able to attribute the data/MC differences to the appropriate systematic uncertainties.

Additional systematic uncertainties related to QE interactions are also pulled on. For instance, Z-expansion eigenvectors EV2/EV4 and EV3 are shifted by $+1.2\sigma$ and $+0.4\sigma$, respectively. These reflect adjustments in the axial form factor.




Figure 7.20: Recovered pulls for the fake data generated without applying the NOvA central value (CV) weight, fit with predictions including this weight. The fitter responds by shifting resonance (MaCCRES), MEC, and QE-like parameters, consistent with the expected mismodeling introduced by omitting the MEC rescaling.

Chapter 8. Results

This chapter presents the PISCES ND+FD fit results considering the full 10 years of NOvA data. This corresponds to a neutrino-mode (FHC) run of 26.61×10^{20} protons on target (POT), and an antineutrino-mode (RHC) run of 12.50×10^{20} POT in the Far detector. The PISCES ND+FD fit is compared to the NOvA standard Extrapolation fit, and the results are discussed in the following sections. First, an ND-only fit is performed to probe the systematic responses, then the full joint ND+FD data fit to the data is presented.

8.1 PISCES ND Fit using NOvA ND data

PISCES is a flexible framework that allows fits to be performed using auxiliary sample configurations, including ND-only fits. In this test, we performed a PISCES fit using only the NOvA ND data. The goal is to evaluate how well PISCES can reproduce the real ND data. Figures 8.1-8.2 show the comparison between the NOvA ND data distribution and the best-fit spectra obtained from the PISCES ND-only fit. The comparison shows great agreement for the ν_μ and low-energy ν_e samples. The agreement remains reasonable for the ν_e Low PID and High PID samples.

8.1.1 Systematic pulls for the NOvA ND-only Data Fit

In the ND-only fit to the NOvA ND data, the pulls indicate that the PISCES framework is using a physically sensible combination of systematic uncertainties to explain the data/MC difference, with all of the shifts staying well within 1σ . The systematics with the largest SHAP importance in the ND-only fit are the ND detector knobs. The ND global calibration, drift, and calibration shape are shifted in the range of ~ -0.3 to $+0.6\sigma$. Together, the




Figure 8.1: Best-fit spectra (blue histograms) obtained by performing a PISCES fit to the real NOvA ND data (black points), using ND samples. PISCES successfully reproduces the data across all ν_μ and ν_e ND samples. There is excellent agreement for the ν_μ and low-energy ν_e samples, and reasonable agreement for the ν_e Low PID and High PID samples.




Figure 8.2: Best-fit spectra (blue histograms) obtained by performing a PISCES fit to the real NOvA ND data (black points), using ND samples. PISCES successfully reproduces the data across all ν_μ and ν_e ND samples. There is excellent agreement for the ν_μ and low-energy ν_e samples, and reasonable agreement for the ν_e Low PID and High PID samples.

detector systematics account for about half of the total SHAP importance weight. These small, coherent adjustments of the visible-energy scale to fine-tune the ND response are within what is expected in an ND-only fit. The fact that there are no extreme pulls shows that the fitter can resolve data/MC tension without over-pulling on a specific systematic uncertainty.

The neutron response is also among the systematic uncertainties with high importance, with a SHAP score of 9%, with a pull of about $+0.2\sigma$. This suggests that how neutrinos deposit energy in the detector is an important lever in describing the ND data distribution, though the actual shift is modest. Additional detector-related systematic uncertainties, such as `key_LightLevel` and the `key_Chrenkov` knobs, also move slightly, $\mathcal{O}(0.1\text{--}0.4\sigma)$, which translates to small corrections in light yield and EM-components.

Flux systematics pulls are less significant in the ND-only fit. The leading PPFX component `key_PPFXcomp1` and the other flux PCs have small pulls (well below 1σ) and relatively low SHAP weight. This indicates that the beam flux prediction is already broadly compatible with the ND data, hence the lower shifts.

Concerning the cross section uncertainties, the fit shifts the MEC, RES, FSI, and DIS systematic knobs, all below 1σ . The first MEC principal component (`key_MECcomp1`) pulls slightly positive, indicating a mild preference for more 2p2h/MEC effects in the ND. Resonance uncertainties (`MaCCRES`, `MvCCRES`, and `RESDeltaScaleSyst`) and the FSI knobs (`FormZone2024`, `hNFSI_MFP_2024`) also shift by a few tenths of a sigma. These uncertainties modify the Δ resonance regions and the intranuclear reinteractions in a way that improves the description of topologies with high pions. DIS hadronization knobs (such as `DISNuHadronQ1Syst` and the 1π DIS modes) also show some modest adjustments, to further improve data/MC agreement.

Overall, these pulls demonstrate that the ND-only PISCES fit effectively reproduces the NOvA ND distributions, without any extreme or unphysical pulls. Instead, it spreads shifts




Figure 8.3: Recovered pulls from the ND-only fit to the NOvA real ND data. The first panel shows the systematic pulls for the most important uncertainties. The second panel represents the subdominant systematic uncertainties.

across detector calibration, neutron response, several interaction and FSI knobs, and cross section. These pull values and the best-fit spectra in Fig. 8.1-8.2, which show great data/MC agreement in all ND samples, indicate that PISCES is able to reproduce the NOvA ND data well and is attributing the residual differences to reasonable, physically motivated systematic uncertainties.

8.2 PISCES ND+FD fit results

A joint ND+FD PISCES fit was performed on the NOvA data. Figure 8.10 shows the 1σ , 2σ , and 3σ contour comparison between the Extrapolation method and the PISCES

ND+FD fit across four parameter spaces.

As a reminder, in the Extrapolation method, the ND data/MC ratio is used to predict the FD (Section 3.1). This cancels or reduces systematic effects that share the same effects between ND and FD. However, below are three cases where non-cancellation remains:

- The ratio is built per sample and per beam mode. This means that systematic uncertainties that affect neutrinos/antineutrinos differently, or ν_μ/ν_e differently do not cancel across samples, as the ratios are formed independently for each sample.
- The ND and FD have different geometries, even though they are functionally similar. They also have different acceptance and containment. These effects that are not identical across detectors and selections leave residuals that can be degenerate with the oscillation parameters.
- Extrapolation uses transfer matrices to map ND information to the FD, but it does not enforce all correlations. Extrapolation does not explicitly carry a full cross-sample covariance information in the fit, so off-diagonal correlations are under-profiled.

By contrast, PISCES builds one covariance across all samples at once and further subdivides the samples into the different oscillation channels. This enforces cross-sample correlations. Furthermore, PISCES can use sample response differences to break degeneracies. This extra information gives PISCES additional constraints over the Extrapolation method.

8.3 PISCES systematic pulls

The systematic pulls corresponding to the PISCES joint ND+FD fit, obtained using the Newton-Raphson method, are presented in Fig. 8.4. Overall, the patterns of pulls in the joint ND+FD fit are very similar to those seen in the ND-only fit. The pulls are all modest with most systematic uncertainties shifted well below $\sim 1\sigma$). Looking at the pull with the highest SHAP importance, the shifts are distributed among several physically motivated

systematic uncertainties. This indicates that the fit is well-behaved and that the fitter is not relying heavily on any single systematic uncertainty to absorb data/MC differences. Detailed below are the primary drivers of the fit and the secondary systematic uncertainties that have noticeable effects.

Calibration (global energy scale)

$\text{key_Calib} = -0.311 \sigma$ (*SHAP* 9.30%). The largest single contribution with the highest SHAP importance comes from the calibration systematic uncertainty, which shifts the visible energy scale. The fit prefers a small downward shift, consistent with what was seen in the ND-only fit.

Flux systematics

- $\text{key_PPFXcomp1} = -0.483 \sigma$ (9.18%)
- $\text{key_PPFXcomp4} = -0.465 \sigma$ (2.03%)
- $\text{key_PPFXcomp5} = +0.447 \sigma$ (1.52%)
- $\text{key_PPFXcomp2} = -0.213 \sigma$ (1.24%)

As detailed in Section 4.1, the hadron production uncertainties are encoded in the set of "PPFX components". Each component modifies the predicted pion and kaon yields in a correlated manner. When a component is shifted, it alters the neutrino flux prediction as a function of energy and changes the relative contribution of pions and kaons across different energy regions. This matters as pions predominantly generate flux at lower energy near the flux peak, whereas kaons generate higher-energy neutrinos. In the joint ND+FD fit, the overall pattern of flux pulls is similar to the ND-only fit, where key_PPFXcomp1 shifts negative, while key_PPFXcomp4 and key_PPFXcomp5 pull with opposite signs to modify the pion-kaon mixture. The main difference is that the magnitudes of the shifts increase once




Figure 8.4: Systematic pulls for the PISCES ND+FD Fit.

the ND data are included, which indicates that the joint fit relies more on the flux model to improve data/MC agreement.

FSI / intranuclear cascade

- $\text{hNFSI_MFP_2024} = -0.880 \sigma \ (5.87\%)$
- $\text{FormZone2024} = -0.217 \sigma \ (1.88\%)$
- $\text{hNFSI_FateFracEV3_2024} = +0.199 \sigma \ (0.37\%)$

hNFSI_MFP_2024 and FormZone2024 are systematic uncertainties related to final state interactions (FSI), as detailed in Section 4.3.7. They modify the hadron mean free path and the hadron formation zone inside the nucleus, which in turn changes the amount of intranuclear reinteraction and absorption experienced by outgoing hadrons. The shift directions are consistent with the ND-only fit. These FSI uncertainties affect ν and $\bar{\nu}$ differently as ν CC interactions tend to produce more π^+ while $\bar{\nu}$ CC interactions produce more π^- , and the intranuclear reinteraction probabilities are charge-dependent. As a result, changes in mean free path and formation zone modify the absorption and scattering of π^+/p^- inside the nucleus. This changes the outgoing hadron energy and thus, the neutrino reconstructed energy. These uncertainties are therefore important, as they help disentangle true ν and $\bar{\nu}$ asymmetry from oscillation from apparent asymmetries introduced by flux or hadronic modeling.

Neutron systematic

The neutron transport systematic ($\text{key_Menate} = +0.321 \sigma$; 5.24%) is the systematic uncertainty associated with how neutrons propagate and deposit energy in the detector (Section 4.4). It affects the amount of energy carried by secondary interactions and captures from neutrons, and therefore changes the contribution of neutrons to the visible hadronic energy. In the joint ND+FD fit, the direction also agrees with the ND-only fit, where there

is a preference for an increase in the modeled neutron activity.

RES (Δ) and DIS systematics

- $\text{MvCCRES} = -0.348 \sigma$ (5.89%)
- $\text{RESDeltaScaleSyst} = -0.383 \sigma$ (4.76%)
- $\text{RESOtherScaleSyst} = -0.257 \sigma$ (1.63%)

This set of pulls adjusts resonant (RES) and deep inelastic scattering (DIS) interactions (Section 4.3). GENIE tends to overestimate the Δ resonance region around 1–2 GeV, hence the downward shift for the RES systematic uncertainties.

The DIS channels are also pulled modestly to adjust how the DIS events are distributed. The joint ND+FD fit pulls DISvnCC1pi_2020 upward ($+0.461 \sigma$; 3.34%) and suppresses the NC 1π components with DISvnNC1pi_2020 ($= -0.572 \sigma$; 2.29%), and $\text{DISvbarpNC1pi_2020}$ ($= -0.665 \sigma$; 1.37%). These shifts are moderate but are important, as RES and DIS modeling directly influence the hadronic final state. Given that the ν reconstructed energy is the sum of the shower energy and the visible hadronic energy, shifts in RES or DIS modeling can change the shape of the ν_e appearance spectrum. As a result, RES and DIS systematics are also among the key uncertainties that influence the reconstructed neutrino energy, and they are particularly important for the ν_e appearance spectrum in the joint ND+FD fit.

MEC / 2p2h response

- $\text{key_MECcomp1} = +0.217 \sigma$ (4.30%)
- $\text{key_MECcomp2} = +0.170 \sigma$ (0.89%)
- $\text{MECEnuShape2020AntiNu} = -0.460 \sigma$ (1.21%)
- $\text{MECEnuShape2020Nu} = -0.016 \sigma$ (0.05%)

The MEC interaction uncertainties used here correspond to the model-based MEC uncertainty described in Section 4.3.5. The fit prefers a mild increase in the overall MEC contribution, consistent with a lower predicted hadronic energy in the MEC region. For the MEC shape systematics, the antineutrino MEC `MECEnuShape2020AntiNu` component pulls negative. This affects the RHC spectra by modifying the spectrum between QE and RES regions. The neutrino MEC shape remains essentially unchanged as the shift on `MECEnuShape2020Nu` is small.

Detector optics / light modeling

- `key__CalibShape` = $+0.207 \sigma$ (3.57%)
- `key__LightLevel` = $+0.087 \sigma$ (0.79%)
- `key__Cherenkov` = $+0.014 \sigma$ (0.16%)

After the global calibration energy shift, the fit uses these other detector systematics knobs to apply extra adjustments to the detector light response. The magnitudes are small, and the signs are coherent with the pattern already seen in the ND-only fit. None of these dominate the fit, but they help correct subtle shower/track light-response differences, which especially matter for ν_e selection.

Muon energy scale

- `CorrMuEScaleSyst2024` = -0.060σ (0.58%)
- `UnCorrMuEScale2024` = -0.212σ (0.56%)

The muon energy scale uncertainties account for how accurately the detector converts the observed muon track length or curvature into reconstructed muon energy. In the ND+FD fit, the muon energy scale shifts are small. Here, the FD spectra do not lean much on the special muon energy scale fix. Instead, the joint fit is resolving the tension with the other physics knobs (flux, RES, FSI, MEC), rather than modifying the muon scale of the FD detector.

Summary (PISCES ND+FD fit)

Overall, the joint ND+FD PISCES fit uses several groups of physically relevant systematic uncertainties to drive the fit. The largest SHAP contributions come from the global energy calibration, the flux components, and the hadronic interaction systematics (RES, FSI, and MEC). These directly impact the reconstructed energy and the hadronic final state, so it makes sense that they influence both the ν_μ and $\nu_e/\bar{\nu}_e$ channels the most. In comparison, the other detector effects, such as light level and Cherenkov, and the muon energy scale, contribute at a secondary level. Their shifts stay well below 1σ , and their SHAP importance is low.

A key point is that the overall pattern of pulls in the joint ND+FD fit matches what was seen in the ND-only fit. Altogether, the joint ND+FD fit makes reasonable, physically motivated adjustments across several systematics, rather than relying on any single parameter to explain the data. The pulls stay compatible with the ND constraints, and none of them show unusually large shifts.

8.4 Systematic Pulls: PISCES vs. Extrapolation

Table 8.1 shows a comparison of the systematic pulls obtained with PISCES and with the Extrapolation method. As a reminder (Section 3.1), in the NOvA Extrapolation method, the ND information enters through a data/MC ratio, which is applied to correct the FD prediction to produce an "extrapolated FD prediction". This modified extrapolated prediction is then used to perform an FD-only oscillation fit. In contrast, PISCES fits the ND and FD samples simultaneously, with the systematic uncertainties encoded in a joint ND+FD covariance matrix.

The advantage of PISCES, in this case, is the direct use of the high-statistics ND samples, which allow the fit to constrain flux shapes, cross sections, final-state interactions FSI, and detector response, instead of compressing the ND information in a one-dimensional ratio,

like in the Extrapolation method. Furthermore, in an FD-only fit, several combinations of the different systematic uncertainties can give very similar FD spectra; therefore, the fit has strong degeneracies. Including the ND directly in the likelihood helps break many of these degeneracies. For example, in the fit to the NOvA real data, the Extrapolation fit prefers a strong -1σ pull on the calibration systematic uncertainty, whereas the PISCES ND+FD fit uses a more moderate calibration shift and redistributes the tension among flux and cross sections systematic uncertainties.

Global calibration response

In the Extrapolation method, the FD-only fit drives the calibration systematic to nearly -1σ . The PISCES ND+FD fit, on the other hand, still prefers a negative calibration, but with a less extreme shift of $\sim -0.33\sigma$. For the Extrapolation analysis, as the ND constraints have been compressed into a ratio, the fit is not directly sensitive to whether such a strong pull might be inconsistent with the ND data. In PISCES, the joint ND+FD fit shifts the calibration in the negative direction, but with a significantly more moderate pull. With the simultaneous ND+FD fit, the ND samples constrain the detector response, and a large global calibration shift is disfavored by the ND. Instead, PISCES reallocates residual tension into a combination of the physics-driven pulls, such as flux, RES interactions systematics, and Final State Interactions. This illustrates a key advantage of PISCES, where it avoids overusing detector-only knobs and instead uses other physically motivated uncertainties that are constrained by the ND.

Flux systematics

Extrapolation pulls on the first flux PPFX component by $\sim -0.26\sigma$, whereas PISCES pulls slightly stronger in the opposite direction with a value of $\sim 0.48\sigma$ (with a SHAP importance of 9.8%). For the remaining flux components, the Extrapolation fit shifts the `ppfx_hadp_pc` modes at the $\sim 0.2\sigma$ level with a mixture of positive and negative pulls. In

PISCES, PPFX components 4 and 5 are both shifted at around 0.46σ , and are among the systematic uncertainties with high SHAP importance. Here, as PISCES pulled moderately on the calibration, the extra tension is distributed into other systematic uncertainties, explaining the slightly higher pull on the first PPFX flux uncertainty. The ND information in PISCES provides strong leverage to constrain flux and cross section uncertainties; and the joint fit makes more use of those systematics. In the Extrapolation method, the flux, cross section, and detector constraints are already folded into the ND/FD ratio, so the FD-only fit does not fully exploit those uncertainties; instead, data/MC differences are mostly absorbed by the large calibration pulls.

Final State Interactions (FSI)

The other systematic uncertainty that is among those with the highest SHAP with PISCES is the `hNFSI_MFP`, which represents the Final State Interactions Mean Free Path (Section 4.3.7). The Extrapolation fit has a modest pull of $\sim -0.34\sigma$ on `hNFSI_MFP`, whereas the PISCES pull is significantly larger, at $\sim -0.88\sigma$. This suggests that `hNFSI_MFP` is one of the systematics where PISCES redistributed the pulls that the Extrapolation method attributed to calibration. PISCES also moves the formation zone systematic uncertainty `FormZone` (Section 4.3.7) by -0.22σ .

Resonance systematic uncertainties: MaCCRES, MvCCRES, RESDelta, RESOther, Low Q^2

`MaCCRES`/`MvCCRES` (Section 4.3.3) are uncertainties associated with the axial and vector mass that set the Q^2 dependence for CC resonant (RES) interactions. `RESOtherScaleSyst`, and `LowQ2RESSupp2020` adjust the overall normalization and the low- Q^2 behavior of the RES region.

The joint PISCES fit prefers a negative `MvCCRES` shift of -0.35σ , with a near-zero `MaCCRES`. PISCES also pulls on the `RESDeltaScaleSyst` and `RESOtherScaleSyst` down by $\sim -0.38\sigma$

and $\sim -0.26\sigma$, respectively.

The Extrapolation method, on the other hand, shifts **MaCCRES** moderately at -0.18σ and barely pulls on **MvCCRES**. Also, both **RESOtherScaleSyst** and **LowQ2RESSupp2020** are close to zero.

What should be noted when comparing these pulls is that the RES systematics (among several other systematic uncertainties) present considerable degeneracy when looking at the fractional response from each individual systematic uncertainty. Table A.2 shows a subset of the systematic uncertainties that are degenerate (with a correlation coefficient above 0.3). It shows that parameters such as **MaCCRES**, **MvCCRES**, **RESDeltaScaleSyst**, **RESOtherScaleSyst**, and **LowQ2RESSupp2020** exhibit strong pairwise degeneracy, for instance:

$$\rho(\text{MaCCRES}, \text{MvCCRES}) \simeq 0.85, \quad (8.1)$$

$$\rho(\text{MaCCRES}, \text{RESOtherScaleSyst}) \simeq 0.81, \quad (8.2)$$

$$\rho(\text{MvCCRES}, \text{RESDeltaScaleSyst}) \simeq 0.84. \quad (8.3)$$

Because of this, many different linear combinations of these uncertainties produce similar changes in the prediction. As a result, a direct comparison between the pulls is difficult. However, the pattern in both methods is that they rely on the RES systematic uncertainties to tune the data/MC agreement, but they do so in slightly different ways. How the pulls are distributed among the different RES systematics are not directly comparable due to the fact that they are degenerate. However, a clear difference is seen in the overall shift magnitudes. In the joint ND+FD fit, PISCES generally pulls more strongly on the RES parameters. This suggests that part of the tension that the Extrapolation method absorbed through its large calibration shift is redistributed into the RES shape and normalization uncertainties once the ND constraints are included.

$\nu/\bar{\nu}$ MEC shape.

PISCES uses a model-spread MEC systematics in this analysis (Section 4.3.5). Therefore, the MEC uncertainties between the two methods are not directly comparable. In PISCES, the first component of the MEC systematics `key_MECcomp1` is shifted by $\sim +0.22\sigma$, with a SHAP importance of 4.30%. `key_MECcomp1` acts as a combined knob affecting both neutrino and antineutrino samples. In addition, PISCES pulls `MECEnuShape2020AntiNu` strongly negative ($\sim -0.46\sigma$), while `MECEnuShape2020Nu` remains close to zero. This suggests that, once the first component of the MEC uncertainty is set, the main remaining MEC adjustment in PISCES is in the antineutrino sector.

On the Extrapolation side, there are two different systematics for the MEC shape uncertainties. `MECShape2024Nu` affects neutrino samples, and `MECShape2024Antinu` for antineutrinos. Extrapolation pulls respectively on those systematics by $\sim -0.44\sigma$ for `MECShape2024Nu`, and $\sim -0.14\sigma$ for `MECShape2024Antinu`.

As the MEC uncertainties in the two methods are fundamentally different, their pull values are not comparable. What is important here, though, is that MEC contributes substantially in both analyses. In PISCES, for instance, the MEC components reach an overall SHAP importance of about 7%, meaning that MEC uncertainties are among the systematic uncertainties driving the fit. Both methods, therefore, use MEC systematics to reduce residual data/MC discrepancies, although they incorporate these adjustments differently.

FD muon energy scale

In the Extrapolation fit, the correlated muon energy-scale parameter `CorrMuEScaleSyst2024` is pulled up by about $\sim +0.47\sigma$.

In the joint ND+FD PISCES fit, the correlated muon-scale parameter is constrained by the ND, and the fit only shifts it by about -0.06σ . A large correlated muon-scale is disfavored by the ND samples, so the fit avoids pulling hard on it. Instead, PISCES makes

small adjustments with uncorrelated muon-scale systematics (e.g., `UnCorrMuEScale2024`, `UnCorrMuCatMuESyst2024`, `PileupMuESyst2024`).

Where both methods agree

Despite the differences in some of the magnitudes of the shifts and their sign, both methods converge on the overall picture:

- Both analyses use mixed-sign flux PPFX principal components to adjust the pion/kaon content.
- Each method pulls on the CC/NC RES systematic uncertainties to adjust the RES interactions in the prediction.
- Both methods adjust DIS hadronization at around 0.25σ level.

It should be noted that the pulls for an FD-only fit with PISCES are presented in Appendix A.3. In that PISCES FD fit, the predictions used are the nominal predictions without the Extrapolation step in the NOvA Standard analysis. The PISCES FD-only fit shows a similar behavior to the Extrapolation case, where the calibration uncertainty dominates the fit, with a pull of order -1σ . This further demonstrates that, in the PISCES joint ND+FD fit, the information from the ND helps to constrain the uncertainties and reduces tendency to rely predominantly on the calibration systematic uncertainties to absorb data/MC differences. Instead, the joint fit distributes the pulls to the other physically relevant systematic uncertainties. This highlights how PISCES makes a broader and direct use of the ND information.

Summary

In summary, both frameworks give a consistent overall picture, where they rely on a mix of detector, flux, cross section, and FSI uncertainties to describe the data. The main difference is in how the shifts are distributed between the two methods. As PISCES performs

a simultaneous ND and FD fit, using the full covariance matrix, the pulls spread across the different systematics, with the high-statistics ND samples providing important additional constraints. In contrast, the FD-only Extrapolation fit tends to lean more on the detector-related uncertainties, such as the global calibration or the correlated muon energy scale.

Overall, the PISCES approach provides a more comprehensive use of the ND information, while the Extrapolation method reflects the limitations of relying solely on FD data.

Table 8.1: Systematic pull values comparison between PISCES ND+FD fit and the Extrapolation method, with the uncertainties ranked by their SHAP importance from PISCES (higher first)

PISCES systematic	PISCES pull	SHAP %	Extrapolated name (match)	Extrap pull
key__Calib	−0.3113	9.30	Calibration	−0.9682
key_PPFXcomp1	−0.4832	9.18	ppfx_hadp_pc00	−0.260156
MvCCRES	−0.3478	5.89	MvCCRES	−0.0638
hNFSI_MFP_2024	−0.8799	5.87	hNFSI_MFP_2024	−0.3442
key_Menate	+0.3212	5.24	NeutronMENATEFileSyst	+0.0273
RESDeltaScaleSyst	−0.3832	4.76	RESDeltaScaleSyst	+0.2218
key_MECcomp1	+0.2167	4.30	MECShape2024Nu/AntiNu	−0.4383 / −0.1393
key__CalibShape	+0.2074	3.57	CalibShape	−0.1744
DISvnCC1pi_2020	+0.4613	3.34	DISvnCC1pi_2020	−0.0134
key__Drift	+0.1462	2.57	CalibDrift	−0.0808
DISvnNC1pi_2020	−0.5724	2.29	DISvnNC1pi_2020	—
key_PPFXcomp4	−0.4646	2.03	ppfx_hadp_pc03	−0.02600
ChargedKaonFlux	+0.2351	1.97	ChargedKaonFlux	−0.2706

Continued on next page

PISCES systematic	PISCES pull	SHAP %	Extrapolated name (match)	Extrap pull
g4_pc01	+0.7748	1.96	g4_pc01	+0.0200
RESOtherScaleSyst	−0.2574	1.63	RESOtherScaleSyst	−0.0109
DISNuHadronQ1Syst	−0.3057	1.60	DISNuHadronQ1Syst	−0.2323
key_PPFXcomp5	+0.4465	1.52	ppfx_hadp_pc04	+0.2251
key_PPFXcomp2	−0.2128	1.24	ppfx_hadp_pc01	+0.2088
MECEnuShape2020AntiNu	−0.4603	1.21	MECEnuShape2020AntiNu	−0.0490
MaCCRES	+0.0480	1.10	MaCCRES	−0.1871
g4_pc00	+0.6411	1.73	g4_pc00	+0.0245
g4_pc02	−0.2988	0.38	g4_pc02	+0.0037
ZNormCCQE	+0.0389	0.63	ZNormCCQE	−0.1124
ZExpAxialFF_EV2	−0.0314	0.34	ZExpAxialFF_EV2	−0.0240
ZExpAxialFF_EV3+EV4	+0.0212	0.28	ZExpAxialFF_EV3/EV4	−0.1294 / +0.0375
ZExpAxialFF_EV1	−0.0309	0.19	ZExpAxialFF_EV1	−0.0374
CorrMuEScaleSyst2024	−0.0604	0.58	CorrMuEScaleSyst2024	+0.4710
UnCorrMuEScale2024	−0.2117	0.56	UnCorrND/FD	+0.0248 /
			MuEScaleSyst2024	+0.0356
key_Chernkov	+0.0144	0.16	Chernkov	−0.2521

Table 8.2: PISCES systematics without a clean counterpart in the Extrapolation list.

PISCES systematic	PISCES pull	SHAP %
FormZone2024	−0.2172	1.88
RDecBR1eta	+0.5762	1.52
DIS exclusive modes (many 2020 variants)	various	$\lesssim 2$
EtaNCEL_MaNCEL	−0.3658	0.33
NormAna24NoBubbles_NormAna24POT	−0.1128	0.40

Table 8.3: Extrapolation systematics not present (as-is) in the PISCES pull list.

Extrapolated systematic	Extrap pull
ProngAngleSystNDXZ/NDYZ/FDXZ/FDYZ2021	$\mathcal{O}(10^{-5} \text{ to } 0.13)$
RelativeCalib, Norm2024	−0.0217, −0.1881
genie_small_pc00..07	mixed (± 0.36 max)
MECShape2024Nu/AntiNu (vs. 2020 basis)	−0.4383 / −0.1393
UnCorrND/FD MuEScaleSyst2024	+0.0248 / +0.0356

8.5 Far Detector Best-Fit Spectra: PISCES vs. Extrapolation

For the FD samples, the PISCES ND+FD best-fit spectra are compared to the standard NOvA Extrapolation best-fit prediction. For each channel, corresponding ratio plots (prediction/data) are also shown for both methods. Overall, both frameworks reproduce the FD data well. The remaining differences between the two are mostly small shape variations. The main features of each FD sample are summarized below.

8.5.1 CC ν_μ Disappearance Samples (FHC and RHC, Q1–Q4)

Across the four quartiles, for both FHC and RHC, the two methods give similar descriptions of the FD ν_μ spectra:

- **Overall pattern:** Both PISCES and Extrapolation reproduce the total event counts well, with the overall scale well within statistical and systematic uncertainties.

- **Oscillation dip region:** The energy region that determines $\sin^2 \theta_{23}$ and Δm_{32}^2 is well captured by both methods. Neither fit overshoots nor undershoots the dip. This indicates that both frameworks can match the disappearance pattern.
- **Low- and high-energy tails:** There are small shape differences in the distribution tails. In the ratio plots showing the ratio of the best-fit/data, both methods stay close to unity, overall. However, PISCES and Extrapolation deviate in slightly different ways. These differences are consistent with how each method distributes its systematic shifts:
 - PISCES spreads the residual data/MC differences across the flux, RES, MEC, and FSI systematic uncertainties.
 - The Extrapolation method relies more heavily on detector knobs such as the FD calibration and the muon energy-scale systematics.

Overall, the disappearance channels show that both methods reproduce the FD data well, with only mild differences in the spectral tails.

8.5.2 CC ν_e Appearance Samples (FHC and RHC)

Since the FD ν_e samples have lower statistics, individual bins fluctuate more, but the main features still agree between the two methods:

- **Appearance peak:** Both predictions reproduce the oscillated ν_e appearance peak in the 1-3 GeV region for FHC and RHC.
- **FHC ν_e sample:** PISCES and Extrapolation agree well in normalization. Small differences at low reconstructed energy and in the high-energy tail are visible in the ratio plots, but they remain within uncertainties.

- **RHC ν_e sample:** The statistics are even lower for the RHC samples, so ratio fluctuations are mostly statistical. PISCES shifts the antineutrino MEC and some FSI parameters differently from Extrapolation (as seen in the pull tables), but the predicted spectra remain consistent with each other and with the data.

In summary, both methods describe the appearance samples well, and the small differences are consistent with the different systematic pull patterns discussed earlier.

8.5.3 Low-Energy ν_e Sample (FHC FD)

The dedicated low-energy ν_e selection has very limited statistics, and the background composition is more varied than in the other samples. As a result, the bin-by-bin ratios fluctuate more, however:

- Both PISCES and Extrapolation reproduce the overall shape and normalization of the ν_e low-energy spectrum.
- The ratio plots show larger statistical fluctuations compared to the main ν_e appearance sample, which is expected given the small event counts.
- The two methods differ slightly in the ratio plots.

Overall Comparison and Interpretation

Both PISCES and the Extrapolation method give a good description of the data in both beam modes and for both the ν_μ disappearance and ν_e appearance channels. The small differences that appear tend to show up in regions that are more sensitive to modeling choices—such as the low-energy ν_e bins or the high and low-energy tails of the ν_μ spectra. These differences are consistent with how each method incorporates the Near Detector information and the resulting difference in systematic pulls.




Figure 8.5: (a): ν_μ Q1 samples - Comparison of NOvA Data (black) and the best-fit spectra using PISCES ND+FD fit with post-fit systematic weights (blue histograms)

Overall, the FD spectra support the conclusion that the PISCES ND+FD fit matches both ND and FD data simultaneously by spreading the shifts across a broader set of physically motivated systematics, rather than relying on a single dominant detector shift.

8.6 Contours comparison: PISCES vs. Extrapolation

8.6.1 $\sin^2 \theta_{23} - \Delta m_{32}^2$ Space:

For the $\sin^2 \theta_{23} - \Delta m_{32}^2$ parameter space, PISCES and the Extrapolation method give broadly consistent contours. However, PISCES produces a tighter contour along $\sin^2 \theta_{23}$ for both mass orderings, while the Δm_{32}^2 precision is very similar between the two methods. The main reason is that $\sin^2 \theta_{23}$ is largely controlled by the overall rate and shape of the $\nu_\mu/\bar{\nu}_\mu$ disappearance channels. Coherent changes in normalization or shape induced by flux,




Figure 8.6: (b): ν_μ Q2 samples - Comparison of NOvA Data (black) and the best-fit spectra using PISCES ND+FD fit with post-fit systematic weights (blue histograms).




Figure 8.7: (c): ν_μ Q3 samples - Comparison of NOvA Data (black) and the best-fit spectra using PISCES ND+FD fit with post-fit systematic weights (blue histograms).




Figure 8.8: (d): ν_μ Q4 samples - Comparison of NOvA Data (black) and the best-fit spectra using PISCES ND+FD fit with post-fit systematic weights (blue histograms).

cross section, calibration, or final state interactions (FSI) uncertainties can therefore produce patterns in the data that can mimic a shift in $\sin^2 \theta_{23}$. In the joint ND+FD fit, PISCES profiles these effects using the full covariance matrix and the complete oscillation channel information. Including the ND information directly helps pin down the appropriate shifts for those systematic uncertainties, which reduces the degeneracies between $\sin^2 \theta_{23}$ and the effects of these uncertainties. By contrast, in the Extrapolation approach, the near-to-far ratio cancels the common effects between ND and FD but does not fully capture the cross-sample correlations, so some of the degeneracies still remain. This is why the PISCES ND+FD contours are slightly tighter along $\sin^2 \theta_{23}$ when compared to the Extrapolation method.

8.6.2 δ_{CP} vs $\sin^2 \theta_{23}$ Space:

The contours are comparable overall, with regions where the Extrapolation method looks slightly tighter in δ_{CP} . PISCES performs a more faithful marginalization because it uses the




Figure 8.9: Figure 8.9(e): ν_e samples - Comparison of NOvA Data (black) and the best-fit spectra using PISCES ND+FD fit with post-fit systematic weights (blue histograms).




Figure 8.9(a): ν_μ Q1 samples - Comparison of the ratio of the Extrapolation method best fit with the NOvA Data (pink histograms) and the ratio of the best-fit spectra using PISCES ND+FD fit with NOvA data (blue histograms).




Figure 8.9(b): ν_μ Q2 samples - Comparison of the ratio of the Extrapolation method best fit with the NOvA Data (pink histograms) and the ratio of the best-fit spectra using PISCES ND+FD fit with NOvA data (blue histograms).




Figure 8.9(c): ν_μ Q3 samples - Comparison of the ratio of the Extrapolation method best fit with the NOvA Data (pink histograms) and the ratio of the best-fit spectra using PISCES ND+FD fit with NOvA data (blue histograms).




Figure 8.9(d): ν_μ Q4 samples - Comparison of the ratio of the Extrapolation method best fit with the NOvA Data (pink histograms) and the ratio of the best-fit spectra using PISCES ND+FD fit with NOvA data (blue histograms).









Figure 8.9(e): ν_e samples - Comparison of the ratio of the Extrapolation method best fit with the NOvA Data (pink histograms) and the ratio of the best-fit spectra using PISCES ND+FD fit with NOvA data (blue histograms).

full ND+FD cross-sample covariance and the different oscillation channel splits. This allows systematic uncertainties that are ν_e appearance-sensitive to stay consistent within the extra ND constraints.

In contrast, the near-to-far ratio tends to over-cancel some of the effects that are shared across samples, and it doesn't fully capture the correlations across samples. Because of that, the extrapolated fit can end up producing slightly sharper δ_{CP} contours. In the Extrapolation fit, for instance, the calibration systematics are driving the fit (as can be seen in Table 8.1). In the joint ND+FD fit case, however, PISCES responds more strongly with the systematic uncertainties that are sensitive to $\nu_e/\bar{\nu}_e$ asymmetry, which is consistent with the above interpretation. Figure 8.4 shows the different systematic pulls from the PISCES joint ND+FD fit. Among the top drivers of the fit are the `h_NFSI_MFP` and `FormZone2024`, which are the final state interaction (FSI) mean free path and hadron formation zone uncertainties. These FSI uncertainties change how hadrons propagate and re-interact inside the nucleus, and since charged hadrons do not behave the same for ν and $\bar{\nu}$ events, shifts in these parameters end up affecting the $\nu/\bar{\nu}$ balance in the CC samples. This means that by changing how hadronic energy is produced and transported inside the nucleus, these parameters influence the $\nu_e/\bar{\nu}_e$ asymmetry. Other systematics that have considerable SHAP importance are the resonance-related (RES) and MEC systematic uncertainties. These are cross section-related systematic uncertainties, and because ν and $\bar{\nu}$ respond differently to changes in the axial, MEC, and RES models, they can shift the predicted ν_e and $\bar{\nu}_e$ appearance rates and energy spectra in different ways. When these parameters are shifted in the fit, it means that PISCES is using these $\nu_e/\bar{\nu}_e$ asymmetry-related systematics to absorb data/MC differences, rather than attributing them to the oscillation parameters. This leads to contours that differ from the Extrapolation fit.

The two fitting approaches yield slightly different best-fit points, but they both sit in a very flat region of the $\Delta\chi^2$ surface. In such regions, the differences in how the parameters

Table 8.4: Best-fit values for 10 years of NOvA data using the PISCES ND+FD joint fit, compared to the best-fit values obtained with the Extrapolation method. The 1σ intervals in the first and second columns are derived from 1D profile scans (using $\Delta\chi^2 = 1$) and are not Feldman-Cousins corrected. The final column shows the official Extrapolation 1σ intervals reported in [2] using the Feldman-Cousins approach.

	PISCES Best-Fit Value (1σ)	Extrapolation Best-Fit Value (1σ)	Official Extrapolation Best-Fit Value (1σ , Feldman-Cousins corrected)
$\sin^2 \theta_{23}$	$0.530^{+0.031}_{-0.065}$	$0.55^{+0.022}_{-0.079}$	$0.55^{+0.02}_{-0.06}$
Δm_{32}^2	$\left(2.421^{+0.029}_{-0.031}\right) \times 10^{-3} \text{ eV}^2$	$\left(2.441^{+0.035}_{-0.040}\right) \times 10^{-3} \text{ eV}^2$	$\left(2.441^{+0.035}_{-0.031}\right) \times 10^{-3} \text{ eV}^2$
δ_{CP}	$\left(0.901^{+0.188}_{-0.213}\right) \pi$	$\left(0.87^{+0.192}_{-0.229}\right) \pi$	$\left(0.87^{+0.30}_{-0.90}\right) \pi$

are profiled can shift the best-fit values without indicating real tension between the two methods. The best-fit spectra distribution comparison for PISCES and Extrapolation can be found in Figures 8.5-8.9.

Although the global best fit favors Normal Ordering, the NOvA data itself sits in a region where the NO and IO ellipses overlap significantly, as shown in the NOvA bi-event plot in Fig. 3.1 discussed in Chapter 3. In this region, mass ordering and δ_{CP} effects mimic each other, leading to eightfold degeneracies preventing the NOvA experiment from providing a decisive determination of the mass ordering. This will be further addressed by next-generation experiments, like DUNE, where the larger matter effects lead to a full separation of the bi-event ellipses, resolving the degeneracies [52].




Figure 8.10: 1σ , 2σ and 3σ contours comparison using the Extrapolation method (dashed lines), and the PISCES ND+FD fit (solid lines) considering the NOvA Real Data, for the $\sin^2\theta_{23}$ vs Δm_{32}^2 and δ_{CP} vs $\sin^2\theta_{23}$ parameter space. The first rows consider the normal ordering for Δm_{32}^2 , the second rows are for the inverted ordering.

Chapter 9. Conclusion

PISCES is a flexible technique that allows for complex fits, such as a joint ND+FD, and can leverage the high statistics of the ND to constrain systematic uncertainties.

We have conducted extensive studies to compare PISCES with the NOvA standard extrapolation analysis, using Asimov fake data. With the same data and the same systematic uncertainties, a fit using Far Detector only information, without any Near Detector input, yields the same results between the two methods. Differences arise when the Near Detector information is included in the fit. In this case, the simultaneous joint ND+FD fit from PISCES produces tighter contours than the standard Extrapolation method. A likely reason is that, at each point in parameter space, PISCES performs a nested minimization, in which it tries to pinpoint the optimal systematic scales that best match the data. In this setup, the inner minimization tends to constrain systematic uncertainties more tightly than the Extrapolation method, where the oscillation parameters and systematic uncertainties are profiled at the same time.

In this work, we have also developed a novel method to better understand the internal mechanism of PISCES, enabling the determination of the PISCES systematic pulls during a fit using the Newton-Raphson method. This technique helps identify which systematic uncertainties correspond to the systematic scale obtained in the PISCES inner minimization when PISCES matches the shifted prediction to the data. The analysis highlighted that many of the systematic uncertainties are degenerate in the fit. In the Newton-Raphson method, these degeneracies prevented the method from recovering the correct injected pulls, due to the non-unique solution where several groups of degenerate systematic uncertainties can explain the same data.

We overcame this limitation by performing a Principal Component Analysis (PCA), where we move from the degenerate original systematics basis to a principal-component (PC) basis in which all PCs are orthogonal and therefore non-degenerate. The solution to the optimal pulls in the PCA basis can then be mapped back to the original systematic basis to obtain the final pulls. Validation tests showed that this Newton-Raphson+PCA method is able to successfully recover the appropriate injected pulls. Another advantage is that PCA also allows us to identify the components with the highest contribution to the variance, and drop components with negligible impact, which makes the minimization more stable. It is worth noting that this PCA treatment of systematic uncertainties can be applied to other analyses as well. Since the systematic uncertainties degeneracy is a common limitation, even for non-PISCES analyses, working in the PCA space and then mapping back to the original systematic uncertainties basis helps mitigate potential degeneracies. Thus, this method can be used for systematic-pull quantification in any other types of fits using covariance matrices to encode systematic uncertainties.

Extensive robustness tests were also performed with PISCES to evaluate the potential bias in the fitter in case of mis-modelling in the Monte Carlo. For that, several fake data analyses were conducted. These include tests in which some underlying parameters in the models were modified. Then the fitter’s response was analyzed, and the shifts in contours and best-fit points were compared to the nominal Asimov case. Any considerable deviation from the Asimov case would hint at a potential bias. Some examples of fake data studies involve tests where the underlying MEC interactions or CRPA modelization models have been modified. Other examples include the case where neutral current backgrounds in the ν_e samples were artificially inflated, and the case where the prediction used in the fit was swapped with the prediction without the ND-data-tuned central value weight. All of these tests introduced significant spectral changes from the Asimov data/Nominal predictions. Yet, in every case, the PISCES ND+FD fit showed minimal shifts in the contours and the best-fit

points, meaning that no significant bias was observed in the tested scenarios. Furthermore, when examining which systematic uncertainties were used by PISCES to resolve data/MC tension, in all of the considered fake data, PISCES pulled on the systematics relevant to the underlying physics targeted by the fake data analysis, as expected. This demonstrates that PISCES is a robust method and that PISCES employs the appropriate physics-motivated systematics to address residual data/MC difference.

A comparison using the NOvA full dataset was also carried out between the PISCES joint ND+FD fit and the NOvA standard Extrapolation method. Overall, in the $\sin^2 \theta_{23}$ – Δm_{32}^2 parameter space, both methods yield broadly consistent contours, with PISCES producing tighter contours in $\sin^2 \theta_{23}$ while the Δm_{32}^2 precision remains comparable. Since $\sin^2 \theta_{23}$ is mainly determined by the overall rate and shape of the $\nu_\mu/\bar{\nu}_\mu$ disappearance channels, several systematic uncertainties such as flux, cross section, calibration, and FSI, can easily mimic shifts in this parameter. When the ND information is used directly in the PISCES ND+FD fit, however, the high-statistics ND samples, especially muon neutrinos, help break the degeneracy between the effects of systematic uncertainties and the actual oscillation patterns. Hence, the improved precision for $\sin^2 \theta_{23}$ with the PISCES joint ND+FD fit.

For the δ_{CP} vs $\sin^2 \theta_{23}$ parameter space, the contours obtained with the PISCES joint ND+FD fit and with the NOvA standard Extrapolation method are comparable overall, with some slight differences in δ_{CP} . Analysis of the systematic pulls obtained from the NOvA Extrapolation method indicates that it relies heavily on the calibration systematic uncertainty, with a strong negative shift of -1σ driving the fit. PISCES, on the other hand, with the direct inclusion of the ND samples, prefers a more moderate pull for the calibration uncertainty and instead distributes the shifts across the other physically relevant systematic uncertainties. More specifically, PISCES attributes more importance to systematic uncertainties that affect ν_e appearance samples, such as the various final state interaction (FSI) uncertainties. This allows PISCES to spread the pulls across a broader

set of physics-motivated systematics, instead of over-pulling on a few detector systematic knobs. These distinctions in the actual systematic uncertainties being used to resolve the data/MC tension are consistent with the slight differences in the obtained contours between both methods.

The best-fit points obtained from the PISCES ND+FD fit and the Extrapolation method are different. However, both points lie in a flat region of the $\Delta\chi^2$ surface, which means that these slight differences in best fit points do not lead to real tension between the two methods. This is supported by the best-fit spectra comparison across the different samples, where the PISCES and Extrapolation best-fit spectra are consistent. Small differences that appear are in the regions that are more sensitive to modeling choices, such as the low and high-energy tails of the spectra. These differences result from the distinct choices of systematic pulls used to drive the fit.

On the practical side, PISCES encodes systematic uncertainties in a covariance matrix, rather than profiling over each individual systematic uncertainty. This gives PISCES an advantage compared to other methods, which rely on treating systematic uncertainties as individual nuisance parameters in the fit, where the optimization typically becomes more challenging once the number of systematic uncertainties grows large. The covariance matrix approach in PISCES also ensures that systematic correlations are handled correctly. This flexibility allows PISCES to naturally accommodate not only simultaneous ND+FD fits, but also potential joint-experiment fits such as a NOvA+T2K analysis. Apart from being robust, PISCES is also computationally efficient even with large numbers of samples and systematic uncertainties.

BIBLIOGRAPHY

- [1] L. Aliaga, G. Brunetti, L. Cremonesi, K. Kaur Maan, and A. Radovich. Beam systematic uncertainties for third analyses. Technical Report NOvA Internal Document 17608-v5, NOvA Collaboration, 2018. NOvA internal document.
- [2] NOvA Collaboration. Precision measurement of neutrino oscillation parameters with 10 years of data from the nova experiment. *Phys. Rev. Lett.*, 136:011802, Jan 2026.
- [3] Fermi National Accelerator Laboratory. Fermilab accelerator complex. <https://www.fnal.gov/pub/science/particle-accelerators/accelerator-complex.html>.
- [4] P. Adamson et al. The numi neutrino beam. *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment*, 806:279–306, January 2016.
- [5] D. S. Ayres et al. The nova technical design report. Technical report, Fermilab, 10 2007.
- [6] A. Norman. The nova data acquisition system. In *Proceedings of the 19th IEEE Real-Time Conference (RT2012)*, 2012. FERMILAB-CONF-12-304-CD, for the NOvA Collaboration.
- [7] F. Reines and C. L. Cowan. Detection of the free neutrino. *Phys. Rev.*, 92:830–831, Nov 1953.
- [8] C. L. Cowan, F. Reines, F. B. Harrison, H. W. Kruse, and A. D. McGuire. Detection of the free neutrino: A Confirmation. *Science*, 124:103–104, 1956.
- [9] F. Reines and C. Cowan. The Reines-Cowan experiments: Detecting the Poltergeist. *Los Alamos Sci.*, 25:4–27, 1997.
- [10] G. Danby, J. M. Gaillard, K. Goulianos, L. M. Lederman, N. Mistry, M. Schwartz, and J. Steinberger. Observation of high-energy neutrino reactions and the existence of two kinds of neutrinos. *Physical Review Letters*, 9:36–44, 1962.
- [11] Hai-Jun Yang, Byron P. Roe, and Ji Zhu. Studies of boosted decision trees for miniboone particle identification. *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment*, 555(1):370–385, 2005.
- [12] K. Kodama, N. Ushida, C. Andreopoulos, N. Saoulidou, G. Tzanakos, P. Yager,

- B. Baller, D. Boehnlein, W. Freeman, B. Lundberg, and et al. Observation of tau neutrino interactions. *Physics Letters B*, 504(3):218–224, Apr 2001.
- [13] K. Kodama, N. Ushida, C. Andreopoulos, N. Saoulidou, G. Tzanakos, P. Yager, B. Baller, D. Boehnlein, W. Freeman, B. Lundberg, and et al. Final tau-neutrino results from the donut experiment. *Phys. Rev. D*, 78:052002, Sep 2008.
- [14] R. Davis Jr., D. S. Harmer, and K. C. Hoffman. Search for neutrinos from the sun. *Physical Review Letters*, 20:1205–1209, 1968.
- [15] B. T. Cleveland et al. Measurement of the solar electron neutrino flux with the Homestake chlorine detector. *Astrophys. J.*, 496:505–526, 1998.
- [16] K. S. Hirata et al. Observation of $\bar{\nu}_\mu$ solar neutrinos in the kamiokande-ii detector. *Physical Review Letters*, 63:16–19, 1989.
- [17] P. Anselmann et al. Solar neutrinos observed by gallex at gran sasso. *Physics Letters B*, 285(4):376–389, 1992.
- [18] A. I. Abazov et al. Search for neutrinos from the sun using the reaction $^{71}\text{Ga}(\nu_e, e^-)^{71}\text{Ge}$. *Phys. Rev. Lett.*, 67:3332–3335, Dec 1991.
- [19] Super-Kamiokande Collaboration. Evidence for oscillation of atmospheric neutrinos. *Phys. Rev. Lett.*, 81:1562–1567, Aug 1998.
- [20] John M. LoSecco. The history of “anomalous” atmospheric neutrino events: A first person account. *Physics in Perspective*, 18(2):209–241, July 2016.
- [21] Q. R. Ahmad et al. Direct evidence for neutrino flavor transformation from neutral-current interactions in the sudbury neutrino observatory. *Physical Review Letters*, 89(1):011301, 2002.
- [22] L. Wolfenstein. Neutrino oscillations in matter. *Phys. Rev. D*, 17:2369–2374, May 1978.
- [23] R. P. Feynman and M. Gell-Mann. Theory of the fermi interaction. *Physical Review*, 109(1), 1958.
- [24] E. C. G. Sudarshan and R. E. Marshak. Chirality invariance and the universal fermi interaction. *Physical Review*, 109(5):1860–1862, 1958.
- [25] S. Navas et al. Review of particle physics. *Phys. Rev. D*, 110:030001, Aug 2024.

- [26] Michael E. Peskin and Daniel V. Schroeder. *An Introduction to Quantum Field Theory*. Addison-Wesley, Reading, MA, 1995.
- [27] Steven Weinberg. A model of leptons. *Phys. Rev. Lett.*, 19:1264–1266, Nov 1967.
- [28] Steven Weinberg. Baryon- and lepton-nonconserving processes. *Phys. Rev. Lett.*, 43:1566–1570, Nov 1979.
- [29] Peter Minkowski. $\mu \rightarrow e\gamma$ at a Rate of One Out of 10^9 Muon Decays? *Phys. Lett. B*, 67:421–428, 1977.
- [30] Tsutomu Yanagida. Horizontal Symmetry and Masses of Neutrinos. *Prog. Theor. Phys.*, 64:1103, 1980.
- [31] Murray Gell-Mann, Pierre Ramond, and Richard Slansky. Complex spinors and unified theories. In P. van Nieuwenhuizen and D. Z. Freedman, editors, *Supergravity*, Amsterdam, 1979. North-Holland.
- [32] Rabindra N. Mohapatra and Goran Senjanović. Neutrino mass and spontaneous parity nonconservation. *Phys. Rev. Lett.*, 44:912–915, Apr 1980.
- [33] Carlo Giunti and Chung W. Kim. *Fundamentals of Neutrino Physics and Astrophysics*. Oxford University Press, Oxford, 2007.
- [34] Ziro Maki, Masami Nakagawa, and Shoichi Sakata. Remarks on the unified model of elementary particles. *Prog. Theor. Phys.*, 28:870–880, 1962.
- [35] S.M. Bilenky and S.T. Petcov. Massive neutrinos and neutrino oscillations. *Reviews of Modern Physics*, 59:671–754, 1987.
- [36] Samoil Bilenky. *Introduction to the Physics of Massive and Mixed Neutrinos*. Springer, 2018.
- [37] S. P. Mikheyev and A. Y. Smirnov. Resonance amplification of oscillations in matter and spectroscopy of solar neutrinos. *Soviet Journal of Nuclear Physics*, 42:913–917, 1985. Originally published in *Yad. Fiz.* **42**, 1441 (1985).
- [38] Alexei Yu. Smirnov. The msw effect and solar neutrinos. *Physica Scripta*, T121:57–64, 2005.
- [39] G. P. Zeller. Neutrino cross section measurements. Review of Particle Physics, 2023, 2023. Particle Data Group.

- [40] J. Nieves, I. Ruiz Simo, and M. J. Vicente Vacas. Inclusive charged-current neutrino–nucleus reactions. *Phys. Rev. C*, 83:045501, 2011.
- [41] Dieter Rein and Lalit M. Sehgal. Neutrino excitation of baryon resonances and single pion production. *Annals of Physics*, 133(1):79–153, 1981.
- [42] L. Alvarez-Ruso, Y. Hayato, and J. Nieves. Progress and open questions in the physics of neutrino cross sections at intermediate energies. *New Journal of Physics*, 16(7):075015, 2014.
- [43] A. A. Aguilar-Arevalo et al. Evidence for neutrino oscillations from the observation of $\bar{\nu}_e$ appearance in a $\bar{\nu}_\mu$ beam. *Phys. Rev. D*, 64:112007, 2001.
- [44] A. A. Aguilar-Arevalo et al. Updated miniboone neutrino oscillation results with increased data and new background studies. *Phys. Rev. D*, 103(5):052002, 2021.
- [45] M. Antonello et al. A Proposal for a Three Detector Short-Baseline Neutrino Oscillation Program in the Fermilab Booster Neutrino Beam. *arXiv preprint*, 2015.
- [46] R. Acciarri et al. Design and construction of the microboone detector. *Journal of Instrumentation*, 12(02):P02017, 2017.
- [47] A. A. Aguilar-Arevalo et al. The ICARUS Experiment at the Fermilab Short-Baseline Neutrino Program. *JINST*, 12:P12004, 2017.
- [48] N. Agafonova et al. Discovery of tau neutrino appearance in the cngs neutrino beam with the OPERA experiment. *Physical Review Letters*, 115(12):121802, 2015.
- [49] M. H. Ahn et al. Measurement of neutrino oscillation by the k2k experiment. *Phys. Rev. D*, 74:072003, 2006.
- [50] L.H. Whitehead. Neutrino oscillations with minos and minos+. *Nuclear Physics B*, 908:130–150, July 2016.
- [51] P. et al. Adamson. First measurement of electron neutrino appearance in nova. *Phys. Rev. Lett.*, 116:151806, Apr 2016.
- [52] B. Abi et al. Deep Underground Neutrino Experiment (DUNE), Far Detector Technical Design Report, Volume II: DUNE Physics. *arXiv preprint*, 2020.
- [53] K. Abe et al. Constraint on the Matter–Antimatter Symmetry–Violating Phase in Neutrino Oscillations. *Nature*, 580:339–344, 2020.

- [54] K. Abe et al. Hyper-Kamiokande Design Report. *arXiv preprint*, 2018.
- [55] J. Ashenfelter et al. First Search for Short-Baseline Neutrino Oscillations at HFIR with PROSPECT. *Phys. Rev. Lett.*, 121:251802, 2018.
- [56] H. Almazan Molina et al. Sterile Neutrino Constraints from the STEREO Experiment with 66 Days of Reactor-On Data. *Phys. Rev. Lett.*, 121:161801, 2018.
- [57] I. Alekseev et al. Search for sterile neutrinos at the DANSS experiment. *Phys. Lett. B*, 787:56–63, 2018.
- [58] F. P. An et al. Observation of electron-antineutrino disappearance at Daya Bay. *Phys. Rev. Lett.*, 108:171803, 2012.
- [59] J. K. Ahn et al. Observation of Reactor Electron Antineutrino Disappearance in the RENO Experiment. *Phys. Rev. Lett.*, 108:191802, 2012.
- [60] Y. Abe et al. Indication for the disappearance of reactor electron antineutrinos in the Double Chooz experiment. *Phys. Rev. Lett.*, 108:131801, 2012.
- [61] F. An et al. Neutrino Physics with JUNO. *J. Phys. G*, 43:030401, 2016.
- [62] J. N. Abdurashitov et al. Measurement of the solar neutrino capture rate by the Russian-American gallium solar neutrino experiment during one half of the 22 year cycle of solar activity. *Phys. Rev. C*, 60:055801, 1999.
- [63] W. Hampel et al. GALLEX solar neutrino observations: Results for GALLEX IV. *Phys. Lett. B*, 447:127–133, 1999.
- [64] K. Abe et al. Solar Neutrino Measurements in Super-Kamiokande-IV. *Phys. Rev. D*, 94:052010, 2016.
- [65] M. Agostini et al. Comprehensive measurement of pp-chain solar neutrinos. *Nature*, 562:505–510, 2018.
- [66] Q. R. Ahmad et al. Direct Evidence for Neutrino Flavor Transformation from Neutral-Current Interactions in the Sudbury Neutrino Observatory. *Phys. Rev. Lett.*, 89:011301, 2002.
- [67] S. Mufson, B. Baugh, C. Bower, T. E. Coan, J. Cooper, L. Corwin, J. A. Karty, P. Mason, M. D. Messier, A. Pla-Dalmau, M. Proudfoot, and the NOvA Collaboration. Liquid scintillator production for the nova experiment. *Nuclear Instruments and*

- Methods in Physics Research A*, 789:28–34, 2015.
- [68] Accelerator Division. Concepts rookie book. *FNAL Operations Department*, 2020.
 - [69] L. Aliaga, A. Griffin, L. Fields, S. Avvakumov, M. Bass, J. Bian, R. Bishop, K. Bujak, A. Cabrera, M. Dalhommé, and the MINERvA Collaboration. Neutrino flux predictions for the numi beam. *Physical Review D*, 94(9):092005, 2016.
 - [70] E.C. Dukes, R. Ehrlich, S. Goadhouse, L. Mualem, A. Norman, and R. Tesarek. The nova power distribution system. *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment*, 902:123–137, September 2018.
 - [71] C. Andreopoulos *et al.* The genie neutrino monte carlo generator. *Nucl. Instrum. Meth. A*, 614:87–104, 2010.
 - [72] C. Andreopoulos *et al.* The genie neutrino monte carlo generator: Physics and user manual (v2.10.0). Technical Report FERMILAB-FN-1004-CD, Fermi National Accelerator Laboratory (FNAL), 2015.
 - [73] R. A. Smith and E. J. Moniz. Neutrino reactions on nuclear targets. *Nucl. Phys. B*, 43:605–622, 1972.
 - [74] Cheryl Berger and Lawrence M. Sehgal. Lepton mass effects in single pion production by neutrinos. *Physical Review D*, 76:113004, 2007.
 - [75] A. Bodek and U. K. Yang. Modeling deep inelastic cross sections in the few gev region. *Nucl. Phys. Proc. Suppl.*, 112:70–76, 2002.
 - [76] A. Bodek and U. K. Yang. Nucleon resonance region and effective pdfs. *Nucl. Phys. Proc. Suppl.*, 139:113–118, 2005.
 - [77] M. A. *et al.* Acero. Adjusting neutrino interaction models and evaluating uncertainties using nova near detector data. *The European Physical Journal C*, 80(12), December 2020.
 - [78] S. Agostinelli *et al.* Geant4 — a simulation toolkit. *Nuclear Instruments and Methods in Physics Research Section A*, 506:250–303, 2003.
 - [79] J. Allison *et al.* Recent developments in geant4. *Nuclear Instruments and Methods in Physics Research Section A*, 835:186–225, 2016.

- [80] Geant4 Collaboration. *Geant4 Physics List Guide*, 2020. Guide For Physics Lists, Geant4 documentation.
- [81] Geant4 Collaboration. Qgsp_bert_hp class reference. Geant4 Class Reference. Implementation in Geant4 physics lists.
- [82] M. Zmeřkal, L. Thulliez, P. Tamagno, and E. Dumonteil. Improvement of geant4 neutron-hp package: Unresolved resonance region description with probability tables. *arXiv preprint*, arXiv:2404.16389, 2024.
- [83] Biswaranjan Behera, Gavin Davies, and Fernanda Psihas. Event reconstruction in the nova experiment, 2017.
- [84] Zakaria Elkarghli. Improvement of the nova near detector event reconstruction and primary vertexing through the application of machine learning methods, 2021.
- [85] James C. Bezdek. *Pattern Recognition with Fuzzy Objective Function Algorithms*. Plenum Press, New York, 1981.
- [86] R. Rabelhofer. *Standard and Non-Standard Neutrino Oscillations at the NOvA Experiment*. PhD thesis, Iowa State University, 2021.
- [87] P. Baldi et al. Improved energy reconstruction in nova with regression methods. *Phys. Rev. D*, 99(1):012011, 2019.
- [88] S. Yu et al. Electron neutrino energy reconstruction in nova using cnn particle ids. *arXiv preprint*, 2019.
- [89] A. Shmakov et al. Interpretable joint event–particle reconstruction for neutrino physics at nova with sparse cnns and transformers. *arXiv preprint arXiv:2303.06201*, 2023.
- [90] Mark Sandler, Andrew Howard, Menglong Zhu, Andrey Zhmoginov, and Liang-Chieh Chen. Mobilenetv2: Inverted residuals and linear bottlenecks. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 4510–4520, 2018.
- [91] João A. B. Coelho. Oscprob: A c++ framework for neutrino oscillation probabilities, 2025. Version 2.2.0 (Oct 2025); GitHub repository.
- [92] A. Cervera, A. Donini, M.B. Gavela, J.J. Gomez Cadenas, P. Hernández, O. Mena, and S. Rigolin. Golden measurements at a neutrino factory. *Nuclear Physics B*, 579(1):17–55, 2000.

- [93] Evgeny K Akhmedov, Robert Johansson, Manfred Lindner, Tommy Ohlsson, and Thomas Schwetz. Series expansions for three-flavor neutrino oscillation probabilities in matter. *Journal of High Energy Physics*, 2004(04):078–078, May 2004.
- [94] F. James and M. Roos. Minuit – function minimization and error analysis reference manual. Technical Report D506, CERN, 1994.
- [95] Steve Baker and Robert D. Cousins. Clarification of the use of chi-square and likelihood functions in fits to histograms. *Nuclear Instruments and Methods in Physics Research*, 221(2):437–442, 1984.
- [96] G. Cowan. *Statistical Data Analysis*. Oxford University Press, 1998.
- [97] M. A. et al. Acero. New constraints on oscillation parameters from ν_e appearance and ν_μ disappearance in the nova experiment. *Phys. Rev. D*, 98:032012, Aug 2018.
- [98] M. A. et al. Acero. Improved measurement of neutrino oscillation parameters by the nova experiment. *Phys. Rev. D*, 106:032004, Aug 2022.
- [99] Gary J. Feldman and Robert D. Cousins. Unified approach to the classical statistical analysis of small signals. *Phys. Rev. D*, 57:3873–3889, Apr 1998.
- [100] Alexandre Sousa, Norman Buchanan, Steven Calvez, Pengfei Ding, Derek Doyle, Alexander Himmel, Burt Holzman, James Kowalkowski, Andrew Norman, Thomas Peterka, and the NOvA Collaboration. Implementation of Feldman–Cousins corrections and oscillation calculations in the HPC environment for the NOvA experiment. In *Proceedings of CHEP 2018 / EPJ Web of Conferences 214*, page 05012, 2019.
- [101] N. Nayak. Technote: Pca flux systematics for the 2020 oscillation analysis. Internal technical note NOvA Internal Document 44049-v1, NOvA Collaboration, 2020.
- [102] Andrzej Czarnecki, Xavier Garcia i Tormo, and William J. Marciano. Michel decay spectrum for a muon bound to a nucleus. *Phys. Rev. D*, 90(9):093002, 2014.
- [103] M. Baird, L. Suter, and J. Wolcott. Executive summary for 2020 3-flavor analysis review. Technical Report NOvA Internal Document 44422-v3, NOvA Collaboration, 2020. NOvA internal document.
- [104] Louise Suter. Extrapolation techniques and systematic uncertainties in the nova muon neutrino disappearance analysis, 2015.
- [105] V. Hewes and O. Samoylov. Nova production 5 detector simulation changes and sys-

- tematics. Internal technical note NOvA Internal Document 43935-v4, NOvA Collaboration, 2021.
- [106] Bhubanjyoti Bhattacharya, Richard J. Hill, and Gil Paz. Model independent determination of the axial mass parameter in quasielastic neutrino-nucleon scattering. *Phys. Rev. D*, 84:073006, 2011.
 - [107] Kirk Bays, Michael Dolce, Shane MacDonald, María Martínez Casales, Aaron Mislivec, Greg Pawloski, and Jeremy Wolcott. NOvA 2020 cross-section tune tech note. Internal technical note NOvA Internal Document 43962, NOvA Collaboration, September 2020. NOvA DocDB 43962.
 - [108] K. M. Graczyk, D. Kielczewska, P. Przewlocki, and J. T. Sobczyk. C_5^A axial form factor from bubble chamber experiments. *Phys. Rev. D*, 80:093001, Nov 2009.
 - [109] C. Andreopoulos et al. *GENIE Physics and User Manual*. GENIE Collaboration, <http://genie-mc.org>, 2016.
 - [110] Michael Dolce. *Constraining neutrino oscillation and interaction parameters with the NOvA Near Detector and Far Detector data using Markov Chain Monte Carlo*. Phd thesis, Tufts University, Medford, MA, USA, August 2023. Adviser: Prof. Hugh Gallagher.
 - [111] P. Stowell, L. Pickering, C. Wilkinson, C. Wret, et al. Tuning the GENIE pion production model with MINER ν a data. *Phys. Rev. D*, 100(7):072005, 2019.
 - [112] V Hewes and A. Norrick. Cross section uncertainties for nux analyses. *NOvA Internal Document*, v1(47802), 2023.
 - [113] Melanie Day and Kevin S. McFarland. Differences in quasielastic cross sections of muon and electron neutrinos. *Physical Review D*, 86(5), September 2012.
 - [114] Oleksandr Tomalak, Qing Chen, Richard J. Hill, and Kevin S. McFarland. Qed radiative corrections for accelerator neutrinos. *Nature Communications*, 13(1), September 2022.
 - [115] K. Abe et al. Measurement of neutrino and antineutrino oscillations by the t2k experiment including a new additional sample of ν_e interactions at the far detector. *Phys. Rev. D*, 96(9):092006, 2017.
 - [116] E. S. Pinzon Guerra, C. Wilkinson, S. Bhadra, S. Bolognesi, J. Calcutt, P. de Perio, S. Dolan, T. Feusels, G. A. Fiorentini, Y. Hayato, K. Ieki, K. B. M. Mahn, K. S.

- McFarland, V. Paolone, L. Pickering, R. Tacik, H. A. Tanaka, R. Terri, M. O. Wascko, M. J. Wilking, C. Wret, and M. Yu. Using world charged pion–nucleus scattering data to constrain an intranuclear cascade model. *Phys. Rev. D*, 99:052007, 2019.
- [117] Andrew Sutton. Technical note: Neutron transport simulation. Internal technical note NOvA Internal Document 56095, NOvA Collaboration, August 2022. NOvA DocDB 56095.
 - [118] Michael Baird, Alex Himmel, Evan Niner, and Erika Catano-Mur. Summary of the need for additional timing cuts. Internal technical note NOvA Internal Document 42952, NOvA Collaboration, January 2020. NOvA DocDB 42952.
 - [119] Alexander Booth. Evaluation of systematic uncertainty in normalisation due to near detector data-mc differences in selection efficiency. Internal technical note NOvA Internal Document 43838, NOvA Collaboration, April 2020. NOvA DocDB 43838.
 - [120] Matthew Strait. Muon energy scale systematic. Internal technical note NOvA Internal Document 20816, NOvA Collaboration, June 2023. NOvA DocDB 20816.
 - [121] V Hewes. Parameter inference with systematic covariance and exact statistics. *NOvA DocDB*, v1(45977), 2023.
 - [122] Kenneth Levenberg. A method for the solution of certain non-linear problems in least squares. *Quarterly of Applied Mathematics*, 2(2):164–168, 1944.
 - [123] Donald W. Marquardt. An algorithm for least-squares estimation of nonlinear parameters. *Journal of the Society for Industrial and Applied Mathematics*, 11(2):431–441, 1963.
 - [124] Jorge J. Moré. The levenberg–marquardt algorithm: Implementation and theory. In G. A. Watson, editor, *Numerical Analysis*, volume 630 of *Lecture Notes in Mathematics*, pages 105–116. Springer, 1978.
 - [125] John E. Dennis and Robert B. Schnabel. *Numerical Methods for Unconstrained Optimization and Nonlinear Equations*. SIAM, 1996.
 - [126] Haw-ren Fang and Dianne P. O’Leary. Modified cholesky algorithms: A catalog with new approaches. *Mathematical Programming*, 115(1):319–348, 2008.
 - [127] Philip E. Gill, Walter Murray, and Margaret H. Wright. *Practical Optimization*. Academic Press, London, 1981.

- [128] Gene H. Golub and Charles F. Van Loan. *Matrix Computations*. Johns Hopkins University Press, 4 edition, 2013.
- [129] Adelman Ben-Israel and Thomas N. E. Greville. *Generalized Inverses: Theory and Applications*. Springer, 2 edition, 2003.
- [130] Nicholas J. Higham. *Accuracy and Stability of Numerical Algorithms*. SIAM, 2 edition, 2002.
- [131] Lloyd N. Trefethen and III David Bau. *Numerical Linear Algebra*. SIAM, 1997.
- [132] Jorge Nocedal and Stephen J. Wright. *Numerical Optimization*. Springer, 2 edition, 2006.
- [133] Åke Björck. *Numerical Methods for Least Squares Problems*. SIAM, 1996.
- [134] Per Christian Hansen. *Rank-Deficient and Discrete Ill-Posed Problems: Numerical Aspects of Linear Inversion*. SIAM, 1998.
- [135] Ian T. Jolliffe. *Principal Component Analysis*. Springer, New York, 2 edition, 2002.
- [136] Ian T. Jolliffe and Jorge Cadima. Principal component analysis: a review and recent developments. *Philosophical Transactions of the Royal Society A: Mathematical, Physical and Engineering Sciences*, 374(2065):20150202, 04 2016.
- [137] Trevor Hastie, Robert Tibshirani, and Jerome Friedman. *The Elements of Statistical Learning: Data Mining, Inference, and Prediction*. Springer, New York, 2 edition, 2009.
- [138] Scott M. Lundberg and Su-In Lee. A unified approach to interpreting model predictions. *arXiv preprint arXiv:1705.07874*, 2017.
- [139] Maria Martinez Casales. *Constraining neutrino interaction uncertainties for oscillation measurements in the NOvA experiment using Near Detector data*. PhD thesis, Iowa State U. (main), Iowa State U., 2023.
- [140] J. Ryckebusch, M. Waroquier, K. Heyde, J. Moreau, and D. Ryckbosch. An rpa model for the description of one-nucleon emission processes and application to $^{16}\text{O}(\gamma, n)$ reactions. *Nuclear Physics A*, 476(2):237–271, 1988.
- [141] S. Dolan, A. Nikolakopoulos, O. Page, S. Gardiner, N. Jachowicz, and V. Pandey. Implementation of the crpa model in the genie event generator and analysis of nuclear ef-

fects in low-energy transfer neutrino-nucleus interactions. *Phys. Rev. D*, 106(7):073001, 2022.

APPENDIX A

A.1 Extrapolation pulls list

Table A.1: Systematic pulls from the Extrapolation method.

Systematic	Shift
Calibration	-0.968172
CorrMuEScaleSyst2024	0.471039
MECShape2024Nu	-0.438311
genie_small_pc02	0.361733
hNFSI_MFP_2024	-0.344229
ppfx_hadp_pc00	-0.260156
Light_Level_FD	-0.256985
Cherenkov	-0.252102
DISNuHadronQ1Syst	-0.232317
ppfx_hadp_pc04	0.225075
RESDeltaScaleSyst	0.221763
ppfx_hadp_pc01	0.208817
genie_small_pc07	-0.212209
Norm2024	-0.18808
MaCCRES	-0.187135
CalibShape	-0.174402
genie_small_pc00	0.15794
ppfx_hadp_pc02	0.147184
MECShape2024AntiNu	-0.139283
ZExpAxialFFSyst2020_EV3	-0.129419
Continued on next page	

Systematic	Shift
ZNormCCQE	-0.112426
MaNCRES	0.110464
Light_Level_ND	0.105186
hNFSI_FateFracEV1_2024	-0.105448
NuTauScale	-0.0899584
MECEnuShape2020Nu	-0.0852667
CalibDrift	-0.0807823
genie_small_pc05	-0.0770497
PileupMuESyst2024	0.077515
UnCorrMuCatMuESyst2024	0.0584536
MvCCRES	-0.0637966
genie_small_pc01	-0.0509531
MECEnuShape2020AntiNu	-0.0490298
ZExpAxialFFSyst2020_EV1	-0.0373849
ZExpAxialFFSyst2020_EV4	0.0374755
UnCorrFDMuEScaleSyst2024	0.0356027
MECInitStateNPfrac2020Nu	0.0335253
LowQ2RESSupp2020	-0.0305282
genie_small_pc06	-0.0304006
ppfx_hadp_pc03	-0.0260024
RPAShapeenh2020	-0.0255242
MvNCRES	0.0247567
UnCorrNDMuEScaleSyst2024	0.024826
g4_pc00	0.0245345
Continued on next page	

Systematic	Shift
NeutronMENATEFileSyst	0.0273136
ZExpAxialFFSyst2020_EV2	-0.0239883
RelativeCalib	-0.0216945
g4_pc01	0.0199923
ProngAngleSystNDXZ2021	0.127539
ProngAngleSystNDYZ2021	-0.015222
DISvnCC1pi_2020	-0.0133558
RESOtherScaleSyst	-0.0109147
DISNuBarHadronQ0Syst	0.0118127
RESvpvnRatioNuXSecSyst	0.040394
genie_small_pc03	0.00598255
g4_pc02	0.00372003
RESvpvnRatioNubarXSecSyst	-0.0025559
ProngAngleSystFDYZ2021	5.18588e-05
ProngAngleSystFDXZ2021	3.77647e-05

A.2 Correlations among leading systematic response vectors

Each systematic uncertainty k can be represented by a response vector $\mathbf{v}_k$ whose components are the fractional changes in the predicted event rates for each ND and FD analysis bin when that parameter is shifted by $+1\sigma$.

To quantify how similar two systematics i and j are, we compute a correlation coefficient

based on the normalized dot product of their response vectors:

$$\rho_{ij} = \frac{\mathbf{v}_i \cdot \mathbf{v}_j}{\|\mathbf{v}_i\| \|\mathbf{v}_j\|} = \frac{\sum_b v_{i,b} v_{j,b}}{\sqrt{\sum_b v_{i,b}^2} \sqrt{\sum_b v_{j,b}^2}}. \quad (\text{A.1})$$

By construction, ρ_{ij} lies between -1 and $+1$. Values close to $+1$ indicate that the two systematics produce very similar fractional shifts across the ND and FD samples, meaning that they form an almost degenerate direction in the fit. Values near -1 correspond to almost opposite response shapes, and $|\rho_{ij}| \lesssim 0.3$ is treated as a weak correlation.

Table A.2 shows some examples of the largest correlations among the leading systematic uncertainties that also appear in Table 8.1. It should be noted that many more correlations exist, but they are not all shown here as with roughly 160 systematics, the number of pairwise correlations is very large.

The RES block (`MaCCRES`, `MvCCRES`, `RESDeltaScaleSyst`, `RESOtherScaleSyst`), the kaon-flux and normalization parameters (`Ana2024ChargedKaonFlux`, `NormAna24Mass`, `key_PPFXcomp1/key_PPFXcomp4`), the FSI block (`hNFSI_MFP_2024`, `FormZone2024`, `MaNCRES`, `g4_pc01`), and the MEC/RPA sector (`key_MECcomp1`, `MECEnuShape2020AntiNu`, `RPAShapeenh2020`) all show strong internal correlations. This means that, in practice, the fit can trade pull among these parameters while mostly staying in the same “shape subspace”. Therefore, the detailed distribution of pulls between PISCES and Extrapolation in Table 8.1 should be interpreted in terms of these correlated groups rather than as completely independent knobs.

Table A.2: Subset of pairwise correlation coefficients ρ between the systematic–uncertainty response vectors used in the PISCES ND+FD fit. The pairs shown here are chosen to highlight the main correlated blocks that also appear in the PISCES/Extrapolation pull comparison (Table 8.1): calibration, RES, kaon flux and normalization, FSI, MEC/RPA, and muon energy scales.

Systematic 1	Systematic 2	ρ
key__Calib	key__CalibShape	+0.863
MaCCRES	MvCCRES	+0.845
MaCCRES	RESDeltaScaleSyst	+0.760
MaCCRES	RESOtherScaleSyst	+0.806
MvCCRES	RESDeltaScaleSyst	+0.841
MvCCRES	RESOtherScaleSyst	+0.804
RESDeltaScaleSyst	RESOtherScaleSyst	+0.559
LowQ2RESSupp2020	RESDeltaScaleSyst	−0.596
LowQ2RESSupp2020	RESOtherScaleSyst	−0.447
Ana2024ChargedKaonFlux	MaCCRES	+0.532
Ana2024ChargedKaonFlux	MvCCRES	+0.534
NormAna24Mass	Ana2024ChargedKaonFlux	+0.762
NormAna24Mass	ZNormCCQE	+0.726
Ana2024ChargedKaonFlux	ZNormCCQE	+0.601
NormAna24Mass	RESDeltaScaleSyst	+0.738
NormAna24Mass	key_PPFXcomp1	+0.651
Ana2024ChargedKaonFlux	key_PPFXcomp4	−0.395
key_PPFXcomp1	key_PPFXcomp4	−0.731
key_PPFXcomp2	key_PPFXcomp6	+0.667

Continued on next page

Systematic 1	Systematic 2	ρ
key_PPFXcomp3	key_MECcomp3	+0.987
		(merged)
key_PPFXcomp10	key_PPFXcomp4	+0.600
key_PPFXcomp14	key_PPFXcomp16	+0.538
key_PPFXcomp22	key_PPFXcomp29	-0.405
key_PPFXcomp35	key_PPFXcomp37	+0.324
FormZone2024	hNFSI_MFP_2024	+0.562
MaNCRES	hNFSI_MFP_2024	+0.457
MaNCRES	g4_pc01	-0.742
MaNCRES	g4_pc02	+0.575
g4_pc01	hNFSI_MFP_2024	-0.337
DISvnCC1pi_2020	hNFSI_MFP_2024	+0.306
key_MECcomp1	RPAShapeenh2020	+0.642
key_MECcomp1	key_PPFXcomp3	+0.706
MECEnuShape2020AntiNu	key_MECcomp1	+0.485
MECEnuShape2020AntiNu	key_PPFXcomp2	+0.563
RPAShapeenh2020	ZNormCCQE	+0.726
RPAShapesupp2020	ZNormCCQE	+0.643
RPAShapeenh2020	key_PPFXcomp1	+0.709
MaNCRES	MvNCRES	+0.723
MaNCRES	g4_pc01	-0.742
CorrMuEScaleSyst2024	UnCorrMuEScaleSyst2024	+0.375

A.3 FD-only fit pulls with the PISCES method

Table A.3: Matched systematics ranked by PISCES FD-only SHAP (higher first). Extrapolation FD-only shifts shown in typewriter.

PISCES systematic	PISCES pull	SHAP %	Extrapolated name (match)	Extrap pull
key__Calib	−1.0756	12.14	Calibration	−0.968172
ZExpAxialFFSyst2020_EV3	−0.3445	10.51	ZExpAxialFFSyst2020_EV3	−0.129419
MvCCRES	−0.4679	7.96	MvCCRES	−0.063797
ZNormCCQE	−0.2849	5.09	ZNormCCQE	−0.112426
MECEnuShape2020Nu	−1.2459	3.00	MECEnuShape2020Nu	−0.085267
CorrMuEScaleSyst2024	+0.2850	2.60	CorrMuEScaleSyst2024	+0.471039
NuTauScale	−0.0460	2.59	NuTauScale	−0.089958
key_PPFXcomp1	−0.2434	2.42	ppfx_hadp_pc00	−0.260156
MaNCRES	+0.5678	2.24	MaNCRES	+0.110464
MaCCRES	−0.0857	1.97	MaCCRES	−0.187135
key__CalibShape	−0.5567	1.68	CalibShape	−0.174402
key_PPFXcomp2	+0.5795	1.62	ppfx_hadp_pc01	+0.208817
MECEnuShape2020AntiNu	+0.6371	1.30	MECEnuShape2020AntiNu	−0.049030
RESOtherScaleSyst	−0.2242	1.28	RESOtherScaleSyst	−0.010915
key_Menate	+0.2919	1.28	NeutronMENATEFileSyst	+0.027314
ZExpAxialFFSyst_EV1	−0.1780	1.27	ZExpAxialFFSyst_EV1	−0.037385
key__LightLevel	−0.4850	1.25	Light_Level_FD	−0.256985
key_Chernikov	−0.4001	1.21	Chernikov	−0.252102

Continued on next page

PISCES systematic	PISCES pull	SHAP %	Extrapolated name (match)	Extrap pull
key_PPFXcomp4	-0.6047	1.14	ppfx_hadp_pc03	-0.026002

A.4 Fake Data Studies

A.4.1 Minerva MEC Fake Data Study

Comparison of MINERvA MEC with the Asimov contours for fake data generated at the 4 other reference oscillation points

Figures A.1 - A.4 show the sensitivity contours comparison of the MINERvA MEC fake data with the Asimov in the case where the fake data were generated at the Asimov B, C, D, and E oscillation points. In all considered cases, the shifts in contours and best fit points were minimal, so no significant bias was found.

A.4.2 CRPA Fake Data Study

Comparison of CRPA with the Asimov contours for fake data generated at the 4 other reference oscillation points

Figures A.5 - A.8 show the sensitivity contours comparison of the CRPA fake data with the Asimov in the case where the fake data were generated at the Asimov B, C, D, and E oscillation points. In all 4 different reference points, the shifts in contours and best fit points were minimal, so no significant bias was found either.




Figure A.1: 1σ, 2σ and 3σ contours comparison considering the MINERvA MEC weighted (dashed lines), and the nominal Asimov (solid lines) fake data, for the $\sin^2\theta_{23}$ vs Δm_{32}^2 and δ_{CP} vs $\sin^2\theta_{23}$ parameter space, for fake data generated at the Asimov B point. Asimov B is the point closest to T2K's 2020 Best-Fit point. The first rows consider the normal hierarchy for Δm_{32}^2 , the second rows are for the inverted hierarchy. The shifts in contours and best fit points are minimal, showing that no significant bias was found and that the fitter is robust against the considered change.




Figure A.2: 1σ , 2σ and 3σ contours comparison considering the MINERvA MEC weighted (dashed lines), and the nominal Asimov (solid lines) fake data, for the $\sin^2 \theta_{23}$ vs Δm_{32}^2 and δ_{CP} vs $\sin^2 \theta_{23}$ parameter space, for fake data generated at the Asimov C point. Asimov C probes for inverted hierarchy and maximum CP. The first rows consider the normal hierarchy for Δm_{32}^2 , the second rows are for the inverted hierarchy. The shifts in contours and best fit points are minimal, showing that no significant bias was found and that the fitter is robust against the considered change.




Figure A.3: 1σ , 2σ and 3σ contours comparison considering the MINERvA MEC weighted (dashed lines), and the nominal Asimov (solid lines) fake data, for the $\sin^2 \theta_{23}$ vs Δm_{32}^2 and δ_{CP} vs $\sin^2 \theta_{23}$ parameter space, for fake data generated at the Asimov D point. Asimov D probes for normal hierarchy, lower octant, and maximum CP. The first rows consider the normal hierarchy for Δm_{32}^2 , the second rows are for the inverted hierarchy. The shifts in contours and best fit points are minimal, showing that no significant bias was found and that the fitter is robust against the considered change.




Figure A.4: 1σ , 2σ and 3σ contours comparison considering the MINERvA MEC weighted (dashed lines), and the nominal Asimov (solid lines) fake data, for the $\sin^2 \theta_{23}$ vs Δm_{32}^2 and δ_{CP} vs $\sin^2 \theta_{23}$ parameter space, for fake data generated at the Asimov E point. Asimov E was chosen as it is located far from NOvA's preferred oscillation point. Asimov E probes for inverted hierarchy, lower octant, and $\delta_{CP} = \pi/2$. The first rows consider the normal hierarchy for Δm_{32}^2 , the second rows are for the inverted hierarchy. The shifts in contours and best fit points are minimal, showing that no significant bias was found and that the fitter is robust against the considered change.




Figure A.5: 1σ , 2σ and 3σ contours comparison considering the CRPA weighted (dashed lines), and the nominal Asimov (solid lines) fake data, for the $\sin^2 \theta_{23}$ vs Δm_{32}^2 and δ_{CP} vs $\sin^2 \theta_{23}$ parameter space, for fake data generated at the Asimov B point. Asimov B is the point closest to T2K's 2020 Best-Fit point. The first rows consider the normal hierarchy for Δm_{32}^2 , the second rows are for the inverted hierarchy. The shifts in contours and best fit points are minimal, showing that no significant bias was found and that the fitter is robust against the considered change.




Figure A.6: 1σ, 2σ and 3σ contours comparison considering the CRPA weighted (dashed lines), and the nominal Asimov (solid lines) fake data, for the $\sin^2 \theta_{23}$ vs Δm_{32}^2 and δ_{CP} vs $\sin^2 \theta_{23}$ parameter space, for fake data generated at the Asimov C point. Asimov C probes for inverted hierarchy and maximum CP. The first rows consider the normal hierarchy for Δm_{32}^2 , the second rows are for the inverted hierarchy. The shifts in contours and best fit points are minimal, showing that no significant bias was found and that the fitter is robust against the considered change.




Figure A.7: 1σ , 2σ and 3σ contours comparison considering the CRPA weighted (dashed lines), and the nominal Asimov (solid lines) fake data, for the $\sin^2 \theta_{23}$ vs Δm_{32}^2 and δ_{CP} vs $\sin^2 \theta_{23}$ parameter space, for fake data generated at the Asimov D point. Asimov D probes for normal hierarchy, lower octant, and maximum CP. The first rows consider the normal hierarchy for Δm_{32}^2 , the second rows are for the inverted hierarchy. The shifts in contours and best fit points are minimal, showing that no significant bias was found and that the fitter is robust against the considered change.




Figure A.8: 1σ , 2σ and 3σ contours comparison considering the CRPA weighted (dashed lines), and the nominal Asimov (solid lines) fake data, for the $\sin^2 \theta_{23}$ vs Δm_{32}^2 and δ_{CP} vs $\sin^2 \theta_{23}$ parameter space, for fake data generated at the Asimov E point. Asimov E was chosen as it is located far from NOvA's preferred oscillation point. Asimov E probes for inverted hierarchy, lower octant, and $\delta_{CP} = \pi/2$. The first rows consider the normal hierarchy for Δm_{32}^2 , the second rows are for the inverted hierarchy. The shifts in contours and best fit points are minimal, showing that no significant bias was found and that the fitter is robust against the considered change.

A.4.3 Increased NC Background for ν_e samples Fake Data Study

Comparison of NC ν_e background doubled with the Asimov contours for fake data generated at the 4 other reference oscillation points

Figures A.9–A.12 present sensitivity contour comparisons between the NC ν_e background fake data and the nominal Asimov, with fake data generated at the Asimov B, C, D, and E oscillation points. Across all cases, shifts in the contours and best-fit points were minimal, indicating no significant bias.

A.4.4 Difference in cross-section central value weight Fake Data Study

Comparison of contours obtained using predictions with/without cross-section CV weight for fake data generated without the weight at the 4 other reference oscillation points

Figures A.13 - A.16 show the comparison of the sensitivity obtained by using the two different predictions for fake data generated at the Asimov B, C, D, and E oscillation points. In all considered cases, the shifts in contours and best fit points were minimal, so no significant bias was found.




Figure A.9: 1σ, 2σ and 3σ contours comparison considering the NC ν_e background doubled (dashed lines), and the nominal Asimov (solid lines) fake data, for the $\sin^2 \theta_{23}$ vs Δm_{32}^2 and δ_{CP} vs $\sin^2 \theta_{23}$ parameter space, for fake data generated at the Asimov B point. Asimov B is the point closest to T2K's 2020 Best-Fit point. The first rows consider the normal hierarchy for Δm_{32}^2 , the second rows are for the inverted hierarchy. The shifts in contours and best fit points are minimal, showing that no significant bias was found and that the fitter is robust against the considered change.




Figure A.10: 1σ , 2σ and 3σ contours comparison considering the NC ν_e background weighted (dashed lines), and the nominal Asimov (solid lines) fake data, for the $\sin^2\theta_{23}$ vs Δm_{32}^2 and δ_{CP} vs $\sin^2\theta_{23}$ parameter space, for fake data generated at the Asimov C point. Asimov C probes for inverted hierarchy and maximum CP. The first rows consider the normal hierarchy for Δm_{32}^2 , the second rows are for the inverted hierarchy. The shifts in contours and best fit points are minimal, showing that no significant bias was found and that the fitter is robust against the considered change.




Figure A.11: 1σ , 2σ and 3σ contours comparison considering the NC ν_e background doubled (dashed lines), and the nominal Asimov (solid lines) fake data, for the $\sin^2\theta_{23}$ vs Δm_{32}^2 and δ_{CP} vs $\sin^2\theta_{23}$ parameter space, for fake data generated at the Asimov D point. Asimov D probes for normal hierarchy, lower octant, and maximum CP. The first rows consider the normal hierarchy for Δm_{32}^2 , the second rows are for the inverted hierarchy. The shifts in contours and best fit points are minimal, showing that no significant bias was found and that the fitter is robust against the considered change.




Figure A.12: 1σ , 2σ and 3σ contours comparison considering the NC ν_e background doubled (dashed lines), and the nominal Asimov (solid lines) fake data, for the $\sin^2\theta_{23}$ vs Δm_{32}^2 and δ_{CP} vs $\sin^2\theta_{23}$ parameter space, for fake data generated at the Asimov E point. Asimov E was chosen as it is located far from NOvA's preferred oscillation point. Asimov E probes for inverted hierarchy, lower octant, and $\delta_{CP} = \pi/2$. The first rows consider the normal hierarchy for Δm_{32}^2 , the second rows are for the inverted hierarchy. The shifts in contours and best fit points are minimal, showing that no significant bias was found and that the fitter is robust against the considered change.




Figure A.13: 1σ , 2σ and 3σ contours comparison considering predictions with (dashed lines)/without (solid lines) cross-section CV weight fitting for a fake data without the weight. The comparisons are for the $\sin^2\theta_{23}$ vs Δm^2_{32} and δ_{CP} vs $\sin^2\theta_{23}$ parameter space, where the fake data are generated at the Asimov B point. Asimov B is the point closest to T2K's 2020 Best-Fit point. The first rows consider the normal hierarchy for Δm^2_{32} , the second rows are for the inverted hierarchy. The shifts in contours and best fit points are minimal, showing that no significant bias was found and that the fitter is robust against the considered change.




Figure A.14: 1σ , 2σ and 3σ contours comparison considering predictions with (dashed lines)/without (solid lines) cross-section CV weight fitting for a fake data without the weight. The comparisons are for the $\sin^2\theta_{23}$ vs Δm_{32}^2 and δ_{CP} vs $\sin^2\theta_{23}$ parameter space, where the fake data are generated at the Asimov C point. Asimov C probes for inverted hierarchy and maximum CP. The first rows consider the normal hierarchy for Δm_{32}^2 , the second rows are for the inverted hierarchy. The shifts in contours and best fit points are minimal, showing that no significant bias was found and that the fitter is robust against the considered change.




Figure A.15: 1σ , 2σ and 3σ contours comparison considering predictions with (dashed lines)/without (solid lines) cross-section CV weight fitting for a fake data without the weight. The comparisons are for the $\sin^2\theta_{23}$ vs Δm_{32}^2 and δ_{CP} vs $\sin^2\theta_{23}$ parameter space, where the fake data are generated at the Asimov D point. Asimov D probes for normal hierarchy, lower octant, and maximum CP. The first rows consider the normal hierarchy for Δm_{32}^2 , the second rows are for the inverted hierarchy. The shifts in contours and best fit points are minimal, showing that no significant bias was found and that the fitter is robust against the considered change.




Figure A.16: 1σ , 2σ and 3σ contours comparison considering predictions with (dashed lines)/without (solid lines) cross-section CV weight fitting for a fake data without the weight. The comparisons are for the $\sin^2\theta_{23}$ vs Δm_{32}^2 and δ_{CP} vs $\sin^2\theta_{23}$ parameter space, where the fake data are generated at the Asimov E point. Asimov E was chosen as it is located far from NOvA's preferred oscillation point. Asimov E probes for inverted hierarchy, lower octant, and $\delta_{CP} = \pi/2$. The first rows consider the normal hierarchy for Δm_{32}^2 , the second rows are for the inverted hierarchy. The shifts in contours and best fit points are reasonable, showing that no significant bias was found and that the fitter is robust against the considered change.