UC Irvine

UC Irvine Electronic Theses and Dissertations

Title

Precision Methods for Long-Baseline Neutrino Experiments: Flux Constraints, Statistical Inference, and Liquid-Argon Detector Monitoring

Permalink

https://escholarship.org/uc/item/1qq6634j

ISBN

9798177346625

Author

Xiao, Yiwen

Publication Date

2026-11-23

Peer reviewed|Thesis/dissertation

eScholarship.org Powered by the California Digital Library
University of California

UNIVERSITY OF CALIFORNIA,
IRVINE

Precision Methods for Long-Baseline Neutrino Experiments: Flux Constraints, Statistical Inference, and Liquid-Argon Detector Monitoring

DISSERTATION

submitted in partial satisfaction of the requirements
for the degree of

DOCTOR OF PHILOSOPHY

in Physics

by

Yiwen Xiao

Dissertation Committee:
Professor Jianming Bian, Chair
Professor Henry Sobel
Professor Juan Pedro Ochoa-Ricoux

2026

© 2026 Yiwen Xiao

Contents

Acknowledgments xviii

Vita xxi

Abstract of the Dissertation xxii

1 Introduction 1

1.1 Neutrinos and the Standard Model 2

1.2 Experimental Evidence for Neutrino Oscillations 3

1.3 Neutrino Sources and Detection 4

1.4 Open Questions in Neutrino Physics 5

1.5 Long-Baseline Accelerator Experiments 6

1.6 Precision in Long-Baseline Measurements 7

1.7 NOvA and DUNE 9

1.8 Scope and Contributions of This Dissertation 9

2 Neutrino Oscillation Physics 11

2.1 Flavor and Mass Eigenstates 11

2.2 PMNS Matrix 12

2.3 Vacuum Oscillations 14

2.3.1 Two-Flavor Limit 16

2.4 Three-Flavor Oscillation Probabilities 19

2.5 Matter Effects . . . 20
2.6 From Oscillation Probabilities to Parameter Inference . . . 23

3 The NOvA Experiment 24
3.1 The NuMI Neutrino Beam . . . 25
3.1.1 Proton Beam, Target, and Secondary Hadrons . . . 25
3.1.2 Magnetic Focusing and FHC/RHC Operation . . . 25
3.1.3 Off-Axis Concept and the NOvA Flux . . . 26
3.2 The NOvA Detectors . . . 28
3.2.1 Detector Technology . . . 28
3.2.2 Near Detector . . . 29
3.2.3 Far Detector . . . 30
3.3 Data Acquisition and Detector Response . . . 30
3.3.1 Signal Readout . . . 30
3.3.2 Light Collection and Calibration . . . 31
3.4 Simulation . . . 32
3.4.1 Beam Simulation and PPFX . . . 32
3.4.2 Neutrino Interactions and GENIE . . . 33
3.4.3 Detector Simulation . . . 34
3.5 Event Reconstruction . . . 34
3.5.1 Hits and Slices . . . 34
3.5.2 Vertex and Prong Reconstruction . . . 36
3.5.3 Tracks and Showers . . . 36
3.6 Event Classification . . . 37
3.7 NOvA Oscillation Program and Flux Systematics . . . 37

4 Neutrino–Electron Elastic Scattering as a Neutrino Flux Standard Candle 39
4.1 Tree-Level Cross Section . . . 40

4.2 Two-Body Kinematics . . . 41
4.3 Relation Between the Event Spectrum and the Flux . . . 42
4.4 Radiative Corrections . . . 43

5 NOvA Neutrino–Electron Elastic-Scattering Analysis and Flux Constraint 47
5.1 Signal and Background Simulation . . . 47
5.2 Event Selection . . . 48
5.2.1 Preselection . . . 49
5.2.2 Fiducial Volume and Containment . . . 49
5.2.3 Single-Particle Requirements . . . 50
5.2.4 Electron-Energy Requirement . . . 50
5.3 Training of the CNN Classifiers . . . 50
5.3.1 Training Samples and Output Categories . . . 51
5.3.2 Training Preprocessing . . . 51
5.3.3 Energy Reweighting . . . 54
5.3.4 Training Configuration and Performance . . . 54
5.4 Final Particle Identification and Kinematic Selection . . . 56
5.5 Background Estimation and Subtraction . . . 58
5.5.1 High-Statistics ν_e CC MEC Sample . . . 59
5.5.2 Sideband Normalization . . . 61
5.5.3 Kinematic Comparison of Signal and Sideband Regions . . . 62
5.5.4 Background-Subtracted Data Sample . . . 63
5.6 Systematic Uncertainties . . . 63
5.6.1 Signal-Selection Efficiency and MR Brem Control Sample . . . 64
5.6.2 GENIE Interaction-Model Uncertainty . . . 66
5.6.3 Beam-Optics Uncertainty . . . 67
5.6.4 Detector-Response Uncertainty . . . 67
5.6.5 Total Covariance Matrix . . . 68

5.7 Bayesian Flux-Constraint Method . . . 68
5.8 Validation of the Flux-Constraint Procedure . . . 71
5.8.1 Asimov Closure Test . . . 72
5.8.2 PPFX Ensemble Stability . . . 73
5.9 Flux-Constraint Results . . . 74
5.9.1 Predicted Neutrino–Electron Event Rate . . . 74
5.9.2 Integrated ν_μ Flux . . . 75
5.9.3 Energy-Dependent Constraint . . . 76
5.10 Interpretation and Propagation of the Constraint . . . 78
5.10.1 Propagation to Other Beam Flavors . . . 78
5.10.2 Downstream Check with a ν_μ CC Selection . . . 78
5.11 Summary . . . 80

6 The NOvA Neutrino Oscillation Analysis 82

6.1 Oscillation-Analysis Overview . . . 82
6.2 Event Samples and Selections . . . 83
6.2.1 Muon-Neutrino Disappearance Samples . . . 84
6.2.2 Electron-Neutrino Appearance Samples . . . 84
6.3 Near-to-Far Extrapolation and the Near Detector Constraint . . . 85
6.3.1 Migration Matrices and True-Energy Predictions . . . 86
6.3.2 Appearance Signal and Background Components . . . 87
6.4 Oscillation Likelihood and Systematic Parameters . . . 88
6.5 Confidence Intervals and the Need for Feldman–Cousins . . . 89
6.6 Connection to the Neutrino–Electron Flux Constraint . . . 90

7 Feldman–Cousins Confidence Intervals in Neutrino Oscillation Measurements 91

7.1 Likelihood-Ratio Inference . . . 91

7.2 The Feldman–Cousins Construction . . . 92
7.3 Nuisance Parameters and the Profile Construction . . . 94
7.4 Application in NOvA Oscillation Measurements . . . 96
7.5 Computational Cost of Grid-Based FC . . . 96
7.6 From Feldman–Cousins to GP-FC . . . 97
7.7 Summary . . . 97

8 Gaussian-Process-Accelerated Feldman–Cousins 99
8.1 Gaussian-process model . . . 99
8.1.1 Training Points and GP Coordinates . . . 99
8.1.2 Kernel and Observation Uncertainty . . . 100
8.1.3 Hyperparameter Optimization and Prediction . . . 101
8.2 Adaptive Selection of FC Calculations . . . 102
8.3 Validation Configuration . . . 103
8.4 Parallel Implementation . . . 103
8.4.1 Runtime Control for Real Fits . . . 104
8.5 Stored-Pseudo-Experiment Tests . . . 105
8.6 Validation Metrics . . . 109
8.7 Tests with Real Likelihood Fits . . . 110
8.8 Summary . . . 112

9 The DUNE Experiment 113
9.1 Experimental Design . . . 114
9.1.1 Neutrino Beam and Near Detector . . . 114
9.1.2 Far Detector . . . 114
9.2 Liquid-Argon Time Projection Chambers . . . 115
9.3 The ProtoDUNE Program . . . 115
9.3.1 ProtoDUNE-SP . . . 116

9.3.2 ProtoDUNE-HD . . . 116
9.4 Liquid-Argon Purity Monitoring . . . 117

10 Liquid-Argon Purity Monitoring in ProtoDUNE 118
10.1 UV Purity-Monitor Measurement Principle . . . 118
10.2 ProtoDUNE-SP Purity Monitors . . . 119
10.2.1 Detector Design and Operation . . . 119
10.2.2 Waveform Analysis . . . 121
10.2.3 Charge Calibration and Corrections . . . 123
10.3 Electron-Lifetime Measurements . . . 124
10.3.1 Measurement Uncertainties . . . 124
10.3.2 ProtoDUNE-SP Performance . . . 125
10.3.3 Comparison with TPC Measurements . . . 126
10.4 Long-Drift UV Purity-Monitor Development . . . 127
10.5 Electron-Lifetime Calibration . . . 129
10.6 ProtoDUNE-HD Automated Monitoring . . . 131
10.7 Summary . . . 132

11 Conclusions and Outlook 134
11.1 NOvA Neutrino–Electron Scattering and Flux Constraint . . . 134
11.2 Parallel Gaussian-Process-Accelerated Feldman–Cousins . . . 135
11.3 ProtoDUNE Purity-Monitor Instrumentation and Operation . . . 136
11.4 Outlook . . . 137

List of Figures

2.1 Energy dependence of the ν_μ survival probability and the corresponding change in the Far Detector energy spectrum. The disappearance feature provides sensitivity to the atmospheric mass splitting and θ_{23}. 18

3.1 Schematic of the NuMI neutrino beamline, showing the graphite target, magnetic focusing horns, decay region, and downstream absorber [14, 15]. 26

3.2 Simulated NuMI neutrino-energy spectra at several off-axis angles, illustrating the narrower energy distribution obtained away from the nominal beam axis [15]. 28

3.3 Schematic of the NOvA detector geometry. Alternating planes of horizontal and vertical cells provide the two views used to reconstruct the three-dimensional event topology [15]. 29

3.4 Overview of the NOvA reconstruction chain, from calibrated detector hits to reconstructed slices, vertices, prongs, and particle objects. [15] 35

4.1 Tree-level Feynman diagrams for neutrino–electron elastic scattering through neutral-current (Z-boson) and charged-current (W-boson) exchange. 40

4.2 True $E_e\theta_e^2$ distribution for simulated neutrino–electron elastic-scattering events. 43

4.3 Radiative-correction weights as functions of E_ν and y_{EM} for the four beam components contributing to the signal. 46

5.1 Architecture of the NuoneID convolutional neural network, modified based on MobileNetV2. 52

5.2 Representative detector pixel maps used to motivate the analysis-specific CNN classification tasks. 53
5.3 Energy distributions of the training samples. Left: True electron-energy distribution before and after reweighting. Right: Shower-energy distributions of signal and background after reweighting. 54
5.4 Training loss and accuracy for NuoneID and Epi0ID. The spikes indicate the beginning of successive training rounds. 55
5.5 Evolution of the NuoneID distribution over successive training rounds. The curves show the classifier response at different stages of training. 56
5.6 Simulated PID distributions with all other selection criteria applied. Panel (a) shows the NuoneID (ν–e PID) distribution, for which the neutrino–electron signal peaks near 0.9 while the background is concentrated predominantly at lower PID values. Panel (b) shows the Epi0ID (e/π^0 PID) distribution, for which the signal is concentrated at high PID values. Background components are stacked as ν_e CC, ν_μ CC, and NC + Other, while the simulated neutrino–electron signal is overlaid. Here, "Other" denotes reconstructed slices without a matched true neutrino interaction in the simulation, including slices arising from mis-grouping of energy deposits during event reconstruction. The black vertical arrows indicate the selection thresholds. The vertical range in panel (a) is restricted to emphasize the high-PID region; the background at low PID extends beyond the displayed range. 57

5.7 Reconstructed $E_e\theta_e^2$ distribution in Near Detector data and simulation with all other selection requirements applied (top), and the ratio of data to simulation (bottom). Data error bars represent statistical uncertainties only. The neutrino–electron signal is concentrated at low $E_e\theta_e^2$, while the background is relatively broad. Backgrounds are separated into ν_e CC, ν_μ CC, and NC plus other interactions, with ν_e CC dominant after selection. The signal region is $0 < E_e\theta_e^2 < 0.005$ GeV rad^2 and the sideband is $0.005 < E_e\theta_e^2 < 0.04$ GeV rad^2. The background prediction is corrected using the sideband procedure described in Sec. 5.5.2. 58

5.8 Primary reconstructed shower-energy distribution in the signal region, $0 < E_e\theta_e^2 < 0.005$ GeV rad^2, for data and simulation. Data error bars represent statistical uncertainties. The background prediction is corrected using the sideband procedure described in Sec. 5.5.2; the NC and Other backgrounds are combined in the figure. The thin red curves show the predictions from the 100 PPFX universes. 60

5.9 True Q^2 distributions of simulated backgrounds in the $E_e\theta_e^2$ signal and sideband regions. 62

5.10 Reweighted rock-MR-Brem control-sample comparisons underlying the signal-selection efficiency uncertainty: ν–e PID (left), e/π^0 PID (center), and $E_e\theta_e^2$ (right). The MC-derived bin-by-bin weights are applied to both MR Brem data and simulation before the efficiencies are compared. 65

5.11 Comparison of calibration variations in the reconstructed $E_e\theta_e^2$ distributions of MR Brem and overlaid neutrino–electron signal samples. The nominal and calibration-shifted predictions are shown for each sample. 66

5.12 GENIE multiverse convergence study used to choose the final 200-universe ensemble. The left panel shows the background RMS and the right panel the corresponding uncertainty on the extracted neutrino–electron rate as the number of universes is increased. Both approach a plateau by approximately 200 universes. 67

5.13 Absolute covariance matrix of the background-subtracted reconstructed electron-energy spectrum used in the flux constraint, retaining the event-yield scale. It contains the data statistical covariance and propagated non-PPFX systematic covariance contributions. This dimensional covariance is the matrix entering the likelihood in Eq. 5.28. 69

5.14 Relative covariance matrix for the background-subtracted reconstructed electron-energy spectrum used in the flux constraint. Each element is defined as $(\Sigma_N)_{ij}/(N_i N_j)$, where $(\Sigma_N)_{ij}$ is the covariance between bins i and j, and N_i and N_j are the corresponding nominal event yields. The matrix includes the statistical uncertainty of the data and the propagated non-PPFX systematic uncertainties described in Sec. 5.6. 70

5.15 Asimov closure test of the Bayesian flux constraint, comparing the constrained and unconstrained distributions of the predicted neutrino–electron event rate and integrated ν_μ flux. 72

5.16 Dependence of the unconstrained and constrained fractional uncertainties on the number of PPFX universes included in the flux constraint. 73

5.17 Probability distributions of the predicted total number of neutrino–electron scattering events across the 100 PPFX universes before (black solid) and after (red dashed) applying the data constraint. The predicted yields correspond to 1.38×10^{21} POT. The fractional spread (RMS/Mean) is reduced from 7.7% to 3.9%. 74

5.18 Probability distributions of the ν_μ flux integrated from 0 to 120 GeV across the 100 PPFX universes before (black solid) and after (red dashed) applying the data constraint. The integrated flux is expressed per 10^{10} POT. The fractional spread (RMS/Mean) is reduced from 8.1% to 4.0%. 75

5.19 Predicted ν_μ flux as a function of neutrino energy before (black solid) and after (red dashed) applying the neutrino–electron constraint (top), and the ratio of the constrained to unconstrained flux predictions (bottom). 76

5.20 Fractional uncertainty (RMS/Mean) of the predicted ν_μ flux before (black solid) and after (red dashed) applying the neutrino–electron constraint (top), and the ratio of the constrained to unconstrained fractional uncertainties (bottom). 77

5.21 Reconstructed neutrino-energy distribution for an independent ν_μ charged-current selection. The data are compared with nominal and constrained MC predictions, with and without neutrino-interaction cross-section weights. . . . 79

5.22 Reconstructed neutrino-energy distributions for the ν_μ charged-current selection before (left) and after (right) applying the flux constraint. The central predictions and the corresponding 100-universe PPFX distributions are shown together with the selected data. 80

6.1 Schematic of the Near-to-Far extrapolation for the ν_μ disappearance signal. The ND reconstructed-energy spectrum is mapped to true energy, transferred to the FD, weighted by the survival probability, and mapped to FD reconstructed energy. 86

7.1 Steps in the Profiled Feldman–Cousins method. At each test point, the observed data determine the nuisance-parameter values used to generate pseudo-experiments. The diagram is based on the procedure described in Ref. [37]. 95

8.1 Distribution of computation times for 39,680 pseudo-experiments on Perlmutter CPU nodes. The mean runtime is 1070 s, with a standard deviation of 130.5 s. 105

8.2 Evolution of the reconstructed FC confidence contours using a uniform allocation of 2000 pseudo-experiments per selected point. The GP fit fraction increases from 15.23% at iteration 0 (left) to 30.36% at iteration 1 (right). Solid curves show the GP-FC reconstruction, while dotted curves show the reference FC contours. 107

8.3 Evolution of the reconstructed FC confidence contours using nonuniform pseudo-experiment allocations. Left: 1200/800 configuration, with 48 points receiving 1200 pseudo-experiments and 16 points receiving 800 pseudo-experiments per iteration. The GP fit fractions are 5.45% at iteration 0 and 32.35% at iteration 5. Right: 1600/800 configuration, with 32 points receiving 1600 pseudo-experiments and 32 points receiving 800 pseudo-experiments per iteration. The corresponding GP fit fractions are 6.56% and 38.40%. Solid curves represent the GP-FC reconstruction, while dotted curves show the reference FC contours. 108

8.4 Comparison of the GP-FC reconstructed confidence contours with the reference FC results at iterations 0 (top left), 2 (top right), and 4 (bottom). The corresponding GP fit fractions are 15.23%, 24.69%, and 53.90%. Solid curves represent the GP-FC reconstruction, while dotted curves show the reference FC contours. 111

10.1 Locations of the three UV purity monitors inside the ProtoDUNE-SP cryostat, at different heights outside the TPC field cage [8]. 120

10.2 Engineering assembly drawing of the short UV purity monitor, showing the electrodes, grids, field rings, and mechanical supports. 120

10.3 Representative fitted anode (left) and cathode (right) waveforms. The fitted amplitudes and timing parameters are used to calculate the charge ratio and drift time [40]. 122
10.4 Charge-ratio calibration measurements for the bottom, middle, and top ProtoDUNE-SP-I UV purity monitors, respectively. The charge-ratio dependence on drift time is used to evaluate the grid-transparency correction [8]. 124
10.5 Corrected Q_a/Q_c history from the three ProtoDUNE-SP UV purity monitors between September 2018 and July 2020. Circulation began in October 2018. Subsequent excursions coincide with documented changes to the recirculation configuration, including dedicated studies and pump interruptions. The bands indicate measurement uncertainties. Reproduced from the DUNE Collaboration ProtoDUNE-SP monitoring record in Ref. [8]. 125
10.6 Electron-lifetime history from the three ProtoDUNE-SP UV purity monitors between September 2018 and February 2020, shown on a logarithmic vertical scale in milliseconds. The early low-lifetime period precedes the start of circulation in October 2018. Later downward excursions occur during documented recirculation tests or pump interruptions. The bands denote the measurement uncertainties. Reproduced from Ref. [8]. 126
10.7 Engineering assembly drawing of a long-drift UV purity-monitor design, showing the extended drift region and field rings. 128
10.8 Electric-field distribution obtained from the three-dimensional electrostatic simulation of the long UV purity monitor. The field magnitude is shown along the drift direction at the center and at transverse offsets of $\pm 20\,\mathrm{mm}$. The field is approximately uniform in the central drift region, with variations near the electrode grids. 128

10.9 Calculated grid transparency and electrostatic inefficiency for different mesh pitches and strip widths in the long-monitor design study. The results compare cathode and anode grid performance under the simulated voltage configuration. 130
10.10 Electron-lifetime histories from the ProtoDUNE-HD automated monitoring system, showing long-term and recent measurements through September 2026. 132
10.11 Corrected cathode-response history from ProtoDUNE-HD automated monitoring, extending through September 2026. 133

List of Tables

4.1 Electroweak couplings in the GENIE tree-level prediction and the radiative-correction calculation used in this analysis. 44

5.1 Event-selection cut flow showing signal efficiency, background suppression, and purity at each stage. ϵ^{step} denotes the efficiency relative to the preceding selection stage, while ϵ denotes the cumulative efficiency relative to the sample before any selection. For backgrounds, ϵ_{bkg} is the cumulative background acceptance relative to the original background sample. All efficiencies and acceptances in the table are reported as percentages. Signal and background yields are obtained from simulation and normalized to 1.38×10^{21} POT, with all event weights and analysis corrections applied. No sideband-derived background normalization is applied to the yields in this table. 59

5.2 Event yields in the signal region, $E_e\theta_e^2 < 0.005$ GeV rad^2, for $0.5 < E_e < 5$ GeV, with and without the sideband background correction. 62

5.3 Uncertainties on the integrated background-subtracted neutrino–electron signal yield. The full energy-binned covariance matrices are used in the flux constraint. 64

5.4 Effect of the neutrino–electron constraint on the integrated flux of each beam component. The means are quoted in the same integrated-flux units as the ν_μ result, per 10^{10} POT. Fractional uncertainties are the RMS/Mean of the 100-universe ensemble before and after posterior reweighting. 78

10.1 Calibration factors and assigned relative uncertainties used in the later NP04 and NP02 purity-monitor processing. 123

Acknowledgments

It has been a very, very long road. Since my undergraduate years, there have been many times when I wanted to give up. Somehow, with a lot of help and a little more determination each time, I kept going. My Ph.D. years were not quite what I expected. During the COVID-19 pandemic, I was stuck in the United States, far from the family I had always traveled so often to see. Looking back, those years were also some of my most productive. I passed my qualifying exam, helped build purity monitors, worked on a compressor for the recycling system, and did a surprising number of interesting things with liquid nitrogen. Later, when my father became seriously ill, I thought about leaving the program many times. I am grateful to Professor Bian, Wenjie, Ben, Alejandro, and the many family members and friends who helped me keep going. I am very lucky to have made it this far.

First, I want to thank my advisor, Professor Jianming Bian. He has supported me from my undergraduate years through my Ph.D. He has answered more of my questions than I can count, including many that must have seemed very basic. More importantly, he has helped me stay connected to my research and keep making progress through some very difficult years. I would not have finished this work without his patience, guidance, and support.

I also want to thank Professors Henry Sobel, Juan Pedro Ochoa-Ricoux, and Michael Smy. They listened to my long talks, asked questions that made me rethink what I was doing, and were always willing to help me work through problems. I learned a great deal from those conversations. I am especially grateful to Jeff Griskevich, the High Bay lab manager. He taught me so much about hands-on research and development: mechanical work, soldering,

wiring, aluminum-extrusion assembly, and how to turn a design into something that actually works. In many ways, he was my engineering mentor.

I am grateful to the NOvA collaboration and to the many people who supported my work through the Near Detector, reconstruction, three-flavor oscillation, and cross-section groups, as well as the analysis and paper review processes. I especially want to thank Gregory Pawloski, Josh Barrow, Matt Wetstein, Mathew Muether, Kathryn Sutton, Jon Paley, Leo Aliaga, Adam Aurisano, Pavel Snopok, Biao Wang, Linda Cremonesi, Zoya, Alex Himmel, Tricia Vahle, and Peter Shanahan for their contributions, advice, and support. I am also grateful to Erika Catano-Mur, Liudmila Kolupaeva, Ryan Nichol, and Artur Sztuc, and to the many other collaborators whose help made this work possible. Thank you for your time, your suggestions, and your understanding. Your feedback helped me develop and refine these analyses, and your support gave me opportunities to test the GP–FC approach at NERSC. I appreciate the patience and trust you showed me throughout this process.

I also want to thank my fellow graduate students and colleagues: Benjamin Jargowsky, Alejandro Yankelevich, Larry, Zhengyang, and Zhongyi. Thank you for answering questions, sharing ideas, and making the work more enjoyable. I owe a special thank-you to Wenjie Wu for always being willing to talk through my neutrino–electron analysis questions. I still remember the time we spent building prototypes together in the High Bay lab. It was hard work, but it was also a lot of fun. I am deeply grateful to Nitish Nayak, who started teaching me how to do research when I was an undergraduate and guided me patiently, step by step. I also want to thank Jiaxi and Bin. I have always admired the care and consistency you bring to your work.

Most of all, I want to thank my parents, Hanping Xiao and Li Han. Thank you for supporting me all these years. I have never been an easy person. I ask too many questions, I am stubborn, and I do not always listen. But you never gave up on me. More than once, you pulled me back when I felt I was standing at the edge of a cliff. I am so lucky to be your daughter.

Dad, thank you for supporting my education even when you were seriously ill. You went through treatments that I know must have been painful, and you rarely let us see how hard they were for you. You wanted me to be able to keep studying without carrying all of your worries. Wherever you are, I hope you can see that I made it to this day. You promised you would move the tassel on my graduation cap. I am still holding you to that promise.

To everyone who supported me, worked with me, cared for me, and did not give up on me along the way: thank you.

This work was supported by the United States Department of Energy under Grant No. DE-SC0009920.

Vita

Yiwen Xiao

EDUCATION

Doctor of Philosophy in Physics **2026**

University of California, Irvine *Irvine, California*

Bachelor of Science in Applied Physics **2019**

University of California, Irvine *Irvine, California*

RESEARCH EXPERIENCE

Graduate Student Researcher **2020–2026**

University of California, Irvine *Irvine, California*

TEACHING EXPERIENCE

Teaching Assistant **2019–2021**

University of California, Irvine *Irvine, California*

ABSTRACT OF THE DISSERTATION

Precision Methods for Long-Baseline Neutrino Experiments:
Flux Constraints, Statistical Inference, and
Liquid-Argon Detector Monitoring

By

Yiwen Xiao

Doctor of Philosophy in Physics

University of California, Irvine, 2026

Professor Jianming Bian, Chair

Long-baseline neutrino oscillation experiments use measurements from near and far detectors to study neutrino flavor transitions and determine oscillation parameters. NOvA uses the NuMI neutrino beam at Fermilab to measure muon-neutrino disappearance and electron-neutrino appearance, together with the corresponding antineutrino oscillation channels. DUNE is a next-generation long-baseline experiment that will use liquid-argon time projection chambers to make precision measurements of neutrino oscillations. This dissertation presents three studies: a neutrino flux constraint using neutrino–electron elastic scattering in the NOvA Near Detector, an accelerated Feldman–Cousins method using Gaussian processes for neutrino oscillation analyses, and contributions to liquid-argon purity monitoring for ProtoDUNE.

The first study constrains the neutrino flux using neutrino–electron elastic scattering in the NOvA Near Detector. The neutrino–electron elastic scattering cross section can be precisely predicted using electroweak theory, making this interaction suitable for neutrino flux measurements. The scattered electron travels nearly along the incident neutrino direction,

providing a useful signature for identifying signal events. Using an exposure of 1.38×10^{21} protons on target (POT), we select 827 neutrino–electron scattering candidates and estimate a signal yield of 740.1 ± 30.2 (stat.) events after background subtraction. We then use the reconstructed energy spectrum of the selected electron candidates to constrain an ensemble of 100 PPFX flux universes through Bayesian reweighting, incorporating statistical and non-PPFX systematic uncertainties. The resulting constraint reduces the fractional uncertainty on the integrated muon-neutrino flux from 8.1% to 4.0%, with a posterior mean of 126.4 $\nu_\mu\,\mathrm{cm}^{-2}/(10^{10}$ POT).

The second study develops an accelerated Feldman–Cousins method using Gaussian processes for NOvA neutrino oscillation analyses. Standard asymptotic approximations for likelihood-ratio test statistics may not provide accurate confidence intervals in the presence of physical parameter boundaries, degeneracies, or limited event statistics. The Feldman–Cousins method uses pseudo-experiments to determine the test-statistic distribution at each point in parameter space. However, generating enough pseudo-experiments requires many likelihood fits and is computationally expensive. GP–FC reduces this cost by using Gaussian-process regression to approximate the significance surface and identify regions of parameter space that require additional pseudo-experiments. We parallelize independent likelihood fits using MPI and the DIY framework, with dynamic load balancing and checkpointing to support large-scale calculations. The intersection-over-union values for iterations 0, 2, and 4 are $(96.56 \pm 0.05)\%$, $(97.88 \pm 0.05)\%$, and $(98.14 \pm 0.05)\%$, respectively, when compared with the reference Feldman–Cousins confidence region.

The third study describes contributions to liquid-argon purity monitoring for the DUNE detector development program. In liquid-argon time projection chambers, electronegative impurities capture drifting ionization electrons and reduce the charge reaching the readout. Purity monitors measure this charge attenuation to determine the electron lifetime in liquid argon. This dissertation describes contributions to the design, construction, operation, waveform analysis, and calibration of purity monitors used in ProtoDUNE. During high-

purity operation of ProtoDUNE-SP, purity monitors in the upper and middle regions measured electron lifetimes exceeding 70 ms, while monitors in the lower region measured lifetimes of at least 30 ms. The dissertation also presents subsequent work on instrumentation, calibration, and automated monitoring for ProtoDUNE-HD.

Chapter 1

Introduction

Neutrinos are electrically neutral leptons that interact through the weak interaction. Their small interaction cross sections make them difficult to detect, despite their abundance in nature. Neutrinos are produced in nuclear reactions, particle decays, and astrophysical processes, and their properties can be studied over a wide range of energies and propagation distances.

The Standard Model was initially formulated with massless neutrinos. Experiments have since established that neutrinos can change flavor as they propagate. These observations show that neutrino flavor states are mixtures of mass eigenstates with nonzero mass-squared differences, providing evidence for physics beyond the Standard Model [1].

Long-baseline accelerator experiments study neutrino oscillations using a neutrino beam and detectors located at different distances from the source. The observed event spectra depend on the neutrino flux, interaction cross sections, detector response, and oscillation probabilities. Oscillation analyses use predictions of these quantities to determine the oscillation parameters and evaluate their uncertainties.

This dissertation examines several aspects of precision measurements in experimental neutrino physics, including neutrino flux constraints, statistical inference, and liquid-argon detector instrumentation. The research was conducted within the NOvA experiment and the

ProtoDUNE detector-development program.

This chapter introduces the neutrino physics and experimental background relevant to these studies. The subsequent chapters present the theoretical framework, experimental methods, analyses, and results.

1.1 Neutrinos and the Standard Model

The Standard Model describes elementary particles and their interactions through the electromagnetic, weak, and strong forces. Matter particles are divided into quarks and leptons, which are spin-1/2 fermions. The lepton sector contains three charged leptons, the electron (e), muon (μ), and tau (τ), and three corresponding neutrino flavors, ν_e, ν_μ, and ν_τ.

Neutrinos have no electric or color charge. Their interactions with matter are mediated by the weak force through the $W^\pm$ and Z bosons. Charged-current (CC) interactions involve the exchange of a $W^\pm$ boson and produce a charged lepton associated with the interacting neutrino flavor. Neutral-current (NC) interactions involve the exchange of a Z boson and leave a neutrino in the final state.

Wolfgang Pauli first proposed the existence of the neutrino in 1930 to explain the continuous energy spectrum observed in nuclear beta decay while preserving energy and momentum conservation. Clyde Cowan and Frederick Reines detected electron antineutrinos from a nuclear reactor in 1956. Later experiments established the existence of the muon neutrino and tau neutrino.

Neutrino flavor states, which participate in weak interactions, are superpositions of states with definite masses. The mass eigenstates evolve with different phases during propagation, and interference between them produces neutrino oscillations. The oscillation probabilities depend on neutrino energy, propagation distance, and the parameters that describe neutrino mixing.

The Pontecorvo–Maki–Nakagawa–Sakata (PMNS) matrix describes the relationship be-

tween neutrino flavor and mass eigenstates. In the three-flavor framework, neutrino oscillations are described by three mixing angles, two independent mass-squared differences, and a CP-violating phase. Oscillation measurements determine the mass-squared differences rather than the absolute neutrino masses. Chapter 2 presents the mixing formalism and oscillation probabilities in detail.

1.2 Experimental Evidence for Neutrino Oscillations

The first experimental evidence for neutrino oscillations came from solar and atmospheric neutrino measurements. Early solar-neutrino experiments observed fewer electron neutrinos than predicted by models of nuclear fusion in the Sun. The deficit raised questions about the solar-neutrino flux and the properties of neutrinos during propagation.

The Sudbury Neutrino Observatory (SNO) provided an important measurement by comparing charged-current interactions, which are sensitive to electron neutrinos, with neutral-current interactions, which are sensitive to all three active neutrino flavors. The total active-neutrino flux agreed with solar-model predictions, while the measured electron-neutrino flux was significantly smaller. The difference showed that some of the electron neutrinos produced in the Sun had changed flavor before reaching Earth.

Atmospheric neutrinos are produced when cosmic rays interact with nuclei in Earth's atmosphere. Their propagation distances vary with the direction of travel. Downward-going neutrinos typically travel tens of kilometers before reaching an underground detector, while upward-going neutrinos can travel through much of the Earth.

Super-Kamiokande observed a deficit of muon-neutrino-like events that depended on the neutrino zenith angle. The dependence on propagation distance was consistent with muon-neutrino disappearance through neutrino oscillations.

Reactor experiments provided additional measurements of electron-antineutrino disappearance. KamLAND observed disappearance consistent with the mass-squared splitting

associated with solar-neutrino oscillations. Daya Bay measured a nonzero value of the mixing angle θ_{13} [2, 3].

Accelerator experiments use neutrino beams to study muon-neutrino disappearance and electron-neutrino appearance, together with the corresponding antineutrino oscillation channels. These measurements provide sensitivity to the atmospheric mixing parameters, the neutrino mass ordering, and the CP-violating phase.

1.3 Neutrino Sources and Detection

Accelerator neutrino beams provide a controlled source for long-baseline oscillation experiments. High-energy protons strike a target and produce secondary hadrons, primarily pions and kaons. Magnetic horns focus charged hadrons of the selected polarity, which subsequently decay in a decay pipe.

The resulting beam contains mostly muon neutrinos or muon antineutrinos, depending on the focusing configuration. Smaller contributions from other neutrino flavors are also present. The beam energy spectrum depends on the momenta and decay kinematics of the parent hadrons, the focusing system, and the beamline geometry.

Beam simulations use models of hadron production, magnetic focusing, particle propagation, and decay to predict the neutrino flux. Uncertainties in hadron production and beamline conditions affect the predicted flux normalization and energy spectrum.

Most neutrino interactions recorded in accelerator oscillation experiments occur on nuclei in the detector. In charged-current interactions, the outgoing charged lepton provides information about the neutrino flavor. The other particles produced in the interaction contribute to the visible energy and event topology.

Nuclear effects influence both the particles produced and the reconstructed neutrino energy. Initial-state nucleon motion, multinucleon interactions, and final-state interactions can change the energy and momentum distributions of particles leaving the nucleus. These

effects introduce uncertainties into neutrino interaction models and energy reconstruction.

Neutral-current interactions do not produce a charged lepton associated with the incoming neutrino flavor. Some neutral-current events produce electromagnetic activity that can resemble an electron-neutrino charged-current interaction, making them an important background in appearance analyses.

Neutrino–electron elastic scattering provides a way to constrain the neutrino flux using a precisely predicted cross section. The outgoing electron travels at a small angle relative to the incident neutrino direction, and its reconstructed energy and direction help identify signal candidates. The cross section is much smaller than typical neutrino–nucleus interaction cross sections, so the expected signal sample is relatively small. Background rejection and detector calibration are important for obtaining a reliable flux constraint.

In the NOvA detectors, charged particles deposit energy in liquid scintillator and produce light. Wavelength-shifting fibers collect the light and guide it to avalanche photodiodes, which convert the optical signals into electrical signals. The reconstruction algorithms use the recorded signals to identify tracks and showers and estimate their energies.

The reconstructed energy spectrum depends on detector calibration, energy resolution, and event-selection efficiency. These effects must be included when comparing the observed events with predictions from neutrino interaction and detector simulations. The NOvA detector and reconstruction methods are described in Chapter 3.

1.4 Open Questions in Neutrino Physics

Several questions about neutrino oscillations remain unresolved. Long-baseline experiments are particularly sensitive to the neutrino mass ordering, the octant of θ_{23}, and the CP-violating phase δ_{CP}.

The neutrino mass ordering describes the relative positions of the three neutrino mass eigenstates. In the normal ordering, the third mass eigenstate is heavier than the other two.

In the inverted ordering, it is lighter.

As neutrinos propagate through matter, coherent forward scattering modifies their oscillation probabilities. The matter effect differs for neutrinos and antineutrinos and provides sensitivity to the mass ordering in long-baseline experiments.

The atmospheric mixing angle θ_{23} controls the leading amplitude of muon-neutrino disappearance. Measurements of this channel primarily constrain $\sin^2(2\theta_{23})$, which has the same value for angles equally displaced above and below 45°. Electron-neutrino appearance measurements provide additional information that can help distinguish the two octants.

The CP-violating phase δ_{CP} affects the interference between oscillation amplitudes. Neutrino and antineutrino appearance probabilities can differ because of intrinsic CP violation and matter effects. The analysis must account for both contributions when interpreting the observed differences.

Oscillation parameters can also be correlated. Different combinations of mass ordering, θ_{23}, and δ_{CP} may produce similar event spectra over the energy range measured by an experiment. Combining appearance and disappearance data helps distinguish these parameter combinations.

1.5 Long-Baseline Accelerator Experiments

Long-baseline neutrino experiments study neutrino oscillations using detectors located at different distances from an accelerator neutrino source. The Near Detector measures neutrino interactions close to the source, where oscillation effects are small. The Far Detector measures the neutrino energy spectrum and flavor composition after the neutrinos have traveled hundreds or thousands of kilometers. The Near Detector measurements help establish the expected event spectra at the Far Detector, where muon-neutrino disappearance and electron-neutrino appearance are studied. These oscillation channels provide sensitivity to different oscillation parameters. Measurements with neutrino and antineutrino beams also

help constrain the CP-violating phase and the neutrino mass ordering through their different oscillation probabilities and matter effects.

The Near Detector helps constrain uncertainties in the neutrino flux, interaction cross sections, and detector response. These constraints improve the predictions used to analyze the Far Detector data. However, the Near Detector measurements cannot be applied directly to the Far Detector. The two detectors observe different neutrino energy distributions because of their positions relative to the beam source. Their geometries and detector responses can also differ, affecting the reconstructed event spectra.

Near-to-far extrapolation uses Near Detector measurements to correct the predicted event spectrum at the Far Detector. The procedure compares the reconstructed energy spectrum in Near Detector data with the simulation and uses the observed differences to correct the prediction. Beam and detector simulations account for the differences in neutrino energy distributions and detector response between the two locations. The corrected prediction is then used to calculate the expected Far Detector spectrum, including neutrino oscillations.

Near-to-far extrapolation reduces uncertainties shared by the two detectors, but it cannot eliminate all uncertainties in the neutrino flux, interaction models, or detector response. Independent measurements can provide further constraints. For example, neutrino–electron elastic scattering provides a complementary constraint on the predicted neutrino flux because its signal cross section does not depend on neutrino–nucleus interaction models.

1.6 Precision in Long-Baseline Measurements

Long-baseline neutrino experiments determine oscillation parameters by comparing observed event spectra with predictions calculated for different values of those parameters. The predictions depend on neutrino oscillation probabilities, the incident neutrino flux, interaction cross sections, and detector response. The flux determines the energy distribution of neutrinos reaching the detector, while the interaction cross sections determine their interaction rates

and the particles produced in those interactions. The detector response determines how these particles are observed and how their energies are reconstructed. Uncertainties in these quantities affect the predicted event rates and reconstructed energy spectra.

Since oscillation parameters are extracted from changes in the event spectra, uncertainties in the predictions can influence their fitted values and precision. For example, variations in the neutrino flux can change the predicted energy spectrum, while an uncertainty in the detector energy scale can shift the reconstructed position of an oscillation feature. Uncertainties in interaction models can also affect the relationship between neutrino energy and reconstructed energy through changes in the predicted final-state particle distributions. These effects can produce spectral changes similar to those associated with different oscillation parameters, making it important to account for them in the analysis.

As discussed in the previous section, the Near Detector helps constrain many of these uncertainties. Near-to-far extrapolation uses its measurements to improve the Far Detector prediction and reduce uncertainties shared by the two detectors. The Near Detector measures event spectra that depend on the combined effects of the flux, interaction cross sections, and detector response, so its data do not necessarily determine each contribution independently. The remaining systematic uncertainties are included in oscillation fits when comparing the Far Detector data with the predictions.

The statistical interpretation of these fits presents a separate consideration. Oscillation parameters are estimated by fitting predicted event spectra to the observed data, and confidence regions are constructed using a test statistic evaluated across parameter space. Under suitable conditions, the distribution of a likelihood-ratio test statistic can be approximated by a chi-squared distribution. This approximation may be inaccurate near physical parameter boundaries, in the presence of parameter degeneracies, or when event statistics are limited. The Feldman–Cousins method instead uses pseudo-experiments to determine the test-statistic distributions and construct confidence regions. Applying this method to multidimensional oscillation analyses requires many pseudo-experiments and likelihood fits,

resulting in substantial computational cost.

1.7 NOvA and DUNE

NOvA and DUNE are two long-baseline neutrino experiments designed to investigate neutrino oscillations using accelerator-produced neutrino beams. Although they share similar physics goals, their experimental configurations and detector technologies differ.

The NOvA experiment uses the NuMI neutrino beam produced at Fermilab and two functionally identical segmented liquid-scintillator detectors separated by an 810 km baseline [4, 5]. Both detectors are positioned off the central beam axis, producing a relatively narrow neutrino-energy spectrum centered near the first oscillation maximum. NOvA studies muon-neutrino disappearance and electron-neutrino appearance using neutrino and antineutrino beam configurations. The Near Detector also provides data for neutrino-interaction measurements and constraints on the beam-flux prediction.

The Deep Underground Neutrino Experiment (DUNE) will use a high-intensity neutrino beam from Fermilab and large liquid-argon time projection chambers at the Sanford Underground Research Facility [6]. These detectors record the ionization produced by charged particles traversing liquid argon, providing detailed three-dimensional information about particle trajectories and energy deposition. Maintaining high liquid-argon purity is important for preserving the ionization charge as electrons drift toward the readout planes. The ProtoDUNE detectors at CERN test detector technologies and operating procedures relevant to DUNE, including liquid-argon purification and purity monitoring [7, 8].

1.8 Scope and Contributions of This Dissertation

This dissertation presents research on neutrino flux constraints, statistical methods for oscillation analyses, and liquid-argon detector instrumentation. The studies were conducted within the NOvA experiment and the ProtoDUNE detector-development program, addressing

different aspects of experimental neutrino physics.

The NOvA study uses neutrino–electron elastic scattering events selected in the Near Detector to constrain uncertainties in the NuMI neutrino flux prediction. The analysis involves event selection, background estimation, and the evaluation of systematic uncertainties. A Bayesian framework is used to reweight an ensemble of flux predictions based on the observed electron-energy spectrum, yielding a constrained flux prediction and its associated uncertainty.

The statistical study develops a parallel implementation of Gaussian-process-accelerated Feldman–Cousins confidence-region construction. It combines adaptive sampling of parameter space with distributed likelihood fits to reduce the computational demands of the Feldman–Cousins procedure. The implementation is evaluated in the context of a NOvA oscillation analysis.

The detector-instrumentation work focuses on liquid-argon purity monitors developed for ProtoDUNE. It includes contributions to detector design and construction, laboratory testing, operation, waveform analysis, calibration, and automated monitoring. These activities support the measurement and monitoring of electron lifetime during liquid-argon detector operation.

The following chapters provide the theoretical, experimental, and analytical background for these studies. Chapter 2 introduces neutrino oscillation theory, followed by the NuMI beam and NOvA detectors in Chapter 3. The neutrino–electron scattering analysis and Bayesian flux constraint are presented in Chapters 4 and 5. The NOvA oscillation-analysis framework in Chapter 6 provides the context for the Feldman–Cousins and Gaussian-process studies in Chapters 7 and 8. The discussion then turns to DUNE and liquid-argon detector technology in Chapter 9, followed by the purity-monitor work in Chapter 10. The dissertation concludes in Chapter 11.

Chapter 2

Neutrino Oscillation Physics

This chapter develops the three-flavor framework for neutrino oscillations, from the relation between flavor and mass eigenstates to oscillation probabilities in vacuum and matter. It concludes by connecting these probabilities to the reconstructed event spectra used in long-baseline analyses.

2.1 Flavor and Mass Eigenstates

Neutrinos are produced and detected through weak interactions. In charged-current interactions, each neutrino flavor is associated with a corresponding charged lepton: an electron, muon, or tau. These interaction states are referred to as flavor eigenstates and are denoted by $|\nu_\alpha\rangle$, where $\alpha = e, \mu, \tau$.

The states that propagate with definite masses are called mass eigenstates, denoted by $|\nu_i\rangle$ with $i = 1, 2, 3$. The flavor and mass eigenstates are related by the unitary Pontecorvo–Maki–Nakagawa–Sakata (PMNS) matrix [9, 10],

$$|\nu_\alpha\rangle = \sum_{i=1}^{3} U_{\alpha i}^{*} |\nu_i\rangle. \tag{2.1}$$

The coefficients $U_{\alpha i}^{*}$ specify the mass-eigenstate components of a neutrino produced with

flavor α.

During propagation, the mass eigenstates evolve with different quantum-mechanical phases. Their relative phases change with the distance traveled, altering the interference between the mass components. As a result, the probability of detecting a neutrino in a particular flavor state can differ from the probability associated with its production flavor. This phenomenon is known as neutrino flavor oscillation.

The oscillation probabilities depend on differences between the squared masses of the neutrino mass eigenstates, defined as

$$\Delta m_{ij}^2 = m_i^2 - m_j^2. \tag{2.2}$$

For three mass eigenstates, only two of these differences are independent: $\Delta m_{31}^2 = \Delta m_{32}^2 + \Delta m_{21}^2$.

2.2 PMNS Matrix

For three active neutrino flavors, the PMNS matrix can be parameterized by three mixing angles, θ_{12}, θ_{13}, and θ_{23}, and one Dirac CP-violating phase, δ_{CP}:

$$U = \begin{pmatrix} 1 & 0 & 0 \\ 0 & c_{23} & s_{23} \\ 0 & -s_{23} & c_{23} \end{pmatrix} \begin{pmatrix} c_{13} & 0 & s_{13}e^{-i\delta_{\mathrm{CP}}} \\ 0 & 1 & 0 \\ -s_{13}e^{i\delta_{\mathrm{CP}}} & 0 & c_{13} \end{pmatrix} \begin{pmatrix} c_{12} & s_{12} & 0 \\ -s_{12} & c_{12} & 0 \\ 0 & 0 & 1 \end{pmatrix}, \tag{2.3}$$

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

The angle θ_{12} determines the relative electron-flavor contributions of the first two mass eigenstates:

$$|U_{e1}|^2 = c_{12}^2 c_{13}^2, \qquad |U_{e2}|^2 = s_{12}^2 c_{13}^2. \tag{2.4}$$

Together with Δm_{21}^2, this angle governs the dominant oscillations observed in solar neutrino

experiments and long-baseline reactor experiments.

The angle θ_{13} determines the electron-flavor contribution of the third mass eigenstate,

$$|U_{e3}|^2 = \sin^2\theta_{13}. \tag{2.5}$$

Its nonzero value allows interference between different oscillation amplitudes in $\nu_\mu \to \nu_e$ appearance, providing sensitivity to the CP-violating phase. The angle θ_{13} also controls the leading oscillation amplitude in short-baseline reactor antineutrino disappearance measurements.

The angle θ_{23} determines the relative muon- and tau-flavor contributions of the third mass eigenstate:

$$|U_{\mu 3}|^2 = s_{23}^2 c_{13}^2, \qquad |U_{\tau 3}|^2 = c_{23}^2 c_{13}^2. \tag{2.6}$$

It plays a central role in atmospheric-scale muon-neutrino disappearance. When $\theta_{23} = 45°$, the muon- and tau-flavor contributions are equal, corresponding to maximal mixing. Values below and above 45° define the lower and upper octants, respectively.

The phase δ_{CP} enters the complex elements of the mixing matrix. For antineutrinos, the mixing amplitudes are complex conjugated, reversing the sign of CP-odd interference terms in vacuum. A nonzero CP-violating contribution can produce different appearance probabilities for neutrinos and antineutrinos.

The solar splitting satisfies $\Delta m^2_{21} > 0$, while the sign of Δm^2_{31} determines the neutrino mass ordering. A positive value corresponds to normal ordering, with ν_3 heavier than ν_1 and ν_2. A negative value corresponds to inverted ordering, with ν_3 lighter than the other two mass eigenstates.

If neutrinos are Majorana particles, two additional phases enter the mixing matrix, but these phases do not affect standard neutrino oscillation probabilities. The PMNS matrix is unitary in the three-active-neutrino framework, so the probabilities summed over all detected active flavors equal one.

2.3 Vacuum Oscillations

Neutrino oscillations arise from the interference of mass eigenstates during propagation. A neutrino produced in a flavor state $|\nu_\alpha\rangle$ is a superposition of mass eigenstates, as defined in Eq. (2.1). In vacuum, each mass eigenstate propagates independently, with its evolution described by

$$|\nu_i(t, L)\rangle = e^{-i(E_i t - p_i L)}|\nu_i\rangle, \tag{2.7}$$

where E_i and p_i are the energy and momentum of the ith mass eigenstate, respectively. Natural units, $\hbar = c = 1$, are used throughout the derivation.

The relativistic energy-momentum relation gives

$$E_i^2 = p_i^2 + m_i^2. \tag{2.8}$$

Since neutrino masses are much smaller than their energies, the energy can be expanded as

$$E_i = \sqrt{p_i^2 + m_i^2} \simeq p_i + \frac{m_i^2}{2p_i}. \tag{2.9}$$

For an ultrarelativistic neutrino traveling a distance L, the propagation time is approximately $t \simeq L$. Consequently,

$$E_i t - p_i L \simeq (E_i - p_i)L \simeq \frac{m_i^2 L}{2E_\nu}, \tag{2.10}$$

where E_ν is the neutrino energy to leading order. This treatment assumes coherent, ultrarelativistic propagation and neglects higher-order mass corrections.

The phase associated with each mass eigenstate contains a contribution that is common to all components. Such a common phase does not affect measurable probabilities. The relative phase between two mass eigenstates is

$$\phi_i - \phi_j = \frac{(m_i^2 - m_j^2)L}{2E_\nu} = \frac{\Delta m_{ij}^2 L}{2E_\nu}. \tag{2.11}$$

This relation shows why neutrino oscillations depend on mass-squared differences rather than on the absolute neutrino masses.

Using the flavor–mass relation, the neutrino state after propagation can be written as

$$|\nu_\alpha(L)\rangle = \sum_i U^*_{\alpha i} \exp\left(-i\frac{m_i^2 L}{2E_\nu}\right)|\nu_i\rangle. \tag{2.12}$$

The amplitude for detecting flavor β is the projection of this state onto $|\nu_\beta\rangle$:

$$\mathcal{A}_{\alpha\to\beta} = \langle\nu_\beta|\nu_\alpha(L)\rangle = \sum_i U_{\beta i}U^*_{\alpha i} \exp\left(-i\frac{m_i^2 L}{2E_\nu}\right). \tag{2.13}$$

The oscillation probability is the squared magnitude of the transition amplitude,

$$P_{\alpha\to\beta} = |\mathcal{A}_{\alpha\to\beta}|^2. \tag{2.14}$$

Expanding the product gives

$$P_{\alpha\to\beta} = \sum_{i,j} U_{\beta i}U^*_{\alpha i}U^*_{\beta j}U_{\alpha j} \exp\left[-i\frac{\Delta m_{ij}^2 L}{2E_\nu}\right]. \tag{2.15}$$

The terms with $i = j$ are independent of the propagation distance, while the terms with $i \neq j$ contain interference between different mass eigenstates. These interference terms are responsible for the oscillatory dependence of the flavor-transition probability on L/E_ν.

Defining the oscillation phase as

$$\Delta_{ij} = \frac{\Delta m_{ij}^2 L}{4E_\nu}, \tag{2.16}$$

and using the unitarity of the PMNS matrix, the probability can be expressed as

$$P_{\alpha\to\beta} = \delta_{\alpha\beta} - 4\sum_{i>j} \mathrm{Re}\left(U^*_{\alpha i}U_{\beta i}U_{\alpha j}U^*_{\beta j}\right)\sin^2\Delta_{ij} + 2\sum_{i>j}\mathrm{Im}\left(U^*_{\alpha i}U_{\beta i}U_{\alpha j}U^*_{\beta j}\right)\sin 2\Delta_{ij}. \tag{2.17}$$

In practical units, the oscillation phase becomes

$$\Delta_{ij} \simeq 1.267\,\frac{\Delta m^2_{ij}[\mathrm{eV}^2]\,L[\mathrm{km}]}{E_\nu[\mathrm{GeV}]}. \tag{2.18}$$

The terms proportional to $\sin^2\Delta_{ij}$ depend on the real parts of products of PMNS matrix elements and are unchanged under complex conjugation of the mixing matrix. The terms proportional to $\sin 2\Delta_{ij}$ depend on the imaginary parts and change sign under complex conjugation. In vacuum, antineutrino oscillation probabilities are obtained by replacing U with U^*. The imaginary terms can therefore produce differences between neutrino and antineutrino appearance probabilities.

For a survival probability, the initial and final flavors are identical, $\alpha = \beta$. In this case, the products of mixing matrix elements become $|U_{\alpha i}|^2|U_{\alpha j}|^2$, which are real. The survival probability is therefore

$$P_{\alpha\alpha} = 1 - 4\sum_{i>j}|U_{\alpha i}|^2|U_{\alpha j}|^2\sin^2\Delta_{ij}. \tag{2.19}$$

2.3.1 Two-Flavor Limit

The two-flavor approximation describes the dominant oscillation behavior when one mass-squared splitting sets the relevant oscillation scale. In the atmospheric-dominated regime, the solar oscillation phase is small, so the first two mass eigenstates evolve with nearly the same phase. The three-flavor probability can then be approximated by an effective two-flavor form.

For two neutrino flavors, ν_α and ν_β, the flavor and mass eigenstates are related by a single

mixing angle θ:

$$\begin{pmatrix} |\nu_\alpha\rangle \\ |\nu_\beta\rangle \end{pmatrix} = \begin{pmatrix} \cos\theta & \sin\theta \\ -\sin\theta & \cos\theta \end{pmatrix} \begin{pmatrix} |\nu_1\rangle \\ |\nu_2\rangle \end{pmatrix}. \tag{2.20}$$

A neutrino initially produced as ν_α is therefore a superposition of two mass eigenstates. After propagating a distance L, its state becomes

$$|\nu_\alpha(L)\rangle = \cos\theta\, e^{-i\phi_1}|\nu_1\rangle + \sin\theta\, e^{-i\phi_2}|\nu_2\rangle, \tag{2.21}$$

where $\phi_i = m_i^2 L/(2E_\nu)$.

Projecting this state onto the initial flavor gives the survival amplitude

$$\mathcal{A}_{\alpha\to\alpha} = \cos^2\theta\, e^{-i\phi_1} + \sin^2\theta\, e^{-i\phi_2}. \tag{2.22}$$

Taking the squared magnitude yields

$$\begin{aligned} P_{\alpha\alpha} &= \cos^4\theta + \sin^4\theta + 2\sin^2\theta\cos^2\theta\cos(\phi_2 - \phi_1) \\ &= 1 - \sin^2(2\theta)\sin^2\left(\frac{\Delta m^2 L}{4E_\nu}\right), \end{aligned} \tag{2.23}$$

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

In practical units, the survival probability is

$$P_{\alpha\alpha} = 1 - \sin^2(2\theta)\sin^2\left(1.267\,\frac{\Delta m^2[\mathrm{eV}^2]L[\mathrm{km}]}{E_\nu[\mathrm{GeV}]}\right). \tag{2.24}$$

The corresponding transition probability is

$$P_{\alpha\beta} = \sin^2(2\theta)\sin^2\left(\frac{\Delta m^2 L}{4E_\nu}\right), \qquad \alpha \neq \beta. \tag{2.25}$$

The two probabilities sum to unity, as expected for a system containing only two flavors.

The oscillation amplitude is determined by $\sin^2(2\theta)$. For maximal mixing, $\theta = 45°$, the

amplitude reaches unity and the survival probability can vanish at an oscillation minimum. For smaller mixing angles, the disappearance is less pronounced.

The oscillation phase determines how the probability varies with propagation distance and neutrino energy. The oscillation length, defined as the distance over which the probability completes one full cycle, is

$$L_{\text{osc}} = \frac{4\pi E_\nu}{|\Delta m^2|}. \tag{2.26}$$

For a fixed baseline, the first disappearance minimum occurs when the magnitude of the oscillation phase equals $\pi/2$, corresponding to

$$E_{\text{min}} = \frac{|\Delta m^2| L}{2\pi} \tag{2.27}$$

in natural units. At a known baseline, the energy of this minimum constrains $|\Delta m^2|$.

This two-state result provides a useful reference for interpreting the dominant oscillation features in long-baseline disappearance spectra.

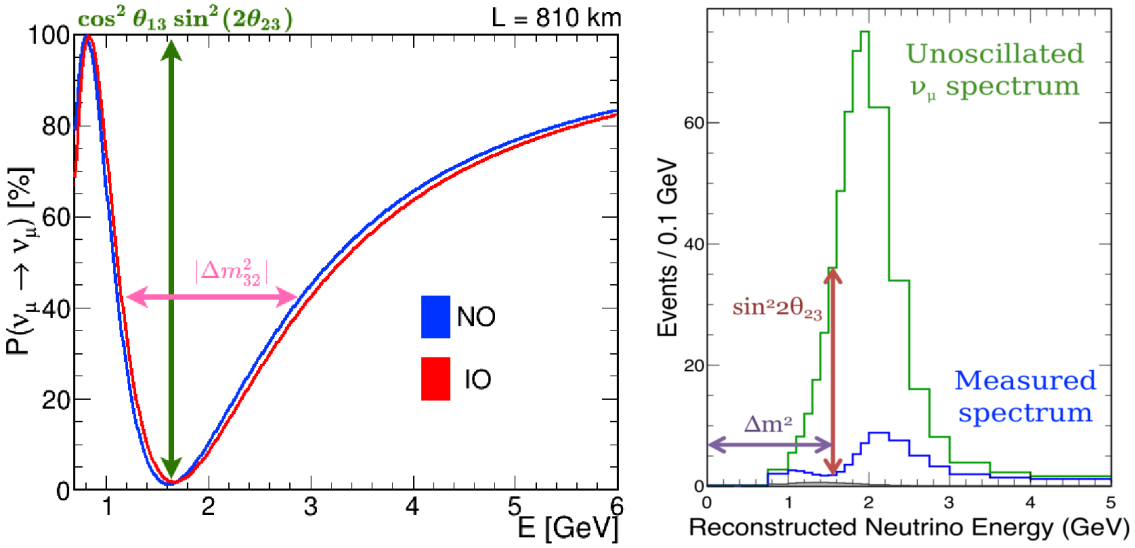

(a) (b)

Figure 2.1: Energy dependence of the ν_μ survival probability and the corresponding change in the Far Detector energy spectrum. The disappearance feature provides sensitivity to the atmospheric mass splitting and θ_{23}.

2.4 Three-Flavor Oscillation Probabilities

The two mass-squared differences generate distinct oscillation scales. At accelerator long-baseline energies, the atmospheric phase is generally much larger than the solar phase, $|\Delta_{31}| \gg |\Delta_{21}|$. The smaller solar contribution nevertheless matters in appearance measurements because it interferes with the atmospheric amplitude. Muon-neutrino disappearance and electron-neutrino appearance consequently probe different combinations of the PMNS parameters.

Muon-Neutrino Disappearance In the limit $\Delta_{21} \to 0$, the two lighter mass eigenstates have the same propagation phase. Their combined contribution can be evaluated using PMNS unitarity, giving

$$P_{\mu\mu} \simeq 1 - 4|U_{\mu 3}|^2(1 - |U_{\mu 3}|^2)\sin^2\Delta_{31}. \tag{2.28}$$

Substituting $|U_{\mu 3}|^2 = s_{23}^2 c_{13}^2$ gives the atmospheric-scale disappearance amplitude. Neglecting corrections associated with the small angle θ_{13}, the familiar approximation is

$$P(\nu_\mu \to \nu_\mu) \simeq 1 - \sin^2(2\theta_{23})\sin^2\Delta_{31}. \tag{2.29}$$

The energy of the disappearance minimum is sensitive to the atmospheric mass splitting, whereas its depth constrains the mixing amplitude. Because $\sin^2(2\theta_{23})$ is invariant under $\theta_{23} \to 90^\circ - \theta_{23}$, the leading disappearance term cannot distinguish the lower and upper octants. The exact three-flavor probability contains corrections that weaken this degeneracy.

Electron-Neutrino Appearance Electron-neutrino appearance involves both the atmospheric transition amplitude, proportional to s_{13}, and the smaller solar amplitude, proportional to Δ_{21}. Their interference introduces sensitivity to the CP phase. Expanding the vacuum probability in the small quantities s_{13} and Δ_{21}, and retaining terms of combined order two,

gives

$$P(\nu_\mu \to \nu_e) \simeq 4s_{13}^2 s_{23}^2 \sin^2 \Delta_{31} + 8s_{13}s_{12}c_{12}s_{23}c_{23} \sin \Delta_{31} \sin \Delta_{21} \cos(\Delta_{31} + \delta_{\mathrm{CP}}) + 4s_{12}^2 c_{12}^2 c_{23}^2 \sin^2 \Delta_{21}. \quad (2.30)$$

The atmospheric term sets the dominant appearance rate and depends on s_{23}^2, rather than $\sin^2(2\theta_{23})$. Appearance data can therefore help distinguish the two θ_{23} octants. Reactor antineutrino measurements independently constrain θ_{13}, reducing the ambiguity between these mixing angles.

The interference term depends on the relative phase of the atmospheric and solar amplitudes. Writing $\cos(\Delta_{31} + \delta_{\mathrm{CP}}) = \cos \Delta_{31} \cos \delta_{\mathrm{CP}} - \sin \Delta_{31} \sin \delta_{\mathrm{CP}}$ separates the CP-even and CP-odd contributions. In vacuum, the antineutrino probability follows by changing δ_{CP} to $-\delta_{\mathrm{CP}}$; the term proportional to $\sin \delta_{\mathrm{CP}}$ changes sign. The last term is the purely solar contribution and is suppressed at atmospheric-scale baselines by the smallness of Δ_{21}.

These expressions expose the leading parameter dependencies but do not replace full three-flavor calculations. Matter modifies the relative propagation phases and can generate an additional neutrino–antineutrino asymmetry.

2.5 Matter Effects

Neutrino propagation through matter differs from propagation in vacuum because neutrinos undergo coherent forward scattering with particles in the medium. Although forward scattering does not change the neutrino momentum, it modifies the phase accumulated during propagation and can therefore affect flavor oscillations.

All three active neutrino flavors experience neutral-current interactions with the particles in ordinary matter. These interactions contribute the same effective potential to each

active flavor and therefore do not change their relative phases. Electron neutrinos also undergo charged-current forward scattering with electrons through W-boson exchange. This interaction introduces an additional potential for the electron flavor,

$$V_{\mathrm{CC}} = \sqrt{2} G_F N_e, \tag{2.31}$$

where G_F is the Fermi constant and N_e is the electron number density [11].

The evolution of the neutrino flavor state is governed by the Schrödinger equation,

$$i \frac{d}{dx} |\nu(x)\rangle = H(x) |\nu(x)\rangle. \tag{2.32}$$

In the flavor basis, the effective Hamiltonian is

$$H = \frac{1}{2E_\nu} U \begin{pmatrix} m_1^2 & 0 & 0 \\ 0 & m_2^2 & 0 \\ 0 & 0 & m_3^2 \end{pmatrix} U^\dagger + \begin{pmatrix} V_{\mathrm{CC}} & 0 & 0 \\ 0 & 0 & 0 \\ 0 & 0 & 0 \end{pmatrix}. \tag{2.33}$$

The neutral-current contribution, proportional to the identity matrix for three active flavors, has been omitted.

The effect of the additional potential can be illustrated using a two-flavor approximation with constant electron density. The effective parameters and the probability derived below apply to this constant-density limit. After removing a term proportional to the identity matrix, the effective Hamiltonian can be written as

$$H_{\mathrm{2f}} = \frac{\Delta m^2}{4E_\nu} \begin{pmatrix} -\cos 2\theta & \sin 2\theta \\ \sin 2\theta & \cos 2\theta \end{pmatrix} + \frac{V_{\mathrm{CC}}}{2} \begin{pmatrix} 1 & 0 \\ 0 & -1 \end{pmatrix}. \tag{2.34}$$

The vacuum mixing angle θ and mass-squared splitting Δm^2 are replaced by effective

quantities in matter. Defining

$$A = 2E_\nu V_{\rm CC}, \tag{2.35}$$

the effective mixing angle satisfies

$$\sin^2(2\theta_m) = \frac{(\Delta m^2 \sin 2\theta)^2}{(\Delta m^2 \cos 2\theta - A)^2 + (\Delta m^2 \sin 2\theta)^2}. \tag{2.36}$$

The corresponding effective mass-squared splitting is

$$\Delta m^2_m = \sqrt{(\Delta m^2 \cos 2\theta - A)^2 + (\Delta m^2 \sin 2\theta)^2}. \tag{2.37}$$

For constant density, the oscillation probability retains the two-flavor form,

$$P^{\rm matter}_{\alpha\beta} = \sin^2(2\theta_m) \sin^2\left(\frac{\Delta m^2_m L}{4E_\nu}\right), \qquad \alpha \neq \beta. \tag{2.38}$$

These expressions show that matter modifies both the effective mixing amplitude and the oscillation frequency. When

$$A = \Delta m^2 \cos 2\theta, \tag{2.39}$$

the effective mixing becomes maximal, $\sin^2(2\theta_m) = 1$, even if the vacuum mixing angle is not maximal. This condition defines the Mikheyev–Smirnov–Wolfenstein (MSW) resonance [12, 13]. For neutrinos in ordinary matter, the resonance condition requires $\Delta m^2 \cos 2\theta > 0$; for antineutrinos, the matter potential changes sign and the corresponding condition is reversed.

For antineutrinos, the charged-current matter potential changes sign, $V_{\rm CC} \to -V_{\rm CC}$, and the mixing matrix is complex conjugated. The resonance condition therefore depends on the sign of the mass-squared splitting and on whether neutrinos or antineutrinos are considered.

In three-flavor oscillations, the atmospheric-scale matter contribution is often characterized by the dimensionless parameter

$$\widehat{A} = \frac{2E_\nu V_{\rm CC}}{\Delta m^2_{31}}. \tag{2.40}$$

For a positive electron density, $\widehat{A}$ has the sign of Δm^2_{31} for neutrinos; the effective matter potential reverses sign for antineutrinos. Matter effects can enhance or suppress electron-neutrino appearance depending on the mass ordering and beam polarity.

Unlike the intrinsic CP-odd term, the matter-induced neutrino–antineutrino difference can persist at $\delta_{\mathrm{CP}} = 0$ or π.

The size of the matter effect depends on neutrino energy, electron density, and propagation distance. It is particularly important for the longer baseline of DUNE, where the accumulated matter contribution is larger than in NOvA. Numerical three-flavor calculations use an appropriate matter-density profile along the neutrino trajectory.

2.6 From Oscillation Probabilities to Parameter Inference

Oscillation probabilities describe flavor transitions as functions of neutrino energy and propagation distance. Experimental analyses instead compare reconstructed event spectra with predictions obtained by combining these probabilities with the beam flux, interaction models, event selection, and detector response. The relationship between true neutrino energy and reconstructed energy is particularly important near oscillation features, where energy resolution and calibration affect the observed spectral shape.

The predicted spectra are fitted to data to estimate oscillation parameters, including Δm^2_{32}, θ_{23}, δ_{CP}, and the mass ordering. External measurements constrain other oscillation parameters, while nuisance parameters describe systematic uncertainties. Correlations and degeneracies can produce multiple solutions with similar likelihood values. The experimental framework used for these fits is described in Chapter 6; the construction of confidence regions is developed in Chapter 7.

Chapter 3

The NOvA Experiment

The NOvA experiment is a long-baseline accelerator neutrino experiment that studies ν_μ and $\bar{\nu}_\mu$ disappearance and ν_e and $\bar{\nu}_e$ appearance [4, 5]. NOvA uses the NuMI neutrino beam at Fermilab and two segmented liquid-scintillator detectors with similar active designs: a Near Detector located close to the beam source and a Far Detector located 810 km away in northern Minnesota.

The detectors are positioned approximately 14.6 mrad off the central NuMI beam axis. This off-axis configuration produces a relatively narrow neutrino-energy spectrum near the energy of the first oscillation maximum and reduces the high-energy tail of the beam. The Near Detector characterizes neutrino interactions close to the beam source, while the Far Detector measures the effects of neutrino oscillations over the 810-km baseline. This chapter describes the beam, detector, simulation, and reconstruction infrastructure relevant to the neutrino–electron scattering analysis.

3.1 The NuMI Neutrino Beam

3.1.1 Proton Beam, Target, and Secondary Hadrons

The Neutrinos at the Main Injector (NuMI) beam is produced at Fermilab by directing high-energy protons from the Main Injector onto a graphite target [14]. The proton interactions produce secondary hadrons, primarily charged pions and kaons. These particles are focused, transported through the decay region, and allowed to decay in flight. The dominant decay

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

produces the main ν_μ component in neutrino-mode running, with corresponding charge-conjugate decays contributing in antineutrino mode.

Kaon decays and secondary muon decays produce additional neutrino components, including intrinsic ν_e and $\bar{\nu}_e$. The resulting beam contains all four relevant electron- and muon-flavor neutrino and antineutrino components, although one muon-flavor component dominates in each focusing configuration. This flavor composition is important for the neutrino–electron analysis because all beam flavors can contribute to the elastic-scattering signal.

The neutrino flux prediction depends on the proton exposure, secondary-hadron production, focusing, transport, and decay. Hadronic interactions in the target and surrounding materials affect the number and momentum distributions of particles contributing to the beam. These processes are modeled in the beam simulation described in Sec. 3.4.1.

3.1.2 Magnetic Focusing and FHC/RHC Operation

The charged mesons emerging from the target have a broad distribution of momenta and production angles. Without focusing, only a small fraction would enter the decay pipe with trajectories capable of producing neutrinos directed toward the detectors. Two pulsed magnetic horns increase this acceptance by bending charged secondaries before they decay.

Current flowing along the horn conductors generates an approximately azimuthal magnetic field in the region between the conductors. A meson traversing this field receives a transverse momentum impulse whose direction depends on its charge and the horn-current polarity. The shape of the conductors and the meson's momentum determine how effectively its trajectory is redirected toward the decay pipe.

In forward horn current (FHC) mode, the magnetic horns focus positively charged mesons, producing a neutrino beam dominated by ν_μ from π^+ and K^+ decays. Reversing the horn current focuses negatively charged mesons, producing a beam dominated by $\bar{\nu}_\mu$ in reverse horn current (RHC) mode.

The magnetic horns do not focus all secondary mesons equally. Their focusing efficiency depends on the momentum and direction of the particles emerging from the target, and mesons of the opposite charge can still contribute to the neutrino beam. As a result, both running modes contain wrong-sign muon-flavor neutrinos and smaller intrinsic electron-flavor components. The focusing conditions also affect the neutrino-energy spectrum. Uncertainties associated with the horn current, target position, and beamline alignment are incorporated into the beam-transport uncertainties of the flux prediction.

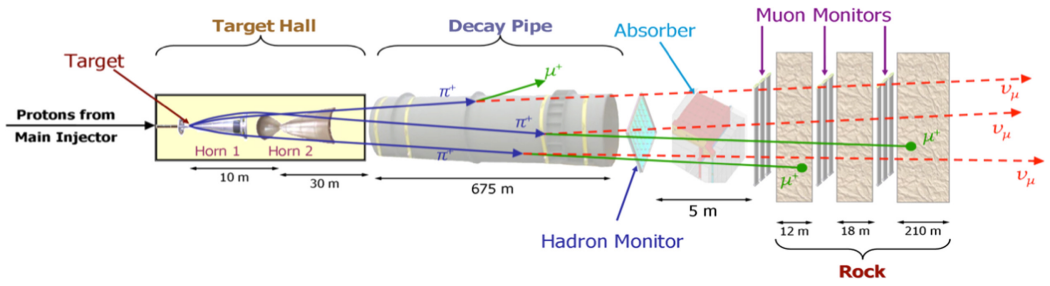

Figure 3.1: Schematic of the NuMI neutrino beamline, showing the graphite target, magnetic focusing horns, decay region, and downstream absorber [14, 15].

3.1.3 Off-Axis Concept and the NOvA Flux

The off-axis technique takes advantage of the kinematics of two-body pion decay. In the pion rest frame, the neutrino has a fixed energy, while in the laboratory frame, its energy

depends on the parent-pion energy and the emission angle. For the decays $\pi^+ \to \mu^+\nu_\mu$ and $\pi^- \to \mu^-\bar{\nu}_\mu$, the ultrarelativistic, small-angle approximation gives

$$E_\nu \simeq \frac{(1 - m_\mu^2/m_\pi^2)E_\pi}{1 + \gamma^2\theta^2} \simeq \frac{0.43E_\pi}{1 + \gamma^2\theta^2}, \tag{3.2}$$

where $\gamma = E_\pi/m_\pi$ is the Lorentz factor of the parent pion, and θ is the angle between the neutrino and the parent-pion directions in the laboratory frame.

For neutrinos emitted along the parent-pion direction ($\theta \simeq 0$), the neutrino energy increases approximately linearly with the pion energy. At a fixed nonzero emission angle, the denominator grows with increasing pion energy, reducing the dependence of E_ν on E_π over a range of parent-pion energies. This kinematic behavior allows an off-axis detector to receive neutrinos within a relatively narrow energy range, even when the parent pions have a broad energy distribution.

In a realistic neutrino beam, the parent pions do not all travel parallel to the nominal beam axis. Their directions and momenta depend on hadron production and magnetic focusing, and the neutrinos are emitted at different angles relative to their parent particles. The energy spectrum observed at a given off-axis location depends on these distributions, together with contributions from kaon decays. A complete flux prediction requires a simulation of hadron production, transport, and decay.

The NOvA detectors are positioned approximately 14.6 mrad off the NuMI beam axis. This configuration produces a neutrino energy spectrum concentrated near 2 GeV, close to the first ν_μ disappearance maximum for the 810 km baseline. The reduced high-energy component also limits contributions from higher-energy neutrino interactions that can enter the oscillation analysis as backgrounds.

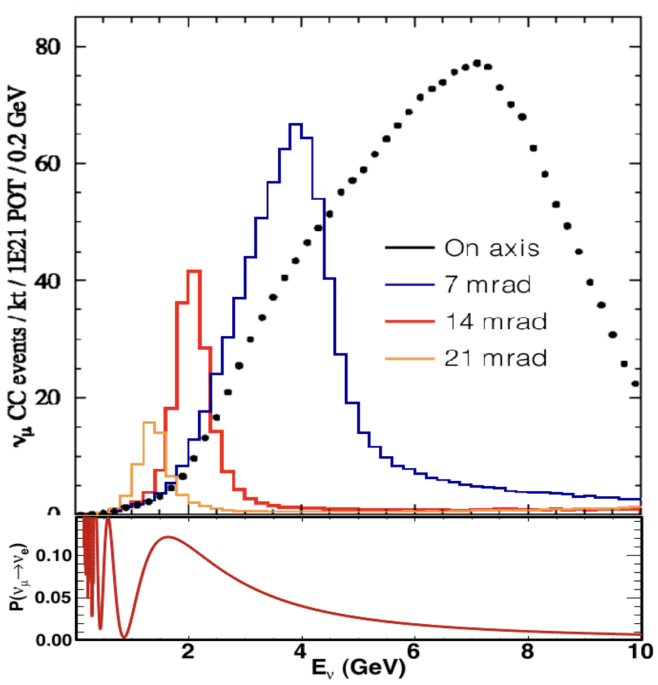

Figure 3.2: Simulated NuMI neutrino-energy spectra at several off-axis angles, illustrating the narrower energy distribution obtained away from the nominal beam axis [15].

3.2 The NOvA Detectors

NOvA uses two functionally identical tracking calorimeters so that detector-response effects can be correlated between the near and far sites. Both detectors are constructed from planes of extruded PVC cells filled with liquid scintillator. Alternating planes are oriented horizontally and vertically, giving two orthogonal views of each interaction and allowing three-dimensional event reconstruction.

3.2.1 Detector Technology

The active detector is assembled from extruded PVC cells filled with liquid scintillator. The PVC provides a mechanically stable, optically separated structure, while the scintillator converts a charged particle's energy loss into detectable light. Each cell contains a wavelength-shifting fiber arranged in a loop. The fiber absorbs part of the scintillation light, re-emits it at longer wavelengths, and guides the shifted light by total internal reflection to an avalanche

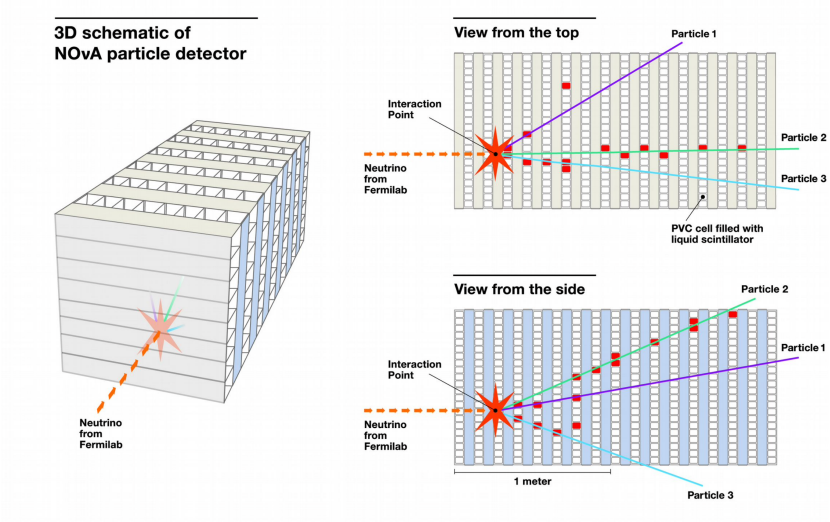

Figure 3.3: Schematic of the NOvA detector geometry. Alternating planes of horizontal and vertical cells provide the two views used to reconstruct the three-dimensional event topology [15].

photodiode (APD). Both ends of the loop are brought to the readout end of the cell, so photons captured in either direction can contribute to the signal. The resulting light yield still depends on the distance traveled along the fiber, which is addressed by the calibration procedure.

Successive planes have alternating horizontal and vertical cell orientations. A single plane localizes an energy deposit along one transverse coordinate and the beam direction, but does not determine the position along the cell. Combining hits from the two orthogonal views allows three-dimensional particle trajectories and shower axes to be reconstructed using spatial and timing correlations. This arrangement provides tracking and calorimetric information in the same volume, with sampling determined by the cell dimensions and plane spacing.

3.2.2 Near Detector

The NOvA Near Detector is located approximately 1 km downstream of the NuMI target and is positioned off the central beam axis [4]. It is a roughly 0.3 kton segmented liquid-scintillator

detector with transverse dimensions of about 4.2 m by 4.2 m and a length of about 15.8 m.

Because of its proximity to the beam source, the Near Detector records a large number of neutrino interactions. These data are used to characterize the unoscillated beam, study neutrino interactions, constrain oscillation predictions, and perform dedicated measurements. The neutrino–electron elastic-scattering measurement in this dissertation uses FHC Near Detector data.

3.2.3 Far Detector

The NOvA Far Detector is located near Ash River, Minnesota, 810 km from Fermilab and at approximately the same off-axis angle as the Near Detector. It has a mass of approximately 14 kton and uses the same basic active detector technology as the Near Detector [4].

The Far Detector is much larger because the neutrino interaction rate is low after the beam has traveled to Minnesota. It is also on the surface, so cosmic-ray rejection is an important part of the event selection in the oscillation analyses. Timing with the NuMI beam, event topology, containment, and dedicated classifiers are used to separate beam-neutrino candidates from this background.

The Far Detector records the energy-dependent depletion of muon-flavor neutrinos and the appearance of electron-flavor neutrinos. Its measurements are interpreted together with the Near Detector prediction, as outlined in Sec. 3.7.

3.3 Data Acquisition and Detector Response

3.3.1 Signal Readout

An APD converts the light emerging from its associated wavelength-shifting fibers into an electrical pulse. Front-end electronics amplify and shape the signal before digitization. The pulse height depends on the number of detected photons, while its time profile also reflects

scintillation emission, light propagation, and the response of the electronics. Signal shaping affects the pulse amplitude and timing, with longer shaping times producing broader pulses.

Channel-level processing estimates the electronic baseline and identifies pulses that satisfy the hit-finding requirements. The resulting records carry channel, time, and signal-size information and are then used for calibration and event reconstruction. Readout thresholds can cause small signals to fall below the detection threshold, while electronic noise affects the identification and measurement of pulses. These effects are especially relevant for low-energy deposits near the edges of electromagnetic showers.

3.3.2 Light Collection and Calibration

The response of a scintillator cell depends on the position of the energy deposit. Light produced farther from the readout end travels a longer distance through the wavelength-shifting fiber and experiences greater attenuation. Because both ends of the fiber connect to the APD, the collected light includes contributions from two propagation paths with different lengths. Differences in scintillator light yield, APD gain, and front-end electronics also contribute to variations among channels.

Track samples allow the detector response to be measured at different positions along the cells. Relative calibration corrects for light attenuation and channel-to-channel variations, bringing the measured signals onto a common scale. Stopping muons provide a reference for absolute energy calibration because their energy loss along the track can be compared with the expected stopping power in the detector material. This comparison establishes the absolute energy scale for the calibrated detector response. The same calibration procedure is used for data and simulated events, with residual differences in detector response accounted for through calibration uncertainties.

Calibration uncertainties can affect both the reconstructed energy and the spatial distribution of hits. An overall energy-scale shift changes the reconstructed shower energy, while variations in the response along individual cells can affect the measured energy of

individual hits and the reconstructed shower shape. The impact of these uncertainties on the neutrino–electron scattering measurement is discussed in the systematic-uncertainty section of Chapter 5.

3.4 Simulation

The NOvA simulation chain models the production of the neutrino beam, neutrino interactions, transport of the final-state particles through the detector, optical and electronic detector response, and the reconstruction of the simulated detector activity. Data and simulation are processed with the same reconstruction and particle-identification algorithms.

3.4.1 Beam Simulation and PPFX

The NuMI beamline is simulated using the G4NuMI framework based on GEANT4 [16]. The simulation begins with the incident proton beam and follows secondary hadron production and transport through the target, focusing horns, decay pipe, and surrounding beamline materials. The decays of the secondary particles determine the neutrino flavor, energy, and direction.

The hadron yields predicted by the beam simulation are corrected using the Package to Predict the Flux (PPFX) [17]. PPFX incorporates external hadron-production measurements and propagates the corresponding uncertainties to the neutrino flux. The hadron-production uncertainties are represented by 100 correlated flux universes, each corresponding to a coherent variation of the hadron-production model. These variations affect the predicted flux across neutrino flavors and energy bins while preserving the correlations among them.

Uncertainties associated with beam focusing and geometry are evaluated separately from PPFX hadron-production uncertainties. These include variations in the horn current, proton-beam position and profile, horn and target alignment, magnetic-field modeling, and beamline material. The beam-transport variations account for changes in the focusing and propagation

of secondary hadrons that are not described by the PPFX hadron-production uncertainties.

3.4.2 Neutrino Interactions and GENIE

Neutrino interactions in the detector are generated with GENIE 3.0.6 [18, 19] using the NOvA production configuration. The generator uses the incident neutrino energy and flavor to simulate the interaction vertex and outgoing particles. At NOvA energies, the relevant neutrino–nucleus processes include quasielastic scattering, resonance production, multinucleon interactions, and deep-inelastic scattering. The predicted final state depends on both the primary interaction and the nuclear environment.

The initial nuclear state is described with a Local Fermi Gas model. Quasielastic scattering uses the Valencia model with a z-expansion axial form factor [20, 21], and the two-particle–two-hole component uses the Valencia meson-exchange-current model [22, 23]. Resonance production is modeled with the Berger–Sehgal model [24], deep-inelastic scattering with the Bodek–Yang model [25], and final-state interactions with the GENIE hN intranuclear cascade [26, 27].

Particles produced at the primary interaction vertex may undergo additional interactions before leaving the nucleus. These final-state interactions can absorb or scatter hadrons and change the visible final state. Nuclear binding, undetected neutrons, and particles below the detection threshold can also affect the reconstructed hadronic energy. Variations in the interaction models influence predicted event rates and the relationship between true and reconstructed energy. NOvA applies additional central-value modifications to improve agreement with Near Detector data [28], with associated uncertainties included in the relevant analyses.

Neutrino–electron elastic scattering occurs on atomic electrons rather than nuclei and produces no hadronic system at the interaction vertex. The signal is generated at tree level in GENIE, with electroweak radiative corrections applied separately in the neutrino–electron analysis, as described in Sec. 4.4.

3.4.3 Detector Simulation

The particles emerging from GENIE are propagated through the detector geometry with GEANT4 [16]. Charged particles lose energy through ionization and may undergo scattering or secondary interactions. Electrons and photons produce electromagnetic showers through processes such as bremsstrahlung and pair production.

The simulation records the energy deposited by particles traversing the detector and models the corresponding optical and electronic response. This includes scintillation light production, light collection, attenuation in the wavelength-shifting fibers, photosensor response, pulse shaping, and electronic noise. The simulated channel signals undergo digitization and follow the same reconstruction chain as detector data.

The reconstructed energy and particle trajectories depend on the detector response. Energy deposited in inactive material, signals below the readout threshold, and incomplete collection of shower hits can affect the reconstructed calorimetric energy. Multiple scattering and finite detector segmentation also influence the reconstructed particle direction. The detector simulation models these effects and provides the response needed to evaluate reconstruction efficiency and resolution.

3.5 Event Reconstruction

Event reconstruction transforms calibrated cell hits into candidate interactions and particle objects. The processing begins by grouping spatially and temporally associated hits, then estimates interaction vertices and identifies tracks and showers. These reconstructed objects provide the energies, directions, and event topologies used in physics analyses.

3.5.1 Hits and Slices

Reconstruction begins with calibrated hits that carry position, time, and visible-energy information. Hits originating from the same interaction are generally correlated in space and

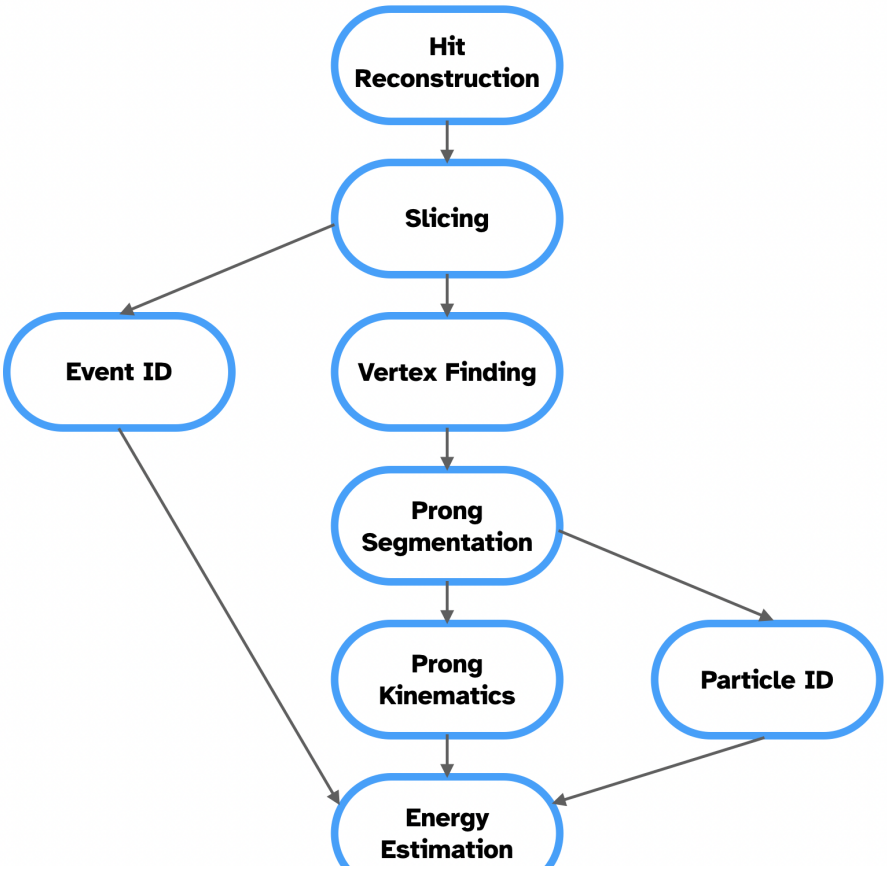

Figure 3.4: Overview of the NOvA reconstruction chain, from calibrated detector hits to reconstructed slices, vertices, prongs, and particle objects. [15]

time. The slicing algorithm uses these correlations to group hits into candidate interactions. Information from the two orthogonal detector views helps associate activity across the detector planes.

At the Near Detector, multiple neutrino interactions can occur within the same readout interval. Spatial and timing information allows the slicing algorithm to separate their detector activity into individual candidates.

3.5.2 Vertex and Prong Reconstruction

After slicing, reconstruction estimates the interaction vertex and identifies the outgoing particle activity. Straight-line structures in each two-dimensional detector view can be found using methods such as the Hough transform. These structures provide initial directions for vertex reconstruction.

The Elastic Arms algorithm fits a set of arms extending from a common vertex to the observed hit distribution. The fit adjusts the vertex position and arm directions to describe the spatial pattern of hits [29].

Information from the two detector views is combined to reconstruct the three-dimensional interaction vertex. Reconstruction then groups hits into particle candidates, or *prongs*, according to their spatial continuity and association with the vertex. Each prong represents a reconstructed particle candidate with an estimated direction, visible energy, and spatial topology.

3.5.3 Tracks and Showers

Different particles produce characteristic patterns of energy deposition in the detector. Muons typically leave long, narrow tracks, while electrons develop electromagnetic showers that spread across multiple cells and planes. Photons can produce similar electromagnetic showers after converting into electron–positron pairs.

Track reconstruction uses the spatial distribution of hits to estimate particle trajectories.

For electromagnetic showers, reconstruction estimates the direction from the shower development, with the initial portion providing the most useful information about its axis. The visible energy is estimated from the calibrated energy deposits associated with the shower.

In the neutrino–electron scattering analysis, the reconstructed energy and direction of the leading electron-like shower are used to determine the kinematic observables for event selection.

3.6 Event Classification

Particle identification in NOvA relies on the spatial patterns of energy deposition produced by different neutrino interactions. Muon tracks, electromagnetic showers, and hadronic activity exhibit different topologies in the segmented detector. NOvA uses the Convolutional Visual Network (CVN) to classify these patterns and identify neutrino interaction types [30].

CVN takes two-dimensional pixel maps of detector activity from the two orthogonal views as input. Convolutional layers extract spatial features from each view, including the shape and extent of tracks and showers. The network combines these features to produce classification scores for different interaction categories. These scores provide additional information for event selection alongside reconstructed kinematic quantities.

The neutrino–electron scattering analysis uses dedicated classifiers based on this approach. Their training and application to event selection are described in Chapter 5.

3.7 NOvA Oscillation Program and Flux Systematics

NOvA uses Near Detector data to improve predictions of the oscillated neutrino spectrum at the Far Detector. The Near-to-Far extrapolation compares the observed and simulated energy spectra at the Near Detector and uses the simulated relationship between true and reconstructed neutrino energy to predict the Far Detector spectrum. The similar detector designs help reduce uncertainties common to both detectors. The prediction also incorporates

differences in the neutrino flux, interaction kinematics, and detector acceptance between the two sites.

Hadron-production and beam-focusing uncertainties affect the neutrino flux as a function of energy and flavor. The PPFX ensemble described in Sec. 3.4.1 represents correlated variations in the hadron-production component of the flux prediction. Neutrino–electron elastic scattering provides an additional way to constrain the flux because its cross section is precisely predicted from electroweak theory, without the nuclear-model uncertainties associated with neutrino–nucleus interactions. The observed electron-energy spectrum, together with the estimated backgrounds and detector acceptance, can be used to constrain the predicted neutrino flux. The analysis presented in Chapter 5 uses a Bayesian reweighting method to update the PPFX flux predictions based on the selected neutrino–electron sample.

Chapter 4

Neutrino–Electron Elastic Scattering as a Neutrino Flux Standard Candle

Neutrino–electron elastic scattering is a useful process for constraining accelerator neutrino fluxes. As a purely leptonic process, its cross section can be calculated precisely within the Standard Model without uncertainties arising from nuclear interactions. The two-body kinematics produce electrons that emerge at small angles relative to the incident neutrino direction. These properties allow the observed electron-energy spectrum to constrain the neutrino flux [31, 32].

Neutrinos and antineutrinos of all three flavors can scatter elastically from electrons through neutral-current interactions mediated by Z-boson exchange. Electron neutrinos and antineutrinos can also interact through charged-current W-boson exchange. The charged-current and neutral-current amplitudes interfere, modifying the scattering cross sections relative to those of muon and tau neutrinos and antineutrinos.

The relevant scattering channels are

$$\nu_e e^- \to \nu_e e^- \qquad \text{CC + NC}, \tag{4.1}$$

$$\bar{\nu}_e e^- \to \bar{\nu}_e e^- \qquad \text{CC + NC}, \tag{4.2}$$

$$\nu_{\mu,\tau} e^- \to \nu_{\mu,\tau} e^- \qquad \text{NC}, \tag{4.3}$$

$$\bar{\nu}_{\mu,\tau} e^- \to \bar{\nu}_{\mu,\tau} e^- \qquad \text{NC}. \tag{4.4}$$

The corresponding tree-level diagrams are shown in Figure 4.1.

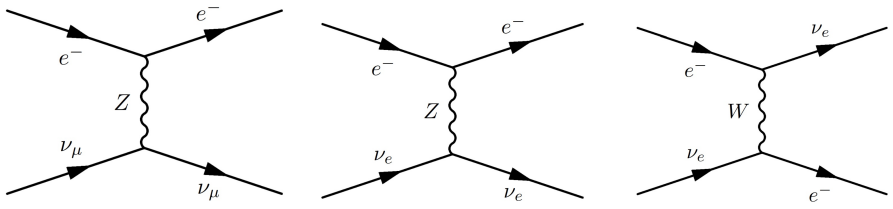

Figure 4.1: Tree-level Feynman diagrams for neutrino–electron elastic scattering through neutral-current (Z-boson) and charged-current (W-boson) exchange.

4.1 Tree-Level Cross Section

For an electron initially at rest, the differential cross section for neutrino–electron elastic scattering can be written in terms of the recoil-electron kinetic energy T_e and the inelasticity $y = T_e/E_\nu$ as [33, 32]

$$\frac{d\sigma(\nu_\ell e^- \to \nu_\ell e^-)}{dy} = \frac{2G_F^2 m_e E_\nu}{\pi}\left[C_{LL}^2 + C_{LR}^2(1-y)^2 - C_{LL}C_{LR}\frac{m_e y}{E_\nu}\right], \tag{4.5}$$

where G_F is the Fermi constant, m_e is the electron mass, and E_ν is the incident neutrino energy. The electroweak couplings at tree level are

$$C_{LL}^{\nu_\ell e} = -\frac{1}{2} + \sin^2\theta_W + \delta_{\ell e}, \qquad C_{LR} = \sin^2\theta_W, \tag{4.6}$$

where θ_W is the weak mixing angle and $\delta_{\ell e} = 1$ for electron neutrinos and zero for muon and tau neutrinos. The charged-current amplitude enters through C_{LL}, including its interference with the neutral-current amplitude.

For antineutrino–electron scattering, the differential cross section takes the same form, with the terms proportional to C_{LL}^2 and C_{LR}^2 interchanged.

4.2 Two-Body Kinematics

In the laboratory frame, the initial electron is assumed to be at rest. The four-momenta of the incoming neutrino, initial electron, outgoing neutrino, and recoil electron are defined as

$$\begin{aligned} p_1 &= (E_\nu, \boldsymbol{p}_\nu), \quad p_2 = (m_e, \mathbf{0}), \\ p_3 &= (E'_\nu, \boldsymbol{p}'_\nu), \quad p_4 = (E_e, \boldsymbol{p}_e). \end{aligned} \tag{4.7}$$

Here, E_e is the total energy of the recoil electron, and θ_e is its scattering angle with respect to the incident neutrino direction.

The Mandelstam variables for this process are

$$s = (p_1 + p_2)^2 = m_e^2 + 2m_e E_\nu, \tag{4.8}$$

$$t = (p_2 - p_4)^2 = -2m_e T_e, \tag{4.9}$$

$$\begin{aligned} u &= (p_1 - p_4)^2 \\ &= m_e^2 - 2E_\nu E_e + 2E_\nu \sqrt{E_e^2 - m_e^2}\cos\theta_e, \end{aligned} \tag{4.10}$$

where $T_e = E_e - m_e$ is the recoil-electron kinetic energy. Energy and momentum conservation give

$$s + t + u = 2m_e^2. \tag{4.11}$$

Substituting the expressions for s, t, and u yields the exact relation

$$\cos\theta_e = \frac{E_\nu + m_e}{E_\nu}\sqrt{\frac{T_e}{T_e + 2m_e}}. \tag{4.12}$$

For $E_\nu, E_e \gg m_e$, the recoil-angle relation can be approximated as

$$\cos\theta_e \simeq 1 - \frac{m_e}{E_e}(1 - y). \tag{4.13}$$

Using $\cos\theta_e \simeq 1 - \theta_e^2/2$ for small scattering angles gives

$$E_e\theta_e^2 \simeq 2m_e(1 - y), \tag{4.14}$$

where $y = T_e/E_\nu$ was defined in Section 4.1. Since $0 \leq y \leq 1$, the kinematic upper bound is approximately

$$E_e\theta_e^2 \lesssim 2m_e. \tag{4.15}$$

The small values of $E_e\theta_e^2$ expected for neutrino–electron elastic scattering provide a useful signature for separating signal events from neutrino–nucleus backgrounds. Figure 4.2 shows the true $E_e\theta_e^2$ distribution for simulated neutrino–electron elastic-scattering events.

4.3 Relation Between the Event Spectrum and the Flux

The neutrino–electron elastic-scattering rate depends on the incident neutrino flux and the interaction cross section. For a given exposure, the total number of scattering events is proportional to

$$N_{\nu e} \propto \sum_\alpha \int \Phi_\alpha(E_\nu)\,\sigma_\alpha(E_\nu)\,dE_\nu, \tag{4.16}$$

where $\Phi_\alpha(E_\nu)$ is the exposure-integrated neutrino fluence for flavor α, and $\sigma_\alpha(E_\nu)$ is the corresponding neutrino–electron elastic-scattering cross section.

The outgoing electron carries only part of the incident neutrino energy, so electrons with

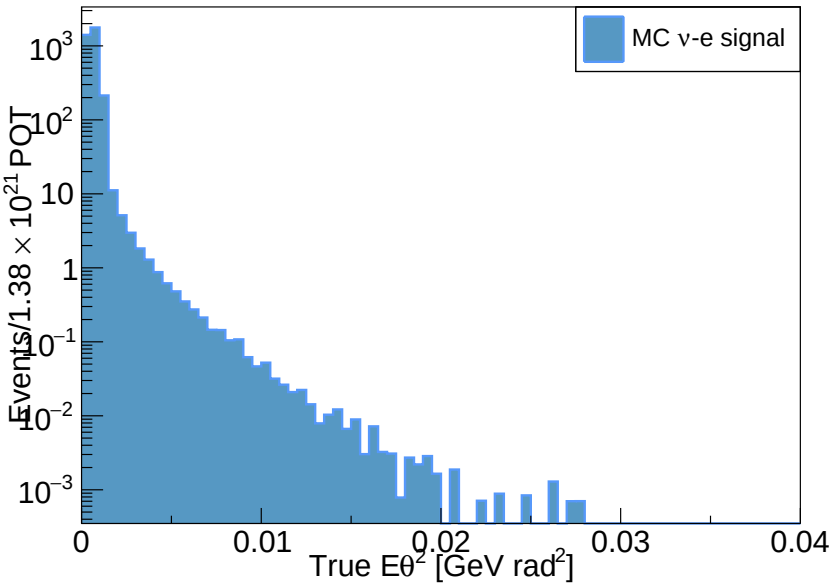

Figure 4.2: True $E_e\theta_e^2$ distribution for simulated neutrino–electron elastic-scattering events.

the same energy can originate from neutrinos of different energies. The recoil-electron energy spectrum is therefore related to the incident neutrino flux through

$$\frac{dN_{\nu e}}{dE_e} \propto \sum_{\alpha} \int \Phi_{\alpha}(E_\nu) \frac{d\sigma_{\alpha}(E_\nu, E_e)}{dE_e}\, dE_\nu. \tag{4.17}$$

Since the neutrino–electron scattering cross section is precisely predicted by the Standard Model, the measured electron-energy spectrum can be used to constrain the neutrino flux. The MINERvA experiment has applied this method to constrain the NuMI neutrino flux [31, 32, 34]. MINERvA reported a reduction in the ν_μ flux-normalization uncertainty from 7.6% to 3.3% and in the $\bar{\nu}_\mu$ uncertainty from 7.8% to 4.7%.

4.4 Radiative Corrections

The GENIE signal simulation uses a tree-level neutrino–electron scattering cross section. The cross section is corrected at the one-loop level, including electroweak and QED radiative corrections. Real-photon emission in $\nu_\ell e^- \to \nu_\ell e^- \gamma$ also contributes to the electromagnetic

final state. These corrections affect the predicted cross section and the electromagnetic-energy spectrum. The correction applied in this analysis follows the prescription used in the MINERvA neutrino–electron measurement, with radiative functions taken from Refs. [32, 35].

Table 4.1 lists the electroweak couplings used in the GENIE tree-level prediction and the radiatively corrected calculation. The one-loop values include electroweak corrections to the tree-level couplings in Eq. 4.6.

Table 4.1: Electroweak couplings in the GENIE tree-level prediction and the radiative-correction calculation used in this analysis.

	$C_{LL}^{\nu_e e}$	$C_{LL}^{\nu_\mu e}$	C_{LR}
GENIE (tree level)	0.7313	−0.2687	0.2313
one-loop	0.7276	−0.2730	0.2334

For the radiative-correction calculation, the electromagnetic-energy fraction is defined as

$$y_{\mathrm{EM}} = \frac{T_e + E_\gamma}{E_\nu}, \tag{4.18}$$

where E_γ is the energy of an emitted photon. This differs from the tree-level definition $y = T_e/E_\nu$ when a photon is emitted. The corresponding total electromagnetic energy in Ref. [35] includes the electron rest mass; the difference is negligible in the high-energy approximation used here. The reconstructed shower energy is determined separately by detector simulation and reconstruction.

Using $K = 2G_F^2 m_e E_\nu/\pi$ and the fine-structure constant α_{EM}, the radiatively corrected neutrino cross section is written as

$$\begin{aligned}\frac{d\sigma^{\mathrm{rad}}(\nu_\ell e^-)}{dy_{\mathrm{EM}}} = K\Bigg[&(C_{LL}^{\nu_\ell e})^2\left(1+\frac{\alpha_{\mathrm{EM}}}{\pi}X_1\right)\\ &+(C_{LR})^2(1-y_{\mathrm{EM}})^2\left(1+\frac{\alpha_{\mathrm{EM}}}{\pi}X_2\right)\\ &-C_{LL}^{\nu_\ell e}C_{LR}\frac{m_e y_{\mathrm{EM}}}{E_\nu}\left(1+\frac{\alpha_{\mathrm{EM}}}{\pi}X_3\right)\Bigg].\end{aligned} \tag{4.19}$$

For antineutrinos, the kinematic factors multiplying the two squared-coupling terms are

exchanged:

$$\frac{d\sigma^{\mathrm{rad}}(\bar{\nu}_\ell e^-)}{dy_{\mathrm{EM}}} = K\Bigg[(C_{LR})^2\left(1+\frac{\alpha_{\mathrm{EM}}}{\pi}X_1\right) + (C_{LL}^{\nu_\ell e})^2(1-y_{\mathrm{EM}})^2\left(1+\frac{\alpha_{\mathrm{EM}}}{\pi}X_2\right) - C_{LL}^{\nu_\ell e}C_{LR}\frac{m_e y_{\mathrm{EM}}}{E_\nu}\left(1+\frac{\alpha_{\mathrm{EM}}}{\pi}X_3\right)\Bigg]. \tag{4.20}$$

For the photon-inclusive electromagnetic-energy spectrum defined in Eq. 4.18, the high-energy radiative functions are [35]

$$X_1(y_{\mathrm{EM}}, E_\nu) = \frac{3y_{\mathrm{EM}}^2 - 30y_{\mathrm{EM}} + 23}{72} - \frac{2}{3}\ln\left(\frac{2y_{\mathrm{EM}}E_\nu}{m_e}\right) - \frac{\pi^2}{6}, \tag{4.21}$$

$$X_2(y_{\mathrm{EM}}, E_\nu) = \frac{-4y_{\mathrm{EM}}^2 - 16y_{\mathrm{EM}} + 23}{72(1-y_{\mathrm{EM}})^2} - \frac{2}{3}\ln\left(\frac{2y_{\mathrm{EM}}E_\nu}{m_e}\right) - \frac{\pi^2}{6}, \tag{4.22}$$

$$X_3(y_{\mathrm{EM}}, E_\nu) = \frac{y_{\mathrm{EM}}^2 + 3y_{\mathrm{EM}} - 3}{4y_{\mathrm{EM}}^2} - \frac{3}{2}\ln\left(\frac{2y_{\mathrm{EM}}E_\nu}{m_e}\right) - \frac{\pi^2}{6}. \tag{4.23}$$

These correspond to the functions f_L, f_R, and f_{LR} in Eqs. (126)–(128) of Ref. [35]. They describe the QED correction to the electromagnetic-energy spectrum in the stated high-energy approximation. The electroweak coupling shifts are accounted for separately through the values in Table 4.1.

The radiative-correction weights are calculated using the one-loop electroweak couplings and the X_1 and X_2 radiative functions. Following the approximation used by MINERvA [32], the X_3 correction is omitted because the corresponding interference term is suppressed by m_e/E_ν at GeV-scale neutrino energies. The interference term itself is retained in the cross-section calculation. The weights are defined as the ratio of the corrected differential cross section to the GENIE tree-level prediction,

$$w_{\mathrm{rad}}(\nu_\ell e) = \frac{(d\sigma/dy_{\mathrm{EM}})_{\mathrm{corrected}}}{(d\sigma/dy_{\mathrm{EM}})_{\mathrm{GENIE}}}, \tag{4.24}$$

with the analogous ratio for antineutrinos. The weights are evaluated separately for ν_μ, $\bar{\nu}_\mu$, ν_e, and $\bar{\nu}_e$, as shown in Figure 4.3.

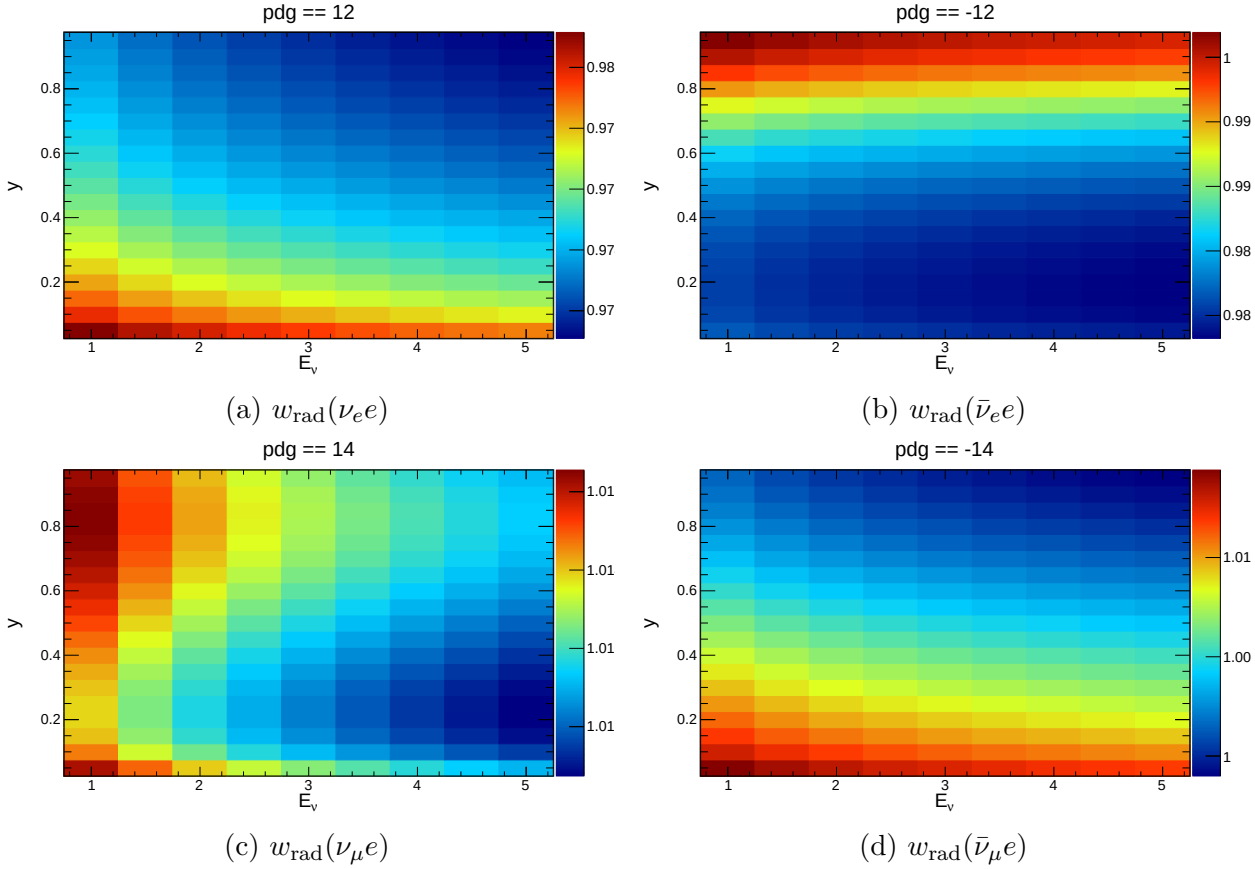

(a) $w_{\text{rad}}(\nu_e e)$

(b) $w_{\text{rad}}(\bar{\nu}_e e)$

(c) $w_{\text{rad}}(\nu_\mu e)$

(d) $w_{\text{rad}}(\bar{\nu}_\mu e)$

Figure 4.3: Radiative-correction weights as functions of E_ν and y_{EM} for the four beam components contributing to the signal.

Chapter 5

NOvA Neutrino–Electron Elastic-Scattering Analysis and Flux Constraint

This chapter presents the NOvA Near Detector neutrino–electron elastic-scattering analysis and its application to the NuMI flux constraint. The background-subtracted reconstructed electron-energy spectrum is used to reweight the standard NOvA PPFX flux ensemble.

The data were collected in neutrino mode between 2014 and 2024, corresponding to NOvA data-taking periods P1–P11a and a total exposure of

$$1.38 \times 10^{21} \text{ POT}. \tag{5.1}$$

The data sample satisfies the standard NOvA data-quality and good-run requirements.

5.1 Signal and Background Simulation

The inclusive Near Detector Monte Carlo sample contains charged-current and neutral-current neutrino–nucleus interactions, neutrino–electron elastic-scattering interactions, and neutrino

interactions in the surrounding rock. The sample is used primarily to model the neutrino–nucleus backgrounds, including ν_e charged-current, ν_μ charged-current, and neutral-current interactions. The inclusive simulation corresponds to an exposure of 5.54×10^{21} POT.

Neutrino–electron elastic scattering has a much smaller cross section than neutrino–nucleus charged-current interactions, resulting in limited signal statistics in the inclusive simulation. A high-statistics enhanced signal sample is used for signal selection and systematic studies.

The enhanced signal sample is produced by overlaying neutrino–electron elastic-scattering interactions on simulated Near Detector spills to account for the effects of additional detector activity on event reconstruction. Hits from different interactions may be grouped into the same reconstructed object, affecting the reconstructed shower properties and event selection. The enhanced signal sample corresponds to an effective exposure of 1.721×10^{24} POT and is normalized to the data exposure of 1.38×10^{21} POT.

The selection was also studied using neutrino–electron elastic-scattering events from the nominal inclusive simulation, which had been used during earlier stages of the analysis. The selection efficiencies obtained from the nominal and enhanced signal samples were consistent, with a small difference associated with the overlay treatment.

The signal interactions are generated at tree level in GENIE, with radiative-correction weights applied as described in Sec. 4.4. The correction increases the selected signal yield by approximately 0.83%.

All simulated events are processed through the standard NOvA detector simulation and reconstruction chain. The same reconstructed variables and event selections are used for data and simulation.

5.2 Event Selection

The selection targets a forward-going electromagnetic shower with little additional activity. The selection requirements were developed using Prod5.1 simulation and the enhanced signal

sample. All selection criteria were fixed before the signal-region data were unblinded.

The selection criteria were evaluated using the figure of merit

$$\mathrm{FOM} = \frac{S}{\sqrt{S + B + (\delta B)^2}}, \tag{5.2}$$

where S and B are the expected signal and background event counts from simulation, respectively, and δB represents the assumed background systematic uncertainty.

5.2.1 Preselection

The preselection removes events with unusually long prongs. The longest reconstructed prong is required to be shorter than 800 cm and span fewer than 120 detector planes. Each slice is limited to fewer than 600 cells. To reduce pileup and incorrect grouping of detector activity, the time separation between the candidate slice and its nearest neighboring slice is required to exceed 25 ns.

5.2.2 Fiducial Volume and Containment

The reconstructed interaction vertex is required to satisfy

$$-185 < x < 175 \text{ cm}, \tag{5.3}$$

$$-175 < y < 175 \text{ cm}, \tag{5.4}$$

$$95 < z < 1095 \text{ cm}. \tag{5.5}$$

Hits associated with reconstructed prongs are required to remain within

$$-190 < X < 180 \text{ cm}, \tag{5.6}$$

$$-180 < Y < 190 \text{ cm}, \tag{5.7}$$

$$105 < Z < 1275 \text{ cm}. \tag{5.8}$$

These requirements reduce backgrounds from neutrino interactions in the surrounding rock and events entering the detector from outside the active volume.

5.2.3 Single-Particle Requirements

The leading shower is required to contain more than 80% of the total reconstructed slice energy. Vertex activity is defined as the summed calibrated hit energy within ±8 detector planes of the reconstructed vertex, excluding hits assigned to the leading shower. The vertex activity is required to satisfy

$$E_{\text{vertex}} < 0.02 \text{ GeV}. \tag{5.9}$$

The distance between the reconstructed interaction vertex and the start of the leading shower is required to be less than 20 cm.

5.2.4 Electron-Energy Requirement

The reconstructed electron energy is required to satisfy

$$0.5 < E_e < 5 \text{ GeV}. \tag{5.10}$$

The lower threshold excludes the energy region where the classifiers have limited background-rejection performance. The upper threshold suppresses higher-energy electron-neutrino and antineutrino charged-current backgrounds.

5.3 Training of the CNN Classifiers

Two convolutional neural networks, NuoneID and Epi0ID, are used to distinguish neutrino–electron elastic-scattering events from neutrino–nucleus backgrounds. The classifiers use the NOvA convolutional neural network architecture, modified based on MobileNetV2. The architecture of NuoneID is shown in Fig. 5.1. Epi0ID uses the same network structure with

two output categories.

5.3.1 Training Samples and Output Categories

Both classifiers are trained using single-electron events as the signal sample and an electron-neutrino-enriched sample as the background. Single-electron events produce electromagnetic showers similar to those from neutrino–electron elastic scattering, but have more uniform distributions in electron energy and shower angle. Using these events for training reduces the risk of overfitting to the small scattering angles characteristic of neutrino–electron elastic scattering. The electron-neutrino-enriched sample provides additional ν_e charged-current events for background training.

NuoneID is trained with four output categories:

1. neutrino–electron signal,
2. π^0,
3. ν_e charged-current interactions,
4. other interactions, including ν_μ charged-current and neutral-current events.

Epi0ID uses the first two categories to distinguish electron showers from π^0 backgrounds.

Examples of detector pixel maps for the different training categories are shown in Fig. 5.2.

5.3.2 Training Preprocessing

A loose selection is applied to the training samples to focus the training on backgrounds that require CNN-based discrimination. The preprocessing includes the preselection requirements, a loose containment cut, and loose single-particle requirements.

The containment region is defined by

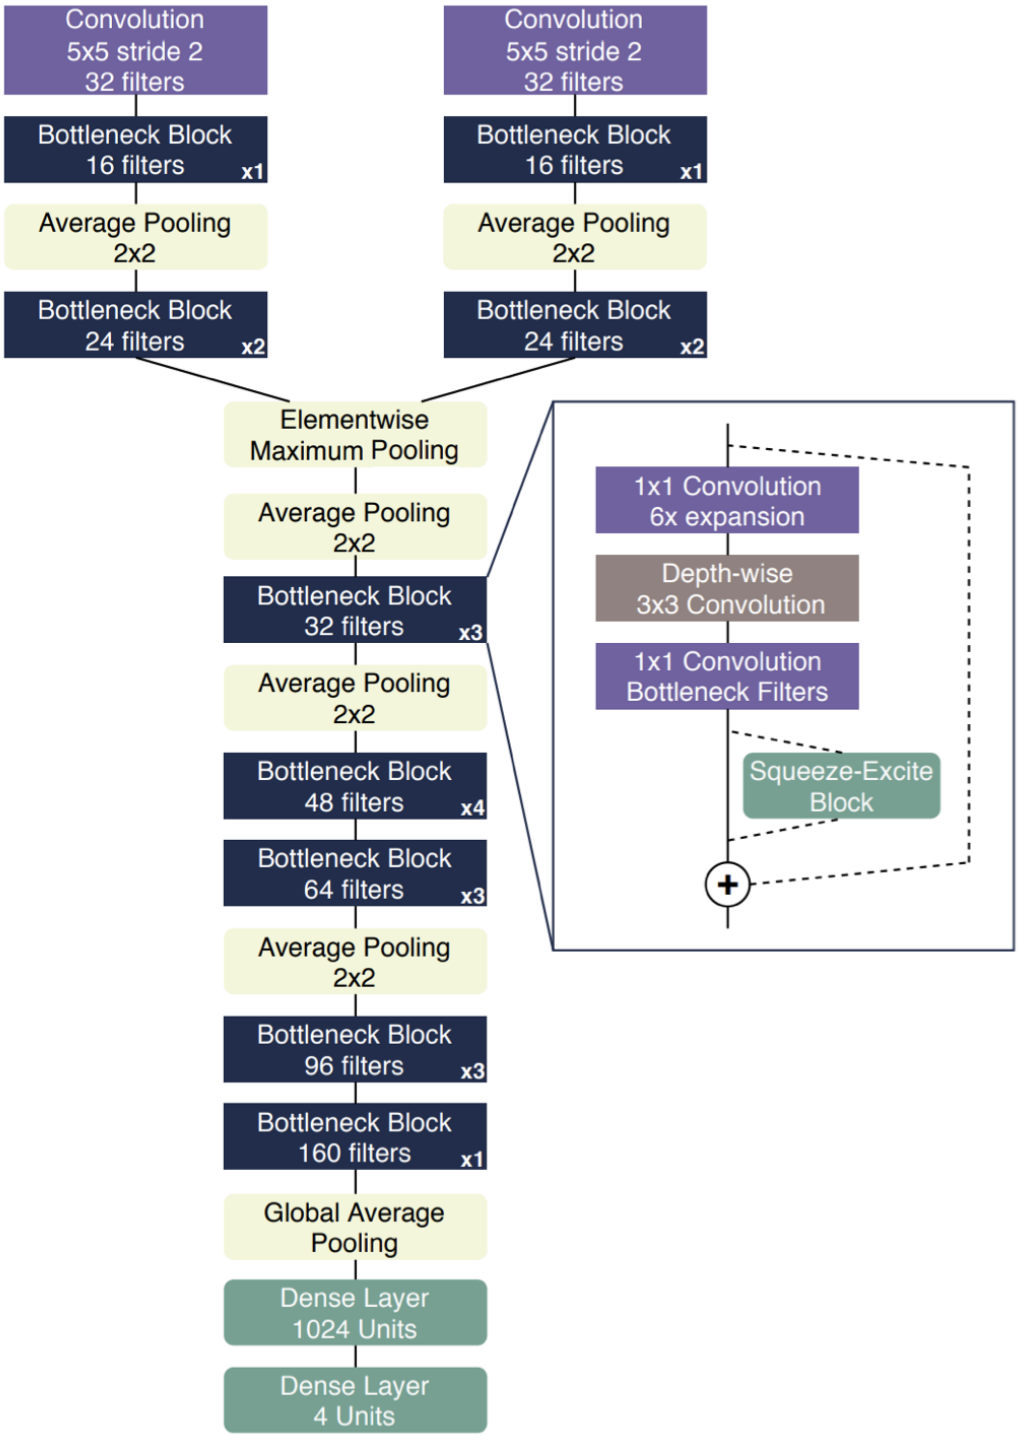

Figure 5.1: Architecture of the NuoneID convolutional neural network, modified based on MobileNetV2.

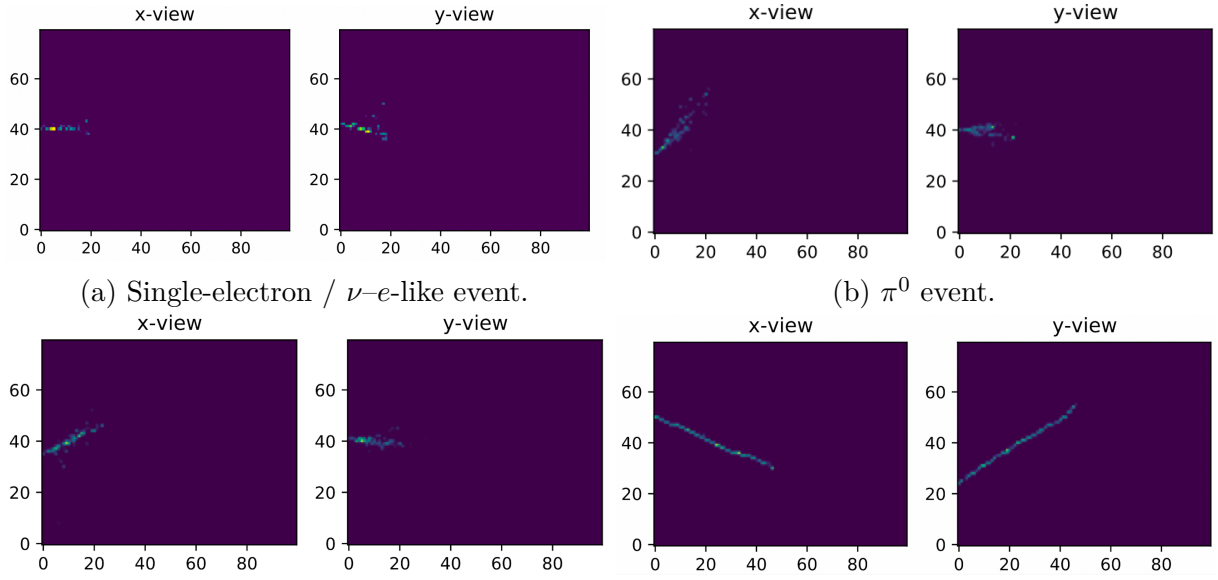

(a) Single-electron / ν–e-like event.

(b) π^0 event.

(c) ν_e charged-current event.

(d) ν_μ charged-current event.

Figure 5.2: Representative detector pixel maps used to motivate the analysis-specific CNN classification tasks.

$$-190 < x < 190 \text{ cm}, \tag{5.11}$$

$$-190 < y < 190 \text{ cm}, \tag{5.12}$$

$$50 < z < 1500 \text{ cm}. \tag{5.13}$$

The single-particle requirements are

$$E_{\text{vertex}} < 0.05 \text{ GeV}, \tag{5.14}$$

$$d_{\text{shower}} < 20 \text{ cm}, \tag{5.15}$$

$$\frac{E_{\text{shower}}}{E_{\text{tot}}} > 0.9, \tag{5.16}$$

where d_{shower} is the distance between the reconstructed vertex and the start of the shower.

5.3.3 Energy Reweighting

The training preprocessing removes high-energy electrons from the signal sample. To reduce the resulting dependence on the training energy spectrum, a true-electron-energy reweighting is applied to flatten this dependence and reduce the possibility that the classifier response is driven by an artificial training-spectrum shape.

The signal normalization relative to the combined backgrounds is also adjusted during training. The configuration adopted for NuoneID uses a signal weight of $c = 2$ to optimize performance in the independent test sample. This reweighting was not applied to Epi0ID, for which the electron-like training events already outnumbered the π^0 events by more than a factor of two.

Figure 5.3 shows the electron-energy distributions before and after reweighting and the resulting shower-energy distributions of the training samples.

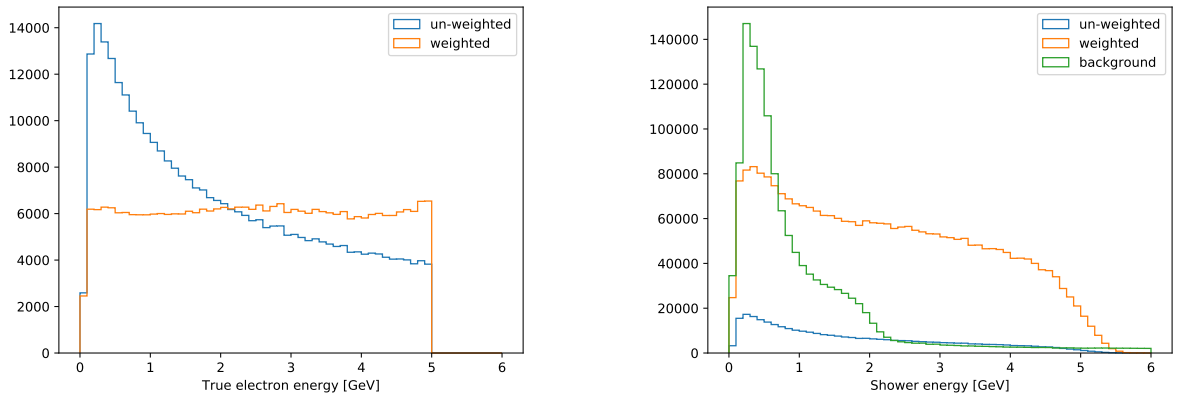

Figure 5.3: Energy distributions of the training samples. Left: True electron-energy distribution before and after reweighting. Right: Shower-energy distributions of signal and background after reweighting.

5.3.4 Training Configuration and Performance

Both classifiers use the NOvA CVN implementation modified from MobileNetV2. The training configuration used a batch size of 128, stochastic-gradient-descent optimization, an

initial learning rate of 0.02, and a decay rate of 0.05. Training and validation were performed in blocks of 1000 and 100 iterations, respectively.

Training was performed on the Wilson Cluster, where each training round was limited to 24 hours. The training loss and accuracy were monitored over successive rounds, as shown in Fig. 5.4. The training accuracies reached approximately 85% for NuoneID and 96% for Epi0ID.

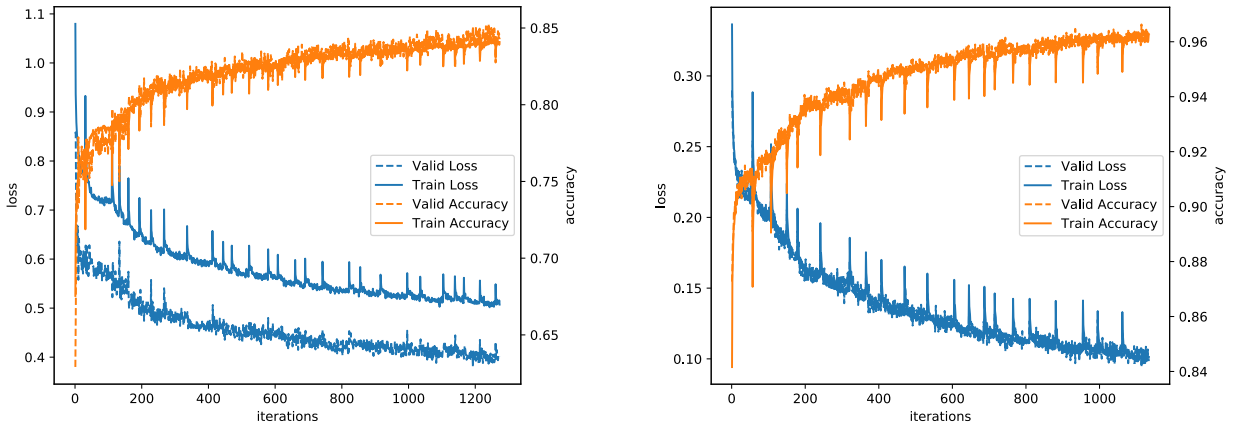

(a) NuoneID. (b) Epi0ID.

Figure 5.4: Training loss and accuracy for NuoneID and Epi0ID. The spikes indicate the beginning of successive training rounds.

After each training round, the saved model was evaluated using independent test samples. The test signal consisted of approximately 30% of the exclusive neutrino–electron sample, corresponding to 7.82×10^{22} POT. The test background consisted of approximately 80% of the Prod5 Near Detector MC sample, corresponding to 1.51×10^{21} POT, with neutrino–electron interactions excluded.

The classifier performance was evaluated using the figure of merit defined in Eq. 5.2. The final models were selected by comparing the maximum FOM obtained after successive training rounds.

Figure 5.5 shows the NuoneID distributions obtained at different stages of training. The distributions shift toward higher NuoneID values during the initial training rounds, with smaller changes observed in later rounds.

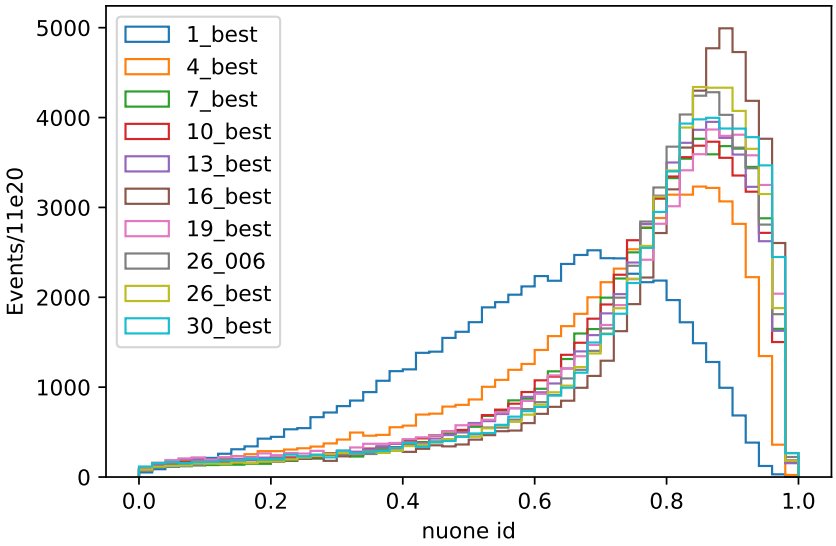

Figure 5.5: Evolution of the NuoneID distribution over successive training rounds. The curves show the classifier response at different stages of training.

5.4 Final Particle Identification and Kinematic Selection

The NuoneID and Epi0ID classifiers are applied after the preceding requirements. The optimized selection thresholds are

$$\text{NuoneID} > 0.73, \qquad \text{Epi0ID} > 0.92. \tag{5.17}$$

Their $N-1$ distributions are shown in Fig. 5.6.

After the particle-identification cuts, the signal region is defined by

$$E_e\theta_e^2 < 0.005 \text{ GeV rad}^2, \tag{5.18}$$

and the background sideband is

$$0.005 < E_e\theta_e^2 < 0.04 \text{ GeV rad}^2. \tag{5.19}$$

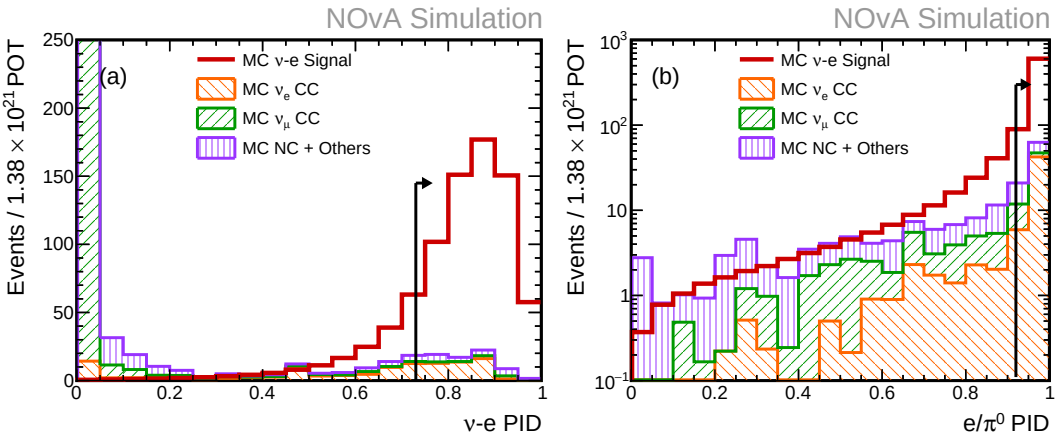

Figure 5.6: Simulated PID distributions with all other selection criteria applied. Panel (a) shows the NuoneID (ν–e PID) distribution, for which the neutrino–electron signal peaks near 0.9 while the background is concentrated predominantly at lower PID values. Panel (b) shows the Epi0ID (e/π^0 PID) distribution, for which the signal is concentrated at high PID values. Background components are stacked as ν_e CC, ν_μ CC, and NC + Other, while the simulated neutrino–electron signal is overlaid. Here, "Other" denotes reconstructed slices without a matched true neutrino interaction in the simulation, including slices arising from mis-grouping of energy deposits during event reconstruction. The black vertical arrows indicate the selection thresholds. The vertical range in panel (a) is restricted to emphasize the high-PID region; the background at low PID extends beyond the displayed range.

The reconstructed $E_e\theta_e^2$ distribution is shown in Fig. 5.7.

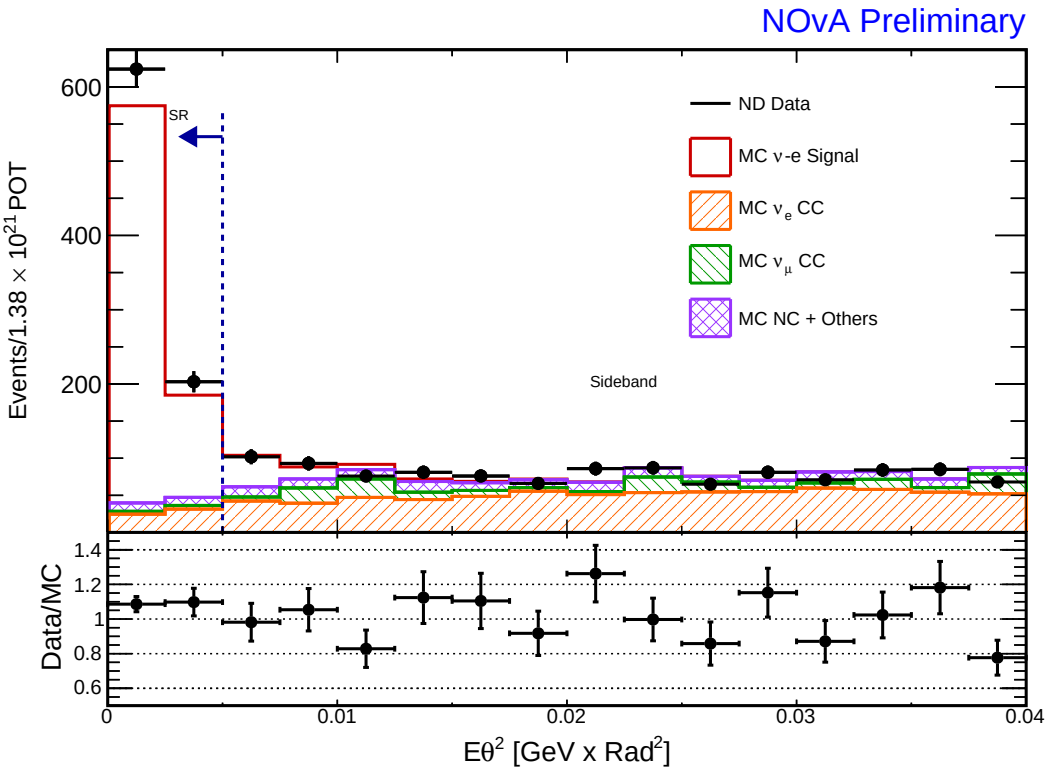

Figure 5.7: Reconstructed $E_e\theta_e^2$ distribution in Near Detector data and simulation with all other selection requirements applied (top), and the ratio of data to simulation (bottom). Data error bars represent statistical uncertainties only. The neutrino–electron signal is concentrated at low $E_e\theta_e^2$, while the background is relatively broad. Backgrounds are separated into ν_e CC, ν_μ CC, and NC plus other interactions, with ν_e CC dominant after selection. The signal region is $0 < E_e\theta_e^2 < 0.005$ GeV rad^2 and the sideband is $0.005 < E_e\theta_e^2 < 0.04$ GeV rad^2. The background prediction is corrected using the sideband procedure described in Sec. 5.5.2.

Table 5.1 gives the selection cut flow. The final sample retains 19.66% of simulated signal events and 2.84×10^{-7} of background events (equivalently 2.84×10^{-5}%), yielding approximately 89% simulated purity before the sideband correction.

5.5 Background Estimation and Subtraction

The simulation predicts 759.58 events in the signal region, including 86.93 background events, for an exposure of 1.38×10^{21} POT. These yields are obtained after the background

Table 5.1: Event-selection cut flow showing signal efficiency, background suppression, and purity at each stage. ϵ^{step} denotes the efficiency relative to the preceding selection stage, while ϵ denotes the cumulative efficiency relative to the sample before any selection. For backgrounds, ϵ_{bkg} is the cumulative background acceptance relative to the original background sample. All efficiencies and acceptances in the table are reported as percentages. Signal and background yields are obtained from simulation and normalized to 1.38×10^{21} POT, with all event weights and analysis corrections applied. No sideband-derived background normalization is applied to the yields in this table.

Selection	N_{sig}	$\epsilon_{\text{sig}}^{\text{step}}$ (%)	ϵ_{sig} (%)	N_{bkg}	$\epsilon_{\text{bkg}}^{\text{step}}$ (%)	ϵ_{bkg} (%)	Purity (%)
No cut	3422	100.0	100.0	2.950×10^8	100.0	100.0	1.160×10^{-3}
Preselection	3126	91.35	91.35	6.910×10^7	23.44	23.44	4.523×10^{-3}
Fiducial volume	3076	98.41	89.89	3.580×10^7	51.85	12.15	8.591×10^{-3}
Containment	2474	80.43	72.30	1.380×10^7	38.41	4.67	1.793×10^{-2}
Single-particle requirement	2080	84.10	60.80	3.250×10^6	23.60	1.10	6.397×10^{-2}
$E_e > 0.5$ GeV	1690	81.25	49.40	6.350×10^5	19.56	0.22	0.2655
$0.5 < E_e < 5$ GeV	1354	80.12	39.58	6.260×10^5	98.45	0.21	0.2159
NuoneID	936.9	69.18	27.38	2.480×10^4	3.97	8.42×10^{-3}	3.640
Epi0ID	747.3	79.76	21.84	1.520×10^4	61.13	5.15×10^{-3}	4.686
$E_e\theta_e^2$	672.6	90.01	19.66	83.91	0.55	2.84×10^{-5}	88.91

normalization using the sideband data. The dominant background is ν_e charged-current scattering. These events contain an electron in the final state and can mimic the signal when the associated hadronic activity is small or not reconstructed. Neutral-current and ν_μ charged-current interactions provide smaller contributions.

Figure 5.8 shows the reconstructed shower-energy distributions in data and simulation after the final selection. The background prediction shown in the figure includes the sideband normalization described in Sec. 5.5.2.

5.5.1 High-Statistics ν_e CC MEC Sample

Few ν_e charged-current MEC events survive the final selection in the nominal inclusive simulation. Their large normalization weights make the predicted contribution statistically unstable. A dedicated high-statistics ν_e CC MEC sample is therefore used to model this contribution in both the signal and sideband regions. This replacement affects only the post-selection background prediction.

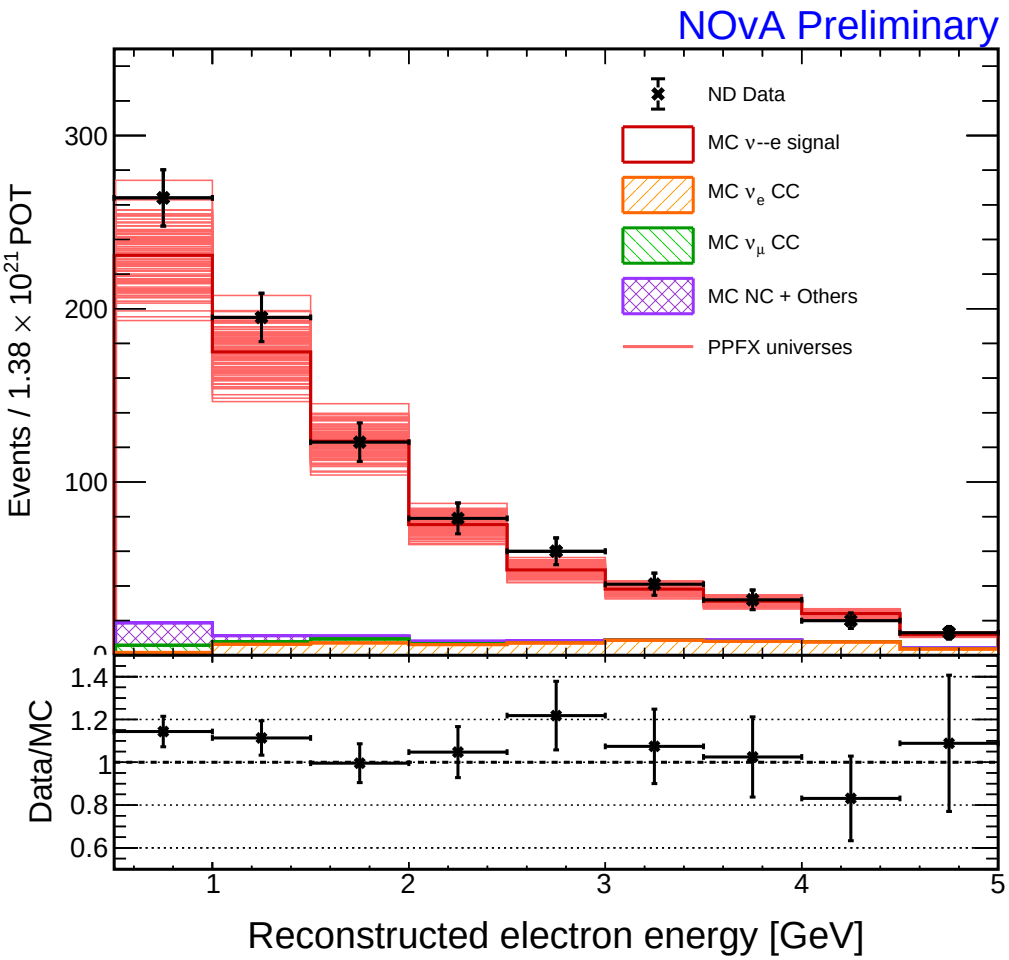

Figure 5.8: Primary reconstructed shower-energy distribution in the signal region, $0 < E_e\theta_e^2 < 0.005$ GeV rad^2, for data and simulation. Data error bars represent statistical uncertainties. The background prediction is corrected using the sideband procedure described in Sec. 5.5.2; the NC and Other backgrounds are combined in the figure. The thin red curves show the predictions from the 100 PPFX universes.

5.5.2 Sideband Normalization

A single normalization factor is obtained from the $E_e\theta_e^2$ sideband:

$$C_{\text{bkg}} = \frac{N_{\text{SB}}^{\text{data}} - N_{\text{sig,SB}}^{\text{MC}}}{N_{\text{bkg,SB}}^{\text{MC}}}. \tag{5.20}$$

The resulting correction is

$$C_{\text{bkg}} = 1.04 \pm 0.03 \text{ (stat.)}, \tag{5.21}$$

where the quoted statistical uncertainty is obtained from the finite sideband data and MC samples. The central value is used to correct the nominal background prediction in the signal region. As a scale estimate, varying C_{bkg} by ± 0.03 changes the total signal-region background by approximately 2.5 events, corresponding to about 0.34% of the background-subtracted signal yield. This contribution is small compared with the statistical uncertainty of the selected data sample and was not included as an additional uncertainty in the reported flux constraint.

For systematic variations evaluated with corresponding shifted MC samples, including the GENIE, beam, and detector/calibration variations, the background-correction factor is recalculated from the corresponding sideband sample and applied to that variation before its effect on the signal-region prediction is evaluated. This recalculation incorporates the response of the sideband normalization into each systematic variation. The correction is applied uniformly to all background components within each sample. The corrected background prediction in the signal region is 86.93 events, compared with 83.91 events before correction. The simulated neutrino–electron signal remains unchanged at 672.65 events. The event yields before and after the sideband correction are summarized in Table 5.2.

The sideband is used to determine an overall background normalization factor. Its limited statistics do not allow for a reliable energy-dependent correction, so the background energy spectrum is taken from simulation. Energy-dependent systematic uncertainties are accounted for separately using covariance matrices that include bin-to-bin correlations.

Table 5.2: Event yields in the signal region, $E_e\theta_e^2 < 0.005$ GeV rad^2, for $0.5 < E_e < 5$ GeV, with and without the sideband background correction.

Sample	With correction	Without correction
Data	827.00	827.00
MC prediction	759.58	756.56
Total MC background	86.93	83.91
ν_e CC background	55.96	54.02
ν_μ CC background	8.84	8.53
NC background	21.62	20.87
Other background	0.52	0.50
ν–e MC signal	672.65	672.65
Background-subtracted data	740.07 ± 30.23	743.09 ± 30.18
Data/MC signal ratio	1.100 ± 0.04	1.105 ± 0.04

5.5.3 Kinematic Comparison of Signal and Sideband Regions

The background composition and kinematic distributions in the signal and sideband regions are compared using simulation. The same background categories dominate both regions. Figure 5.9 compares the true Q^2 distributions of the simulated backgrounds in the two regions. Reconstructed electron energy and Bjorken x and y distributions were also examined.

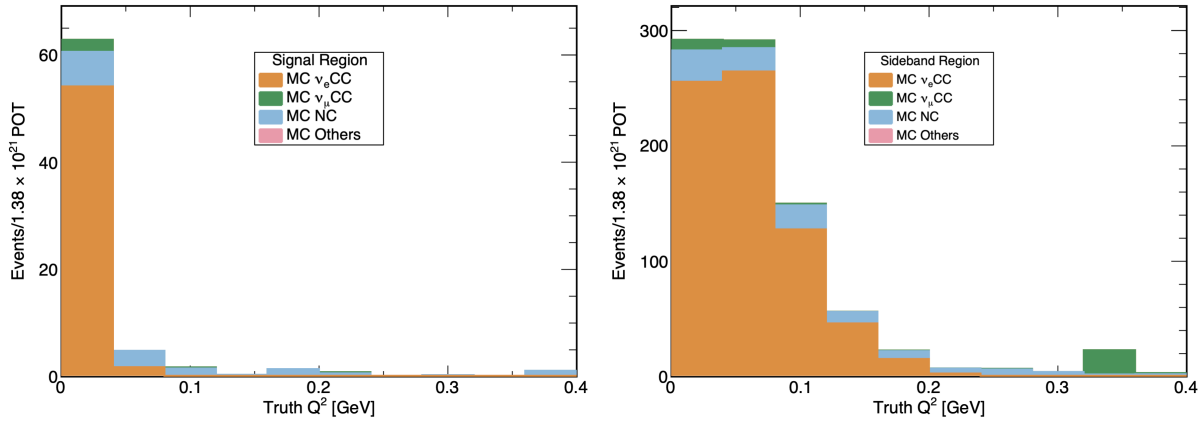

(a) Signal region.

(b) Sideband region.

Figure 5.9: True Q^2 distributions of simulated backgrounds in the $E_e\theta_e^2$ signal and sideband regions.

5.5.4 Background-Subtracted Data Sample

A total of 827 events are observed in data in the signal region. Applying the sideband-derived normalization and subtracting the corrected background prediction gives

$$N_{\nu e}^{\text{data}-\text{bkg}} = 740.1 \pm 30.2 \text{ (stat.)}. \tag{5.22}$$

The quoted uncertainty here is statistical; the non-PPFX systematic covariance used in the flux constraint is introduced separately in Sec. 5.6. The sample includes elastic-scattering contributions from all neutrino flavors in the beam. Relative to the nominal signal prediction, the background-subtracted yield corresponds to

$$\text{Data/MC} = 1.10 \pm 0.04 \text{ (stat.)}. \tag{5.23}$$

The background-subtracted reconstructed electron-energy spectrum used in the Bayesian constraint is obtained bin by bin from the selected spectrum in Fig. 5.8 after subtraction of the sideband-corrected background prediction. The 100 PPFX universe predictions shown in that figure define the prior hadron-production ensemble that is reweighted by the measurement.

5.6 Systematic Uncertainties

Non-PPFX uncertainties affecting the background-subtracted reconstructed electron-energy spectrum are propagated through covariance matrices. The dominant categories are signal-selection efficiency, GENIE interaction modeling of the background, beam optics, and detector response. PPFX hadron-production variations are not included in this covariance matrix because they form the ensemble being constrained.

Table 5.3 summarizes the systematic uncertainties on the integrated background-subtracted signal yield.

Table 5.3: Uncertainties on the integrated background-subtracted neutrino–electron signal yield. The full energy-binned covariance matrices are used in the flux constraint.

Source	Uncertainty (%)
Signal selection	1.80
GENIE interaction modeling	1.15
Beam optics	0.13
Detector response	1.98
Total systematic	2.91

5.6.1 Signal-Selection Efficiency and MR Brem Control Sample

The signal-selection uncertainty is evaluated with muon-induced bremsstrahlung electromagnetic showers, referred to as MR Brem. These showers originate from muons produced by neutrino interactions in the rock surrounding the Near Detector. After the parent-muon activity is removed, the event contains an isolated electromagnetic shower and provides a high-statistics proxy for the detector response and reconstruction of the neutrino–electron signal.

MR Brem data and simulation are processed with the same reconstruction and signal-selection criteria. Both samples correspond to an exposure of 1.38×10^{21} POT. Because the control sample and the neutrino–electron signal have different production mechanisms and kinematic spectra, the MR Brem distributions are reweighted to match the corresponding simulated neutrino–electron distributions.

The reweighting is performed separately for NuoneID, Epi0ID, and $E_e\theta_e^2$. For each variable, the weight in bin i is defined as

$$W_i = \frac{N_i^{\nu e,\mathrm{MC}}}{N_i^{\mathrm{MRBrem,MC}}}, \tag{5.24}$$

where $N_i^{\nu e,\mathrm{MC}}$ and $N_i^{\mathrm{MRBrem,MC}}$ are the simulated neutrino–electron and MR Brem event counts in the corresponding bin. The weights are derived from simulation and applied to both MR Brem data and simulation.

For each of the three final discriminating requirements, the selection efficiency is measured

in the reweighted MR Brem data and simulation with the other two requirements applied. The relative efficiency difference is calculated as

$$\delta_i = \frac{\epsilon_i^{\text{data}} - \epsilon_i^{\text{MC}}}{\epsilon_i^{\text{data}}}, \tag{5.25}$$

where ϵ_i^{data} and ϵ_i^{MC} are the corresponding selection efficiencies.

The MC/data efficiencies are 0.901/0.908 for $E_e\theta_e^2$, 0.834/0.823 for Epi0ID, and 0.830/0.837 for NuoneID. The corresponding absolute relative data–simulation differences used in the study are 0.8%, 1.4%, and 0.9%, respectively. Their quadrature sum is 1.8%, which is assigned as the signal-selection uncertainty. Figure 5.10 shows the reweighted distributions used in these comparisons.

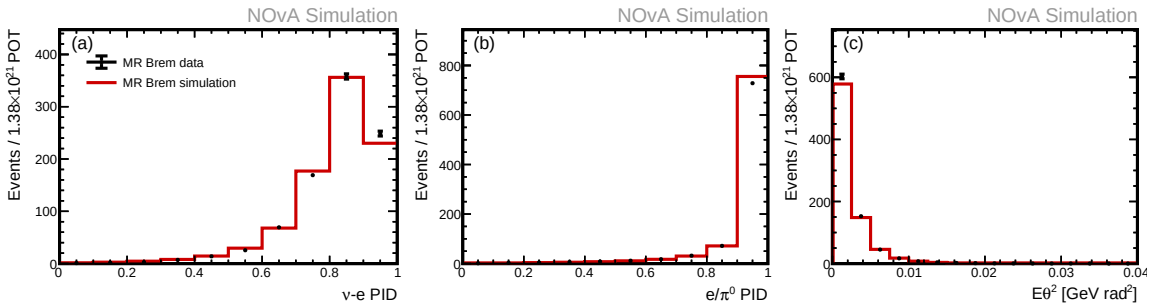

Figure 5.10: Reweighted rock-MR-Brem control-sample comparisons underlying the signal-selection efficiency uncertainty: ν–e PID (left), e/π^0 PID (center), and $E_e\theta_e^2$ (right). The MC-derived bin-by-bin weights are applied to both MR Brem data and simulation before the efficiencies are compared.

The use of MR Brem as a proxy is validated by comparing its response to systematic variations with that of the overlaid neutrino–electron signal sample. Calibration, light-level, and Cherenkov-response variations are evaluated in both samples.

Figure 5.11 shows the comparison for calibration variations. The two samples share similar responses to the applied shifts, with a larger effect observed in MR Brem. Similar behavior is observed for the light-level and Cherenkov-response variations.

The validation also includes beam-related uncertainties. The neutrino–electron distri-

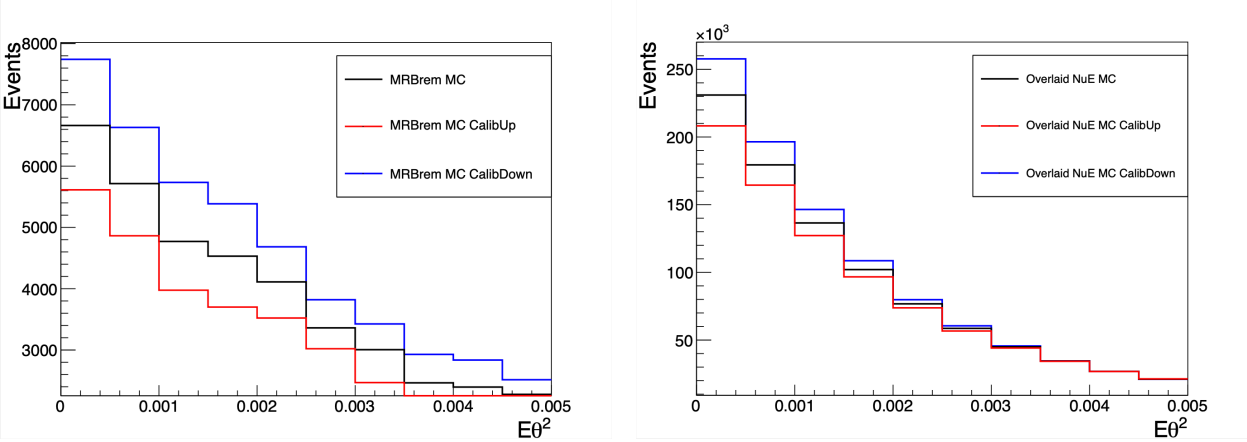

(a) MR Brem.

(b) Overlaid ν–e signal.

Figure 5.11: Comparison of calibration variations in the reconstructed $E_e\theta_e^2$ distributions of MR Brem and overlaid neutrino–electron signal samples. The nominal and calibration-shifted predictions are shown for each sample.

butions predicted from MR Brem data and simulation are compared with 24 variations obtained by shifting 12 beam parameters up and down by one standard deviation. The difference between the MR-Brem-based data and simulation predictions covers the spread of the beam-parameter variations.

These comparisons support the use of MR Brem as a proxy for estimating the signal-selection uncertainty. The detector-variant MR Brem samples are used for validation only; the assigned 1.8% uncertainty is obtained from the MR Brem data–simulation selection-efficiency differences.

5.6.2 GENIE Interaction-Model Uncertainty

GENIE model uncertainties enter through the neutrino–nucleus background. The GENIE reweighting knobs vary cross-section, hadronization, and final-state-interaction model parameters within their assigned $\pm 1\sigma$ ranges. Up to 1000 trial universes were studied during method development. As shown in Fig. 5.12, the background RMS and the uncertainty on the extracted signal rate stabilize at approximately 200 universes. The final GENIE uncertainty is evaluated using 200 universes. In each universe the signal-region background and the

sideband-derived background normalization are recalculated. The resulting background-subtracted reconstructed electron-energy spectrum in each universe is compared with the nominal spectrum, and the ensemble of bin shifts defines the GENIE covariance matrix, retaining correlations between energy bins. The integrated rate uncertainty obtained from this ensemble is 1.15%.

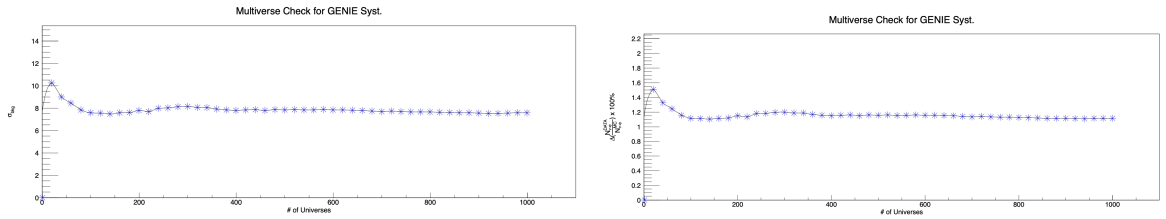

Figure 5.12: GENIE multiverse convergence study used to choose the final 200-universe ensemble. The left panel shows the background RMS and the right panel the corresponding uncertainty on the extracted neutrino–electron rate as the number of universes is increased. Both approach a plateau by approximately 200 universes.

5.6.3 Beam-Optics Uncertainty

Beam-optics uncertainties include horn current, beam position and profile, horn and target alignment, magnetic-field modeling, and beamline material descriptions. For each beamline parameter, the corresponding $+1\sigma$ and -1σ beam-simulation variations are expressed relative to the nominal neutrino flux and used to reweight the simulated event samples. The reweighted samples are then propagated through the event selection and sideband background-normalization procedure. When the two shifts are asymmetric, the larger absolute deviation from the nominal result is assigned to that parameter. The per-parameter contributions are then combined in quadrature; the resulting integrated rate uncertainty is 0.13%.

5.6.4 Detector-Response Uncertainty

Detector-response systematic uncertainties enter this analysis through the simulated neutrino–nucleus background prediction that is subtracted from the selected data sample. Dedicated detector-variation Monte Carlo samples probe the absolute energy calibration, calibration

shape, scintillation-light level, Cherenkov response, detector aging, and neutron-interaction modeling. Scintillation-light and Cherenkov-response variations produce the largest individual effects. After propagation through the background prediction, sideband normalization, and subtraction procedure, the combined detector-response uncertainty on the extracted neutrino–electron signal yield is 1.98%.

This uncertainty is distinct from the MR Brem uncertainty described in Sec. 5.6.1. The detector-response variations affect the background prediction, whereas the MR Brem control sample constrains data–simulation differences in the neutrino–electron signal-selection efficiency.

5.6.5 Total Covariance Matrix

The non-PPFX uncertainties are propagated to the reconstructed electron-energy spectrum to construct energy-binned covariance matrices. For the MR Brem selection uncertainty, the data–MC efficiency differences are used to generate varied electron-energy spectra, including statistical fluctuations, from which the selection covariance is constructed. The combined covariance matrix is

$$\Sigma_N = \Sigma_{\text{stat}} + \Sigma_{\text{sel}} + \Sigma_{\text{GENIE}} + \Sigma_{\text{beam}} + \Sigma_{\text{det}}. \tag{5.26}$$

The absolute covariance matrix used in the Bayesian flux constraint is shown in Fig. 5.13. Figure 5.14 shows the corresponding relative covariance matrix, $(\Sigma_N)_{ij}/(N_i N_j)$, for comparison. The likelihood uses the absolute covariance matrix Σ_N.

5.7 Bayesian Flux-Constraint Method

The flux constraint procedure follows the Bayesian reweighting strategy used in the MINERvA neutrino–electron measurements [31, 32]. MINERvA's reported uncertainty reductions provide context, but the two measurements are not directly comparable because their NuMI flux

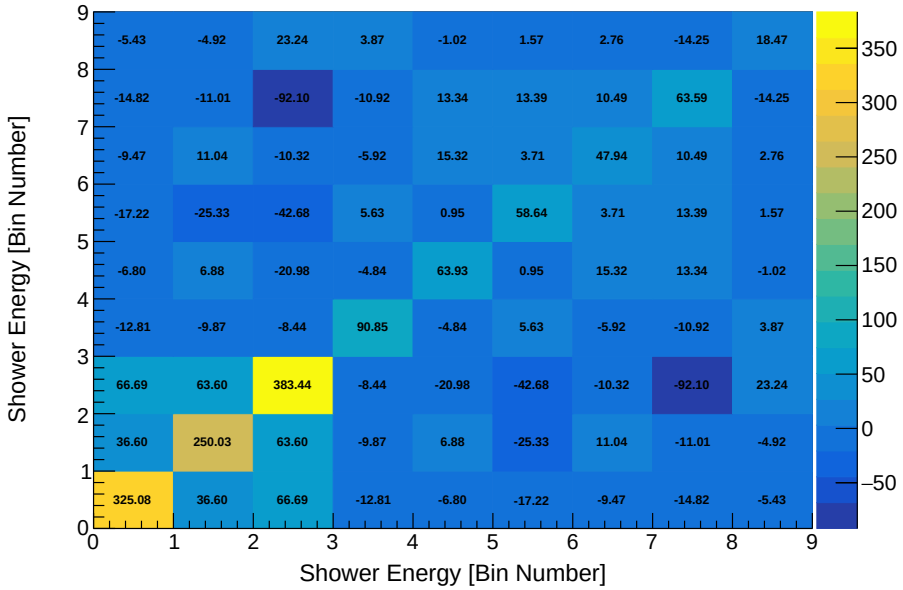

Figure 5.13: Absolute covariance matrix of the background-subtracted reconstructed electron-energy spectrum used in the flux constraint, retaining the event-yield scale. It contains the data statistical covariance and propagated non-PPFX systematic covariance contributions. This dimensional covariance is the matrix entering the likelihood in Eq. 5.28.

spectra, flavor compositions, and detector off-axis angles differ. For NOvA, the input data vector is the background-subtracted reconstructed electron-energy spectrum and the prior ensemble is the set of PPFX flux universes. For a flux model M,

$$P(M|N_{\nu e}) \propto \pi(M)\, P(N_{\nu e}|M), \tag{5.27}$$

where $\pi(M)$ is the prior probability of the flux model.

For a spectrum with K bins, the likelihood is a multivariate Gaussian,

$$P(N_{\nu e}|M) = \frac{1}{(2\pi)^{K/2}|\Sigma_N|^{1/2}} \times \exp\left[-\frac{1}{2}(\mathbf{N}-\mathbf{M})^T\Sigma_N^{-1}(\mathbf{N}-\mathbf{M})\right], \tag{5.28}$$

where $\mathbf{N}$ is the background-subtracted data spectrum, $\mathbf{M}$ is the reconstructed signal prediction for one PPFX universe, and Σ_N is the covariance matrix containing the statistical uncertainty

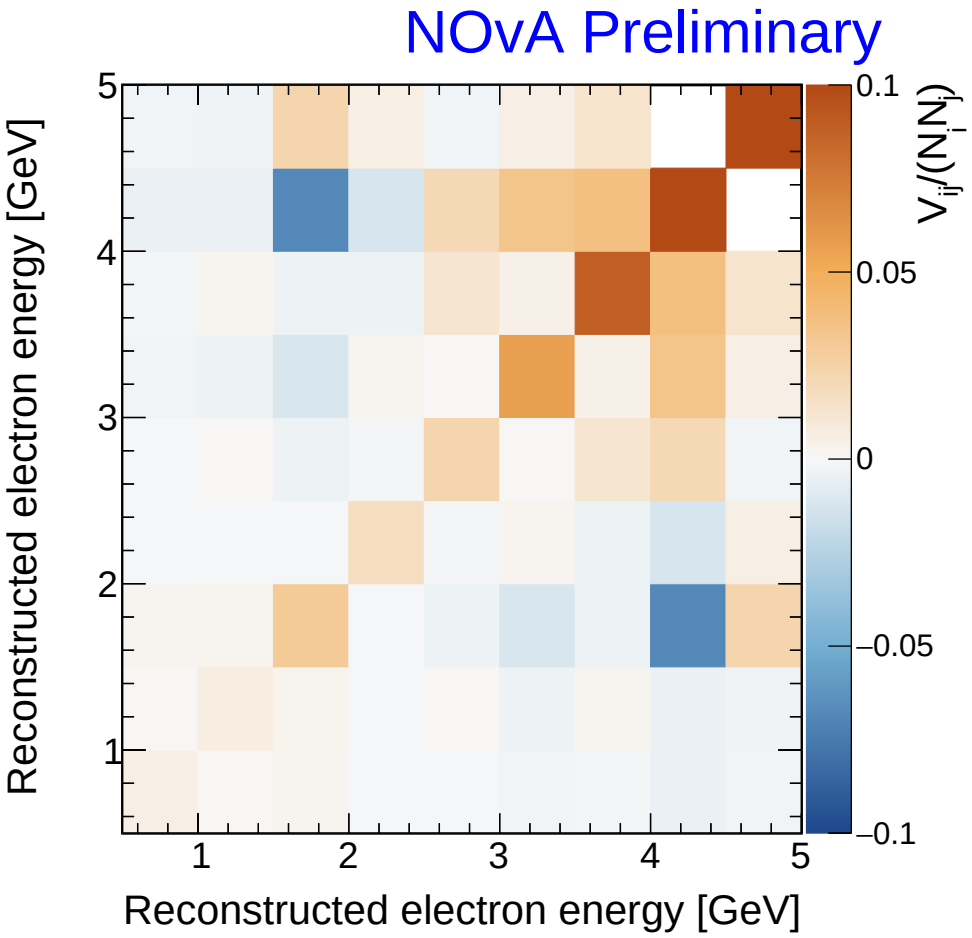

Figure 5.14: Relative covariance matrix for the background-subtracted reconstructed electron-energy spectrum used in the flux constraint. Each element is defined as $(\Sigma_N)_{ij}/(N_iN_j)$, where $(\Sigma_N)_{ij}$ is the covariance between bins i and j, and N_i and N_j are the corresponding nominal event yields. The matrix includes the statistical uncertainty of the data and the propagated non-PPFX systematic uncertainties described in Sec. 5.6.

of the measured spectrum and the propagated non-PPFX systematic uncertainties.

The 100 PPFX universes form a discrete prior ensemble with equal prior probabilities. The individual PPFX parameters are not fitted. Instead, Eq. 5.28 assigns a posterior weight to each complete flux universe. With equal priors, the universe weights are proportional to their likelihoods, $w_u \propto P(\mathbf{N}|M_u)$.

Within each PPFX universe, the correlated flux variations of ν_μ, $\bar{\nu}_\mu$, ν_e, and $\bar{\nu}_e$ are propagated simultaneously to the predicted neutrino–electron scattering spectrum. The likelihood therefore assigns a weight to the complete flux universe, and the same posterior weight is applied to all four flavor-dependent flux predictions associated with that universe. This preserves the correlations among the flavor components in the flux constraint. The PPFX universe weights are applied to the nominal selected-event sample to obtain the shifted signal predictions; the event selection is not rerun for each universe.

For an observable X predicted in universe u, the posterior mean is

$$\langle X \rangle_{\text{post}} = \frac{\sum_u w_u X_u}{\sum_u w_u}. \tag{5.29}$$

The posterior fractional uncertainty is quantified by the weighted RMS divided by the weighted mean.

5.8 Validation of the Flux-Constraint Procedure

The Bayesian flux-constraint procedure is validated using simulated data before its application to the measured neutrino–electron spectrum. The studies examine the closure of the constraint, its response to statistical and systematic variations, and its stability with respect to the number of PPFX universes.

5.8.1 Asimov Closure Test

An Asimov sample is used to evaluate whether the constraint reproduces the underlying flux prediction while reducing its uncertainty.

For the predicted neutrino–electron event rate, the constrained mean is 675.5 events, compared with an unconstrained mean of 672.8 events. The difference is approximately 0.4%, while the fractional uncertainty decreases from 8.1% to 4.2%.

For the ν_μ flux integrated over 0–120 GeV, the constrained mean is 154.9, compared with an unconstrained mean of 153.9. The difference is approximately 0.65%, and the fractional uncertainty decreases from 8.0% to 4.7%.

Figure 5.15 compares the constrained and unconstrained distributions for both quantities. The constrained means remain close to the unconstrained predictions, while the posterior uncertainties are substantially reduced.

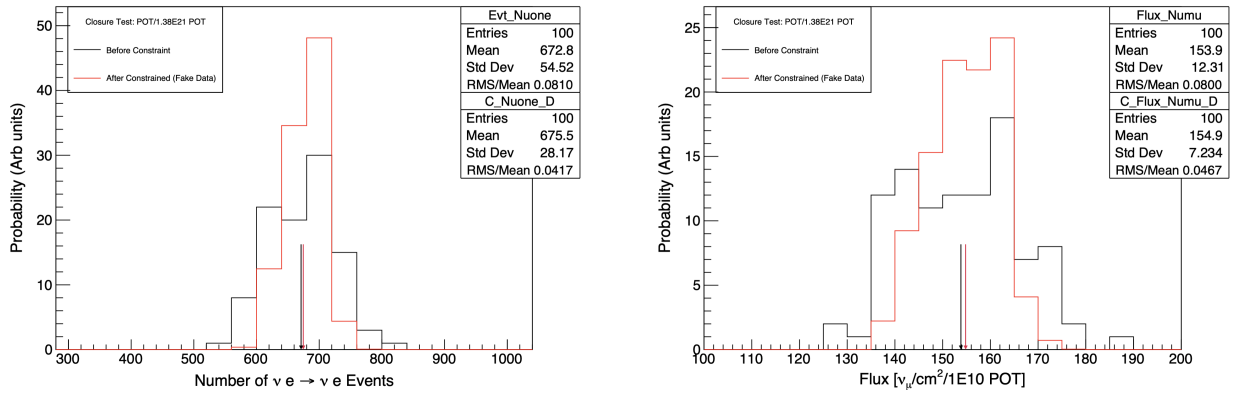

(a) Predicted neutrino–electron event count.

(b) Integrated ν_μ flux.

Figure 5.15: Asimov closure test of the Bayesian flux constraint, comparing the constrained and unconstrained distributions of the predicted neutrino–electron event rate and integrated ν_μ flux.

Potential correlations between signal and background systematic uncertainties were examined using a calibration-variation test. Nominal, calibration-up, and calibration-down fake-data samples were constructed with consistent variations applied to the signal and background simulations. The Bayesian flux-constraint procedure was repeated for each sample, and the resulting constrained ν_μ flux predictions were compared with the central-

value prediction.

The maximum change in the constrained flux relative to the nominal fake-data result was compared with the calibration uncertainty estimated from the signal and background contributions added in quadrature. Over the integrated energy range, the correlation-induced variation was 5.9%, compared with an estimated calibration uncertainty of 6.3%, giving a ratio of 0.94. The calibration uncertainty also exceeded the observed variation in most energy bins. These results indicate that the adopted covariance treatment adequately accounts for the correlated effects examined in this calibration study.

5.8.2 PPFX Ensemble Stability

The stability of the constraint is evaluated using progressively larger subsets of the 100-universe PPFX ensemble. Each subset contains the first N universes in their original ordering.

Figure 5.16 shows the dependence of the fractional uncertainty on the number of included universes. The unconstrained RMS/Mean continues to change gradually as more universes are included. For the constrained neutrino–electron event rate and integrated ν_μ flux, the step-to-step changes fall below approximately 0.25% after about 55 universes.

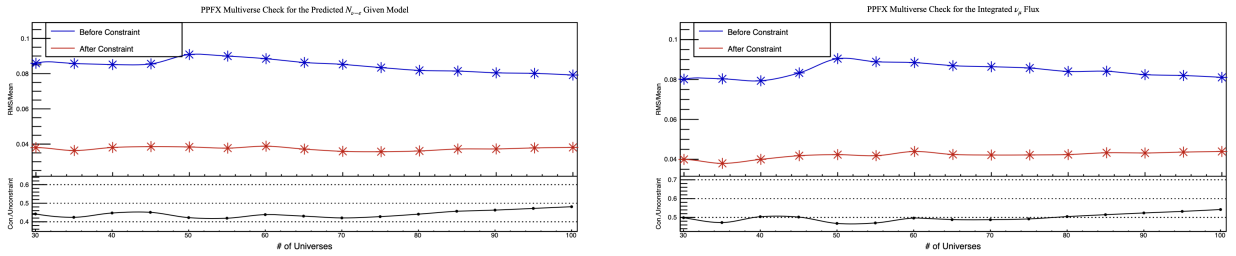

(a) Predicted neutrino–electron event rate.

(b) Integrated ν_μ flux.

Figure 5.16: Dependence of the unconstrained and constrained fractional uncertainties on the number of PPFX universes included in the flux constraint.

The final constraint uses all 100 PPFX universes. The convergence study indicates that the constrained uncertainties are stable within the tested cumulative sequence.

5.9 Flux-Constraint Results

The Bayesian flux constraint is applied to the P1–P11a Near Detector data using the likelihood and covariance matrix described in the preceding sections.

5.9.1 Predicted Neutrino–Electron Event Rate

The prior PPFX ensemble gives a 7.7% fractional uncertainty on the predicted total selected neutrino–electron event rate. After reweighting, the posterior MC prediction has a mean of approximately 722 events, with the fractional uncertainty reduced to 3.9%:

$$7.7\% \rightarrow 3.9\%. \tag{5.30}$$

The corresponding prior and posterior distributions are shown in Fig. 5.17.

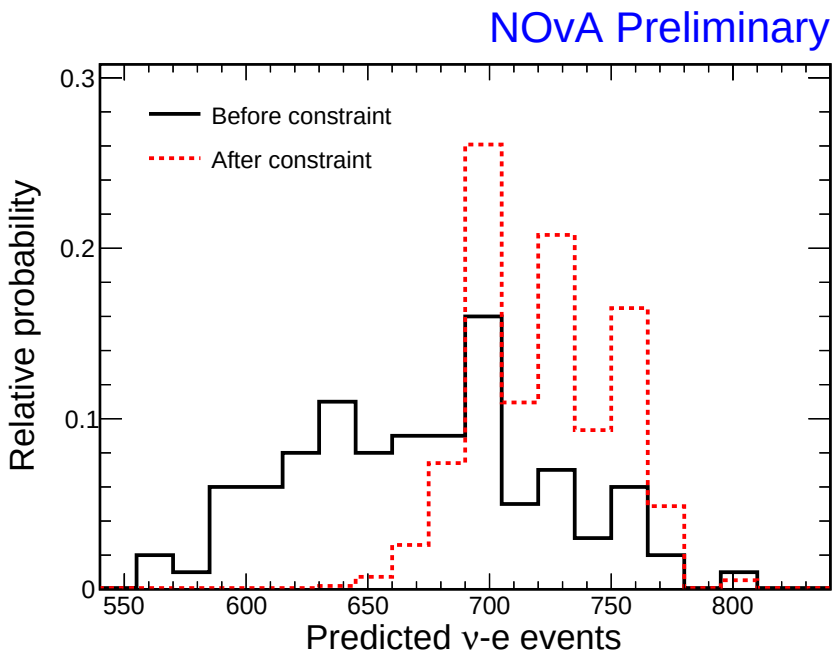

Figure 5.17: Probability distributions of the predicted total number of neutrino–electron scattering events across the 100 PPFX universes before (black solid) and after (red dashed) applying the data constraint. The predicted yields correspond to 1.38×10^{21} POT. The fractional spread (RMS/Mean) is reduced from 7.7% to 3.9%.

5.9.2 Integrated ν_μ Flux

The posterior weights are applied to the corresponding ν_μ flux prediction. Integrated over the full simulated range from 0 to 120 GeV, the constrained mean is

$$\Phi_{\nu_\mu}^{\text{post}} = 126.4\ \nu_\mu\,\text{cm}^{-2}/10^{10}\,\text{POT}, \tag{5.31}$$

and the fractional uncertainty is reduced from

$$8.1\% \rightarrow 4.0\%. \tag{5.32}$$

The 120 GeV upper edge refers to the range of the simulated flux prediction, rather than the reconstructed electron-energy range of the selected sample. The prior and posterior distributions of the integrated ν_μ flux are shown in Fig. 5.18.

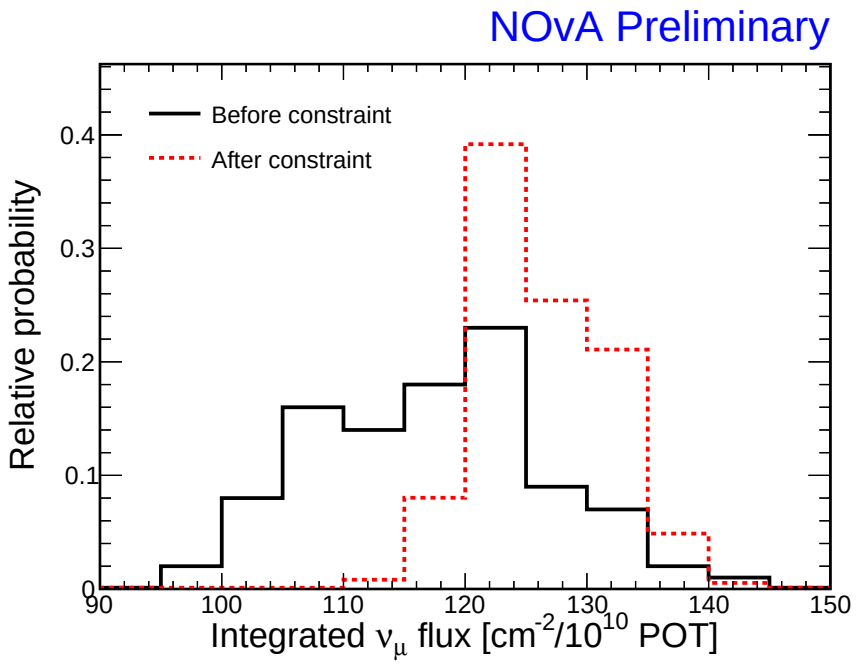

Figure 5.18: Probability distributions of the ν_μ flux integrated from 0 to 120 GeV across the 100 PPFX universes before (black solid) and after (red dashed) applying the data constraint. The integrated flux is expressed per 10^{10} POT. The fractional spread (RMS/Mean) is reduced from 8.1% to 4.0%.

5.9.3 Energy-Dependent Constraint

The posterior weights alter both the central value and fractional uncertainty of the ν_μ flux as a function of neutrino energy. Figure 5.19 compares the constrained and unconstrained flux predictions. Across the energy range with appreciable NOvA flux, the constrained central value differs from the unconstrained prediction by approximately 3–15%.

Figure 5.20 shows the corresponding fractional uncertainties. The largest relative uncertainty reduction occurs near the 2 GeV flux peak, where the fractional uncertainty is reduced by as much as approximately 50%.

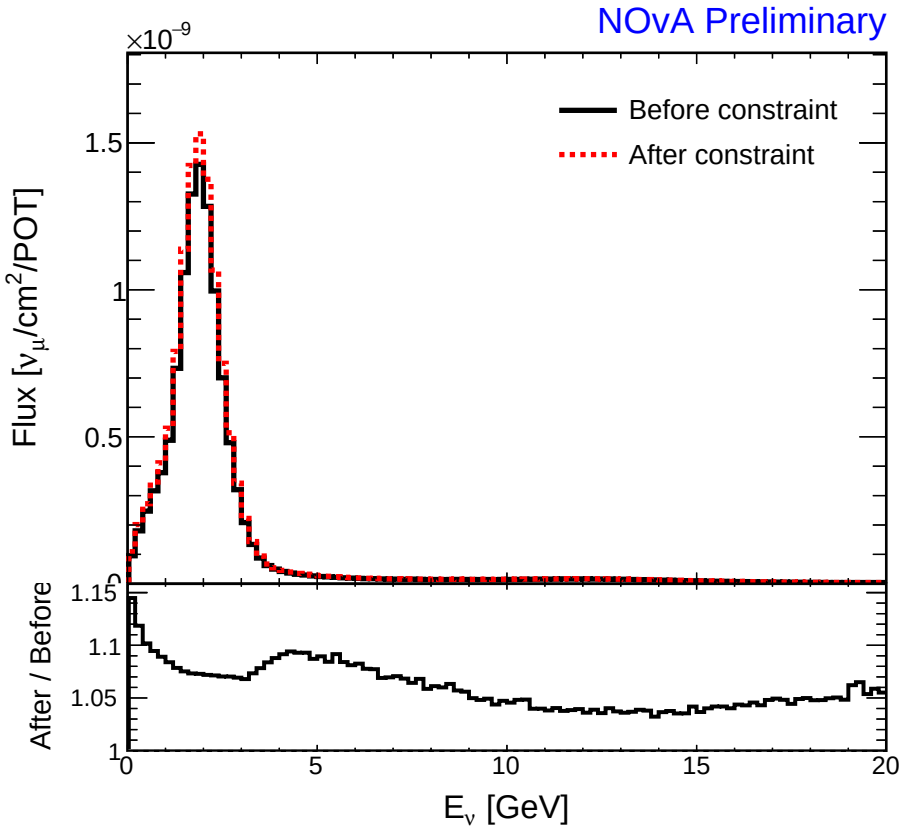

Figure 5.19: Predicted ν_μ flux as a function of neutrino energy before (black solid) and after (red dashed) applying the neutrino–electron constraint (top), and the ratio of the constrained to unconstrained flux predictions (bottom).

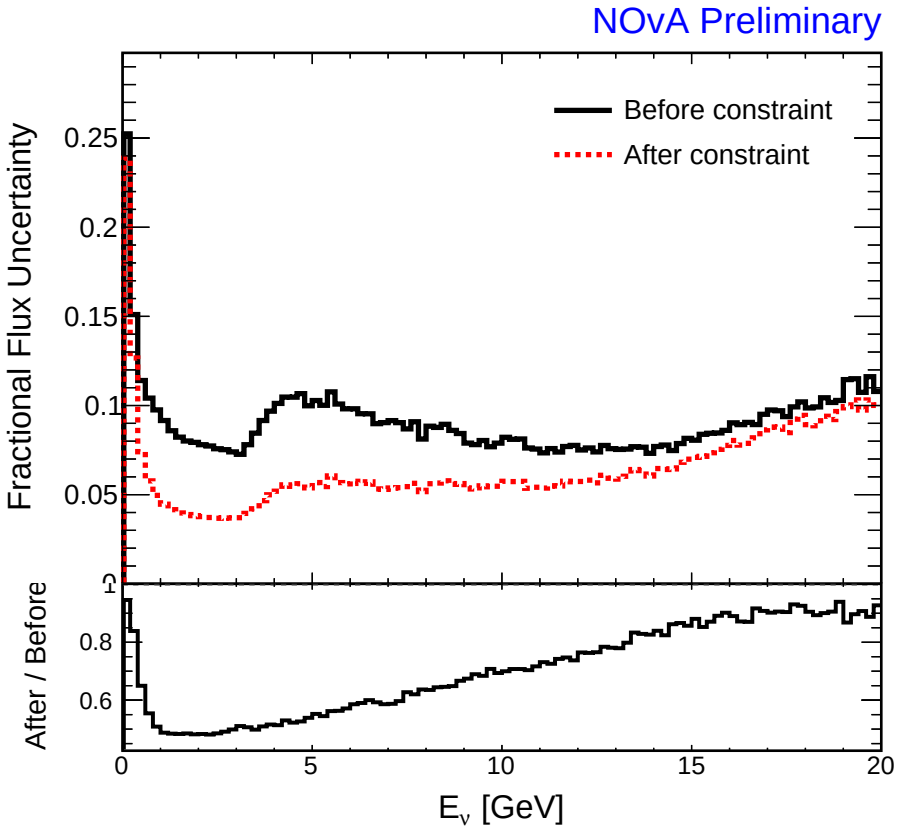

Figure 5.20: Fractional uncertainty (RMS/Mean) of the predicted ν_μ flux before (black solid) and after (red dashed) applying the neutrino–electron constraint (top), and the ratio of the constrained to unconstrained fractional uncertainties (bottom).

5.10 Interpretation and Propagation of the Constraint

5.10.1 Propagation to Other Beam Flavors

The posterior PPFX weights are also applied to the $\bar{\nu}_\mu$, ν_e, and $\bar{\nu}_e$ flux predictions to obtain the corresponding constrained distributions.

Table 5.4 summarizes the integrated-flux response of the four beam components over the same simulated flux range used for the ν_μ result. The wrong-sign $\bar{\nu}_\mu$ component is constrained as well: its fractional uncertainty decreases from 14.0% to 9.7%. The reduction is smaller than for the dominant ν_μ component. The ν_e and $\bar{\nu}_e$ components also change under the same posterior reweighting. The uncertainty reductions for the other neutrino flavors arise from correlations within the PPFX ensemble and the posterior weights determined by the neutrino–electron data.

Table 5.4: Effect of the neutrino–electron constraint on the integrated flux of each beam component. The means are quoted in the same integrated-flux units as the ν_μ result, per 10^{10} POT. Fractional uncertainties are the RMS/Mean of the 100-universe ensemble before and after posterior reweighting.

Component	Prior mean	Posterior mean	Prior RMS/Mean	Posterior RMS/Mean
ν_μ	117.35	126.44	8.09%	4.01%
$\bar{\nu}_\mu$	10.56	11.67	13.97%	9.67%
ν_e	15.66	16.70	7.65%	4.52%
$\bar{\nu}_e$	27.98	30.29	11.77%	9.62%

5.10.2 Downstream Check with a ν_μ CC Selection

The posterior PPFX weights are implemented in the NOvA CAFAna framework and applied to an independent ν_μ charged-current selection.

Figure 5.21 compares the nominal and constrained predictions for the reconstructed neutrino-energy distribution, with and without neutrino-interaction cross-section weights. The constrained predictions show an increase in the expected event rate, consistent with the upward shift in the ν_μ flux obtained from the neutrino–electron constraint.

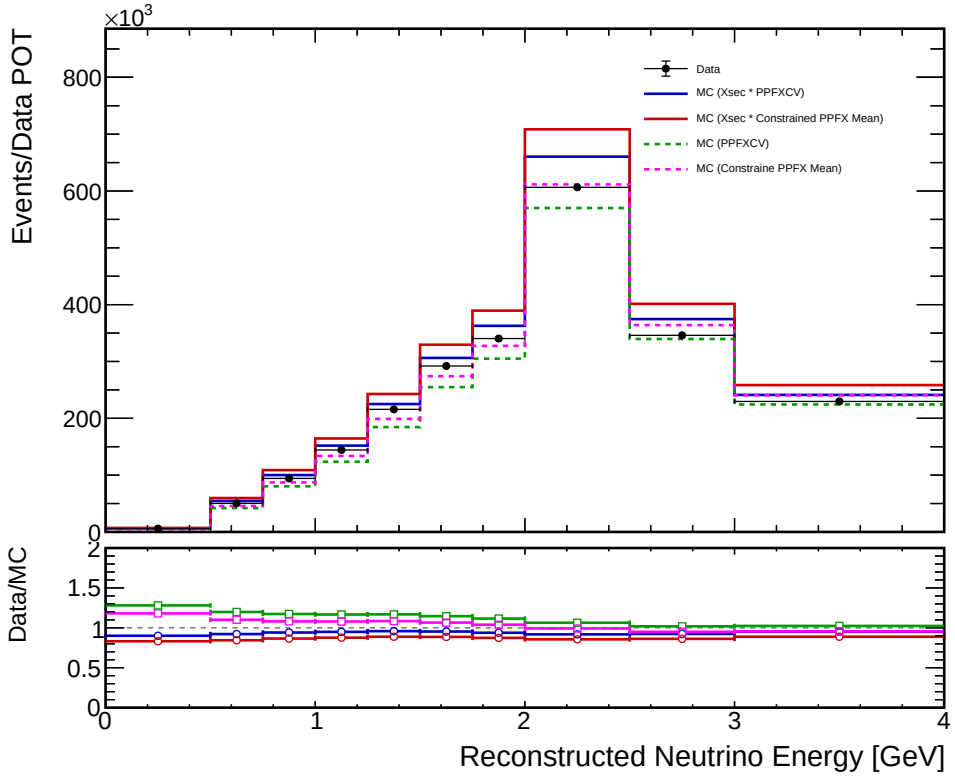

Figure 5.21: Reconstructed neutrino-energy distribution for an independent ν_μ charged-current selection. The data are compared with nominal and constrained MC predictions, with and without neutrino-interaction cross-section weights.

Figure 5.22 shows the corresponding PPFX multiverse predictions before and after applying the flux constraint. The posterior weights modify the central prediction and the spread of the PPFX universes in the ν_μ CC reconstructed-energy distribution.

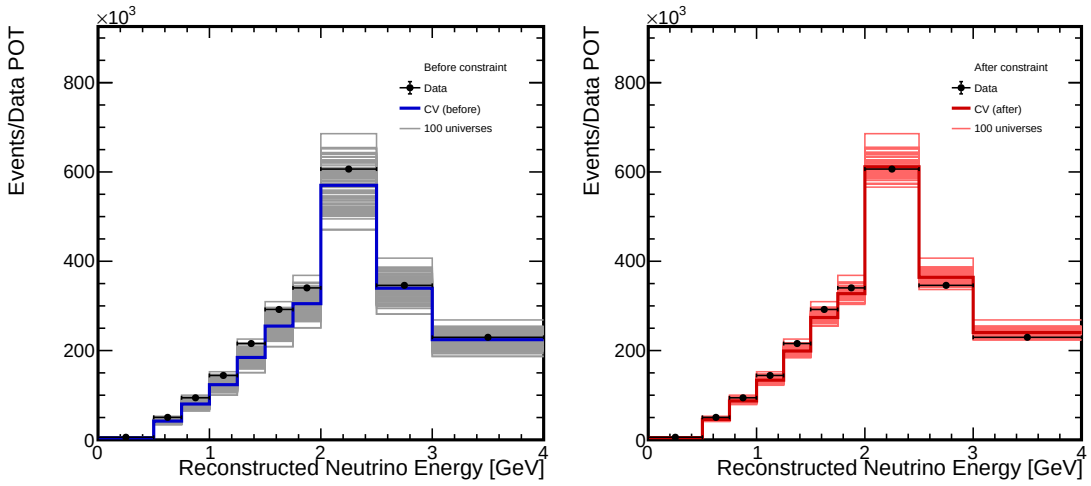

Figure 5.22: Reconstructed neutrino-energy distributions for the ν_μ charged-current selection before (left) and after (right) applying the flux constraint. The central predictions and the corresponding 100-universe PPFX distributions are shown together with the selected data.

These comparisons demonstrate the propagation of the posterior PPFX weights to an independent NOvA event selection.

5.11 Summary

This chapter presents a measurement of neutrino–electron elastic scattering using 1.38×10^{21} POT of NOvA Near Detector neutrino-mode data and its application to constraining the NuMI neutrino flux. The final signal region contains 827 observed events. A background normalization factor of $C_{\text{bkg}} = 1.04$ is obtained from the $E_e\theta_e^2$ sideband, yielding a background-subtracted signal of 740.1 ± 30.2 (stat.) events.

The background-subtracted electron-energy spectrum is used to constrain the neutrino flux by reweighting 100 PPFX universes within a Bayesian framework. The fractional uncertainty

on the predicted neutrino–electron event rate is reduced from 7.7% to 3.9%. The constrained ν_μ flux, integrated over 0–120 GeV, has a mean of 126.4 $\nu_\mu\,\mathrm{cm}^{-2}/10^{10}\,\mathrm{POT}$, with its fractional uncertainty reduced from 8.1% to 4.0%.

The constraint procedure is examined using Asimov tests, together with a study of the stability of the PPFX ensemble. The posterior PPFX weights are also implemented in CAFAna and applied to an independent ν_μ charged-current selection.

Chapter 6

The NOvA Neutrino Oscillation Analysis

NOvA measures muon-neutrino disappearance and electron-neutrino appearance using Far Detector (FD) event spectra constrained by Near Detector (ND) data. Forward-horn-current (FHC) running produces a neutrino-enriched beam, while reverse-horn-current (RHC) running produces an antineutrino-enriched beam. The joint analysis combines ν_μ and $\bar{\nu}_\mu$ disappearance with ν_e and $\bar{\nu}_e$ appearance to constrain oscillation parameters and the neutrino mass ordering.

The detector, simulation, and reconstruction methods are described in Chapter 3. This chapter introduces the selected oscillation samples, the ND-to-FD prediction, and the likelihood used to extract oscillation parameters. These elements provide the experimental setting for the statistical methods discussed in Chapters 7 and 8.

6.1 Oscillation-Analysis Overview

At the ND baseline, neutrino oscillations are negligible within the standard three-flavor framework. The ND event spectra are used to constrain uncertainties in the neutrino flux, interaction modeling, and detector response. At the FD, 810 km from the source, the energy spectra are modified by ν_μ disappearance and ν_e appearance. To predict these spectra, the

Near-to-Far extrapolation uses ND measurements together with simulation to account for differences in neutrino flux and detector acceptance between the two detectors.

The analysis begins with simulated FHC and RHC interactions in both detectors. Reconstructed tracks, showers, and event-identification quantities define samples enriched in muon- or electron-neutrino charged-current interactions. ND data constrain the predicted event populations in reconstructed observables. Simulation is used to map the ND constraints from reconstructed energy to true neutrino energy. The corrected true-energy spectrum is then extrapolated to the FD, with oscillation probabilities applied at each energy. The resulting FD prediction is expressed in reconstructed energy and compared with the FD data.

The oscillation fit combines the selected samples while varying oscillation parameters and systematic nuisance parameters. Correlated uncertainties in the beam flux, interaction model, and detector response can affect several samples and energy bins at once. The ND constraint reduces some of these effects, while the remaining variations are propagated to the FD prediction.

6.2 Event Samples and Selections

The analysis uses separate ν_μ-like and ν_e-like samples in FHC and RHC running. Event reconstruction and CVN-based classification are introduced in Sec. 3.5 and Ref. [30]. The oscillation selections identify candidate interactions based on their reconstructed topology, with additional requirements on event containment, reconstruction quality, and timing relative to the NuMI beam spill. Their output is a set of reconstructed-energy spectra with simulated signal and background components. The different event topologies determine both the energy estimator and the dominant backgrounds.

6.2.1 Muon-Neutrino Disappearance Samples

The ν_μ disappearance selection identifies charged-current interactions through a reconstructed muon track consistent with a minimum-ionizing particle. Selected events must satisfy containment and reconstruction-quality requirements, with their timing consistent with the NuMI beam spill. The selection is applied to both FHC and RHC data to obtain neutrino- and antineutrino-enriched samples, respectively. Both samples contain wrong-sign interactions, which are included in the corresponding simulated predictions.

The disappearance measurement uses the shape of the reconstructed neutrino-energy spectrum. The neutrino energy is estimated from the reconstructed muon energy and the visible hadronic energy. The muon track provides a relatively well-constrained contribution, while the hadronic system can contain energy that is not detected or is reconstructed with poorer resolution. The resulting energy smearing affects how the reduction in event rate due to neutrino oscillations appears in the reconstructed-energy spectrum.

Events with different hadronic-energy contributions have different neutrino-energy resolutions. The disappearance sample is divided into categories based on the fraction of reconstructed energy carried by the hadronic system. Each category has its own energy spectrum and simulated true-to-reconstructed-energy distribution, allowing the oscillation fit to account for differences in energy resolution.

6.2.2 Electron-Neutrino Appearance Samples

The ν_e appearance selection identifies charged-current interactions with an electromagnetic shower consistent with an electron. Because the appearance signal is small relative to the unoscillated muon-neutrino beam, the selection must distinguish electron-like signal candidates from neutral-current interactions, misidentified ν_μ charged-current events, and the intrinsic electron-neutrino component of the beam. CVN-based classification provides event-topology information used for this separation.

The FD is located at the surface, where cosmic-ray activity is an additional background.

Events are selected based on their containment and topology, with additional requirements to reject cosmic rays and ensure consistency with the NuMI beam spill. The remaining cosmic background is estimated from detector data recorded outside the beam timing window. Beam-related backgrounds are treated separately because they arise from different interaction and flux components and respond differently to oscillation and systematic variations.

For selected appearance candidates, the neutrino energy is reconstructed from the electromagnetic shower and visible hadronic energy. The resulting energy spectrum, together with the event classification, defines the information compared with the FD appearance prediction. The FHC and RHC appearance samples have different signal rates and background compositions and are included separately in the joint oscillation fit. Together, they provide sensitivity to δ_{CP} and matter effects.

6.3 Near-to-Far Extrapolation and the Near Detector Constraint

The ND and FD observe different neutrino energy spectra because of their positions relative to the NuMI decay region. Their event spectra also depend on differences in detector acceptance and energy reconstruction. The Near-to-Far extrapolation uses ND measurements and simulation to construct the FD prediction [28].

The ND data/MC ratio provides a correction to the reconstructed-energy spectrum. Simulation maps this correction to true neutrino energy. The corrected spectrum is extrapolated to the FD and weighted by the oscillation probability before applying the FD detector response.

For the ν_μ disappearance signal, the ND reconstructed-energy correction can be expressed as

$$C_{\mathrm{ND}}(E_{\mathrm{reco}}) = \frac{N_{\mathrm{ND}}^{\mathrm{data}}(E_{\mathrm{reco}}) - N_{\mathrm{ND}}^{\mathrm{bkg}}(E_{\mathrm{reco}})}{N_{\mathrm{ND}}^{\mathrm{sig,MC}}(E_{\mathrm{reco}})}. \tag{6.1}$$

Here, $N_{\mathrm{ND}}^{\mathrm{bkg}}$ is the estimated ND background, and $N_{\mathrm{ND}}^{\mathrm{sig,MC}}$ is the simulated disappearance signal.

Simulation provides the reconstructed-to-true-energy distributions used in the extrapolation.

Figure 6.1 illustrates the extrapolation procedure for the ν_μ disappearance signal. Systematic variations are used to evaluate uncertainties in the flux, interaction modeling, and detector response.

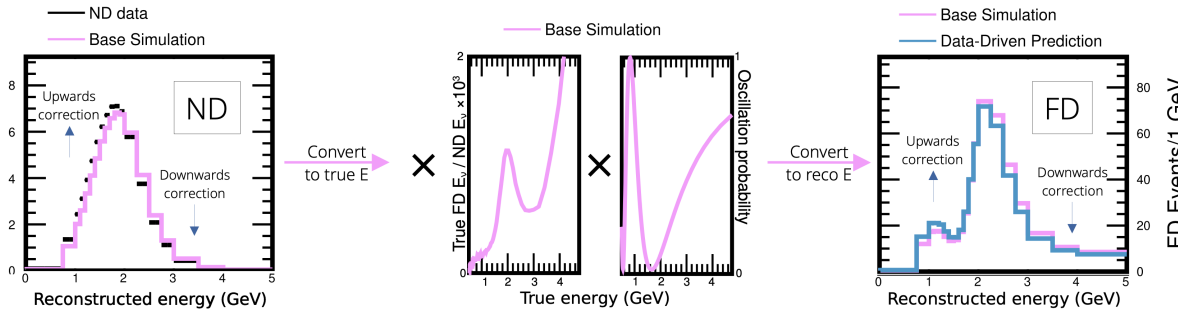

Figure 6.1: Schematic of the Near-to-Far extrapolation for the ν_μ disappearance signal. The ND reconstructed-energy spectrum is mapped to true energy, transferred to the FD, weighted by the survival probability, and mapped to FD reconstructed energy.

6.3.1 Migration Matrices and True-Energy Predictions

The reconstructed-to-true-energy mapping uses the simulated joint distribution of reconstructed and true neutrino energy. Let $N_{ij}^{\mathrm{ND,MC}}$ denote the simulated ND disappearance-signal count in reconstructed-energy bin i and true-energy bin j, with $N_i^{\mathrm{ND,MC}} = \sum_j N_{ij}^{\mathrm{ND,MC}}$. For the disappearance signal, the mapping can be written schematically as

$$\widetilde{N}_j^{\mathrm{ND}} = \sum_i N_{ij}^{\mathrm{ND,MC}} \frac{N_i^{\mathrm{ND,data}}}{N_i^{\mathrm{ND,MC}}}, \tag{6.2}$$

where $N_i^{\mathrm{ND,data}}$ is the ND data count after subtracting the estimated background contribution. The simulated distribution determines the true-energy composition of each reconstructed-energy bin.

The FD disappearance-signal prediction can then be represented as

$$N_k^{\mathrm{FD,pred}} = \sum_j R_{kj}^{\mathrm{FD}}\, T_j^{\mathrm{FD/ND}}\, P_{\mu\mu}(E_j)\, \widetilde{N}_j^{\mathrm{ND}}, \tag{6.3}$$

where $T_j^{\mathrm{FD/ND}}$ transfers the ND true-energy population to the FD before oscillations and accounts for differences in flux, exposure, and acceptance. The FD response matrix R_{kj}^{FD} describes the reconstruction and selection of events in reconstructed-energy bin k for a given true-energy bin j. The survival probability $P_{\mu\mu}(E_j)$ is evaluated in true energy.

Uncertainties in the detector response and interaction modeling affect the mapping between true and reconstructed neutrino energy. For example, variations in the hadronic energy response change the reconstructed-energy distribution for a given true neutrino energy. These effects are propagated through the Near-to-Far extrapolation and included in the oscillation fit.

6.3.2 Appearance Signal and Background Components

The FD ν_e appearance signal comes from ν_μ oscillating into ν_e during propagation. The prediction starts with the ND ν_μ sample and uses simulation to extrapolate the neutrino flux to the FD. The $\nu_\mu \to \nu_e$ oscillation probability is then applied to obtain the appearance signal. The FD detector response and appearance selection determine its reconstructed energy spectrum.

Several backgrounds contribute to the FD appearance sample. The beam contains an intrinsic ν_e component that is present before oscillations. Neutral-current interactions and ν_μ charged-current events can also be misidentified as electron-neutrino candidates. These beam-related backgrounds are predicted using their corresponding fluxes, interaction models, and selection efficiencies. Cosmic-ray events form an additional background at the FD and are estimated using data collected outside the NuMI beam spill.

Both FHC and RHC samples contain wrong-sign interactions. Neutrino and antineutrino

contributions are modeled separately to account for differences in their fluxes, cross sections, and oscillation probabilities. The individual signal and background components are then combined to form the predicted reconstructed-energy spectra. Systematic variations are applied to the corresponding components to evaluate uncertainties in event rates, energy distributions, and selection efficiencies.

6.4 Oscillation Likelihood and Systematic Parameters

The oscillation parameters are determined by comparing the observed FD energy spectra with the predicted spectra. The fit includes both ν_μ disappearance and ν_e appearance samples in FHC and RHC modes.

For each analysis bin i, let n_i be the observed number of events and μ_i the predicted number. The prediction depends on the oscillation parameters $\boldsymbol{\omega}$ and the systematic parameters $\boldsymbol{\eta}$. Assuming Poisson statistics, the likelihood is

$$\mathcal{L}_{\text{stat}}(\boldsymbol{\omega}, \boldsymbol{\eta}) = \prod_i \frac{\mu_i(\boldsymbol{\omega}, \boldsymbol{\eta})^{n_i} e^{-\mu_i(\boldsymbol{\omega}, \boldsymbol{\eta})}}{n_i!}, \tag{6.4}$$

where the index i runs over all energy bins and analysis samples.

The predicted spectra also depend on systematic uncertainties. These include uncertainties in the predicted neutrino flux, neutrino interaction cross sections, detector calibration, and energy reconstruction. They can affect both the number of predicted events and the shapes of the energy spectra.

Systematic uncertainties are included in the fit through nuisance parameters. Some uncertainties affect several samples or energy bins in a correlated way. For nuisance parameters with Gaussian constraints, the systematic penalty is

$$\chi^2_{\text{syst}}(\boldsymbol{\eta}) = (\boldsymbol{\eta} - \boldsymbol{\eta}_0)^T V^{-1} (\boldsymbol{\eta} - \boldsymbol{\eta}_0), \tag{6.5}$$

where $\boldsymbol{\eta}_0$ contains the nominal parameter values and V is their covariance matrix.

The fit minimizes the following statistic:

$$\chi^2(\boldsymbol{\omega}, \boldsymbol{\eta}) = 2 \sum_i \left[\mu_i - n_i + n_i \ln \left(\frac{n_i}{\mu_i} \right) \right] + \chi^2_{\text{syst}}(\boldsymbol{\eta}) + \chi^2_{\text{ext}}, \tag{6.6}$$

where the first term measures the agreement between the observed and predicted event counts. The second term accounts for systematic uncertainties, and χ^2_{ext} includes constraints from other experiments. The logarithmic term is defined to be zero when $n_i = 0$.

The oscillation fit includes $\sin^2 \theta_{23}$, the atmospheric mass-squared splitting, δ_{CP}, and the neutrino mass ordering. Other oscillation parameters are fixed or constrained according to the fit configuration.

For each tested oscillation hypothesis, the fit minimizes χ^2 over the remaining free parameters, including the nuisance parameters. The resulting χ^2 is compared with the overall minimum to determine how well the hypothesis agrees with the data. This comparison forms the basis for the confidence intervals discussed in the next chapter.

6.5 Confidence Intervals and the Need for Feldman–Cousins

Confidence intervals are often calculated using Wilks' theorem, which approximates the likelihood-ratio test statistic with a chi-square distribution. However, this approximation is not always valid for neutrino oscillation analyses. Some oscillation parameters have physical boundaries, δ_{CP} is periodic, and the neutrino mass ordering is discrete. Different combinations of θ_{23}, δ_{CP}, and mass ordering can also produce multiple minima in the fit [36, 37].

The approximation can also be inaccurate when the number of observed events is small, particularly in the ν_e appearance samples. As a result, the test statistic may not follow the expected chi-square distribution. Using fixed thresholds from this distribution can lead to

incorrect confidence intervals.

The Feldman–Cousins method uses simulated experiments to determine the distribution of the test statistic [38]. For each tested oscillation hypothesis, many pseudo-experiments are generated and fitted in the same way as the observed data. The resulting distribution is used to determine the critical value for a given confidence level. Chapter 7 describes the profiled Feldman–Cousins method used by NOvA [37].

The Feldman–Cousins method requires a large number of pseudo-experiments and repeated fits, making it computationally expensive. Chapter 8 introduces a Gaussian-process method to reduce this computational cost.

6.6 Connection to the Neutrino–Electron Flux Constraint

The oscillation analysis described in this chapter is based on previously published NOvA results. The neutrino–electron elastic-scattering measurement in Chapter 5 provides an additional constraint on the NuMI neutrino flux using a process with a precisely predicted cross section.

This constraint was not available for the earlier oscillation analysis. Its application to current and future NOvA oscillation analyses can provide additional information on the neutrino flux and help reduce flux-related uncertainties. Incorporating the constraint requires accounting for correlations with the flux uncertainties already included in the oscillation fit.

Chapter 7

Feldman–Cousins Confidence Intervals in Neutrino Oscillation Measurements

The NOvA oscillation analysis uses confidence intervals to report the allowed values of oscillation parameters. A common approach uses Wilks' theorem to calculate these intervals from a likelihood-ratio test statistic. This approximation does not always work for neutrino oscillation measurements. The number of events can be small, some parameters have physical boundaries, and different oscillation parameters can give similar predictions [36, 37].

The Feldman–Cousins (FC) method uses simulated experiments to determine confidence intervals [38]. NOvA uses a version called Profiled Feldman–Cousins to account for other oscillation parameters and systematic uncertainties [37]. This chapter explains the method and its computational cost. Chapter 8 describes how Gaussian processes can reduce that cost.

7.1 Likelihood-Ratio Inference

Let $\boldsymbol{\theta}$ denote the oscillation parameters being tested and $\boldsymbol{\eta}$ denote the nuisance parameters. Nuisance parameters describe systematic uncertainties and include oscillation parameters

that are not being tested. For observed data x, the likelihood is

$$\mathcal{L}(\boldsymbol{\theta}, \boldsymbol{\eta}; x). \tag{7.1}$$

The best-fit values maximize the likelihood:

$$(\hat{\boldsymbol{\theta}}, \hat{\boldsymbol{\eta}}) = \arg\max_{\boldsymbol{\theta},\boldsymbol{\eta}} \mathcal{L}(\boldsymbol{\theta}, \boldsymbol{\eta}; x). \tag{7.2}$$

To test a specific point $\boldsymbol{\theta}_0$, the fit fixes $\boldsymbol{\theta}$ at that point and finds the best values of the nuisance parameters:

$$\hat{\hat{\boldsymbol{\eta}}}_0 = \arg\max_{\boldsymbol{\eta}} \mathcal{L}(\boldsymbol{\theta}_0, \boldsymbol{\eta}; x). \tag{7.3}$$

The likelihood-ratio test statistic compares this fit with the overall best fit:

$$\Delta\chi^2(\boldsymbol{\theta}_0) = -2 \ln \frac{\mathcal{L}(\boldsymbol{\theta}_0, \hat{\hat{\boldsymbol{\eta}}}_0; x)}{\mathcal{L}(\hat{\boldsymbol{\theta}}, \hat{\boldsymbol{\eta}}; x)}. \tag{7.4}$$

A larger $\Delta\chi^2$ means the tested point agrees less well with the data than the best-fit point.

Wilks' theorem predicts a chi-square distribution for this statistic under certain conditions [36]. Those conditions do not always hold in oscillation analyses. Physical boundaries, multiple fit minima, and small event samples can change the distribution. A fixed chi-square threshold can then give confidence intervals with incorrect coverage. Coverage is the fraction of repeated experiments whose intervals include the true parameter value.

7.2 The Feldman–Cousins Construction

The FC method finds the distribution of the test statistic using pseudo-experiments. It does not assume a chi-square distribution. The calculation is repeated for each tested point $\boldsymbol{\theta}_0$.

For a given test point, the method generates many pseudo-experiments under that hypothesis. Each pseudo-experiment is fitted in the same way as the observed data. The

result is a set of test-statistic values:

$$\Delta\chi^2_k(\boldsymbol{\theta}_0), \qquad k = 1, \ldots, N_{\mathrm{PSE}}. \tag{7.5}$$

These values give the distribution

$$f\Big(\Delta\chi^2 \mid \boldsymbol{\theta}_0\Big). \tag{7.6}$$

For confidence level $1-\alpha$, the critical value $c_\alpha(\boldsymbol{\theta}_0)$ satisfies

$$P\Big[\Delta\chi^2 \leq c_\alpha(\boldsymbol{\theta}_0) \mid \boldsymbol{\theta}_0\Big] = 1-\alpha. \tag{7.7}$$

The tested point is included in the confidence region if

$$\Delta\chi^2_{\mathrm{obs}}(\boldsymbol{\theta}_0) \leq c_\alpha(\boldsymbol{\theta}_0). \tag{7.8}$$

The same pseudo-experiments can be used to calculate a p-value:

$$p(\boldsymbol{\theta}_0) = P\Big[\Delta\chi^2_{\mathrm{PSE}} \geq \Delta\chi^2_{\mathrm{obs}}(\boldsymbol{\theta}_0) \mid \boldsymbol{\theta}_0\Big], \tag{7.9}$$

where the p-value is the fraction of pseudo-experiments with a test statistic at least as large as the observed value. The point is included in the confidence region when $p(\boldsymbol{\theta}_0) \geq \alpha$, with the boundary convention set by the finite pseudo-experiment sample.

FC calculates a separate threshold at each test point. This improves the confidence interval when the chi-square approximation is inaccurate. It also requires many pseudo-experiments and fits.

7.3 Nuisance Parameters and the Profile Construction

To generate pseudo-experiments, the FC method needs predicted event spectra for each tested hypothesis. These predictions also depend on parameters that are not being tested. For example, a test of $\sin^2\theta_{23}$ still depends on other oscillation parameters and parameters describing systematic uncertainties. Their true values are unknown. Profiled FC uses the observed data to choose their values.

NOvA uses Profiled Feldman–Cousins to choose the nuisance-parameter values [37]. At each test point $\boldsymbol{\theta}_0$, the observed data are fitted with $\boldsymbol{\theta}$ fixed. The fit gives the nuisance-parameter values

$$\hat{\hat{\boldsymbol{\eta}}}_0 = \hat{\hat{\boldsymbol{\eta}}}(\boldsymbol{\theta}_0; x_{\text{obs}}). \tag{7.10}$$

The pseudo-experiments are generated using these values:

$$x_k \sim P\Big(x \mid \boldsymbol{\theta}_0, \hat{\hat{\boldsymbol{\eta}}}_0\Big)\,. \tag{7.11}$$

Each pseudo-experiment is then fitted to calculate its likelihood-ratio test statistic. These results give the distribution used to test $\boldsymbol{\theta}_0$.

This approach uses the nuisance-parameter values that best describe the observed data at the tested point. It does not simply set every nuisance parameter to its nominal value. NOvA compared different approaches using simplified models. The profiled method gave more accurate coverage in those studies [37].

Profiled FC does not guarantee exact coverage for every possible nuisance-parameter value. For example, a validation sample may be generated at $(\boldsymbol{\theta}_0, \boldsymbol{\eta}_0)$, while the FC pseudo-experiments use $(\boldsymbol{\theta}_0, \hat{\hat{\boldsymbol{\eta}}}_0)$. Coverage can be tested with simulated data for representative parameter values.

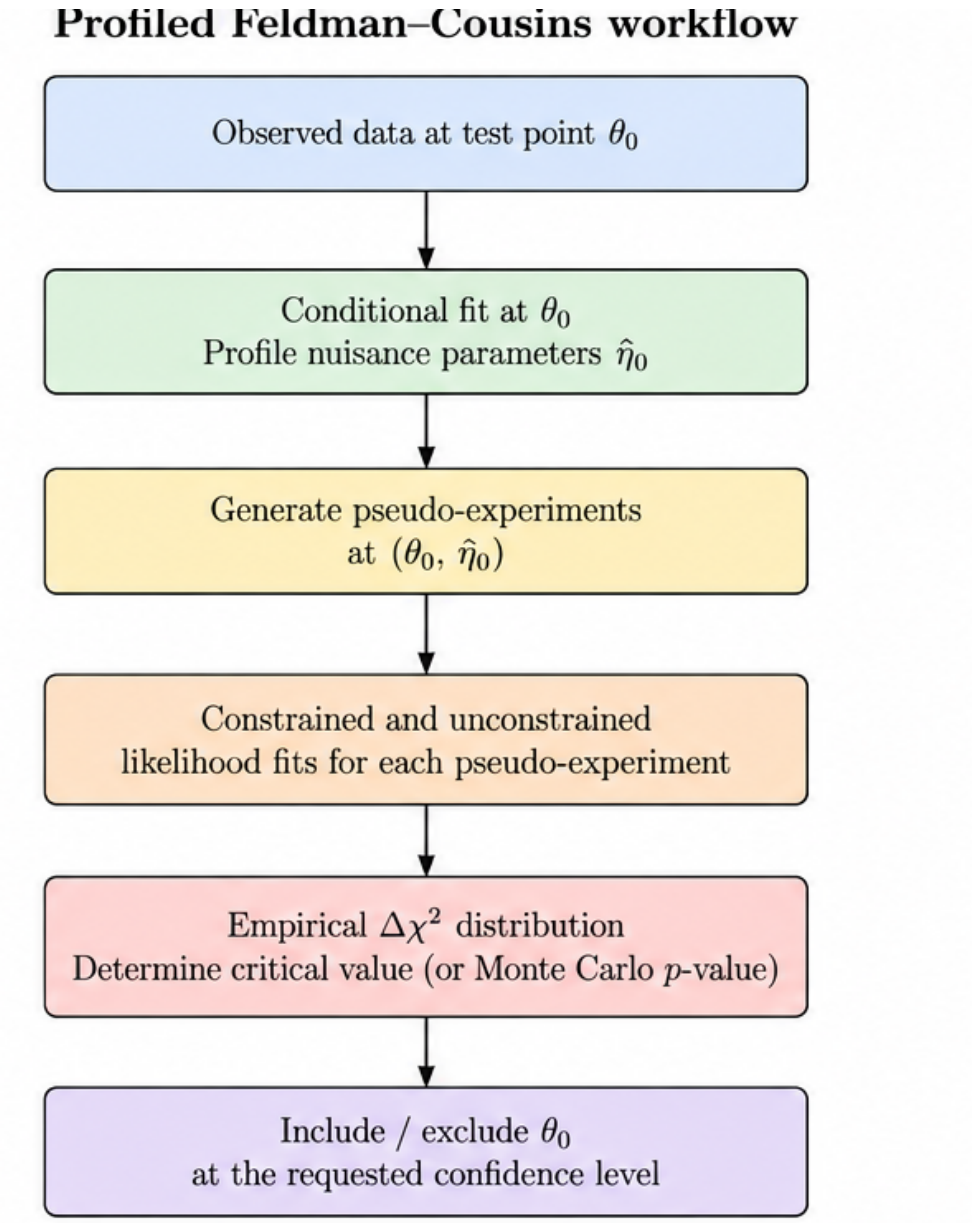

Figure 7.1: Steps in the Profiled Feldman–Cousins method. At each test point, the observed data determine the nuisance-parameter values used to generate pseudo-experiments. The diagram is based on the procedure described in Ref. [37].

7.4 Application in NOvA Oscillation Measurements

NOvA uses FC to calculate confidence intervals for oscillation parameters and to test hypotheses such as the neutrino mass ordering. A one-dimensional interval and a two-dimensional confidence region require different fits. Each calculation uses its own pseudo-experiments. A result calculated for two parameters cannot simply be reduced to a one-parameter result by minimizing the FC-corrected surface.

At each test point, the Profiled FC calculation follows these steps:

1. Fit the observed data with the tested parameters fixed. Use the fit to find the nuisance-parameter values.

2. Generate pseudo-experiments at the tested point using those values.

3. Fit each pseudo-experiment twice: once with the tested parameters fixed and once with all parameters free.

4. Use the results to form the distribution of the test statistic.

5. Compare the observed test statistic with this distribution to calculate a critical value or p-value.

Repeating the calculation across a parameter grid gives FC-corrected confidence regions. The resulting p-value surface can be used to draw confidence contours. It is not a likelihood surface. A different statistical question, such as a one-dimensional interval, needs its own FC calculation.

7.5 Computational Cost of Grid-Based FC

A direct FC calculation generates pseudo-experiments at every point in a parameter grid. For N_{grid} grid points and N_{PSE} pseudo-experiments per point, the number of pseudo-experiments

is approximately

$$N_{\text{eval}} \sim N_{\text{grid}} N_{\text{PSE}}. \tag{7.12}$$

Each pseudo-experiment also needs several fits. The total computing time can be large, especially for a grid with many points.

There are two ways to reduce the cost. First, the calculation can focus on points near the boundary of the confidence interval. Points far from the boundary may need fewer pseudo-experiments. Second, independent pseudo-experiments can run in parallel on multiple processors.

The first idea motivates Gaussian-Process-accelerated Feldman–Cousins (GP-FC). The second is used in the parallel GP-P implementation. Both methods are discussed in Chapter 8.

7.6 From Feldman–Cousins to GP-FC

GP-FC uses the same FC definition of the p-value:

$$p(\boldsymbol{\theta}) = P\Big[\Delta\chi^2_{\text{PSE}} \geq \Delta\chi^2_{\text{obs}}(\boldsymbol{\theta}) \mid \boldsymbol{\theta}\Big]. \tag{7.13}$$

The difference is how the method chooses where to run pseudo-experiments.

A standard grid-based FC calculation generates many pseudo-experiments at every grid point. GP-FC uses a Gaussian process to predict p-values at unseen points based on previously calculated results. It then runs more pseudo-experiments near the boundary of the confidence interval. This reduces the total number of pseudo-experiments needed. Chapter 8 describes the method, its parallel implementation, and its validation against a 30×30 reference grid.

7.7 Summary

The FC method uses pseudo-experiments to calculate confidence intervals when the chi-square approximation is unreliable. NOvA uses Profiled FC to choose nuisance-parameter values

from fits to the observed data. The method requires many pseudo-experiments and repeated fits. GP-FC reduces this cost by focusing calculations on points that help determine the confidence-region boundaries. Chapter 8 compares GP-FC with a reference grid calculation to check their numerical agreement. This comparison does not test statistical coverage.

Chapter 8

Gaussian-Process-Accelerated Feldman–Cousins

The Feldman–Cousins (FC) construction described in Chapter 7 determines confidence regions using pseudo-experiments. The calculation becomes expensive when repeated over a multidimensional oscillation-parameter grid. GP-FC uses a Gaussian process (GP) to predict FC significance at unseen points and select points for additional FC calculations [39]. This reduces the number of full FC calculations needed to determine the confidence contours.

For NOvA oscillation fits, the implementation combines GP regression, adaptive sampling, and parallel pseudo-experiment generation. The method is evaluated using stored FC pseudo-experiments and likelihood fits performed within the adaptive GP-FC workflow.

8.1 Gaussian-process model

8.1.1 Training Points and GP Coordinates

The GP regression uses the coordinates and FC significance estimates at evaluated points to predict the significance at unseen points. The FC significance is expressed as the confidence content associated with the pseudo-experiment test-statistic distribution.

In this implementation, the GP uses grid indices as coordinates for the parameter scan. For the two-dimensional $(\sin^2\theta_{23}, \Delta m^2_{32})$ study, the squared distance between two points with indices (i, j) and (i', j') is

$$D_{ij,i'j'} = (i - i')^2 + (j - j')^2. \tag{8.1}$$

8.1.2 Kernel and Observation Uncertainty

The GP uses a kernel to describe the correlation between significance values at different points in the parameter space. We use a rational-quadratic (RQ) kernel because it can describe variations over multiple length scales. This is useful for FC significance surfaces, which may change at different rates across the parameter space.

The RQ kernel is defined as

$$k(\mathbf{x}, \mathbf{x}') = \left[1 + \frac{D(\mathbf{x}, \mathbf{x}')}{2\alpha\ell^2}\right]^{-\alpha}, \tag{8.2}$$

where D is defined in Eq. 8.1. The parameter ℓ controls the correlation length, and α controls the mixture of length scales. The initial values are $\ell = 4.5$ and $\alpha = 0.5$.

The FC significance at each evaluated point is estimated from a finite number of pseudo-experiments. The statistical uncertainty is estimated from the width of a binomial proportion confidence interval. Including this uncertainty prevents the GP from treating the estimated significance values as exact measurements.

For an estimated significance p from N pseudo-experiments, the interval width used as the observation uncertainty is

$$\sigma_p = \frac{2z}{1 + z^2/N}\sqrt{\frac{p(1-p)}{N} + \frac{z^2}{4N^2}}, \qquad z = 1. \tag{8.3}$$

The corresponding variance σ_p^2 is added to the diagonal of the GP covariance matrix.

8.1.3 Hyperparameter Optimization and Prediction

The kernel hyperparameters control how the GP describes correlations between points. Their values affect both the predicted significance and its uncertainty.

The hyperparameters are determined by maximizing the marginal likelihood of the evaluated FC significance values. For a GP with zero prior mean, the negative log marginal likelihood is

$$-2 \log L = \mathbf{y}^T K^{-1} \mathbf{y} + \log |K| + n \log(2\pi), \tag{8.4}$$

where $\mathbf{y}$ contains the FC significance estimates, K is the training covariance matrix, and n is the number of evaluated points. The covariance matrix includes the observation uncertainties described in the previous section.

The hyperparameters are optimized each time the GP is updated with new FC results.

Once the hyperparameters are determined, the GP predicts the significance and its uncertainty at other points in the parameter space. For a zero-mean GP, the predicted mean is

$$\boldsymbol{\mu}_* = K_* K^{-1} \mathbf{y}, \tag{8.5}$$

and the predicted covariance is

$$\Sigma_* = K_{**} - K_* K^{-1} K_*^T, \tag{8.6}$$

where K_* is the covariance between prediction points and evaluated points, and K_{**} is the covariance between prediction points.

The predicted mean gives the estimated FC significance at each point. The diagonal entries of Σ_* give the corresponding prediction variances.

The GP predictions and their uncertainties are used to select points for additional FC calculations, as described in the following section.

8.2 Adaptive Selection of FC Calculations

The GP predicts the FC significance and its uncertainty at each grid point. These predictions are used to select points for additional FC calculations.

The selection is based on an acquisition function that gives higher priority to points near the confidence-region boundaries. For a target confidence level C, the priority is defined as

$$A_i(C) = \left| \frac{\sigma_i}{\mu_i - C} \right|, \tag{8.7}$$

where μ_i and σ_i are the GP predicted significance and its standard deviation at grid point i.

A point receives higher priority when its predicted significance is close to the target confidence level or when its prediction has a large uncertainty. This allows the algorithm to focus on points that can help determine the confidence-region boundaries.

The acquisition priorities are evaluated for the 1σ, 2σ, and 3σ confidence boundaries. The corresponding confidence contents are

$$C_{1\sigma} = 0.6827, \tag{8.8}$$

$$C_{2\sigma} = 0.9545, \tag{8.9}$$

$$C_{3\sigma} = 0.9973. \tag{8.10}$$

The priorities associated with these boundaries guide the selection of points for additional FC calculations. At each iteration, additional pseudo-experiments are generated at the selected grid points. The new FC significance estimates are added to the existing results, and the GP is updated. The acquisition function is then recalculated to select points for the next iteration.

This process concentrates FC calculations on the points that are most useful for determining the confidence-region boundaries.

8.3 Validation Configuration

The GP-FC reconstruction is tested using a scan of the $(\sin^2\theta_{23}, \Delta m^2_{32})$ parameter space under the normal hierarchy. The scan covers

$$0.42 \leq \sin^2\theta_{23} \leq 0.61 \tag{8.11}$$

and

$$2.30 \leq \frac{\Delta m^2_{32}}{10^{-3}\,\mathrm{eV}^2} \leq 2.70. \tag{8.12}$$

The reference FC calculation contains a 30×30 grid, corresponding to 900 points in the parameter space.

The reference FC results are used to evaluate how well the GP-FC method reconstructs the confidence contours using fewer pseudo-experiments. Since the number of available pseudo-experiments varies across the grid, the fraction used by GP-FC is defined as

$$f_{\mathrm{GP}} = \frac{\sum_i N_{i,\mathrm{used}}}{\sum_i N_{i,\mathrm{reference}}}, \tag{8.13}$$

where $N_{i,\mathrm{used}}$ is the number of pseudo-experiments used by GP-FC at grid point i, and $N_{i,\mathrm{reference}}$ is the number available in the reference FC calculation.

This fraction measures the total pseudo-experiment usage relative to the reference calculation. For example, $f_{\mathrm{GP}} = 30\%$ means that GP-FC uses 30% of the available reference pseudo-experiments to reconstruct the confidence contours.

8.4 Parallel Implementation

Although GP-FC reduces the number of pseudo-experiments needed to construct confidence regions, the remaining FC calculations can still be computationally expensive. Each pseudo-experiment requires likelihood fits to determine the test statistic, and many pseudo-

experiments are needed to estimate the significance at a given parameter point.

To reduce the computation time, the GP-FC method is implemented using distributed parallel computing with MPI and DIY. The GP regression identifies the parameter-space points that require additional FC calculations, while the parallel framework distributes these calculations across available computing resources.

The work is divided into independent batches of pseudo-experiments. Multiple batches can be assigned to the same parameter point, allowing the computational resources to be concentrated in regions selected by the GP. The pseudo-experiments are evaluated independently, and their results are combined to update the FC significance estimates.

After each iteration, the updated FC results are used to retrain the GP and select points for the next iteration. The accumulated results are also saved, allowing the calculation to resume without repeating completed pseudo-experiments.

This approach combines two sources of computational efficiency. Adaptive sampling reduces the number of FC calculations, while parallel processing reduces the time needed to perform the remaining calculations.

8.4.1 Runtime Control for Real Fits

The computation time of individual pseudo-experiments can vary considerably because each requires numerical likelihood minimization. Most fits finish within a relatively short time, but a small number can take much longer. These long-running fits can delay the completion of a parallel iteration, even after most of the assigned work has finished.

To study this effect, the runtime of 39,680 pseudo-experiments was measured using Perlmutter CPU nodes. Figure 8.1 shows the resulting distribution.

Based on this study, a runtime limit of 1800 s was selected. This limit allows most pseudo-experiments to finish while reducing delays caused by unusually long-running fits.

A runtime limit of 1800 s is applied to each pseudo-experiment to reduce delays from unusually long-running likelihood fits. Only completed pseudo-experiments are included in

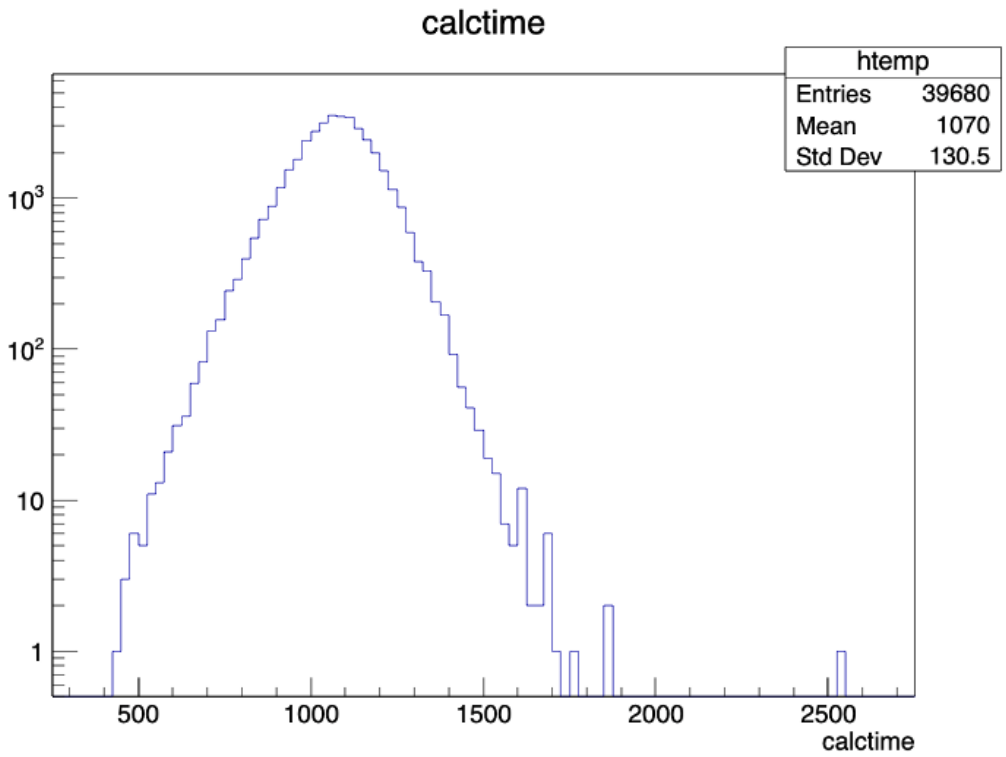

Figure 8.1: Distribution of computation times for 39,680 pseudo-experiments on Perlmutter CPU nodes. The mean runtime is 1070 s, with a standard deviation of 130.5 s.

the FC significance estimate.

8.5 Stored-Pseudo-Experiment Tests

The GP-FC method reduces the computational cost of FC confidence interval construction by concentrating pseudo-experiments on selected points in the parameter space. However, the number of pseudo-experiments assigned to each selected point can also affect the efficiency of the reconstruction.

In a uniform allocation, all selected points receive the same number of pseudo-experiments. Although this approach is straightforward, the selected points do not necessarily contribute equally to determining the confidence-region boundaries. Points with higher GP priorities may benefit from additional pseudo-experiments, while fewer may be sufficient at points with lower priorities.

The parallel implementation allows the number of pseudo-experiments to vary among selected points. This provides additional flexibility in distributing computational resources according to the GP predictions. Different allocations can be explored by changing both the number of pseudo-experiments per point and the proportion of selected points receiving a larger allocation.

To investigate these strategies, we performed tests using stored FC pseudo-experiments. Reusing existing FC results allows different configurations to be compared without repeating the computationally expensive likelihood fits. This is particularly useful given the limited computing resources available for different analyses.

The tests use the reference FC sample described in Sec. 8.3. At each iteration, the GP selects 64 points from the 900-point parameter grid for additional FC calculations. The stored-pseudo-experiment tests examine how different allocations affect the reconstruction of the confidence contours as more pseudo-experiments are incorporated.

Figure 8.2 shows the early iterations of a uniform allocation of 2000 pseudo-experiments per selected point. At iteration 0, the GP fit fraction is 15.23%. The reconstructed 1σ and 2σ contours already reproduce the main features of the reference FC result, while larger differences remain along the 3σ contour.

At iteration 1, the GP fit fraction increases to 30.36%, and the reconstructed contours become closer to the reference result. Further iterations mainly improve local differences along the confidence-region boundaries.

Two nonuniform allocation strategies were also investigated. In both configurations, the selected points are ranked according to their GP priorities. The higher-priority points receive more pseudo-experiments than the remaining selected points.

In the 1200/800 configuration, the 48 highest-priority points receive 1200 pseudo-experiments each, while the remaining 16 points receive 800 each. This corresponds to 70,400 pseudo-experiments per iteration.

In the 1600/800 configuration, the 32 highest-priority points receive 1600 pseudo-experiments

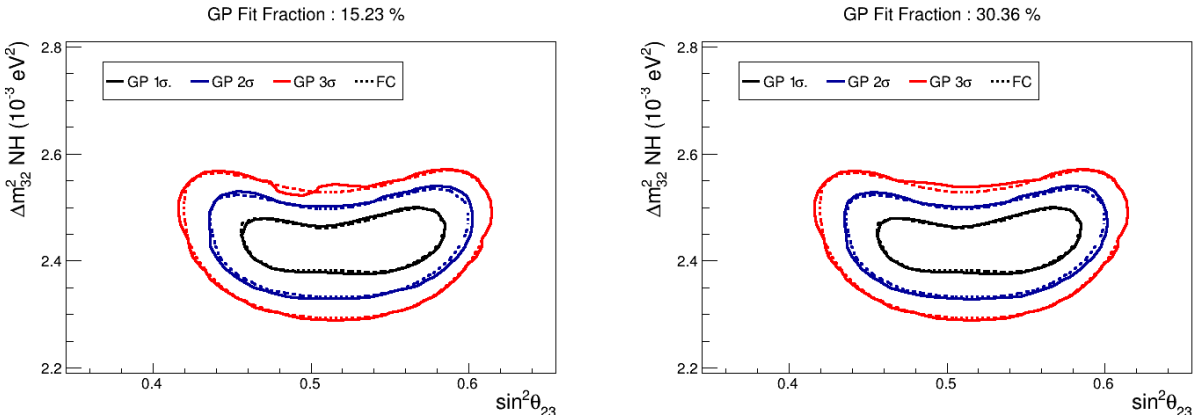

Figure 8.2: Evolution of the reconstructed FC confidence contours using a uniform allocation of 2000 pseudo-experiments per selected point. The GP fit fraction increases from 15.23% at iteration 0 (left) to 30.36% at iteration 1 (right). Solid curves show the GP-FC reconstruction, while dotted curves show the reference FC contours.

each, while the remaining 32 points receive 800 each. This corresponds to 76,800 pseudo-experiments per iteration.

These configurations explore different ways of concentrating computational resources on high-priority points. The 1200/800 configuration distributes additional pseudo-experiments across a larger group of selected points, while the 1600/800 configuration assigns more pseudo-experiments to a smaller group.

Figure 8.3 compares the reconstructed contours at iterations 0 and 5 for the two nonuniform allocations.

For the 1200/800 configuration, the GP fit fraction increases from 5.45% at iteration 0 to 32.35% at iteration 5. For the 1600/800 configuration, the corresponding fractions are 6.56% and 38.40%.

Both configurations reproduce the main confidence-region structure in the early iterations. As additional pseudo-experiments are incorporated, the reconstructed contours become closer to the reference FC result, particularly along the outer boundary.

These tests show that the GP-FC method can reconstruct the main features of the FC confidence regions using only a fraction of the available reference pseudo-experiments. They also demonstrate the flexibility of the parallel implementation in assigning different numbers

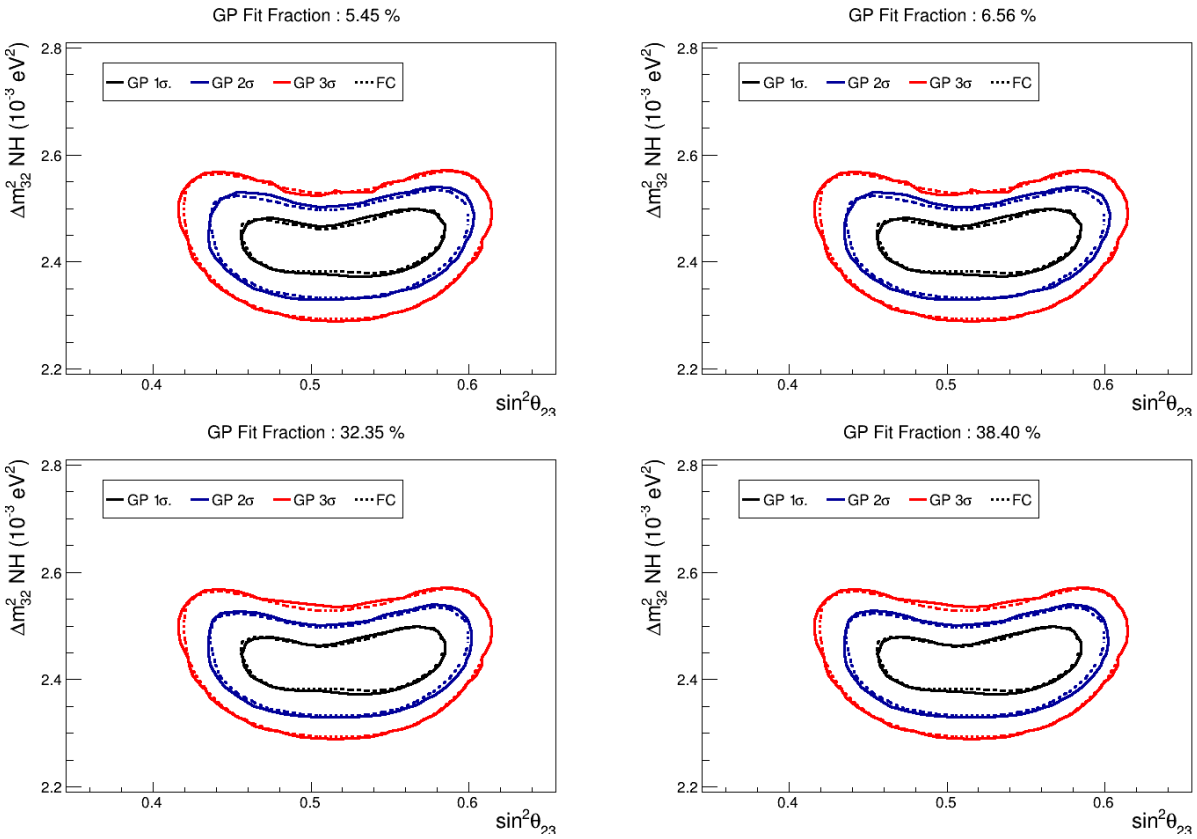

Figure 8.3: Evolution of the reconstructed FC confidence contours using nonuniform pseudo-experiment allocations. Left: 1200/800 configuration, with 48 points receiving 1200 pseudo-experiments and 16 points receiving 800 pseudo-experiments per iteration. The GP fit fractions are 5.45% at iteration 0 and 32.35% at iteration 5. Right: 1600/800 configuration, with 32 points receiving 1600 pseudo-experiments and 32 points receiving 800 pseudo-experiments per iteration. The corresponding GP fit fractions are 6.56% and 38.40%. Solid curves represent the GP-FC reconstruction, while dotted curves show the reference FC contours.

of pseudo-experiments according to the GP priorities.

The allocation strategies use different numbers of pseudo-experiments per iteration. The GP fit fraction provides a common measure of sample usage when comparing their reconstruction performance.

The GP fit fraction measures the number of pseudo-experiments used relative to the reference FC sample. It does not directly represent a reduction in elapsed computation time, which also depends on the cost of individual likelihood fits and the overhead of parallel processing.

8.6 Validation Metrics

The accuracy of the GP-FC reconstruction is evaluated by comparing its confidence regions with those obtained from the reference FC calculation.

For a confidence content C, let F_C and G_C denote the sets of grid points inside the reference FC and reconstructed GP-FC confidence regions, respectively.

The reference-region overlap fraction is defined as

$$O_C = \frac{|F_C \cap G_C|}{|F_C|}, \tag{8.14}$$

which measures the fraction of the reference FC region recovered by GP-FC.

The intersection over union (IoU) is defined as

$$J_C = \frac{|F_C \cap G_C|}{|F_C \cup G_C|}. \tag{8.15}$$

The overlap fraction measures how completely GP-FC reproduces the reference region, while IoU also accounts for additional bins included in the reconstructed region. Together, these metrics provide a quantitative comparison of the confidence regions at the 1σ, 2σ, and 3σ confidence levels.

The contour comparison also reports overall overlap and IoU values obtained without selecting an individual 1σ, 2σ, or 3σ histogram. These overall values are evaluated separately from the confidence-level-specific comparisons.

The GP posterior uncertainty provides an additional diagnostic of the reconstructed contours. The predicted significance surface is shifted by ± 1 posterior standard deviation, and the corresponding contours are recalculated. These variations illustrate the sensitivity of the reconstructed boundaries to the GP prediction uncertainty. They do not represent frequentist confidence bands.

The comparison is performed on the 30×30 parameter grid, where small changes in the contour boundaries can affect the overlap and IoU. These metrics quantify agreement with the reference FC regions but do not independently establish frequentist coverage.

8.7 Tests with Real Likelihood Fits

The GP-FC method was further tested with NOvA likelihood fits performed during the adaptive iterations. At each iteration, the GP identifies parameter-space points for additional FC calculations, and the corresponding pseudo-experiments are generated and evaluated in parallel. The resulting FC significance estimates are then used to update the GP model for the next iteration.

The calculation was performed on Perlmutter at NERSC using 100 CPU nodes, providing 25,600 parallel execution slots per iteration. Five iterations were completed in three successive runs. The first two runs each required approximately 220 node-hours, and the final run required approximately 110 node-hours. The total computational cost was approximately 550 node-hours. The calculation was resumed between runs without repeating the completed FC calculations.

Figure 8.4 compares the reconstructed confidence contours with the reference FC results at three stages of the calculation.

At iteration 0, GP-FC uses 15.23% of the reference pseudo-experiments. The reconstructed confidence regions already reproduce the main features of the reference FC contours, with an overall bin overlap of 98.65% and an intersection over union (IoU) of 96.56%.

At iteration 2, the GP fit fraction increases to 24.69%, while the overall bin overlap reaches 100% and the IoU increases to 97.88%. The reconstructed regions thus show close agreement with the reference using a relatively small fraction of its pseudo-experiments.

Further iterations refine the reconstructed contours. At iteration 4, the GP fit fraction reaches 53.90%, with an overall bin overlap of 100% and an IoU of 98.14%. The 3σ region reaches an IoU of 99.49%, showing close agreement along the outer confidence-region boundary.

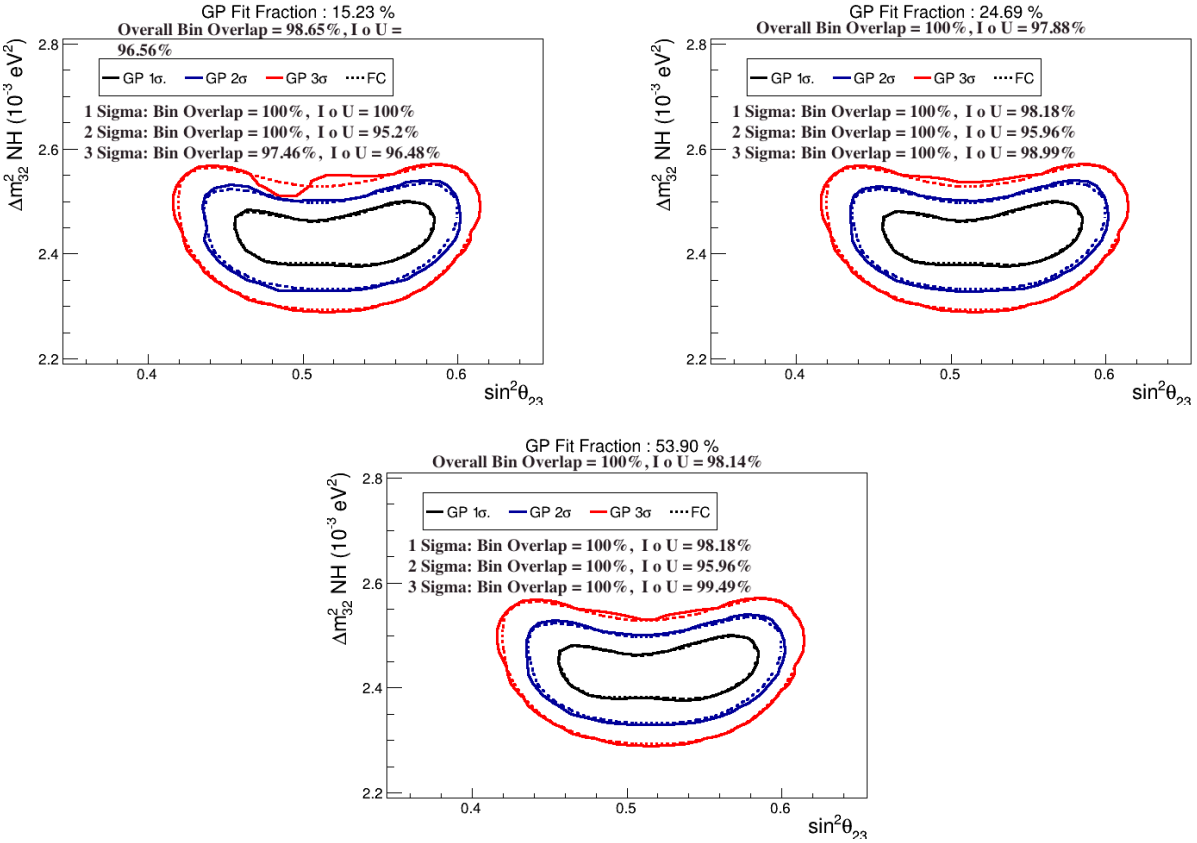

Figure 8.4: Comparison of the GP-FC reconstructed confidence contours with the reference FC results at iterations 0 (top left), 2 (top right), and 4 (bottom). The corresponding GP fit fractions are 15.23%, 24.69%, and 53.90%. Solid curves represent the GP-FC reconstruction, while dotted curves show the reference FC contours.

The overall IoU remains close to 98% in the later iterations, indicating that the recon-

structed confidence regions remain consistent with the reference FC results as additional pseudo-experiments are incorporated.

These tests demonstrate that the GP-FC method can reproduce the reference FC confidence regions using substantially fewer pseudo-experiments than the full reference calculation. The successful execution of the adaptive iterations also validates the integration of GP-based point selection with distributed FC likelihood fits on a large-scale computing system.

8.8 Summary

This chapter presented a parallel implementation of the GP-accelerated Feldman–Cousins method for neutrino oscillation analyses. Gaussian-process regression guides the adaptive selection of parameter-space points, while distributed computing enables FC pseudo-experiments to be evaluated in parallel.

The implementation was validated using NOvA oscillation fits. Tests with stored pseudo-experiments demonstrated the flexibility of different allocation strategies, while real likelihood fits verified the complete adaptive workflow. Using 24.69% of the reference pseudo-experiments, GP-FC achieved an overall intersection over union of 97.88% with the reference FC confidence regions.

The real-fit calculation was performed on NERSC using 100 CPU nodes per iteration, completing five iterations with approximately 550 node-hours. These results demonstrate that GP-FC can reproduce the reference FC confidence regions using substantially fewer pseudo-experiments. The parallel implementation provides a practical approach to reducing the computational demands of FC confidence interval construction in neutrino oscillation analyses.

Chapter 9

The DUNE Experiment

The Deep Underground Neutrino Experiment (DUNE) is a long-baseline neutrino experiment designed to study neutrino oscillations with a high-intensity accelerator beam and large underground detectors. Its main physics goals include determining the neutrino mass ordering, searching for CP violation in neutrino oscillations, and improving the precision of oscillation parameter measurements. DUNE will also study neutrinos from core-collapse supernovae and search for rare processes such as proton decay [6].

DUNE will direct a neutrino beam from Fermilab toward the Sanford Underground Research Facility (SURF) in South Dakota, approximately 1300 km away. The long baseline provides sensitivity to matter effects in neutrino propagation, which differ between the two mass orderings. Measurements in neutrino and antineutrino beam modes will help constrain the CP-violating phase δ_{CP}. The experiment combines a near detector at Fermilab with a far detector based on liquid-argon time projection chambers (LArTPCs) [6].

The far-detector technology is being developed and tested through the ProtoDUNE program at CERN. These full-scale prototypes provide experience with detector construction, cryogenic operation, charge readout, and calibration before deployment in DUNE. This chapter introduces the experimental design and the liquid-argon technology relevant to the purity-monitor work discussed in Chapter 10.

9.1 Experimental Design

9.1.1 Neutrino Beam and Near Detector

DUNE will use the Long-Baseline Neutrino Facility (LBNF) beam at Fermilab, directed toward the far detector at SURF. Unlike the off-axis NuMI beam configuration used by NOvA, DUNE is designed to measure an on-axis neutrino beam with a broad energy spectrum. This allows the experiment to study neutrino oscillations across the first and second oscillation maxima over its 1300 km baseline [6].

The near detector, located at Fermilab, will characterize the unoscillated neutrino beam and measure neutrino interactions on argon. These measurements provide constraints on the flux and interaction models used in the oscillation analysis [6].

9.1.2 Far Detector

The DUNE far detector will be located approximately 1.5 km underground at SURF. It is designed to contain four liquid-argon detector modules, each with a fiducial mass of approximately 10 kt. The large target mass provides sensitivity to accelerator neutrino oscillations, supernova neutrinos, and rare processes such as proton decay. The underground location reduces cosmic-ray backgrounds, which is particularly important for non-beam physics measurements [6].

The far detector uses liquid-argon time projection chambers (LArTPCs) to reconstruct charged particles produced in neutrino interactions. Compared with the segmented liquid-scintillator detectors used in NOvA, LArTPCs provide finer spatial information about particle tracks and electromagnetic showers. This enables detailed event reconstruction and particle identification. The large liquid-argon volume also requires stable cryogenic operation and high liquid-argon purity.

9.2 Liquid-Argon Time Projection Chambers

Liquid argon serves as both the neutrino interaction target and the detection medium in a LArTPC. Charged particles traveling through the liquid produce ionization electrons and scintillation light. An applied electric field transports the ionization electrons toward the anode, where segmented electrodes measure the collected charge. The readout geometry provides the coordinates along the anode plane, while the electron drift time, combined with an interaction-time reference, determines the position along the drift direction [6, 8].

The DUNE horizontal-drift design uses cathode structures and anode plane assemblies to form large drift regions. Ionization electrons travel through the liquid argon under an electric field of approximately 500 V/cm. The combination of charge and scintillation-light measurements allows the reconstruction of particle trajectories and energy deposition.

Electronegative impurities, including oxygen and water, can capture drifting electrons and reduce the collected charge. Under approximately uniform conditions, the surviving charge is described by

$$Q(t) = Q_0 \exp\left(-\frac{t}{\tau_e}\right), \tag{9.1}$$

where Q_0 is the charge before attachment, t is the drift time, and τ_e is the effective electron lifetime. A longer electron lifetime corresponds to less charge loss during drift.

The large drift distances in DUNE make liquid-argon purity important for maintaining detector performance. Cryogenic circulation and filtration reduce impurity concentrations, while dedicated purity monitors measure electron lifetime during detector commissioning and operation.

9.3 The ProtoDUNE Program

The ProtoDUNE program at the CERN Neutrino Platform tests full-scale detector components developed for the DUNE far detector. The prototypes provide an environment for studying

LArTPC performance, detector calibration, and the operation of large cryogenic systems. The program includes ProtoDUNE-SP, the original single-phase prototype, and ProtoDUNE-HD, which tests updated horizontal-drift detector technology developed for DUNE [7, 8, 40].

9.3.1 ProtoDUNE-SP

ProtoDUNE-SP was constructed to demonstrate the single-phase LArTPC technology at a scale comparable to the DUNE far-detector components. The detector contained approximately 770 t of liquid argon, with an active volume of approximately $7 \times 6 \times 7.2$ m^3. A central cathode divided the TPC into two drift regions, with ionization electrons collected by anode plane assemblies on either side. The detector operated at the CERN Neutrino Platform during its Phase-I program [7, 8].

ProtoDUNE-SP collected data using charged particle beams and cosmic rays to study detector response and reconstruction. These measurements provided information about charge collection, particle identification, and calorimetric performance in a large LArTPC. The detector also demonstrated the operation of high-voltage, cryogenic, and purification systems at full scale [7].

Three ultraviolet (UV) purity monitors were installed at different heights inside the cryostat, outside the TPC field cage. They measured electron lifetime during liquid-argon purification and detector operation, providing information about liquid-argon purity at different locations within the cryostat [8, 40].

9.3.2 ProtoDUNE-HD

ProtoDUNE-HD uses the ProtoDUNE cryostat with detector components developed for the DUNE horizontal-drift far detector. The detector follows the same basic LArTPC principle demonstrated by ProtoDUNE-SP, with updated charge-readout and photon-detection systems. Its operation provides further tests of detector integration, cryogenic performance, and the technologies planned for DUNE [40].

The horizontal-drift configuration uses a central cathode and anode structures to form two drift regions. The detector tests the performance of these components under operating conditions relevant to the DUNE far detector. The measurements support the development of detector calibration and reconstruction methods for the full-scale experiment.

Liquid-argon purity is monitored throughout detector commissioning and operation. The UV purity monitors provide regular electron-lifetime measurements, allowing changes in argon purity to be followed over time.

9.4 Liquid-Argon Purity Monitoring

The electron lifetime characterizes charge attenuation during drift and is an important quantity for calibrating the response of a LArTPC. It can be measured using reconstructed TPC tracks or dedicated purity monitors, which provide measurements independent of TPC event reconstruction [40].

The UV purity monitors used in ProtoDUNE measure electron survival over a short drift distance. Their measurements depend on the local liquid-argon conditions and the electric field within the monitor. These measurements complement TPC-based lifetime measurements and provide a means of monitoring liquid-argon purity during detector operation. The design, calibration, and operation of the ProtoDUNE UV purity monitors are discussed in Chapter 10.

Chapter 10

Liquid-Argon Purity Monitoring in ProtoDUNE

The charge collected in a liquid-argon time projection chamber depends on the survival of ionization electrons during drift. Electron-lifetime measurements are used to assess liquid-argon purity during detector commissioning and operation. UV purity monitors provide these measurements independently of TPC track reconstruction, allowing the purity conditions to be monitored throughout detector operation.

This chapter presents my work on the ProtoDUNE UV purity-monitor program, including hardware design and preparation, waveform analysis, charge calibration, long-drift monitor development, and automated monitoring for ProtoDUNE-HD.

10.1 UV Purity-Monitor Measurement Principle

A UV purity monitor uses light from a xenon flash source to release photoelectrons from an Al–Ti–Au photocathode. The electrons pass through a cathode grid, drift through liquid argon under a controlled electric field, and pass an anode grid before collection. Field rings define the drift region, while the grids shield the electrode signals from charge moving through the central volume [8, 40, 41].

The cathode and anode signals provide the initial and surviving charge, Q_{c} and Q_{a}. After instrumental corrections, their ratio is

$$R \equiv \frac{Q_{\mathrm{a}}}{Q_{\mathrm{c}}} = \exp\left(-\frac{t_{\mathrm{d}}}{\tau_e}\right), \tag{10.1}$$

where t_{d} is the measured drift time. The electron lifetime follows from

$$\tau_e = -\frac{t_{\mathrm{d}}}{\ln R}. \tag{10.2}$$

The charge ratio is measured directly, and the electron lifetime is calculated from the ratio and drift time. The measurement precision depends on charge extraction, grid transparency, and the response of the readout channels.

10.2 ProtoDUNE-SP Purity Monitors

10.2.1 Detector Design and Operation

Three short UV purity monitors were installed inside the ProtoDUNE-SP-I cryostat, outside the TPC field cage, at heights of approximately 1.8, 3.7, and 5.6 m above the cryostat floor. Their locations allowed electron-lifetime measurements at different heights in the liquid argon. Each monitor was approximately 25 cm long and consisted of a photocathode, a cathode grid, field rings, an anode grid, and an anode. The distance between the two grids was 15.74 cm for the top and bottom monitors and 14.75 cm for the middle monitor [8].

CAD drawings were prepared to specify the dimensions and assembly of the electrodes, grids, field rings, and mechanical supports. These drawings guided component fabrication and standardized the assembly and testing procedures. Figure 10.2 shows the assembly drawing of the short UV purity monitor.

The monitors were operated at different voltages to adjust the electron drift time as the liquid-argon purity changed. Anode and cathode voltages of $(250, -50)$ V and $(500, -50)$ V

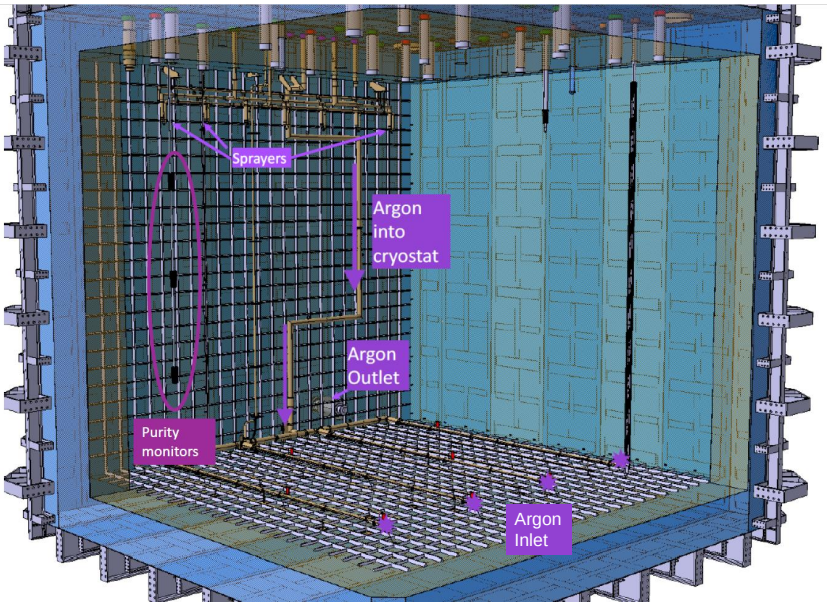

Figure 10.1: Locations of the three UV purity monitors inside the ProtoDUNE-SP cryostat, at different heights outside the TPC field cage [8].

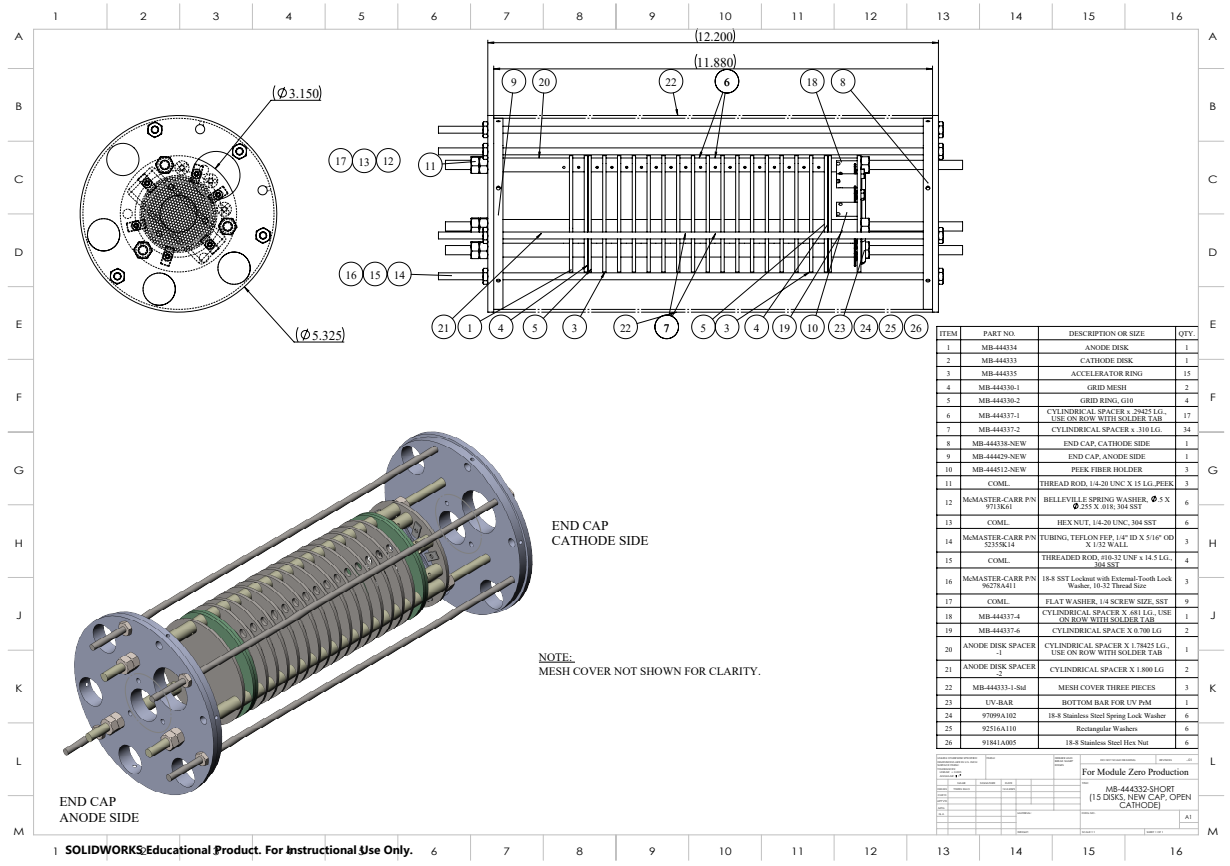

ITEM	PART NO.	DESCRIPTION OR SIZE	QTY.
1	MB-444334	ANODE DISK	1
2	MB-444333	CATHODE DISK	1
3	MB-444335	ACCELERATOR RING	15
4	MB-444330-1	GRID MESH	2
5	MB-444330-2	GRID RING, G10	4
6	MB-444337-1	CYLINDRICAL SPACER x .29425 LG., USE ON ROW WITH SOLDER TAB	17
7	MB-444337-2	CYLINDRICAL SPACER x .310 LG.	34
8	MB-444338-NEW	END CAP, CATHODE SIDE	1
9	MB-444429-NEW	END CAP, ANODE SIDE	1
10	MB-444512-NEW	PEEK FIBER HOLDER	3
11	COML.	THREAD ROD, 1/4-20 UNC X 15 LG.,PEEK	3
12	McMASTER-CARR P/N 9713K61	BELLEVILLE SPRING WASHER, Ø.5 X Ø.255 X .018, 304 SST	6
13	COML.	HEX NUT, 1/4-20 UNC, 304 SST	6
14	McMASTER-CARR P/N 52355K14	TUBING, TEFLON FEP, 1/4" ID X 5/16" OD X 1/32 WALL	3
15	COML.	THREADED ROD, #10-32 UNF x 14.5 LG., 304 SST	4
16	McMASTER-CARR P/N 96278A411	18-8 SST Lockout with External-Tooth Lock Washer, 10-32 Thread Size	3
17	COML.	FLAT WASHER, 1/4 SCREW SIZE, SST	9
18	MB-444337-4	CYLINDRICAL SPACER X .681 LG., USE ON ROW WITH SOLDER TAB	1
19	MB-444337-6	CYLINDRICAL SPACE X 0.700 LG	2
20	ANODE DISK SPACER -1	CYLINDRICAL SPACER X 1.78425 LG., USE ON ROW WITH SOLDER TAB	1
21	ANODE DISK SPACER -2	CYLINDRICAL SPACER X 1.800 LG	2
22	MB-444333-1-Std	MESH COVER THREE PIECES	3
23	UV-BAR	BOTTOM BAR FOR UV PrM	1
24	97099A102	18-8 Stainless Steel Spring Lock Washer	6
25	92516A110	Rectangular Washers	6
26	91841A005	18-8 Stainless Steel Hex Nut	6

Figure 10.2: Engineering assembly drawing of the short UV purity monitor, showing the electrodes, grids, field rings, and mechanical supports.

produced drift fields of approximately 15 and 30 V/cm, corresponding to drift times of approximately 2.3 and 1.2 ms. The longer drift time increased charge attenuation and improved sensitivity when the electron lifetime was large. At the lower voltage ratio, the cathode grid was not fully transparent to drifting electrons, requiring an additional correction to the measured charge ratio [8].

The optical system used an external xenon lamp to illuminate the photocathode through quartz fibers. Eight fibers delivered light to each monitor, increasing the photoelectron signal compared with a single-fiber configuration [8].

During ProtoDUNE-SP-I operation, changes in the cathode and anode responses were examined alongside the drift time and high-voltage settings to distinguish changes in liquid-argon purity from instrumental effects. Following a pump-related interruption in 2019, the middle monitor showed unstable timing and anode response. A resistance measurement gave approximately 250 MΩ, compared with approximately 696 MΩ at installation.

The investigation suggested that moisture near the high-voltage feedthrough could affect the monitor response. Although the available measurements did not establish humidity as the cause of the instability, the investigation motivated additional high-voltage feedthrough testing during the development of subsequent purity monitors.

10.2.2 Waveform Analysis

Each ProtoDUNE-SP-I purity-monitor run consisted of 200 xenon-lamp flashes recorded over approximately 40 s. The cathode and anode waveforms were averaged separately to reduce fluctuations before charge extraction. The signals were amplified by separate integrating amplifiers and digitized at 2 MS/s [8].

The amplifier response modifies the shape of the recorded signals. In particular, the waveform maximum depends on both the collected charge and the rise time of the signal. The charge amplitude was obtained by fitting a model that describes the signal development and the RC response of the readout electronics.

The waveform model includes the signal amplitude V_0, start time t_{start}, rise time t_{rise}, and the RC time constant. It also accounts for the secondary-peak contribution associated with the signal induced on the anode before the electrons pass through the anode grid, together with a residual baseline. The fitted amplitude V_0 is proportional to the charge measured by the corresponding readout channel.

The uncorrected charge ratio is calculated from the fitted amplitudes,

$$R_{\mathrm{raw}} = \frac{V_{0,\mathrm{a}}}{V_{0,\mathrm{c}}}, \tag{10.3}$$

and the drift time is obtained from the fitted timing parameters,

$$t_{\mathrm{d}} = t_{\mathrm{start,a}} - t_{\mathrm{start,c}} - \frac{t_{\mathrm{rise,c}}}{2} + \frac{t_{\mathrm{rise,a}}}{2}. \tag{10.4}$$

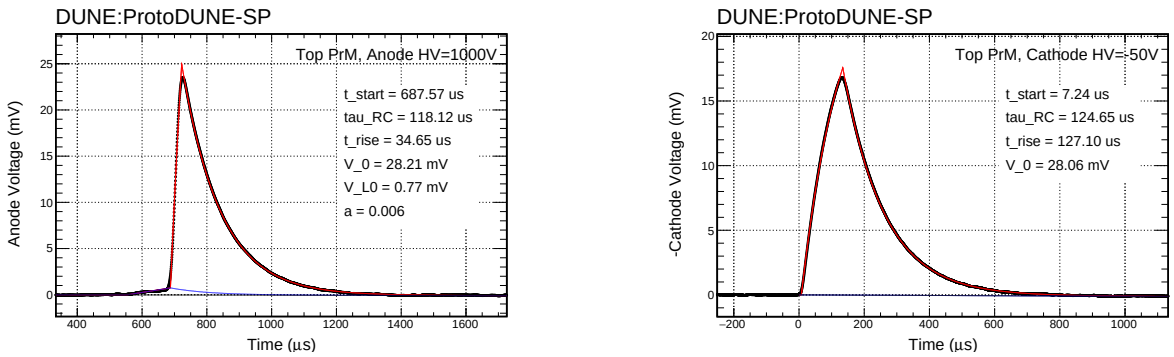

Figure 10.3: Representative fitted anode (left) and cathode (right) waveforms. The fitted amplitudes and timing parameters are used to calculate the charge ratio and drift time [40].

The offline processing uses zero-voltage runs to characterize and subtract the readout baseline. An FFT filter is applied to an identified noise component in the middle-monitor configuration. These steps reduce electronic contributions to the fitted amplitudes and timing parameters.

10.2.3 Charge Calibration and Corrections

The raw charge ratio is affected by differences between the cathode and anode readout channels, residual secondary-peak effects, and incomplete grid transparency. The corrected charge ratio is written as

$$R_{\mathrm{corr}} = R_{\mathrm{raw}}\, C_{\mathrm{gain}}\, C_{\mathrm{2nd}}\, C_{\mathrm{grid}}, \tag{10.5}$$

where C_{gain} corrects the relative readout response, C_{2nd} accounts for the residual secondary-peak contribution after waveform fitting, and C_{grid} corrects incomplete grid transparency when applicable. The correction factors depend on the monitor and readout configuration.

The relative gain was measured by swapping the cathode and anode readout channels and comparing their responses. The secondary-peak contribution was included in the waveform fit, but a small residual effect remained. This residual was evaluated through calibration measurements and corrected using C_{2nd}. The uncertainty associated with the correction was included in the charge-ratio uncertainty.

Table 10.1 summarizes the calibration factors used in the later NP04 and NP02 processing.

Table 10.1: Calibration factors and assigned relative uncertainties used in the later NP04 and NP02 purity-monitor processing.

Configuration	Gain correction	Gain unc.	Secondary-peak correction (unc.)
NP04 Bottom	0.9633	∼1.3%	0.9690 (0.79%)
NP04 Middle	0.9750	∼1.3%	0.9905 (0.96%)
NP04 Top	1.0018	∼1.3%	0.9783 (0.94%)
NP02	0.97984	∼1.3%	0.9863 (0.12%)

The NP02 calibration also includes a readout-box correction associated with a configuration change in December 2024. This correction is applied to the charge ratio for the corresponding data-taking period.

The grid correction addresses charge loss at the cathode grid. In ProtoDUNE-SP-I, the voltage settings used to extend the drift time did not provide full electron transmission through this grid. Calibration measurements at voltage settings with high grid transparency were used to determine the dependence of the charge ratio on drift time.

The charge ratios measured at different drift times were fitted with the expected exponential attenuation. The resulting relationship was used to estimate the charge ratio at the normal operating drift time. Comparing this estimate with the measured ratio at the normal operating voltages provided the grid-transparency correction [8]. Figure 10.4 shows the calibration measurements for the three monitors.

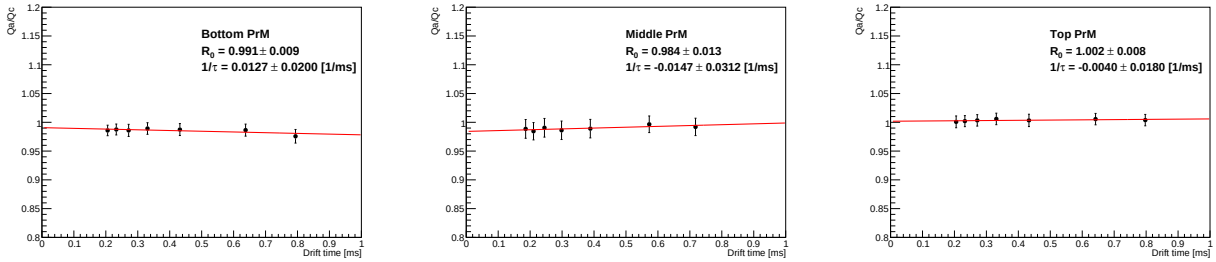

Figure 10.4: Charge-ratio calibration measurements for the bottom, middle, and top ProtoDUNE-SP-I UV purity monitors, respectively. The charge-ratio dependence on drift time is used to evaluate the grid-transparency correction [8].

After applying the relevant corrections, the electron lifetime is calculated from R_{corr} and the measured drift time using Eq. 10.2.

10.3 Electron-Lifetime Measurements

10.3.1 Measurement Uncertainties

The lifetime is sensitive to uncertainties in the corrected charge ratio, particularly when the liquid argon is very pure. For fixed drift time, propagation of the ratio uncertainty gives

$$\frac{\sigma_{\tau_e}}{\tau_e} \simeq \left|\frac{1}{\ln R}\right| \frac{\sigma_R}{R} = \frac{\tau_e}{t_{\mathrm{d}}} \frac{\sigma_R}{R}. \tag{10.6}$$

This approximation describes the charge-ratio contribution when the drift-time uncertainty is small. Timing uncertainty is also relevant to the complete measurement.

At high purity, $t_{\mathrm{d}}/\tau_e \ll 1$ and

$$R \simeq 1 - \frac{t_{\mathrm{d}}}{\tau_e}. \tag{10.7}$$

As R approaches unity, the same fractional uncertainty in charge ratio produces a much larger fractional uncertainty in lifetime. The ProtoDUNE-SP Phase-I analysis reported total statistical and systematic charge-ratio uncertainties of 1.9%, 2.2%, and 3.9% for the top, middle, and bottom monitors, respectively, at a drift time of approximately 2.3 ms [8].

10.3.2 ProtoDUNE-SP Performance

The ProtoDUNE-SP monitoring record follows the evolution of liquid-argon purity during purification and detector operation. Figure 10.5 shows the corrected charge-survival ratios from the three monitors. The response improved after circulation began in October 2018, with additional variations during changes in cryogenic operation [8].

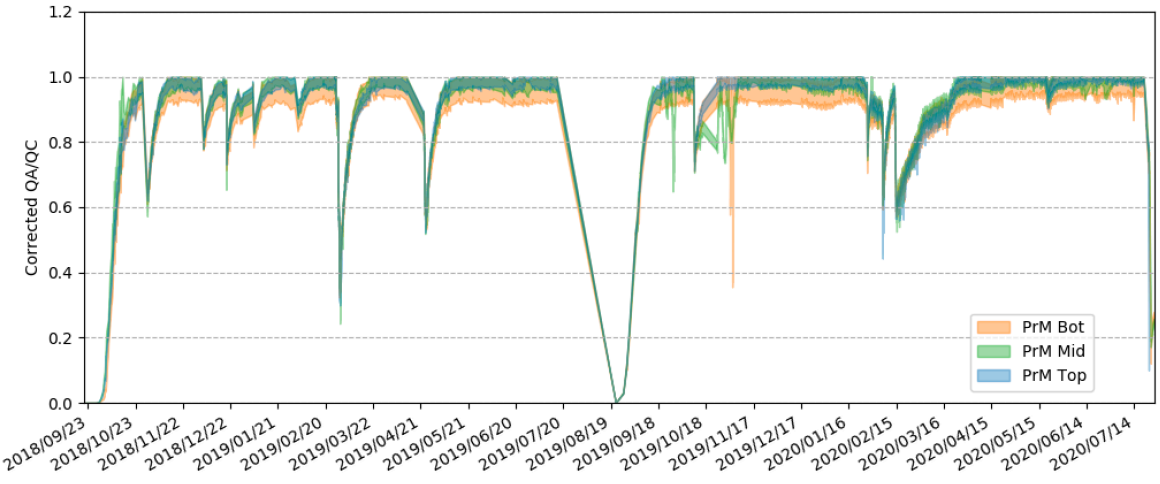

Figure 10.5: Corrected $Q_{\mathrm{a}}/Q_{\mathrm{c}}$ history from the three ProtoDUNE-SP UV purity monitors between September 2018 and July 2020. Circulation began in October 2018. Subsequent excursions coincide with documented changes to the recirculation configuration, including dedicated studies and pump interruptions. The bands indicate measurement uncertainties. Reproduced from the DUNE Collaboration ProtoDUNE-SP monitoring record in Ref. [8].

Figure 10.6 shows the corresponding electron-lifetime history. During sustained high-purity operation, the monitors reported lifetimes above 30 ms [8, 40]. In this regime the charge-survival ratio is close to unity, so the large inferred lifetime values carry substantial

uncertainty. The measurements nevertheless show that attenuation over the TPC drift interval was small.

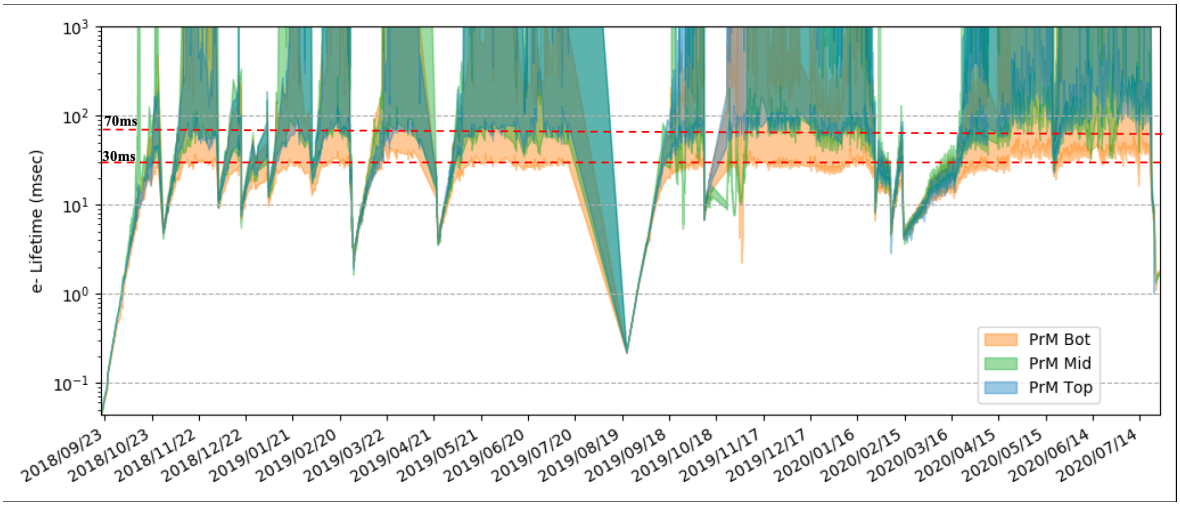

Figure 10.6: Electron-lifetime history from the three ProtoDUNE-SP UV purity monitors between September 2018 and February 2020, shown on a logarithmic vertical scale in milliseconds. The early low-lifetime period precedes the start of circulation in October 2018. Later downward excursions occur during documented recirculation tests or pump interruptions. The bands denote the measurement uncertainties. Reproduced from Ref. [8].

10.3.3 Comparison with TPC Measurements

An independent electron-lifetime estimate can be obtained from cosmic-ray tracks crossing the TPC cathode. The TPC method depends on track reconstruction and charge calibration, while the UV monitors depend on their electrode response, grid transmission, and waveform calibration. Comparing the measurements tests their consistency under different instrumental conditions.

In the published ProtoDUNE comparison, TPC and UV-monitor measurements were matched within six hours. The collaboration found consistency in the central region of the second anode plane assembly, reporting $\chi^2/\mathrm{ndf} \simeq 1$ [40]. This comparison provides an independent detector-level check of the UV-monitor measurements.

10.4 Long-Drift UV Purity-Monitor Development

The short ProtoDUNE-SP-I purity monitors become less sensitive to electron lifetime at high liquid-argon purity because the charge attenuation over their drift distance is small. Increasing the drift time by reducing the operating electric field also affects grid transparency, introducing additional uncertainty in the charge-ratio measurement. A longer drift region was developed to increase the measurable attenuation while maintaining high electron transmission through the grids.

The long-monitor design studied for ProtoDUNE-SP-II had an overall length of approximately 75 cm and a nominal electron drift distance of approximately 64 cm. The design included 60 field rings with a spacing of 0.31 in. and an aperture of 2.35 in. The later NP04 long middle monitor has an anode-disk–to–cathode-grid distance of 62.774 cm, compared with 17.054 cm for the short top and bottom monitors.

Figure 10.7 shows an engineering drawing of the extended purity-monitor design. The mechanical geometry defines the positions of the electrodes, grids, and field rings, which determine the electric-field distribution along the drift region.

The electric-field distribution was evaluated using three-dimensional electrostatic simulations in ANSYS Maxwell3D. The monitor geometry was constructed from the mechanical design, with voltages assigned to the electrodes and field rings. The electric field was calculated using a finite-element model.

For the simulated configuration, the cathode and anode voltages were -50 and 4000 V, respectively. The resulting electric field in the central drift region was approximately 62 V/cm.

Figure 10.8 compares the electric-field magnitude along the drift direction at the center of the monitor and at transverse positions displaced by ± 20 mm. The field remains approximately uniform through most of the drift region, with localized variations near the electrode grids. The similar field profiles at different transverse positions indicate that the field-ring geometry provides a nearly uniform drift field across the studied region.

The grid geometry was also studied to evaluate electron transmission and electrostatic

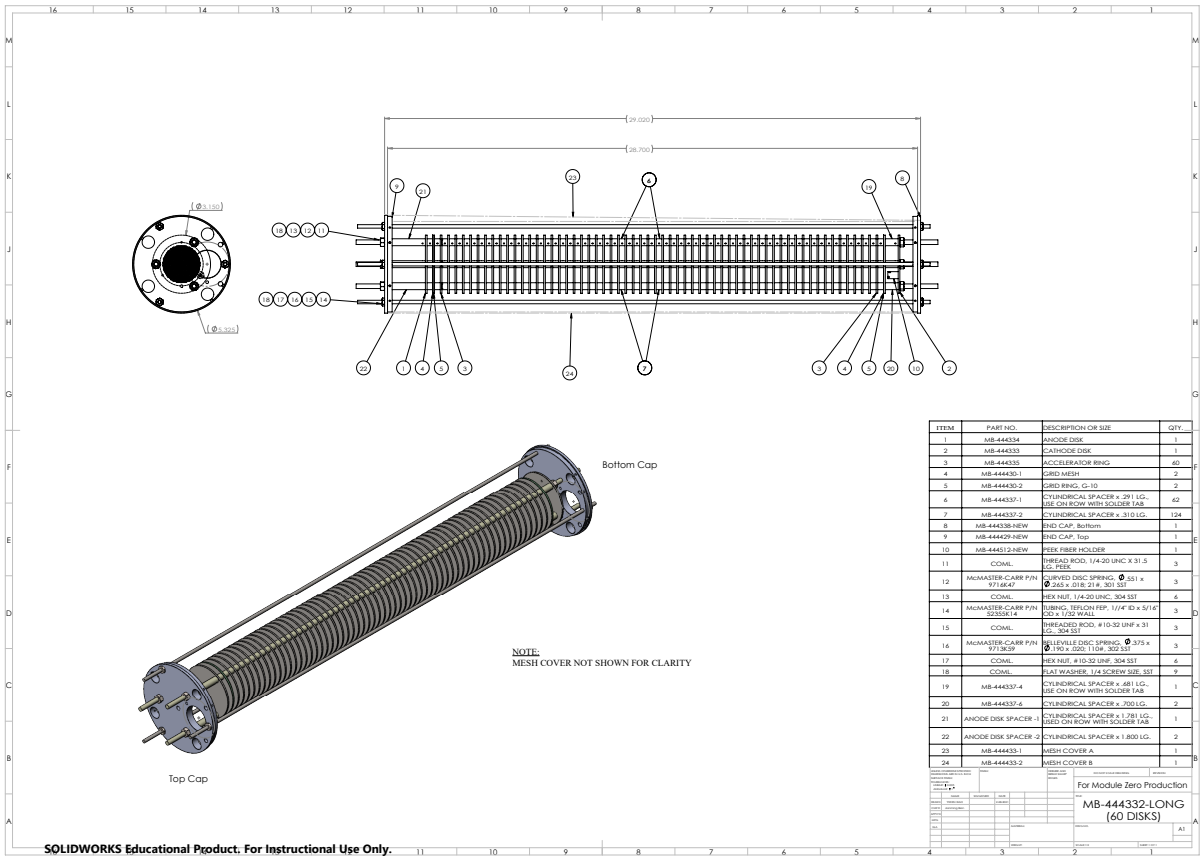

Figure 10.7: Engineering assembly drawing of a long-drift UV purity-monitor design, showing the extended drift region and field rings.

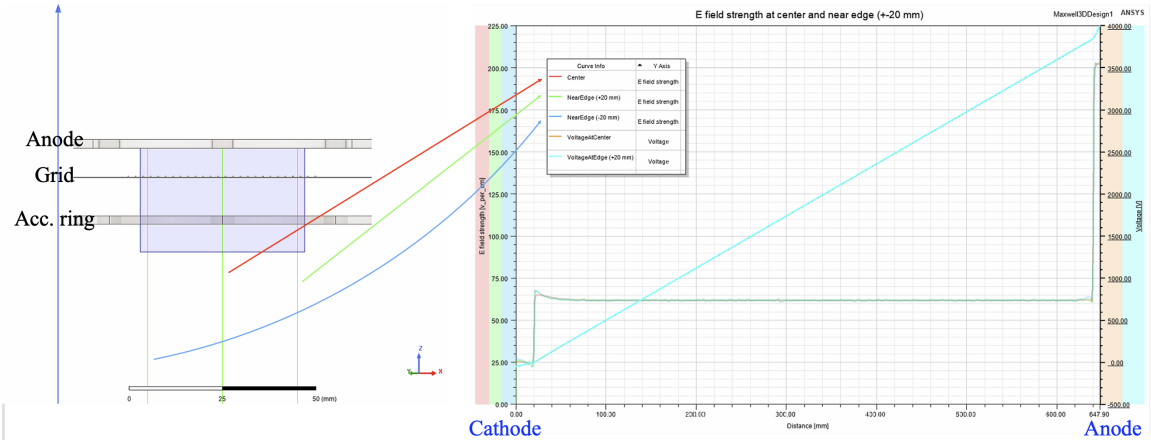

Figure 10.8: Electric-field distribution obtained from the three-dimensional electrostatic simulation of the long UV purity monitor. The field magnitude is shown along the drift direction at the center and at transverse offsets of ±20 mm. The field is approximately uniform in the central drift region, with variations near the electrode grids.

shielding. The grids shield the electrode signals from electrons moving through the drift region while allowing those electrons to pass toward the collecting electrodes. Their performance depends on the spacing and dimensions of the grid elements and on the electric fields on either side of the grid.

A stainless-steel hexagonal mesh was considered as an alternative to the wire-grid design used in the short monitors. Different mesh pitches and strip widths were evaluated using the long-monitor voltage configuration. The calculations included the electron transparency of the cathode and anode grids and their electrostatic inefficiencies.

Figure 10.9 summarizes the calculated grid performance. For a fixed strip width, increasing the mesh pitch generally improves electron transmission but reduces electrostatic shielding. Increasing the strip width can improve shielding while reducing electron transparency, particularly at smaller pitches.

The calculations show that high electron transparency alone is insufficient to characterize grid performance. Both electron transmission and electrostatic shielding must be considered when selecting the grid geometry.

In the later NP04 measurements, the short purity monitors had drift times of approximately 0.23–0.29 ms, while the long monitor had drift times of approximately 4.4–5.2 ms. The longer drift interval increased the measurable charge attenuation at high purity.

As shown in Eq. 10.6, a longer drift time reduces the relative lifetime uncertainty associated with a given charge-ratio uncertainty. The electrostatic simulations and grid studies provided a basis for evaluating the extended drift geometry and its measurement performance.

10.5 Electron-Lifetime Calibration

The electron lifetime measured by a UV purity monitor depends on the electric field in its drift region. Since the electron-attachment rate varies with electric field, a lifetime measured at the purity-monitor field cannot be directly applied to the TPC operating field.

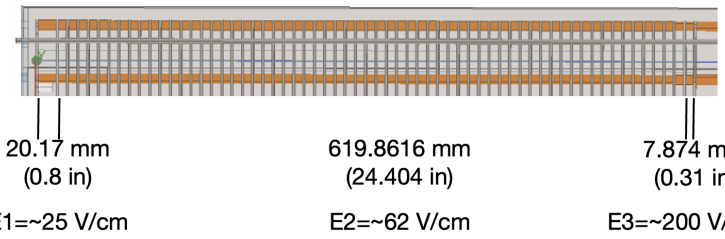

Tc/Ta/Ec/Ea		Pitch						
Strip width (um)	Strip thickness (inch)	1 mm	1.5 mm	1.8 mm	2.0 mm	4.0 mm	5.0 mm	8.0 mm
32	0.001	100%/100%/0.00796/0.0204	100%/100%/0.0147/0.0377	100%/100%/0.0192/0.0491	100%/100%/0.0222/0.057	100%/100%/0.0571/0.146	100%/100%/0.0765/0.196	100%/100%/0.139/0.357
50	0.0016	100%/100%/0.00593/0.0152	100%/100%/0.0117/0.0299	100%/100%/0.0155/0.0397	100%/100%/0.0182/0.0466	100%/100%/0.049/0.125	100%/100%/0.0663/0.17	100%/100%/0.123/0.316
63.5	0.002	100%/100%/0.00484/0.0124	100%/100%/0.01/0.0257	100%/100%/0.0135/0.0347	100%/100%/0.016/0.041	100%/100%/0.0446/0.114	100%/100%/0.0609/0.156	100%/100%/0.115/0.293
100	0.003	94.1%/99.6%/0.00277/0.0071	100%/100%/0.00693/0.0178	100%/100%/0.00981/0.0251	100%/100%/0.0119/0.0304	100%/100%/0.0364/0.0931	100%/100%/0.0505/0.129	100%/100%/0.098/0.251
200	0.0063		79.4%/86.6%/0.00219/0.00562	89.6%/96.4%/0.00413/0.0106	94.1%/99.6%/0.00554/0.0142	100%/100%/0.0237/0.0608	100%/100%/0.0347/0.089	100%/100%/0.0727/0.186
500	0.016					83.4%/90.6%/0.00702/0.018	94.1%/99.6%/0.00139/0.0355	100%/100%/0.0393/0.101

High inefficiency

Low transparency

Figure 10.9: Calculated grid transparency and electrostatic inefficiency for different mesh pitches and strip widths in the long-monitor design study. The results compare cathode and anode grid performance under the simulated voltage configuration.

For an impurity concentration n_s, the electron lifetime can be expressed as

$$\tau_e = \frac{1}{k_{\mathrm{a}}(E) n_s}, \tag{10.8}$$

where $k_{\mathrm{a}}(E)$ is the field-dependent electron-attachment coefficient.

For the same impurity concentration, the electron lifetime at the TPC field is related to the purity-monitor measurement by

$$\tau_e^{\mathrm{TPC}} = \tau_e^{\mathrm{PrM}} \frac{k_{\mathrm{a}}^{\mathrm{PrM}}}{k_{\mathrm{a}}^{\mathrm{TPC}}}, \tag{10.9}$$

where the attachment coefficients are evaluated at the corresponding electric fields.

The purity-monitor electric field is determined from the operating voltages and drift geometry. The attachment coefficients are then evaluated at the purity-monitor field and the nominal TPC field using the field-dependent attachment-rate parameterization of Ref. [42].

For the NP04 calibration, the TPC field is 500 V/cm, while the NP02 calibration uses

460 V/cm. The ratio of the corresponding attachment coefficients provides the conversion from the measured purity-monitor lifetime to the TPC-field-equivalent lifetime.

The conversion is applied after the instrumental corrections to the charge ratio. Each measurement is associated with its operating voltages, waveform parameters, charge-ratio corrections, and electron-lifetime result. These quantities provide the information needed to apply the appropriate calibration for different monitor configurations and data-taking periods.

The resulting TPC-field-equivalent electron lifetimes provide a measure of charge attenuation at the detector operating field for use in TPC calibration.

10.6 ProtoDUNE-HD Automated Monitoring

Continuous monitoring of liquid-argon purity is important during ProtoDUNE-HD commissioning and operation, as changes in the cryogenic conditions can affect the electron lifetime. An automated monitoring system was developed to process UV purity-monitor measurements and provide updated lifetime results throughout detector operation.

The system combines data acquisition, waveform analysis, charge calibration, and electron-lifetime calculation. UV flashes are generated at scheduled intervals, and the corresponding cathode and anode waveforms are recorded. The acquired data are transferred to an analysis computer, where the waveforms are processed to extract the charge amplitudes and drift times. The calibration procedures described in Sec. 10.2.3 are then applied to determine the corrected charge ratios and electron lifetimes. The resulting measurements are organized chronologically to track the evolution of liquid-argon purity.

The monitoring system also checks the availability of new measurements. If no data are received within the expected interval, an email notification is generated to indicate a possible interruption in acquisition or data transfer. Updated monitoring plots are distributed through the purity-monitor operations channel, allowing changes in the measurements to be examined

together with the detector operating conditions.

Figure 10.10 shows the electron-lifetime histories obtained from the automated monitoring system. The measurements extend through September 2026 and provide a record of liquid-argon purity during ProtoDUNE-HD operation. The time series allows changes in electron lifetime to be examined over both extended periods and shorter operating intervals.

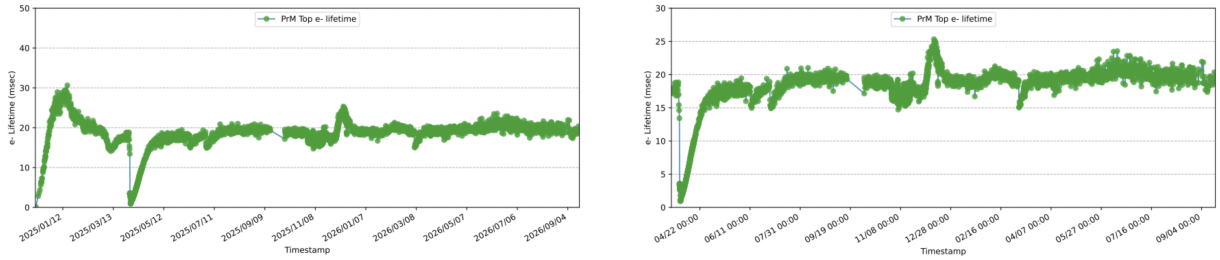

Figure 10.10: Electron-lifetime histories from the ProtoDUNE-HD automated monitoring system, showing long-term and recent measurements through September 2026.

The cathode response is also monitored to evaluate the stability of the instrument. Figure 10.11 shows the corrected cathode signal over the same operating period. Changes in the cathode amplitude can indicate variations in the light intensity, photocathode response, or readout electronics. Examining the cathode response together with the electron lifetime helps identify instrumental variations that could otherwise affect the interpretation of the purity measurements.

The automated system provides regular electron-lifetime measurements while reducing the need for manual data processing. The combination of lifetime measurements, instrument-response monitoring, and data-availability checks supports the identification of changes in liquid-argon purity and the continued operation of the purity-monitor system.

10.7 Summary

The ProtoDUNE UV purity monitors provide electron-lifetime measurements through the attenuation of ionization charge over a controlled drift distance. Waveform fitting and instrument-specific calibrations are used to determine the corrected charge ratios and electron

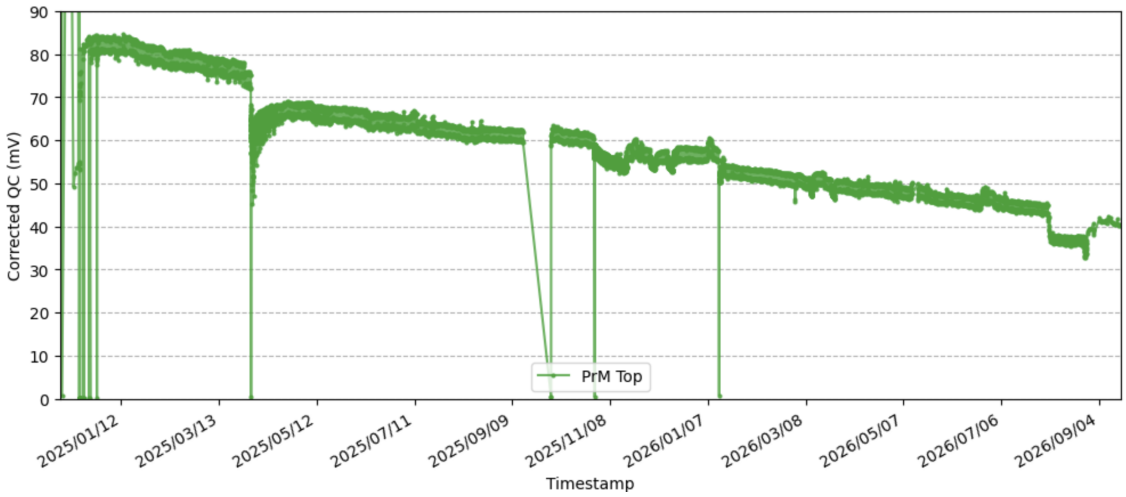

Figure 10.11: Corrected cathode-response history from ProtoDUNE-HD automated monitoring, extending through September 2026.

lifetimes. Comparisons with TPC-based measurements provide an independent cross-check of the results.

The ProtoDUNE-SP measurements illustrate the increasing lifetime uncertainty at high liquid-argon purity, where charge attenuation over the short drift distance becomes small. The development of longer-drift monitors and improved grid geometries addressed this limitation by increasing the measurable attenuation while maintaining a controlled electric field. Field-dependent lifetime corrections support the application of purity-monitor measurements to TPC calibration, while automated monitoring in ProtoDUNE-HD provides regular measurements of liquid-argon purity and instrument response.

Chapter 11

Conclusions and Outlook

This dissertation presents three studies in experimental neutrino physics: a neutrino-flux constraint using neutrino–electron elastic scattering in NOvA, a parallel Gaussian-process-accelerated Feldman–Cousins (GP-FC) method for oscillation-parameter inference, and UV purity-monitor development and electron-lifetime measurements for ProtoDUNE. Together, they address uncertainties in neutrino-flux predictions, the computational demands of statistical inference, and the measurement of ionization-charge attenuation in liquid argon.

11.1 NOvA Neutrino–Electron Scattering and Flux Constraint

Neutrino–electron elastic scattering provides a measurement of the neutrino flux through its precisely predicted Standard Model cross section. The NOvA Near Detector analysis uses 1.38×10^{21} POT of neutrino-mode data. Events are selected using a forward electromagnetic-shower topology, two CNN-based particle-identification variables, and the characteristic $E_e\theta_e^2$ distribution. The final signal region contains 827 data events. After applying the sideband-derived background normalization and subtracting the predicted background, the

measured signal yield is

$$740.1 \pm 30.2 \text{ (stat.) events.} \tag{11.1}$$

The background-subtracted reconstructed electron-energy spectrum constrains a 100-universe PPFX flux ensemble through Bayesian reweighting. The covariance matrix incorporates statistical and non-PPFX systematic uncertainties, while the PPFX universes describe correlated flux variations across neutrino flavors and energies. Closure tests, statistical-fluctuation studies, systematic variations, and finite-ensemble checks were used to validate the reweighting procedure.

The flux constraint reduces the fractional uncertainty on the predicted neutrino–electron event rate from 7.7% to 3.9%. The posterior mean of the integrated ν_μ flux is

$$126.4 \ \nu_\mu \, \mathrm{cm}^{-2}/10^{10} \, \mathrm{POT}, \tag{11.2}$$

with a fractional uncertainty of 4.0%, compared with 8.1% before the constraint. The posterior also changes the energy-dependent flux prediction, with the strongest constraints in regions connected to the selected electron-energy spectrum through the interaction kinematics and PPFX correlations.

The posterior weights can be applied to the same PPFX ensemble in other NOvA analyses using the corresponding beam simulation. This provides a data-driven constraint on the flux prediction while preserving correlations among neutrino flavors and energies.

11.2 Parallel Gaussian-Process-Accelerated Feldman–Cousins

The Feldman–Cousins construction provides confidence regions in oscillation analyses where asymptotic likelihood approximations are unreliable. Its computational cost arises from the large number of pseudo-experiments and likelihood fits needed throughout the parameter

space. The GP-FC approach uses a Gaussian Process to approximate the FC p-value surface and select additional evaluation points near confidence boundaries.

A parallel GP-FC implementation was developed and tested using realistic NOvA oscillation fits. The Gaussian-Process calculation is coordinated centrally, while independent pseudo-experiment fits are distributed among MPI workers. Dynamic scheduling and checkpoint/restart support calculations with widely varying fit runtimes. Timing studies involving 39,680 pseudo-experiments characterized the runtime distribution, and stored-pseudo-experiment tests informed the GP-bin and pseudo-experiment settings used in subsequent real-fit studies.

For the $(\sin^2\theta_{23}, \Delta m^2_{32})$ scan, the reconstructed GP-FC confidence regions achieved an intersection over union of 96.56%–98.14% with the reference FC regions. Overall bin overlap reached 100% after the second iteration.

These results demonstrate that the parallel GP-FC method can reproduce the reference FC confidence regions for the tested NOvA configurations while evaluating a subset of the available pseudo-experiment sample. The contour comparisons establish agreement for the studied scans. A direct assessment of frequentist coverage and a complete accounting of computational savings remain separate validation steps.

11.3 ProtoDUNE Purity-Monitor Instrumentation and Operation

UV purity monitors measure the survival of photoelectrons drifting through liquid argon and provide electron-lifetime information independently of TPC track reconstruction. The work presented here spans monitor design and testing, waveform analysis, charge calibration, long-drift development, and automated monitoring during ProtoDUNE operation.

Cathode and anode waveforms are fitted to extract the charge amplitudes and drift time. Corrections for channel response, waveform effects, and grid transmission are applied before

the charge-survival ratio is converted to electron lifetime. In ProtoDUNE-SP Phase I, the total fractional uncertainties on $Q_{\mathrm{a}}/Q_{\mathrm{c}}$ at a drift time of approximately 2.3 ms were 1.9%, 2.2%, and 3.9% for the top, middle, and bottom monitors, respectively.

The ProtoDUNE-SP measurements recorded electron lifetimes above 30 ms during extended periods of high-purity operation. As the charge-survival ratio approaches unity, the inferred lifetime becomes increasingly sensitive to small uncertainties in the measured charge ratio. Comparisons with TPC-based measurements provide an independent cross-check of the purity-monitor results.

Longer-drift UV monitors were developed to increase measurable charge attenuation at high purity. Electrostatic simulations evaluated the field uniformity of the extended drift region, while grid studies examined the balance between electron transmission and electrostatic shielding. Field-dependent attachment-rate corrections relate purity-monitor lifetimes to the TPC operating field. For ProtoDUNE-HD, automated acquisition and analysis provide regular lifetime and cathode-response histories, together with checks for interruptions in data availability.

11.4 Outlook

The NOvA flux constraint can be propagated to other analyses through the posterior PPFX-universe weights. Applying the approach to additional exposures or beam configurations will require corresponding event selections, background estimates, systematic uncertainties, and validation studies.

The parallel GP-FC framework can be extended to additional oscillation-parameter scans. Comparisons with reference FC constructions will remain important for assessing reconstructed confidence boundaries. Repeated-ensemble studies can establish the coverage properties of the accelerated construction, while wall-clock and node-hour measurements can quantify its computational performance.

The ProtoDUNE purity-monitor studies provide a basis for continued electron-lifetime calibration and long-term monitoring. The extended drift geometry, grid-performance studies, and field-dependent calibration methods are relevant to future purity-monitor designs, while the automated ProtoDUNE-HD measurements support the assessment of liquid-argon purity during detector operation.

The three studies contribute complementary methods to experimental neutrino physics: a data-driven neutrino-flux constraint, an adaptive approach to computationally intensive confidence-region construction, and instrumentation for monitoring charge attenuation in liquid argon.

Bibliography

[1] S. Navas et al. "Review of Particle Physics". In: *Phys. Rev. D* 110 (2024), p. 030001. DOI: 10.1103/PhysRevD.110.030001.

[2] K. Eguchi et al. "First Results from KamLAND: Evidence for Reactor Antineutrino Disappearance". In: *Phys. Rev. Lett.* 90 (2003), p. 021802. DOI: 10.1103/PhysRevLett.90.021802.

[3] F. P. An et al. "Observation of Electron-Antineutrino Disappearance at Daya Bay". In: *Phys. Rev. Lett.* 108 (2012), p. 171803. DOI: 10.1103/PhysRevLett.108.171803.

[4] D. S. Ayres et al. *The NOvA Technical Design Report*. Tech. rep. FERMILAB-DESIGN-2007-01. Fermilab, 2007. DOI: 10.2172/935497.

[5] M. A. Acero et al. "New constraints on oscillation parameters from ν_e appearance and ν_μ disappearance in the NOvA experiment". In: *Phys. Rev. D* 98 (2018), p. 032012. DOI: 10.1103/PhysRevD.98.032012. arXiv: 1806.00096 [hep-ex].

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

[7] B. Abi et al. "First Results on ProtoDUNE-SP Liquid Argon Time Projection Chamber Performance from a Beam Test at the CERN Neutrino Platform". In: *JINST* 15 (2020), P12004. DOI: 10.1088/1748-0221/15/12/P12004.

[8] A. Abed Abud et al. “Design, construction and operation of the ProtoDUNE-SP Liquid Argon TPC”. In: *JINST* 17.01 (2022), P01005. DOI: 10.1088/1748-0221/17/01/P01005. arXiv: 2108.01902 [physics.ins-det].

[9] B. Pontecorvo. “Mesonium and Antimesonium”. In: *Sov. Phys. JETP* 6 (1957), p. 429.

[10] Z. Maki, M. Nakagawa, and S. Sakata. “Remarks on the Unified Model of Elementary Particles”. In: *Prog. Theor. Phys.* 28 (1962), pp. 870–880.

[11] L. Wolfenstein. “Neutrino Oscillations in Matter”. In: *Phys. Rev. D* 17 (1978), pp. 2369–2374. DOI: 10.1103/PhysRevD.17.2369.

[12] S. P. Mikheyev and A. Yu. Smirnov. “Resonance Enhancement of Oscillations in Matter and Solar Neutrino Spectroscopy”. In: *Soviet Journal of Nuclear Physics* 42 (1985). English translation of Yad. Fiz. 42, 1441 (1985), pp. 913–917.

[13] S. P. Mikheyev and A. Yu. Smirnov. “Resonant Amplification of Neutrino Oscillations in Matter and Solar-Neutrino Spectroscopy”. In: *Il Nuovo Cimento C* 9 (1986), pp. 17–26. DOI: 10.1007/BF02508049.

[14] P. Adamson et al. “The NuMI Neutrino Beam”. In: *Nucl. Instrum. Meth. A* 806 (2016), pp. 279–306. DOI: 10.1016/j.nima.2015.08.063.

[15] Alejandro Jaime Yankelevich. “Optimizing Neutrino Flavor Conversion Measurements through Machine Learning”. Dissertation dated 9 March 2026. PhD thesis. University of California, Irvine, 2026.

[16] S. Agostinelli et al. “GEANT4—a simulation toolkit”. In: *Nucl. Instrum. Meth. A* 506 (2003), pp. 250–303. DOI: 10.1016/S0168-9002(03)01368-8.

[17] L. Aliaga et al. “Neutrino flux predictions for the NuMI beam”. In: *Phys. Rev. D* 94 (2016), p. 092005.

[18] C. Andreopoulos et al. “The GENIE Neutrino Monte Carlo Generator”. In: *Nucl. Instrum. Meth. A* 614 (2010), pp. 87–104.

[19] J. Tena-Vidal et al. “Neutrino-nucleon cross-section model tuning in GENIE v3”. In: *Phys. Rev. D* 104 (2021), p. 072009. DOI: 10.1103/PhysRevD.104.072009.

[20] R. Gran et al. “Neutrino-nucleus quasielastic and 2p2h interactions up to 10 GeV”. In: *Phys. Rev. D* 88 (2013), p. 113007. DOI: 10.1103/PhysRevD.88.113007.

[21] A. S. Meyer et al. “Deuterium target data for precision neutrino-nucleus cross sections”. In: *Phys. Rev. D* 93 (2016), p. 113015. DOI: 10.1103/PhysRevD.93.113015.

[22] J. Nieves, J. E. Amaro, and M. Valverde. “Inclusive quasielastic charged-current neutrino-nucleus reactions”. In: *Phys. Rev. C* 70 (2004), p. 055503. DOI: 10.1103/PhysRevC.70.055503.

[23] J. Nieves, I. Ruiz Simo, and M. J. Vicente Vacas. “Inclusive charged-current neutrino-nucleus reactions”. In: *Phys. Rev. C* 83 (2011), p. 045501. DOI: 10.1103/PhysRevC.83.045501.

[24] Ch. Berger and L. M. Sehgal. “Lepton mass effects in single pion production by neutrinos”. In: *Phys. Rev. D* 76 (2007), p. 113004. DOI: 10.1103/PhysRevD.76.113004.

[25] A. Bodek and U. K. Yang. “Higher twist, ξ_w scaling, and effective LO PDFs for lepton scattering in the few GeV region”. In: *J. Phys. G* 29 (2003), pp. 1899–1906. DOI: 10.1088/0954-3899/29/8/369.

[26] S. Dytman. “Neutrino Event Generators”. In: *AIP Conf. Proc.* Vol. 896. 2007, pp. 178–184. DOI: 10.1063/1.2720454.

[27] S. Dytman. “Final State Interactions in Neutrino–Nucleus Experiments”. In: *Acta Phys. Polon. B* 40 (2009), pp. 2445–2460.

[28] M. A. Acero et al. “Improved measurement of neutrino oscillation parameters by the NOvA experiment”. In: *Phys. Rev. D* 106 (2022), p. 032004. DOI: 10.1103/PhysRevD.106.032004. arXiv: 2108.08219 [hep-ex].

[29] M. Baird et al. “Event reconstruction techniques in NOvA”. In: *J. Phys. Conf. Ser.* 664 (2015), p. 072035. DOI: 10.1088/1742-6596/664/7/072035.

[30] A. Aurisano et al. “A Convolutional Neural Network Neutrino Event Classifier”. In: *JINST* 11 (2016), P09001. DOI: 10.1088/1748-0221/11/09/P09001.

[31] MINERvA Collaboration. “Measurement of Neutrino Flux from Neutrino-Electron Elastic Scattering”. In: *Phys. Rev. D* 93 (2016), p. 112007.

[32] MINERvA Collaboration. “Constraint of the MINERvA Medium Energy Neutrino Flux Using Neutrino-Electron Elastic Scattering”. In: *Phys. Rev. D* 100 (2019), p. 092001.

[33] J. Erler and S. Su. “The Weak Neutral Current”. In: *Prog. Part. Nucl. Phys.* 71 (2013), pp. 119–149. DOI: 10.1016/j.ppnp.2013.03.004.

[34] D. Ruterbories et al. “Improved Constraint on the MINERvA Medium Energy Neutrino Flux Using Antineutrino–Electron Scattering Data”. In: *Phys. Rev. D* 107 (2023), p. 012001. DOI: 10.1103/PhysRevD.107.012001.

[35] Oleksandr Tomalak and Richard J. Hill. “Theory of elastic neutrino-electron scattering”. In: *Physical Review D* 101 (2020), p. 033006. DOI: 10.1103/PhysRevD.101.033006. arXiv: 1907.03379 [hep-ph].

[36] S. S. Wilks. “The Large-Sample Distribution of the Likelihood Ratio for Testing Composite Hypotheses”. In: *Ann. Math. Stat.* 9 (1938), pp. 60–62. DOI: 10.1214/aoms/1177732360.

[37] M. A. Acero et al. “The Profiled Feldman–Cousins Technique for Confidence Interval Construction in the Presence of Nuisance Parameters”. In: *JINST* 20 (2025), T02001. DOI: 10.1088/1748-0221/20/02/T02001. arXiv: 2207.14353 [hep-ex].

[38] G. J. Feldman and R. D. Cousins. “Unified Approach to the Classical Statistical Analysis of Small Signals”. In: *Phys. Rev. D* 57 (1998), pp. 3873–3889. DOI: 10.1103/PhysRevD.57.3873.

[39] L. Li et al. “Efficient Neutrino Oscillation Parameter Inference Using Gaussian Processes”. In: *Phys. Rev. D* 101 (2020), p. 012001. DOI: 10.1103/PhysRevD.101.012001.

[40] S. Abbaslu et al. “Spatial and temporal evaluations of the liquid argon purity in ProtoDUNE-SP”. In: *JINST* 20.09 (2025), P09008. DOI: 10.1088/1748-0221/20/09/P09008. arXiv: 2507.08586 [physics.ins-det].

[41] Y. Xiao. “Purity Monitoring for ProtoDUNE-SP”. In: *J. Phys. Conf. Ser.* 2156 (2022), p. 012211. DOI: 10.1088/1742-6596/2156/1/012211.

[42] Y. Li et al. “Parameterization of electron attachment rate constants for impurities in LArTPC detectors”. In: *Journal of Instrumentation* 17.11 (2022), T11007. DOI: 10.1088/1748-0221/17/11/T11007. arXiv: 2205.06888 [physics.ins-det].