

A Measurement of ν_e Appearance and ν_μ Disappearance Neutrino Oscillations with the NOvA Experiment

Thesis by
Daniel Pershey

In Partial Fulfillment of the Requirements for the
Degree of
Doctor of Philosophy

The logo for the California Institute of Technology (Caltech), featuring the word "Caltech" in a bold, orange, sans-serif font.

CALIFORNIA INSTITUTE OF TECHNOLOGY
Pasadena, California

2018
Defended May 25, 2018

© 2018

Daniel Pershey

ORCID: 0000-0003-4298-0337

All rights reserved

ACKNOWLEDGEMENTS

Before starting off, I'd like to thank all those who have helped me reach this finish line of my graduate career. Looking back, my days in California have been great, and I've learned more than I expected at Caltech. First, my advisor, Ryan Patterson, was instrumental in mentoring me through this thesis and was always patient while showing me how to attack problems experimental physicists deal with everyday. It was also a great help to have Chris Backhouse, a convener of the appearance working group, working down the hall from me, and he could field my questions whenever Ryan was busy. I'm happy to have had the chance to work alongside other great scientists in the Caltech group; here's to Kirk Bays, Leon Mualem, and Varun Raj. Special thanks go to my friend and office mate, Joseph Lozier, for the camaraderie earned through several years of grad school.

There were several people from my college days who provided a bedrock of excitement for elementary particle physics and helped steer me toward studying neutrino physics at Caltech. I have to call out Gary Feldman in particular, whose subdued optimism for NOvA's potential over the relative "largeness" of θ_{13} convinced me to stay in the experiment for graduate school. I'd also like to thank Alex Sousa for getting my feet wet with research.

I would like to thank the NOvA collaboration, in general. Everyone has played their part in making NOvA, and this thesis, a success. And, through omnipresent meetings, they've taught me defend my ideas in front of an audience of scientists. At some time during my time at Caltech, Tricia Vahle, Mayly Sanchez, and Alex Himmel have all held the reins of the nue appearance group; thanks for all the effort coordinating a string of successful results and the personal guidance. Jon Paley and Mat Muether as conveners for the near detector cross section group have been vital for both pushing my $CC\pi^0$ cross section through to results and teaching me how to make a cross section measurement. A special thanks also goes out to Matt Judah and Junting Huang on their great job tightening up TDSlicer. Finally, I'd like to name a few of the scientists I've had the pleasure of working with on NOvA: Evan Niner, Michael Baird, Gavin Davies, Jeremy Wolcott, Erika Cataño Mur, Jeff Hartnell, Peter Shanahan, Alex Radovic, Mark Messier, Kanika Sachdev, and Louise Suter.

ABSTRACT

NOvA is a two-detector, neutrino oscillation experiment on an 810 km baseline in the NuMI beam at Fermilab. NuMI is an off-axis, narrow-band beam centered near 1.8 GeV and is configurable to send a high-purity ν_μ or $\bar{\nu}_\mu$ neutrino flux which passes through the NOvA far detector, sensitive to ν_μ and $\bar{\nu}_\mu$ disappearance and ν_e and $\bar{\nu}_e$ appearance. The NOvA near and far detectors are functionally identical, liquid scintillator trackers. They are finely segmented relative to the radiation length in the detectors allowing for the detailed reconstruction of electromagnetic shower necessary to identify ν_e CC events.

Through a joint analysis of the appearance and disappearance oscillation channels using neutrino and anti-neutrino data, NOvA constrains the allowed parameter space in δ_{CP} , the neutrino mass hierarchy, $\sin^2 \theta_{23}$, and Δm_{32}^2 . We observe 58 and 18 appearance candidates on an expected background of 15.06 and 5.32 events in neutrino and anti-neutrino mode, respectively. This observation gives a 1.77σ preference for the normal mass hierarchy and 1.7σ preference for $\theta_{23} > \pi/4$. NOvA will maintain its position as a long-baseline oscillation experiment with leading sensitivity for several more years in which time analysis will further disambiguate the fundamental neutrino parameters that determine oscillation behavior.

TABLE OF CONTENTS

| | |
|--|-----|
| Acknowledgements | iii |
| Abstract | iv |
| Table of Contents | v |
| List of Illustrations | ix |
| List of Tables | xv |
| Chapter I: Neutrino Physics | 1 |
| 1.1 The Free Neutrino and Neutrino Mass | 1 |
| 1.2 Neutrino Oscillations | 3 |
| 1.3 Neutrino Oscillations in Matter | 6 |
| 1.4 Neutrino Interactions | 8 |
| The Weak Force | 9 |
| Chapter II: Experimental Measurement of Neutrino Mixing Parameters | 12 |
| 2.1 Solar Mixing Parameters | 12 |
| 2.2 Atmospheric Mixing Parameters | 15 |
| 2.3 Measuring θ_{13} : The Final Mixing Angle | 17 |
| 2.4 Long-Baseline Accelerator Experiments | 19 |
| Chapter III: The NOvA Experiment | 22 |
| 3.1 The NOvA Detector Design | 22 |
| Liquid Scintillator | 22 |
| WLS Fiber | 24 |
| Avalanche PhotoDiode (APD) | 25 |
| The NOvA Cell | 26 |
| Detector Assembly | 26 |
| 3.2 The NOvA Physics Program | 30 |
| Chapter IV: The NuMI Beam | 33 |
| 4.1 Current Fermilab Accelerator Program | 33 |
| 4.2 The Neutrinos at the Main Injector (NuMI) Beam | 34 |
| Linac and Booster Accelerators | 34 |
| Proton Stacking in the Recycler Ring | 35 |
| The Main Injector, Acceleration and Injection into NuMI | 35 |
| The NuMI Beamline | 35 |
| Intrinsic Beam Backgrounds | 37 |
| Notable Improvements with the ANU | 38 |
| 4.3 Off-Axis Neutrino Beams | 38 |
| Chapter V: The NOvA Data AcQuisition (DAQ) System | 42 |
| 5.1 Channel Readout | 42 |
| Front End Electronics | 42 |
| Single-point and Multi-Point Readout | 43 |
| 5.2 Collecting Hits from All Cells and Calibrating Timing | 45 |

| | | |
|---|--|-----|
| 5.3 | Event Triggering | 47 |
| 5.4 | Measuring Channel Thresholds | 48 |
| 5.5 | APD Sag | 50 |
| Chapter VI: Simulation of the NOvA Experiment | | 52 |
| 6.1 | Detector Simulation | 52 |
| 6.2 | NuMI Beam Simulation | 53 |
| | Package to Predict the FluX (PPFX) | 53 |
| 6.3 | Event Generation and Particle Propagation | 55 |
| | GENIE | 55 |
| | CRY | 57 |
| | Geant4 | 58 |
| 6.4 | Simulation of Čerenkov Light | 58 |
| 6.5 | Scintillator Nonlinearity and a Birks Parameterization | 60 |
| Chapter VII: Calibration of the NOvA Detectors | | 62 |
| 7.1 | Attenuation Calibration | 62 |
| 7.2 | Absolute Calibration | 66 |
| Chapter VIII: Event Reconstruction | | 69 |
| 8.1 | Event Clustering | 71 |
| 8.2 | Vertex Seeding with a Hough Transform | 73 |
| 8.3 | Vertex Identification | 75 |
| 8.4 | Generating Prongs | 77 |
| 8.5 | Reconstructing Tracks with a Kalman Filter | 79 |
| 8.6 | Energy Reconstruction for ν_μ CC and ν_e CC events | 80 |
| 8.7 | Event Classification with a Convolution Neural Network | 81 |
| 8.8 | Reconstructing Michel Electrons | 84 |
| 8.9 | An Alternative Clustering Algorithm: TDSlicer | 87 |
| | TDSlicer Algorithm | 87 |
| | A Note on Properly Splitting Distinct Physics Slices | 90 |
| | TDSlicer Performance | 92 |
| Chapter IX: $\text{CC}\pi^0$ Cross Section Analysis | | 94 |
| 9.1 | Definition of Signal | 95 |
| 9.2 | Event Reconstruction | 96 |
| 9.3 | Kinematic Estimation | 98 |
| 9.4 | Event Selection | 100 |
| | Data Quality and Spill Cuts | 100 |
| | Fiducial Cuts | 101 |
| | Containment Cuts | 101 |
| | Background Rejection with CVN | 102 |
| | Kinematic Restrictions | 104 |
| 9.5 | Background Decomposition | 107 |
| 9.6 | Signal Estimation | 108 |
| 9.7 | Unfolding | 110 |
| 9.8 | Assembling the Cross Section | 116 |
| 9.9 | Systematic Uncertainties | 119 |
| | Neutrino Cross Section Systematics | 120 |

| | |
|---|-----|
| Particle Tracking Cross Section Uncertainties | 121 |
| Flux Modeling | 126 |
| Detector Response | 126 |
| Normalization Systematics | 130 |
| Total Error Budget | 131 |
| Validation of Computed Systematic Uncertainties | 131 |
| 9.10 Signal Estimation with Data | 135 |
| 9.11 Analysis Results | 139 |
| Muon Kinematics | 139 |
| π^0 Kinematics | 139 |
| Q^2 and W | 140 |
| Spread in $\langle\sigma_x\rangle_\Phi$ | 142 |
| 9.12 Summary | 143 |
| Chapter X: Overview of the Oscillation Analysis | 144 |
| 10.1 Contributing Datasets | 146 |
| 10.2 Blindness Policy for the Oscillation Analysis | 146 |
| Chapter XI: Far Detector Event Selection | 148 |
| 11.1 ν_μ Selection | 148 |
| Basic Quality Cuts | 148 |
| Preselection | 150 |
| PID | 150 |
| Cosmic Rejection | 151 |
| Hadronic Energy Fraction Bins | 151 |
| 11.2 ν_e Selection | 153 |
| Basic Quality Cuts | 153 |
| Core Preselection | 154 |
| Core Cosmic Rejection | 155 |
| Core PID Cut and Bins | 156 |
| Peripheral Selection | 157 |
| Chapter XII: Constraining the Far Detector Prediction with Near Detector Data | 160 |
| 12.1 Near Detector ν_μ CC Selection | 160 |
| 12.2 Near Detector ν_e CC Selection | 168 |
| 12.3 Near Detector ν_e Decomposition | 170 |
| 12.4 BEam Nue Decomposition (BENDecomp) | 171 |
| 12.5 The Michel Decomposition | 173 |
| Michel Reconstruction | 174 |
| MichelDecomp Algorithm | 175 |
| Restricting Decomposition to Bins with Robust Fit | 177 |
| Decomposition Result | 177 |
| 12.6 Data/MC Comparison for the Decomposed sample | 180 |
| 12.7 Near-to-Far Extrapolation | 181 |
| Strategies for Constraining Individual Flavor Transitions | 182 |
| 12.8 The Prediction for the Disappearance Spectrum | 184 |
| 12.9 The Prediction for the Appearance Spectrum | 187 |
| Chapter XIII: Systematic Uncertainties | 191 |

| | |
|---|-----|
| 13.1 Flux Systematic Uncertainty | 191 |
| 13.2 MC Tuning for the Oscillation analysis | 194 |
| 13.3 Neutrino Interaction Systematic Uncertainty | 200 |
| 13.4 Acceptance Systematic | 201 |
| 13.5 Detector Response Systematics | 203 |
| 13.6 Normalization Systematics | 204 |
| 13.7 Overview of Systematic Effects | 204 |
| Chapter XIV: Fitting, Sensitivity, and Confidence Interval Construction . . . | 206 |
| 14.1 Log-Likelihood Fit | 206 |
| 14.2 Sensitivities | 207 |
| 14.3 Feldman-Cousins Pseudo-experiments | 210 |
| Chapter XV: Oscillation Results | 212 |
| 15.1 Observed Far Detector Spectra | 212 |
| 15.2 Oscillation Fits and Constraints on Physics Parameters | 220 |
| Chapter XVI: Summary | 225 |
| Bibliography | 227 |
| Appendix A: Neutrino-Nucleus Scattering Phenomenology | 240 |
| Leptonic Cross Sections | 242 |
| Quasi-Elastic Scattering | 243 |
| Deep-Inelastic Scattering | 244 |
| Resonant Scattering | 244 |
| Meson Exchange Current (MEC) Scattering | 245 |
| $\nu - A$ Coherent π Production | 246 |
| Appendix B: $CC\pi^0$ ID Fits for Cross Section Analysis | 247 |
| B.1 Measurement in p_μ | 247 |
| B.2 Measurement in $\cos \theta_\mu$ | 250 |
| B.3 Measurement in p_π | 253 |
| B.4 Measurement in $\cos \theta_\pi$ | 257 |
| B.5 Measurement in Q^2 | 260 |
| B.6 Measurement in W | 264 |
| Appendix C: $CC\pi^0$ Cross Section and Covariance Tables | 268 |

LIST OF ILLUSTRATIONS

| <i>Number</i> | <i>Page</i> |
|---|-------------|
| 1.1 Schematic of the theory of neutrino oscillations | 5 |
| 1.2 Summary of ν_μ charged current cross section measurements. | 9 |
| 1.3 Three allowed weak interactions on fundamental fermions | 10 |
| 2.1 The solar neutrino flux, by fusion channel | 13 |
| 2.2 The L/E oscillation pattern observed by KamLAND | 14 |
| 2.3 Combined solar oscillation contours, with KamLAND and solar ex- periments. | 14 |
| 2.4 Prediction for matter effects on solar survival probability with exper- imental data | 15 |
| 2.5 The atmospheric neutrino flux at the surface | 16 |
| 2.6 Observation of atmospheric neutrino disappearance from Super-K | 17 |
| 2.7 $\bar{\nu}_e$ survival probability from Daya Bay | 18 |
| 2.8 Contour from reactor $\bar{\nu}_e$ disappearance from Daya Bay | 18 |
| 2.9 Comparison of several ν_μ disappearance measurements | 19 |
| 2.10 Allowed δ_{CP} ranges for each mass hierarchy from T2K | 21 |
| 3.1 Emission spectra for fluors in NOvA scintillator | 23 |
| 3.2 Absorption and emission spectra for NOvA WLS fiber | 24 |
| 3.3 Quantum efficiency of the NOvA APD vs wavelength | 25 |
| 3.4 Schematic view of the fundamental NOvA cell | 27 |
| 3.5 Cross section of the NOvA cell | 27 |
| 3.6 A diagram of the NOvA dual orientation used for 3D reconstruction | 28 |
| 3.7 Example event display of observed neutrino | 29 |
| 3.8 Closeup of example event display | 30 |
| 3.9 ν_e appearance on the NOvA baseline | 31 |
| 3.10 NOvA sensitivity to mass hierarchy and no CP violation rejection | 31 |
| 4.1 An overview of the Fermilab accelerator complex | 34 |
| 4.2 The process of building a NuMI spill with the Fermilab accelerator | 36 |
| 4.3 A schematic of the NuMI beamline | 37 |
| 4.4 The neutrino energy as a function of parent meson energy for various off-axis angles | 40 |
| 4.5 The neutrino flux for various off-axis detector locations | 41 |

| | | |
|------|--|----|
| 5.1 | Schematic of the NOvA front end electronics | 43 |
| 5.2 | Two ADC-vs-time traces for simulated hits | 44 |
| 5.3 | Timing resolution in near and far detectors | 45 |
| 5.4 | Structure of recorded NOvA hit. | 45 |
| 5.5 | Top-view of the far detector, showing layout of electronics | 46 |
| 5.6 | A schematic of the DCM and timing chain layout in the detector | 47 |
| 5.7 | Time traces for 64 channels used for evaluating thresholds | 48 |
| 5.8 | Distribution of changes in ADC value used to set thresholds | 49 |
| 5.9 | An illustration of channel cross-talk due to charge depletion | 50 |
| 5.10 | An event display of an FEB flasher event | 51 |
| 6.1 | Simulated NOvA flux | 54 |
| 6.2 | Data-driven corrections to π^+ production in $p + {}^{12}\text{C}$ collisions | 54 |
| 6.3 | Error bands on hadron production cross sections used for constraining the NuMI flux | 55 |
| 6.4 | Čerenkov light released as a function of wavelength | 59 |
| 6.5 | Photons released per cm vs particle speed for charged particles | 60 |
| 6.6 | Birks-Chou parameter tuning used in NOvA simulation | 61 |
| 7.1 | A schematic of the tricell topology used in the calibration procedure | 63 |
| 7.2 | The attenuation of light output over the length of a NOvA cell. | 64 |
| 7.3 | Simulated threshold biases for each view | 65 |
| 7.4 | LOWESS attenuation fit to discrepancies between modeled shape and observed data | 66 |
| 7.5 | The mean light output vs position in the cell, before and after atten- uation calibration | 66 |
| 7.6 | Muon dE/dx distribution vs the distance to track endpoint | 67 |
| 7.7 | Calibration scale calculation and application to muon dE/dx | 68 |
| 8.1 | Common NOvA event topologies | 70 |
| 8.2 | Diagram of the NOvA reconstruction chain | 71 |
| 8.3 | Nomenclature for the DBSCAN algorithm | 73 |
| 8.4 | Event clustering performance | 73 |
| 8.5 | Vertex seeding with a Hough transform | 75 |
| 8.6 | Performance of Hough vertex seed algorithm | 75 |
| 8.7 | Resolution on each coordinate of the reconstructed vertex | 77 |
| 8.8 | Primary lepton reconstruction efficiency with NOvA | 79 |
| 8.9 | ν_μ CC energy reconstruction | 81 |
| 8.10 | ν_e CC energy reconstruction and performance | 81 |

| | | |
|------|---|-----|
| 8.11 | CVN extracted feature maps for a ν_e CC event | 83 |
| 8.12 | CVN extracted feature maps for a ν_μ CC event | 83 |
| 8.13 | Event display of inelastic ν_μ CC event selected by ν_e PID | 85 |
| 8.14 | Simulated MID input distributions for true Michels and background | 86 |
| 8.15 | Data/MC comparison for MID, our Michel electron identifier | 87 |
| 8.16 | A sample near detector simulated spill clustered with nominal slicer | 88 |
| 8.17 | A sample far detector simulated spill clustered with nominal slicer | 88 |
| 8.18 | A centroid-clustering algorithm based on finding density peaks | 90 |
| 8.19 | A sample near detector simulated spill clustered with alternative slicer | 91 |
| 8.20 | A sample far detector simulated spill clustered with alternative slicer | 91 |
| 8.21 | Illustration of TDSlicer's power for separating distinct physics interactions | 92 |
| 8.22 | Effect of near detector pileup with Slicer4D and TDSlicer | 93 |
| 9.1 | Accessible $CC\pi^0$ phase space in p_π and p_μ | 96 |
| 9.2 | Prong inputs to photon $\Delta \log \mathcal{L}$ score | 98 |
| 9.3 | Resolution on muon and π^0 kinematics. | 99 |
| 9.4 | Data/MC comparison for NC rejection PID, CVNm | 102 |
| 9.5 | Simulated distributions for events estimated as QE and Coherent by CVN | 103 |
| 9.6 | Efficiency and purity as a function of neutrino energy | 105 |
| 9.7 | Efficiency and purity as a function of Q^2 | 105 |
| 9.8 | Efficiency and purity as a function of p_π | 106 |
| 9.9 | Efficiency and purity as a function of $\cos \theta_\pi$ | 106 |
| 9.10 | PID for the $CC\pi^0$ analysis, $CC\pi^0$ ID | 107 |
| 9.11 | Background $CC\pi^0$ ID spectrum for various selected particle types | 108 |
| 9.12 | Systematic fake-data test of signal estimation procedure | 110 |
| 9.18 | A closure test of the unfolding procedure | 112 |
| 9.13 | Kinematic migration matrices | 113 |
| 9.14 | Unfolding response to systematically shifted fake data | 114 |
| 9.15 | Unfolding response to systematically shifted fake data | 114 |
| 9.16 | Unfolding response to systematically shifted fake data | 115 |
| 9.17 | Unfolding response to systematically shifted fake data | 115 |
| 9.19 | Event Efficiency vs each kinematic variable | 117 |
| 9.20 | The NuMI flux through the near detector, with error band | 118 |
| 9.21 | GENIE predicted cross section for each kinematic variable | 119 |

| | | |
|------|--|-----|
| 9.22 | Systematic error calculation from varying the simulated DIS π^- production cross section | 122 |
| 9.23 | Fitted error band for the $\pi^\pm \rightarrow \pi^0$ cross section | 123 |
| 9.24 | Effect of modifying π^\pm charge exchange cross section on CC π^0 ID | 125 |
| 9.25 | dE/dx for protons used to tune Birks-Chou parameters | 127 |
| 9.26 | Effect of updated light model on CC π^0 ID | 128 |
| 9.27 | Calorimetric energy of hits from muon tracks | 129 |
| 9.28 | Calorimetric energy of reconstructed Michel electrons | 129 |
| 9.29 | Reconstructed π^0 mass in data and MC | 130 |
| 9.30 | Systematic error budget as a function of each kinematic variable | 132 |
| 9.31 | CC π^0 ID distribution in data and MC for control photon sample | 133 |
| 9.32 | Angular variables used for selecting high-purity proton sample | 134 |
| 9.33 | CC π^0 ID distribution in data and MC for control proton sample | 135 |
| 9.34 | Example fits used for data-driven signal estimate | 136 |
| 9.35 | Data/MC comparison of nominal CC π^0 ID and kinematic variables | 137 |
| 9.36 | Data/MC for each kinematic variable before and after signal constraint | 138 |
| 9.37 | Derived cross sections in muon kinematics | 140 |
| 9.38 | Derived cross sections in π^0 kinematics | 141 |
| 9.39 | π^0 kinematic cross sections broken into channel of intra-nuclear π^0 production | 141 |
| 9.40 | Derived cross sections in Q^2 and W | 142 |
| 10.1 | A schematic overview of the oscillation analysis | 145 |
| 11.1 | Schematic of ν_μ CC selection | 149 |
| 11.2 | Cosmic rejection BDT for ν_μ CC selection | 151 |
| 11.3 | Fully selected far detector ν_μ CC events | 152 |
| 11.4 | Far detector ν_μ selected events in quartile HadEFrac bins | 153 |
| 11.5 | Schematic of ν_e CC selection | 154 |
| 11.6 | The p_T/p vs distance to the top of the detector for cosmic and signal events | 156 |
| 11.7 | CVNe tuning distributions for ν_e CC selection | 157 |
| 11.8 | Cosmic rejection BDT for the ν_e peripheral sample | 158 |
| 11.9 | CVN-BDT cut tuning for the ν_e peripheral sample | 159 |
| 12.1 | Data/MC for near detector ν_μ CC neutrino energy | 162 |
| 12.2 | Area-normalized data/MC for near detector ν_μ CC neutrino energy | 163 |
| 12.3 | Data/MC for near detector ν_μ CC hadronic energy fraction | 164 |
| 12.4 | Data/MC for near detector ν_μ CC muon energy | 165 |

| | | |
|-------|--|-----|
| 12.5 | Data/MC for near detector $\bar{\nu}_\mu$ CC neutrino energy | 166 |
| 12.6 | Data/MC for near detector $\bar{\nu}_\mu$ CC hadronic energy fraction | 167 |
| 12.7 | Data/MC for near detector $\bar{\nu}_\mu$ CC muon energy | 168 |
| 12.8 | Data/MC for selected near detector ν_e CC events | 170 |
| 12.9 | Ancestor particle type for near detector selected ν_e CC events | 171 |
| 12.10 | Ancestor particle type of near detector contained and uncontained ν_μ CC events | 172 |
| 12.11 | Ancestral p_T/p_z distribution for neutrinos with pion ancestors | 173 |
| 12.12 | Event display of inelastic ν_μ CC event selected by ν_e PID with a Michel decay | 174 |
| 12.13 | Data/MC comparison for inputs to Michel ID | 175 |
| 12.14 | Data/MC comparison for the Michel ID | 176 |
| 12.15 | Far detector background uncertainty under different decomposition techniques | 178 |
| 12.16 | Data/MC distribution of ν_e selected events before decomposition | 179 |
| 12.17 | Data/MC distribution of ν_e selected events after decomposition | 179 |
| 12.18 | Near detector electromagnetic energy distribution in data and MC | 180 |
| 12.19 | Near detector hadronic energy distribution in data and MC | 181 |
| 12.20 | Schematic of data-driven $\nu_\mu \rightarrow \nu_\mu$ estimation procedure | 183 |
| 12.21 | Illustration of systematic error mitigation with HadEFrac extrapolation | 186 |
| 12.22 | Predicted ν_e and $\bar{\nu}_e$ appearance events | 189 |
| 12.23 | Schematic illustration of ν_e and $\bar{\nu}_e$ selected samples | 190 |
| 13.1 | Covariance matrix for flux uncertainties | 192 |
| 13.2 | The flux predictions in ten individual systematic universes | 193 |
| 13.3 | Ordered eigenvalue distribution | 194 |
| 13.4 | Flux uncertainty coverage by the first five PCA components | 194 |
| 13.5 | Effect of RPA suppression on $\cos \theta_\mu$ distribution | 196 |
| 13.6 | Fitted ν_μ CC MEC kinematics | 196 |
| 13.7 | Fitted $\bar{\nu}_\mu$ CC MEC kinematics | 197 |
| 13.8 | Fitted MEC E_{had} for determining MEC systematic uncertainty | 198 |
| 13.9 | Derived ν_μ MEC contribution with error band | 199 |
| 13.10 | Derived $\bar{\nu}_\mu$ MEC contribution with error band | 199 |
| 13.11 | Distortion in ν_e selected events from a systematic distortion | 201 |
| 13.12 | Differences in $\cos \theta_{\text{lep}}$ between near and far detectors | 202 |
| 13.13 | Differences in Q^2 between near and far detectors | 203 |
| 13.14 | Summary of errors on total predicted ν_e appearance events | 205 |

| | | |
|-------|---|-----|
| 14.1 | NOvA sensitivity to physics parameters without systematics | 208 |
| 14.2 | NOvA sensitivity to physics parameters with systematics | 209 |
| 14.3 | NOvA sensitivity to physics milestones without systematics | 209 |
| 14.4 | NOvA sensitivity to physics milestones with systematics | 210 |
| 15.1 | Observed ν_e Appearance Spectra | 213 |
| 15.2 | Observed $\bar{\nu}_e$ Appearance Spectra | 213 |
| 15.3 | Observed ν_e and $\bar{\nu}_e$ event counts vs expectations | 214 |
| 15.4 | Observed E_ν distribution for selected ν_e and $\bar{\nu}_e$ events | 215 |
| 15.5 | Observed inelasticity distribution for selected ν_e and $\bar{\nu}_e$ events | 215 |
| 15.6 | Observed CVN distribution for selected ν_e and $\bar{\nu}_e$ events | 215 |
| 15.7 | Observed cosmic rejection BDT distribution for selected ν_e and $\bar{\nu}_e$ events | 216 |
| 15.8 | Observed ν_μ disappearance spectra in HadEFrac quartiles | 217 |
| 15.9 | Observed $\bar{\nu}_\mu$ disappearance spectra in HadEFrac quartiles | 217 |
| 15.10 | Observed ν_μ Disappearance Spectra | 218 |
| 15.11 | Observed $\bar{\nu}_\mu$ Disappearance Spectra | 218 |
| 15.12 | Data / MC muon energy distribution for far detector ν_μ events | 219 |
| 15.13 | Data / MC HadEfrac distribution for far detector ν_μ events | 219 |
| 15.14 | Observed far detector events as a function of accumulated exposure | 219 |
| 15.15 | Allowed contours in $\sin^2 \theta_{23}$ and δ_{CP} | 221 |
| 15.16 | Rejection significance of δ_{CP} | 222 |
| 15.17 | Allowed contours in Δm_{32}^2 and $\sin^2 \theta_{23}$ | 223 |
| 15.18 | Rejection significance of $\sin^2 \theta_{23}$ | 224 |
| 15.19 | Rejection significance of Δm_{32}^2 | 224 |
| 16.1 | Projected NOvA sensitivity to leptonic CV violation and mass hierarchy | 226 |

LIST OF TABLES

| <i>Number</i> | <i>Page</i> |
|---|-------------|
| 9.1 CC π^0 cutflow | 104 |
| 9.2 Fiducial volume mass, by element | 117 |
| 9.3 Comparison of systematic error sizes for CC π^0 analysis | 131 |
| 9.4 Total cross sections derived by integrating each kinematic distribution, compared to GENIE prediction | 143 |
| 12.1 Total number of simulated events predicted after each decomposition scheme | 180 |
| 12.2 Predicted ν_μ disappearance events | 185 |
| 12.3 Predicted $\bar{\nu}_\mu$ disappearance events | 186 |
| 12.4 Predicted ν_e appearance events | 188 |
| 12.5 Predicted $\bar{\nu}_e$ appearance events | 189 |
| 15.1 Observed events in the appearance and disappearance samples | 212 |
| 15.2 Expected ν_e and $\bar{\nu}_e$ event counts at different oscillation parameters | 214 |
| 15.3 Comparison of observed and predicted disappearance events, by quartile | 216 |

Chapter 1

NEUTRINO PHYSICS

Neutrinos were theoretically postulated prior to detection in 1930 by Wolfgang Pauli[1] to explain the continuous spectrum of positron energies observed in contemporary β decay experiments[2]. Shortly thereafter, Enrico Fermi developed his famed four-fermion theory of nuclear decay[3] in which the neutrino played a central role. The theory predicts the reaction $\bar{\nu} + p \rightarrow e^+ + p$ should occur, but, with the coupling constant measured in nuclear β decay, the cross section was so small this was initially considered un-detectable. The neutrino was first detected in 1954 by Cowan and Reines through observing this inverse β decay reaction outside a nuclear reactor[4]. In 1957, Goldhaber showed the neutrino was only found in left-handed helicity states[5] by observing the polarization of γ decay of excited states produced by nuclear β decay.

In the following years, the ν_μ [6] and ν_τ [7] neutrino flavors were discovered and each grouped with an analogous charged lepton in the standard model. Meanwhile, the long-standing “solar neutrino flux problem” was identified with oscillations between these three neutrino flavor types. This was later confirmed experimentally and gave the first evidence that neutrinos have a finite mass. The problem of neutrino mass is non-trivial. As neutrinos are only found in the left-handed state, some modification of the standard model must be made to accommodate this observation.

1.1 The Free Neutrino and Neutrino Mass

In analogy to massive leptons, a free Lagrangian for the neutrino can be written as

$$L = \bar{\nu} (i\gamma^\mu \partial_\mu) \nu - m\bar{\nu}\nu. \quad (1.1)$$

This spinor can be decoupled into right and left-handed components via a sum of projection operators so that the mass term can be re-written as

$$m\bar{\nu}\nu = m(\bar{\nu}_R\nu_L + \bar{\nu}_L\nu_R) \quad (1.2)$$

and thus, given $m \neq 0$, both ν_L and ν_R must exist by requiring hermiticity. Experimentally, only ν_L is known to participate in weak interactions, and thus ν_R is termed a “sterile” neutrino.

With the inclusion of this sterile neutrino, the mass of the neutrino would be determined from coupling to the Higgs field given by $g_V \langle \phi_H \rangle / \sqrt{2}$. With current limits on $m_\nu < 0.1$ eV, this framework gives an unnaturally small coupling strength between the neutrino and the Higgs field compared to other fermions. This presents a fine-tuning problem.

The charge conjugation operator, C , transforms a spinor into its associated anti-particle, sending $\psi \xrightarrow{C} \psi^c$. C also changes the chirality of a spinor, satisfying $(\psi_L)^c = (\psi^c)_R$. A separate state can be formed as

$$\phi = \frac{1}{\sqrt{2}} (\psi + \psi^c) . \quad (1.3)$$

As defined, ϕ is an eigenvector of C , suggesting that ϕ represents a particle which is its own anti-particle. The neutrino is the only fundamental fermion known which could be represented this way as all other fermions are charged, and thus fermion/anti-fermion pairs are experimentally known to be distinct. Again breaking ϕ into left and right-handed components expands the mass term into

$$\bar{\phi}\phi = (\bar{\phi}_L\phi_R + \bar{\phi}_R\phi_L) \quad (1.4)$$

$$= \frac{1}{2} (\bar{\psi}_L\psi_R + \bar{\psi}_L\psi_R^c + (\bar{\psi}^c)_L\psi_R + (\bar{\psi}^c)_L(\psi^c)_R) + L \leftrightarrow R \quad (1.5)$$

$$= (\bar{\psi}_L\psi_R + \bar{\psi}_R\psi_L) + \frac{1}{2} (\bar{\psi}_L(\psi_L)^c + (\bar{\psi}_L)^c\psi_L) + \frac{1}{2} (\bar{\psi}_R(\psi_R)^c + (\bar{\psi}_R)^c\psi_R) . \quad (1.6)$$

This is the sum of three independent, hermitian mass terms which can be parameterized as

$$\bar{\psi} M \psi = m_D \bar{\psi}_L \psi_R + m_L \bar{\psi}_L (\psi_L)^c + m_R \bar{\psi}_R (\psi_R)^c + h.c. \quad (1.7)$$

$$= \begin{pmatrix} \bar{\psi}_L & (\bar{\psi}_R)^c \end{pmatrix} \begin{pmatrix} m_L & m_D \\ m_D & m_R \end{pmatrix} \begin{pmatrix} \psi_L \\ (\psi_R)^c \end{pmatrix} . \quad (1.8)$$

Here, m_D is an analogy to the Dirac mass shared by other fermions. But, in the Majorana case, there are the additional m_L and m_R parameters allowed to influence the mass eigenstates. The eigenvalues of this matrix are

$$m_{1,2} = \frac{m_L + m_R}{2} \pm \frac{1}{2} \sqrt{(m_L - m_R)^2 - 4m_D^2} . \quad (1.9)$$

In the scenario when $m_R \gg m_L$, the two observable masses are m_R and m_D^2/m_R . This limiting case is the seesaw mechanism. Assuming m_D is near the Higgs vev,

≈ 100 GeV, and a right-handed neutrino state at the GUT scale, $\approx 10^{16}$ GeV, the light neutrino mass is predicted as ≈ 0.01 eV.

Each of these models predicts a neutrino state that is not currently incorporated into the standard model. If the neutrino is a Dirac particle, there is a predicted sterile neutrino eigenstate that does not participate in weak interactions but contributes to a mass term with a unnaturally smallness mass. In the Majorana case, a GUT scale neutrino naturally explains the small neutrino masses observed in nature. Though we only considered neutrino mass in a one-neutrino theory, these scenarios can naturally be extended to incorporate the three neutrino flavor eigenstates that we see in nature, or more.

Several approaches to directly measure the neutrino mass are actively being pursued in the laboratory. There are searches for distortions in the β decay spectrum expected from a finite neutrino mass. With this approach, KATRIN will push limits to 0.2 eV[8] with an enormous conventional spectrometer. Project 8 is similarly searching for endpoint kinematic distortions by measuring the frequency of cyclotron radiation emitted from decay electrons[9] with eventual sensitivity of 0.1 eV[10]. Additionally, $0\nu\beta\beta$ decay searches are sensitive to a Majorana mass. In these experiments, the Majorana mass is measured from the $0\nu\beta\beta$ decay rate. These experiments study the handful of isotopes whose primary decay process is the $2\nu\beta\beta$ process in order to reduce backgrounds for detecting the very rare $0\nu\beta\beta$ process. They have currently probed down to a $\langle m_{\beta\beta} \rangle \approx 0.1$ eV[11] with improved sensitivity in the future.

1.2 Neutrino Oscillations

Neutrinos interact via the weak force in flavor eigenstates of ν_e , ν_μ , and ν_τ . However, neutrinos have been observed to oscillate between flavor eigenstates while traveling through space, suggesting that these flavor eigenstates are not exactly the free neutrino mass states. This was the first indicator that showed neutrinos states have non-degenerate masses. Thus, there must be at least two neutrino states with finite mass.

The PMNS matrix[12] gives an explicit representation of the relationship between flavor and mass eigenstates of neutrinos. The matrix elements are defined numerically as

$$\langle \nu_i | \nu_\alpha \rangle = U_{\alpha i} . \quad (1.10)$$

As both sets form an orthonormal basis, any flavor eigenstate can be written in terms

of PMNS matrix elements as

$$|\nu_\alpha\rangle = \sum_{i=1,2,3} U_{\alpha i} |\nu_i\rangle. \quad (1.11)$$

For the remainder of this work, Greek(Latin) subscripts on a neutrino will refer to the flavor(mass) eigenstates as in this equation. But, as each mass state travels, it accumulates a phase, so that the neutrino wavefunction can be written as a function of time in terms of each mass eigenstate as

$$|\nu(t)\rangle = \sum U_{\alpha i} e^{-iE_i t} |\nu_i\rangle, \quad (1.12)$$

where $|\nu(0)\rangle = |\nu_\alpha\rangle$. As described earlier, the experimental bounds on neutrino mass are sub-eV, so that neutrinos produced in the laboratory can be treated as relativistic with $p_\nu \approx E_\nu$, in which case, the energy for a specific mass eigenstate with a given momentum is

$$E_i \approx p_\nu \left(1 + \frac{m_\nu^2}{2E_\nu^2} \right), \quad (1.13)$$

so that

$$|\nu(t)\rangle = e^{-ip_\nu t} \sum U_{\alpha i} e^{-i(m_i^2/E_\nu)t} |\nu_i\rangle. \quad (1.14)$$

If the matrix $U_{\alpha i}$ is non-diagonal, the flavor composition of this expression after some fixed propagation time will have changed. Ignoring the un-observable $p_\nu t$ phase in front of the sum, the different phases on each of the initial coefficients suggest $\nu(t) \neq \nu(0)$, giving rise to a time-dependent flavor transition. Further, t is just a measure of the oscillation baseline, L , since the neutrinos are relativistic. The oscillation probability to a flavor state β , is given by the projection of $|\nu_\beta\rangle$ with the expression in Eqn. 1.14:

$$P_{\alpha \rightarrow \beta} = |\langle \nu_\beta | \nu(t) \rangle|^2 = \left| \sum U_{\beta i}^\dagger U_{\alpha i} e^{-i(m_i^2/E_\nu)L/c} \right|^2. \quad (1.15)$$

In the two-neutrino case, the structure of U must be

$$\begin{pmatrix} \cos \theta & \sin \theta \\ -\sin \theta & \cos \theta \end{pmatrix} \quad (1.16)$$

by hermiticity. With some simple algebra, the oscillation probabilities between each two state pairs can be written succinctly, assuming the two relevant flavors are ν_e and ν_μ as

$$P_{e \rightarrow \mu} = P_{\mu \rightarrow e} = \sin^2 2\theta \sin^2 \frac{\Delta m^2 L}{4E_\nu}, \quad (1.17)$$

$$P_{e \rightarrow e} = P_{\mu \rightarrow \mu} = 1 - P_{e \rightarrow \mu}. \quad (1.18)$$

The relevant kinematic variable is L/E_ν , with $\sin^2 2\theta$ and Δm^2 as free parameters. As \sin^2 is an even function, the system is degenerate under transformations of $\Delta m^2 \rightarrow -\Delta m^2$. In general, only neutrino oscillations in matter are sensitive to the sign of the mass splitting, as detailed in the final section of this chapter. This transition probability prediction can be readily tested by placing a detector at a fixed baseline and measuring the proportion of ν_e and ν_μ as a function of E_ν . The reliance on Δm^2 , in particular, suggests that if neutrino oscillations are observed, the neutrino masses must be non-degenerate.

A schematic diagram of a neutrino oscillation is shown in Figure 1.1. The setup mirrors the $\nu_\mu \rightarrow \nu_e$ process that is measured in this thesis. Physically, neutrinos are produced as a flavor eigenstate, which is in turn a superposition of mass eigenstates. As shown above, these mass eigenstates accumulate phase at different rates, so that a different physical state passes through a detector, giving the neutrino a certain probability to be observed as a different flavor eigenstate.

The three-neutrino picture is in essence exactly the same. To calculate a transition probability, the bra, $\langle \nu_\beta |$, is contracted with the time-evolved ket. The algebra, however, becomes much more complicated. For n neutrino flavors, the complex unitary transition matrix has $(n - 1)^2$ degrees of freedom, assuming neutrinos are Dirac particles. If neutrinos are Majorana particles, this introduces an additional $n - 1$ degrees of freedom. For $n = 2$, this simplifies to a single mixing angle as in Eqn. 1.17. For $n = 3$, there are three mixing angles, labeled θ_{13} , θ_{23} , and θ_{12} , along with a CP violating phase, δ . Additionally, there are two independent Δm^2 's that set the frequency of oscillations. The PMNS matrix encapsulates these neutrino

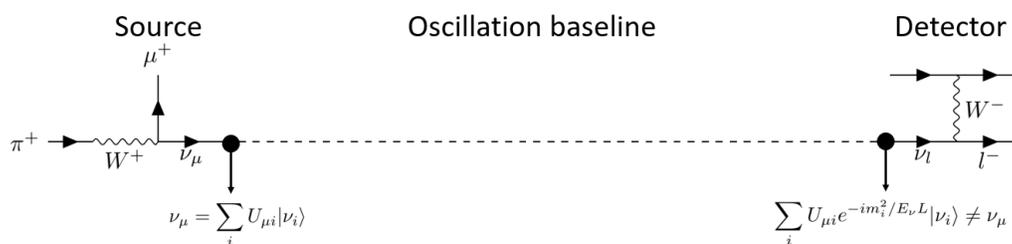


Figure 1.1: A diagram of an experiment that can measure a neutrino oscillation. A ν_μ , a flavor eigenstate, is created from a π^+ decay. After the decay, this is a coherent sum of mass eigenstates. During propagation, the mass eigenstates accrue a phase which depends on the mass of the eigenstate. An average, unobservable phase has been removed from the ket. Thus, if the masses are non-degenerate, the neutrino state entering the detector is no longer the ν_μ flavor state superposition. This leads to non-vanishing probability to observe another flavor eigenstate in the detector.

mixing parameters as

$$U_{\text{PMNS}} = U_1 U_2 U_3, \quad (1.19)$$

$$U_1 = \begin{pmatrix} 1 & 0 & 0 \\ 0 & c_{23} & s_{23} \\ 0 & -s_{23} & c_{23} \end{pmatrix} \quad U_2 = \begin{pmatrix} c_{13} & 0 & s_{13}e^{-i\delta} \\ 0 & 1 & 0 \\ -s_{13}e^{i\delta} & 0 & c_{13} \end{pmatrix} \quad U_3 = \begin{pmatrix} c_{12} & s_{12} & 0 \\ -s_{12} & c_{12} & 0 \\ 0 & 0 & 1 \end{pmatrix}. \quad (1.20)$$

where s_{ij} and c_{ij} are a shorthand for $\sin \theta_{ij}$ and $\cos \theta_{ij}$. The factorization in this way highlights the historical development of neutrino oscillation experiments. U_3 describes ‘‘solar’’ neutrino oscillations with $\Delta m^2 \sim 10^{-5} \text{ eV}^2$ while U_1 describes ‘‘atmospheric’’ oscillations with $\Delta m^2 \sim 10^{-3} \text{ eV}^2$. U_2 is also at the atmospheric Δm^2 and depends on the CP-violating phase. The connection of each of these matrices to fundamental neutrino parameters and the experiments that measured these parameters is the main focus of Chapter 2.

1.3 Neutrino Oscillations in Matter

The above description is enough to entirely describe the neutrino oscillation picture for three neutrino flavors in vacuum. For most experiments, however, this is not an applicable assumption. For instance, solar neutrinos are produced during nuclear fusion in the solar core, and long baseline accelerator neutrino beams pass through hundreds of km of rock before passing through a detector. A framework for calculating transition probabilities was developed by Mikheyev Smirnov[13] and Wolfenstein[14].

An asymmetry in cross section with matter is apparent, as both charged and neutral currents contribute to the leptonic processes $\nu_e + e \rightarrow \nu_e + e$ while $\nu_\alpha + e \rightarrow \nu_\alpha + e$ is only possible through a neutral current interaction if $\alpha = \mu, \tau$. This asymmetry can be treated with a perturbative term to the free Hamiltonian. In the two-neutrino (ν_e and ν_μ) scenario, the algebra is sufficiently simple to show here:

$$i \frac{d}{dt} \begin{pmatrix} \nu_1 \\ \nu_2 \end{pmatrix} = \left(\frac{m_i^2}{2E} \delta_{ij} + \Delta H \right) \begin{pmatrix} \nu_1 \\ \nu_2 \end{pmatrix}. \quad (1.21)$$

H' can be written in the flavor basis:

$$\Delta H = \frac{1}{2E} \begin{pmatrix} A & 0 \\ 0 & 0 \end{pmatrix}, \quad (1.22)$$

$$A = \pm \sqrt{8} G_F E N_e. \quad (1.23)$$

Here, the electron density, N_e , forms the basis for an effective potential which only affects the ν_e flavor eigenstate. The potential is positive for $\nu_e e$ scattering and negative for $\bar{\nu}_e e$ scattering, suggesting the perturbation affects neutrino and anti-neutrino oscillations differently. This explains the form $\Delta H \propto G_F N_e \bar{\nu}_e \nu_e$ with the dimensionless pre-factor determined by a precise calculation[15].

The equation of motion, Eqn. 1.21, however, is written in terms of the mass eigenstates. So, ΔH must have a change-of-basis transformation applied which is just a rotation through θ , where θ is exactly the neutrino mixing parameter. Thus, ΔH can be written in the vacuum mass eigenstate basis as

$$\Delta H' = U^\dagger \Delta H U = \frac{A}{2E} \begin{pmatrix} \cos^2 \theta & \cos \theta \sin \theta \\ \cos \theta \sin \theta & \sin^2 \theta \end{pmatrix}, \quad (1.24)$$

so that the perturbed Schroedinger equation becomes

$$i \frac{d}{dt} \begin{pmatrix} \nu_1 \\ \nu_2 \end{pmatrix} = \frac{1}{2E} \begin{pmatrix} m_1^2 + A \cos^2 \theta & A \cos \theta \sin \theta \\ A \cos \theta \sin \theta & m_2^2 + A \sin^2 \theta \end{pmatrix} \begin{pmatrix} \nu_1 \\ \nu_2 \end{pmatrix}. \quad (1.25)$$

In the mass basis, $\Delta H'$ is no longer diagonal, which shows the vacuum mass eigenstates are no longer exactly the mass eigenstates while traveling through matter. Calculating the eigenvalues and eigenstates is straightforward. For the rest of this section, m_i will denote the mass eigenstate in vacuum while $m_{i,m}$ will represent the state in matter. The adjusted effective masses of the states are

$$m_{1,m/2,m}^2 = \frac{1}{2} \left(m_1^2 + m_2^2 + A \right) \pm \sqrt{(\Delta m^2 \cos 2\theta - A)^2 + (\Delta m^2 \sin 2\theta)^2}. \quad (1.26)$$

The eigenstates still exhibit neutrino oscillations, but with the physical parameters modified:

$$\Delta m_m^2 = \sqrt{(\Delta m^2 \cos 2\theta - A)^2 + (\Delta m^2 \sin 2\theta)^2}, \quad (1.27)$$

$$\sin 2\theta_m = \frac{\sin 2\theta}{\sqrt{\left(\cos 2\theta - \frac{A}{\Delta m^2}\right)^2 + (\sin 2\theta)^2}}. \quad (1.28)$$

We can now estimate the effect this has on neutrino oscillations for NOvA. With appropriate addition of \hbar 's and c 's, the perturbation potential works out numerically as

$$A = 7.6 \times 10^{-5} \text{ eV}^2 \times \frac{E_\nu}{\text{GeV}} \frac{\rho}{\text{g/cm}^3}, \quad (1.29)$$

assuming the medium is electrically neutral with equal proton and neutron densities. For a neutrino energy $E_\nu = 2 \text{ GeV}$, near the NOvA oscillation peak, and a density

of 2.84 g/cm^3 , the mean density of the earth's crust, this gives $A = 4.3 \times 10^{-4} \text{ eV}^2$. As we will cover in the next chapter, $\sin^2 2\theta = 0.082$ and $\Delta m^2 = 0.0025 \text{ eV}^2$ are appropriate for $\nu_\mu \rightarrow \nu_e$ oscillations so that Eqn. 1.28 gives

$$\sin 2\theta_m \approx 1.17 \sin 2\theta \quad (1.30)$$

for the $\nu_\mu \rightarrow \nu_e$ appearance case and

$$\sin 2\theta_m \approx 0.86 \sin 2\theta \quad (1.31)$$

for the $\bar{\nu}_\mu \rightarrow \bar{\nu}_e$ appearance case. On inspection of Eqn. 1.28, it is clear that the neutrino and anti-neutrino adjustments from this matter effect switch if $\Delta m^2 \rightarrow -\Delta m^2$. Thus, the relative enhancement and suppression of ν_e and $\bar{\nu}_e$ appearance is sensitive to the sign of Δm^2 , with a $\sim 36\%$ difference in overall appearance probability between the two polarities in a setup similar to the NOvA experiment.

1.4 Neutrino Interactions

In an experiment, a neutrino scatter is used to measure the scattered neutrino's flavor. Comparing an ensemble of observed neutrino flavors relative to the expected composition at the source is then used to infer an oscillation probability. Thus, a brief overview of neutrino scattering is below. For a more in-depth description, see Appendix A for details on phenomenological scattering relevant at NOvA's beam energies along with production models. This appendix contains useful background for reading Chapter 9 but can be safely bypassed for following the remainder of the thesis.

In the standard model, neutrinos have no color or electromagnetic charge and thus only interact weakly. When speaking of neutrino interactions, it is implicit that we are referring to interactions off nuclear hadrons. Neutrino cross sections are far too low for studying scatters between two beams. There are purely leptonic interactions when scattering off atoms, such as $\nu_\mu + e \rightarrow \nu_e + \mu$ scattering off orbital electrons, but these processes represent a small fraction of the total cross section, as we'll soon see.

Coupling low cross sections, relatively broad-band beams, and poor modeling of nuclear structure, neutrino cross sections on nuclei is a historically challenging problem. Qualitatively scattering off of nuclei depends on the neutrino energy and the four-momentum transferred to the hadronic part. Below $E_\nu \approx 1 \text{ GeV}$, a nucleus is effectively a collection of point-like nucleons. At high energies, $E_\nu \gtrsim 5 \text{ GeV}$,

the four-momentum transfer is large enough to probe nucleon substructure, and scattering off of partons dominates the total cross section. The transitional energy range between 1 and 5 GeV exhibits a mix of processes contributing to the total cross section, as shown in Figure 1.2. In this energy range, resonant events, that produce an intermediate baryon such as Δ_{1232} , are important and lead to pion production. The competing processes here make this range particularly troublesome to measure and model. Exclusive cross section measurements here, particularly looking at final state pion production, are valuable for understanding experimental neutrino samples at this energy.

The Weak Force

Before describing neutrino interactions in detail, it is instructive to first to lay the foundations of the weak force. The remainder of standard model interactions are not fundamental to this work and are described in[17]. The weak force contributes to interactions of each of the fundamental fermions in the standard model. These 12 are grouped into three generations of two leptons and quarks. The leptons are

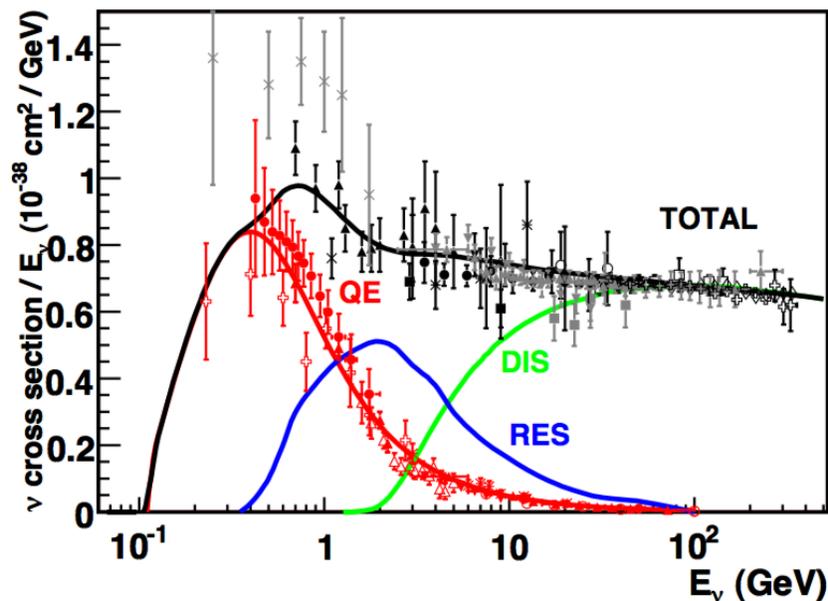


Figure 1.2: The charged current neutrino cross section prediction, as a function of neutrino energy, compared to several measurements. In the 1-5 GeV energy range, QE, resonant, and DIS scattering channels are all significant players in the nuclear scattering environment. Figure from[16].

arranged according to

$$\begin{pmatrix} e \\ \nu_e \end{pmatrix} \quad \begin{pmatrix} \mu \\ \nu_\mu \end{pmatrix} \quad \begin{pmatrix} \tau \\ \nu_\tau \end{pmatrix}. \quad (1.32)$$

The weak force allows interactions between leptons of the same generation. The quarks can similarly be grouped into three generations as

$$\begin{pmatrix} u \\ d \end{pmatrix} \quad \begin{pmatrix} c \\ s \end{pmatrix} \quad \begin{pmatrix} t \\ b \end{pmatrix}. \quad (1.33)$$

Unlike in relation to lepton interactions, the weak force allows for intergenerational mixing for quarks. For example $u - s$ coupling is allowed, but suppressed by a factor of $\sin \theta_C$ where $\theta_C \approx 13^\circ$ [18]. The relationship between couplings of any two quarks is given by the CKM matrix[19]. Both charged current (CC) and neutral current (NC) interactions are possible. These are mediated through the W^\pm (CC) and Z (NC) bosons. Diagrams for three simple processes mediated by the weak force are shown in Figure 1.3.

An interaction Lagrangian can be formulated for a two body CC scattering process with two final states as

$$L = \frac{g_w^2}{8} \frac{i(g_{\mu\nu} - q_\mu q_\nu / M_W^2)}{q^2 - M_W^2} J_\alpha^\mu J_\beta^\nu, \quad (1.34)$$

where q_μ is the four-momentum carried by the W , g_w is a coupling constant, M_W is the W mass, and J_i^μ refers to a interaction current defined as

$$J^\mu = \bar{f}_1 \gamma^\mu (1 - \gamma^5) f_2, \quad (1.35)$$

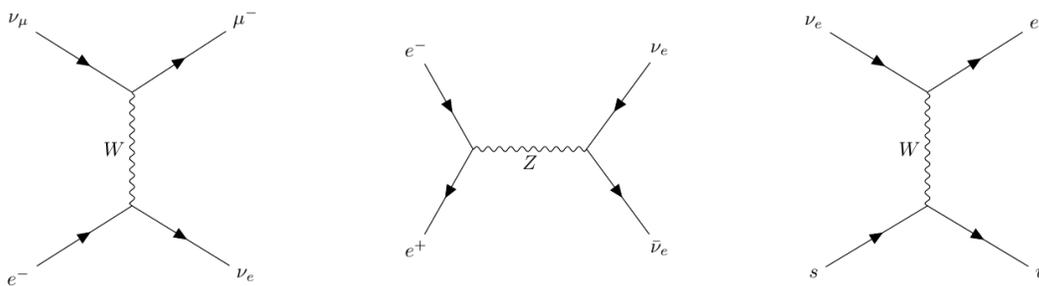


Figure 1.3: Selected Feynman diagrams for a few, simple weak scattering processes. The left gives a purely leptonic CC interaction. The middle calculates the rate for electron-positron annihilation to two neutrinos. The right diagram shows a possible interaction between leptons and quarks, only possible through the Cabibbo mixing.

with f_1 and f_2 representing fermion spinors. This current has a pure $V - A$ coupling. The $1 - \gamma^5$ factor is (twice) the projection operator onto the left-chiral spinor component. With some matrix algebra and using $\{\gamma^5, \gamma^\mu\} = 0$, this can be rewritten as

$$J^\mu = \frac{1}{2} \bar{f}_1 \gamma^\mu (1 - \gamma^5) (1 - \gamma^5) f_2 = \frac{1}{2} \bar{f}_1 (1 + \gamma^5) \gamma^\mu (1 - \gamma^5) f_2 = \frac{1}{2} \bar{f}_{1,R} \gamma^\mu f_{2,L}, \quad (1.36)$$

That is, the weak force can be seen as an interaction between the left-chiral fermions and right-chiral anti-fermions. For high energy scattering, this chiral asymmetry predicts neutrino cross sections to be a factor of three higher than the analogous anti-neutrino scattering process from conservation of angular momentum[17]. Also, this stark asymmetry is responsible for the weak force's parity violation, first observed by decay of polarized ^{60}Co atoms[20]. The $V - A$ coupling form was theoretically formulated shortly afterwards[21].

For a more detailed description of neutrino scattering phenomenology at energies relevant for NOvA and their modeling and theoretical descriptions, see Appendix A.

Chapter 2

EXPERIMENTAL MEASUREMENT OF NEUTRINO MIXING PARAMETERS

As described in Chapter 1, the flavor and mass eigenstates of neutrinos are separate, so neutrinos may oscillate between flavors during propagation. The first experimental signature came from the Homestake experiment[22], which calculated the solar neutrino flux by observing neutrino capture on chlorine: $^{37}\text{Cl} + \nu_e \rightarrow ^{37}\text{Ar} + e^-$. The observed flux was significantly lower than was predicted by solar fusion models. This was long termed the solar neutrino problem, and the neutrino deficit put our understanding of nuclear physics into dangerous territory. The discrepancy was later identified with more sensitive detection techniques as an oscillation of solar-produced ν_e 's oscillating into other neutrino flavors.

Since the discovery of neutrino oscillations, there have been several efforts to directly measure the various parameters of the PMNS mixing matrix. Many of the landmark analyses are detailed in this chapter.

2.1 Solar Mixing Parameters

Fusion processes in the sun produce a large flux of $\sim\text{MeV}$ neutrinos. The solar flux is shown in Figure 2.1 as predicted by solar models. The majority of the flux comes from the $p + p \rightarrow d + e^+ + \nu_e$ process. However, the maximum neutrino energy for this process is less than 1 MeV, and undetectable by most experiments.

The Homestake mine first measured a deficit of about a factor of three of ν_e flux from the sun with a threshold energy of 0.8 MeV[22], posing the solar neutrino problem. This was later confirmed by measuring the ν_e flux using neutrinos from the pp process with capture on gallium by GALLEX[24]. This deficit was clarified with SNO which, unlike previous experiments, was sensitive to ν_e CC, ν_x NC, and $\nu_x + e$ interactions[25]. SNO determined the total neutrino flux using their NC observation, which was consistent with solar flux predictions, but inferred that the ν_e flux was only one-third of what was expected from calculations. There was no oscillation observed in L/E , suggesting that the baseline is several times larger than the oscillation period. This gave the first conclusive evidence for solar neutrinos oscillating between flavors.

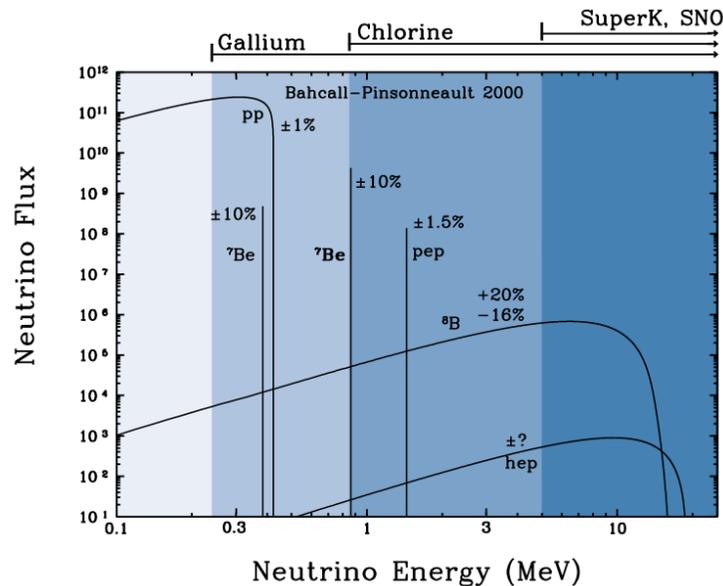


Figure 2.1: The solar neutrino flux, differentiated by the fusion production channel. The top shows the detection threshold for certain classes of experiments. By far the most common neutrinos are from pp fusion, which are below detection threshold for most experiments. Figure taken from [23].

The KamLAND experiment observed anti-neutrinos from over 50 nuclear reactors[26]. The reactors were concentrated in a belt between 140 and 210 km from the detector site in the Japanese Alps. With a few kton-yrs of exposure, multiple oscillation periods in L/E were seen, as shown Figure 2.2. The period of this oscillation in L/E gives a precise measure of $|\Delta m_{12}^2| = 7.59 \pm 0.21 \times 10^{-5} \text{ eV}^2$. This constraint on the oscillation frequency, when including a constraint on $\sin^2 \theta_{12}$ from solar neutrino experiments, focused the allowed parameter space into a single region, as shown in Figure 2.3.

The solar core is quite dense, so that the resonance condition detailed in Section 1.3 is satisfied for $E_\nu \sim 5 \text{ MeV}$ [27]. Thus, solar neutrinos offer a test of the MSW effect, and can determine the sign of Δm_{21}^2 rather than just the magnitude. At energies below the resonance condition, the observed ν_e survival probability should follow the vacuum oscillation prediction when averaged over many periods:

$$P_{e \rightarrow e} = 1 - \frac{1}{2} \sin^2 2\theta_{12}. \quad (2.1)$$

But, at higher energies, solar neutrinos pass through this resonance, and are predicted[27] to have the survival probability

$$P_{e \rightarrow e} = \sin^2 \theta_{12}. \quad (2.2)$$

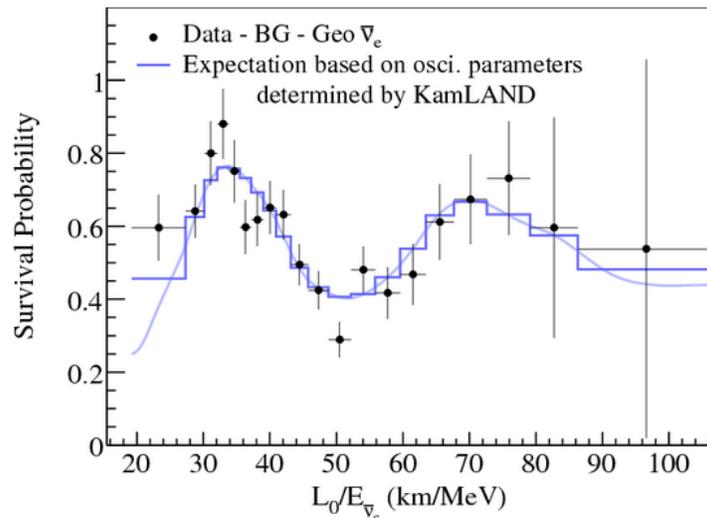


Figure 2.2: The reconstructed L/E distribution for observed $\bar{\nu}_e$ events at KamLAND. Roughly two oscillation periods of the characteristic oscillation curve are observed, thanks to the large energy range and a concentrated belt of nuclear reactors around the detector. Figure taken from [23]

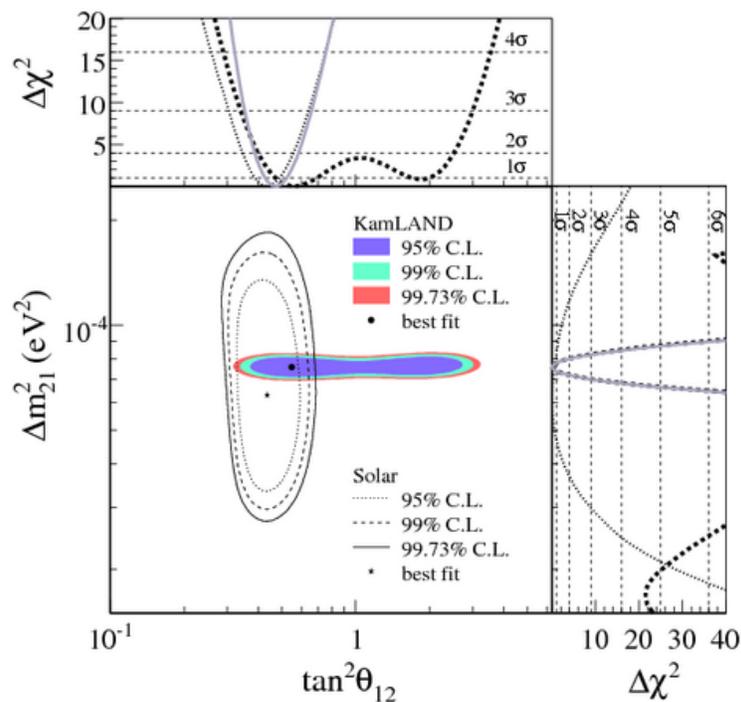


Figure 2.3: A combined fit of KamLAND and solar disappearance data. Solar experiments give a tight constraint on $\tan^2 \theta_{12}$, but their measurement of Δm_{21}^2 spans an order of magnitude. KamLAND complements this data nicely with a precise measurement of Δm_{21}^2 , determined from observing the oscillation frequency. Figure taken from [23]

As, shown in Figure 2.4, the ν_e survival probability agrees with the MSW prediction as a function of energy assuming that ν_2 is more massive than ν_1 . The current best fits to world data give $\Delta m_{21}^2 = 7.59 \times 10^{-5} \text{ eV}^2$ and $\theta_{12} = 33.0^\circ$.

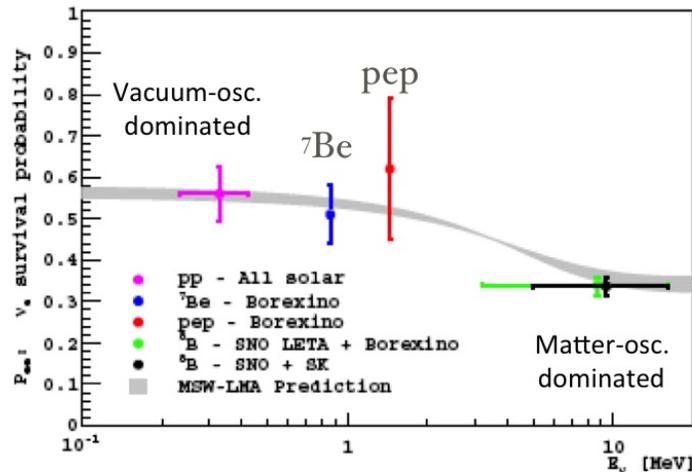


Figure 2.4: The survival probability, as a function of neutrino energy, for the ν_e flux generated in the solar core according to the MSW prediction. The MSW prediction does a good job of explaining disappearance probabilities measured by experiments at different characteristic energies. Figure taken from [28]

2.2 Atmospheric Mixing Parameters

Apart from solar neutrino oscillations, the Super-Kamiokande (Super-K) detector observed oscillations of atmospheric neutrinos. Neutrinos are produced in the upper atmosphere as charged solar particles, primarily protons, scatter. This in turn produces a cascade of mesons, primarily pions, which decay to neutrinos. A survey of measurements of the ν_e and ν_μ fluxes is shown in Figure 2.5 which shows a peak near 1 GeV.

The scale height of the atmosphere is ~ 10 km, while the diameter of the earth is $\sim 10^4$ km. Since the neutrino cross section is so low, the earth is transparent to a neutrino beam and neutrinos produced anywhere in the atmosphere that pass through a detector volume can be studied. This gives a range of baselines between 10 and 10^4 km for analyzing atmospheric neutrinos with energies ~ 1 GeV. Super-K used directional information to determine an oscillation baseline for every atmospheric event tagged. Observed rates of ν_e CC and ν_μ CC interactions were discrepant with the predicted flux of each flavor. Further, when binned in the observed lepton azimuthal direction, the discrepancy is most apparent for upward-going GeV neutrino

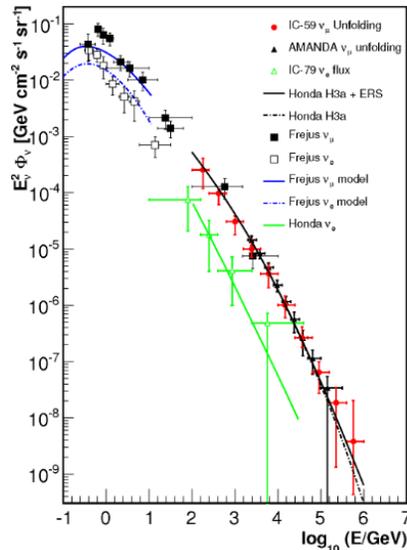


Figure 2.5: The flux of atmospheric neutrinos, as a function of energy. Both ν_e and ν_μ fluxes are represented. The fluxes peak near $E = 1$ GeV. Figure taken from <http://icecube.wisc.edu/news/view/259>

candidates. Sub-GeV samples were less discrepant, suggesting that ν_μ disappear on a baseline of $L/E \sim 10\text{km} / 0.1\text{GeV}$, the scale height of the atmosphere over the typical energy of events in the sub-GeV sample. There was a slight excess of ν_e CC events compared to the flux predictions, but not enough to compensate for the large ν_μ CC deficit. The results are shown in Figure 2.6, and together show evidence for $\nu_\mu \rightarrow \nu_\tau$ oscillations. This was confirmed by directly detecting ν_τ appearance[29] in the neutrino beam between CERN and Gran Sasso.

Analysis of Super-K data suggests that the ν_μ disappearance is nearly maximal, in other words, $\theta_{23} \approx 45^\circ$. Also, $|\Delta m_{32}^2| \sim 2.5 \times 10^{-3} \text{ eV}^2$ was determined, suggesting atmospheric oscillations happen on a baseline much shorter than solar oscillations.

The IceCube collaboration has provided a similar analysis, looking at the disappearance of muon neutrinos between 6 and 56 GeV interacting in Antarctic glacial ice[31]. The two experiments agree on the allowed $\Delta m_{32}^2 - \sin^2 \theta_{23}$ parameter space.

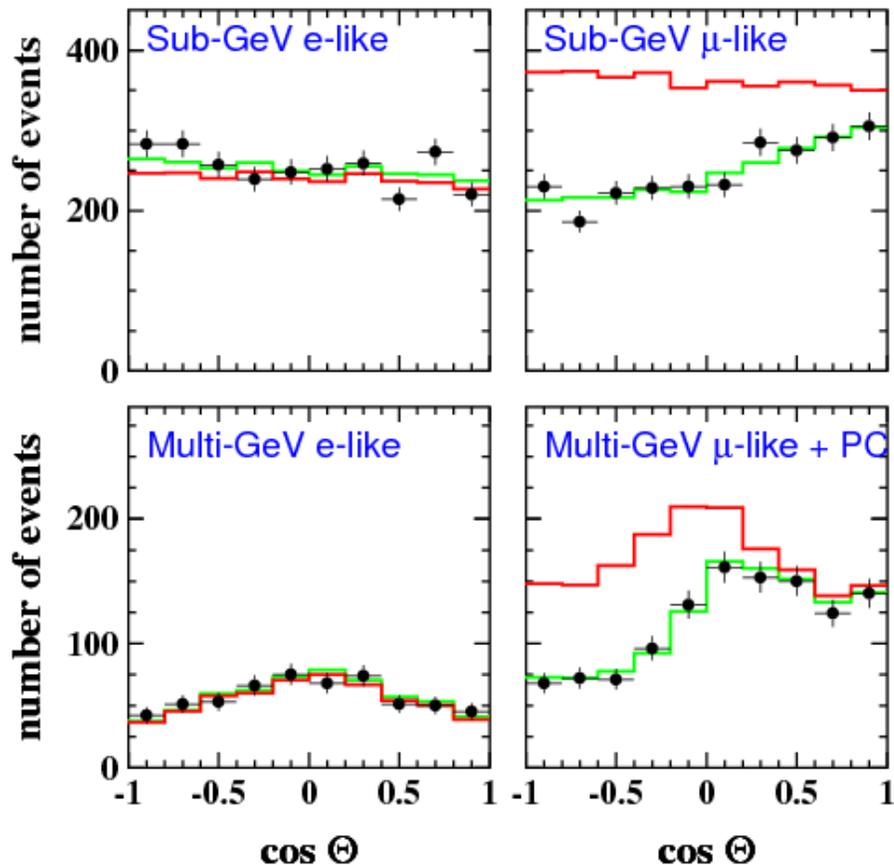


Figure 2.6: Data from Super-K. The left column shows tagged ν_e CC events, while the right shows the corresponding ν_μ CC tagged events. The prediction with the atmospheric flux without oscillations is shown in red. A best fit oscillation prediction is given in green. For the ν_μ CC events, there is a characteristic deficit of upward-going events at high energies. This deficit is nearly maximal for $\cos \theta = 1$. Figure taken from [30]

2.3 Measuring θ_{13} : The Final Mixing Angle

Since the atmospheric Δm_{32}^2 is much larger than the solar Δm_{21}^2 , solar oscillations have little effect on oscillation probabilities over the atmospheric baseline. Thus, the $\bar{\nu}_e$ survival probability is approximately

$$P_{e \rightarrow e} \approx 1 - \sin^2 2\theta_{13} \sin^2 \frac{1.267 \Delta m_{13}^2 L}{E}, \quad (2.3)$$

with Δm_{13}^2 measured in eV^2 , E in GeV, and L in km. Using nuclear reactors as a $\bar{\nu}_e$ source, Daya Bay[32], RENO[33], and Double Chooz[34] have measured the $\bar{\nu}_e$ survival probability to deduce $\sin^2 2\theta_{13}$. All three experiments implemented a two detector approach where a near detector, situated close to the reactor source

compared to the oscillation baseline, is used to measure the un-oscillated prediction and constrain the predicted spectrum at a far detector, placed at the oscillation maximum to measure a $\bar{\nu}_e$ deficit. It was only demonstrated that $\theta_{13} \neq 0$ in 2011, but in the interim, this has become the most precisely known parameter in the PMNS matrix with $\sin^2 2\theta_{13} = 0.082 \pm 0.004$. The observed far detector spectrum is shown in Figure 2.7 for Daya Bay, with a near detector constraint, and the corresponding contour is shown in Figure 2.8.

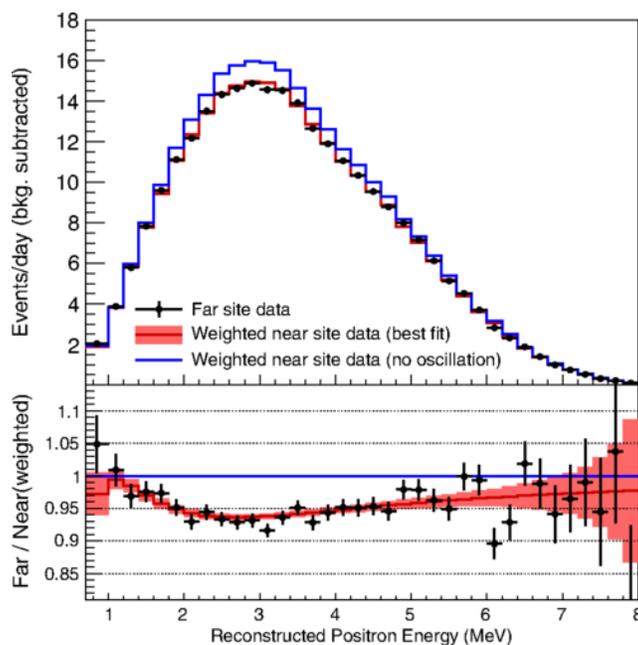


Figure 2.7: The observed data compared to the no-oscillation hypothesis, blue, and the best fit oscillation, red, for the Daya Bay 2012 analysis. Figure taken from [32].

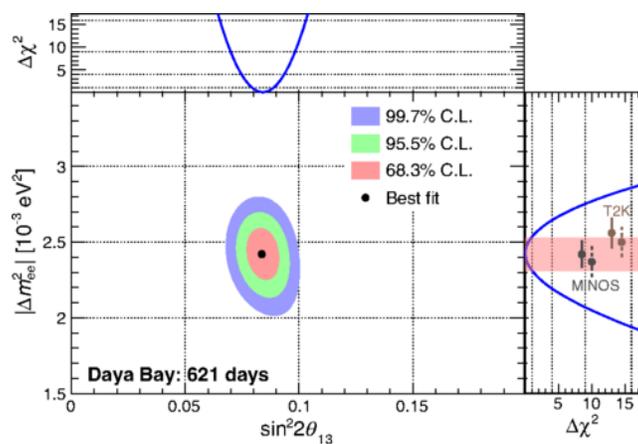


Figure 2.8: Resulting contour from the Daya Bay 2012 analysis in Δm_{13}^2 and $\sin^2 \theta_{13}$. Figure taken from [32].

2.4 Long-Baseline Accelerator Experiments

With the Δm_{32}^2 measured by Super-K, it was clearly possible to study atmospheric neutrino oscillations in a human-engineered neutrino beam. Such a beam can be tuned in energy to produce only neutrinos in the energy range relevant for oscillations. The first maximum oscillation occurs at $L/E \approx 500$ km/GeV. The two-detector K2K and MINOS experiments were built for this measurement. Both confirmed ν_μ disappearance and inferred oscillation parameters consistent with the Super-K result[35][36], again with oscillation probability approximately given by the two flavor scenario. MINOS gave much more precise measurement of $|\Delta m_{32}^2|$ consistent with the Super-K allowed range of $\sin^2 \theta_{23}$. More recently, the T2K and NOvA experiments have similarly studied this disappearance channel with a terrestrial beam. A comparison of the region allowed by MINOS, T2K, and NOvA along with Super-K and IceCube is shown in Figure 2.9.

MINOS first measured the rate of $\nu_\mu \rightarrow \nu_e$ with their detectors whose data preferred a non-zero value of θ_{13} [37]. The vacuum oscillation probability for this channel can be exactly calculated from elements of the PMNS matrix, but this excludes all matter effects. A calculation involving matter interactions as calculated by MSW is

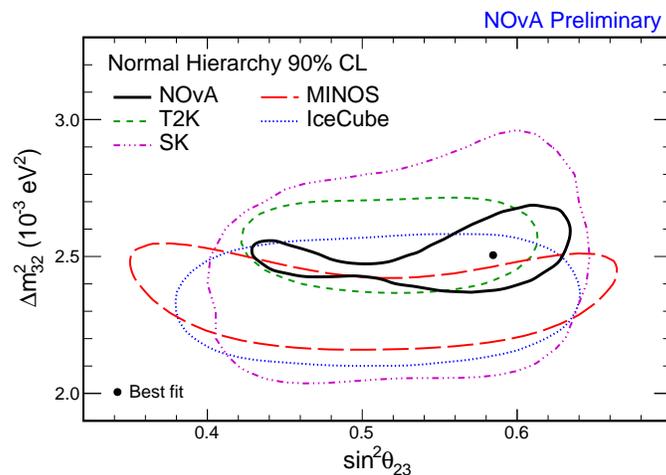


Figure 2.9: A comparison of allowed atmospheric parameter space for the several experiments sensitive to measuring ν_μ disappearance. Both accelerator and atmospheric-based experiments are represented. In this plot, NOvA data represents the results of the analysis shown in this thesis.

given by[38]:

$$\begin{aligned}
P_{\mu \rightarrow e / \bar{\mu} \rightarrow \bar{e}} \approx & \sin^2 \theta_{23} \sin^2 2\theta_{13} \left(\frac{\Delta_{13}}{B_{\mp}} \right) \sin^2 \left(\frac{B_{\mp} L}{2} \right) + \\
& \cos^2 \theta_{23} \sin^2 2\theta_{12} \left(\frac{\Delta_{12}}{A} \right)^2 \sin^2 \left(\frac{AL}{2} \right) + \\
& J_{CP} \frac{\Delta_{12}}{A} \frac{\Delta_{13}}{B_{\mp}} \sin \left(\frac{AL}{2} \right) \sin \left(\frac{B_{\mp} L}{2} \right) \cos \left(\pm \delta - \frac{\Delta_{13} L}{2} \right),
\end{aligned} \tag{2.4}$$

where in the above we have used the following abbreviations:

$$\Delta_{ij} = \Delta m_{ij}^2 / E_{\nu} \tag{2.5}$$

$$A = \sqrt{2} G_F n_e \tag{2.6}$$

$$B_{\mp} = \sqrt{(\Delta_{13} \cos 2\theta_{13} \mp A)^2 + (\Delta_{13} \sin 2\theta_{13})^2} \tag{2.7}$$

$$J_{CP} = \cos \theta_{13} \sin 2\theta_{12} \sin 2\theta_{23} \sin 2\theta_{13} . \tag{2.8}$$

This expression for the oscillation probability involves many fundamental neutrino parameters. Due to matter effects, it depends on the sign of Δm_{32}^2 , while vacuum probabilities only depend on the magnitude. In other words, measuring such a channel can determine the mass hierarchy: whether ν_3 is heavier or lighter than $\nu_{1,2}$. Numerically, the first term in this expression is the largest. This term involves a factor of $\sin^2 \theta_{23}$ rather than the factor of $\sin^2 2\theta_{23}$ which governs the two-flavor oscillation scenario suitable for measuring $\nu_{\mu} \rightarrow \nu_{\mu}$ survival. Since $\theta_{23} \approx \pi/4$, measuring this channel also gives sensitivity to the octant problem: whether θ_{23} is greater than, less than, or exactly $\pi/4$. Lastly, the final term depends on the CP violating phase, showing this channel is sensitive to CP violation in the leptonic sector.

The $\nu_{\mu} \rightarrow \nu_e$ channel is a crucial measurement for studying fundamental neutrino properties. With sensitivity to the mass hierarchy and octant problems and potential for a discovery of leptonic CP violation, the NOvA experiment was built specifically for this measurement. To break degeneracies between unknown parameters in the appearance probability, NOvA has collected large datasets of neutrino and anti-neutrino running. This thesis describes NOvA's first oscillation analysis to include anti-neutrino data, thus measuring both ν_e and $\bar{\nu}_e$ appearance. This inclusion of anti-neutrino data adds to NOvA's history in successfully measuring $\nu_{\mu} \rightarrow \nu_e$ appearance and $\nu_{\mu} \rightarrow \nu_{\mu}$ disappearance.

Leptonic CP violation would be a suggestive clue in studying the nature of neutrino mass along with the evolution of our universe. Within the seesaw model of neutrino

mass, necessitating that neutrinos be Majorana particles, leptogenesis may occur in the early universe. Leptogenesis gives a mechanism for a natural preference for lepton production over anti-lepton production by CP-violating exchanges of right-handed neutrinos[39]. Under the model, these interactions were abundant when the universe was much hotter, on the scale of the mass of the lightest right-handed neutrino eigenstate participating in the seesaw. This may also contribute to a baryon asymmetry in the current universe through $B + L$ violating processes. The degree of CP-violation in this early stage of the universe can be connected with the CP-violation observed in neutrino oscillation experiments through various models. The value of $\sin^2 \theta_{13}$ determined by reactor experiments is large enough that the baryon asymmetry observed in the current universe can be explained through the leptogenesis model if the CP-violating angle in the PMNS matrix has $|\sin \delta_{CP}| \sim 0.5$ or greater. Accelerator oscillation experiments can directly test this condition which, if confirmed, would reinforce an important interplay between cosmology and elementary particle physics.

The T2K experiment is also actively measuring the appearance and disappearance channels in a $\nu_\mu/\bar{\nu}_\mu$ beam. $\delta \sim 3\pi/2$ and the normal mass hierarchy best describe the T2K data[40], though the statistical significance remains relatively low. The 90% allowed ranges of δ for both hierarchies are shown in Figure 2.10.

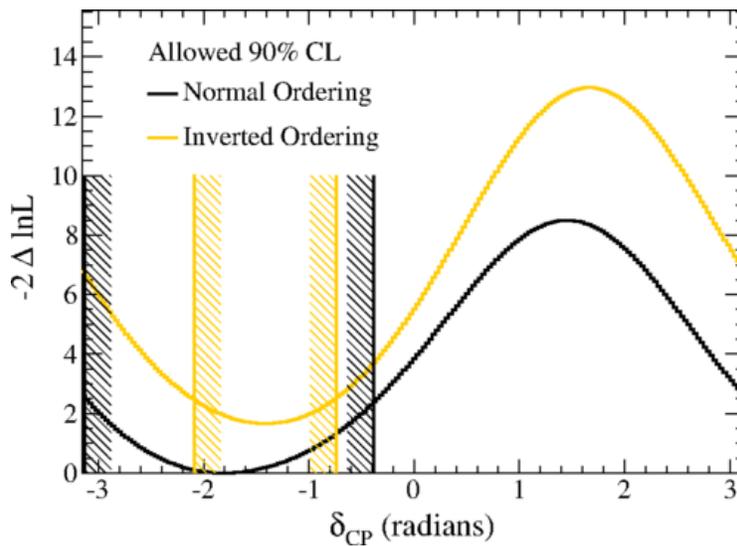


Figure 2.10: The constraint on allowed δ_{CP} from the T2K experiment. Both ν_e and $\bar{\nu}_e$ appearance measurements are included in the fit. The exclusion bands are separated according to assumptions on the true hierarchy. Figure taken from [40].

Chapter 3

THE NOVA EXPERIMENT

The NuMI Off-axis ν_e Appearance experiment (NOvA) is a two-detector oscillation experiment using the Neutrinos from the Main Injector (NuMI) beam from Fermilab. A far detector, 810 km downstream of the NuMI beam in Ash River, MN, is used to measure an oscillated neutrino spectrum, and a near detector is situated at Fermilab 1 km downstream of NuMI. Near detector data is used to constrain systematic uncertainties in the predicted far detector spectrum with a large dataset unaffected by neutrino oscillations at the atmospheric or solar Δm^2 . Both detectors are functionally identical and were designed to efficiently identify ν_e CC events for studying $\nu_\mu \rightarrow \nu_e$ oscillation in the beam. This is done by sampling electromagnetic showers several times per radiation length.

3.1 The NOvA Detector Design

A liquid-scintillator tracker was chosen as the design technology to image final state electrons in ν_e CC events. The detector samples an event's topology several times per radiation length by finely segmenting the detector and using low-Z materials during construction. The far detector is 14 kton while the near detector is only 280 ton as it is much closer to the beam source. A liquid scintillator is the active material in the detector. The scintillator is segmented by PVC cells. This PVC is the primary structural support for the detector in an effort to keep the active mass fraction as high as possible while maintaining a low-Z medium. Additionally, a specially designed optical fiber is used to trap scintillation light and carry it to the light-detecting element, an avalanche photodiode.

Also of note, the far detector is on the surface with only 3.76 mwe of overburden[41] which gives a very high rate of cosmic ray particles tracked through the detector. The near detector is approximately 100 m underground with a smaller surface area giving a negligible cosmic rate during beam spills.

Liquid Scintillator

The liquid scintillator used in NOvA is the experiment's active material and accounts for 69% of the detector mass[42]. The scintillator is a mixture of three components: the scintillant, a WaveLength-Shifting (WLS) agent that absorbs light in the range

emitted by the scintillator and re-emits in the absorption band of the optical fiber, and a solvent that maintains an even concentration of the fluors.

The NOvA scintillant is pseudocumene which decays in the ultra-violet at wavelengths of 270-320 nm. The WLS agent consists of two fluors: PPO and bis-MSB. Scintillation light from the pseudocumene excites the PPO which subsequently emits light with wavelengths of 340-380 nm. The second WLS fluor, bis-MSB, absorbs light emitted by the PPO and re-emits in the 390-440 nm range[42]. The emission spectra for the WLS fluors and the NOvA blended scintillator are shown in Figure 3.1.

The solvent used is a mineral oil. The choice of mineral oil only does not affect the light production in the fiber absorption band and was thus set by requiring the attenuation be >3.25 m[42] for light at a wavelength of 420 nm. As the NOvA cell cross section is 3.7 cm by 5.6 cm this allows the light to traverse the cell several times to increase the probability of being absorbed by the fiber.

The NOvA scintillator is an even mixture of, by mass, 94.6% mineral oil, 5.2% scintillant, and 0.1% WLS compounds. There are also a trace amounts of stadis-425 to increase the electrical conductivity, preventing an accumulation of static charge and a sudden, catastrophic discharge, and vitamin E to prevent degradation of the scintillator transparency over the lifetime of the experiment.

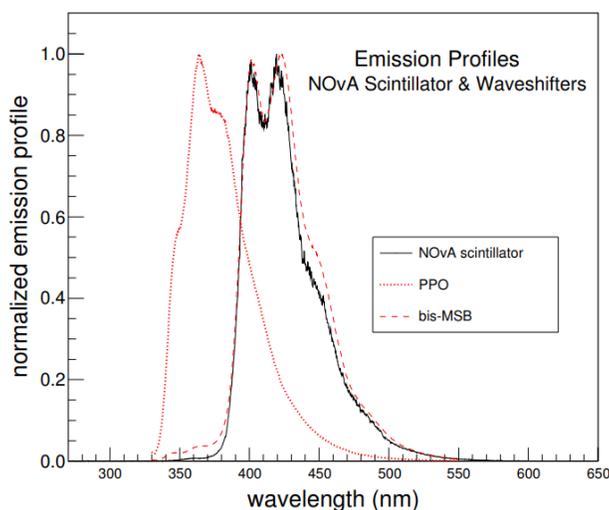


Figure 3.1: The shape of the emission spectrum for the WLS fluors shown in red. The black curve shows the emission spectrum of the NOvA scintillator blend. The black curve has a very similar shape to the bis-MSB WLS fluor, the last emitter chain.

The index of refraction of the liquid scintillator is high enough to emit Cerenkov radiation as fast particles pass through. This Cerenkov light is in ultraviolet, and a portion of the emission spectrum excites the PPO fluor and subsequently the bis-MSB. This Cerenkov light increases the light yield of fast particles by a few percent[43].

WLS Fiber

A custom optical fiber was designed to collect light from the liquid scintillator and carry it to a readout. The fiber is doped with fluorescent dye whose emission and absorption are shown in Figure 3.2. The absorption region falls in the wavelength range of 370-480 nm while photons are emitted at lower energies with wavelengths of 470-570 nm.

Light in the lower portion of the emission spectrum overlaps with the absorption spectrum of the dye. This will attenuate light at these lower wavelengths. The simulated attenuation after a variety of fiber lengths is shown as a function of wavelength in Figure 3.2.

The optical fiber is constructed of a plastic core with two acrylic claddings. The index of refraction of the core is 1.59 and 1.49(1.42) for the inner(outer) coatings[44]. The differences in index of refraction in the coatings give a relatively broad range of incident light that is re-scattered into the fiber.

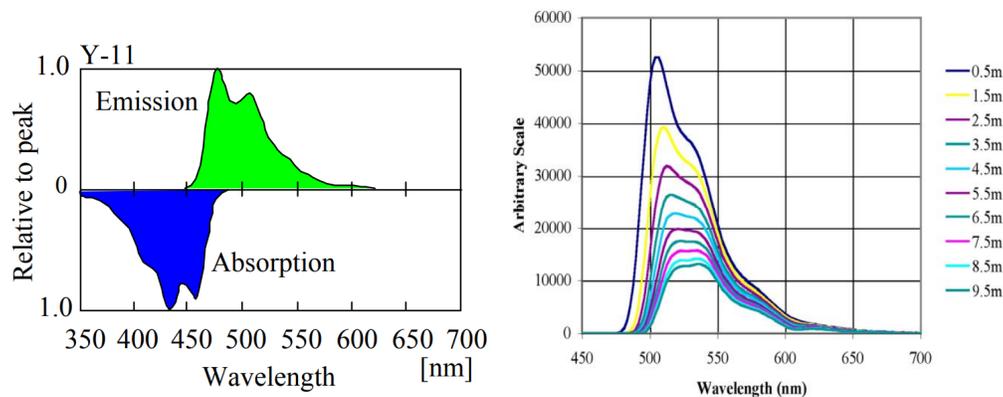


Figure 3.2: On the left, the absorption and emission shape of the fiber dye is shown as a function of wavelength. The right shows the effect of attenuation within the fiber for different lengths traveled through the fiber. At lower wavelength, the attenuation is greater as there is more overlap with the dye absorption spectrum precipitating re-scattering.

The dye concentration was set to 300 ppm to maximize the light output at the far end.

Avalanche PhotoDiode (APD)

The NOvA detectors employ APD's to amplify small photon signals received from the optical fiber. Each APD houses an array of small, single-pixel silicon detectors. In each silicon detector, an incident photon will free an electron from a silicon atom making an electron/hole pair. A high reverse bias voltage, near 425 V, is applied across the silicon wafer creating an electric field ahead of a p-n junction. This causes freed electrons to drift past the p-n junction where the electron signal avalanches, leading to a high-gain amplification of the signal. The voltage applied to each APD is set to ensure a consistent amplification across all channels in the detector[44]. During the commissioning of the detector, the gain was set to 100 and was subsequently increased to 150 to increase the signal-to-noise in the detector.

The quantum efficiency for NOvA's APD's is over 80% over the entire dye emission spectrum. This is shown in Figure 3.3, along with a typical quantum efficiency curve for a PhotoMultiplier Tube (PMT) and the NOvA scintillator emission spectrum.

Thermal noise in the silicon can stochastically make an electron-hole pair. This can lead to random noise amplified in the APD which may fake a physics hit.

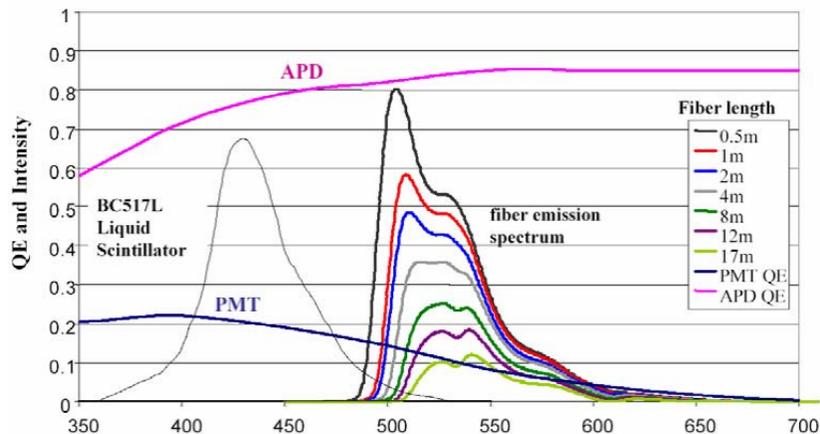


Figure 3.3: The quantum efficiency for NOvA APD's is shown in pink as a function of wavelength. The emission spectrum of the WLS dye is shown in the 480-600 nm range for a variety of attenuation distances. In the overlap region with the dye, the quantum efficiency is over 80%. For comparison, the PMT quantum efficiency is shown in blue with the scintillator emission spectrum in gray. The quantum efficiency for a PMT is significantly lower over all pictured wavelengths.

To reduce the thermal noise expected, each APD is cooled to -15° C. Test-stand measurements show that, at this temperature, the electronic noise was equivalent to 4 signal photoelectrons. For reference, a typical charged particle at the far end of a cell in the far detector is expected to give 40 photoelectrons.

To maintain proper environmental conditions, each APD is hermetically sealed from the ambient environment. A ThermoElectric Cooler (TEC) maintains the APD at a steady temperature, and a heat sink is attached to the TEC to remove waste heat from the cooling process. Humidity is also a concern as the APD operates at a much lower temperature than the detector hall. Therefore, dry gas is pumped into the compartment housing each APD.

The APD's used in the experiment are mass-produced by Hamamatsu, a private vendor that manufactured over 12000 units for NOvA. Each APD unit has 32 distinct pixels used for reading separate channels in the detector. The optical fiber that collects scintillation light connects directly to the surface of each of these pixels.

The NOvA Cell

The NOvA cell is the fundamental detection unit of the detectors. Cells are 15.7(4) m long in the far (near) detector. The average interior cross section of a cell is 5.64 cm parallel to the beam direction and 3.73 cm in the transverse direction[45]. A cell only measures the horizontal or vertical position component, depending on the cell's orientation. Additionally, the cells are rounded near the edges to reduce structural strain[45]. Cells are made of PVC to make the detector more structurally secure while maintaining the low-Z design goals. A schematic of a cell is shown in Figure 3.4 and a technical drawing of the cell cross section is shown in 3.5.

A cell is filled with liquid scintillator and a WLS fiber is placed in the cell. The PVC segmentation isolates light deposited locally in the scintillator in each cell. The PVC is mixed with TiO_2 allowing light to bounce several times within a cell and increasing the absorption probability in the fiber. The fiber is connected directly to a channel on an APD.

Cells are produced by extruding PVC. This process creates an extrusion containing 16 cells. The extrusions are capped at the far end to contain liquid scintillator.

Detector Assembly

The detectors are built with large collections of cells. A module is made of gluing two 16-cell extrusions together. A single 32-channel APD reads out cells by module

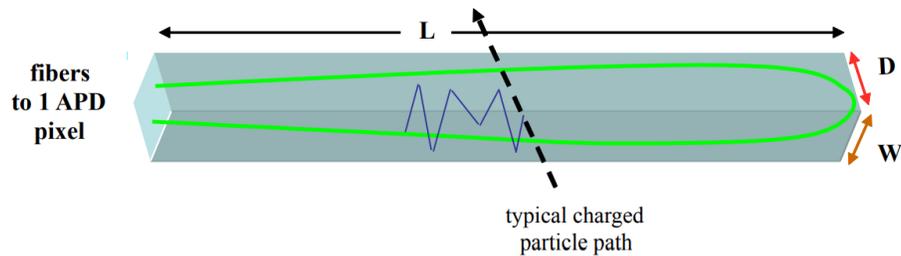


Figure 3.4: A schematic view of a NOvA cell, filled with liquid scintillator and a wavelength-shifting fiber. As a charged particle passes through the cell, the scintillator releases photons that are collected by the fiber and taken to a single pixel on an APD. Note the example photon bounces off the edges of the cell several times before capture, allowed by the reflectivity of the TiO_2

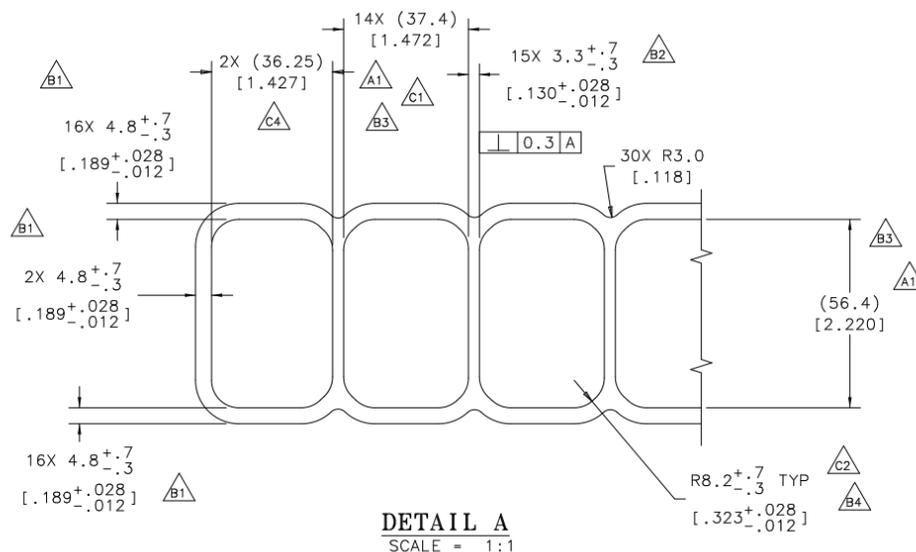


Figure 3.5: The cross section design for the NOvA cells. The interior dimensions are 3.7 cm by 5.6 cm, and the edges are rounded off to mitigate structural strain.

so that there is a one-to-one correspondence between APD and module. After a module is glued, a WLS fiber is placed in each cell. The fiber is looped so that the two ends are accessible at the top of each cell to connect to the APD.

As the far and near detectors are very different sizes, the remaining procedure is different for the two. In the far detector, a plane is made by gluing 12 modules together. A plane is approximately as tall as it is wide (≈ 15.7 m) and is the total extent of the detector perpendicular to the beam direction.

A block is then constructed by gluing 32 planes together in the beam direction. The planes alternate between a horizontal and vertical orientation as shown in

Figure 3.6. Planes with cells running horizontally (vertically) measure the vertical (horizontal) position of the cell to the cell width (3.7 cm) while the other coordinate is only measured to the cell length (15.7 m). Because of this, vertically-oriented planes are also referred to as "X-View" planes while horizontals are called "Y-View" planes. Correlating hits from planes with alternating orientation then allows a 3D reconstruction of events.

Two consecutive blocks make a diblock with a 1 kton mass. The full detector comprises 14 diblocks placed consecutively. The data acquisition electronics for each diblock is entirely separate. This allowed for an incremental commissioning phase where data was read from a partial detector before all blocks had been placed and commissioned. This proved invaluable as it gave quick access to cosmic data as a test of detector and performance calibration procedures. Data from these partial detector periods is also included in physics analyses so that relevant data can still be collected as maintenance work is done on the detector.

The near detector construction is quite similar though much smaller. Each plane is only three modules, or 4 m, across. And the detector is only three diblocks long. In the beam direction, the near detector is only 12.5 m long, which corresponds to an energy loss of 2.2 GeV for a Minimum Ionizing Particle (MIP) such as a muon from

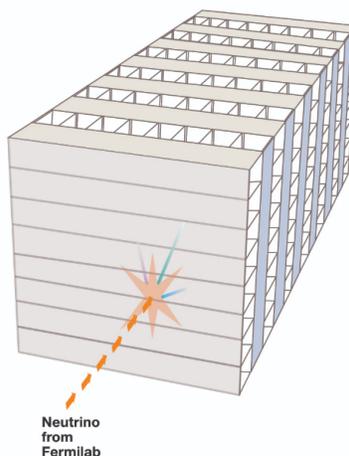


Figure 3.6: A diagram showing the alternating orientation of NOvA planes. The most upstream plane shows cells oriented horizontally and gives a detailed measurement of a particle's vertical location during a physics interaction. The next plane is oriented vertically and gives horizontal information on the particle trajectory. Graphic shown courtesy of Fermilab, <http://news.fnal.gov/2014/10/fermilabs-500-mile-neutrino-experiment-up-and-running/>

a ν_μ CC event. As this is near the characteristic energy of the neutrino oscillations on an 810 km baseline, an additional muon ranger is added to the back of the detector to aid event containment. The muon ranger is constructed by alternating 4-inch planes of steel with two outfitted NOvA planes, one in each orientation[46]. There are 11 groups of steel and NOvA planes which allows the near detector to contain muons with energies up to 5 GeV. As a cost-saving measure, the muon ranger is only $2/3$ as high as the rest of the detector in order to re-use steel planes constructed for the NOvA surface prototype detector, now de-commissioned.

For both detectors, it is interesting to note that the PVC extrusions provide most structural support for the detector. This keeps the active fraction of the detector high which in turn increases the identification efficiency of electron showers in the detectors.

For example, an example neutrino event recorded in the far detector is shown in Figure 3.7. Figure 3.8 shows the same event with a closeup on the event to show the detail recorded by the NOvA detectors. The figure shows the top- and side-views of the event which are used in tandem for three dimensional reconstruction.

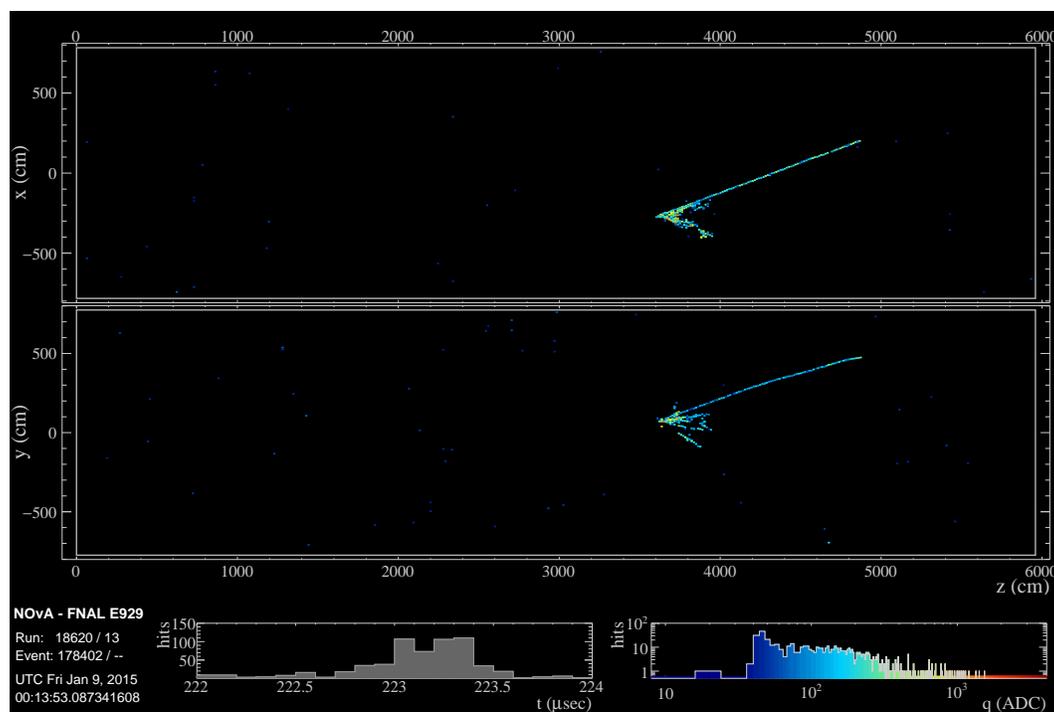


Figure 3.7: A high-energy neutrino event observed in far detector data. The top pane shows a view of the event from the top while the bottom pane shows a side view.

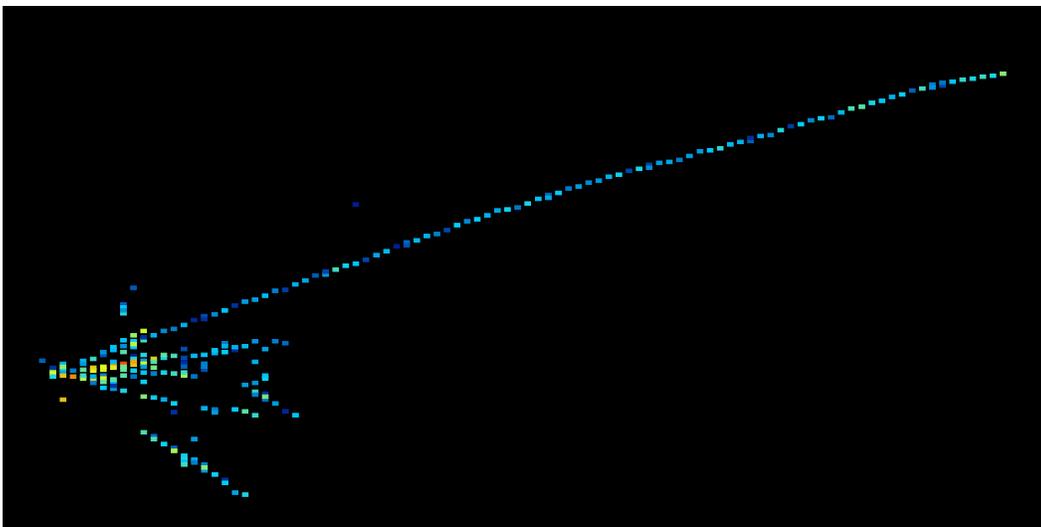


Figure 3.8: A closeup of the same event, illustrating the detailed reconstruction allowed by the fine-granularity of the NOvA detectors.

3.2 The NOvA Physics Program

The primary physics goal is to measure neutrino oscillations in a $\nu_\mu/\bar{\nu}_\mu$ beam, with the sign configurable. By analyzing $\nu_\mu \rightarrow \nu_\mu$ disappearance together with $\nu_\mu \rightarrow \nu_e$ appearance in both neutrino and anti-neutrino beam running, NOvA has sensitivity to place leading constraints on neutrino mixing parameters. The disappearance analysis constrains $\sin^2 2\theta_{23}$ and Δm_{32}^2 which drive the atmospheric disappearance analyses pioneered by Super-K. As described in Chapter 1, the lowest order contribution to the appearance probability depends on the mass hierarchy, $\sin^2 \theta_{23}$, Δm_{32}^2 , and δ_{CP} . These several unknown parameters complicate analyses but give the experiment rich sensitivity to understanding neutrino mixing properties.

Additionally, the NuMI beam is configurable to run in both neutrino and anti-neutrino mode so that the experiment can analyze oscillations in a high purity beam of ν_μ and $\bar{\nu}_\mu$ separately. This configurable beam sign gives NOvA sensitivity to test both the CP symmetry by comparing $\nu_\mu \rightarrow \nu_e$ and $\bar{\nu}_\mu \rightarrow \bar{\nu}_e$ oscillation probabilities and the CPT symmetry by comparing $\nu_\mu \rightarrow \nu_\mu$ and $\bar{\nu}_\mu \rightarrow \bar{\nu}_\mu$ oscillations without relying on a particular neutrino oscillation model. The oscillation probability curves for both the ν_e and $\bar{\nu}_e$ oscillation probabilities are shown in Figure 3.9 in the NOvA setup.

The projected sensitivity of the NOvA oscillation analyses depends on the particular true oscillation parameters. The experiment's sensitivity to rejecting CP symmetry and the inverted mass hierarchy is shown in Figure 3.10 as a function of year with

proposed beam running for the choice of oscillation parameters favored by NOvA's 9.49×10^{20} neutrino mode data.

Additionally, the NOvA near detector will collect enormous datasets of neutrino interactions, several million $\nu_\mu/\bar{\nu}_\mu$ CC fiducial events. These events will be helpful in advancing the understanding of neutrino interactions in the ~ 2 GeV region. As described in Section 1.4, the few GeV region is the most poorly understood region for neutrino scattering as there are multiple phenomenological processes that contribute. This region also dominates several long-baseline neutrino oscillation experiments.

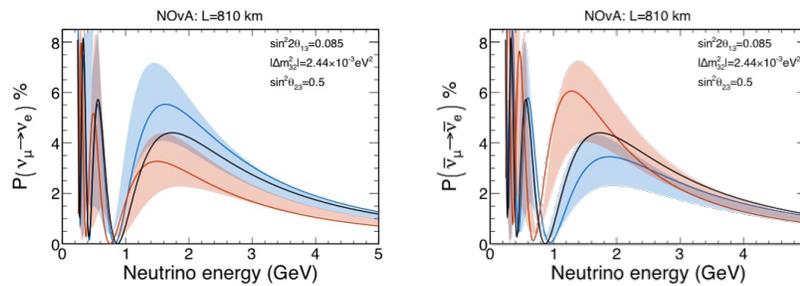


Figure 3.9: The ν_e (left) and $\bar{\nu}_e$ (right) appearance probabilities as a function of energy for the NOvA experiment. The blue (red) shaded regions give the spread in oscillation probability expected by varying δ_{CP} for the normal (inverted) mass hierarchies. The solid colored lines show the oscillation probability for each hierarchy in the case of no CP violation, $\delta_{CP} = 0$, and the black shows the vacuum probability curve with $\delta_{CP} = 0$ and is the same for both neutrinos and anti-neutrinos.

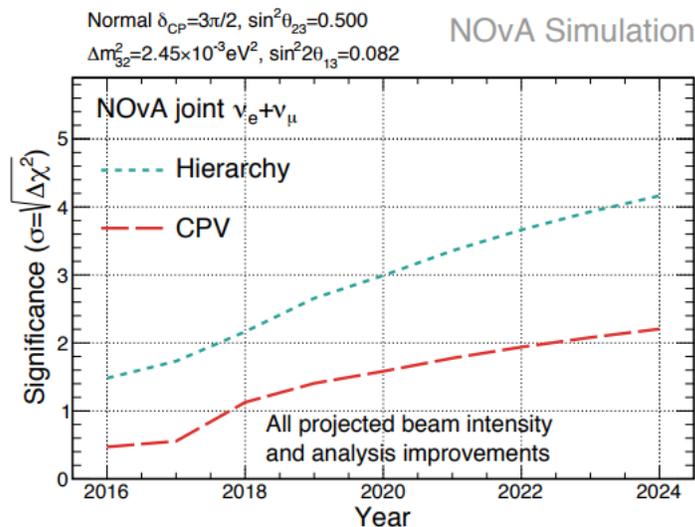


Figure 3.10: NOvA sensitivity to rejecting the no CP violation hypothesis, red, and the wrong mass hierarchy, blue, vs time.

Of particular note, NOvA's detection technology affords a detailed reconstruction of final state particles in the hadronic final state. There is still limited data available to test generation models for hadronic states, but models typically do not accurately predict the hadronic states across a large band of neutrino energy. Observing outgoing hadronic momenta distributions will be helpful for characterizing backgrounds for oscillation analyses. For instance, π^0 production in ν_μ CC and NC events is known to mimic the primary electron in a ν_e CC event. This is an important background for $\nu_\mu \rightarrow \nu_e$ oscillation measurements.

Though not discussed further in this thesis, the large far detector exposure and fine-grained calorimetry allowed by NOvA cells affords precise event classification for many different event topologies. This gives NOvA sensitivity to search for BSM physics in previously unconstrained phase space with many physics measurements. There are active searches for magnetic monopoles[47][48], the ν_μ magnetic moment[49], and $n - \bar{n}$ oscillations[50]. The near detector, with its close proximity to the NuMI beam and low cosmic background, is also sensitive to exotic particles. For instance, NOvA is sensitive to a search for elastic scattering of dark matter particles with masses near 300 MeV[51] which may be produced in the high intensity proton-nucleus collisions used to generate the neutrino beam.

Chapter 4

THE NUMI BEAM

NOvA is an accelerator-based oscillation experiment. As such, a high-intensity neutrino beam is vital to its success. NOvA uses the NuMI beamline which was recently upgraded to double the intensity to stably run at 700 kW. The beamline is configurable to produce a high purity beam of either ν_μ or $\bar{\nu}_\mu$. The NOvA far detector has accumulated data from 9.49×10^{20} protons-on-target (POT) for neutrino running and 6.91×10^{20} POT for anti-neutrino running, over a nominal year of exposure for both configurations.

The NOvA detectors are placed 14 mrad off-axis of the NuMI beam. This has the effect of creating a narrow-band neutrino flux centered near the oscillation maximum of 1.6 GeV for an 810 km baseline.

4.1 Current Fermilab Accelerator Program

After decommissioning the Fermilab Tevatron, the lab heavily invested in repurposing the accelerator campus for several medium-energy, high-intensity experiments through the Accelerator and NuMI Upgrades (ANU). There are three principal beams produced at Fermilab. The NuMI beam was constructed for the MINOS experiment[52] and is currently in use by the NOvA and MINERvA experiments. The Booster Neutrino Beam (BNB) creates a beam of lower energy neutrinos studied with experiments in the short-baseline program such as MicroBooNE and SBN. All three of these beams share the source and initial beam stages within the Fermilab accelerator complex. The third beam fuels two precision muon experiments, g-2 and Mu2e. An overview of the Fermilab accelerator is shown in Figure 4.1.

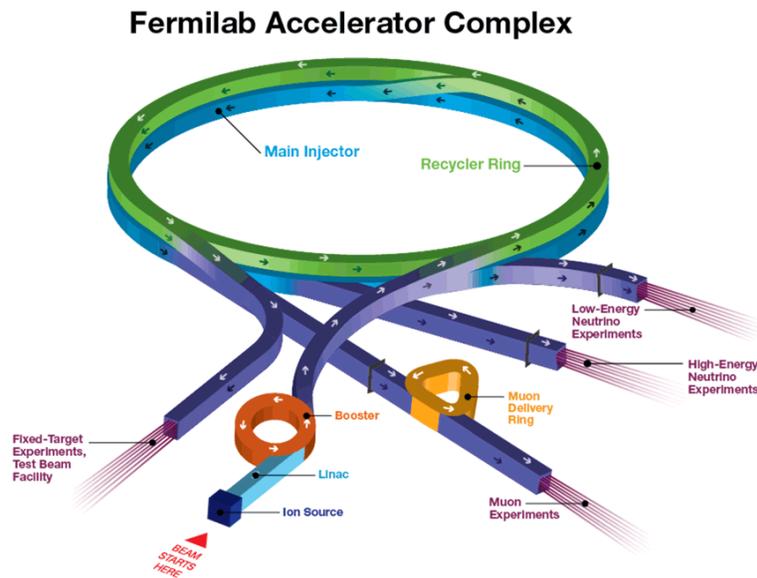


Figure 4.1: An schematic layout of the current Fermilab accelerator program. Graphic shown courtesy of Fermilab, <http://www.fnal.gov/pub/science/particle-accelerators/accelerator-complex.html>.

4.2 The Neutrinos at the Main Injector (NuMI) Beam

As is typical for accelerator-based neutrinos beams, the NuMI beam is a tertiary beam. A primary beam of protons is built and fed into the main injector ring. These protons are then steered into a target which produces a secondary meson beam which in turn decays into a neutrino beam used by the NOvA experiment.

Linac and Booster Accelerators

The proton beamline uses a bottle of hydrogen gas as fuel. A linac accelerates H^- ions to 400 MeV and directs groups of protons into the proton booster, which accelerates the proton beam to 8 GeV.

The beam is segmented into buckets which are regions of phase space which lock into resonance with the RF cavities and can stably orbit the Fermilab rings. A bunch refers to the group of protons that lie in a single bucket. The booster defines batches as 84 consecutive proton batches. This is the proton unit that can be injected into subsequent stages of the accelerator. Each batch is $1.6 \mu s$, and the booster generates batches at 15 Hz. The main injector can accommodate six proton batches. These

six batches are collectively referred to as a spill.

Proton Stacking in the Recycler Ring

The anti-proton recycler, formerly used by the Tevatron, has been modified as a proton stacker. Protons from the booster are injected into the recycler ring. Protons orbit in the recycler as the booster injects the requisite proton batches into the spill before being kicked into the main injector.

In the recycler, proton batches are slip-stacked to increase the intensity of each proton batch. Initially, six batches are inserted into the recycler from the booster frequency of 15 Hz. The first six proton batches from the booster are injected into the recycler ring and immediately and slightly slowed causing them to drift relative to their initial injection phase. The next six batches from the booster are then injected into the recycler. The RF cavities in the recycler are tuned to slightly different frequencies to allow the first six batches to continue regressing in phase while the subsequent six are maintained at a higher speed[53]. The frequency difference used to orbit the first and second set of six batches are tuned so that after insertion of the twelfth batch, the initial six batches have drifted in phase to lock up exactly with the second set of six batches. At this point, the beam is locked into a 6+6 slip-stacked configuration and kicked into the main injector. The slip-stacking process doubles the number of protons per batch, and thus spill, delivered to the main injector to achieve 5×10^{13} protons per spill. The design benchmark rate for NuMI spills is 4/3 s[54] which is often surpassed in practice.

The Main Injector, Acceleration and Injection into NuMI

Once the beam has been slip-stacked in the recycler ring, it is kicked into the main injector which accelerates the proton batches to 120 GeV over 0.5 s[44]. The six slip-stacked batches are then directed towards the NOvA far detector, NNW and 3° downwards, creating a spill for the NuMI beamline with a beam-spill timing window of 10 μ s.

An illustrative chart of the lifetime of protons in the Fermilab accelerator prior to insertion in the NuMI beam is shown in Figure 4.2.

The NuMI Beamline

Upon exiting the main injector, the proton spill reaches the NuMI target hall. The target itself consists of a series of 48 graphite fins, each 24 mm long with a small gap between consecutive fins for a total target length of 122.5 cm[55]. To combat

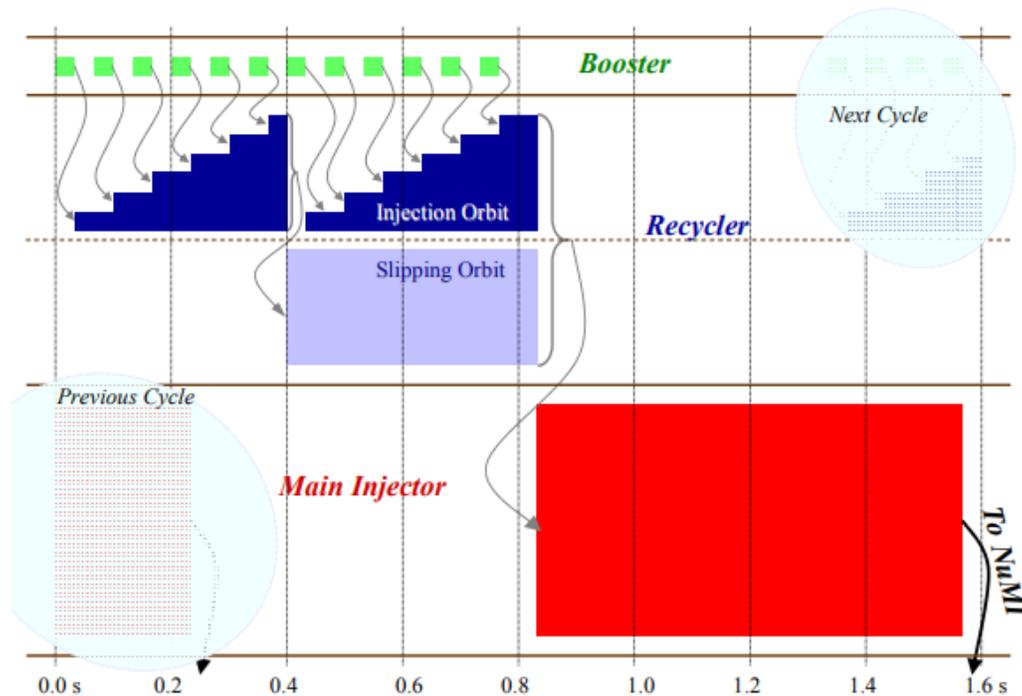


Figure 4.2: The life of a proton batch in the Fermilab accelerator before entering the NuMI beam. The top pane shows 12 proton batches in the booster kicked into the proton stacker, represented in the middle pane. In the stacker, the batches are slip-stacked to increase beam intensity. The stacker then transfers the slip-stacked protons into the main injector where they are accelerated to 120 GeV and sent to the NuMI beam.

the high intensity of the beam, the target is water cooled by two heat exchangers that run parallel with the target. In the target, protons interact with the carbon to produce a secondary meson beam, primarily composed of π^\pm . There is also a contamination of K^\pm and K_L , but π^\pm dominate the meson beam as the K production cross sections are lower. K mesons are further suppressed as, due to differences in mass, the K beam is more diffuse and contributes significantly less to the overall neutrino flux. The length of the target was limited to reduce the number of interactions mesons in the secondary beam have with the downstream fins of the target.

Two magnetic horns are placed downstream of the end of the target, which focus the secondary meson beam. Additionally, the horns deflect mesons of the wrong sign away from the beam and thus isolate mesons of like charge. The current through the horns is 200 kA which is far too much to run in steady state. The horns are therefore pulsed and water-cooled to avoid overheating.

The sign of the current passed through the horn gives NuMI its ability to produce a high purity beam of either neutrinos or anti-neutrinos. In forward horn current mode, positively charged mesons are focused, decaying to neutrinos, e.g. $\pi^+ \rightarrow \mu^+ \nu_\mu$. In reverse horn current mode negatively charged mesons are focused leading to an anti-neutrino beam.

Downstream of the horns is a 675 m vacuum decay pipe. In this volume, the NuMI beam finally appears as the secondary meson beam decays producing a tertiary neutrino beam.

The decay hall ends with a water-cooled hadron dump. Additionally, there are four muon monitors installed in the rock downstream of the hadron dump. These are simple ionization chambers that monitor the beam position and optical properties of the horns. A cartoon illustrating all stages of the NuMI beamline is shown in Figure 4.3.

Intrinsic Beam Backgrounds

Neutrinos are neutral and are not bent by magnetic fields. Thus, there are intrinsic backgrounds introduced by imperfect focusing of the parent meson beam. There is a flux of $\bar{\nu}_\mu$ in the neutrino beam and ν_μ in the anti-neutrino beam. This is referred to as the wrong-sign component. The magnetic horns effectively suppress the wrong-sign component in the beam by deflecting wrong-sign parent mesons, but those parents that leave the target very parallel to the beam are not deflected and contribute to the overall NuMI flux. In the 1-3 GeV energy region, relevant for neutrino oscillations, the wrong sign component is 1.7(11.2)% of the total far detector event rate for (anti-)neutrino running[56].

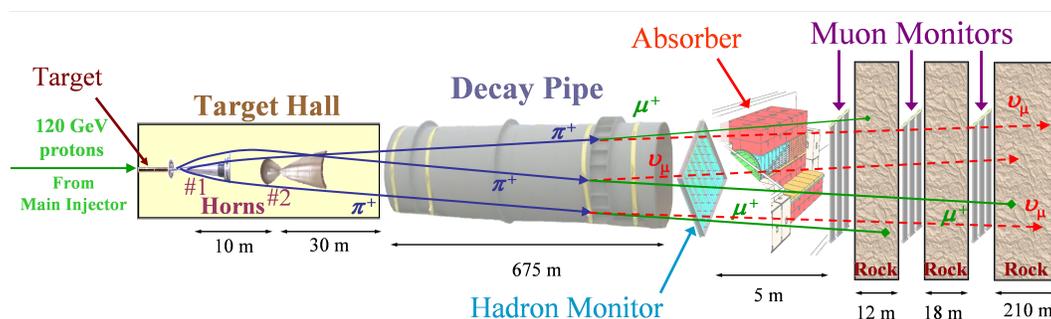


Figure 4.3: An overview of the NuMI beamline. A 120 GeV proton beam is incident on a target on the left producing a secondary meson beam. This secondary beam is focused by two magnetic horns and allowed to decay into the desired tertiary neutrino beam.

Additionally, there is an intrinsic beam- ν_e component from a $\nu_e/\bar{\nu}_e$ flux. In the 1-3 GeV range, this ν_e component primarily comes from the $\mu^+ \rightarrow \bar{\nu}_\mu e^+ \nu_e$ decay of the daughter muons from the $\pi^+ \rightarrow \mu^+ \nu_\mu$ process that makes the dominant ν_μ component. This component is 0.6(0.7)% of the total flux for (anti-)neutrino running[56].

These components are particularly troublesome for the $\nu_\mu \rightarrow \nu_e$ appearance analysis. Techniques have been developed to constrain these components of the beam with near detector data that will be discussed later.

Notable Improvements with the ANU

After the ANU, the beam intensity delivered to NuMI has increased by nearly a factor of two to 700 kW[44]. As the main injector can only accommodate six batches of protons, the only potential ways to increase intensity is to increase the number of protons delivered in each batch or to reduce the overall accelerator period between delivering spills to NuMI. There were gains on both of these fronts.

Full 6+6 slip-stacking the proton batches before inserting into the main injector increases the number of protons per batch. Meanwhile, parallelizing the injection of proton bunches into the main injector reduced the accelerator period from 2.2 to less than 1.33 s[52]. The lifetime of protons in the Fermilab accelerators is roughly 1.5 s. After 1.33s, 20 cycles of the 15 Hz booster, the booster begins injecting proton batches into the recycler to build the next NuMI spill. As each spill sent to NuMI takes 12 cycles from the booster, this gives 40% of the remaining protons available for other Fermilab beamlines including the BNB and muon campus.

4.3 Off-Axis Neutrino Beams

The current generation of long-baseline accelerator-based neutrino oscillation experiments, consisting of NOvA and T2K, place their detectors slightly off the center of the neutrino beam axis. For NOvA, the detector is 14 mrad off-axis. This has the effect of tightening the resulting neutrino energy spectrum around the characteristic energy for the oscillations. The effect comes directly from analyzing the Lorentz-boosted kinematics of π^\pm/K^\pm decay in the Center Of Momentum (COM) frame. A short overview of the kinematics is presented below.

Consider a $\pi^+ \rightarrow \mu^+ \nu_\mu$ decay. As this is a two body decay, in the COM frame, the neutrino energy is monochromatic and can be calculated by manipulating the particle four vectors. Taking the convention that all kinematic variables in the COM

frame are starred, this is

$$\mu = \pi - \nu \implies m_\mu^2 = m_\pi^2 - 2m_\pi E_\nu^* \implies E_\nu^* = \frac{m_\pi^2 - m_\mu^2}{2m_\pi}. \quad (4.1)$$

Numerically, this is 29.8 MeV. In the analogous $K^+ \rightarrow \mu\nu_\mu$ decay, E_ν^* is 237.6 MeV in the COM. To translate this to the lab frame, let θ^* denote the angle between the ν_μ and the boost direction in the COM frame. Then, the four-vector for the ν_μ in the lab frame is

$$\nu = (\gamma_\pi E_\nu^*(1 + \beta \cos \theta^*), \gamma_\pi E_\nu^*(\beta + \cos \theta^*), E_\nu^* \sin \theta^*, 0). \quad (4.2)$$

From this, the angle of the decay neutrino in the lab is determined from its angle of decay in the COM frame and γ_π as

$$\theta \approx \tan \theta = \frac{E_\nu^* \sin \theta^*}{\gamma_\pi E_\nu^*(\beta + \cos \theta^*)} \approx \frac{1}{\gamma_\pi} \frac{\sin \theta^*}{1 + \cos \theta^*}. \quad (4.3)$$

It then follows that

$$1 + (\gamma_\pi \theta)^2 = 1 + \frac{\sin^2 \theta^*}{(1 + \cos \theta^*)^2} = \frac{2 + 2 \cos \theta^*}{(1 + \cos \theta^*)^2} = \frac{2}{1 + \cos \theta^*}. \quad (4.4)$$

Applying this to Eqn. 4.2 and using $\gamma_\pi = E_\pi/m_\pi$, the neutrino energy can be rewritten as

$$E_\nu = \frac{2E_\nu^*}{m_\pi} \frac{E_\pi}{1 + (\gamma_\pi \theta)^2}. \quad (4.5)$$

There is a linear relationship between E_π and E_ν for on-axis neutrinos. Therefore, the on-axis neutrino distribution is as broad as the E_π distribution coming out of the magnetic horns. However, for non-zero values of θ , $E_\nu \sim 1/E_\pi$ for very large values of E_π . This implies the neutrino energy must turn over at some characteristic E_π and is therefore concentrated in a bounded interval.

The relation in Eqn. 4.5 is shown in Figure 4.4 for several choices of off-axis angle, including the 14 mrad scenario used for NOvA. To determine the flux of neutrinos, this curve must be convolved with the E_π distribution. But, as shown in the figure, nearly all of the E_π contributes neutrinos of energy relevant for neutrino oscillations compared to just the very lowest energy π^\pm in the on-axis case.

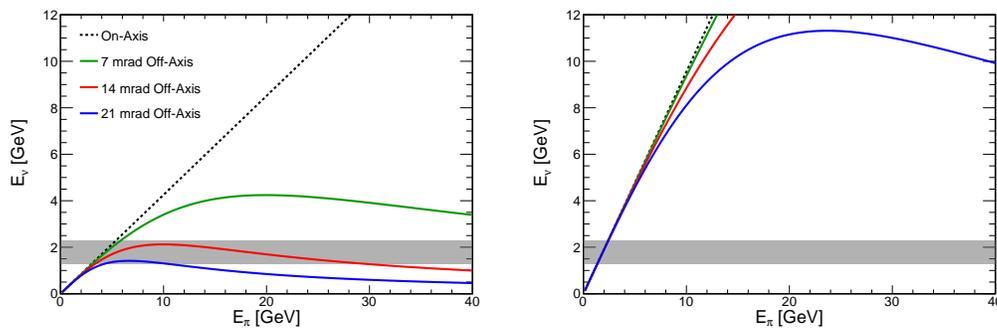


Figure 4.4: The neutrino energies for a variety of off-axis angles determined for $\pi^\pm(K^\pm)$ decay on the left(right). In each plot, the black shows the on-axis relationship which is linear. The green, red, and blue curves are the relations for a detector at different off-axis angles with red corresponding to NOvA. The gray shaded region shows the region where NOvA will search for oscillations. The red curve is nearly flat and within the gray region for a large portion of the available E_π range.

To accurately predict the NOvA flux, the π^\pm distribution coming out of the NuMI target must be known. The beamline can then be simulated to accommodate optical aberrations produced by the horn. Results of such a simulation are shown in Figure 4.5. As the simplified model suggests, at NOvA's 14 mrad off-axis angle, the beam is sharply peaked near the 2 GeV oscillation peak.

This peak sharpening is invaluable for reducing NC background in the oscillation analyses. As only the hadronic part of a NC event is visible in the detector, a selected NC background event may have a true visible energy in the oscillation peak even if the neutrino energy is quite high. This pileup of NC events at the oscillation maximum is mitigated by removing the large neutrino flux at high energies.

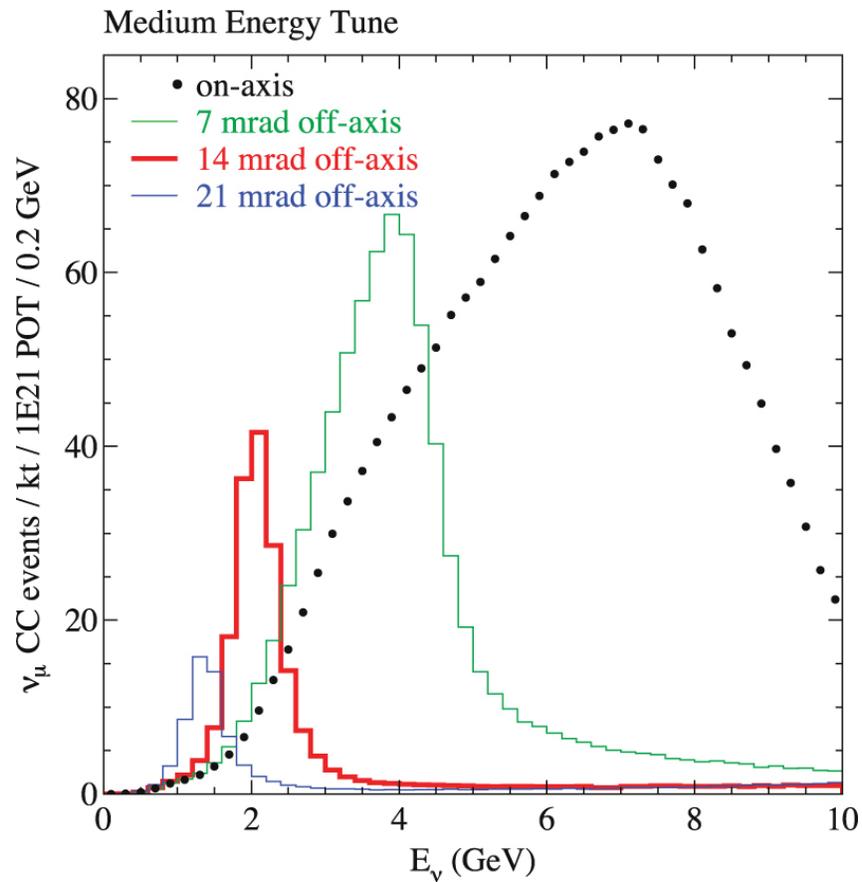


Figure 4.5: The simulated neutrino event rate at the far detector for various off-axis angles. The black distribution shows the on-axis spectrum while the red shows the prediction for the NOvA detector position. While the curve for NOvA has fewer total events than on-axis, there are significantly more near the 2 GeV oscillation peak.

Chapter 5

THE NOVA DATA ACQUISITION (DAQ) SYSTEM

As discussed in Section 3.1, the detector uses a silicon avalanche photodiode as the fundamental and singular detection element. The DAQ system refers to the process of detecting physics hits from the APD's, stamping them with appropriate metadata, and sending them to a temporary server where data that pass any trigger are pushed to disk.

5.1 Channel Readout

The channel readout is responsible for collecting and digitizing the raw voltage output from the APD's. Hits are identified from these APD traces by looking at sudden jumps in voltages. After detecting a hit above threshold, the readout is responsible for tagging each hit with a header containing channel and time information and outputting the hits to the next processes in the DAQ. All channels are read out continuously without any input from an external trigger. All hits are destined for a temporary server farm which will select hits that pass any triggers.

Front End Electronics

As described earlier, an APD amplifies photon signals from 32 individual cells in the detector. Each NOVA APD is connected to a Front End Board (FEB) which houses four specialized circuits. The voltage trace reported by the APD is sent to an ASIC custom-designed by NOVA. This ASIC's primary function is to shape the APD output so that a physics hit can be read out on a timescale comparable to the clock-tick. The ASIC output is then sent to an ADC which digitizes the signal and tags signal hits above threshold. Those hits that pass threshold are sent to a FPGA where they are time-stamped and sent forward through the DAQ. The configurable parameters, such as rise and fall-time in the ASIC and thresholds in the ADC, are programmed with the FPGA for every channel. The FEB's also hold a TEC regulator that measures the APD temperature and ensures it is kept at a constant -15 C, configurable with the FPGA, to maintain stable gain. A schematic of the FEB is shown in Figure 5.1.

The ASIC is shaped to a 460(140) ns rise-time and 7000(4500) ns fall-time for the far(near) detector. These time constants are programmable for each FEB and were

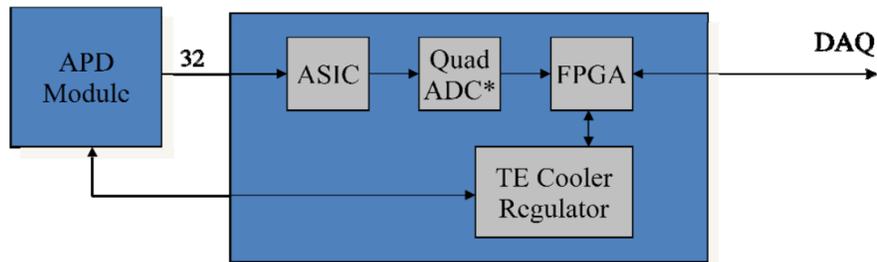


Figure 5.1: A schematic of the NOvA front end electronics. An APD, furthest left, is the basic detection element. The APD drives an ASIC which shapes signal photon pulses into a waveform that can be read on time-scales similar to the clock tick. The shaped pulse then travels through an ADC which digitizes the signal. That then passes through an FPGA, which compares the ADC to a threshold, deciding whether to pass the signal along to the DAQ. There is also a thermo-electric cooler which interacts with the APD, keeping it at a constant temperature for a stable gain.

chosen to easily detect signal hits over statistical noise. In the far detector, FEB's are outfitted with four ADC's which run at 16 MHz. The ADC's are mass produced and were not modified for the experiment. Each ADC digitizes the voltage from eight channels. The ADC uses 8:1 multiplexing so that it reads every channel out at 2 MHz, with 500 ns spacing, very similar to the programmed rise-time. In the near detector, the proximity to the beam causes high event pile-up in the $10 \mu\text{s}$ beam spill so that better timing resolution is required. In an effort to keep the far and near detector FEB's as similar as possible, this is accomplished by placing 16 ADC chips on each FEB, each using 2:1 multiplexing. This quadruples the readout rate to 8 MHz.

The ADC reads the ASIC's output for each channel continuously and uses Dual Correlated Sampling (DCS) to trigger hits from detector activity. The ADC reads a sequence of measurements in real-time which can be labeled $\{m_i\}$. The ADC then computes the value of $m_{i+3} - m_i$ and triggers if the difference is greater than the channel threshold. The DCS negates the effect of any gradual changes in the baseline input voltage and rather looks at a sudden increase in ADC value. As an example, three ADC traces for physics hits are shown in Figure 5.2.

Single-point and Multi-Point Readout

With the first data collected in the far detector, the DAQ employed a single-point readout scheme where, after a hit triggers with the DCS value of $m_{i+3} - m_i$ above threshold, the DCS records this DCS value and the time of m_i to describe the hit. In Oct 2014, multipoint readout was adopted in the far detector. The near detector

began recording data after this and always operated in multipoint mode. In this readout scheme, the DCS sends forward the four ADC values from m_i to m_{i+3} [57]. These four points are fit offline to infer an initial hit time and a pulse height. With singlepoint timing, the resolution, $500/\sqrt{12}$ ns, is driven by the sample rate with a geometric factor for drawing uniformly across this time window. Propagating these four values increases the precision of the pulse timing and height by fitting these four values to the expected pulse shape. Two simulated ADC traces are shown in Figure 5.2 which show the curve that was fit through the four points recorded in the multipoint scheme. The improved timing resolution with multipoint timing is shown in Figure 5.3.

The switch to multipoint has a minimal effect on raw data size. Due to the header and timestamp, this switch only increases the bandwidth required for passing hits from the DAQ onward through the DAQ by one third. The structure of the memory required for each hit in both single and multipoint readout is shown in Figure 5.4.

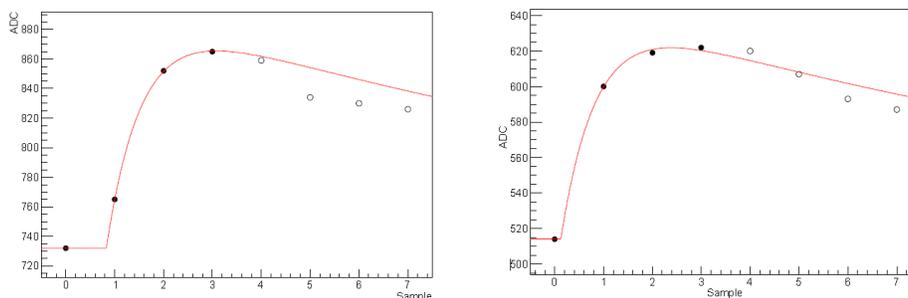


Figure 5.2: Two time traces for simulated physics hits. The solid black points give the four data points read out in the multipoint DAQ scheme. Singlepoint readout would only save the DCS value of $m_{i+3} - m_i$. The red curve is fit through the four saved points offline which gives a significantly improved resolution on the time and height of the triggered hit.

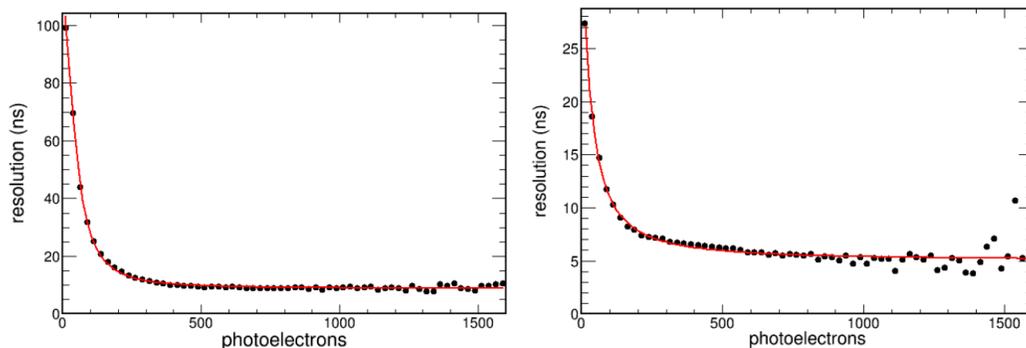


Figure 5.3: The timing resolutions, in the multipoint readout scheme, in the far detector, left, and near detector, right. The resolution depends on the photoelectrons deposited in the hit, x -axis, as low energy hits are more likely to pass threshold on a delayed sample.



Figure 5.4: The anatomy of a single hit, as prepared by the FPGA on the FEB for single(multi)point timing on the left(right). Due to the overhead of the header and timestamp to specify the channel and time of the hit, including the three additional ADC values for offline fitting only increases the size of the hit by a third.

5.2 Collecting Hits from All Cells and Calibrating Timing

There are 344064 channels in the far detector which need to be collected and sent to the buffer farm to be made available to detector triggers. Each FEB digitizes and triggers hits from 32 channels. An intermediary computer called a Data Concentrator Module (DCM) collects output from 64 individual channels[44]. NOvA's far detector holds 168 DCM's with 12 DCM's servicing each of the 14 diblocks. There are six DCM's on the top of the detector which service the vertical modules and six on the west side of the detector which service the horizontal modules. A picture of the top of the detector with shows these six DCM's per diblock on the top-view of the detector in Figure 5.5. Each DCM collects data from all FEB's it services into 50 μ s slices called microslices[58]. Rather than transfer each microslices to the temporary server, or buffer farm, immediately, the DCM's transfer sets of 100 microslices, covering 5 ms of data, at a time. After collecting a millislice of data, the data is transferred to the buffer farm. Smaller amounts of data are not transferred

as the accumulated data is too small to warrant a ping to the buffer farm.

Once transferred to the buffer farm, microslices from all DCM's are sorted and grouped by the time of the microslice. A collection of the microslices from all DCM's in the detector for one particular $50 \mu\text{s}$ block of time is called a microblock and is the smallest piece of data the NOvA detectors can permanently record. These microblocks are stored on the buffer farm with a lifetime of about 30 minutes. Those that pass a trigger are copied from the buffer node to a permanent storage disk.

A timing chain was developed to synchronize all DCM's, both detectors, and the accelerator. Each detector has a Master Timing Distribution Unit (MTDU) which reads the time from a GPS antenna and allowing for easy syncing between both detectors.

For the far detector, a MTDU drives 14 slave TDU's (STDU)'s each responsible for synchronizing all DCM's on one diblock[59]. The MTDU is connected to the first STDU and subsequent STDU's are connected to the last in the chain. This collection of MTDU's and 14 STDU's is referred to as the timing backbone. From the backbone, each STDU is connected to each of the 12 DCM's it synchronizes.

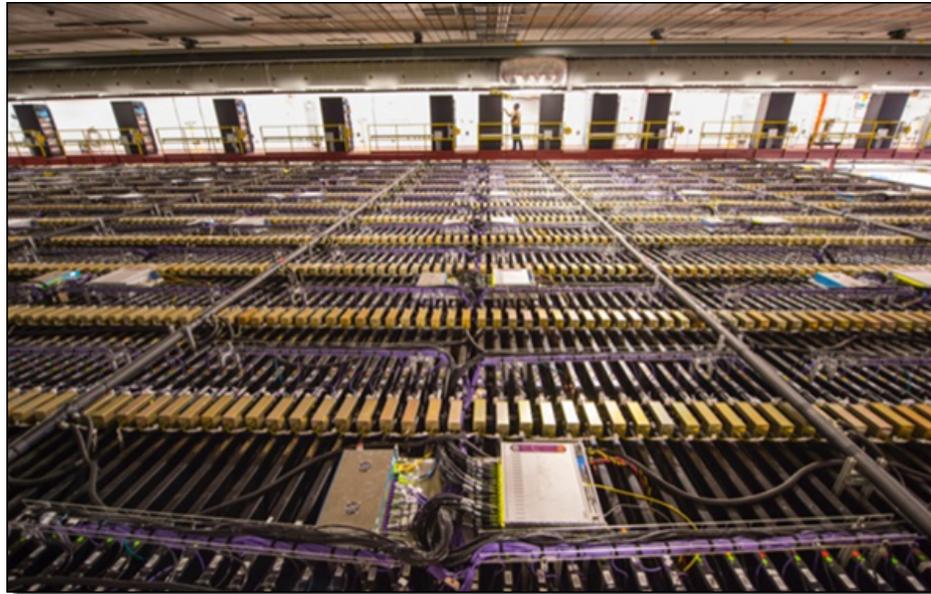


Figure 5.5: A view of the top of the far detector, perpendicular to the beam. At the bottom of the picture there is DCM on the left paired with a power distribution unit, driving the DCM and all FEB's and APD's it services, on the right. There are six pairs along the center of the picture showing the six DCM's that service the vertical modules in this diblock.

These are separated into two groups of six daisy-chained DCM's. All DCM's in each group are mounted on the top or side of the detector. Together with the backbone, this gives the detector's timing chain. As a fail-safe, a redundant set of MTDU and 14 STDU's are also connected to each DCM. A schematic of the timing chain is shown in Figure 5.6. A timing calibration is applied to each DCM to correct for the different cable lengths required to push a synchronized time to each channel.

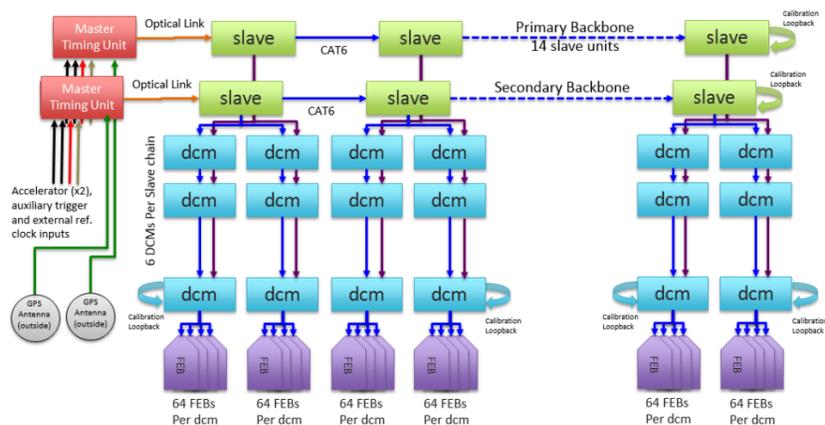


Figure 5.6: A schematic of the DCM and timing chain layout in the detector.

The near detector is very similar to the far detector both in terms of the data collecting and timing calibration but with fewer DCM's and STDU's.

5.3 Event Triggering

Data is stored in microblocks in the buffer farm for thirty minutes. While in the buffer farm, a trigger may select a contiguous group of microblocks and write these to permanent disk storage. The cosmic rate in the far detector is high, roughly 100 kHz, and the neutrino interaction rate is low with only a couple per day expected. An activity-based neutrino trigger is impractical, but there is a more accurate and easier solution. Therefore, the NuMI beam spills are triggered by a timestamp sent from the Fermilab accelerator with a time correction for the neutrinos' time of flight. Whenever a spill signal is received by the detector, the buffer farm selects 1000 contiguous microblocks, or $500 \mu\text{s}$ of continuous data for the entire detector and writes the data to disk. The timing is calibrated so that the $10 \mu\text{s}$ beam spill is in the 218-228 μs range within this recorded data. The near detector is underground and much closer to the beam, so a beam trigger based on detector activity is possible, but the same trigger applied to the far detector data is used based on the accelerator's signal for consistency in exposure counting.

The detector also writes data from a cosmic trigger which selects $500 \mu\text{s}$ slices of data at regular intervals. The primary purpose of this data is to calibrate the detector with cosmic muons. The detectors also take data from a number of triggers based on detector activity such as finding Michel electrons[60] from stopped muon tracks or searching for exotic particles like magnetic monopoles[47]. But these triggers are not used for any beam physics and are thus outside the focus of this thesis.

5.4 Measuring Channel Thresholds

The FEB FPGA allows for a separate programmed threshold ADC value for every channel. In an effort to reduce the statistical noise from each channel while maintaining triggering efficiency for physics hits, the statistical noise for each channel is measured in situ with the current detector configuration. To do this, every channel's ADC value is read out continuously for over several thousand contiguous readings. This results in an ADC trace vs time which gives a direct measure of the statistical variation observed in each channel. A collection of ADC traces from all 32 channels of a particular FEB in such a scan is shown in Figure 5.7.

After the trace is taken, the DCS value $m_{i+3} - m_i$ for a given channel is calculated for each clock tick and is plotted. As there are 344064 channels in the far detector,

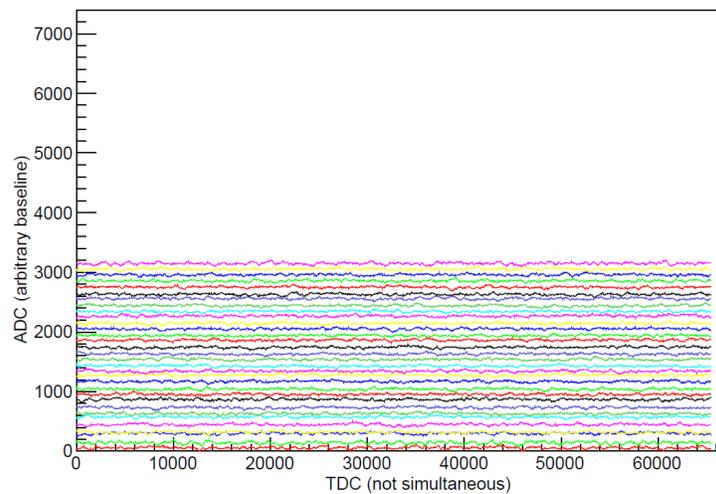


Figure 5.7: Example data from a threshold scan from an FEB at the far detector. Each curve gives a continuous ADC reading for a given channel as a function of time. A threshold is set for each channel by analyzing the spread in the DCS value of $m_{i+3} - m_i$. It's notable that, though the baseline may change appreciably over several hundred ticks, the DCS value only looks at the difference in nearby hits, making the triggering insensitive to this drifting. The curves are separated by an arbitrary offset for ease of viewing.

the threshold is set aggressively so that the noise triggering rate is 100 Hz. With this benchmark, only a few hundred noise hits are expected across the entire detector, in the 10 μ s beam spill.

An example distribution of the DCS values for a given channel during a threshold scan is shown in Figure 5.8. A Gaussian parameterization fits the model well for ≈ 3 orders of magnitude, but underestimates the noise rate at very high DCS values. A procedure was developed to overestimate the noise trigger rate based solely on hits within the Gaussian portion of the distribution. An exponential is fit in the DCS range where the distribution is 1-2 orders of magnitude down from the peak[61]. This procedure was developed to mitigate relying on the non-Gaussian tails which are difficult to understand and correctly simulate. After evaluating a threshold for every channel, the thresholds are then pushed to each FEB where the FPGA appropriately sets the thresholds.

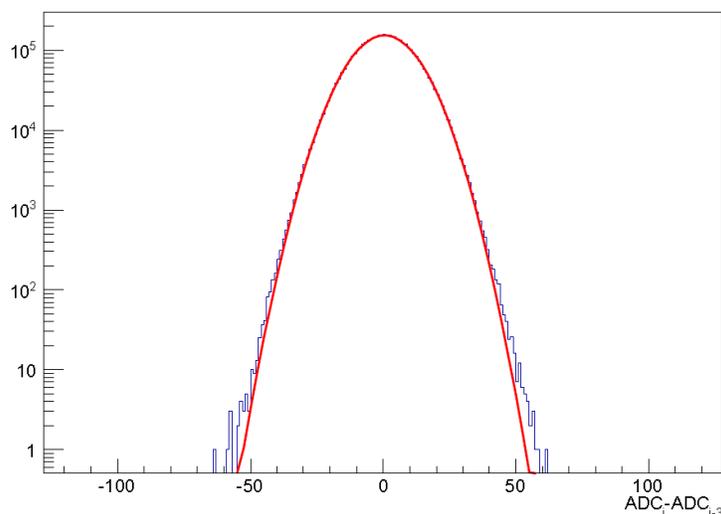


Figure 5.8: An example distribution for the DCS values acquired during a DSO scan for a particular channel. The distribution is fit to a Gaussian, a parabola in the log-plot, which fits the data very well for three orders of magnitude. The thresholds are set based on this Gaussian part of the distribution so that the noise trigger rates in the detector are acceptably low.

The thresholds can not be measured while the detector is reading out hits because the ADC's read out a very long trace for each cell. The process is bandwidth limited so that additional hits may not be recorded and sent to a DCM during the procedure, which takes about thirty minutes to complete at the far detector. But, the thresholds

are re-calculated every time the detector configuration is changed as the downtime is a reasonable investment for insuring the rates of noise hits remain acceptably low.

5.5 APD Sag

A notable non-linearity in the FEB readout was observed with test stand measurements[62], ahead of the first detector data collected. In the measurement, an LED illuminated a single pixel of a NOvA APD, and the voltage of each channel was read out by an oscilloscope. The results, after reducing statistical noise by averaging multiple measurements, is shown in Figure 5.9. For channels not lit by the LED, a characteristic dip of the same shape as the LED signal is clear. The amplitude of the dip is 1.86(3)% of the total charge deposited in the APD. This number is reproducible with different channels on the APD, different tested APD's, and in the case where multiple channels are triggered[63]. Sag has been integrated into the detector simulation and was shown to have minimal effect on physics measurements.

This effect was later explained as a manifestation of charge depletion in the APD circuit. The NOvA APD's are built with a bypass capacitance of 100 pF. Adding an additional capacitor of 10 nF connecting the voltage rail to ground would supply enough available charge to effectively eliminate the issue[64]. This was not identified

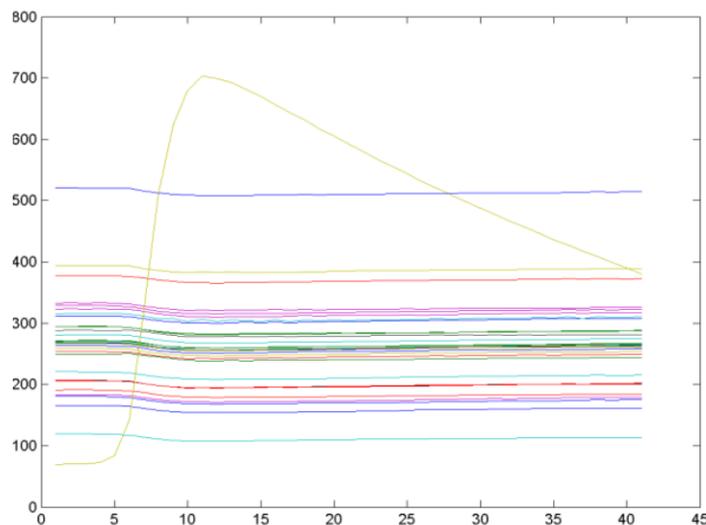


Figure 5.9: The voltage off the 32 channels of an APD after a particular channel, the yellow trace which ticks upwards at 5, is periodically lit by an LED. The x-axis shows the time in clock-ticks while the y-axis gives the average voltage reading off the channel. All other channels in the APD dip down in phase with the lit APD's trigger. The amplitude of this dip is 1.86(3)% of the total charge deposited on the APD.

until after both detectors were fully commissioned and, as the cause was known, and the amount of sag was shown to be the same for all APDs, there was no effort to retrofit the detector.

The sag phenomenon is directly related to FEB flasher events observed in both detectors. In the case of a very high energy deposit across the channels read by a single APD, in excess of 20 MIP-equivalents, the dip induced in those channels that were not hit is deep enough that they can trigger themselves as they return to their baseline voltage. An FEB flasher lights up the APD channels not hit by the initial particle at a characteristic few μs after the initial particles were recorded. An example of such a flasher in a cosmic brem event is shown in Figure 5.10. As these events are naturally predicted by sag, FEB flashers are incorporated into simulation.

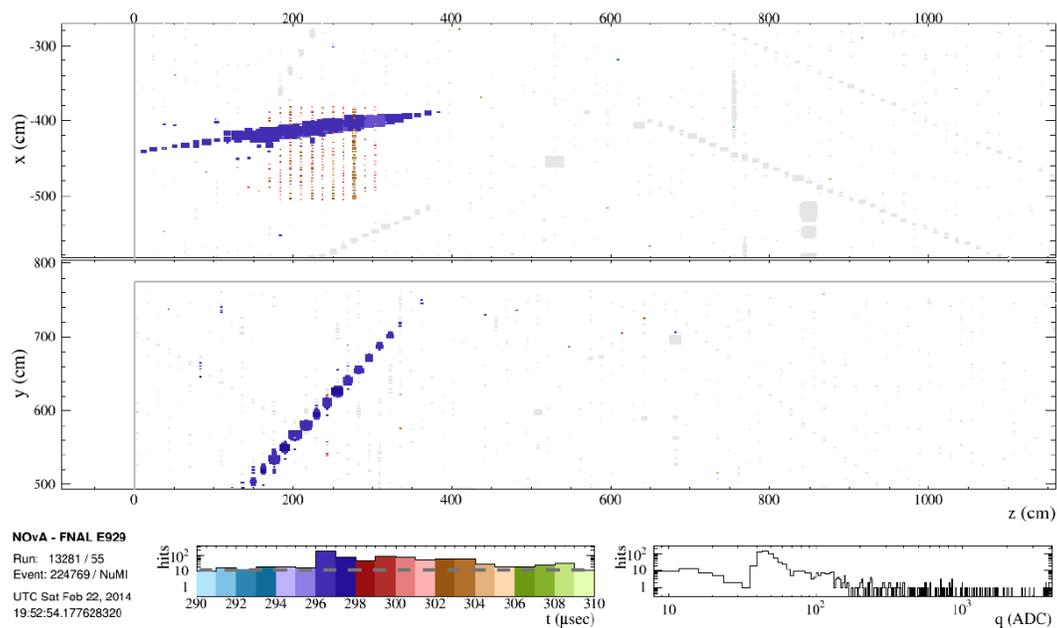


Figure 5.10: A high-energy cosmic brem event leaves a very large amount energy. In each plane, this energy is spread across a few channels recorded by the same APD which induces enough sag that the remaining hits in these FEB's trigger a few μs later. The colors of the hits show the recorded time of each hit.

Chapter 6

SIMULATION OF THE NOVA EXPERIMENT

To draw conclusions based on data from the NOvA detectors, a sophisticated and realistic simulation of the detectors has been developed. NOvA has developed a description of detector geometry, light production in response to tracked particles, and the optical model determining the light delivered to each APD. Using event generator and particle tracking software developed by outside collaborations, the NOvA simulation provides raw predictions for event rates and topologies in the NOvA detectors.

Event simulation is innately imperfect, and so a thorough understanding of the current status and limitations of the simulation is important when considering how systematic effects may affect analyses. Efforts continue after several years of work to incorporate improvements to the detector model in an effort to reduce systematic uncertainties. As described in later chapters, physics analyses have developed error bands on simulated error predictions based on the limitations of the software described below, often constrained with in-situ and external measurements.

6.1 Detector Simulation

The detector geometry is simulated using the Geometry Description Markup Language (GDML)[65] developed at CERN in conjunction with ROOT. The detector description includes a detailed description of the detector halls and a realistic description of the detector shape, from modeling of a single cell to large-scale assembly of the detectors. An extensive comparison of the simulated geometry and that detailed by the project has produced a robust systematic error stemming from geometry mis-modeling[66].

Within individual cells, the photon transport system is simulated, relating particle energy depositions to photon signals seen by the APD. This takes into account the light yield spectrum from the scintillator, reflectivity of TiO₂-laced PVC cells, and the absorption spectrum of the WLS fiber. As the WLS fiber is arranged in a loop, these re-emitted photons on the fiber have a chance of going each way through the fiber, depending on the photon's angle of incidence on the fiber. In the far detector, brightness of fibers drifted $\approx 5\%$ systematically with time during production. Before

outfitting, each fiber's brightness was measured and cataloged in a database so that the simulation can incorporate these differences.

The readout simulates the quantum efficiency for the APD's at a flat 85% with a 5% variation in APD gain. The signal from APD's is shaped by the rise and falltimes programed into the FEB ASIC. This is sampled at 2(8) MHz for the far(near) detector as is done in the physical detector and the DCS value $m_{i+3} - m_i$. If the DCS value is above threshold, a hit is stored. The thresholds for each channel are stored for data for each subrun so that the MC can use appropriate thresholds from this database.

Additionally, as both the cosmic rate and statistical noise trigger rate for hits is high in the far detector, a few 100 Hz during the NuMI spill window, simulation in the far detector is overlaid on data taken from the cosmic trigger. This supplies real conditions for both noise hits and cosmic rays passing through the far detector without potential beam contamination. Simulating beam events over this base noise gives a precise measure of how excess event activity will impact signal efficiency without calling on any specific model.

6.2 NuMI Beam Simulation

The NuMI beam simulation for the NOvA detectors is done using the G4NuMI package, a specific adaptation of Geant4[67] for application in the NuMI beamline. The beam line simulation includes a detailed description of proton-nucleus, subsequent meson-nucleus interactions in the target, and meson focusing by the magnetic field generated by the horns.

The mesons are then passed through a simulation of the decay pipe where they decay. The flux of neutrinos at each detector is then given by the neutrino spectrum passing through the solid angle occupied by the NOvA detectors. The near detector, whose face has a cross section of 16 m² from a baseline of ≈ 1 km, spans a few mrad² of solid angle while the far detector can be treated as a point detector at 14 mrad off-axis. The flux of neutrinos through the far detector is shown in Figure 6.1.

Package to Predict the FluX (PPFX)

The MINERvA collaboration, in an effort precisely measure the NuMI flux for cross section studies, has developed the PPFX framework[68] to predict the flux through detectors recording data from the NuMI beamline. This analysis constrains the interaction modeling for various ancestor hadrons in the NuMI beam with an ambitious survey of hadron production channels from MIPP[69], NA49[70], and

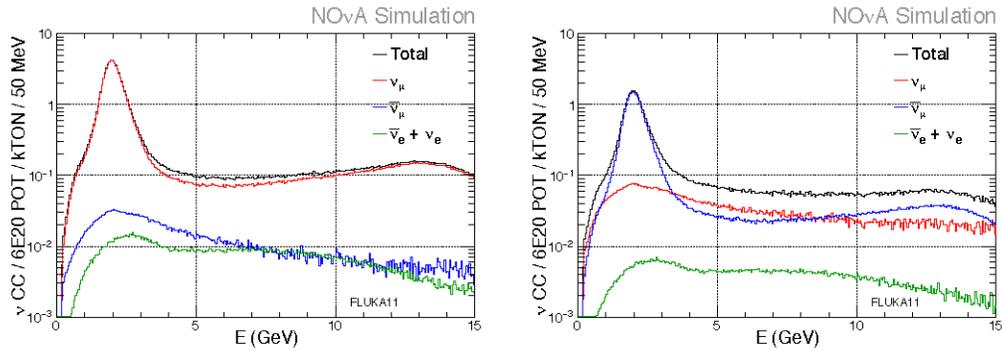


Figure 6.1: The far detector ν_μ CC predicted at the far detector, with no oscillations applied. The left(right) shows the distributions in (anti-)neutrino running. Red and blue give the histograms for ν_μ and $\bar{\nu}_\mu$ while green is the sum of $\nu_e + \bar{\nu}_e$ events.

NA61[71] datasets. Using these datasets, the cross sections for various hadron-carbon interactions possible in the target are constrained in the phase space most relevant for neutrino production in the NuMI beam. Figure 6.2 shows the adjustment applied to the G4NuMI simulation for π^+ scattering on carbon as a function of the incident π^+ transverse momentum and Feynman scaling variable.

The resulting error band on hadron production for various channels after the PPFX

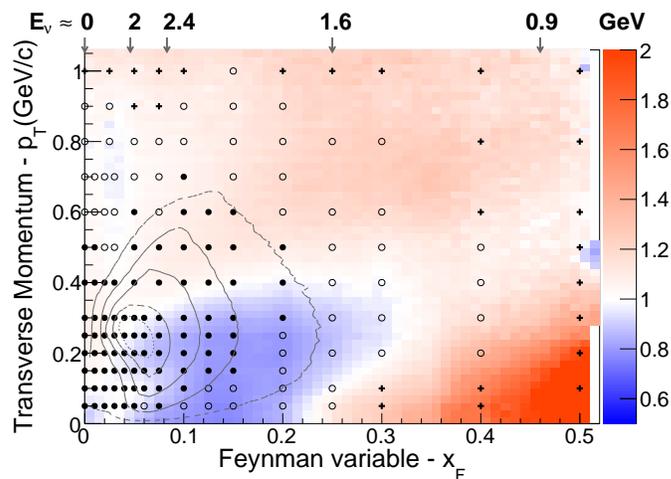


Figure 6.2: An overview of the NA49 constraint on $p + {}^{12}\text{C} \rightarrow \pi^+ + X$ production from MINERvA's PPFX analysis. The axes give transverse momentum vs x_F , which is a proxy for the decay neutrino's energy. Color shows the normalization correction applied to the simulation. The contours represent the typical phase space encountered in the NOvA beam. The marker gives the size of the statistical error on the normalization: filled dots have an error less than 2.5%, open dots have an error between 2.5 and 5%, and crosses have a statistical error greater than 5%.

constraints is shown in Figure 6.3. Uncertainties on hadron production drive the systematic uncertainties on NOvA analyses arising from flux modeling. The PPFX data-driven constraint reduces these uncertainties, which is particularly helpful for cross section analyses in the near detector for which the flux directly influences the measurement.

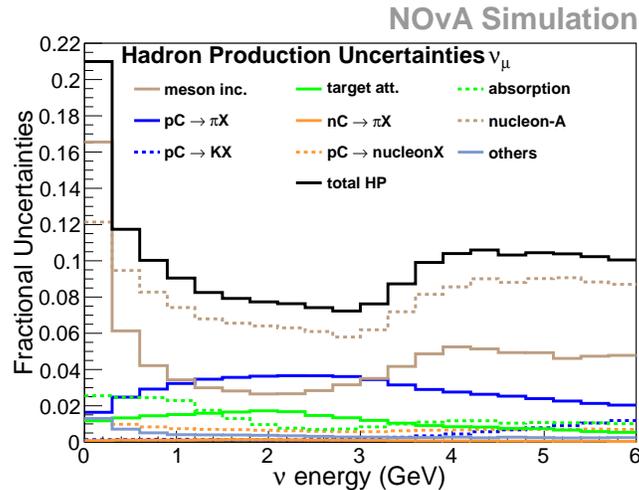


Figure 6.3: The fractional uncertainty for various hadron production channels necessary for understanding our target.

6.3 Event Generation and Particle Propagation

NOvA analyses use the GENIE neutrino event generator[72] to predict final state topologies for neutrino interactions from the beam in the NOvA detector. The Cosmic RaY (CRY) generator is used to simulate the decay chain of cosmic ray particles through the detector. CRY is principally used for the calibration of the detectors. Each of these software packages is described in more detail below. The GEANT4 software package[67] is used to propagate these final state particles from both the GENIE and CRY event generators through the detectors.

GENIE

The GENIE event generator[72] produces events described by final state topology and kinematics based on an input neutrino flux and target geometry, for an accounting of composite materials. Individual nuclei are simulated as a Relativistic Fermi Gas (RFG) of composite, non-interacting nucleons. The cross sections off RFG nucleons are amended to account for correlations between nucleons. Events are first thrown at a specified energy, given the flux and target material. Once a neutrino of a given energy is determined to have scattered, event kinematics are determined

using differential cross section in $d^2\sigma/dQ^2dW$ for different phenomenological interaction types as discussed in more detail below. The initial neutrino interaction is then propagated out of the nucleus using the INTRANUKE software to re-scatter generated hadrons. The GENIE final particle list is given as the final state that emerges from the nucleus.

As described in Appendix A, the 1-5 GeV range is a transition region for neutrino-nucleus scattering. Within this region, ν_μ CC scattering can be, at its simplest, described as a mix of QE, MEC, Res, DIS, and Coh processes. Genie generates interactions from each of these phenomenological interaction types in a way that both maintains the inclusive ν_μ CC cross section while predicting hadronic state multiplicities according to global fits with exclusive hadron production data[72], such as single- π producing interactions.

QE events are easily described by a set of hadronic form factors[73] which, through applications of CVC and PCAC, is reduced to a single degree of freedom determined by neutrino scattering measurements in bubble chambers, m_A^{QE} . Coherent events are also relatively simple to model. Similar to QE events, they lie in a low Q^2 range where expansion around the PCAC symmetry predicts a single dipole form factor.

Resonant and DIS interactions are much more difficult to separate as they share similar final state topologies and both lie in a Q^2 and W region much higher than populated by the QE process. The production mechanism, however, is radically different as resonant interactions cover processes such as $\nu + p \rightarrow l^- + \Delta^{++} \rightarrow l^- + p + \pi^+$ and DIS interactions are described by $\nu + q \rightarrow l^- + q'$. Genie includes 16 individual baryon resonances in its cross section model. For $W < 1.7$ GeV, the total DIS contribution to the cross section is scaled down so that GENIE's predictions for both the inclusive cross section and various exclusive hadronic final states agree with measurements[72].

MEC events, dominated by the 2p2h sub-process, are also notoriously difficult to model as they involve calculating correlations of nucleon wavefunctions in the nucleus, considerably more complicated than the RFG simulation of the nucleus GENIE employs. GENIE therefore uses an empirical MEC model used to explain discrepancies seen with data near 1 GeV in the Q^2 region between QE and resonant events.

After determining interaction type, Q^2 , and W from the differential cross sections, particles from the hadronic current are translated into final state particles using the

a hadronization models. For $W < 2.3(> 3.0)$ GeV, the cascade is calculated using the AGKY(PYTHIA) model. In the transition region, a weighted average is used. In generating final state hadronic particles, the average expected number of hadrons of each flavor is calculated given the event kinematics, and number for each particle type is drawn subject to overall charge conservation. Kinematics of these final state hadrons are then determined by the p_T and x_F of the parent nucleon and past measurements. After this process has finished, a hadronic final state, in conjunction with the lepton kinematics determined from Q^2 and W , gives the final state particles for the event whose four-momenta are passed to Geant4 for propagation through the detector.

In addition to predicting cross sections for all constituent interaction modes, GENIE also predicts error bands for the various parameters involved in the models according to global fits. These affect both the shape and normalization of the outgoing neutrino interactions. These error bands will be used to calculate systematic uncertainties on various physics analysis.

Though the GENIE collaboration continues to incorporate new data into its cross section fits, there is always a latency period before GENIE can incorporate new measurements into the cross section model. Therefore, the NOvA oscillation analyses have adopted a slightly modified tune to GENIE in order to incorporate recent measurements and improve the GENIE predictions for neutrino interactions in the detectors. Additionally, this procedure amended systematic error bands to more appropriately reflect our current understanding of neutrino scattering in the 1-5 GeV energy range. This will be discussed in detail in section 13.2.

CRY

The CRY package is used to simulated cosmic rays through the detector. A flux window is defined above the detector, parameterized by latitude and elevation above sea level, which contains information about primary particle type and initial kinematics. Primaries are limited to energies between 1-1000 GeV. These particles are allowed to decay en-route to the detector. All secondary particles with energy greater than 1 MeV are tracked to the detector. Once the simulated cosmic rays have reached the top of the detector, the resulting distribution of cosmic particles is passed to Geant4 for tracking through the detector.

Geant4

Geant4 is a particle tracking toolkit that, given a detector geometry, set of initial particle four-momenta, and a configurable list of physics processes to consider, will propagate and decay the particle tree through the detector until only stable, stationary particles remain in the detector. The NOvA simulation uses the QGSP_BERT_HP physics list[74]. This is a recommended physics list for high energy physics interactions at the few GeV scale. It simulates particle decay, hadronic interactions, and electromagnetic processes such as energy loss through ionization.

For tracking hadronic particles, a Bertini cascade[75] is used to simulate hadronic interactions in the detector. In this model, nuclei are treated as a nucleon gas and hadron-nucleus cross sections are summed from hadron-nucleus cross section measurements. This approximation works well for energies below 10 GeV and above k_{fc} , approximately 250 MeV for ^{12}C . Hadrons produced in relevant neutrino interactions fall comfortably below this upper limit, while in the lower limit, ionization energy loss dominates.

Geant4 calculates trajectories by stepping particle tracks through the detector. To determine the path length, the mean free path for each physics process is calculated given the local density. The path length is taken as the lowest mean free path for all constituent physics processes. Since the mean free path often depends on the local density, a new step must also be generated whenever a particle crosses a material boundary in the geometry. For NOvA, the later almost always sets the step length, as particles are constantly passing through planes subdivided into scintillator and PVC.

There are several standard physics libraries for users to load. NOvA has tested a few of these and shown that the choice of physics library has a sub-percent effect on physics analyses, negligible compared to other systematic uncertainties.

6.4 Simulation of Čerenkov Light

Čerenkov radiation is released when a particle travels faster than the speed of light through a medium[76]. It was subsequently described theoretically with the energy spectrum released per unit of track length given by the Frank-Tamm formula[77]:

$$\frac{d^2N_\gamma}{dx d\lambda} = \frac{2\pi\alpha z^2}{\lambda^2} \left(1 - \frac{1}{\beta^2 n^2(\lambda)}\right), \quad (6.1)$$

in terms of the particle speed, β , wavelength, λ , and index of refraction, $n(\lambda)$. If

the index of refraction varies only slowly, this formula behaves like $\sim 1/\lambda^2$ so that most of the radiation is emitted well below the region sensitive to the scintillator and electronics. But, for wavelengths between 300-400 nm, the scintillation light can be absorbed and re-emitted by the WLS fluors and absorbed directly by the WLS fiber in the 400-500 nm band. This makes Čerenkov light indistinguishable from scintillation light in this region and enhances our light yields for fast particles. The photon yield for various values of β is shown as a function of wavelength in Figure 6.4.

Čerenkov light is included in the photon transport simulation of the DAQ. A bench-top measurement determined $n(\lambda) = 1.47$ for wavelengths relevant for NOvA in mineral oil. With this parameter, Čerenkov light accounts for roughly 4% of the light collected by fast particles as shown in Figure 6.5.

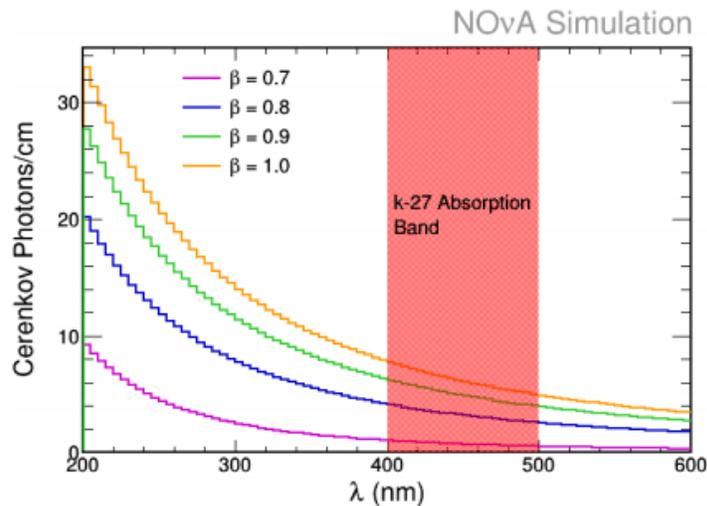


Figure 6.4: The number of Čerenkov photons/cm released by a particle of various speeds for different wavelengths. The red shaded region shows the absorption spectrum of the k-27 dye in the optical fiber. Additionally, light in the 300-400 nm range is absorbed by the WLS fluors in the scintillator mix and re-absorbed in the fiber absorption band.

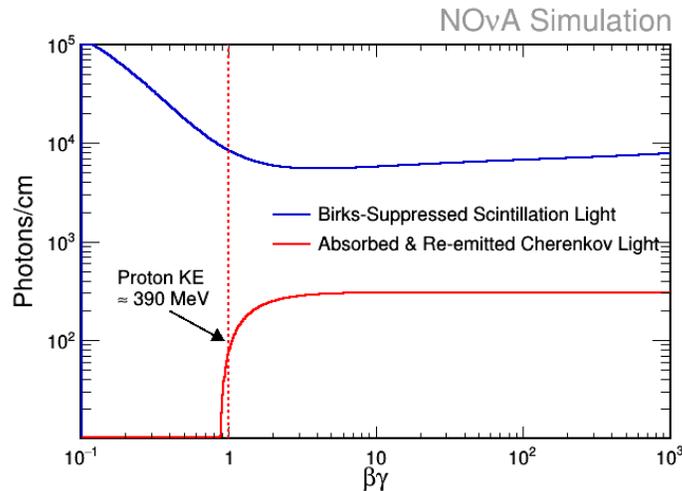


Figure 6.5: A comparison of light yield, in photons/cm, from scintillation, blue, and Čerenkov radiation, red as a function of $\beta\gamma$. For fast particles, Čerenkov supplements the scintillation light by $\approx 4\%$.

6.5 Scintillator Nonlinearity and a Birks Parameterization

There is a saturation effect for scintillators where, if particle dE/dx is high enough, there is not enough local available charge to fluoresce to keep the relationship between light yield and dE/dx linear. An empirical description of this effect[78][79] is described by

$$LY = A \frac{1}{1 + k_B \frac{dE}{dx} + k_C \left(\frac{dE}{dx}\right)^2} \frac{dE}{dx}, \quad (6.2)$$

which is implemented in our simulation. Both k_B and k_C were tuned using tagged protons in the near detector. Protons were chosen since they are massive and exhibit high and varying dE/dx while tracking, giving sensitivity to estimate these parameters.

A MC was generated for various values of the Birks-Chou parameters. In data and each MC sample, protons were selected by tagging quasi-elastic ν_μ CC events with a proton track length of at least 50 cm and a cut on activity off the muon and proton tracks. For each MC sample, the dE/dx distributions for selected protons in the final six planes were compared to those selected in data. A χ^2/dof between data and MC was calculated for each set. A χ^2/dof map was constructed using all of the MC's generated. There is a clear minimum at $k_B = 0.01 \text{ gcm}^2/\text{MeV}$ and $k_C = 0$ where $\chi^2/dof \approx 0.9$. This χ^2/dof map is shown in Figure 6.6.

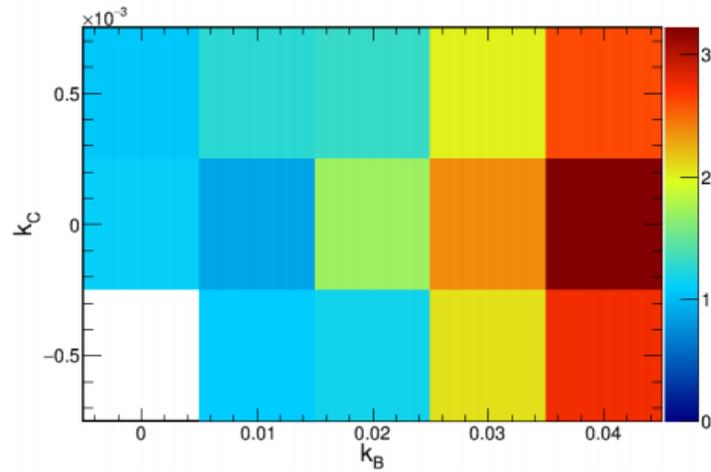


Figure 6.6: A map of χ^2/dof of proton dE/dx between MC sets generated at different Birks-Chou parameters and the observed data. There is a minimum at $k_B = 0.01 \text{ gcm}^2/\text{MeV}$ and $k_C = 0$ which is applied to the nominal NOvA simulation. There is no entry in the histogram at $k_B = 0$ and $k_C = -0.0005$, as these parameters would always give a Birks “enhancement”, which is unphysical.

Chapter 7

CALIBRATION OF THE NOVA DETECTORS

The NOvA DAQ offers a prescription for recording hits above some threshold. But before using the NOvA detectors to extract physics, they must first be calibrated so that neutrino energies can be reliably reconstructed and energy deposition patterns can be used for event classification.

NOvA uses cosmic muons to perform this calibration. For most of the length of a muon track, energy deposition is a MIP, whose energy deposit per unit of tracklength can be used as a standard candle. At high muon energies, tracks can deposit more than a MIP-worth of energy due to the relativistic rise in the Bethe-Bloch equation, which is simulated by GEANT4 in the calibration procedure. The rate of cosmic particles through the far detector is ~ 100 kHz, so there is more than sufficient data to calibrate in this way.

The calibration procedure involves two steps. The first involves correcting for various geometric effects which influence the measured energy deposit for a MIP. Afterwards, an absolute calibration is applied so that this energy measure consistent throughout the detector is related to a human-interpretable value, GeV. Data and MC in both detectors are calibrated in parallel using the same methods. Additionally, a number of other standard candles, such as the Michel energy distribution and reconstructed π^0 mass, are compared in data and MC to assess the systematic uncertainty involved in detector calibration. These comparisons are discussed in Section 9.9.

7.1 Attenuation Calibration

NOvA cells are quite long, 15.7 m in the far detector, while the WLS fiber bridging between the liquid scintillator and the APD has an attenuation length of a few meters. So, photoelectrons collected by the fiber from muons that pass through the far end of a cell will be naturally attenuated by the fiber length. To correct for this effect, an attenuation calibration procedure is applied translating ADC values into a new unit, PECorr, that corresponds to the same true energy regardless of where in the detector it was deposited.

The variable W is central to the calibration. This refers to the length along the cell at which a hit was reconstructed. So for vertical cells which measure the x and z

component of the hit, W gives the y coordinate. W runs from -780 to 780 cm in the far detector, with more positive numbers closer to the readout.

The calibration relies on tricell hits. This refers to hits that have three cells in a row that were lit up following a particle trajectory. In this case, the tracklength through the middle hit is easily determined by simply correcting for the angle of attack of the particle. This angle of attack is calculated by performing a quick least-squares linear fit to the cosmic track. A diagram of a tricell hit is shown in Figure 7.1. A tricell hit refers specifically to the interior hit of the triplet, and the calibration will consider only these tricell hits.

The calibration begins after selecting a set of tricell hits from cosmic tracks. For each tricell, the plane and cell numbers are recorded along with the W reconstructed by the track and the ADC value of the hit. For each cell, a two dimensional histogram of energy deposit per unit tracklength, in PE/cm, vs W is made and the mean PE/cm value as a function of W is stored as a profile histogram. For example, this 2D histogram is shown in Figure 7.2 with the mean profile overlaid.

Before correcting for the fiber attenuation, two corrections are applied to account for threshold effects. Thresholds bias the true dE/dx distributions for recorded hits. This is most visible in the far end of cells in the far detector, where attenuation in the fiber can make the PE count from a MIP less discernible from statistical noise

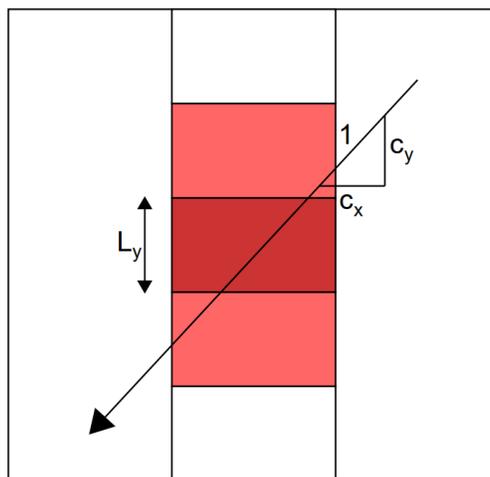


Figure 7.1: A diagram of a tricell hit in the detector where three consecutive cells in a plane have been hit. The track length through the inner cell is easy to reconstruct and is given simply by $L_y/\cos \theta_y$ using the diagram's nomenclature. This relation does not hold for the outer cells in the triplet, so only cells like the middle, dark red are used for calibration

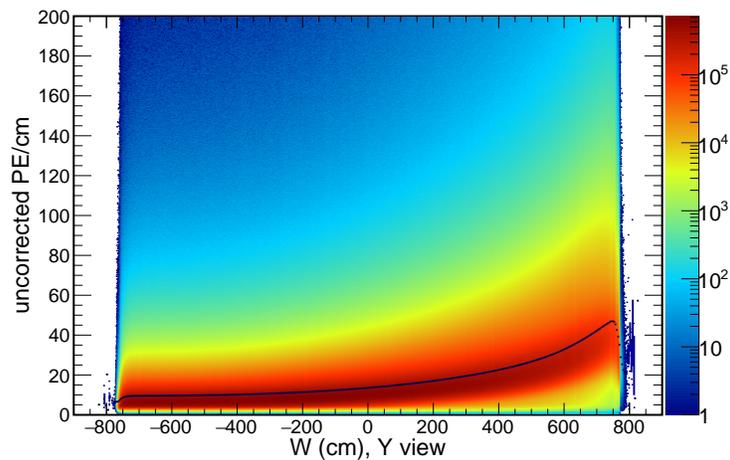


Figure 7.2: The 2D distribution of PE/cm vs W for all through-going cosmic ray muons in data for far detector Y-view cells is shown by color while the mean PE/cm value for each value of W is shown by the black data points. There is an obvious bias as a function of W dominated by light attenuation in the fiber that the attenuation calibration will correct.

than near the readout. The threshold correction also corrects for self-shielding in the detector. As cosmic particles lose energy while traveling through the detector, the flux of particles has a different energy distribution at the bottom of the detector and thus further affect the rate that hits fall below threshold.

To correct for these threshold effects in each cell, PE deposits measured in data and MC are adjusted by the bias induced by the thresholds, as predicted by simulation. Both threshold effects grow with the size of the detector and so the threshold corrections are larger in the far detector. This procedure corrects PE counts by 10-30% across the face of the detector. This correction is shown for X-view and Y-view cells in Figure 7.3 as a function of cell number within the plane and W .

To correct for the attenuation in the fiber, the central portion of a cell is fit to the empirical function

$$y = C + A \left(e^{\frac{W}{X}} + e^{-\frac{L+W}{X}} \right) \quad (7.1)$$

In this fit function, L gives the length of the NOvA Cell, and C , A , and X are fit parameters. The fiber is looped so that there are two paths photons may take to reach the readout. The first term gives the expected PE from the direct route through the fiber while the second gives the PE expected from looping around the bottom of the

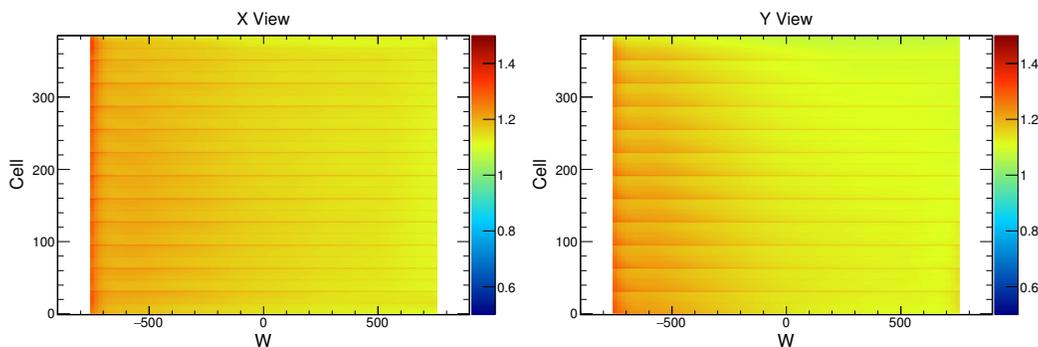


Figure 7.3: The correction to the mean PECorr/cm caused by threshold effects, as calculated by MC, as a function of cell number and W . The left(right) plot shows the threshold calculation for X(Y)-view cells

cell and returning to the APD. X is immediately recognized as the fiber attenuation length. Rather than fix this to teststand measurements, this is allowed to float in the fit for an in-situ measurement that accounts for cell-to-cell variations.

This functional form is only valid in the center of cells with $-750 < W < 750$ cm at the far detector, $-150 < W < 50$ for vertical cells in the near detector muon ranger, and $-150 < W < 150$ cm for the rest of the detector. The parameterization in Eqn. 7.1 has been shown to fit the near and far ends quite poorly with each exhibiting a characteristic "rolloff" shape. To capture this shape, a non-parametric, LOcally WEighted Scatter plot Smoothing (LOWESS) algorithm[80] is applied to capture the rolloffs, along with any residual shape in the central part of the cell not captured by Eqn. 7.1. This smooths the mean PE/cm profile with nearby hits weighted as

$$w_i = \begin{cases} \left(1 - \left|\frac{W - W_i}{30cm}\right|^3\right)^3 & |W - W_i| < 30cm \\ 0 & |W - W_i| \geq 30cm . \end{cases} \quad (7.2)$$

An example of the attenuation fitting procedure is shown in Figure 7.4. After the LOWESS smoothing, the attenuation is complete and has constructed a curve across the W range of a cell relating to the same true mean energy deposit for cosmic muons, called PECorr. After constructing this curve in every cell of this detector, hits may be calibrated to a consistent energy scale across the detector. The simulated mean energy reconstructed per true energy deposited is shown in Figure 7.5.

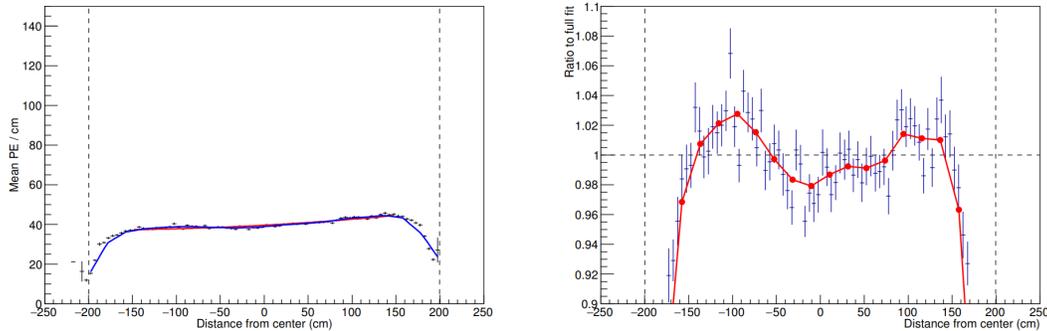


Figure 7.4: The attenuation correction curve in the near detector. On the left, the mean PE/cm vs W profile is plotted. The red curve gives the fit to the attenuation parameterization while the blue shows the attenuation curve after LOWESS smoothing. The right shows the LOWESS correction through the residuals from the initial attenuation fit.

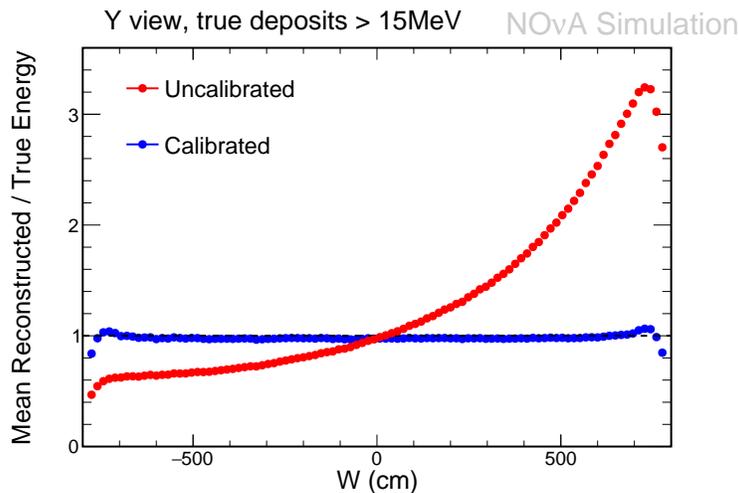


Figure 7.5: The mean PECorr per true energy deposit before(after) the attenuation calibration in red(blue). Before the attenuation calibration, there is a factor of six change in detector response as a function of W . After the correction is applied, the distribution is relatively flat with only slight discrepancies very near the cell ends.

7.2 Absolute Calibration

After the attenuation calibration, there is a prescription for translating a measured ADC value to a consistent energy scale that is valid for energy deposits left anywhere in the detector. The final step in the calibration is relating this to a human-readable unit, GeV.

This is again done with tricell hits from cosmic muons, but with additional restrictions on the endpoint of the muon. In the far detector, muons are required to stop with $200 < W < 600$ [81]. This removes reliance on threshold modeling at low W

and the issues with the attenuation curve shape very near the readout. In addition to an endpoint cut, each track is required to have a reconstructed Michel electron near the track endpoint to eliminate any proton contamination in the sample.

After selecting muon tracks, the dE/dx is plotted in PECorr/cm as a function of the distance to the endpoint of the track. This is shown in Figure 7.6. A Bragg peak is clearly visible, and hits more than 200 cm from the endpoint, the relativistic rise of the Bethe-Bloch equation, affecting the mean dE/dx but not the most common value, begin to influence dE/dx . Thus, hits in the 100-200 cm range from the track endpoint are used as the standard candle to calibrate to the expected energy loss of a MIP in the detector.

The PECorr distribution of hits in the calibration window 100-200 cm from track endpoint for both data and MC are then used to translate energy deposits into a physically meaningful unit. A PECorr to MeV conversion factor is determined by scaling the mean PECorr/cm of this sample to the MC prediction for MeV/cm in the MC sample. After determining this ratio for both data and MC in both detectors, the calibration procedure is complete. The correction procedure is shown in Figure 7.7 for far detector data and MC.

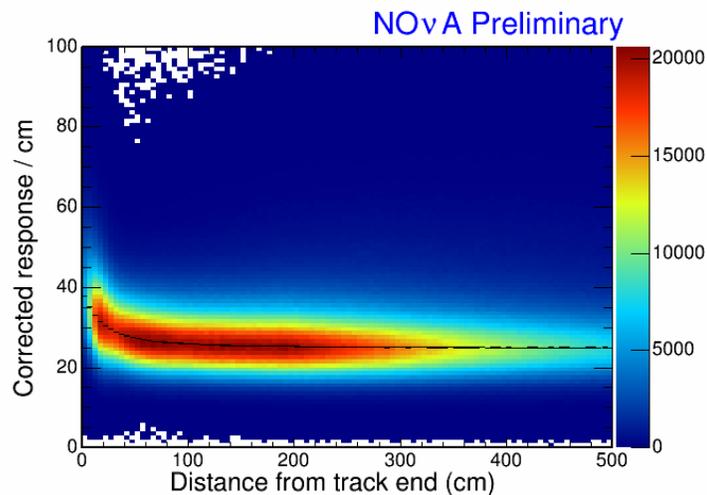


Figure 7.6: The distribution of PECorr/cm for all hits on muon tracks as selected for the absolute calibration as a function of the distance to the track endpoint. Hits in the calibration window from 100-200 cm to the track endpoint are used to calibrate absolute energy by scaling average PECorr/cm to the simulation's prediction of average MeV/cm.

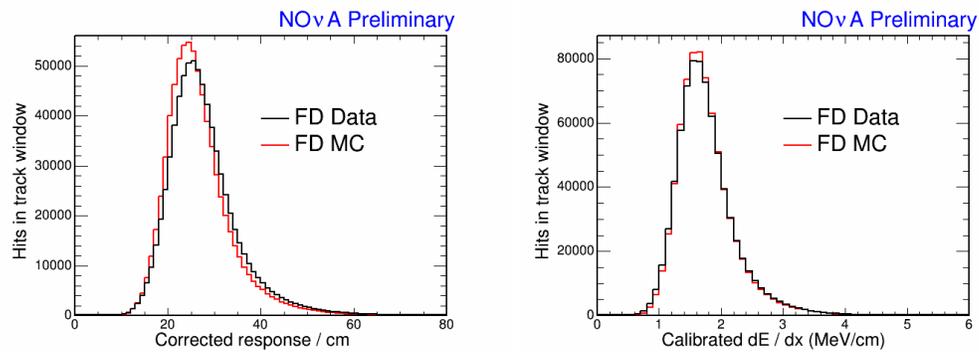


Figure 7.7: The dE/dx distribution for hits in the calibration window in PECorr/cm, left, and calibrated MeV/cm, right. Before the absolute calibration is applied, the plots on the left don't perfectly align.

EVENT RECONSTRUCTION

The DAQ and calibration procedures produce a set of hits in the detector, termed CellHits. These hits contain a time-stamp and an energy deposit, calibrated to GeV. From these CellHits, the event reconstruction chain aims to properly identify neutrino interactions and approximate particle trajectories within each neutrino interaction. Both the ν_e appearance and ν_μ disappearance analysis rely on reconstruction as an input into event classifiers and energy estimates used in the analysis.

In the NOvA detectors, ν_e CC and ν_μ CC events are topologically very different and thus require different reconstruction tools to describe the event. Typical event topologies are characterized in Figure 8.1. A ν_μ CC event produces a muon in the final state which leaves a long track of constant dE/dx , while the electron in a ν_e CC event builds an electromagnetic shower. Figure 8.1 also shows an NC interaction that produces a high energy π^0 , which decays into two photons with a branching fraction of 98.6%[82].

NC π^0 production is a particularly problematic background since the finite granularity of the NOvA detectors merges both decay photons into a single indistinguishable shower for nearly parallel decay photons. The photon conversion length in the detectors is about 40 cm, while cells are about 4 cm wide. In order for there to be a “gap” cell with no activity between the two photon showers so that the two showers can be reconstructed separately, this geometry suggests that the photons must have $\theta_{12} > 0.28$. Due to relativistic kinematics of π^0 events, there are negligibly few π^0 decays with $E_\pi > 1.25$ GeV[83] that satisfy this angular constraint.

The overall chain begins with a clustering algorithm identifying neutrino interactions. After running a clustering algorithm, a Hough transform[84] finds an approximate event vertex. A fuzzy k-means algorithm[85] then reconstructs the vertex by minimizing the spread of energy deposits in angular space in both views simultaneously. Finally, prongs, NOvA’s fundamental reconstructed representation of a particle trajectory, are first drawn in each view separately by finding clusters of energy deposits in angular space. Three-dimensional prongs are then constructed by pairing x - and y -view prongs. These are the base reconstruction algorithms used by all NOvA analyses. A diagrammatic overview of this reconstruction chain

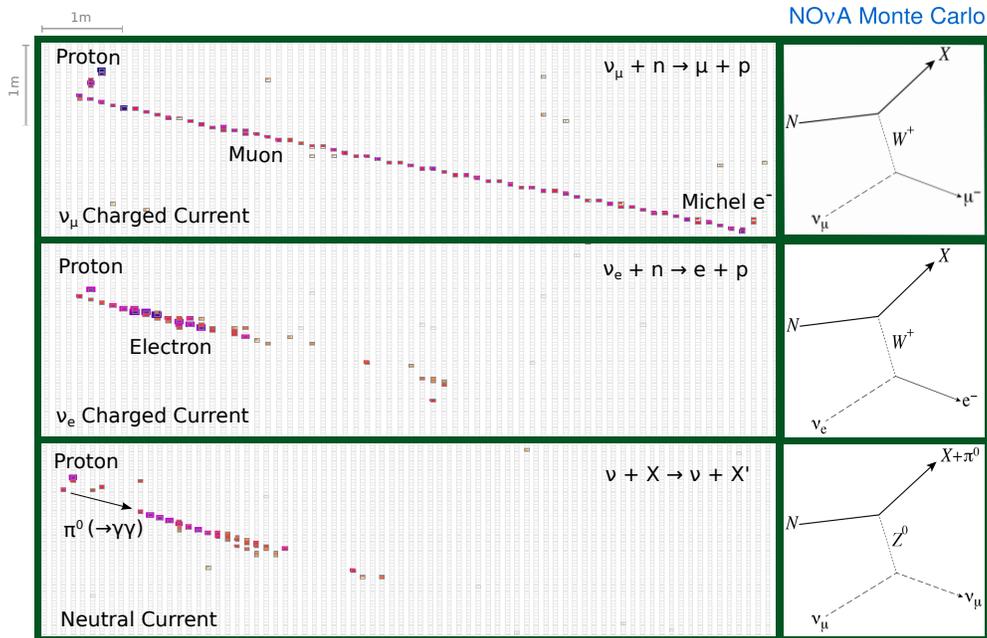


Figure 8.1: An example event display for an interacting ν_μ CC (top), ν_e CC (middle), and NC+ π^0 (bottom). These are the most important event topologies to reconstruct as the ν_μ disappearance and ν_e appearance use ν_μ CC and ν_e CC events as signal, respectively. Additionally, NC events with a π^0 are an important background to the ν_e appearance analysis as the electromagnetic decay of the π^0 can mimic an electron.

is shown in Figure 8.2. Additionally for ν_μ candidate events, tracks are used to reconstruct muon candidates. Specially developed tracking algorithms can more appropriately follow scatters along the muon trajectory than the general prong reconstruction. The last step in the reconstruction chain used in the oscillation analysis is a Michel electron-finding algorithm. Each of these processes is described in detail in the following.

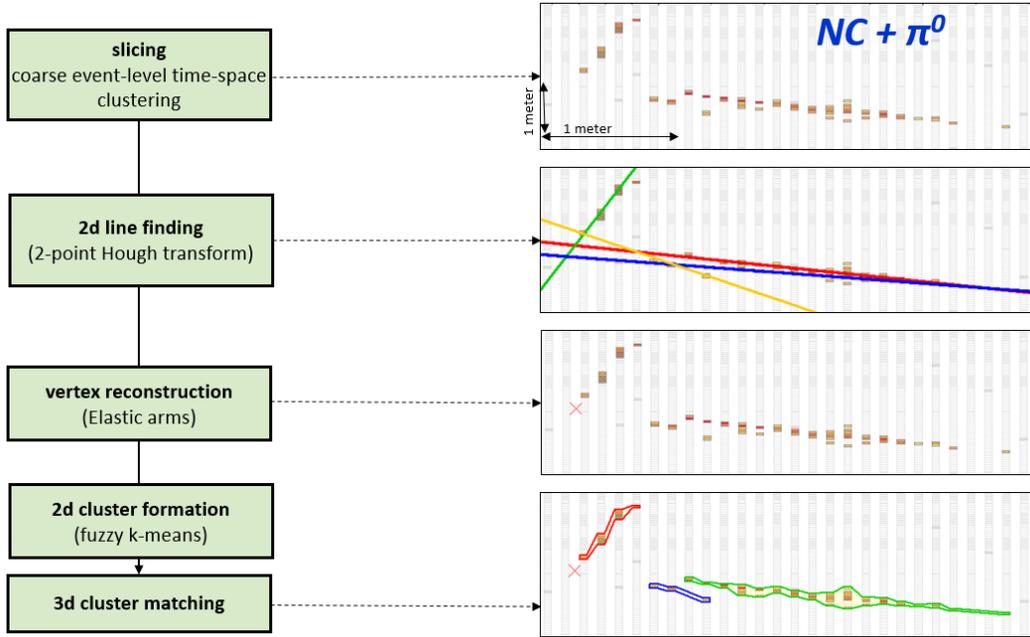


Figure 8.2: A schematic of the NOvA reconstruction chain. From clustering hits together into a slice, a vertex is reconstructed from which we draw prongs, our reconstructed representation of a particle.

8.1 Event Clustering

CellHits recorded by the DAQ are clustered into “slices” that represent a distinct reconstructed physics interaction in NOvA analyses. Clusters are produced with an implementation of the DBSCAN algorithm[86]. This algorithm will generate clusters of varying density while also rejecting hits consistent with noise. The DBSCAN algorithm is convenient in that it can generate clusters of arbitrarily complex shapes. NOvA uses Slicer4D, a custom-tuned DBSCAN algorithm, to group the hits triggered by the DAQ into reconstructed neutrino interactions[87].

The distance between two CellHits is fundamental to DBSCAN and is determined using the spatial and temporal “score” between hits as an indicator of correlation. The spatial score between two hits is complicated by the fact that each hit only measures the xzt or yzt coordinates. If $u_{x/y}(i)$ is 1 if and only if hit i is in a x/y -view plane, the score between two points is:

$$r_{ij} = (x_i - x_j)^2 u_x(i)u_x(j) + (y_i - y_j)^2 u_y(i)u_y(j) + (z_i - z_j)^2. \quad (8.1)$$

Strictly speaking, this is not a metric function as it is possible for $r_{ij} > r_{ik} + r_{kj}$ if hit k measures a different x, y component than i and j . But, as each hit only measures three dimensions, xzt or yzt , this is a reasonable score function to use

when clustering in all four dimensions simultaneously. Using this, the overall score between two hits is

$$d_{ij} = \frac{(|(t_i - t_j)| - r_{ij}/c)^2}{\tau^2} + \frac{r_{ij}^2}{dr^2}, \quad (8.2)$$

where τ is the quadrature sum of the timing resolutions of the two hits, c is the speed of light correcting for time-of-flight, and dr defines the unit spatial distance. In the near (far) detector, dr is set to 75(100) cm while τ depends on the energy of the hit but is approximately 10(60) ns. As noted previously, the ADC reading on the APD's are read out four times as frequently in the near detector, in anticipation for dealing with high event pileup.

With this score function, DBSCAN calculates scores between each pair of CellHits. A neighborhood around a point, i , is defined as all points, j , that have $d_{ij} < \epsilon$ where $\epsilon = 5(2)$ in the near(far) detector. A hit in the detector is said to be a core hit of a cluster if there are at least four hits that lie within its neighborhood. DBSCAN builds a cluster by first identifying a core point with at least N neighbors and collecting the neighboring hits into the cluster. In NOvA, we take $N = 4$. It then iterates through each of the remaining cluster hits. If the queried hit is also a core hit, its neighbors are added to the cluster if they have not already been. The clustering process is shown schematically in Figure 8.3 with all connected core hits in red. Border hits are defined as the hits that are not core hits, but lie within the neighborhood of a core point. These are shown in green in the figure and are included in the cluster. Noise hits are defined as any hit in the space that is not either a core or border hit of some cluster.

Slicer4D's performance was tuned on neutrino simulation with two metrics. The purity of a slice is defined as the fraction of energy deposited in a reconstructed slice that came from the dominant neutrino interaction and the completeness is the fraction of energy deposited in the detector by the dominant neutrino that was captured in the slice. As shown in Figure 8.4, Slicer4D is very effective at generating slices of high purity and high completeness indicating Slicer4D does well with isolating neutrino interactions from nearby physics. The reconstruction chain will subsequently build upon these clusters generated by Slicer4D.

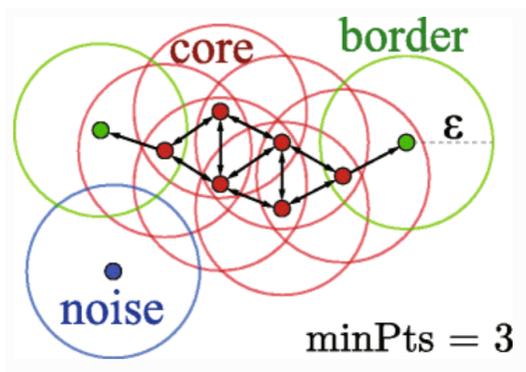


Figure 8.3: An illustration of the DBSCAN algorithm. Points in red are labeled as core points and included in the cluster since they have at least N , three in this case, hits in their neighborhood. The green points are border points that lie within the neighborhood of a core point but are not core themselves. The blue hit is labeled as noise as it is sufficiently isolated to not be in the neighborhood of some core point. Figure from[88]

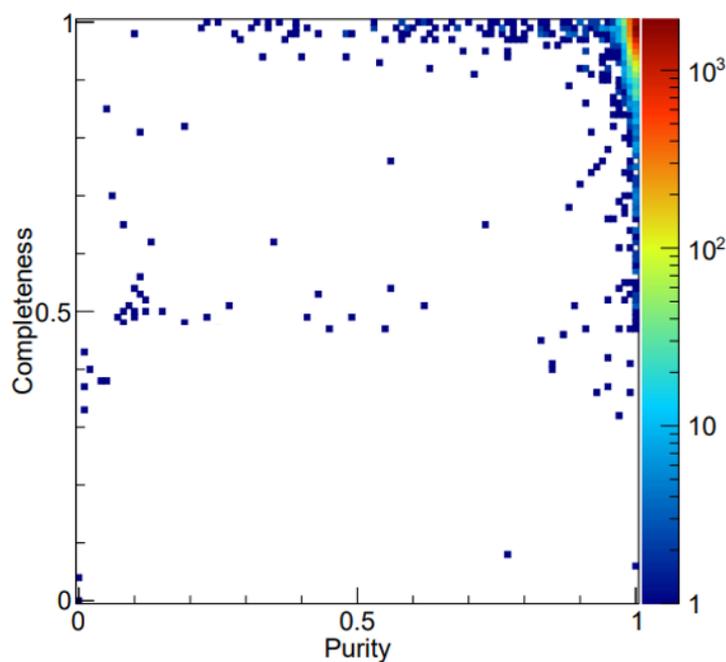


Figure 8.4: The purity and completeness for far detector simulated neutrino interactions determined by the event slicer. The purity is a measure of how much external light entered the neutrino interaction while the completeness is a measure of what fraction of the neutrino's visible energy was deposited in the slice.

8.2 Vertex Seeding with a Hough Transform

NOvA reconstructs an event vertex with the ElasticArms algorithm described below. But, this is a very computationally expensive algorithm. Thus, the vertex is first

seeded with the output from a modified Hough transform, a simple algorithm for detecting and reconstructing lines in data[84]. These lines give a quick and simple reconstructed representation of a particle trajectory through the detector. Because of our detector geometry, the algorithm runs in each view separately.

Within a Hough transform, a Hough map is first constructed by taking pairs of hits in the slice, calculating the line between them, and plotting in (ρ, θ) space where ρ gives the distance of shortest approach to the line and θ is the angle made between the x -axis and the line[89]. A Gaussian smoothing is applied to the determined (ρ, θ) pair according to

$$w_{ij} = e^{-(\rho-\rho_0)^2/\sigma_\rho^2} e^{-(\theta-\theta_0)^2/\sigma_\theta^2}, \quad (8.3)$$

with σ_ρ tuned to $3/\sqrt{12}$ and σ_θ set to $3/d\sqrt{6}$. Only hits that are within $\sqrt{15000} \approx 120$ cm contribute to the Hough map. The upper limit was tuned to decrease the computational requirement while maintaining performance.

Once the Hough map has been populated with all eligible hit pairs, lines are evident as peaks in the (ρ, θ) space. During development, there was no clear prescription for generating a set of robust Hough lines outright as there were often additional lines drawn by statistical noise originating in cross-talk between different physical lines in the event. Thus, an iterative approach is employed where first the most prominent peak in the Hough map is taken as a reconstructed line. Its associated hits are subsequently masked from the event. The procedure then continues until either ten lines are drawn or there is no sufficiently prominent peak in the Hough map. This dramatically emphasizes sub-dominant peaks in the Hough map as illustrated in Figure 8.5.

Two metrics were evaluated to test the performance. First, the distance of closest approach between the true vertex and each Hough line was calculated. The first(second) Hough lines in ν_μ CC events were on average 4.068(8.178) cm away from the true vertex. The distribution of this metric is shown in Figure. 8.6. Additionally, the reconstructed direction of the Hough line was compared to the true direction of the corresponding particle. For ν_μ CC events, the average cosine between these two directions is 0.991(0.9481) for the first(second) Hough lines.

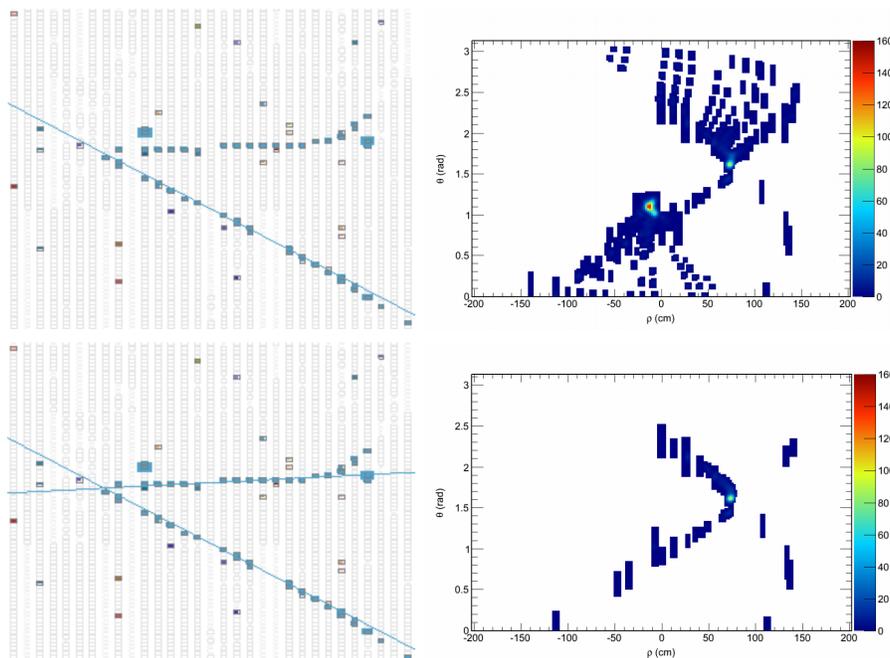


Figure 8.5: An example event showing the iterative procedure for reconstructing Hough lines. In the top row, the most dominant line is drawn according to the maximum in the Hough map. On the bottom, the Hough map is recalculated, but with hits from the dominant line vetoed which makes the sub-dominant Hough line much clearer.

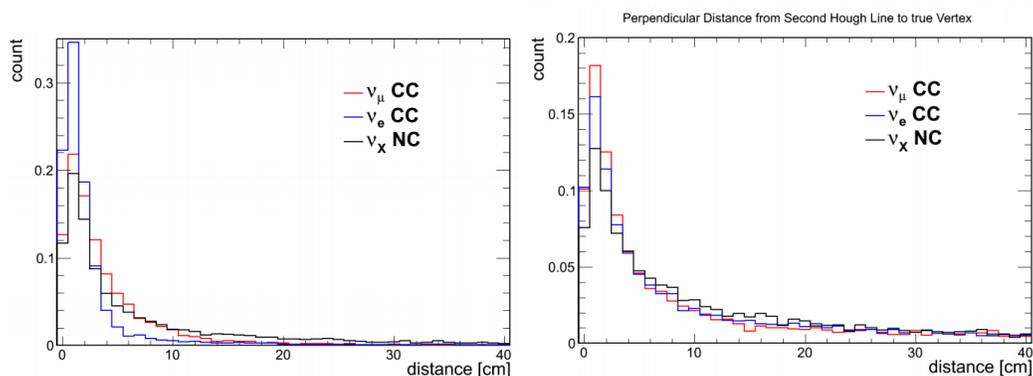


Figure 8.6: The distribution of distance of closest approach between the true vertex and first(second) Hough lines on the left(right) for ν_μ CC (red), ν_e CC (blue), and NC (black) events.

8.3 Vertex Identification

An ElasticArms algorithm was tailored to NOvA's specific needs to precisely reconstruct an event vertex. Broadly, for a given vertex position hypothesis, this method assumes that particles should be emitted from the vertex in straight lines, or arms, and computes an energy function for how far observed hits are from these arms.

Each hit can be shared among multiple arms governed by a weight parameter which must not be greater than 1 for any hit. The vertex is reconstructed by minimizing this energy function.

The NOvA energy function involves three terms:

$$E = \sum_{i=1}^N \sum_{a=1}^M V_{ia} M_{ia} + \lambda \sum_{i=1}^N \left(\sum_{a=1}^M V_{ia} - 1 \right)^2 + \frac{2}{7X_0/9} \sum_a^n D_a, \quad (8.4)$$

where a runs over the M postulated arms and i runs over all N hits in the slice. The number of arms, M , is set by the number of lines found by the Hough transform algorithm. Here, V_{ia} gives the weight of hit i in arm a and M_{ia} gives the distance of closest approach between the hit and arm. The first term is a measure of how far away the hits associated with each arm are from the postulated arm, with λ a tunable strength parameter. The second term is a penalty for hits that are not fully associated with a given arm. This term prevents the trivial maximization scenario where all hits are totally uncoupled from all arms. In the third, D_a is the distance between the vertex and the arm start position. This term is included specifically for NOvA, and is tied to X_0 , the radiation length in the detector. This is a penalty for reconstructing arms further from the origin than photons are expected.

Hit association with the given arms is parameterized by the distance of closest approach to each of the arms, M_{ia} :

$$V_{ia} = \frac{e^{-\beta M_{ia}}}{e^{-\beta \lambda} + \sum_b e^{-\beta M_{ib}}}, \quad (8.5)$$

where $\beta = 1/T$, an abstract temperature which governs how free hits are to be shared among multiple prongs that may be quite far away in angular space. At very high temperatures, this blurs out the event, and only the large scale structure of the event is visible to the algorithm. We then anneal the system, bringing the temperature of the hit associations down while seeding each step with the previous iteration, so that hits are associated more rigidly with nearby prongs[90]. This procedure prevents reconstructing a vertex at a local minimum as the high temperature computation by construction settles near the global minimum with lower temperature iterations migrating to a more precise vertex using finer scale details relevant to the event.

The procedure involves tuning several parameters, which is described in [90]. The resolution in x and y is roughly 4 cm, the width of a cell, according to simulation, while the resolution in z is about 6 cm, the length of a cell. These resolutions for ν_μ

CC interactions are shown in Figure 8.7. In three dimensions the vertex is biased by an average of 11.56(10.65) for $\nu_\mu(\nu_e)$ CC interactions.

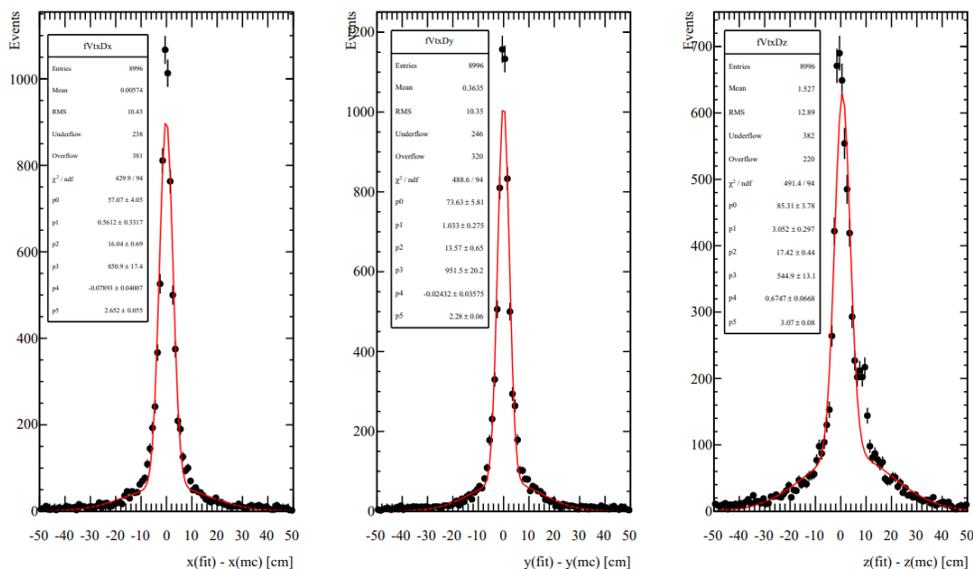


Figure 8.7: The resolution of the x (left), y (middle), and z (right) components of the vertex for true ν_μ CC interactions.

8.4 Generating Prongs

Prongs serve as NOvA's fundamental reconstructed representation of a particle. Though ElasticArms outputs a set of arms roughly translating to such a construct, the algorithm in practice was shown to often overestimate the number of particles reconstructed. Thus, a dedicated prong-reconstructing algorithm inspired by Fuzzy-K Means[85] was developed to accurately reconstruct prongs corresponding to physically distinct particles. The Fuzzy-K Means algorithm requires both the number of prongs and that the association weight for each hit summed over all prongs must be one. The later is a problem since, in our detector, there are often noise hits in reconstructed slices and very short tracks that are too small to reconstruct a full prong. Ideally, both of these classes should not be associated with any prong. To address these issues, we employ the FuzzyKVertex algorithm. This is a slightly tailored application of the Probabilistic Clustering Algorithm, an extension of the classic Fuzzy-K Means algorithm[91].

In NOvA, the vertex is fixed to the output of ElasticArms. Prongs are generated in each view separately. Prongs are assumed to be straight lines directed radially outward from the origin. In this framework, a particle will be fundamentally represented by an angle, a single number. Prongs are identified by peaks in the

angular distribution of energy deposits.

After determining appropriate prong centers in angular space, hits are associated between prongs according to the angular separation between hit i and prong a ,

$$d_{ia} = \frac{(\theta_i - \theta_a)^2}{\sigma_i^2}, \quad (8.6)$$

where σ_i is an uncertainty on the true angle from the origin of the given hit arising from discreteness of the detector. For hits near the vertex, σ_i can be quite large, while further away this becomes smaller due to a much longer lever arm. Hit membership in each prong is then

$$\mu_{ia} = e^{-\alpha d_{ia} \sqrt{N}}, \quad (8.7)$$

where α is a tunable parameter, set to 0.5, and N gives the number of prongs reconstructed in the event.

Once prong centers have been found and hits have been associated with them, the final step is to merge the two-dimensional prongs in each view into a set of three-dimensional prongs. This is done by comparing the cumulative energy deposition as a function of z in both views. Two two-dimensional prongs from separate views are married according to the maximum discrepancy between these two CDF's.

In all, the FuzzyKVertex algorithm reconstructs the primary lepton in 92(88)% of $\nu_\mu(\nu_e)$ CC events. The efficiency is higher at high lepton momentum, characteristic of events in our beam peak. This efficiency, shown as a function of the true lepton energy, is shown in Figure 8.8.

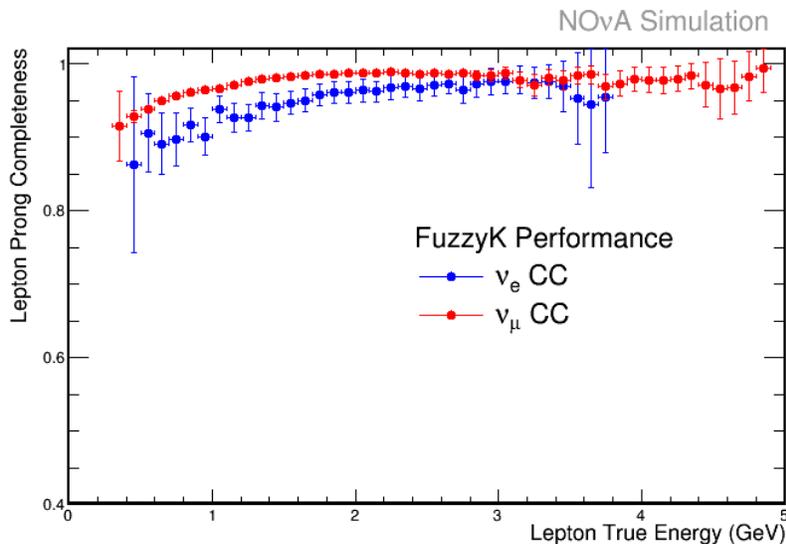


Figure 8.8: The efficiency for reconstructing the primary lepton in CC events with the FuzzyKVertex algorithm as a function of the true lepton energy. Muons are shown in red while electrons are shown in blue.

8.5 Reconstructing Tracks with a Kalman Filter

The FuzzyKVertex algorithm is very efficient at identifying muon prongs in the detectors, but the detector shows more information about long tracks than a prong can capture. Namely, as a muon undergoes multiple scattering, its direction can change slightly while tracking through the detector. While FuzzyKVertex can not accommodate these scatters, a tailored tracking algorithm can.

For reconstructing muons in ν_μ CC events, we use a Kalman filter algorithm[92]. The KalmanTrack algorithm seeds tracks by looking for pairs of hits that are at least four planes away and adding colinear hits with these seeds[93]. From the initial seed hits, the track is propagated forward, gradually adding hits that are colinear with the seeds. The expected change in angle from multiple scattering is about 5 mrad per plane of the detector[93]. When accruing hits along the track, hits are allowed to deviate from the forward propagated track prediction by eight units of χ^2 with a spread set to this 5 mrad. Constructing the full track in this way allows the track reconstruction to follow the slight angular deviations expected for muon tracks in the detector. In this way, the Kalman tracker more efficiently collects hits from the primary muon into a single reconstructed object than the prong finder.

In addition to the reconstructed track, a muon PID, the Reconstructed Muon Identifier (ReMId)[94], was developed. Rather than being an event-level classifier, this is a

measure of how consistent an individual track is with a muon trajectory. ReMId is a BDT based on four input variables. This uses a likelihood for the dE/dx hypothesis expected for a muon track and a scattering likelihood that aims to reject π^\pm tracks that undergo more dramatic scatters. Additionally, the BDT uses the tracklength and the fraction of planes along the track that contain no hadronic energy.

8.6 Energy Reconstruction for ν_μ CC and ν_e CC events

NOvA is tasked with measuring oscillation probabilities on a fixed baseline as a function of energy to provide a range of L/E to analyze. To do this, separate energy estimators were developed for both ν_μ CC and ν_e CC events. The reconstruction methods used are the same for both neutrino and anti-neutrino samples, but separate tunings are run for the ν and $\bar{\nu}$ estimators.

For ν_μ CC events, the muon and hadronic energies are estimated separately and summed to give the neutrino energy estimate:

$$\hat{E}_\nu = \hat{E}_\mu + \hat{E}_{\text{had}}. \quad (8.8)$$

The muon in the event is identified as the track with the highest ReMId score. Muons have a low, MIP dE/dx giving quite a long track-length. Thus, E_μ can be much more precisely measured from track-length that has an uncertainty determined by the length of a plane than from the calorimetric energy, which is sensitive to electronic noise. In the near detector, the track-length is separated into the portion in the fully-instrumented detector and the portion traveling through the muon catcher. The energy passing through each region is reconstructed separately and then added together. The hadronic energy is then reconstructed from the calorimetric energy not on the muon track. Both \hat{E}_μ and \hat{E}_{had} are reconstructed with a spline fit to optimize the neutrino energy resolution, whose results are shown in Figure 8.9. These fits give a 3% resolution for E_μ and a 30% resolution for E_{had} .

With ν_e CC events, the electron is not measurable by track length. The length of an electromagnetic shower on average grows like $\ln E_e$, while stochastic bremsstrahlung activity towards the end of the shower can vary the true shower length dramatically on an event-by-event basis. Thus, both the E_{EM} and E_{had} must be measured from calorimetry. A calorimetric electromagnetic energy is determined using prongs that are electron- or photon-like while an uncorrected hadronic is given by the calorimetric energy not captured in the electromagnetic part. Then, a fit is performed to minimize the resolution spread assuming a quadratic:

$$E_\nu = aE_{\text{EM}} + bE_{\text{had}} + cE_{\text{EM}}^2 + dE_{\text{had}}^2. \quad (8.9)$$

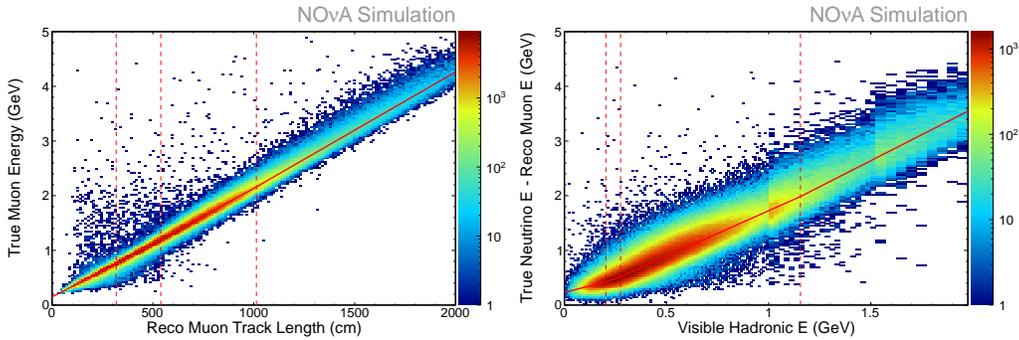


Figure 8.9: Spline fits used in reconstructing both the \hat{E}_μ , left and \hat{E}_{had} , right. The migration matrix is approximately linear. To account for the nonlinearities, both quantities are reconstructed in a piecewise linear fit.

There is no constant term allowed. Such a term can be readily identified with differences in binding energies of initial and final state nuclei, but these energies are necessarily much lower than the neutrinos studied. Additionally, the fit is performed on a distribution of events that is flat in energy. This is to remove the effective Bayesian prior which biases events to reconstruct in the beam peak, where events are very concentrated in true energy. This flat spectrum training significantly reduces the residual bias in mean reconstructed energy observed as a function of the true energy. This fit gives a neutrino energy resolution of 11%. The surface that is fitted along with the neutrino energy resolution is shown in Figure 8.10.

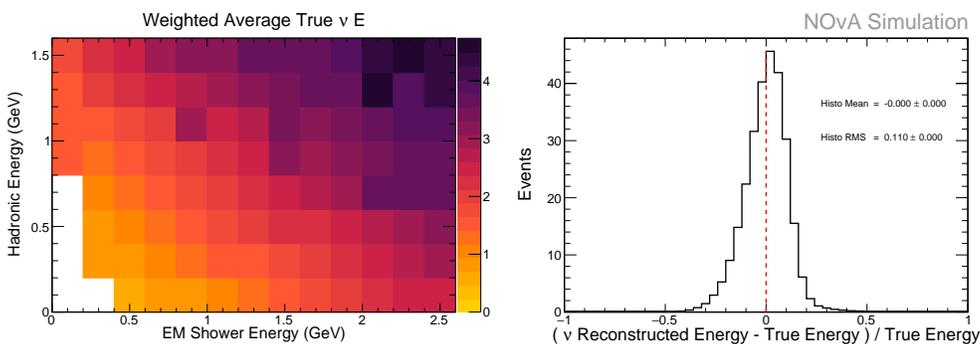


Figure 8.10: The left shows the mean true energy, in GeV, as a function of reconstructed electromagnetic and hadronic energies. A quadratic fit is performed on this surface, giving our energy estimator. The right shows the relative resolution of our far detector sample.

8.7 Event Classification with a Convolution Neural Network

The neural network has long been a common strategy employed in high energy physics to classify interaction types according to an event's reconstructed topology.

Such an identifier was developed for the NOvA ν_e appearance analysis[95]. This PID, however, only selected 35% of interacting ν_e CC signal in the far detector. In an experiment with limited event counts like NOvA, there is always an emphasis on maximizing the efficiency that can be reasonably achieved. In a search for maximizing the selection efficiency, a convolution neural network, CVN, was developed that achieved an efficiency increase for selecting ν_e CC events equivalent to a 30% increase in accumulated exposure[96].

Our implementation follows the GoogLeNet architecture[97]. The eventual goal of the algorithm is to produce a feature map, roughly equivalent to conventional reconstruction in high energy physics, through a series of refining edge detecting algorithms. The feature maps are then fed through a conventional neural network producing a classification score.

The algorithm uses a three dimensional pixel map for each view as input encoding the cell, plane, and charge of hits within a slice. Features are then extracted from the feature extraction operational unit, the inception module. For each instance of an inception module, several image detection kernels are applied in parallel with different kernel sizes and pooling. Afterwards, the output of each process is condensed into the next iteration of the feature map. The events are passed through three inception modules in each view separately before being joined together in a final inception module immediately upstream of the final neural network identifier. Treating both views independently until the last inception module gives the PID increased sensitivity by allowing it to see telling features in one view that may have been eclipsed by other particles in the other view such as a photon conversion gap.

The pixel map for a ν_e CC and ν_μ CC event is shown along with its feature map in Figures 8.11 and 8.12, respectively. The feature map shows the response of several image detection algorithms for the pictured event. For pedagogy, three of these feature maps have been highlighted and emphasized to show differences between the two displayed events.

NOvA's implementation of CVN is trained to identify ν_e CC, ν_μ CC, ν_τ CC, NC, and cosmic events. Additionally, there is an "other" category to catch all other events. Each interaction mode is further subdivided into final state topology, determined by the multiplicities of nucleons and pions. An overall CVN value for each interaction channel is achieved by summing the CVN output for all possible final state topologies for the given channel. In this way, CVN is a versatile classifier algorithm used in several NOvA analyses, including all those analyses presented in this work. The

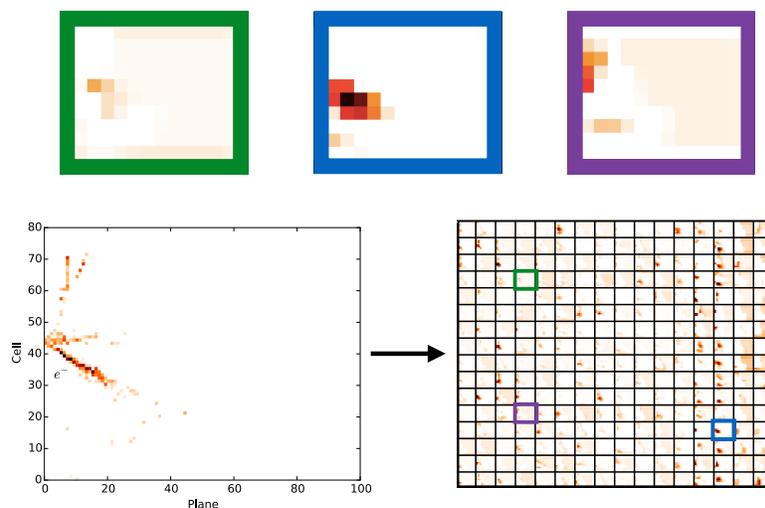


Figure 8.11: A raw pixel map for an ν_e CC event, left, and the derived feature map, right. The response to each of the features shows the response to an individual image detection algorithm. These responses are quite different for different interaction types.

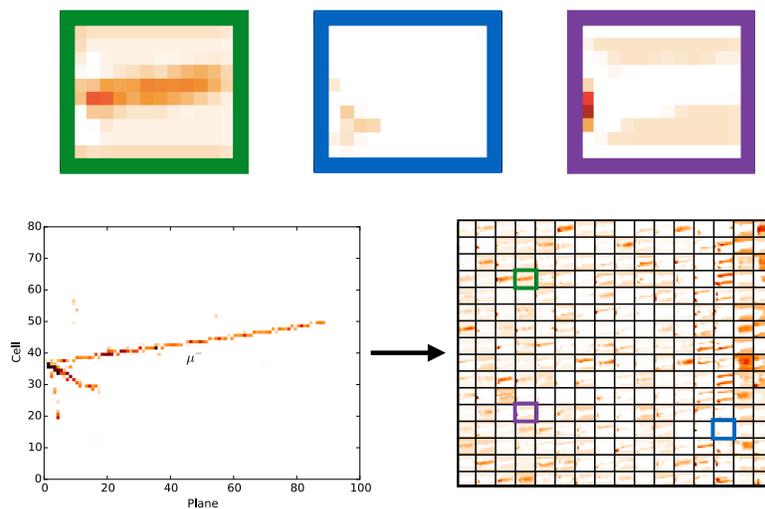


Figure 8.12: A raw pixel map for an ν_μ CC event, left, and the derived feature map, right. The response to each of the features shows the response to an individual image detection algorithm. These responses are quite different for different interaction types.

specifics of the PID performance as well as selected backgrounds will be discussed later.

8.8 Reconstructing Michel Electrons

Free muons decay to electrons with a lifetime of $2.2 \mu\text{s}$ [82]. This is significantly longer than the timing resolution of hits, roughly 10(60) ns in the near (far) detector, so that CellHits from Michel electrons are rarely sliced in with the parent slice. As such, Michel electrons are an indicator of a muon in a neutrino interaction entirely disjoint from reconstruction within the slice. And thus, the presence of a Michel electron is not correlated with other reconstruction and classification algorithms. NOvA uses Michels for three purposes. As discussed earlier, they are used to tag stopping muons in the detector as a part of the calibration procedure. And, as the Michel energy spectrum is known theoretically, the energy spectrum is used as a cross-check of the calibration procedure. As will be discussed later, Michel electrons are used in the oscillation analysis to constrain the far detector ν_μ CC background selected by the ν_e selection.

The immediately obvious algorithm for detecting Michel electrons in a scintillator tracker is to look for time delayed CellHits near the end of a reconstructed track. This is exactly the procedure required for calibration as the muon track contains the physics of interest while the Michel merely tags those muons that stop in the detector. However, this is not a suitable strategy for identifying Michel electrons in slices that were selected as ν_e -like as, if there were an easily trackable muon, it would be rejected by the selection. An illustration of this is shown in Figure 8.13. A simulated muon did not leave enough energy in the detector to be reconstructed, but there is evidence of the muon from the decay Michel electron. Thus, we must be able to detect Michel electrons caused by muons too low in energy to track. To satisfy both needs, the MEFinder algorithm[98] was developed to identify Michel electrons found near any hit within the parent physics slice. It produces two Michel qualities: Michels found near the end of a track endpoint, called TrkME's, and Michels found elsewhere in the slice, called SlcME's.

The algorithm begins by selecting all Michel candidate CellHits in a spill that were not already associated with a physics slice. To be a candidate hit, the hit must have registered at least 50 ADC of energy. This cut essentially removes all electronic noise while rejecting a negligible fraction of Michel hits. It must also be within 40 cm of some hit in a physics slice and within $10 \mu\text{s}$ of the mean time of the parent physics slice. Further, if the potential Michel candidate hit is a retrigger of a cell that was lit in the parent slice, it must have occurred more than $1.2 \mu\text{s}$ after the parent slice. This is due to poor modeling of the APD dead-time, observed to be $\approx 0.8 \mu\text{s}$

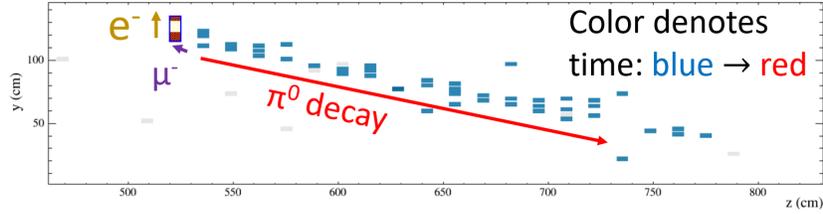


Figure 8.13: A simulated ν_μ CC event reconstructed in the ν_e CC sample due to a prominent electromagnetic shower. The muon energy was so low it did not leave a hit in the detector. The majority of energy goes into a π^0 decay which here looks like a single electromagnetic shower. Though this slice is difficult to reject, the time-delayed Michel cluster shown by the red hits gives an indication that this event is likely a true ν_μ CC while highlighting the need to search for Michels throughout the slice.

for simulation and $\approx 1.1 \mu s$ for data[99].

Upon identifying a set of Michel candidate hits for all slices, they are clustered together using DBSCAN[86], as in the Slicer4D algorithm. The distance score is given by the difference in plane number, ΔP , cell number, ΔC , and time, Δt as

$$d = \sqrt{\left(\frac{\Delta P}{\sigma_P}\right)^2 + \left(\frac{\Delta C}{\sigma_C}\right)^2 + \left(\frac{\Delta t}{\sigma_t}\right)^2}, \quad (8.10)$$

where $\sigma_P = 1.85$, $\sigma_C = 1.4$, and $\sigma_t = 40$ ns are tunable parameters selected by optimizing the Michel reconstruction efficiency and purity. The DBSCAN algorithm uses $\epsilon = 5$ and minimum $N_{pts} = 1$, allowing for Michel clusters with as few hits as one. These are relatively loose selection cuts which allows other background time-delayed physics activity, such as neutron captures, to be reconstructed. To reduce non-Michel noise, Michel clusters are required to have at least 0.01 GeV of calorimetric energy and have a time delay of at least $0.8 \mu s$. The time delay cut reduces the selection efficiency by 30%, while the energy cut removes very few true Michels. The Michel electron is then associated with the physics slice that is nearest in time and the minimum distance between a hit in the Michel cluster and a hit in the physics slice. The reconstructed Michel cluster is saved as a TrkME if it lies within 20 cm of the endpoint of some reconstructed Kalman track in the event and as a SlcME if it found elsewhere in the event.

Since Michel electrons are not the only time-delayed physics process associated with neutrino interactions and MEFinder reconstructs Michels found near any hit in the event, there is a substantial non-Michel background reconstructed. This is predominantly composed of neutron capture events which have a lifetime of

$53.6 \pm 0.7 \mu\text{s}$ in NOvA[100]. Fortunately, neutron captures have a very different energy spectrum compared to Michel electrons so that the two can be separated on an event-by-event basis. This neutron capture background has been serendipitously helpful as a test of our neutron simulation and as a constraint on the NC background in the ν_μ CC selected sample[100].

The Michel IDentifier (MID) is a $\Delta \log \mathcal{L}$ PID variable designed to distinguish Michel clusters from reconstructed background. This PID uses four input variables: the Michel calorimetric energy, the number of hits, the minimum distance to the parent physics slice, and the time delay. The energy and number of hits are very correlated. The distance to the parent slice and time delay are also correlated due to the Δt cut for re-triggered hits based on the APD dead-time modeling. Therefore, the $\Delta \log \mathcal{L}$ is implemented as the product of two two-dimensional likelihoods. The signal and background templates histograms are shown in Figure 8.14.

The MID distributions in data and MC for all Michel clusters observed in the ν_e selected sample is shown in Figure 8.15. As expected, the non-Michel background is much more prominent for SlcME clusters. For both cases, the simulation does quite well predicting the MID shape everywhere that Michels dominate the spectrum.

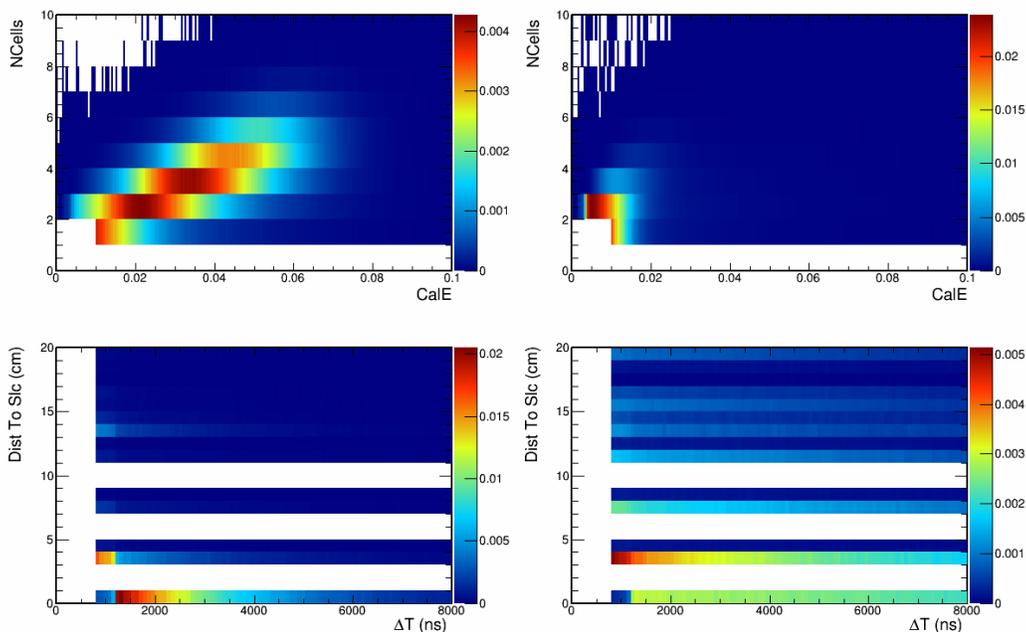


Figure 8.14: The MID input histograms for Michel electrons, left, and background, right. The top row shows the calorimetric energy vs number of cluster hits. The bottom shows the distance to the parent slice vs the time delay.

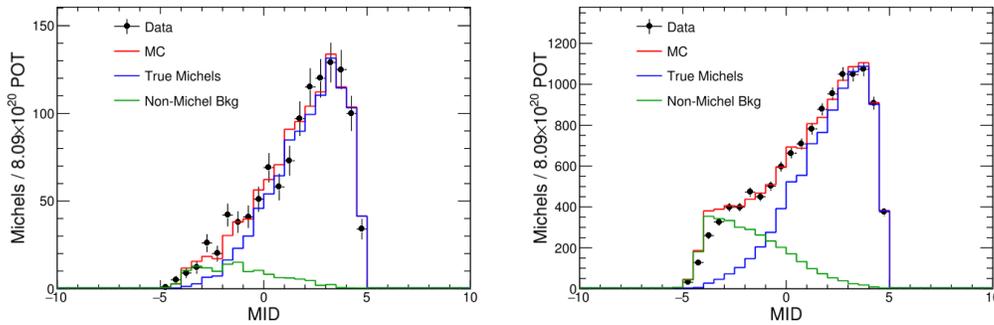


Figure 8.15: The MID distributions in data and MC for Michel clusters reconstructed in the near detector ν_e selected sample. TrkME's are shown on the left with SlcME's on the right. We see good agreement in the MID region where true Michels dominate.

8.9 An Alternative Clustering Algorithm: TDSlicer

Though Slicer4D, described above, has proven very successful and has been the base reconstruction for all NOvA analyses thus far, a new slicer has recently become available, the Time Density Slicer (TDSlicer). As the name suggests, the clustering algorithm at the heart of TDSlicer involves finding local maxima in the density of time-of-flight corrected hit times.

The new slicer was developed to mitigate a few failure modes common in Slicer4D. First, the slicer is very sensitive to pileup effects in the near detector. An attenuation of 12% is observed during 700 kW beam running in neutrino mode. Also, the metric Slicer4D uses internally does not satisfy the triangle inequality. Since each hit in the detector only measures xzt or yzt , the full 4D Euclidean distance is not defined, and thus Slicer4D only uses zt in comparing distances between hits in planes of opposite views. This can lead to physics clusters that are separated in xy coordinates, but sharing similar zt phase space being sliced together. The last involves including noise hits that are on the periphery of the definition of a neighborhood that can contribute to a neutrino slice failing containment cuts and bias the energy estimate. TDSlicer has been shown to more aggressively trim these excess hits from its reconstructed clusters.

TDSlicer Algorithm

There are three steps involved in generating TDSlicer clusters. First, centroids are identified by finding local maxima in the density of the hits. Then, a group two sets of 3D clusters in xzt and yzt are made by grouping nearby hits with these centroids.

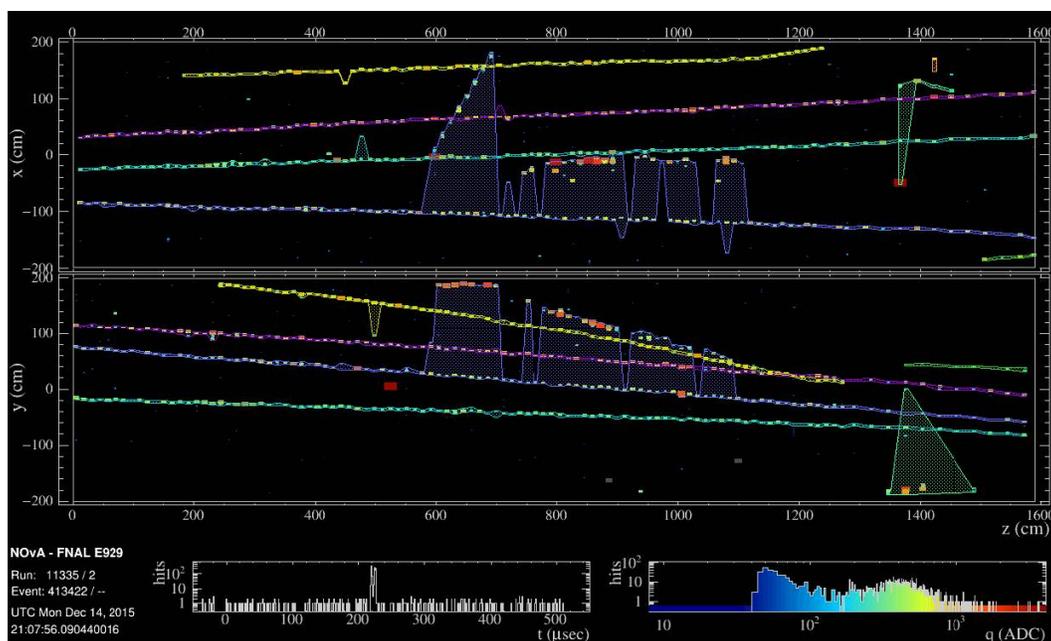


Figure 8.16: A simulated near detector spill with slices constructed by Slicer4D drawn. In this example, there are three independent physics interactions sliced together.

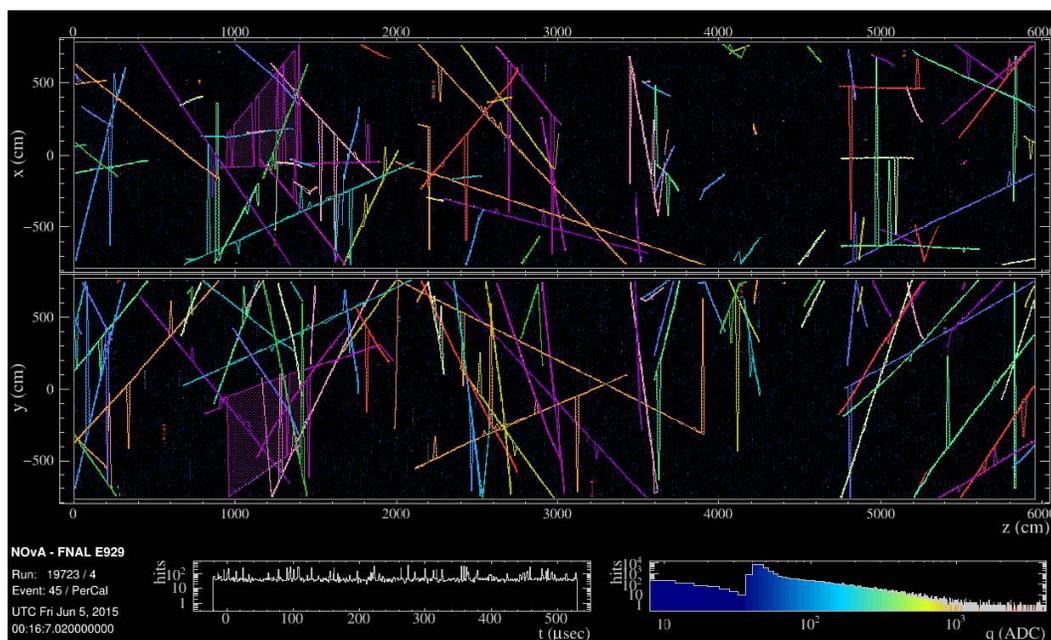


Figure 8.17: A simulated far detector spill, dominated by cosmic overlaid data, with slices constructed by Slicer4D drawn. The sudden spikes in the slice boundaries come from noise hits sliced in with the physics. There is a simulated neutrino that is sliced together with a cosmic in purple in the 900-1900 cm range of z .

Then, the 3D cluster groups are merged into a set of 4D clusters in $xyzt$. To identify the centroids, hits are clustered in each view separately, to ensure the metric satisfies the triangle inequality, by finding maxima in the time density distribution. A centroid-finding algorithm is applied as was described by Rodriguez and Liao[101]. As prescribed, first the density of points around each point is calculated as

$$\rho_i = \sum_j \exp\left(-d_{ij}^2/\tau\right) \quad (8.11)$$

where d_{ij} gives the Euclidean distance between two points in the space and τ is a configurable scale parameter. In the near (far) detector, τ is set to 16(80) ns, directly related to the timing resolution for hits. For NOvA, d_{ij} is defined as

$$d_{ij} = \frac{||dt| - R/c|}{\tau} \quad (8.12)$$

with dt giving the time difference between the hits, R gives the spatial separation.

After evaluating the density for every hit, an isolation score is evaluated according to

$$\delta_i = \min_{j|\rho_j > \rho_i} (d_{ij}) \quad (8.13)$$

which is just Euclidean distance to the closest point in the space with a higher density. To make this well defined, a maximum hit isolation is set to a number greater than the maximum possible isolation between two hits in a spill. Cluster centroids are identified with hits of high density and isolation as is shown in Figure 8.18.

To finish the time density slicing, centroid hits are found in each view using cuts on $\rho > 3(10)$ and $\delta > 8(6)$ in the near(far) detector.

After identifying cluster centroids, 3D clusters in xzt and yzt are created using Prim's algorithm[102], a minimally connected spanning tree. From the centroid, this builds a cluster by successively adding the hit outside the cluster that is nearest to some hit inside the cluster. There is a cut on the minimum distance of $d_{ij} = 8(5)\tau$ for the near(far) detector beyond which points will not be clustered.

Once the 3D clusters have been made in each view, a simple view-merging step is run. This compares the average zt values for all possible pairs of 2D slices in

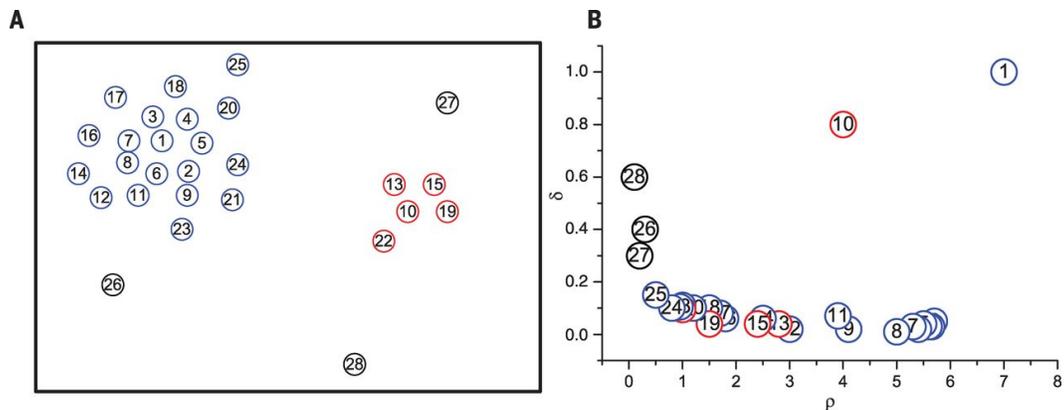


Figure 8.18: The left shows a scatter plot of potential data to be clustered. A density is then calculated for every point in the dataset. The numbers correspond to the density-orderings of the data points. Then, an isolation is computed, which is the minimum distance to a point of higher density. Though points 2-9 have a higher density than point 10, they are quite close to point 1, and thus score low in isolation. Point 10 is the most dense point in the smaller cluster. Points 1 and 10 are clearly identified as centroids in this example featuring high density and isolation scores. Points 26, 27, and 28 have too great an isolation to be included with any cluster and too low a density to form their own cluster and are thus treated as noise. Figure taken from Rodriguez and Laio[101].

opposite views. The length scale in z is 200(60) cm and the time scale is 6(60) ns in the near(far) detector. Iteratively, the best pair of 3D slices is combined to make a 4D slice and then deleted from the list of available 3D slices to pair. Those remaining slices that are more than 6(15) units away in the near(far) detector from their nearest-matched slice in the opposite view are counted as noise.

The two spills whose Slicer4D-determined clusters that were shown in Figs. 8.16-8.17 are shown in Figs. 8.19-8.20 with TDSlicer's clusters.

A Note on Properly Splitting Distinct Physics Slices

As the DBSCAN algorithm used by Slicer4D is constructed, its two large physics slices that are clearly separate may be joined by the presence of “bridge” hits where, though the bulk of the two populations are separated, there may be a few outlier hits from each distribution lying between the two cores. If there were enough hits caused by the two physics slices that the density of outliers is sufficient for them to be propagated to the DBSCAN algorithm, they will cause the two clearly distinct populations to be clustered together. TDSlicer ambitiously tries to split these two interactions at the cost of improperly grouping some fraction of bridge hits. In NOvA, bridge hits are dominated by hits with poor timing resolution, and thus low

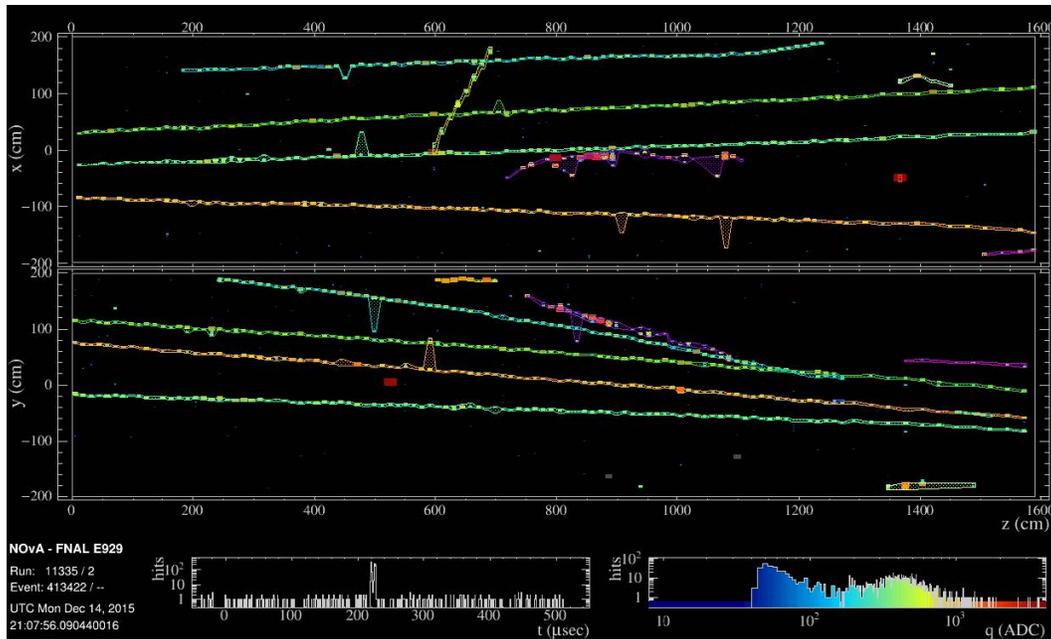


Figure 8.19: The simulated far detector spill shown in Figure 8.16, but with TDSlicer output drawn. The three distinct physics interactions that Slicer4D merged are now clustered separately.

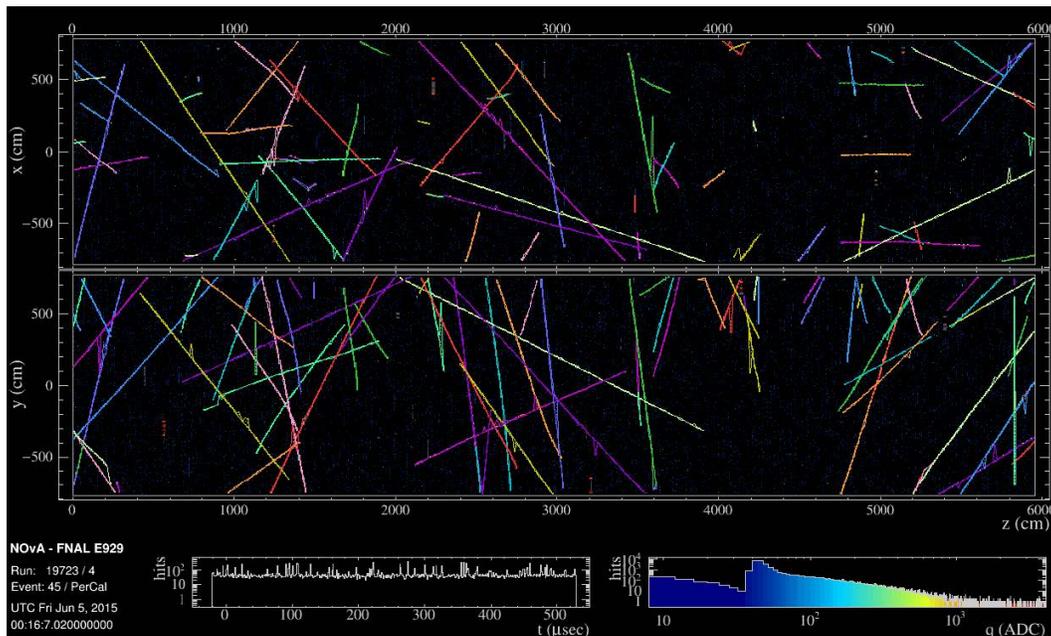


Figure 8.20: The simulated far detector spill shown in Figure 8.17, but with TDSlicer output drawn. There are far fewer noise hits included with the slice. Also, the neutrino that Slicer4D merged with a cosmic is now clustered on its own and clearly contained.

energy, so that this mis-grouping will not significantly impact the reconstruction or energy estimation of the event. An illustration of this is shown in Figure 8.21 for the event shown in Figs. 8.16 and 8.19.

TDSlicer Performance

The TDSlicer algorithm was designed specifically to make the near detector less susceptible to a loss of events at high beam intensity. The performance of the algorithm was tested against our nominal slicing algorithm, Slicer4D. The main metric involves the number of “good” slices reconstructed. These are slices with at least 90% of the slice’s calorimetric energy coming from the dominant interaction in the slice and that contain at least 90% of the light deposited in the detector by that interaction.

In the near detector, it is particularly interesting to look at the effect pileup has on

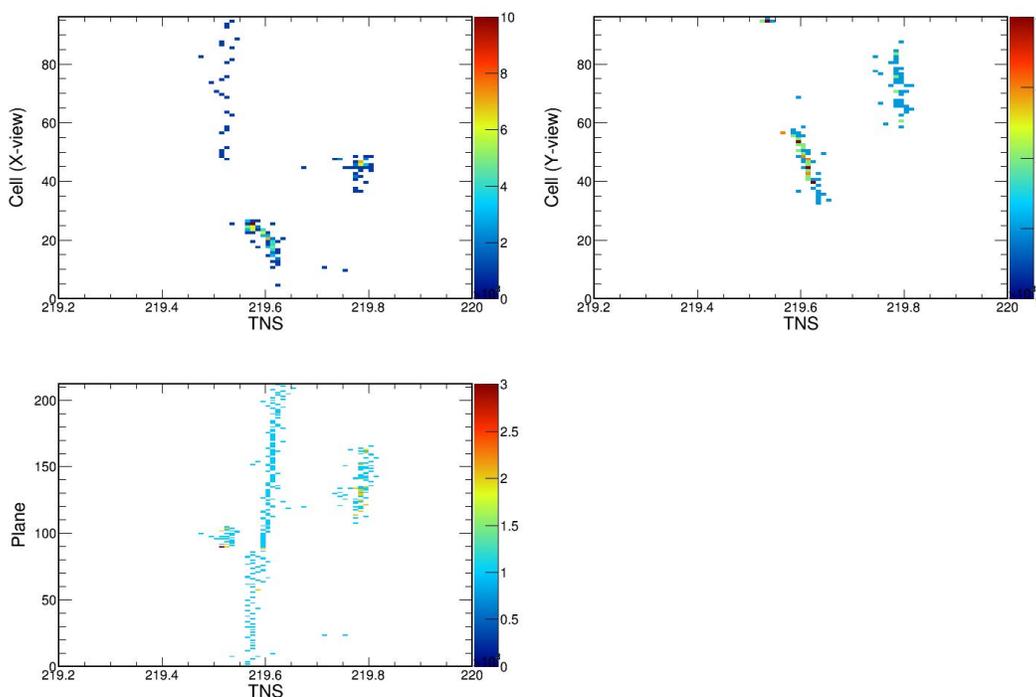


Figure 8.21: Time traces for the three merged physics interactions shown in Figure 8.16. The top plots show the cell number for X(Y)-View planes on the left(right) as a function of the time in μs while the bottom shows the plane number vs time. Though the three interactions are clearly separable by eye, there is a spattering of ambiguous hits in each of the separating regions. TDSlicer splits the core, separable pieces of these interactions at the expense of potentially improperly clustering these ambiguous hits.

slicing efficiency. Plotting the average number of reconstructed good slices in a spill vs the spill intensity shows a clear linear attenuation due to pileup as shown in Figure 8.22. Extrapolating these linear fits to an intensity near 0, Slicer4D produces nearly 2% more slices than TDSlicer. But the slope of the linear fit, representing slice loss from pileup, for Slicer4D is 250% of the analogous loss from TDSlicer so that TDSlicer produces more good slices over the entire range of data taking intensities. Currently, the NuMI beam intensity is averaging about 50×10^{12} POT/spill. At these intensities, TDSlicer produces 5.5% more good slices than Slicer4D.

In the far detector, the performance metrics are different as neutrino pileup is effectively impossible but neutrinos can be merged with cosmic rays. A collision with a cosmic ray is highly likely to make the event fail containment cuts. Thus, a more relevant measure is the fraction of simulated neutrinos that pass containment cuts when simulated including cosmic data overlay. TDSlicer scores well in this metric, too, as it reconstructs roughly 5% more contained ν_e CC and ν_μ CC slices compared to Slicer4D.

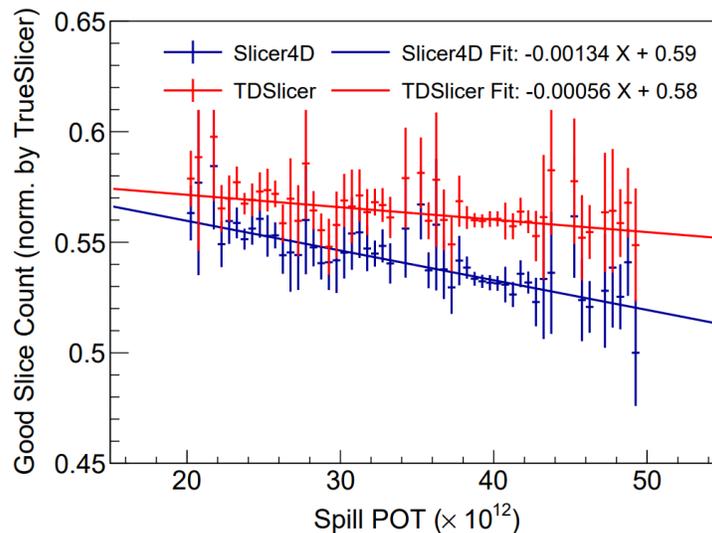


Figure 8.22: The number of good slices, with at least 90% purity and completeness, reconstructed by TDSlicer, red, and Slicer4D, blue as a function of the spill intensity. TDSlicer is significantly more resilient to event pileup.

Chapter 9

CC π^0 CROSS SECTION ANALYSIS

This chapter details a neutrino-nucleus cross section analysis developed concurrent with the oscillation analysis. This measurement and the oscillation analysis are independent, both relying on the previous chapters. For following the oscillation analyses, the central theme of this thesis, the reader may immediately continue to Chapter 10.

There is a long history of interest in pion production measurements in neutrino experiments. There are a number of measurements at bubble chambers aimed at comparing the cross section ratios for resonant pion and nucleon final states in both neutrino and anti-neutrino interactions. These were motivated by testing the isospin relations predicted by Δ_{1232} decay, as described in Section 1.4. There has been a resurgence of interest in the subject coincident with exploration of neutrino oscillations, as understanding neutrino interactions is important for accurately predicting selection and reconstruction efficiencies for various event topologies. The majority of recent work has focused on the transition region between QE and DIS scattering in the neutrino energy range between ~ 0.5 and 5 GeV.

In particular, π^0 production is of importance for experiments studying $\nu_\mu \rightarrow \nu_e$ flavor oscillations, as a π^0 decays electromagnetically and can be mistaken for the oscillation signal, ν_e CC scattering. This scattering channel was measured in bubble chambers[103][104] and more recently at MiniBooNE[105], K2K[106], and MINERvA[107]. Within this chapter, we present a measurement of the $\nu_\mu + A \rightarrow \mu\pi^0 X$, or CC π^0 , cross section, where X is any hadronic fragment including states with additional mesons. We report the results as flux-averaged differential cross sections in terms of the final state lepton and pion kinematics along with Q^2 and W . The dataset uses 3.72×10^{20} POT of exposure, only $\sim 10\%$ of planned NOvA running, and is dominated by systematic uncertainty. For this analysis, near detector data was blinded from the analysis until reconstruction and analysis techniques had been sufficiently studied with fluctuated and systematically distorted MC.

The base event generator used in the simulation is GENIE2.10.2 with particle propagation in GEANT4.9.6.

9.1 Definition of Signal

$CC\pi^0$ events are defined as any ν_μ CC interaction with a final state π^0 , of any energy, determined after final state interactions within the target nucleus. The analysis signal is further restricted to a region of kinematic phase space that we can reliably reconstruct in the near detector. Selected $CC\pi^0$ events that lie outside of this kinematic region will be treated as a background and subtracted from the observed data, as described below.

There is a cut on reconstructed $E_\nu > 1$ GeV because of a very limited event rate. The NuMI beam decays rapidly below the beam peak at 1.8 GeV while the cross section for π^0 production grows with increasing E_ν . Combining these two effects, only 0.2% of the $CC\pi^0$ interactions produced in the detector are expected to have $E_\nu < 1$ GeV. A high cut of $E_\nu < 5$ GeV was also determined. Though the analysis would otherwise select an appreciable number of events above this energy, these events are more susceptible to systematics. The NuMI flux is much less precisely determined at these high energies, so that including these high energy events increases the overall uncertainty on the integrated flux, the dominant systematic uncertainty in the analysis.

The size of the near detector suggests a maximum analyzable p_μ . A calculation using MIP dE/dx shows the near detector is 4.7 GeV long in the z direction. This number accounts for both the fully-active and muon catcher portions of the detector. This analysis will therefore be blind to events with higher p_μ . The maximum p_μ is lowered to 4 GeV/c to allow for a sizable fiducial volume for containable events at these momenta. As shown in Fig 9.1, there are indeed few events above $p_\mu = 4$ GeV/c.

For the π^0 kinematics, the detector is large enough to contain electromagnetic showers as the length of a shower grows logarithmically. A cut of $p_\pi < 3$ GeV/c is applied as there are too few events outside this region to make a measurement with sufficient statistics, as shown in Figure 9.1. Extrapolating the tail with an exponential predicts 0.008% of otherwise selected signal events in the $p_\pi > 3$ GeV/c region at the data exposure. A negligible portion of the selected sample passing this cut on reconstructed p_π has true energy above this.

Additionally, the four-momentum transfer to the hadronic system, Q^2 , is restricted

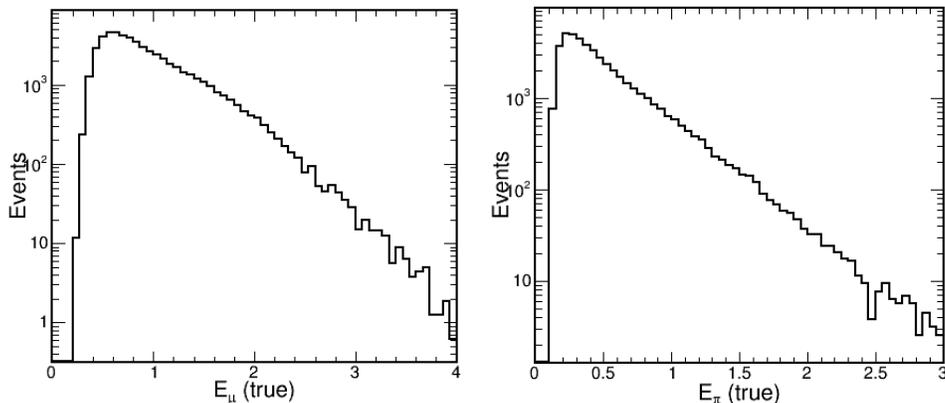


Figure 9.1: The distribution of true E_π , right, and E_μ , left. There are only a handful of events predicted with $p_\pi > 3 \text{ GeV}/c$ or $p_\mu > 3 \text{ GeV}/c$. As such, the analysis treats events in these overflow ranges as background.

with $Q^2 < 4 \text{ GeV}^2/c^2$. Additionally, the invariant mass of the hadronic system, W , is required to have $1 < W < 3 \text{ GeV}/c^2$. These are both required due to negligible event rates outside the bounds. There is very little spillover of events whose true kinematics lie outside of this range that reconstructs into the sample.

In summary, the signal definition for this analysis is further constrained by requiring: $1 < E_\nu < 5 \text{ GeV}$, $p_\pi < 3 \text{ GeV}/c$, $p_\mu < 4 \text{ GeV}/c$, $Q^2 < 4 \text{ GeV}^2/c^2$, and $1 < W < 3 \text{ GeV}/c^2$. $CC\pi^0$ events whose true kinematics are outside this region are treated as a background. Further, events whose reconstructed estimates for these variables fall outside this range are cut from the analysis.

9.2 Event Reconstruction

The $CC\pi^0$ selection is based on identifying π^0 decay photons within ν_μ CC events. This analysis developed specific tools to identify photon candidates within events. We will require selected events to have at least two prongs: one representing the muon and one representing a π^0 decay photon. After running the reconstruction described in chapter 8, the next step is to select which prong is associated with the muon. This is done by selecting the prong most colinear with the highest ReMID track. For signal events, the muon track and prong are within 12° in over 96% of events while fewer than 1% of simulated prongs from other particles are this colinear with the muon track.

A $\Delta \log \mathcal{L}_\gamma$ score between photon and non-photon hypotheses was developed and calculated for every reconstructed prong in each selected event. This identifier relies on four inputs:

- The Bragg Peak Identifier (BPI) is the ratio of the energy deposited in the furthest $\min(6, N_{\text{hit}}/2)$ hits from the prong start to the mean energy deposited on the rest of the prong, giving a measure of the characteristic increase in dE/dx of tracks at low momentum. As a ratio, this variable is insensitive to several systematics such as the azimuthal angular distribution for simulated particles and calibration uncertainties
- The Energy Per Hit (EPH) which is a measure of particle dE/dx
- The gap between the reconstructed vertex and the prong start
- The Number of Missing Planes (NMP) gives the largest number of consecutive planes without any deposited energy, an indication of large gaps in the prong without any energy deposited

Each of these input variables is shown in Figure 9.2. The photon candidate is taken as the prong with the highest $\Delta \log \mathcal{L}_\gamma$, excluding the prong that was tagged as the muon. This affords a slice-level PID, $\text{CC}\pi^0\text{ID}$, that is the $\Delta \log \mathcal{L}_\gamma$ of the photon candidate.

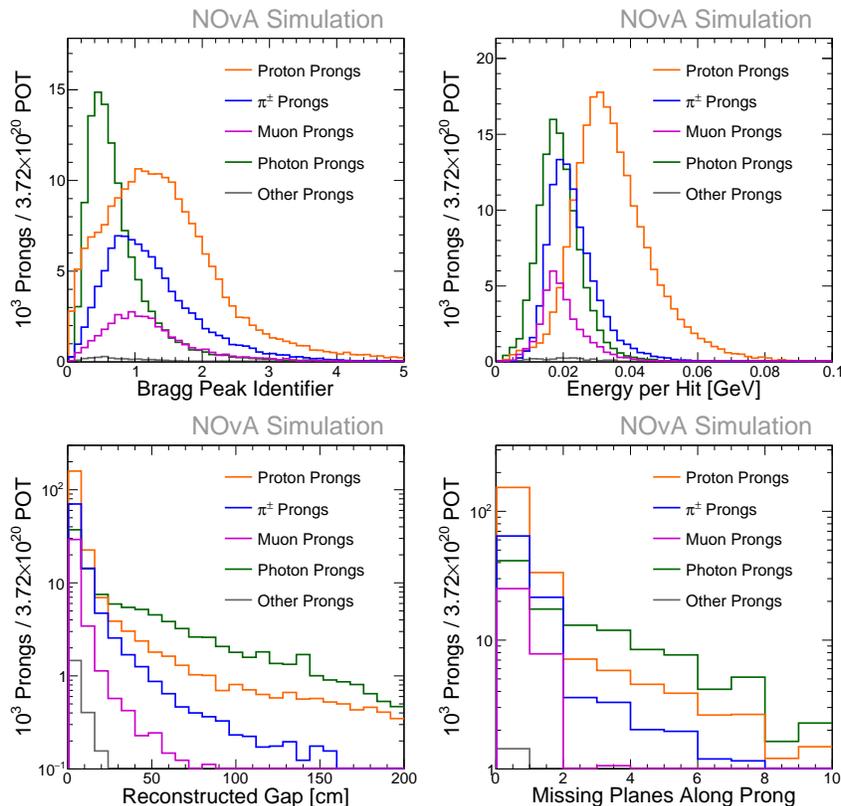


Figure 9.2: The four variables used as input to the $\Delta \log \mathcal{L}_\gamma$ selector. Photon prongs are drawn in green showing separation power in each variable.

9.3 Kinematic Estimation

The analysis delivers flux-averaged cross sections, differential in six variables: p_μ , $\cos \theta_\mu$, p_π , $\cos \theta_\pi$, Q^2 , and W . We need an automated way to reconstruct each of these variables from the reconstructed kinematics of the event. Both $\cos \theta_\mu$ and $\cos \theta_\pi$ are taken from the reconstructed muon and π^0 prong direction. The muon momentum is a linear function of the track length through the Fully Instrumented (FI) portion of the detector and the Muon Catcher (MC) as

$$\hat{p}_\pm = A \times L_{\text{FI}} + B \times L_{\text{MC}}, \quad (9.1)$$

where A and B are constants determined from simulation. The muon momentum is estimated with a resolution of 3.5% above 0.6 GeV/c. The π^0 momentum is estimated from the calorimetric energy of the photon candidate prong with a scale factor to account for deposited energy in inactive material. The absolute uncertainty on p_π is greatest between 1 and 1.5 GeV/c. Above this momentum, however, the opening angle is too slight to have both photons from the π^0 decay reconstructed separately. Thus, all the energy from the π^0 is deposited on the single selected

photon prong. The resolution of the muon and π^0 momentum and azimuthal angle are shown in Figure 9.3.

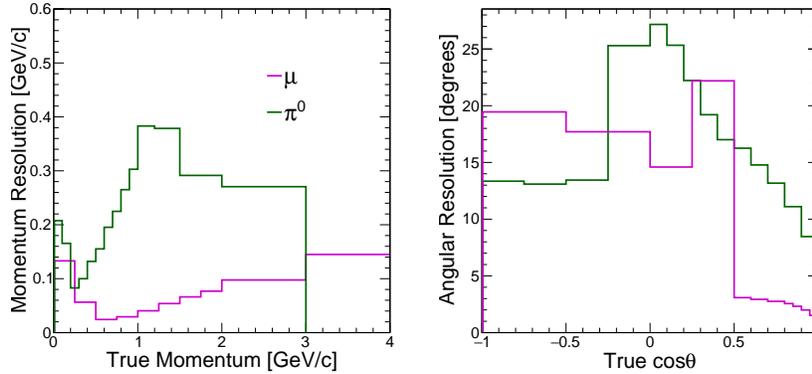


Figure 9.3: The absolute kinematic resolutions on the π^0 , green, and the muon, purple. The left plot shows the momentum resolution while the right shows the angular resolution. Counter-intuitively, the absolute resolution improves for π^0 's with more than 1 GeV of momentum. This comes from kinematic constraints which force the decay photons to be nearly parallel at high momentum so that a high fraction of the π^0 energy is reconstructed on the single selected prong.

An event energy is reconstructed as

$$\hat{E}_\nu = \hat{E}_\mu + \hat{E}_{\text{Had}}, \quad (9.2)$$

where \hat{E}_μ is reconstructed as described above. \hat{E}_{Had} is estimated by a spline fit to the calorimetric energy of the hadronic shower. When averaged over the sample, this estimator gives a 9.5% energy resolution.

The variables Q^2 , the four-momentum transfer to the hadronic shower, and W , the invariant mass of the hadronic shower, are important kinematic variables though not directly observable. They rely on the reconstructed energy and muon kinematics as

$$\hat{Q}^2 = -(p_\mu - p_\nu)^2 = \hat{E}_\nu \hat{E}_\mu (1 - \cos \hat{\theta}_{\mu\nu}) + m_\mu^2, \quad (9.3)$$

$$\hat{W} = |p_N + p_\mu - p_\nu| = \sqrt{m_N^2 - \hat{Q}^2 + 2m_N (\hat{E}_\nu - \hat{E}_\mu)}. \quad (9.4)$$

The presented cross sections are unfolded to Q^2 and W , calculated in this way from the true final state kinematics in lieu of GENIE's prediction for these variables.

9.4 Event Selection

The selection for the $CC\pi^0$ analysis is separated into five classes: spill and data quality cuts, fiducial cuts, containment, background rejection with CVN, and kinematic restrictions. The constituent cuts in each of these five classes are listed in detail in the following.

Data Quality and Spill Cuts

The first step in the selection is to identify which NuMI spills are suitable for analysis. These are referred to as spill cuts, and ensure that the beam was running stably and the detector electronics were functioning properly. The spill cuts require:

- Fraction of hits occurring outside the beam peak < 0.45
- Number of missing DCMs = 0
- Time delay between NuMI timestamp and spill < 0.5 ns
- Spill POT $> 2 \times 10^{12}$
- $-202 < \text{Horn current} < -198$
- $-2.0 < \text{Horn } x \text{ and } y \text{ position on target} < 2.0$
- $0.57 < \text{Beam width in } x \text{ and } y < 1.58$

Additionally, there are data quality cuts applied which primarily remove low energy slices that are difficult to reconstruct:

- At least one 3D kalman track with a reconstructed energy
- $N_{\text{hit}} > 20$
- Longest stretch of contiguously hit planes > 4
- At least one reconstructed cosmic track
- At least two reconstructed prongs
- Hits on the photon candidate prong ≥ 10
- Calorimetric energy on the photon candidate prong ≥ 0.1 GeV

Fiducial Cuts

After applying basic quality cuts, there is a cut on the reconstructed vertex of the event.

- $-100 < x$ coordinate of vertex < 100 cm
- $-100 < y$ coordinate of vertex < 100 cm
- $100 < z$ coordinate of vertex < 1000 cm

giving a fiducial region with a mass of 35 tons. There is both a leakage of true fiducial events that are reconstructed outside of the fiducial volume and non-fiducial events that reconstruct within the fiducial volume. As these two currents are very similar, 2.8% for fiducial leakage out and 2.6% for non-fiducial leakage in, the selected non-fiducial events are treated as analysis signal. Roughly 0.04% of selected events are expected to be from interactions in the surrounding cavern that pass through and reconstruct within the fiducial volume.

Containment Cuts

Containment cuts are separated into primary μ and selected π^0 containment. This analysis uses the same μ containment cuts as the oscillation analyses.

- There are no hits in the first two planes of the detector
- There are no hits in the last two planes of the detector
- The z coordinate of the start of the μ track < 1150 cm
- The z coordinate of the stop of the μ track < 1270 cm **or** the distance traveled through the air-gap between active volume and muon catcher < 55 cm
- The calorimetric energy recorded in the muon catcher < 30 MeV
- The distance between the detector wall and the track endpoint projected forward along the track > 4 cm
- The distance between the detector wall and the track endpoint projected backward along the track > 8 cm

For the hadronic containment, there is a containment cut on the prong endpoint of the photon candidate, roughly corresponding to a veto of four cells. With either containment strategy, there will be a leakage of events into the sample whose true $E_\nu > 5$ GeV. As already discussed, these events are treated as a background.

- $-185 < x$ and y coordinate of photon endpoint < 185 cm
- $20 < z$ coordinate of photon endpoint < 1230 cm

Background Rejection with CVN

The $CC\pi^0$ analysis uses CVNm, the ν_μ CC PID determined by CVN, to reject the NC component of the selected sample. A cut of $CVNm > 0.5$ was determined to maximize the ν_μ CC selection efficiency while keeping the NC contamination in the selected sample negligible. Above this cut, the NC contamination is only 1.4% of the total sample. The CVNm distribution is shown Figure 9.4 at this stage of the cutflow. As shown, the simulation follows the data well.

In addition to discriminating between ν_μ CC and NC events, CVNm also calculates a PID for the most likely true process of a selected interaction: QE, Res, DIS, or other. The other category is principally populated by coherent events which can not produce a primary π^0 , and quasi-elastics can only through intra-nuclear scattering.

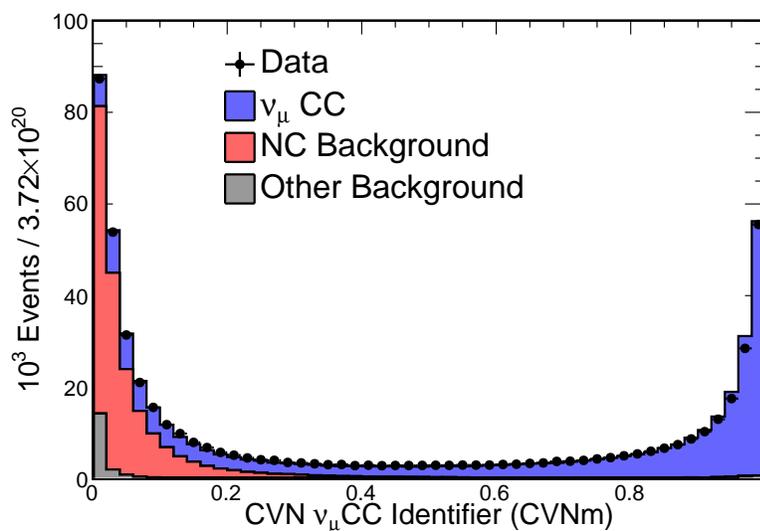


Figure 9.4: The CVNm, the NC-rejecting PID used in this analysis, distribution in data and MC. The plot shows good agreement. The ν_μ CC is shown in blue, and is a superset of $CC\pi^0$ signal.

Due to unsatisfactory modeling of MEC events, CVN is not trained to identify these separately. Rather, simulated MEC events are usually classified as QE. This is based on the event topology. CC events with only a charged lepton in the final state along with some number of nucleons are tagged as QE while events with shower topologies that include mesons are classified as Res or DIS. The outgoing proton in a QE reaction can undergo a hadronic scatter within the nucleus, producing a π^0 in the final state and thus making the event a signal event. This is a potentially grave situation as it would force CVN to train on the kinematics of the event for deciding whether to tag an event as QE. But simulation predicts that, before this cut, only 0.4% of signal comes from QE events.

Figure 9.5 shows that the signal is largely in the Res-like and DIS-like CVNm types, with QE and Other types mostly background. The MC matched to QE in CVNm is 2.2% signal, Res matched is 26.8%, DIS matched is 45.2% signal, and Other matched is 2.8% signal. In addition to the CVNm > 0.5 cut, there is a cut on CVNm's best matched true process since the signal is effectively clustered in the Res-like and DIS-like bins. So, in summary we have the following cuts using the CVN PID:

- CVNm > 0.5
- Best CVNm match is ν_μ CC Res **or** ν_μ CC DIS

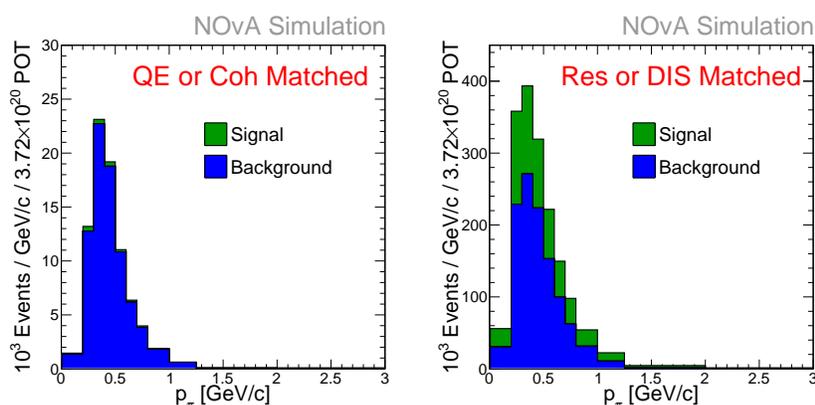


Figure 9.5: The reconstructed p_π distribution for signal, green, and background, blue. The left plot shows the distribution for CVN matches with QE or Coh hypotheses. These events are topologically distinct from signal events, which include a π^0 . Thus, almost all of the signal reconstructs as either Res or DIS according to CVN.

Kinematic Restrictions

As described earlier, the analysis signal is all $CC\pi^0$ events whose true kinematics lie within the region that can be robustly reconstructed and analyzed by NOvA. $CC\pi^0$ events lying outside of this true phase space are treated as a background. The final step of the selection aims to remove the selected $CC\pi^0$ outside the true kinematic region. This involves cutting on the analogous reconstructed variables at the same values used in the truth restriction. This gives

- $1 < E_\nu < 5 \text{ GeV}$
- $p_\mu < 4 \text{ GeV}/c$
- $p_\pi < 3 \text{ GeV}/c$
- $Q^2 < 4 \text{ GeV}^2/c^2$
- $1 < W < 3 \text{ GeV}/c^2$

The $CC\pi^0$ background is only 0.2% of the selected sample after this last restriction.

A table of the selected events in data and MC at each stage in the cutflow is shown in Table 9.1. In the 3.72×10^{20} POT analyzed, there are 165813 data events selected, which is 1.3% lower than the total MC, within the systematic error band described below. This gives a signal efficiency and purity of 21.2% and 34.8%, respectively. Figures 9.6 - 9.9 show the selection efficiency and purity as a function of several reconstructed variables.

| Cut | N_{evts} | N_{GENIE} | Efficiency | Purity |
|----------------------|---------------------|--------------------|------------|--------|
| Fiducial | 3.18×10^6 | 2.34×10^6 | 100% | 11.8% |
| Basic reconstruction | 2.192×10^6 | 1.44×10^6 | 85.1% | 16.3% |
| Containment | 517317 | 400797 | 32.8% | 22.6% |
| NC Rejection | 213376 | 197433 | 22.2% | 31.0% |
| QE/Coh Rejection | 197858 | 186779 | 22.0% | 32.7% |
| Prong Quality | 188158 | 175105 | 21.3% | 33.6% |
| Kin. Restriction | 165813 | 167935 | 21.2% | 34.8% |

Table 9.1: Analysis cutflow for cuts described earlier in the text. The event counts are as predicted by MC, scaled to the data exposure of 3.72×10^{20} . Signal efficiencies are calculated relative to interactions within the true fiducial volume.

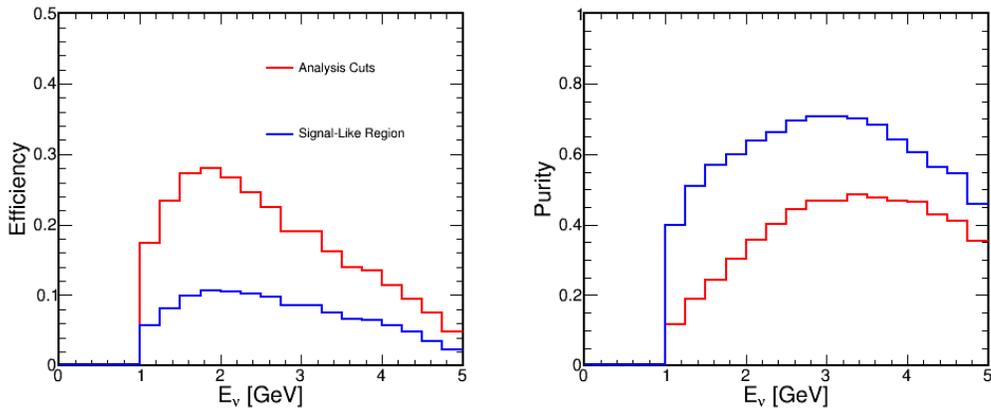


Figure 9.6: The efficiency, left, and purity, right, of the analysis cuts as a function of reconstructed neutrino energy. The red curves show the high-efficiency, low-purity cuts defining the analysis sample. The blue curves represent the region of the sample with an additional $CC\pi^0ID > 3$ cut. Though this cut is not applied in the analysis, it offers a convenient signal-rich subsample.

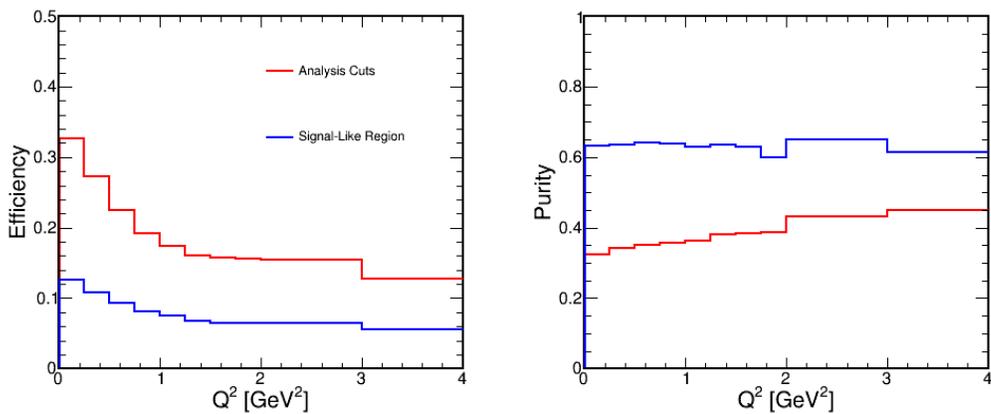


Figure 9.7: The efficiency, left, and purity, right, of the analysis cuts as a function of reconstructed Q^2 . The red curves show the high-efficiency, low-purity cuts defining the analysis sample. The blue curves represent the region of the sample with an additional $CC\pi^0ID > 3$ cut. Though this cut is not applied in the analysis, it offers a convenient signal-rich subsample.

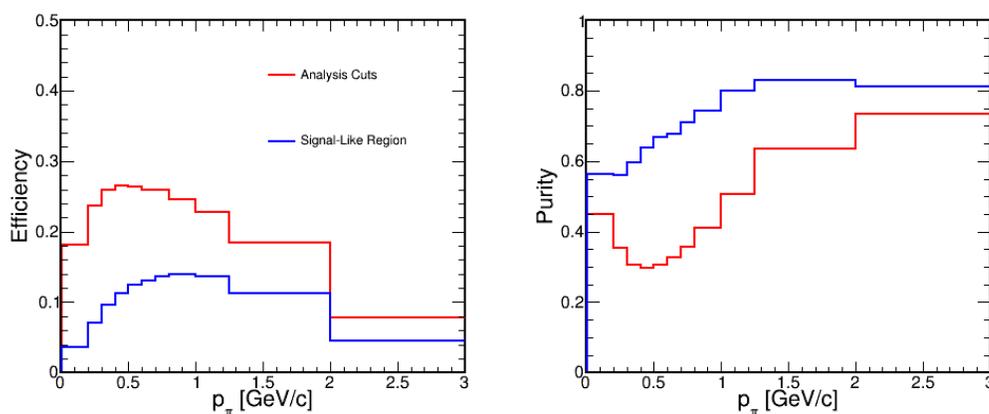


Figure 9.8: The efficiency, left, and purity, right, of the analysis cuts as a function of reconstructed p_π . The red curves show the high-efficiency, low-purity cuts defining the analysis sample. The blue curves represent the region of the sample with an additional $CC\pi^0ID > 3$ cut. Though this cut is not applied in the analysis, it offers a convenient signal-rich subsample.

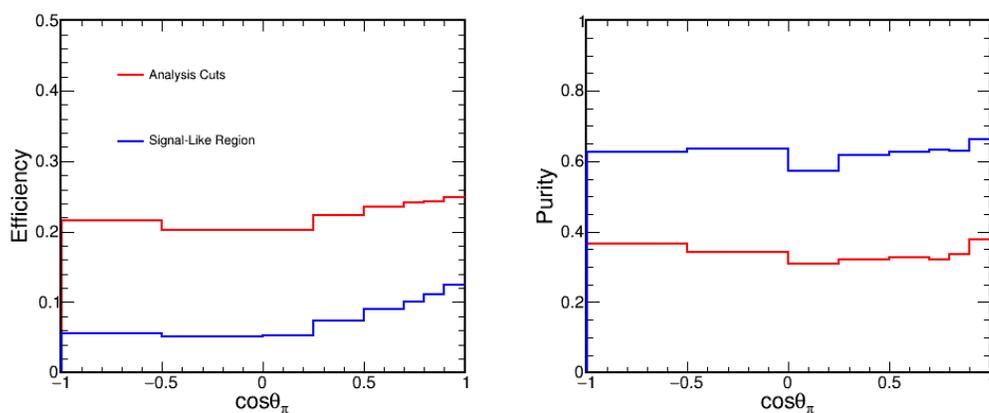


Figure 9.9: The efficiency, left, and purity, right, of the analysis cuts as a function of reconstructed $\cos\theta_\pi$. The red curves show the high-efficiency, low-purity cuts defining the analysis sample. The blue curves represent the region of the sample with an additional $CC\pi^0ID > 3$ cut. Though this cut is not applied in the analysis, it offers a convenient signal-rich subsample.

9.5 Background Decomposition

The analysis considers three classes of background events. First, there are selected events which are not ν_μ CC. These comprise only 3.2% of the sample, split between NC and $\bar{\nu}_\mu$ CC events. There are also $CC\pi^0$ events whose true kinematics lie outside of the outline analysis region, which only accounts for 0.2% of the sample. The remainder, ν_μ CC events without a π^0 in the final state are by far the largest component of the background. The $CC\pi^0$ ID distribution for signal and these three background classes is shown in Figure 9.10. The shape of this distribution is used to fit the signal and background in the observed data in order to constrain the estimated signal in each kinematic bin.

This third, largest, class can be subdivided further according to which particle produced the photon candidate prong. The particle flavor composition of these prongs is important to understand when assessing systematic uncertainties for the analysis. This distribution is shown in Figure 9.11 as a function of $CC\pi^0$ ID. Prongs from protons, 40.9%, and π^\pm , 33.0%, make up the majority while 10.7% come from muons and 13.4% come from photons.

The significant photon component is a concern as the $CC\pi^0$ ID distribution is similar to the analysis signal. Therefore there is no clear way to constrain this background component using data. The most common production mechanism for these selected

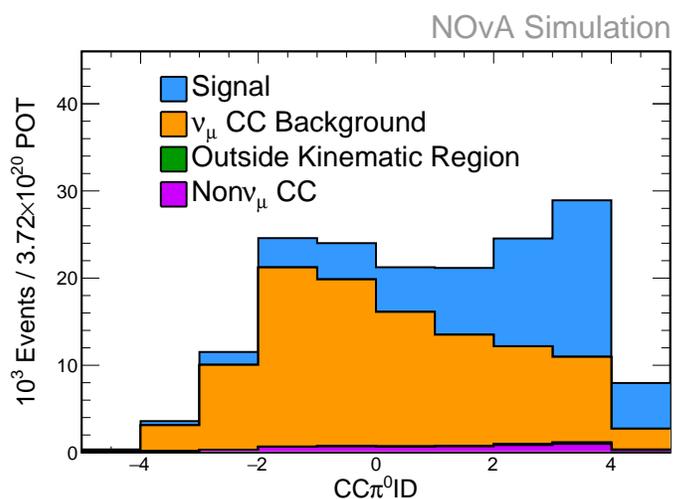


Figure 9.10: The $CC\pi^0$ ID distribution for all selected events showing clear discrimination between signal in blue and ν_μ CC background without a final state π^0 in orange. Apart from these two dominant interaction types, there are two small background components: events that are not ν_μ CC, purple, and $CC\pi^0$ events that fail the true kinematic restrictions, green.

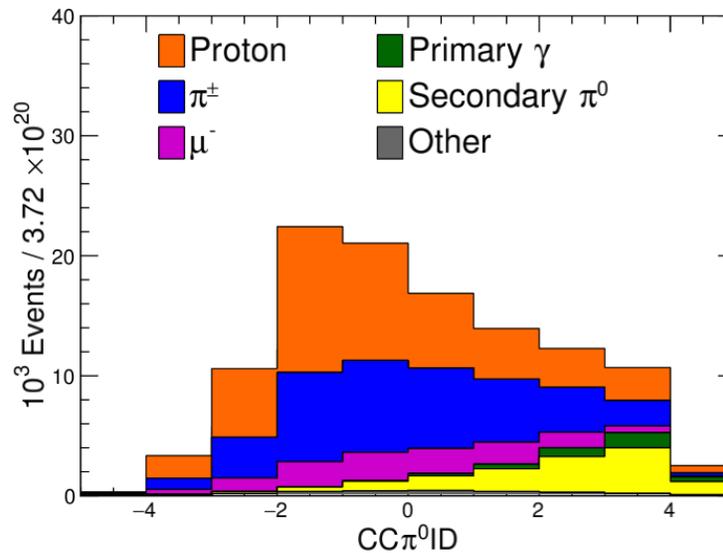


Figure 9.11: The selected ν_μ CC background without a final state π^0 decomposed by the particle type of the selected photon prong as a function of the $CC\pi^0ID$ selector. Proton and π^\pm prongs are most common with a notable pileup of photons at high values of $CC\pi^0ID$.

photons is from the $\pi^\pm \rightarrow \pi^0$ Charge eXchange (CX). This CX is also poorly measured. Therefore, as described below, we will assess a systematic uncertainty on the cross section analysis due to the uncertainties in the CX cross section.

9.6 Signal Estimation

As with any PID, the $CC\pi^0ID$ distribution separates signal $CC\pi^0$ from different background types as shown in Figure 9.10. The signal estimate is determined by fitting MC signal and background templates of the $CC\pi^0ID$ to data. As this is a differential measurement, the signal and background are constrained in each bin of the measured kinematic variable separately. Therefore, there is a separate fit in each bin of the kinematic variable to the observed $CC\pi^0ID$ spectrum.

The background is decomposed into non- ν_μ CC, $CC\pi^0$ whose final state kinematics fail signal definition requirements, and all other ν_μ CC background events. The non- ν_μ CC and the $CC\pi^0$ outside the true kinematic region are fixed to the MC prediction since they are such a small populations in our sample, at 3.2% and 0.2% respectively.

The remaining background template and the signal template, which, together, account for 95.6% of the sample, are fitted to the data or fake data. The signal and background normalization float independently and without a pull term. In each

fit, the $CC\pi^0ID$ variable is divided into ten equal bins from -5 to 5. Since the fit determines two normalization parameters, the fit has eight degrees of freedom. The signal estimate in any bin is the integral of the signal histogram deduced from this fit while the background content in each bin is also allowed to float as a nuisance parameter.

The fit is performed in each analysis bin separately, and completely independently of neighboring bins. As an example, the $CC\pi^0ID$ distribution is shown in Figure 9.12 for events with $0.5 < p_\pi < 0.6$ GeV before and after a fit to fake-data where the M_A^{Res} parameter in GENIE, shifting the axial mass in the resonant neutrino production model, has been shifted.

Apart from adjusting the total MC shape to match data, the fit also correctly predicts the true signal and background in this systematically shifted universe. The difference between the fitted signal estimate and the true signal in the systematic universe will be related to a systematic uncertainty on the analysis, as will be detailed below.

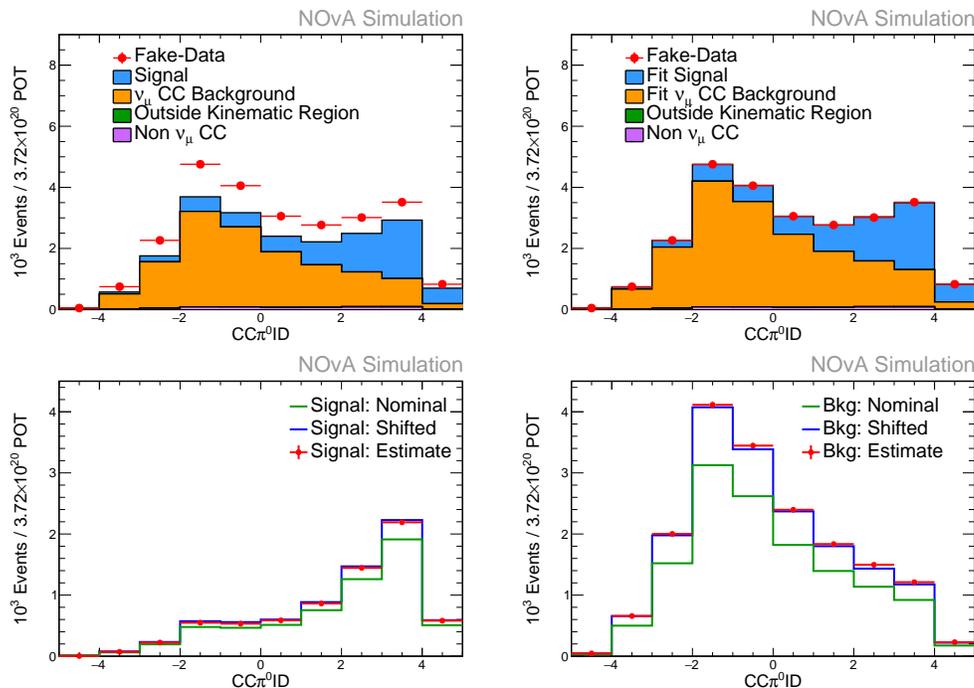


Figure 9.12: An illustrative example of the signal extraction via the fitting procedure. Shown are MC templates for events with $0.5 < p_\pi < 0.6$ GeV/c. The non- ν_μ CC and $CC\pi^0$ outside of the true kinematic region components are held fixed as they are so small while the other two are fit to the fake data. Here, data is fake data, determined by systematically shifting M_A^{Res} up in the MC by 1σ . The top left plot shows the nominal MC histograms with the shifted fake data separated into signal, blue, and backgrounds, orange, green, and violet. The top right plot shows the same fake data with the MC templates after fitting the MC components to the fake data. The estimated signal $CC\pi^0$ ID distribution is shown in red on the bottom for signal, left, and background, right, compared to the nominal distributions in blue and the systematically shifted fake data in red. The bottom plots demonstrate that this procedure properly estimates the signal and background normalizations in the fake data.

9.7 Unfolding

The reco-to-true migration matrices for each of the measured variables are shown in Figure 9.13. These plots are the input to an unfolding procedure. Unfolding is a model-dependent algorithm which deconvolves various detector and resolution effects. The output is the best estimate of the distribution in the true variable of interest.

Three unfolding methods were tested with fake-data in relation to this analysis. Tikhonov regularization[108], regularized by the second derivative predicted by simulation, suffered from violent features born from MC fluctuations and was quickly

dropped. D’Agostini[109] and SVD[110] unfolding were studied further.

D’Agostini is an iterative method where the signal estimate in reco space re-weights the migration matrix a set number of times and then is projected onto the true variable in the weighted migration matrix as shown in Figure 9.13. The number of times the signal estimate reweights the migration matrix is the regularization parameter.

SVD, Singular Value Decomposition, unfolds by finding the principle eigenvectors of the migration matrix. This unfolding technique is regularized by the number of eigenvalues considered in the decomposition. The unfolded result is a weighted sum of the N eigenvectors with largest eigenvalues. The weight of each eigenvector in the final result depends on the signal estimate. Eigenvectors with low eigenvalues tend to be noisy, with rapid oscillations. Thus, regularization can be thought of as setting a frequency cut-off.

Both of these methods were studied as a function of regularization strength. Figures 9.14-9.17 show how each of these unfolding methods behave for different regularization strengths when estimating the signal of systematically shifted fake-data. These four plots show the behavior for the four largest GENIE ν -interaction systematics shifts, M_A^{Res} , M_V^{Res} , NonRES1pi, and NonRes2pi. These four knobs adjust the axial and vector mass of resonant events and the pion production rates in DIS events. Fake-data is generated with a systematic shift applied, and the resulting unfolded distribution is compared to the GENIE prediction with the systematic shifted. Results of this procedure are shown in 9.14-9.17.

From Figs. 9.14-9.17, D’Agostini unfolding with a regularization parameter of 1 or 2 works, in that the ratio of the unfolded distribution to the fake-data is relatively flat near 1, indicating that the unfolded result approximately matches the true fake-data distribution. Similarly, SVD unfolding with regularization of 4 works relatively well. D’Agostini with a regularization of 1 gives the flattest ratio of unfolded to shifted prediction for each of these four systematics across the p_π range. Since these four GENIE systematics have a large effect on the true cross section, and the same unfolding procedure works well with all of them, this unfolding procedure was selected for use in the analysis. The residuals between the shifted true fake-data distributions and the unfolded distributions are the basis of a systematic calculation, discussed below.

As a cross-check of the unfolding procedure, an in-n-out test was performed; the

result is shown in Figure 9.18. The reconstructed and true p_π distributions were plotted. The reconstructed histogram was then pushed through the unfolding procedure. The result of the unfolding perfectly matches with the true histogram, giving a cross-check of the mechanics.

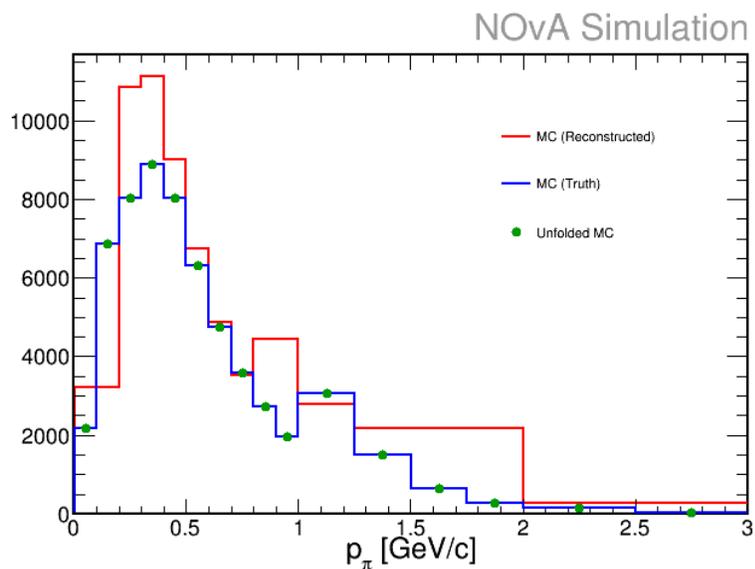


Figure 9.18: An In-N-Out test showing the simulated reconstructed distribution, red, and the true distribution, blue. The Unfolded reconstructed distribution is shown in green, and is identical to the MC truth distribution.

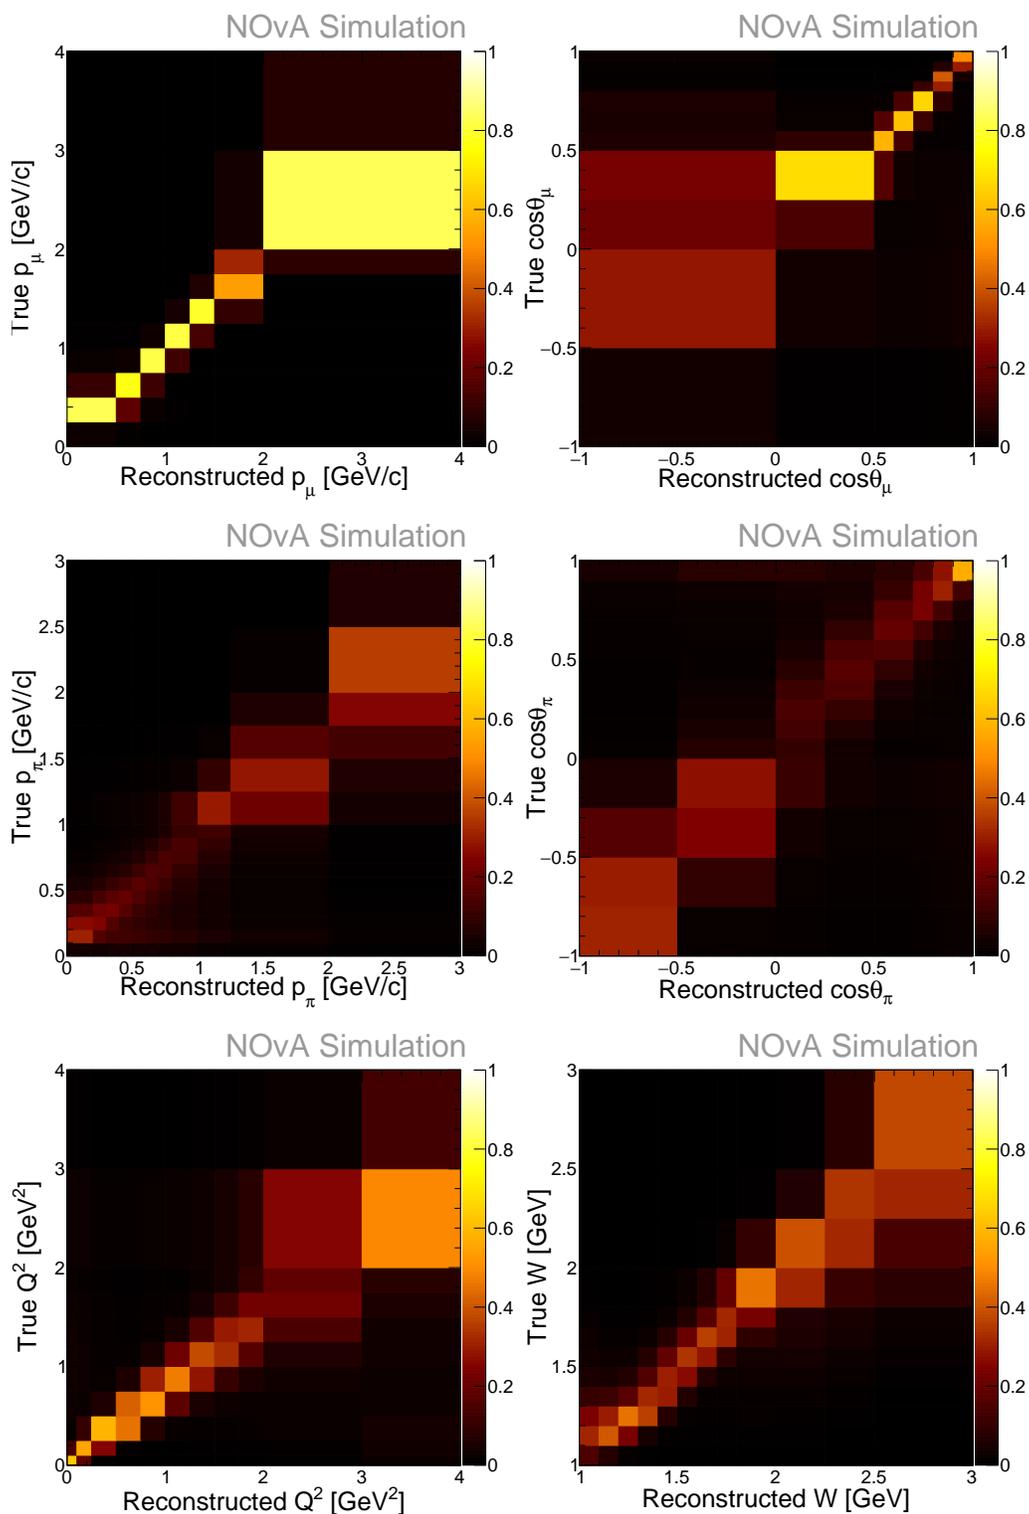


Figure 9.13: The migration matrices between reco and true space for each analyzed variable. All signal events that pass the analysis selection are represented in each plot.

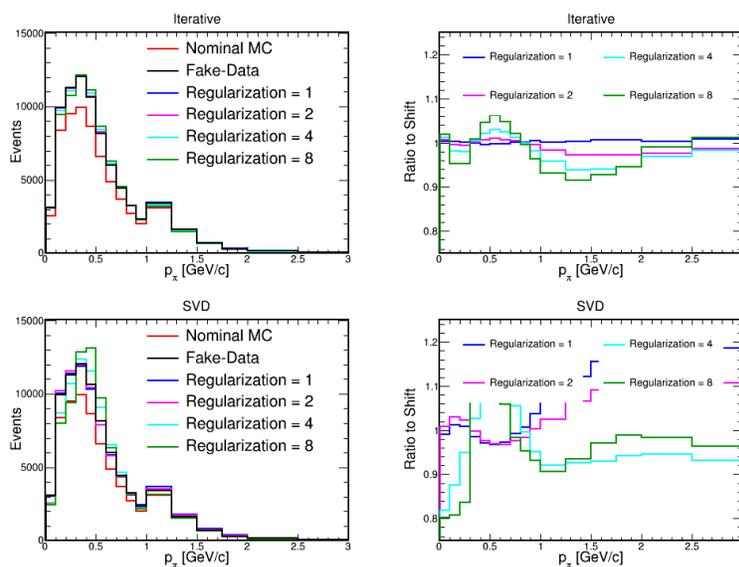


Figure 9.14: A comparison of how variously regularized unfolding procedures resolve systematic shape distortions. The left shows the raw unfolded distribution compared to the systematically shifted fake-data. The right plots show the unfolded ratio to fake-data. The systematic here shifts the axial mass for resonant production up by 1σ . Iterative unfolding with a regularization of two is used in the $CC\pi^0$ analysis.

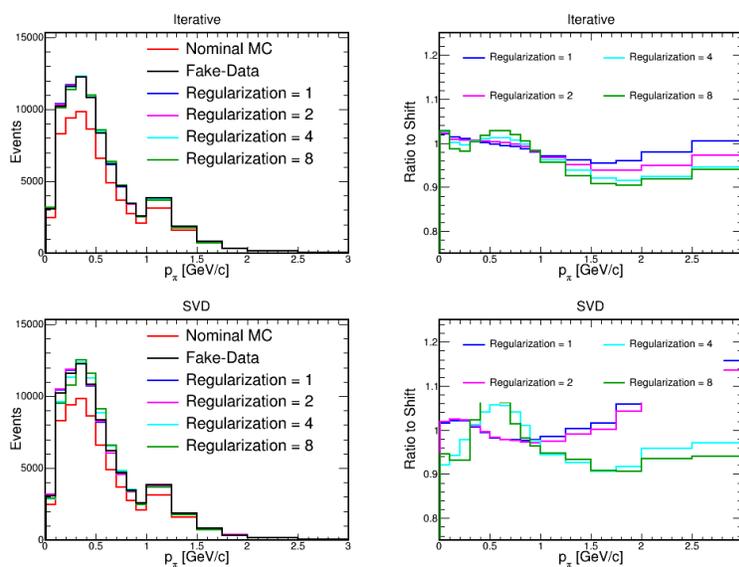


Figure 9.15: Same as Figure 9.14 with the response to a different systematic. The systematic here shifts the vector mass for resonant production up by 1σ .

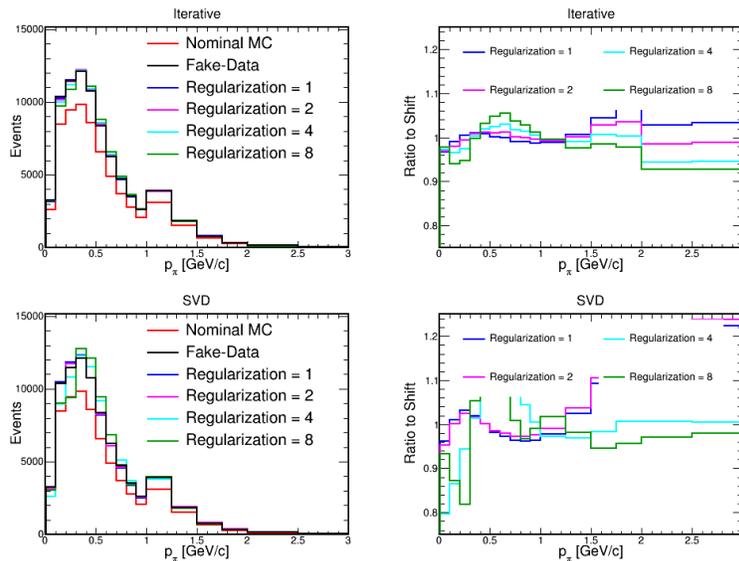


Figure 9.16: Same as Figure 9.14 with the response to a different systematic. The right plots show the unfolded ratio to fake-data. The systematic here shifts up the non-resonant single- π production cross section by 1σ .

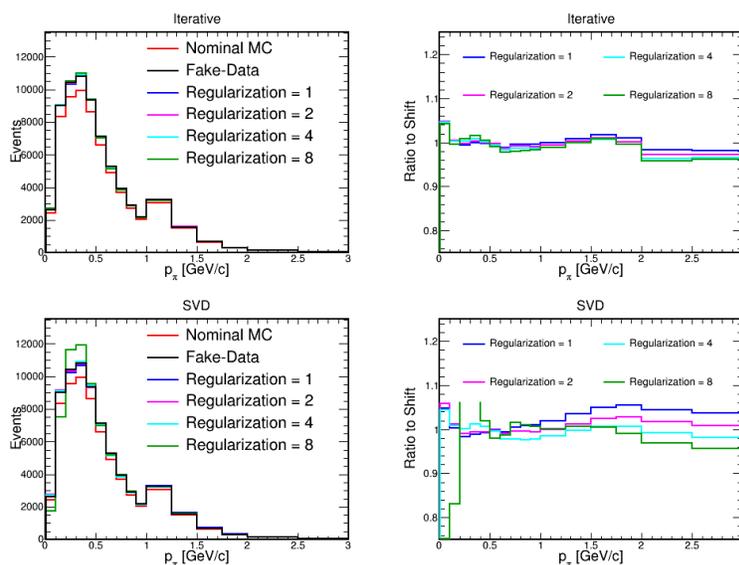


Figure 9.17: Same as Figure 9.14 with the response to a different systematic. The right plots show the unfolded ratio to fake-data. The systematic here shifts up the non-resonant $2\text{-}\pi$ production cross section by 1σ .

9.8 Assembling the Cross Section

This analysis produces flux-averaged $CC\pi^0$ cross sections, differentiated in final state kinematics and Q^2 and W . Any flux-averaged differential cross section is put together with its constituent parts via the equation:

$$\left(\frac{d\sigma}{dx}\right)_i = \frac{1}{\Delta x_i N_n \times \Phi \times \epsilon_i} U(\hat{S})_i \quad (9.5)$$

In this equation, x is the variable of interest, i refers to a particular bin, Δx_i is that bin's width, N_n is the number of interaction targets, Φ is the integrated flux through the analysis's fiducial volume, and ϵ_i is the detection efficiency in the given bin. \hat{S} is a histogram that gives the signal estimate, using simulation constrained by data, in reconstructed bins of x . U refers to the unfolding procedure which corrects for smearing effects in the reconstruction process.

After estimating signal from observed data and unfolding detector effects, a flux-averaged differential cross section is nearly ready to be drawn. Only three ingredients are still needed, which are determined from the simulation: an efficiency correction, the total number of targets in the fiducial volume, and the integrated flux.

The efficiency is calculated by comparing the number of signal events selected by the analysis and generated by GENIE in each kinematic bin. The efficiency depends on the particular kinematic bin and is the ratio of these two histograms. The average efficiency is 21.2%, and the shape in each kinematic variable of interest is shown in Figure 9.19

Since the $NO\nu A$ detector is not made of any single material, the result is presented as a cross section per nucleon of $NO\nu A$ detector soup, along with the list of constituent ingredients in the $NO\nu A$ soup. The number of nucleons is determined in MC using random sampling within the fiducial volume[111]. For each trial, a point is selected, and the material at that point is determined from the geometry file. Doing this a large number of times accurately calculates the occurrence of nuclei of each Z in the fiducial volume which directly gives the total number of nucleons in the defined fiducial volume. This is repeated one million times, which was shown to be sufficiently many to have a negligible error on any cross section analysis[111].

The results of the target count for the $CC\pi^0$ fiducial volume are tabulated in Table 9.2. This gives a total fiducial mass of 35427 ± 9 kg and a nucleon count of $2.12e31$. The error quoted is statistical and comes from only sampling the fiducial volume a finite number of times. As this error is only 0.03%, it is neglected in the systematic error budget.

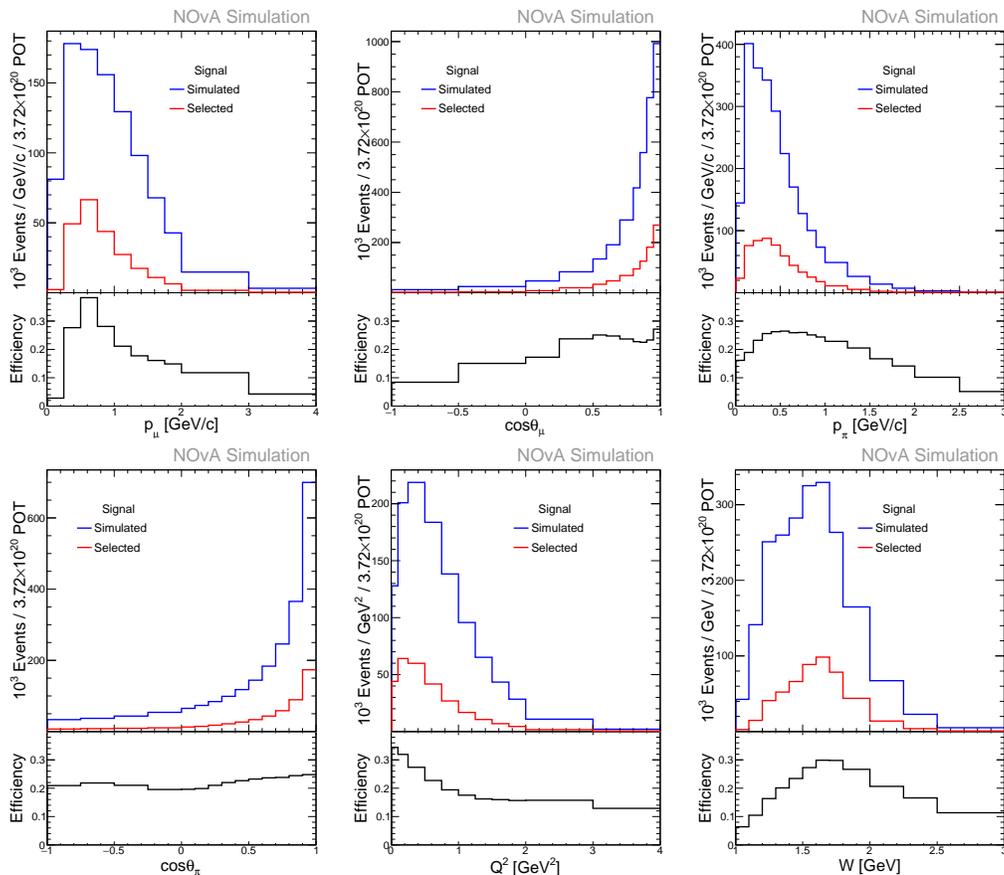


Figure 9.19: The calculated efficiencies for each of the six kinematic variables analyzed. The efficiency is calculated as the ratio of the red, the total signal selected, and blue, the total signal generated, histograms and is shown in the bottom pane of each plot.

| Element | Atomic Z | Mass [kg] | N_{nuc} | Mass Fraction |
|---------|----------|-----------|------------------|---------------|
| C | 6 | 23651.4 | 1.41e31 | 0.667 |
| Cl | 17 | 5685.4 | 3.40e30 | 0.160 |
| H | 1 | 3814.5 | 2.28e30 | 0.108 |
| Ti | 22 | 1139.0 | 6.81e29 | 0.032 |
| O | 8 | 1053.2 | 6.30e29 | 0.030 |
| Sn | 50 | 42.1 | 2.52e28 | 0.001 |
| S | 16 | 33.9 | 2.03e28 | 9.56e-4 |
| Ca | 20 | 9.3 | 5.56e27 | 2.62e-4 |
| N | 7 | 9.2 | 5.50e27 | 2.59e-4 |
| Na | 11 | 0.9 | 5.38e26 | 2.59e-5 |

Table 9.2: The derived mass of the fiducial volume, broken down by atom type, for one calculation of the detector mass. Elements are sorted by total mass fraction.

The last piece needed for the cross section is the integrated flux. As described in Section 6.2, the beam is simulated with G4NuMI. Simulated events are then reweighted according to MINERvA's PPFX calculation which constrains the simulated hadron production uncertainties in the base simulation according to outside data. As the signal definition requires events to have $1 < E_\nu < 5$ GeV, only the flux in this energy range is included in the calculation of the integrated flux. This is a primary motivation for truncating the upper energy of signal events, as the PPFX flux prediction is much less precise at higher energies. The flux histogram used in the analysis is shown in Figure 9.20.

Now, all ingredients for the flux-averaged differential cross section have been described, and Eqn. 9.5 can now be calculated. The GENIE cross section prediction in the NuMI beam is shown in Figure 9.21 for each of the kinematic differential variables. Additionally, the figure shows the total cross section for the signal definition broken into single- and multi- π components and according to the GENIE resonant and DIS labels. Integrated over each kinematic variable, the total cross section is predicted to be 3.28×10^{-39} cm².

Statistical errors for the analysis were calculated using MC-based mock data, in which each bin content has been replaced with a random Poisson fluctuation of the MC prediction in the given bin. Specifically, distributions are fluctuated at the data exposure, 3.72×10^{20} POT. The statistical error on the total cross section is 0.7%. The error in each bin is somewhat higher. Also, the unfolding algorithm correlates

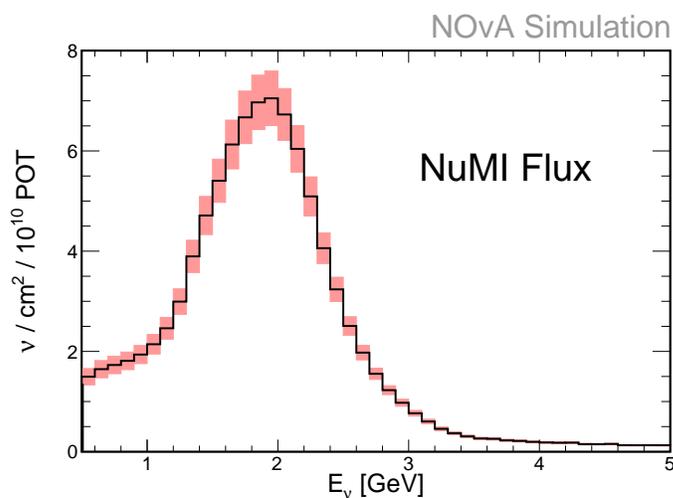


Figure 9.20: The flux used for the analysis, with the PPFX constraint. Only the region with $1 < E_\nu < 5$ GeV is considered. The red band gives the systematic error band on the total flux prediction.

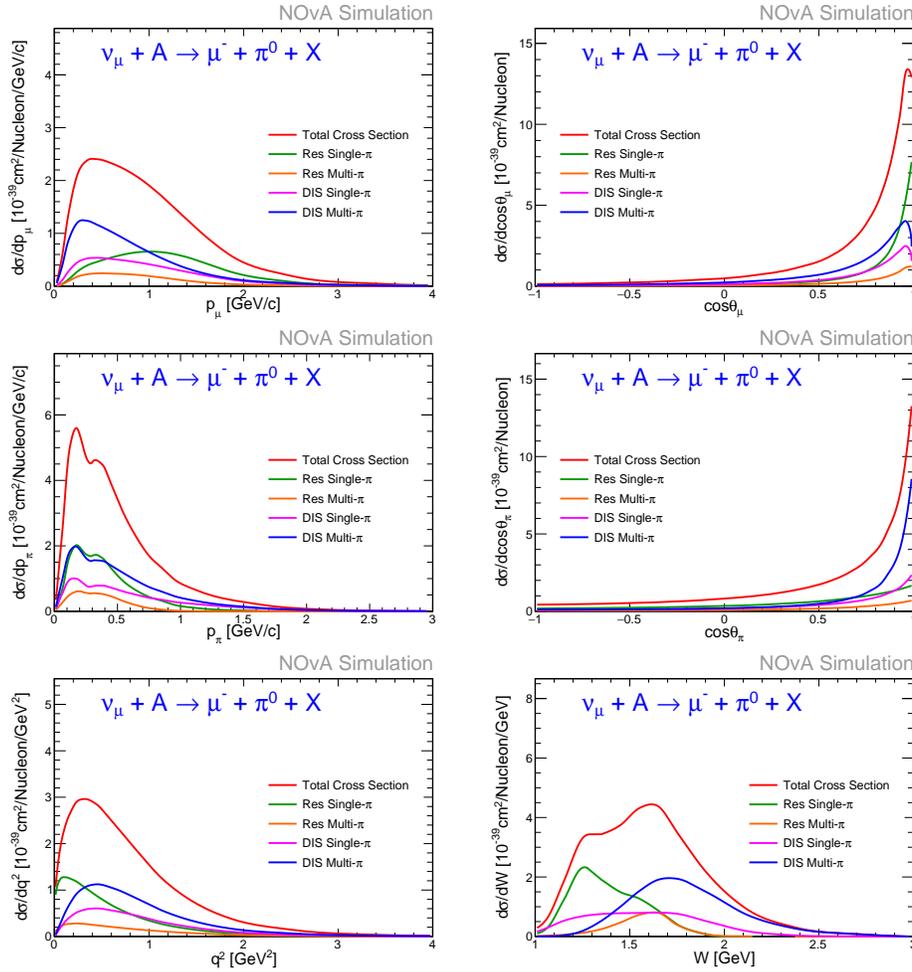


Figure 9.21: The central value cross section predictions from GENIE. The total cross section is shown in red. Other curves break down the cross section into Res and DIS interactions and by single and multi- π production. Note the large multi- π component. Previous $CC\pi^0$ analyses have only analyzed single- π events, but NOvA's relatively high energy allows the study of this large component.

the statistical errors to some degree so that each analysis bin is no longer statistically independent. This is accounted for when calculating the covariance matrix.

9.9 Systematic Uncertainties

A rigorous set of fake-data systematic studies were employed before un-blinding data. Sources of systematic errors were studied using MC fake-data and errors were determined for all non-negligible sources of systematic uncertainty studied. Final systematic uncertainties are calculated by comparing a series of extracted cross sections using systematically shifted MC to the cross section determined with the nominal MC.

These sources have been separated into five classes of systematic uncertainty: neutrino cross section, particle tracking cross section, flux modeling, detector response, and normalization. These are each detailed in the following.

Neutrino Cross Section Systematics

GENIE includes with it a set of orthonormal systematic shifts which tune neutrino cross sections within the current experimental uncertainties for the processes that GENIE simulates[112]. The analysis considers all systematics that cause an average residual $> 0.5\%$ in the reconstructed p_π distribution. In total, 10 systematic knobs were identified, which can be factored into three classes according to the size of shift in the reco p_π distribution. The systematic knobs included in the analysis adjust both the resonant and DIS pion production rates and the hadronization model within the struck nucleus.

There are two large effects that change the number of events observed by at least 10% when averaged over reconstructed p_π bins.

- M_A^{Res}
- NonRes- 2π

There are six noticeable effects that change the number of events observed by 1-10% when averaged over reconstructed p_π bins.

- M_V^{Res}
- NonRes- 1π
- MFP_pi
- MFP_N
- FrAbs_pi
- FrElas_N

There are two noticeable effects that change the number of events observed by 0.5-1% when averaged over reconstructed p_π bins.

- FrCEX_pi

- FrInel_pi

One shortcoming of the GENIE ν -interaction is a failure to properly cover DIS events with high W^2 . These events are at a high enough center of mass energy that the neutrino interactions are constrained by proton-electron scattering data via the PDF theory. GENIE calculates a systematic for DIS events with $W < 2 \text{ GeV}/c^2$. To cover the remaining DIS events, a 15% uncertainty on the DIS normalization for events outside this range is taken.

- DIS_HighW

The quadrature sum of these 11 systematic parameters gives a 4.6% error on the total cross section. It should be noted that the remaining systematic knobs were tested. The quadrature sum of the systematic effect from these small contributions was everywhere a small effect for the cross section differentiated in p_π . As such, entirely dropping these knobs from the analysis is reasonable.

To calculate the systematic uncertainty for each of these 10 sources, a cross section is calculated with the shifted MC used as fake-data for both the $+1\sigma$ and -1σ shifts. The $+1\sigma$ and -1σ shifted cross sections are compared to the nominal MC cross section, and the systematic error in each bin is determined as

$$err_i = \frac{1}{2} \left(|\sigma_i^{+1} - \sigma_i^{nom}| + |\sigma_i^{-1} - \sigma_i^{nom}| \right). \quad (9.6)$$

To illustrate this method, Figure 9.22 shows the process applied for a shift in M_A^{Res} by $\pm 1\sigma$.

Particle Tracking Cross Section Uncertainties

As discussed earlier, there is a notable fraction of background events whose photon candidate was made by a true photon. The majority of these come from the charge eXchange (CX) reaction, $\pi^\pm \rightarrow \pi^0$ when tracking through the detector. This CX cross section, unfortunately, is not perfectly known and affects the signal prediction in the analysis. Thus, a systematic uncertainty is calculated by adjusting this cross section in MC. This was the largest source of systematic uncertainty on the MiniBooNE result[105] for this cross section channel.

This analysis is sensitive to the CX cross section. This cross section is only known to the 20% level by the 1981 Ashery analysis[113] and the 2017 DUET analysis

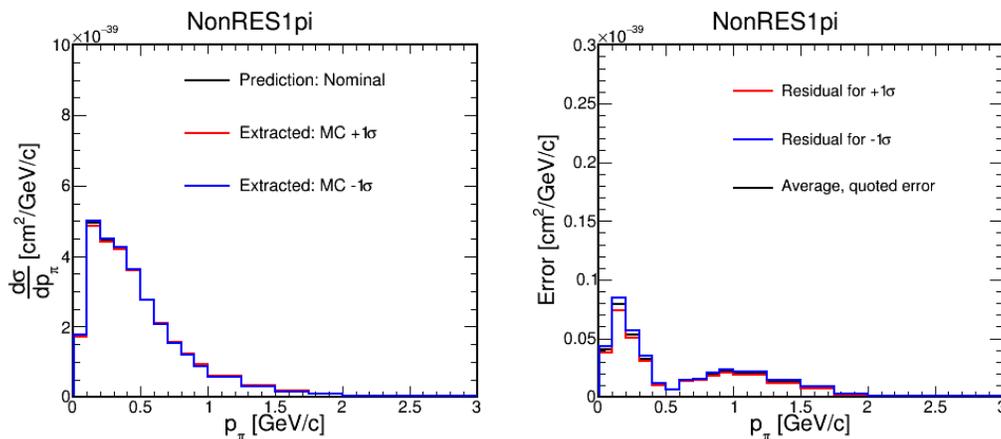


Figure 9.22: An example showing the determination of a systematic uncertainty. The left plot compares the cross section measured in data extracted using three sets of MC: the GENIE predictions for nominal, $+1\sigma$, and -1σ values of the GENIE DIS pion production cross rates. The difference between the cross section determined with the shifted MC and the nominal MC is taken as the systematic, when averaged over the $\pm 1\sigma$ shifts.

[114] at each data point of interest. Both of these measurements are taken around the π^\pm momentum of 300 MeV/c which is at the peak of the Δ_{1232} resonance for π^\pm -nucleon scattering, and thus cover the most relevant region of phase space for π^\pm scattering off nuclei.

The first step towards calculating a proper systematic is studying how consistent the NO ν A simulation is with experimental measurements of this cross section. A MC π^+ scattering experiment off a thin foil of carbon in the NO ν A simulation determined the exclusive cross sections for π^+ -carbon interactions within the NO ν A simulation framework[115].

The DUET result gives a modern and relatively precise measurement. The collaboration also published the covariance matrix with their results[114] while the Asher result does not. Therefore, the simulation was tuned to this DUET result. The cross section curve for π^\pm charge exchange determined by NO ν A simulation was tested according to the DUET data. Using the covariance matrix, a $\Delta\chi^2$ fit is performed to determine what range of scale factors of the simulation curve are compatible with the data[116]. This resulted in the symmetric interval, 1.061 ± 0.146 , as shown in Figure 9.23.

The confidence interval constructed suggests the central value is slightly higher than the NO ν A simulation. As the NO ν A MC set was generated before the DUET result

was published, the nominal MC used for the $CC\pi^0$ analysis involves a reweight to increase the CX cross section by 6.1%. This allows us to take advantage of the narrowest error band allowed given the uncertainty on the DUET measurement. All numbers and plots in this chapter apply this 1.061 scale factor to the π^\pm CX cross section.

Weighting MC events to a new value of the CX cross section is not as simple as weighting events that have a $\pi^\pm \rightarrow \pi^0$ interaction. This would change the number of background neutrino interactions with a π^\pm in the final state which is clearly wrong for this effect. It is also not as simple as weighting background events with a CX, and adjusting the remaining background events with a π^\pm to impose unitarity. In the case when the probability for a π^\pm to CX becomes very close to 1, this system would lead to negative weights applied to the remaining background events with a π^\pm .

A model was developed[116] which assumes the number of CX interactions in an

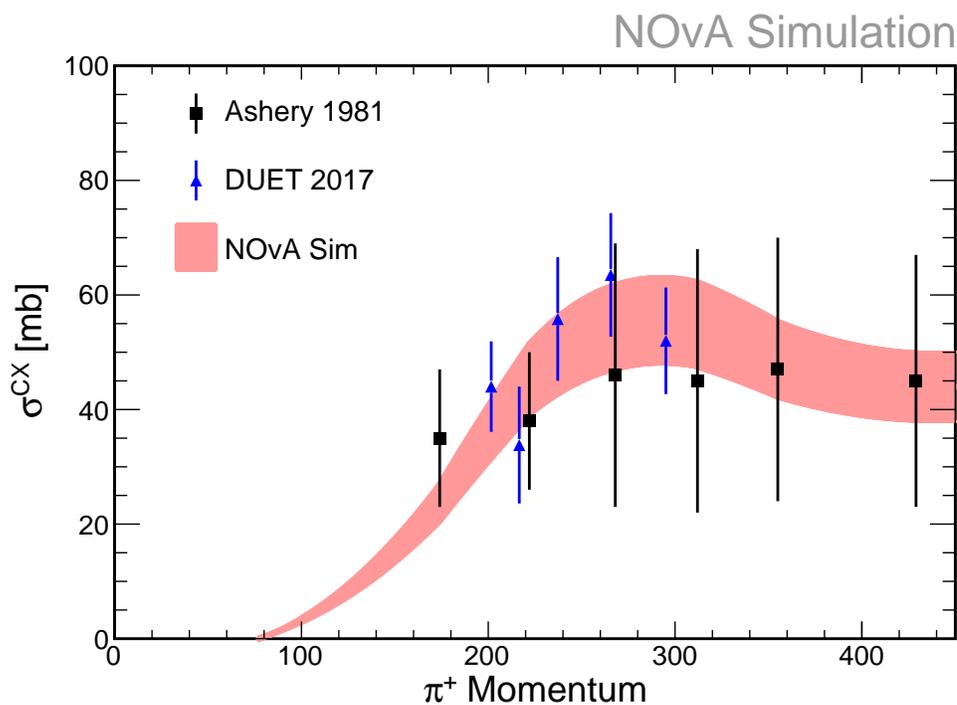


Figure 9.23: A comparison of the Ashery[113] and DUET[114] measurements of the $\pi^+ \rightarrow \pi^0$ CX cross section and the spread in cross section of the NOvA simulation consistent with these results. A 68.3% confidence interval for the range of scale factors applied to the nominal simulation was calculated as 1.061 ± 0.146 when comparing to the DUET data.

interaction follows a Poisson distribution. Eqn. 9.7 shows the probability for an event with a π^\pm in the final state to have 0 CX interactions is simply $e^{-\lambda}$:

$$P(\lambda, N) = \frac{\lambda^N e^{-\lambda}}{N!} \implies P(\lambda, 0) = e^{-\lambda}. \quad (9.7)$$

To weight MC to a scaled CX cross section, the expected number of CX interactions, λ in the Poisson distribution, must be scaled by the same scale factor in the CX cross section. Writing the desired scale factor as $1 + \delta$, the weight needed to be applied to background events with a π^\pm in the final state but no CX interaction is

$$\frac{P((1 + \delta)\lambda, 0)}{P(\lambda, 0)} = \frac{e^{-(1+\delta)\lambda}}{e^{-\lambda}} = e^{-\delta\lambda} = e^{\delta \ln P(\lambda, 0)}. \quad (9.8)$$

Then, enforcing unitarity on the number of background events with π^\pm , the scale for events with a CX interaction is

$$\left(1 - P(\lambda, 0)e^{\delta \ln P(\lambda, 0)}\right) \left(\frac{1}{1 - P(\lambda, 0)}\right). \quad (9.9)$$

As $P(\lambda, 0)$ can be empirically measured in MC as the fraction of background events with a π^\pm in the final state that do not have a CX, this model gives a quick and easy weight in closed form for background events with a π^\pm in the final state. $P(\lambda, 0)$ was measured as a function of true p_π^+ and the weights are calculated separately in each analysis bin of true p_π . Over the sample for the central value weight of 1.061, this procedure increases the fraction of background events with a π^\pm in the final state that have a CX interaction by 5.54% and decreases those that do not have a CX by 2.75%, leaving the total background rate constant. To give a sense of the scale of this effect, Figure 9.24 shows the amount the background $CC\pi^0ID$ template can change under the allowed CX cross section region.

As there is now a procedure for weighting MC events to adjust the value of the CX cross section, the systematic uncertainty is readily calculated. This is calculated by taking the difference in extracted cross sections of the nominal MC. Doing this gives a 3.8% bias on the total cross section, comparable to the summed effect of neutrino cross section systematic uncertainties. Without the DUET result, the analysis would have used the error band quoted by Ashery which would have made this systematic the dominant uncertainty in the analysis.

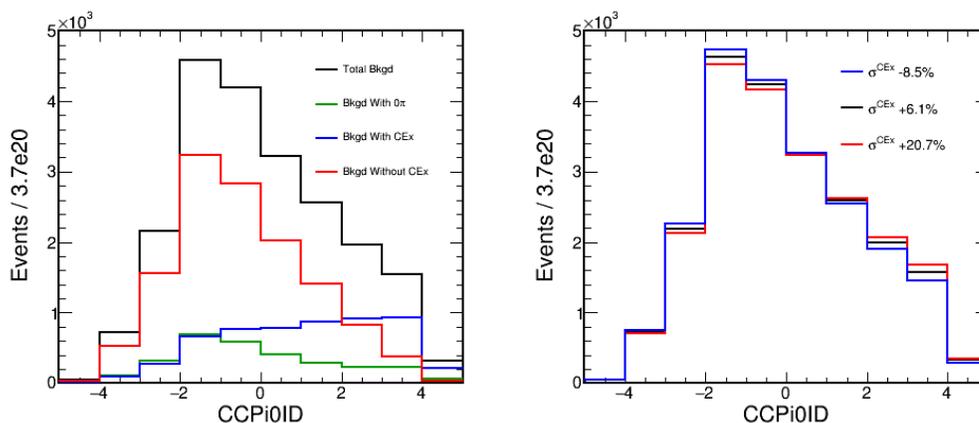


Figure 9.24: To give a sense of the size of effect weighting across the DUET-allowed cross section for π^\pm CX, the ν_μ CC background is decomposed into 0π in green, events with a π^\pm but no CX in red, and events with a π^\pm CX in blue in the left plot. The 0π component represents 13% of the background, while events with a π^\pm CX represent 26% leaving 61% of the background for events with a π^\pm but no CX. The right plot shows the effect on the $CC\pi^0ID$ template shape for background caused by changing this CX cross section to the adjusted central value (6.1%) and within the allowed error band.

The kinematics of the outgoing π^0 in CX interactions are not well known. In fact, there is no differential cross section available at the interesting energy range for NOvA.

As there is not published data to provide an error band on the kinematics of this distribution, a very conservative smearing procedure was done to the simulation prediction of outgoing π^0 momenta. This smearing is a conservative and rough adjustment to the simulation. It was performed to demonstrate that the total CX cross section was more important to the $CC\pi^0$ analysis sensitivity than the CX cross section shape so that these shape uncertainties can be dropped. In this smearing, a Δp_π was thrown from a uniform distribution between $\pm 300/\sqrt{3}$ MeV in each x , y , and z direction, giving a maximum magnitude slightly greater than the Fermi momentum in carbon. In the framework, a square distribution was much easier to implement. The expectation value of each x , y , z component is $300/\sqrt{3}/2$ MeV in this system, so that the quadrature sum of the expectation values is 150 MeV, half the maximum Δ_π , meaning this is the right size box to fit the spherical model.

This Δp is then added to the reconstructed p_π in simulation and the resulting cross section is compared to the nominal. After applying this procedure, the resulting systematic error was in the sub-percent level. As this smearing is not well motivated

and the effect is negligible compared to changing the normalization, the kinematic effect is dropped from the systematic budget.

Flux Modeling

Two sources of flux systematic uncertainty were assessed. The hadronization component comes from the poor theoretical understanding and limited experimental data for hadron production for the proton-nucleon scattering in the NuMI target. The Minerva PPFX error band prediction is used in the systematic calculation.

A systematic uncertainty from the beam transport and focusing component was also assessed. This systematic covers variations in the state parameters of the beam such as horn current, beam spot position on the target, beam spot size, and bending from the magnetic field in the decay pipe.

The strategy for calculating the systematic error is similar to the neutrino cross section systematics. The cross section is extracted with systematically shifted MC and compared to the nominal. The total flux error is dominated by hadronic production uncertainty, giving a 8.3% uncertainty on the total cross section. As with many neutrino cross section analyses, this is the largest source of systematic uncertainty.

Detector Response

NOvA has studied several sources of uncertainties coming from mis-modeling the detector response. The three largest sources of error were shown to significantly impact this analysis and are included in the systematic error budget. The uncertainty is calculated in a similar way to the neutrino cross section systematics, with comparing the nominal cross section with the cross section extracted with a shifted MC, giving

$$err_i = |\sigma_i^{shift} - \sigma_i^{nom}|. \quad (9.10)$$

The first relates to the simulated light level. The Birks-Chou formula[78][79] is an empirical relation giving a parameterization of dE/dx non-linearity of light yield in scintillator, given by

$$LY = A \frac{\frac{dE}{dx}}{1 + k_B \frac{dE}{dx} + k_C \left(\frac{dE}{dx}\right)^2}. \quad (9.11)$$

The simulation used in this $CC\pi^0$ analysis predates the inclusion of Čerenkov light emitted in the scintillator in NOvA's simulation.

The Birks-Chou parameters were determined by a tune to selected protons in the near detector data sample[117]. As shown in Figure 9.25, the tuned MC much better matches our observed energy deposition for slow protons with high dE/dx in the regime where Birks-Chou effects are large. Without Čerenkov light in the simulation, this tune suggested an anomalously high value $k_B = 0.40 \text{ g/cm}^2\text{MeV}$. Organic liquid scintillators like pseudocumene have a typical value of $k_B = 0.10 \text{ g/cm}^2\text{MeV}$.

To compensate for this known short-coming of the MC dataset, a separate systematically shifted MC sample was generated with $k_B = 0.10 \text{ g/cm}^2\text{MeV}$. The extracted cross section is then compared to the nominal, and a systematic error from scintillator non-linearity mis-modeling is calculated. The signal template in $CC\pi^0$ is not significantly affected by this distortion, as electromagnetic showers are composed of several small MIP particles whose light level is fixed by construction in the calibration procedure. The background template, dominated by proton and π^\pm prongs, is significantly adjusted. Further, the background template within the updated simulation framework, including Čerenkov light and typical Birks-Chou parameters, is found to lie within the error band derived from this systematic. The effect on the signal and background $CC\pi^0$ ID templates from the updated simulation model is shown in Figure 9.26. Though the updated simulation lies much closer to the Birks-shifted simulation used for the analysis, it lies within the difference

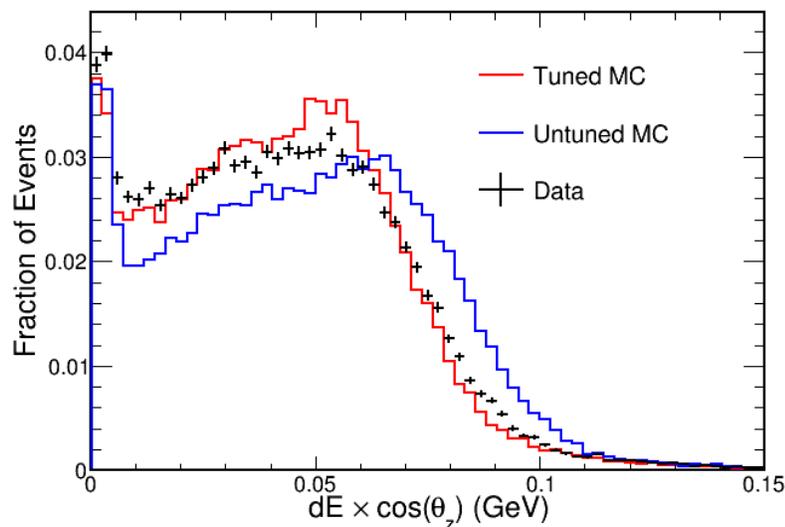


Figure 9.25: The direction-corrected energy deposit for the last plane of a high purity proton track sample.

between the shifted and nominal simulations used in this analysis.

Apart from flux uncertainties, this is the dominant error source for the analysis. It causes a 6.8% error on the total extracted cross section.

The effect of absolute calibration scale modeling is also included as a source of systematic uncertainty. Studies show that observed dE/dx is higher in the MC used for this analysis than data[118][119][120]. There are several studies pointing to this. Figure 9.27 shows that calorimetric energy in hits on μ tracks from selected ν_μ CC events is higher in MC than data. Another study, looking at the calorimetric energy of tagged Michel electrons[118] at the end of μ tracks in selected ν_μ CC events, shows reconstructed energy is higher in MC than data, as shown in Figure 9.28. Yet another study looked at the invariant mass peak for NC π^0 events and found that the invariant mass was typically found to be higher in MC than data, as plotted in Figure 9.29. Most discrepant was a study of proton dE/dx , which found observed MC dE/dx 5.4% higher than MC[120].

A negative shift in the MC energy scale sufficiently covers all discrepant studies found. There is no physical scenario where the MC calibration can be shifted both up and down relative to data, so including both of these effects fundamentally over-estimates the systematic uncertainty from calibration scale. Instead, since all studies involving dE/dx have shown observed dE/dx high in MC, only the negative shift is considered for this analysis. The energy on true electromagnetic prongs from electrons and photons is shifted down by 2%, determined by data / MC discrepancies in the Michel electron spectrum. All other hadronic energy is shifted down by 5%, driven by the proton dE/dx differences in data and MC, the largest discrepancy

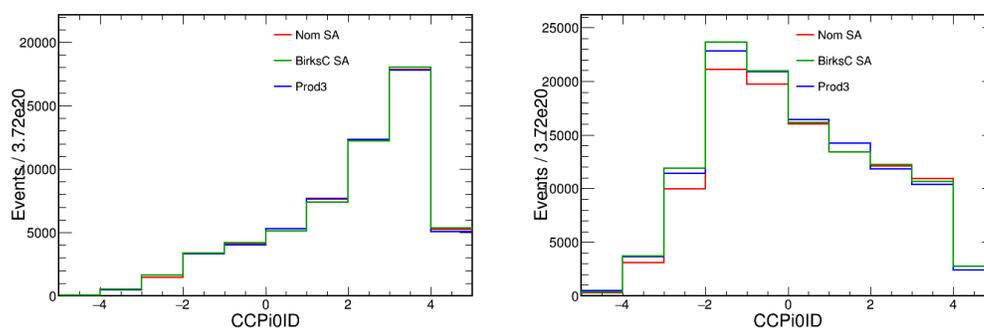


Figure 9.26: Differences in the signal, left, and background, right, $CC\pi^0$ ID templates. The nominal MC used in the $CC\pi^0$ is shown in red while the Birks-shifted MC is shown in green. The blue curve shows the template in the updated simulation that includes Čerenkov light.

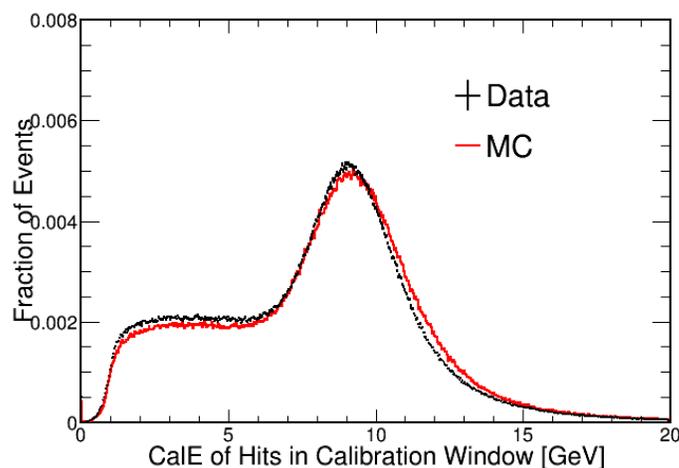


Figure 9.27: The calorimetric energy of hits within the calibration window of tagged muon tracks in the near detector ν_μ CC sample as selected by the oscillation analyses used to determine the absolute energy scale. The MC peaks at 1.1% higher than the data. The broad shoulder below the peak comes from "clipping" hits that do not traverse the entire z -range of the cell.

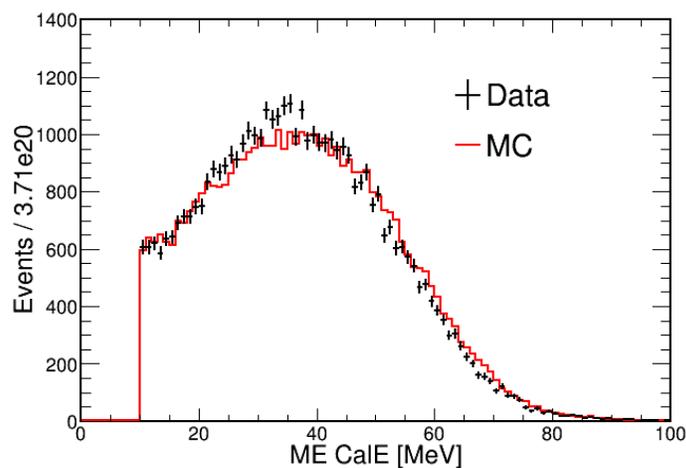


Figure 9.28: The calorimetric energy of Michel electrons tagged at the end of rock muons in the near detector. The MC peaks 1.5% higher than data.

between any of our calibration test-samples.

The last detector response systematic that is accounted for modifies the functional form of the light yield across the length of the cell, a way of modifying the attenuation calibration. This effect is most relevant for large discrepancies between data and MC cosmics at the far detector, but shifted files were generated for the near detector and evaluated for the $CC\pi^0$ analysis. The systematic uncertainties are modest, but

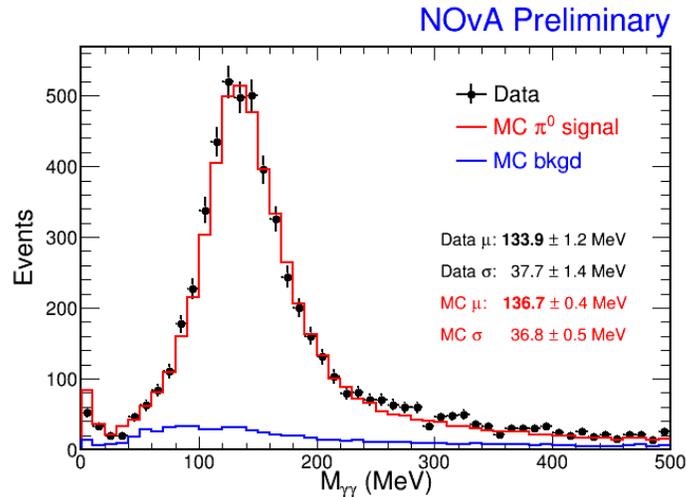


Figure 9.29: The reconstructed π^0 invariant mass peak for a high purity sample of tagged NC π^0 events. The MC is peaked at 2.8 ± 1.3 MeV higher than the data.

noticeable in some regions of the kinematic variables studied. The discrepancy between data and MC is different for X and Y -view cells, so the two systematics are treated separately. The size of this systematic uncertainty is smaller than others included, but not negligibly so. It is therefore included in the error budget. Summed in quadrature, the absolute and attenuation calibration systematic gives a 2.6% uncertainty on the total cross section.

Normalization Systematics

Three additional sources of systematic uncertainty are included in the analysis. They have been shown to be small and not affect the derived shapes of kinematic variables, and are thus treated as normalization systematics. A 0.7% uncertainty on the fiducial mass of the detector is included[121]. Additionally, all MC was simulated at a constant spill intensity of 25×10^{12} POT/spill. The data, however was taken at a distribution of spill intensities. This can bias the number of slices reconstructed as higher pileup can lead to multiple interactions being sliced together which the selection rejects. This effect was shown to bias the number of events selected by 0.5%[122].

The final normalization systematic estimates how well the containment efficiency is modeled in the simulation. The data-MC ratio was calculated for events passing selection cuts in two subdivision of the fiducial volume based on the vertex x and y coordinates. Region 1 is defined as the subregion of the fiducial volume divided by the reconstructed vertex: $-70 < x$ and $y < 70$ cm, while Region 2 is the remainder of

the fiducial volume. Geometrically, Region 1 is a tighter fiducial volume in x and y . The data to MC ratio is 90.8% in Region 1 and 92.7% in Region 2. The difference, 1.9%, is taken as normalization systematic as a containment mis-modeling effect. Summing with the two other normalization systematics, this gives a 2.1% total bias on the measured cross section.

Total Error Budget

The total error budget, as a function of the measured kinematic variables, is shown in Figure 9.30. The flux systematic is a large, but not everywhere dominant systematic. The light level, and calibration scale systematic uncertainties are comparable to the effect from flux. As these systematic effects are likely to be smaller in future analyses, there is certainly room for improvement in making a more precise measurement of this channel in the future. Other sources of error are often small compared to the total quadrature sum, but are included for consistency.

Though the total flux averaged cross section is not the main goal of this analysis, the systematic uncertainty on the total cross section is illustrative of the precision of the analysis. These errors, broken down by each source of systematic are shown in Table 9.3.

Validation of Computed Systematic Uncertainties

Analyses should strive to study as many disjoint control samples as possible before analyzing data. The $CC\pi^0$ analysis found two such samples – based on selections that hone in on a high-purity sample of photons and protons. Within these samples, we will test the agreement of $CC\pi^0ID$ between data and simulation. The ratio of data to simulation is shown along with an error band. Since these are single particle

| Systematic Source | Rel. Error |
|-------------------------|------------|
| Neutrino cross sections | 4.6% |
| π^\pm CX | 3.8% |
| Flux | 8.3% |
| Light Level | 6.8% |
| Calibration | 2.6% |
| Normalization | 2.1% |
| Quadrature Sum | 12.5% |

Table 9.3: The effect of each systematic on the extracted total cross section. The flux uncertainties are the largest source of systematic error, with large contributions from the light level systematic, calibration, and cross section uncertainties.

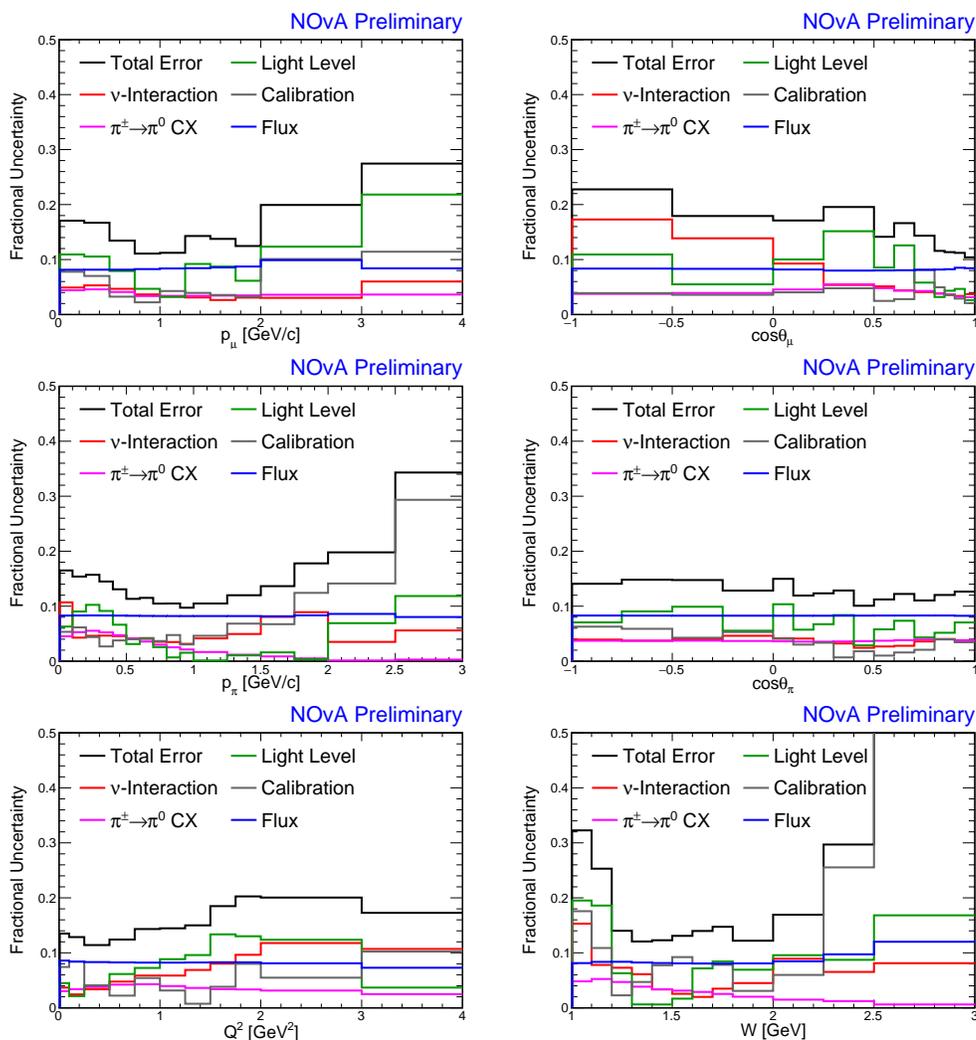


Figure 9.30: The systematic error budget for each of the measured kinematic variables.

samples, particle production cross sections play a minor role, and thus only detector response systematics are included in the error band. The results have been area-normalized to further reduce the effect of cross section systematics. Photon and proton prongs are the most commonly selected particles in the analysis for signal and background events, respectively, so this study lends confidence to the validity of our MC template shapes, up to systematic distortions that we quote.

The photon sample involves looking at constituent prongs within a sample of two-prong NC- π^0 events[123]. The selection is entirely topological on the preselection of two-prong events. First, the ReMID score for the event must be below 0.5 to reject any CC background. Then, each prong is required to have at least two planes along the prong that did not have any energy deposited in them; this is put in place to

highlight brems activity in the prongs. Finally, the average dE/dx must be below 3 MeV/cm, to further reduce the dominant proton background.

From the resulting sample, the higher photon $\Delta \log \mathcal{L}_\gamma$ scores of the two prongs is plotted as the event $CC\pi^0ID$. There are over 2000 events selected, and the selected sample is predicted to be 82.1% pure in photons. Further, most non-photon prongs come from π^\pm and protons whose $CC\pi^0ID$ distribution is skewed low. The area normalized comparison between data and simulation is shown in Figure 9.31. The detector response systematic errors are also plotted. The data / MC ratio lies within the estimated error band, but the error band is not much wider than the ratio. Thus, it seems the error band is well-estimated.

For the proton sample, the goal is the same, but the selection strategy is very different. The selection strategy entirely relies on the use of ν_μ CC QE events as an in-situ “proton gun”. Using two-prong ν_μ CC selected events, the muon prong is selected by the highest ReMID value, as in the main $CC\pi^0$ analysis. The second

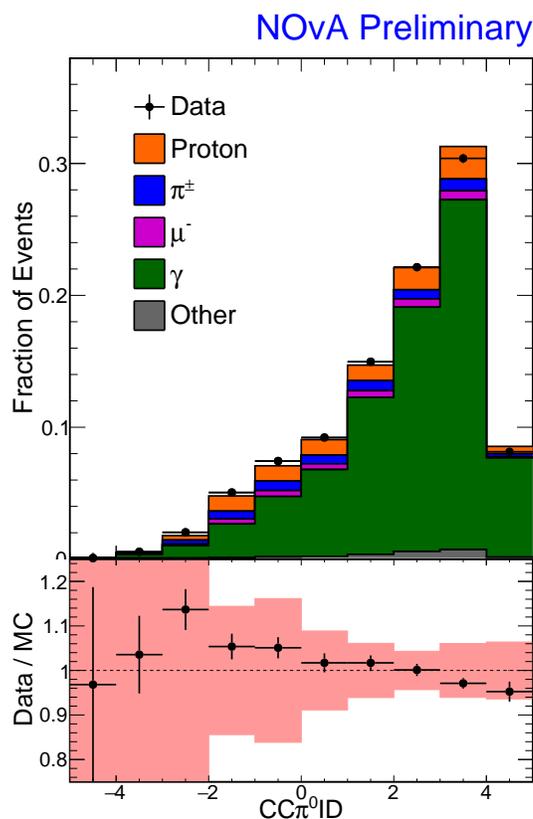


Figure 9.31: The area-normalized $CC\pi^0ID$ distribution in data and MC for a high-purity photon sample of $NC-\pi^0$ events. The data/MC ratio lies comfortably within the detector response systematic error band.

prong is then tested for compatibility with the two-body QE hypothesis.

From this preselection stage, there are two cuts on the reconstructed angle of the second prong. First, the muon and proton candidate prongs must have $\cos \theta_{\mu p} > -0.8$. This cut is imposed to reject a reconstruction failure where the vertex can be reconstructed a few planes upstream of the true vertex, splitting a single muon trajectory into two prongs, both produced by the muon. The kinematic variable, along with the particle type that produced each proton candidate prong, is shown in Figure 9.32.

The second cut relies on reconstructing the proton direction in two distinct ways. First, it is directly measured from the reconstructed prong direction. Then, the muon kinematics and reconstructed neutrino energy are, together with the quasielastic formula for the neutrino energy, used to infer the direction in which a proton is expected. The second and final kinematic cut requires $\cos \theta_{pp} > 0.9$. This variable is also shown in Figure 9.32. It is worthwhile noting that there is no direct cut applied to any dE/dx variables in this sample, which are correlated with $CC\pi^0$ ID.

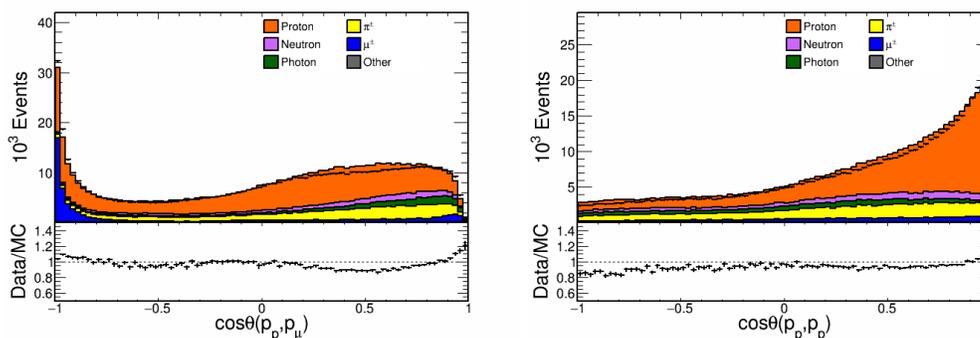


Figure 9.32: The kinematic variables used for selecting the high-purity proton sample, both depending on the reconstructed angle of the proton candidate. It must first be not anti-parallel with the muon prong, and then consistent with the proton direction inferred from the QE formula given the reconstructed neutrino energy and muon kinematics.

After these two directional cuts, there are over 30000 events selected in data, with a predicted proton purity of 82.9%, with an even higher purity at $CC\pi^0$ ID. As in the photon case, the area-normalized $CC\pi^0$ ID spectrum is drawn in data and MC. The ratio is then compared to the simulated detector-response error band. As in the photon sample, the data / MC ratio lies within, but near, the 1σ error band, validating the systematic error cocktail that we've developed. The distribution is shown in Figure 9.33.

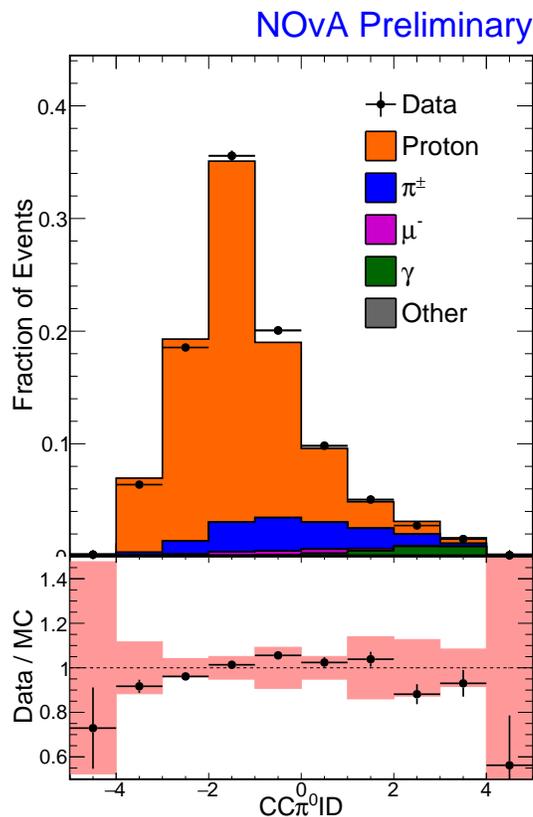


Figure 9.33: The area-normalized $CC\pi^0ID$ distribution in data and MC for a high-purity proton sample identified from events consistent with ν_μ CC QE kinematics. The data/MC ratio lies near the 1σ error band determined from detector response systematic errors.

9.10 Signal Estimation with Data

Now, the signal and background are fit to the observed data. This happens in every analysis bin for each measured kinematic variable separately. For the six variables, this gives 56 individual fits. As examples, two of these fits in bins of reconstructed p_π are shown in Figure 9.34.

Of course, the measurement in each variable is built of several of these fits, one in each analysis bin. The overall data excess is shown in Figure 9.35 as a function of each reconstructed variable and $CC\pi^0ID$. In each of these plots, a row corresponds to the $CC\pi^0ID$ distribution that is fit in the particular analysis bin. The most notable discrepancy is at high p_π , where there is a large deficit of low $CC\pi^0ID$ events, suggesting the simulation significantly overestimated the background in this region. The $\cos\theta_\pi$ distribution also has an interesting feature where the excess over simulation is confined to high $CC\pi^0ID$ that are neither very forward or very

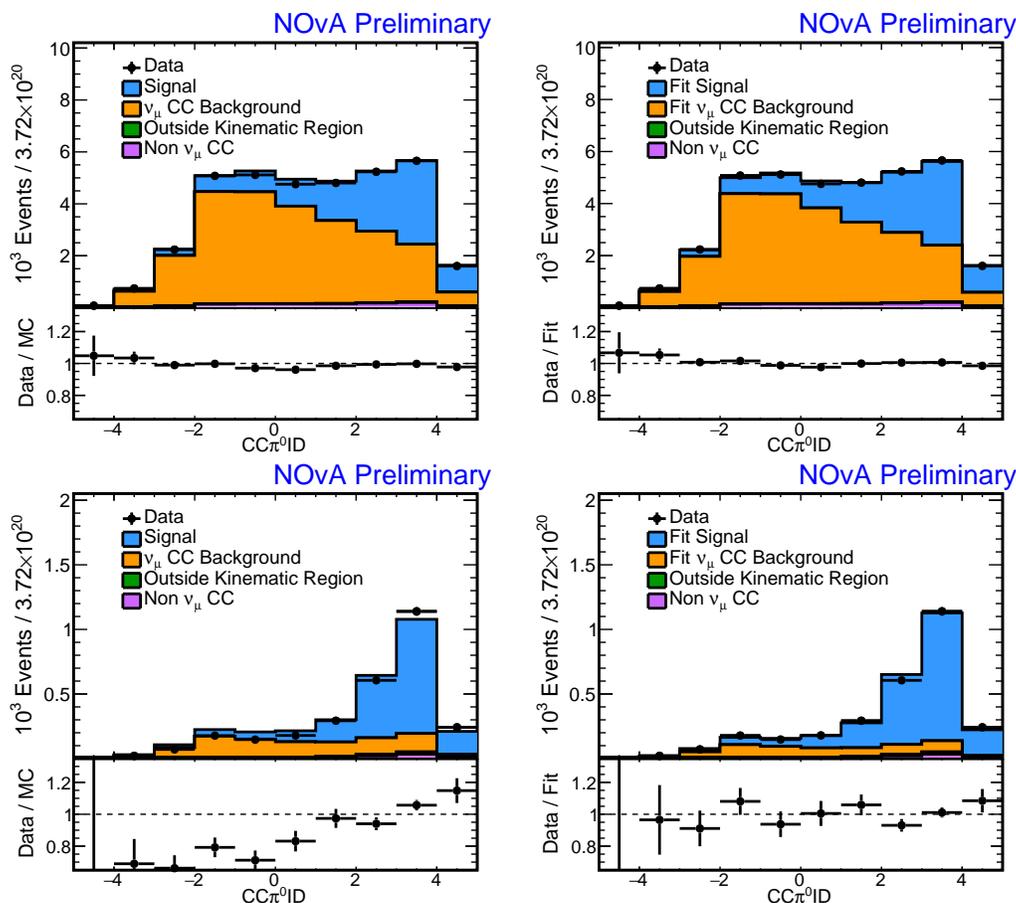


Figure 9.34: The results of the fit to data for events with $0.3 < p_\pi < 0.4$ GeV/c, top, and $1 < p_\pi < 1.25$ GeV/c, bottom. In each row, the left shows the comparison of data and raw simulation while the signal and background normalization on the right has been fit to the data.

backward. And, in W , there is a sharp overall deficit between 1.5 and 1.7 GeV, but the deficit for signal-like events is much less dramatic.

Next, the signal and background normalizations in each analysis bin are collated and used to adjust the simulated signal and background histogram. For each measured variable, this is shown in Figure 9.36. These estimated signal histograms, constrained by data through the fit in $CC\pi^0ID$, are then handed off to the unfolding algorithm and used to produce a cross section.

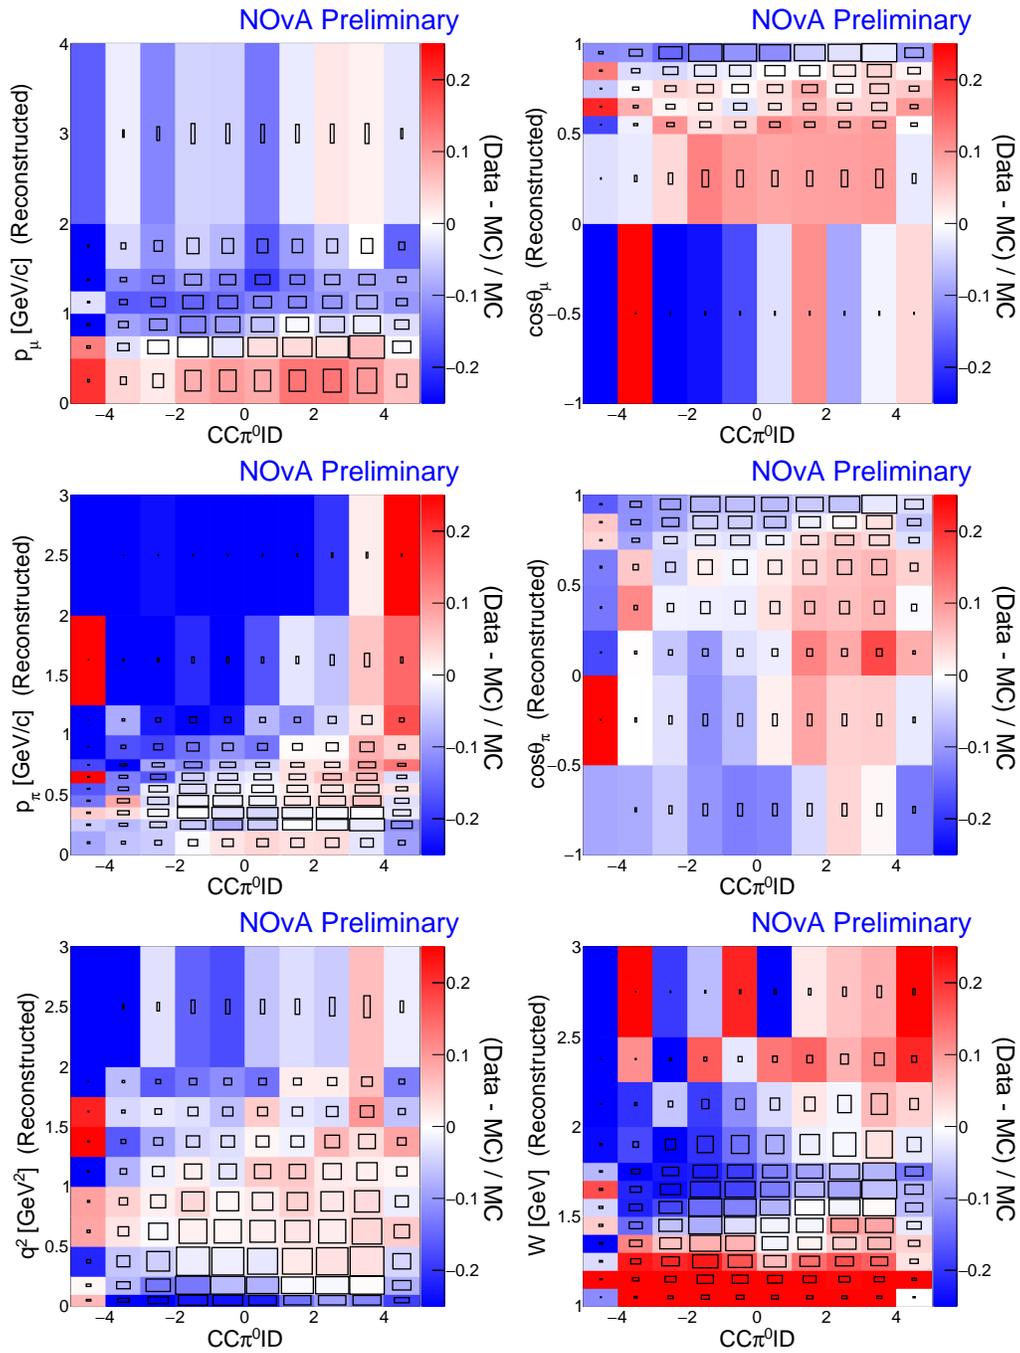


Figure 9.35: The state of agreement between our simulated prediction and data. The colors give the relative excess, red, or deficit, blue of data events as a function of each kinematic variable and $CC\pi^0ID$. The boxed distributions show the population of selected events.

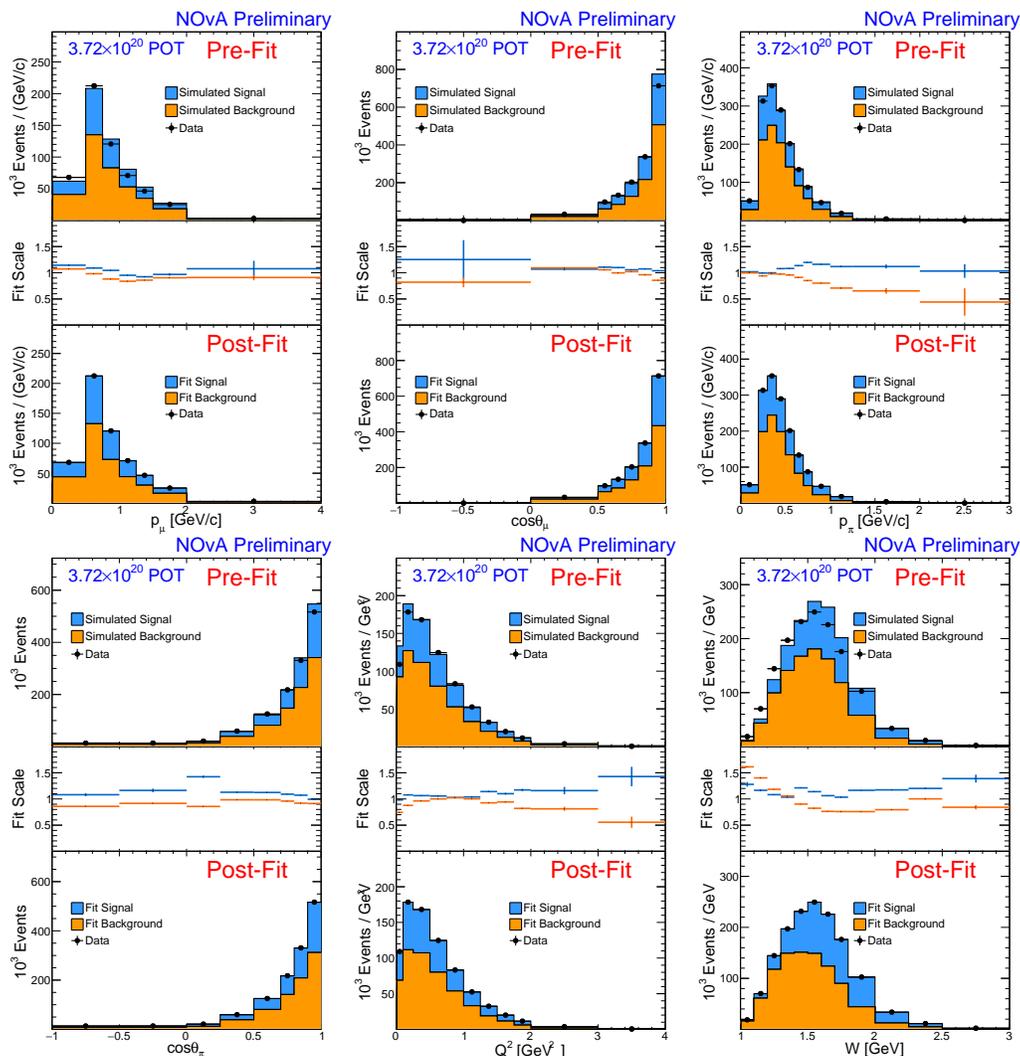


Figure 9.36: For each of the six measured kinematic variables, the unconstrained data and MC comparison is shown on the top pane of each plot. The template fit gives a signal and background scale factor in every analysis bin, which is shown in the middle. These normalization scales are then applied to the simulated distributions in the top panes, resulting in the constrained signal and background, shown in the bottom panes.

9.11 Analysis Results

After unfolding and efficiency correction, the cross section is easily computed. This is done separately for measurements in p_μ , $\cos \theta_\mu$, p_π , $\cos \theta_\pi$, Q^2 , and W . To reiterate, the analysis signal is any ν_μ CC event with a π^0 in the final state that lies within the true kinematic region as defined.

Muon Kinematics

The flux-averaged differential cross sections in the muon final state kinematics are shown in Figure 9.37.

The observed p_μ distribution is biased towards lower p_μ . In the shape-only cross section comparison, discrepancies at 1σ in each bin exist both below 0.75 GeV/c and between 1.25 and 2 GeV/c. The region of excess observed strongly correlates to the range where multi- π interactions dominate. Below 0.75 GeV/c, 61.1% of simulated events have multiple π 's in the final state while only 46.2% of events between 1.25 and 2 GeV/c and 33.7% of events above 2 GeV/c do.

In the azimuthal angle distribution, the simulated shape agrees well with the observed cross section. Only a slight deficit for the most forward-going muons is seen. There is a turnover in the predicted cross section at very slight angles relative to the beam, characteristic of the nuclear shielding effect in DIS scattering at low Q^2 . This follows the observed data in the bin with $\cos \theta_\mu > 0.95$.

π^0 Kinematics

The resulting cross sections in Q^2 and W are shown in Figure 9.38.

In the p_π differential cross section GENIE overpredicts the cross section near $p_\pi = 0.3$ GeV/c. The disagreement would be more exaggerated in a model without FSI interactions as initially produced hadrons could not transmute and move lower from this momentum. This is precisely the resonance momentum for $p + \pi \rightarrow \Delta_{1232}$ production. This channel features heavily in hadronization models within the nucleus. The GENIE prediction in the momentum region below this dip is dominated by inelastic hadronic scatters that produce a π^0 . In this very low p_π region, the shape agrees as a test of the outgoing momenta and flavor of hadrons scattering in the Δ_{1232} resonance. The p_π differential cross section is separated into FSI interaction channels in Figure 9.39.

The $\cos \theta_\pi$ distribution shows a clear excess at large angles relative to the beam, with $\cos \theta_\pi < 0.8$. For very forward-going π^0 's, the cross section prediction is

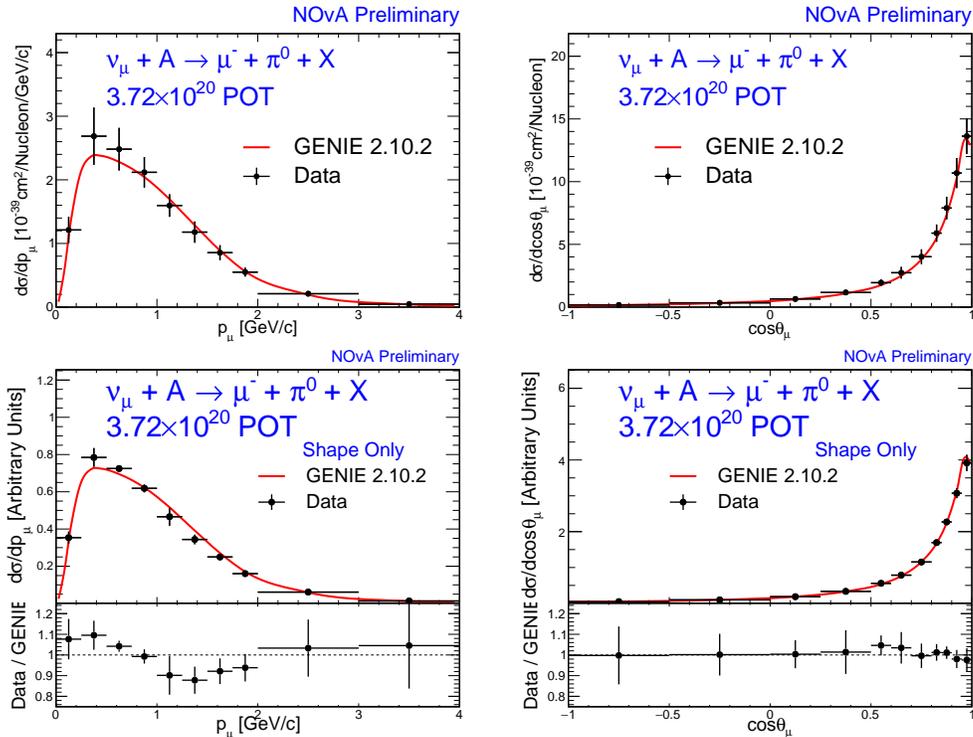


Figure 9.37: The measured cross sections, differentiated in the muon momentum, left, and azimuthal direction, right. The top row show the measured variables – the absolute, flux-averaged cross sections while the bottom plots have been area-normalized with shape-only uncertainties drawn.

dominated by multi- π events, predominantly from GENIE’s DIS model. This points to a hardening of the azimuthal angle for these multi- π events as a potential way to improve the generator’s prediction. This relationship is diametrically opposed to the MINERvA result[107] which studied single meson π^0 production in ν_μ CC events at a higher mean neutrino energy and with $\theta_\mu < 25^\circ$.

Q^2 and W

Cross sections differential in Q^2 and W are shown in Figure 9.40.

There is a deficit of signal events at $Q^2 < 0.1 \text{ GeV}^2$. As alluded to before, nuclear shielding makes this region inaccessible to DIS scatters so that the overall deficit is likely related to a deficit in low- Q^2 resonant events. For the remainder of the distribution, the observed tail is slightly harder than the prediction, as foreshadowed by the excess at low p_μ .

The shape of the W distribution is relatively well modeled, particularly compared to other W measurements in meson production measurements[105][107]. GENIE

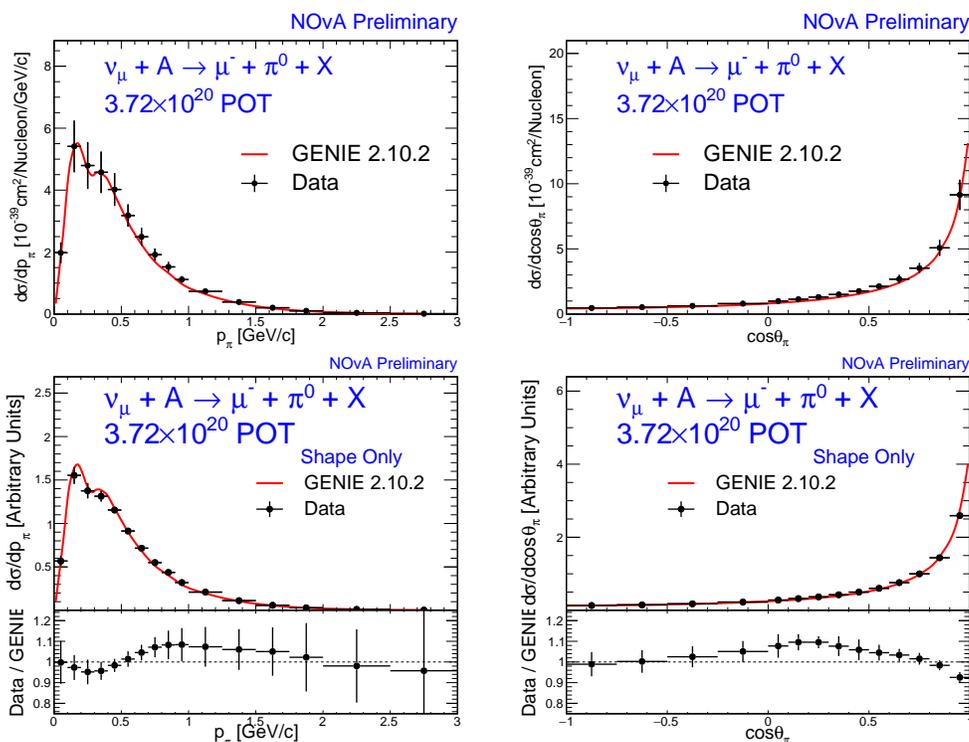


Figure 9.38: The measured cross sections, differentiated in the π^0 momentum, left, and azimuthal direction, right. The top row shows the measured variables – the absolute, flux-averaged cross sections – while the bottom plots have been area-normalized with shape-only uncertainties drawn.

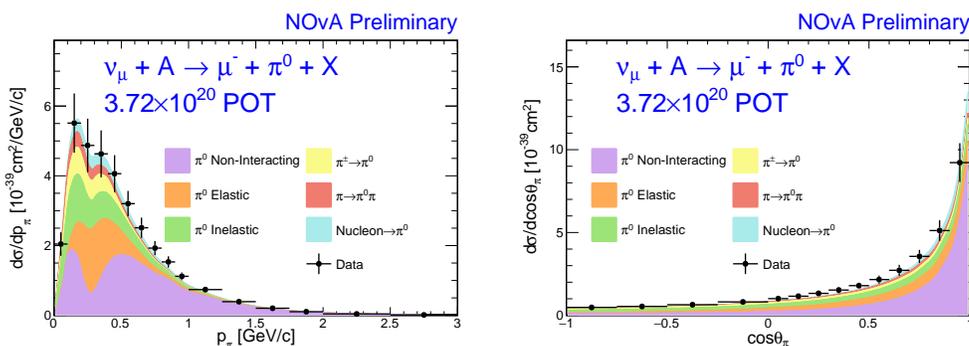


Figure 9.39: The GENIE prediction for the CC π^0 cross section differentiated in p_π , left, and $\cos\theta_\pi$, right. The cross sections are further subdivided into final state interaction scattering types. The purple curve gives the prediction for π^0 's that do not re-scatter within the nucleus after the initial neutrino interaction.

under-predicts between the Δ_{1232} resonance peak and the distribution maximum at $W = 1.6$ GeV. This is precisely the region where higher mass resonances lie, pointing to this sample's potential for focus in future work tuning event generators.

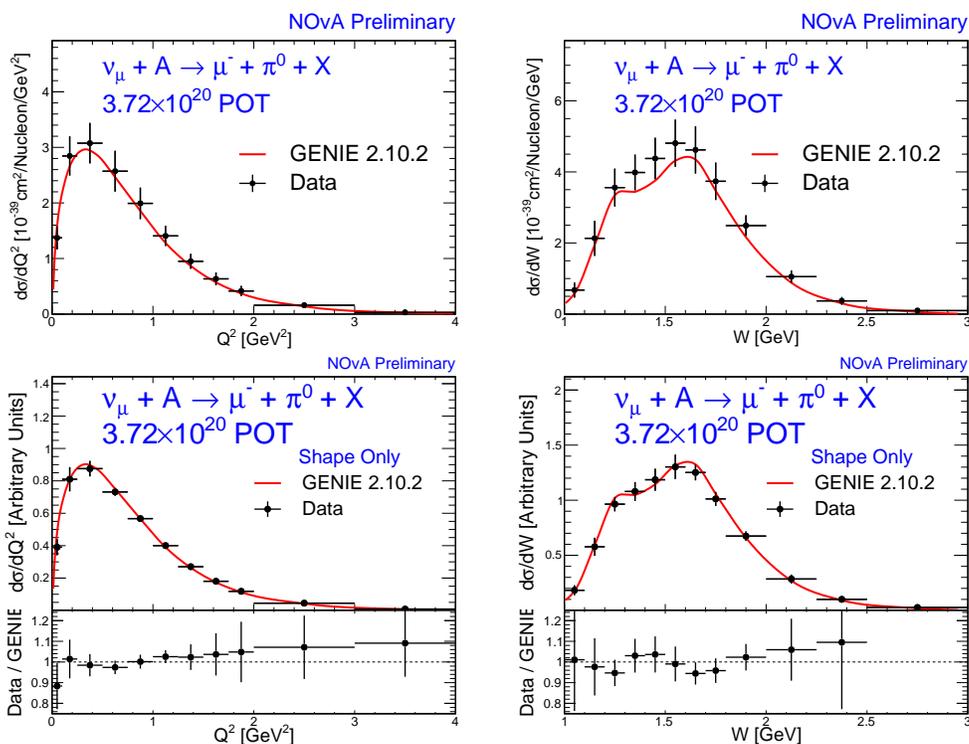


Figure 9.40: The measured cross sections, differentiated in Q^2 , left, and W , right. The top row show the measured variables – the absolute, flux-averaged cross sections while the bottom plots have been area-normalized with shape-only uncertainties drawn.

Spread in $\langle \sigma_x \rangle_\Phi$

Relative excesses in data compared to MC will happen in different regions for each of the differential variables studied since there are modeling issues with the underlying GENIE prediction used in determining the result. This, along with the fact that efficiency curves have shape in each of the differential variables, leads to a different number of predicted signal events for every differential variable and propagates to a different integrated flux-averaged total cross section for each kinematic variable, $\langle \sigma_x \rangle_\Phi$. These integrals are shown in Table 9.4.

| Kinematic Variable | $\langle\sigma_x\rangle_\Phi[10^{-39}\text{cm}^2]$ |
|--------------------|--|
| p_π | 3.49 ± 0.42 |
| $\cos\theta_\pi$ | 3.53 ± 0.42 |
| p_μ | 3.43 ± 0.43 |
| $\cos\theta_\mu$ | 3.48 ± 0.43 |
| Q^2 | 3.51 ± 0.44 |
| W | 3.69 ± 0.43 |
| GENIE | 3.28 |

Table 9.4: Due to discrepancies in the underlying GENIE model, the integrated cross section is different for every differential variable. The center of these measurements is $3.52 \times 10^{-39}\text{cm}^2$, which is 7.3% higher than the GENIE prediction of $3.28 \times 10^{-39}\text{cm}^2$.

9.12 Summary

As described, the $\text{CC}\pi^0$ analysis produced a set of systematically-limited cross section measurements of π^0 production in ν_μ CC events, differential in six kinematic variables. The mean $\langle\sigma_x\rangle_\Phi$ is $3.52 \pm 0.44 \times 10^{-39}\text{cm}^2$ per nucleon, 7.3% higher than the GENIE prediction. The studied energy region directly overlaps the transitional energy range between QE and DIS dominated scattering regimes. This energy region is very relevant for future and current oscillation measurements. Special attention is made to include multi- π events in the measurement signal, as these events have been shown to cause the majority of π^0 background events in $\nu_\mu \rightarrow \nu_e$ oscillation measurements in NOvA.

Kinematic distributions broadly agree with the generated prediction, particularly in the high-level descriptive variables Q^2 and W . However, discrepancies described above show promise for refining neutrino generator models' treatment of nuclear structure, meson production at low Q^2 , and FSI interaction models.

OVERVIEW OF THE OSCILLATION ANALYSIS

The remainder of this work details the combined analysis of $\nu_\mu \rightarrow \nu_\mu$ disappearance and $\nu_\mu \rightarrow \nu_e$ appearance with NOvA. This is the first NOvA analysis where anti-neutrino data is available so that the $\bar{\nu}_\mu \rightarrow \bar{\nu}_e$ probability is incorporated into the analysis, or more generally, the first NOvA analysis using far detector anti-neutrino data. As already described, the fundamental oscillation parameters influence the neutrino and anti-neutrino appearance probabilities differently so that the combined analysis is significantly more sensitive than either channel. Additionally, we use outside data to constrain θ_{13} . This parameter directly factors into the ν_e appearance probability and is precisely determined by reactor experiments. There is also a constraint on the solar oscillation parameters, which weakly influence our predictions and are determined from solar oscillation experiments and KamLAND.

The primary goal of the oscillation analysis is to determine constraints on the dominant physical parameters that govern these oscillation channels and are poorly constrained: $\sin^2 \theta_{23}$, Δm_{32}^2 , and δ_{CP} . The targeted observables for the analysis are the spectra of ν_μ CC, $\bar{\nu}_\mu$ CC, ν_e CC, and $\bar{\nu}_e$ CC events measured in the far detector. These spectra produce constraints on the physics parameters by fitting to a data-driven prediction at various values of oscillation parameters.

The analysis is organized into two stages which build the measurements in parallel. First, the direct observable spectra are measured at the far detector. Then, near detector data is used to constrain the predicted spectra at the far detector. These two are exactly the inputs needed for the likelihood fit we will use to place limits on the physically relevant parameters. The two stages are shown schematically in Figure 10.1, with the blue shaded area corresponding to the far detector measurement while the yellow details the steps used to constrain the far detector prediction with near detector data. These two procedures are outlined below and described in detail in the following two chapters. This will be followed by a description of systematic uncertainties, the log-likelihood fit, and finally the results of the analysis.

The first stage, blue in Figure 10.1, is as simple as applying two selections for isolating ν_μ CC and ν_e CC events in neutrino and $\bar{\nu}_\mu$ CC and $\bar{\nu}_e$ CC events in anti-neutrino data. Often, we will simply refer to ν events corresponding to neutrino

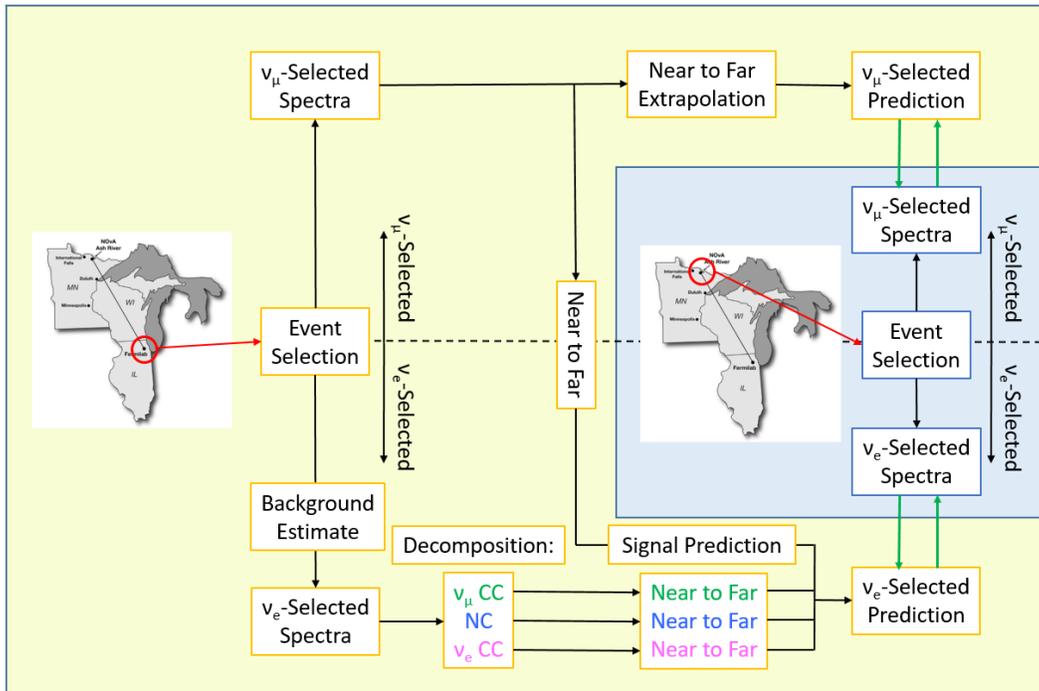


Figure 10.1: An overview of the analysis, showing schematically how the ν_μ disappearance and ν_e appearance spectra are made. Boxes in the blue field relate to the direct far detector observation while those in the yellow relate to building the predictions. The green arrows represent the final fit between the prediction and observation which will place limits on physical parameters.

data while implying an analogous $\bar{\nu}$ sample in the anti-neutrino data. The second stage, yellow, involves a number of specially designed algorithms to constrain the far detector predictions using both ν_μ - and ν_e -selected data. Additionally, the approach to building the far detector ν_μ CC prediction is quite different from the ν_e CC approach. For the ν_μ CC prediction, the process begins by applying a selection to isolate ν_μ CC events in the near detector. The resulting sample is then ultra-pure in ν_μ CC events, so that observed discrepancies between data and simulation in the near detector are propagated forward to the far detector prediction. This process is called the “Near-to-Far Extrapolation”.

For estimating the ν_e appearance spectrum, a ν_e selection in the near detector gives an estimate of the analysis background as the baseline between the beam and near detector is too small to allow oscillations. The situation is complicated by the fact that relatively large components of ν_μ CC, NC, and beam ν_e CC coexist in the sample. The survival rates at the far detector of each of these components is quite different; the ν_μ CC component will near-maximally oscillate in the beam peak so

that they are less represented in the far detector. A decomposition procedure is applied to statistically separate these three true interaction types. Then, a near-to-far extrapolation is applied to separately predict the different true components which are summed to give a beam background prediction.

As the appearance measurement is the $\nu_\mu \rightarrow \nu_e$ oscillation, signal is manifest as $\nu_\mu(\nu_e)$ CC events in the near(far) detector. Thus the ν_μ -selected sample in the near detector is used to constrain the simulated far detector ν_e signal prediction. This is again done with a near-to-far extrapolation. Last, the signal and background predictions are summed, giving the total prediction for the $\nu_\mu \rightarrow \nu_e$ appearance channel, or the CP-conjugate process for the anti-neutrino case.

10.1 Contributing Datasets

Analysis data has accumulated over seven periods over 2014-2018. For data-taking periods 1, 2, 3, and 5, the beam was configured to run neutrinos. These periods give a total of 9.48×10^{20} POT, giving nearly 1.5 nominal-years of running. During much of period 1, and occasional runs afterwards, portions of the detector were inactive due to commissioning and maintenance. The accumulated exposure was calculated as 8.85×10^{20} POT-equiv[124], where POT-equiv translates the analysis exposure into an equivalent exposure of full-detector running.

Anti-neutrino running was split between periods 4, 6, and 7. The early anti-neutrino running in period 4 gave quite a small dataset that has been available but not analyzed for some time. This was run to test the MC's performance on simulating anti-neutrino events in the near detector. Integrated over all periods, 6.91×10^{20} POT-equiv of anti-neutrino data was recorded for this analysis.

10.2 Blindness Policy for the Oscillation Analysis

NOvA enacted a set of guidelines to ensure far detector analyses are properly blinded before analysis procedures have been finalized. This safeguards analyses against potential ethical issues relating to adapting the analysis to conform with any personal bias.

Within this policy, all near detector data is unrestricted, as understanding this data is vital to constraining far detector simulated predictions. Reconstructed slices at the far detector are considered “in-the-box” if they lie within $10 \mu\text{s}$ of the NuMI spill window and have a calorimetric energy between 0.8 and 3 GeV. Before an analysis effort may open the box, all outstanding issues or concerns brought up by

collaboration members must be addressed. And, upon opening the box, the observed far detector spectra of interest can not be modified.

Chapter 11

FAR DETECTOR EVENT SELECTION

For each oscillation channel, there is a separate selection leading to four observed neutrino spectra at the far detector: ν_μ survival, $\bar{\nu}_\mu$ survival, ν_e appearance, and $\bar{\nu}_e$ appearance. Each of these consists of five cut stages: timing cuts, spill and data quality, preselection, cosmic rejection, and PID. The ν_μ and $\bar{\nu}_\mu$ selections are identical while the ν_e and $\bar{\nu}_e$ only differ by small changes in the tuned cut values. As such, the $\nu_\mu/\bar{\nu}_\mu$ selections will simply be referred to as the ν_μ selection, and the ν_e selection will be adopted as a shorthand for the $\nu_e/\bar{\nu}_e$ with the properly tuned parameters modified for each dataset.

11.1 ν_μ Selection

The ν_μ selection strategy is illustrated in Figure 11.1. There are four basic cut stages. The last step in the chain is not a cut, but rather sorts the events into bins of the fraction of neutrino energy belonging to the hadronic shower, which will be justified later.

Basic Quality Cuts

First, events are required to lie within the beam window. Each beam spill records 500 μs of data. The beam spill lies between 218 and 228 μs . Therefore, the time of the slice must have

- $217 < t < 229 \mu\text{s}$.

An extra μs has been added as padding to ensure that all neutrinos reconstruct within the selected window. According to simulation, this is a very conservative boundary.

The large time chunk recorded in the 500 μs spill that is not in the beam window is the “cosmic window” which features all events that have

- $25 < t < 217 \mu\text{s}$ or $229 < t < 475 \mu\text{s}$.

Events in the cosmic window are used for an in-situ estimate of the cosmic background selected in the beam window.

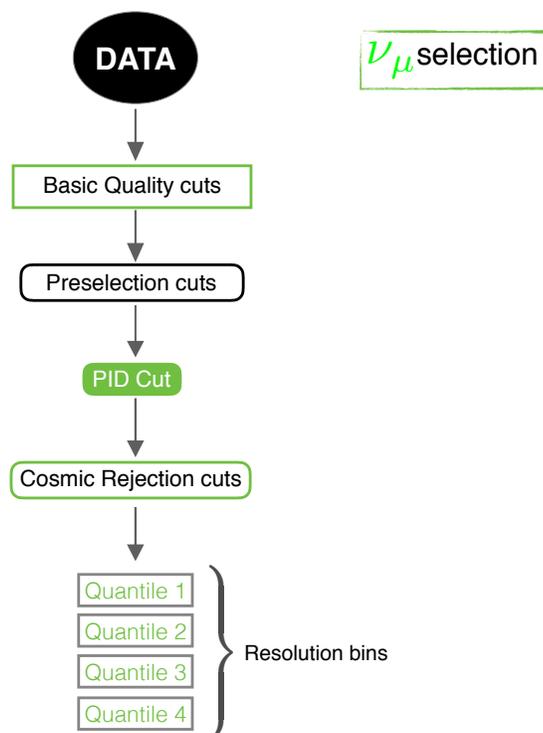


Figure 11.1: A schematic of the cutflow for the ν_μ selection in either neutrino or anti-neutrino running. After passing basic quality cuts, the events are then required to pass preselection cuts. Two PID's are then applied to remove, first, beam backgrounds and, second, cosmic contamination in the sample. After this, selected events are grouped into quartiles according to the fraction of energy observed to be in the hadronic shower.

Spill cuts identical to the spill cuts applied for the $CC\pi^0$ analysis, as listed in Section 9.4, are then applied. Again, these are included to insure that the beam configuration lies within normal operational ranges in which we are confident we can suitably simulate the resulting flux. For the data quality cuts, events must satisfy

- $N_{\text{hit}} > 20$
- Number of planes > 4
- Number of tracks ≥ 1
- Valid ReMId score with reconstructed energy
- Reconstructed neutrino energy < 5 GeV

Preselection

Preselection cuts are applied to loosely reject cosmic backgrounds without significantly reducing the signal efficiency. These are based on reconstructing the distance between the muon track endpoints and the detector faces when projected along the reconstructed track direction and the distances between all reconstructed prong endpoints and each of the detector edges.

- The projected distance along the tagged muon between the track endpoint and the detector edge > 6 cm
- The projected distance along the tagged muon between the track start and the detector edge > 6 cm
- All prongs must reconstruct > 60 cm from the top detector face
- All prongs must reconstruct > 12 cm from the bottom detector face
- All prongs must reconstruct > 16 cm from the east detector face
- All prongs must reconstruct > 12 cm from the west detector face
- All prongs must reconstruct > 18 cm from the upstream detector face
- All prongs must reconstruct > 18 cm from the downstream detector face

Preselection cuts reduce the number of cosmic events by a factor of 1000 from about 10^7 to about 10^4 events. Again, the cosmic background is estimated from out-of-time recorded data events. Further reduction is achieved through PID cuts.

PID

A PID cut on two variables is applied to reduce the beam background, dominated by NC events. As a side-effect, this cut also reduces the otherwise-selected cosmic background. ReMId, the track-based muon identifier is included, as is CVN's match to ν_μ CC events, CVNm. The cut requires

- ReMId > 0.7
- CVNm > 0.7

Cosmic Rejection

An additional BDT was trained to further reduce the cosmic background. This is based on six variables:

- $\cos \theta_y$ on the muon track
- The track length of the muon track
- The maximum y value achieved by the track
- CVN's match to cosmic events
- The fraction of slice hits on the muon track
- The minimum distance to a detector face from the muon track, projected along the track

The resulting BDT must reconstruct to be greater than 0.53. The BDT distribution for beam and cosmic events is shown in Figure 11.2. After the PID and cosmic rejection cuts, there are only a handful of beam and cosmic background events expected in the sample, as shown in Figure 11.3.

Hadronic Energy Fraction Bins

The ν_μ distribution is grouped into four different analysis bins, separated by the fraction of the neutrino energy that goes into the hadronic shower, HadEFrac. As

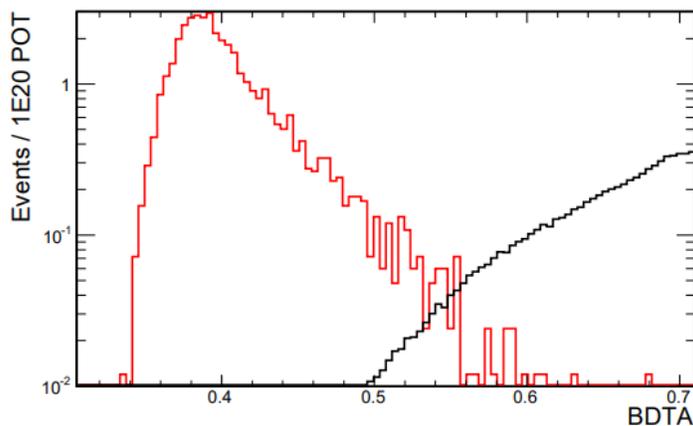


Figure 11.2: The trained cosmic-rejection BDT, properly scaled to exposure. The response to cosmic data is shown in red while simulated beam events are shown in black.

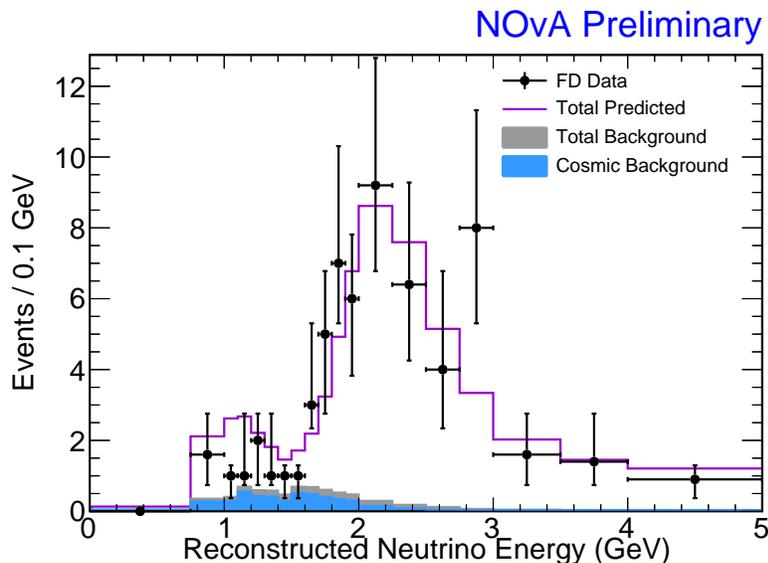


Figure 11.3: ν_μ -Selected events in the neutrino sample. The blue, gray, and purple stacked histograms give the cosmic, beam background, and signal components, respectively.

the muon energy is much more precisely reconstructed than the hadronic energy, this is a proxy for the energy resolution of events. The variable is separated into four bins, whose boundaries depend on the reconstructed neutrino energy. Within each bin of neutrino energy, the bin boundaries are taken as the quartile markers in HadEFrac so that the sample is split into four equally-sized populations. Figure 11.3 shows the selected far detector ν_μ and $\bar{\nu}_\mu$ events summed over all HadEFrac bins. As illustrated, the $\bar{\nu}_\mu$ events lie at lower HadEFrac. This is due to the anti-neutrino cross section being skewed, preferring more elastic. The neutrino energy resolution in the various quartiles range from 6% to 12%. Separating the selected events in this way in the fit gives the analysis more statistical power; it effectively improves the neutrino energy resolution which makes the characteristic dip near the oscillation maximum better resolved. The distinction between the two different sets of oscillation parameters is starkest in the lowest HadEFrac bin. Thus, the analysis has more power-per-event for those events that reconstruct in this bin.

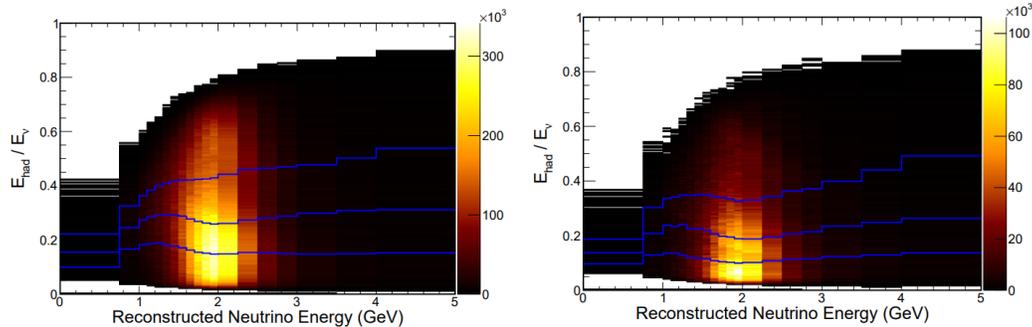


Figure 11.4: The distribution of HadEfrac as a function of reconstructed neutrino energy for ν_μ CC, left, and $\bar{\nu}_\mu$ CC events, right. The blue lines give the quartile boundaries, which define the energy dependent HadEfrac bin boundaries.

11.2 ν_e Selection

The ν_e selection strategy is shown in Figure 11.5. Since there are so few events, the selection strategy is two-pronged. There is a “core” sample which uses a selection stream very similar to that used to select ν_μ CC events. These events pass a strict containment and preselection cut. To regain some of the signal events lost with these cuts, the “peripheral” sample sends those events with a high PID value that fail the containment or preselection cuts and passes them through a less strict containment cut. Each of these is described below, first running through the core selection followed by the peripheral.

Basic Quality Cuts

The selection begins with the same timing and spill cuts used in the ν_μ selection. Additionally, basic reconstruction quality cuts are applied:

- Average number of hits per plane < 8
- Event reconstructed a vertex
- At least one prong reconstructed from vertex

Where the first cut is added to cut FEB flasher events.

After basic quality cuts, the cutflow bifurcates into “core” and “peripheral” selection streams. The core sample is analogous to the ν_μ selection that applies preselection, cosmic rejection, and a PID cut. The “peripheral” sample considers all events that fail the preselection or cosmic rejection cuts, and applies a looser cosmic rejection

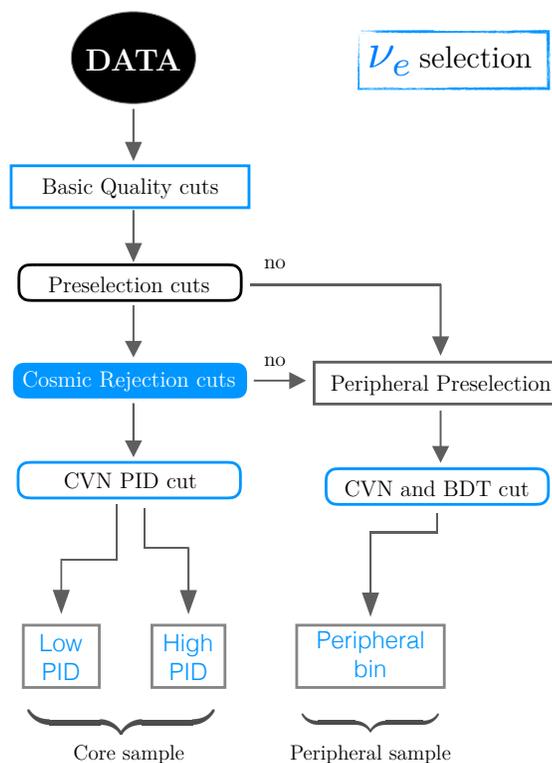


Figure 11.5: The two-pronged selection strategy used by for the ν_e appearance result. The left portion of the cutflow is analogous to the ν_μ selection strategy shown in Figure 11.1. The bottom operation is not a cut, but rather sorts events into a high and low PID bin. The right shows the peripheral selection. Events that fail the initial preselection or cosmic rejection cuts are given a second chance for selection to optimize the event efficiency.

cut in a tighter PID region to regain much of the signal initially lost to containment. These two branches of the cutflow are described below.

Core Preselection

The first preselection cuts restrict the event kinematics to the relevant region for ν_e appearance. They constrain the reconstructed number of hits, neutrino energy, and the longest prong reconstructed in the event:

- $30 < N_{\text{hit}} < 150$
- $1 < E_\nu < 4$
- $150 < \text{longest reconstructed prong} < 150 \text{ cm}$

Core Cosmic Rejection

Next, a containment cut is applied, very similar to the containment used in the ν_μ selection, based on reconstructed distance between prongs and the detector walls:

- All prongs must reconstruct > 60 cm from the top detector face
- All prongs must reconstruct > 12 cm from the bottom detector face
- All prongs must reconstruct > 16 cm from the east detector face
- All prongs must reconstruct > 12 cm from the west detector face
- All prongs must reconstruct > 18 cm from the upstream detector face
- All prongs must reconstruct > 18 cm from the downstream detector face

Next, there is a cut on the reconstructed transverse momentum fraction, p_T/p . In a CC event, this ought to be 0, and any departure indicates a measure of “mis-reconstructedness” of the event. But, there is no such restriction for cosmic events. The cut depends on the maximum y coordinate reached by the slice, as the cosmic activity rate is much higher near the top of the detector. The cosmic distribution is shown in Figure 11.6, overlaid with the simulated beam prediction. The cut is determined to be

- $p_T/p < 0.58$ if $\max-y > 590$
- $p_T/p < 0.80$ if $590 > \max-y > 350$
- $p_T/p < 1.00$ if $350 > \max-y$

There is one final cut designed to remove cosmic photons that scatter into the downstream portion of the detector that are traveling anti-parallel to the beam. The cut uses the “sparseness asymmetry”, which calculates the density of hits in the first and second halves of the most energetic prong. Due to electromagnetic shower structure, there should be a higher hit density in the second half of the shower. So, for events that are within 200 cm of the downstream wall, we apply

- $(\text{Hit density in downstream half} - \text{hit density in upstream half})/(\text{hit density}) > -10\%$

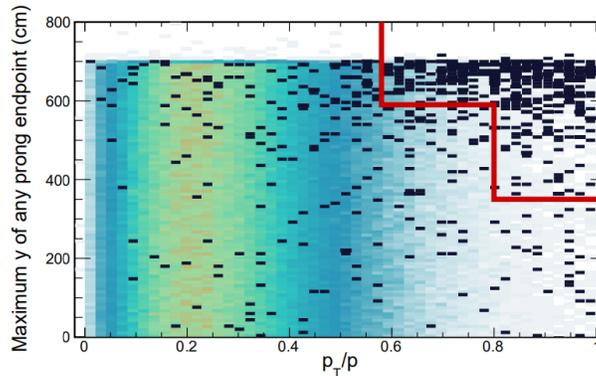


Figure 11.6: The distribution of the reconstructed distance to the top of the detector vs p_T/p for signal, the colored distribution, and cosmic events, the black data points. The distribution is shown for neutrino data and looks very similar in anti-neutrino mode with the same tuned cut.

Core PID Cut and Bins

The CVN match to ν_e CC events, CVNe, is used as the principle PID for the analysis. Using a FOM to optimize the statistical power of appearance signal in simulation, analysis cuts are placed on the CVNe value. Additionally, the analysis sample is separated into two analysis bins according to the reconstructed CVNe value. This is similar to the HadEFrac binning in the disappearance measurement. Binning in CVNe groups together samples of similar purity which offers more statistical power for measuring the appearance probability over grouping the entire sample together. The curves used in tuning the analysis bin boundaries in CVNe are shown in Figure 11.7. This gives

- $0.84(0.89) < \text{CVNe} < 0.96(0.98)$

for the low-PID bin and

- $0.96(0.98) < \text{CVNe} < 1.00$

for the high-PID bin for $\nu_e(\bar{\nu}_e)$ appearance.

This gives the full list of cuts for the core selection. Additionally, within each PID bin, the events are plotted as a function of reconstructed neutrino energy to mitigate shape-dependent systematic effects when applying the near detector constraint.

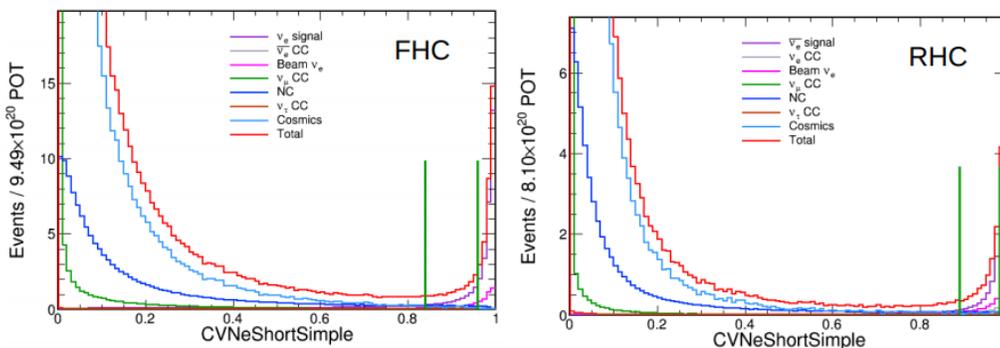


Figure 11.7: The cut tunings for ν_e , left, and $\bar{\nu}_e$, right, appearance. In the plot, the purple shows appearance signal, with other beam and cosmic backgrounds overlaid. The PID cuts for the sample are tuned to 0.84 and 0.89 for neutrino and anti-neutrino running separately.

Peripheral Selection

As illustrated in Figure 11.5, the peripheral sample selection begins with all events that pass basic quality cuts but fail at least one core preselection and cosmic rejection cut. As expected, the events passing this stage of the peripheral sample are overwhelmingly dominated by cosmic events. But, the CVNe distribution measured on cosmic data falls sharply, so that at very high values of CVN, the simulation predicts several signal events on a background of tens of cosmic events.

To isolate the beam component under this large cosmic component, a BDT was trained on five variables used in the core preselection and cosmic rejection cuts:

- N_{hit}
- The minimum reconstructed distance from a prong endpoint to the detector top face
- The minimum reconstructed distance from a prong endpoint to the detector face other than the top
- p_T/p
- The sparseness asymmetry

The BDT for simulated beam and measured cosmic distribution, when applied to peripheral and core preselected events for neutrino events, is shown in Figure 11.8. In this figure, both basic quality cuts have been applied along with a relatively loose $\text{CVNe} > 0.75$ restriction.

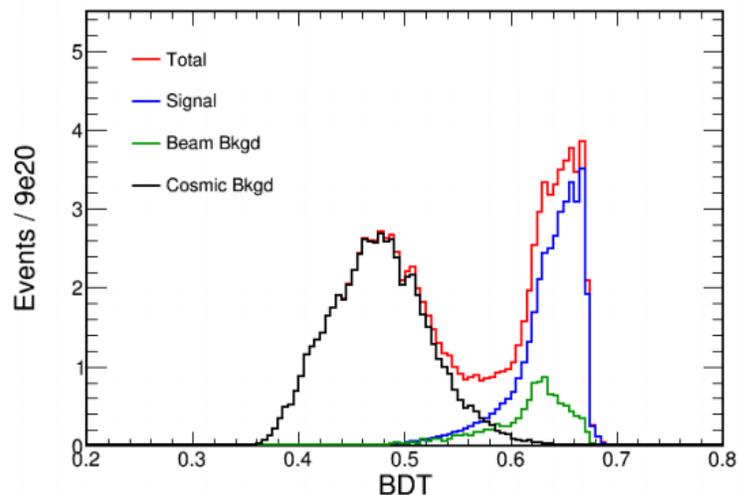


Figure 11.8: The cosmic BDT’s response to cosmic data, black, and simulated signal, blue, and beam background, green. Beam events in neutrino-running are represented here.

As shown in Figure 11.9, signal events are only noticeable at the very highest CVNe values. So, the peripheral sample requires a harsher cut on CVNe to compensate a looser cosmic rejection restriction, only taking events with $CVNe > 0.96(0.98)$ for $\nu_e(\bar{\nu}_e)$ appearance. This corresponds exactly to the high-PID bin restriction in the core sample. After this restriction, a BDT cut was tuned to optimize the significance of the signal in the peripheral sample. The tuned cut is shown in Figure 11.9 for both neutrino and anti-neutrino data. For the neutrino tuning, we require

- $BDT > 0.53$ **or** $CVNe > 0.99$

and for anti-neutrinos,

- $BDT > 0.57$

Unlike in the core sample, the peripheral sample is not analyzed in bins of energy, rather all selected events are grouped into a single bin. This is principally a precaution as peripheral events fail some core cut, and thus are more likely to accumulate near the edges of the detector where calibration uncertainties are more significant.

Including the peripheral sample in the analysis is a way to increase the effective fiducial volume of the detector. The relative effective fiducial volume between two setups can be compared with the statistical $FOM^2 = s^2/(s + b)$. Doing this, the

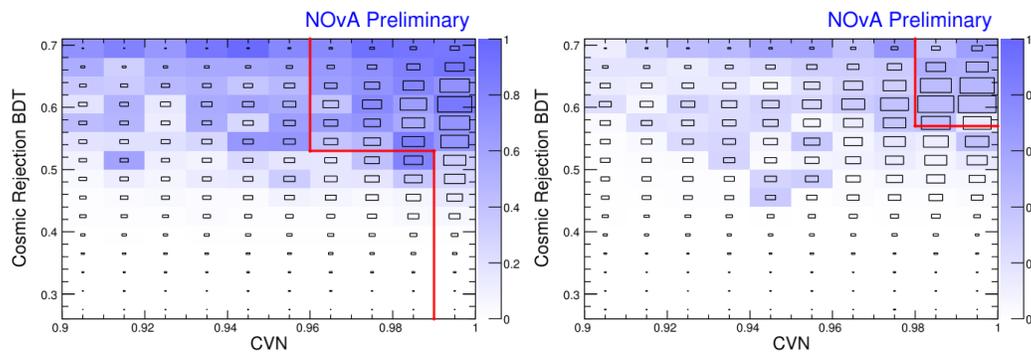


Figure 11.9: The tuning of the BDT cuts for the neutrino, left, and anti-neutrino, right. The color represents the signal purity in each bin, while the hollow boxes indicate where our signal events lie.

peripheral sample adds 15% fiducial mass to the core-only selection for neutrino running 11% for anti-neutrino running.

Chapter 12

CONSTRAINING THE FAR DETECTOR PREDICTION WITH NEAR DETECTOR DATA

Due to relatively imprecise knowledge of neutrino cross sections and fluxes at the relevant energies, NOvA has a near detector near the beam source to constrain systematic uncertainties in the far detector prediction. The near detector measures the response of reconstruction and PID techniques on a large sample of neutrino events. Observed discrepancies between data and simulation are propagated forward to the far detector where they are used to adjust the oscillated spectra predicted by the base simulation.

The ν_μ disappearance and ν_e appearance channels use slightly different extrapolation techniques. The specifics of each will be detailed in the following. The process starts with adapting the far detector ν_e and ν_μ selection for use in the near detector. Then, comparisons are drawn between data and simulation that are then used to constrain signal and backgrounds in a channel-specific way.

For both the disappearance and appearance predictions, ν_μ CC events selected in the near detector are used to constrain the signal prediction in the far detector. Though the appearance analysis is searching for ν_e CC events in the far detector, the signal is a $\nu_\mu \rightarrow \nu_e$ oscillation and the ν_μ CC observed in the near detector gives a constraint on this signal.

12.1 Near Detector ν_μ CC Selection

The spill quality, data quality, and PID cuts for the near detector ν_μ selections are identical to those used in the far detector. There is no cosmic rejection cut in the near detector due to ~ 100 m rock overburden. The containment cuts are, however, significantly different as the two detectors have different geometries. In the near detector, this is separated into a muon and hadronic containment cut. For the muon containment, we require:

- No activity in two most-upstream and most-downstream planes of the detector
- Reconstructed muon track must start in first 11 m of the detector

- Either the muon does not enter the muon catcher, or the air gap, defined below, is < 55 cm
- The projected distance along the tagged muon between the track endpoint and the detector edge > 5 cm
- The projected distance along the tagged muon between the track start and the detector edge > 10 cm

For reference, the boundary between the fully-instrumented portion of the detector and the muon catcher is at 12.7 m. The air gap arises because the muon catcher is not as tall as the fully-instrumented detector so that a track with $p_y < 0$ may originate in the fully instrumented portion of the detector, exit the detector, and re-enter in the muon catcher. The tracklength outside of the detector is the air gap.

To contain the hadronic shower, we require

- All prongs must reconstruct > 25 cm from the downstream detector face
- All prongs must reconstruct > 20 cm from any other detector face
- The muon track is the only track allowed to enter the muon catcher

After this, a $CVNm > 0.7$ and $ReMid > 0.7$ cut is applied as in the far detector. As in the far detector, the sample is subsequently binned in HadEFrac. The near detector sample uses the same energy-dependent bins that separate the far detector sample into quartiles. Correcting the simulation in HadEFrac bins is an important feature, as it implicitly reduces systematic uncertainties that affect neutrino interactions in a localized region of the HadEFrac range.

Relevant kinematic variables for the selected sample are plotted in Figures 12.2-12.7. The overall normalizations agree quite well across HadEFrac bins, as shown in Figure 12.1. There are 1.90×10^6 and 0.39×10^6 events selected in neutrino and anti-neutrino data, respectively[125]. The background, almost entirely NC, represents only 0.34% of the simulated sample. In the plots below, the shape-only comparison has been drawn to highlight the observed distortions in shape which are covered by our systematic uncertainty.

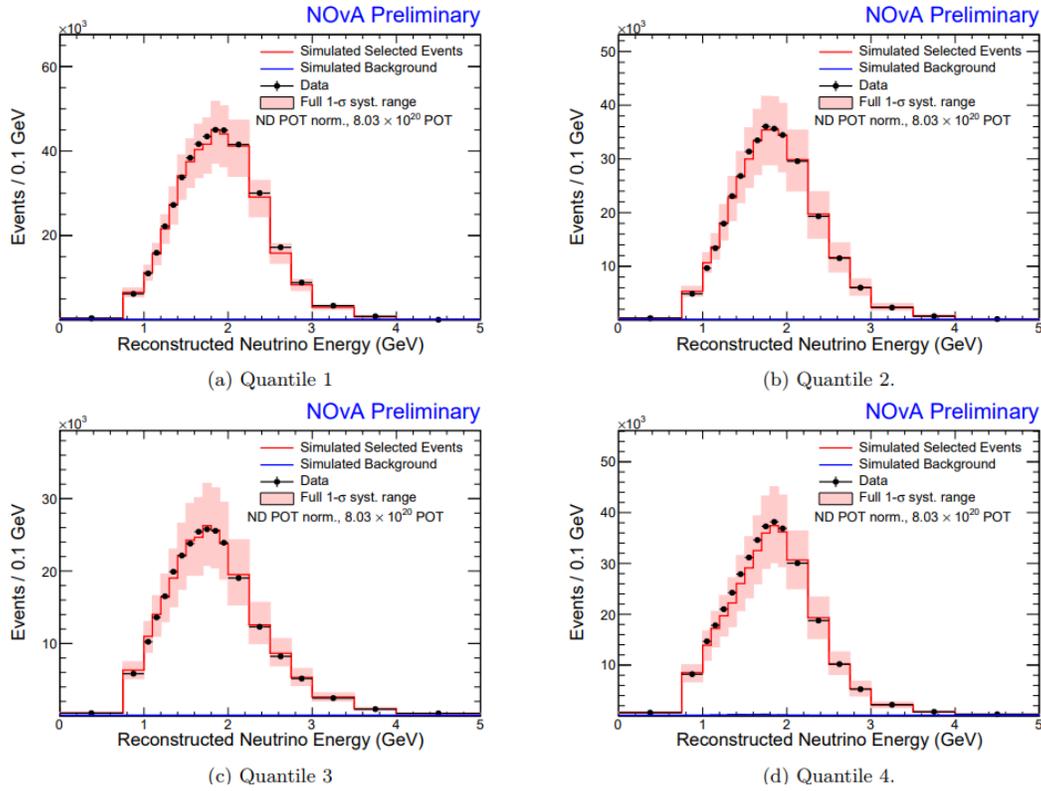


Figure 12.1: A comparison of data and simulation for the reconstructed neutrino energy for selected ν_μ CC. The individual plots give the comparison in each HadEfrac bin.

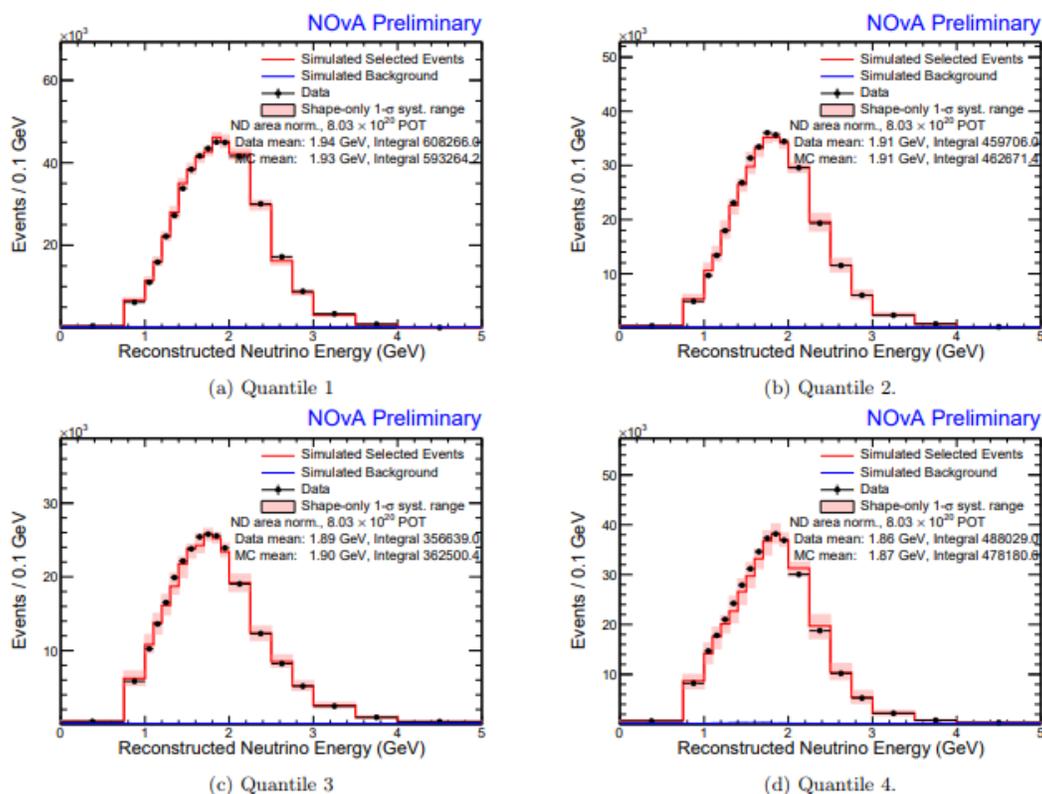


Figure 12.2: A comparison of data and simulation for the reconstructed neutrino energy for selected ν_μ CC. The individual plots give the comparison in each HadEfrac bin. The error bars on the simulation give the shape-only error.

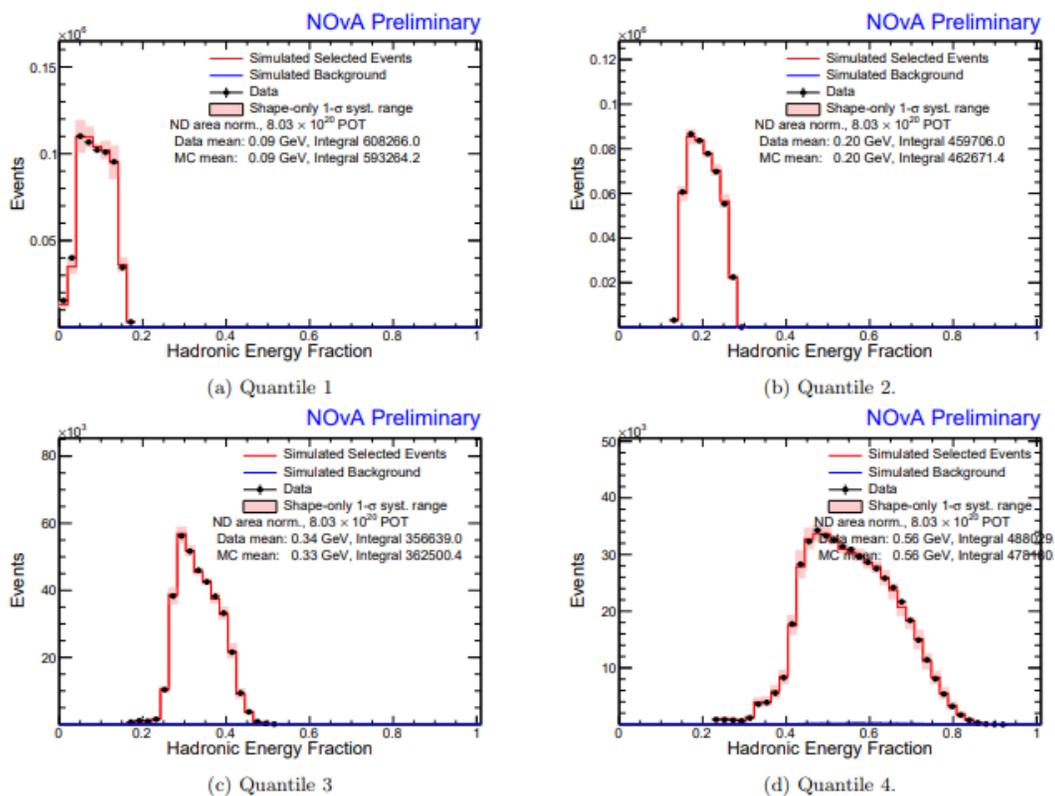


Figure 12.3: A comparison of data and simulation for the reconstructed HadEfrac for selected ν_μ CC. The individual plots give the comparison in each HadEfrac bin. The error bars on the simulation give the shape-only error.

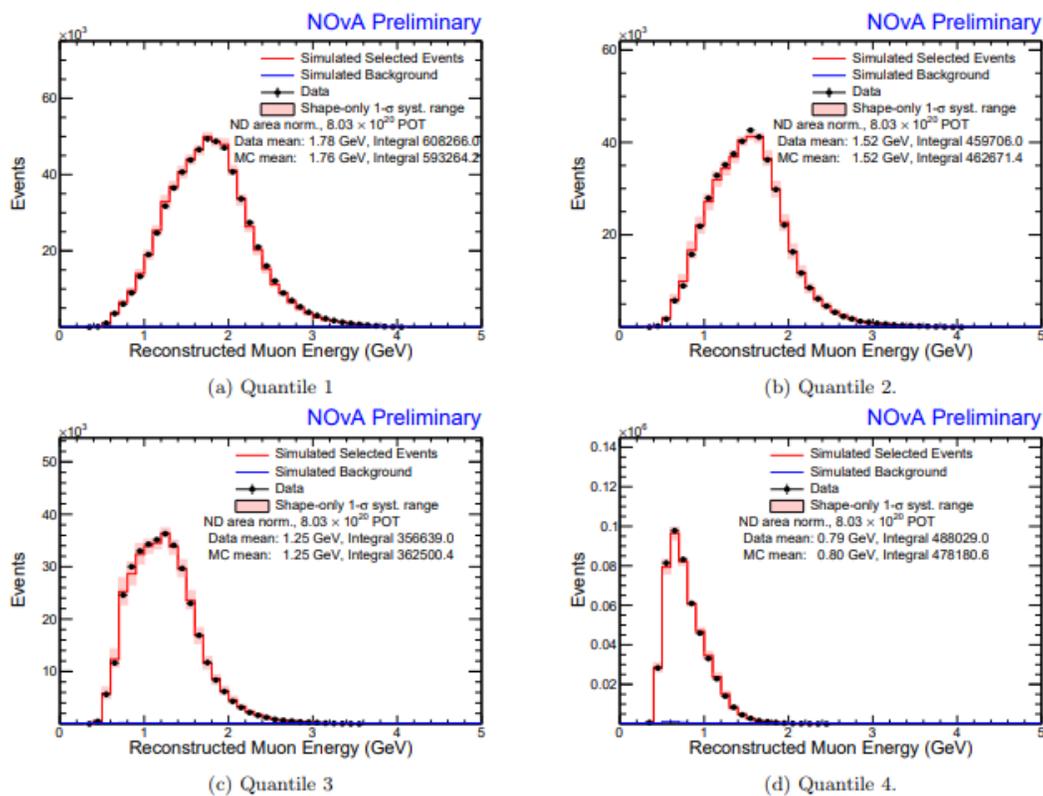


Figure 12.4: A comparison of data and simulation for the reconstructed muon energy for selected ν_μ CC. The individual plots give the comparison in each HadEFrac bin. The error bars on the simulation give the shape-only error.

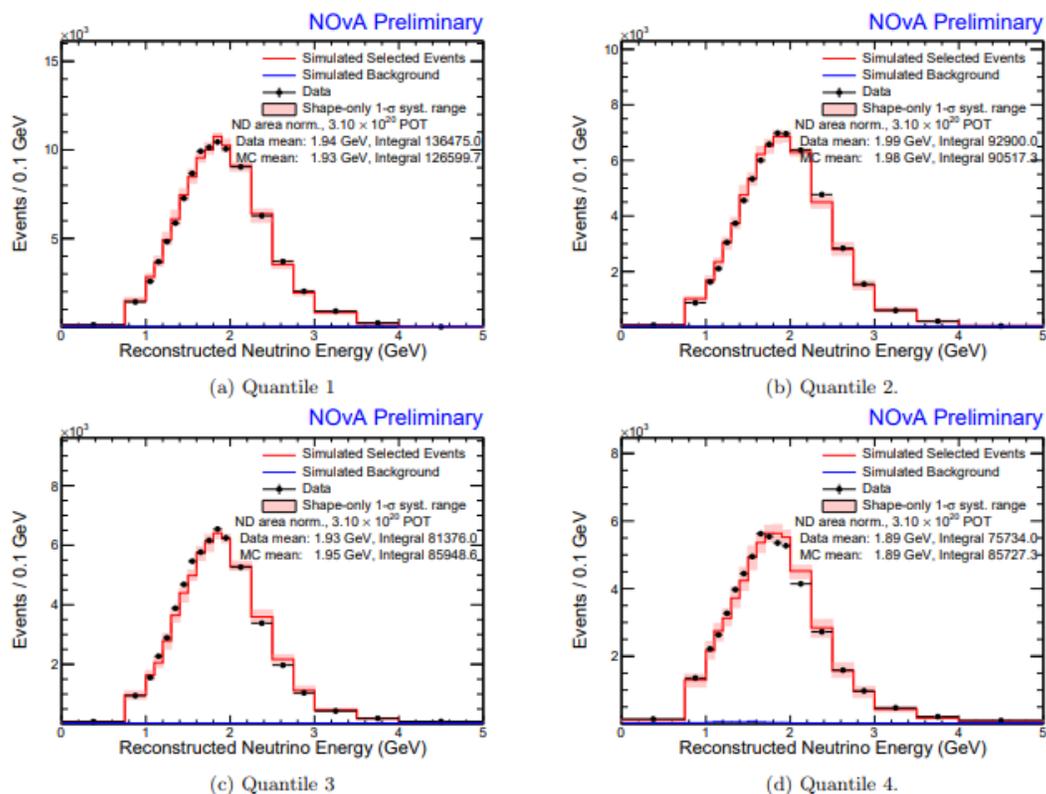


Figure 12.5: A comparison of data and simulation for the reconstructed neutrino energy for selected $\bar{\nu}_\mu$ CC. The individual plots give the comparison in each HadEfrac bin. The error bars on the simulation give the shape-only error.

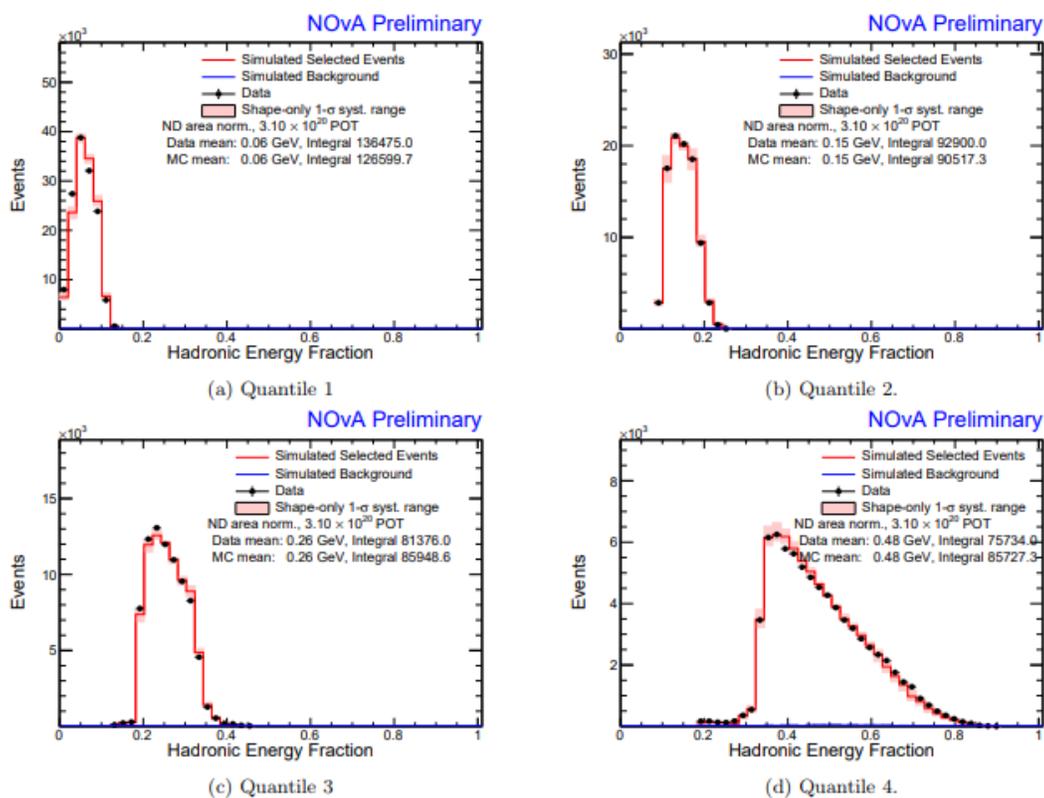


Figure 12.6: A comparison of data and simulation for the reconstructed HadEFrac for selected $\bar{\nu}_\mu$ CC. The individual plots give the comparison in each HadEFrac bin. The error bars on the simulation give the shape-only error.

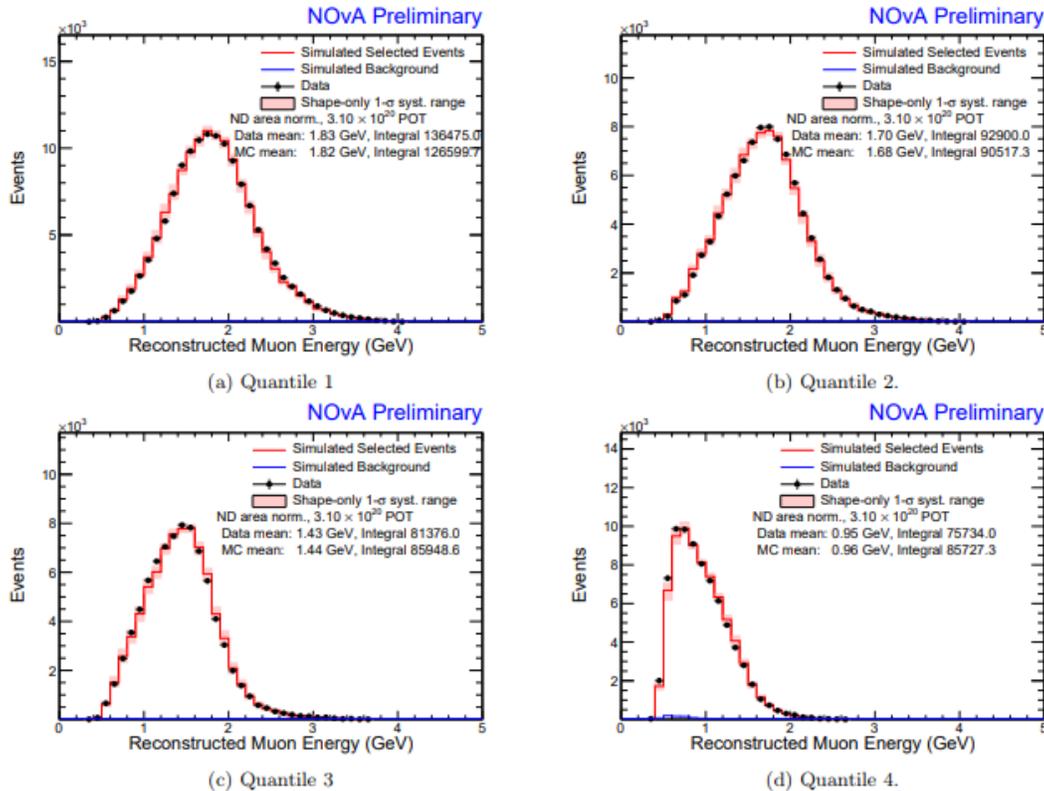


Figure 12.7: A comparison of data and simulation for the reconstructed muon energy for selected $\bar{\nu}_\mu$ CC. The individual plots give the comparison in each HadEfrac bin. The error bar on the simulation gives the shape-only error.

12.2 Near Detector ν_e CC Selection

The near detector ν_e selection principally differs from the far detector analogue in containment strategy. As such, the spill and basic quality cuts are identical to those used at the far detector.

The near detector includes a fiducial cut on the reconstructed vertex. This is added to reduce the selection rate of rock-induced slices[126]. These rock-induced interactions are dominated by single photon showers originating from π^0 decays in high energy neutrino interactions that convert within the fiducial volume. The fiducial volume ensures that there is self-shielding from these events. As the total flux is lower further off-axis, the event rate is significantly higher in the rock at low x and high y in detector coordinates. Therefore, the fiducial volume is off-center from the detector so that there is more shielding in the direction where the flux is increased. We require the vertex be reconstructed with

- $-100 < x < 160$ cm

- $-160 < y < 100$ cm
- $150 < z < 900$ cm

For reference, the detector cross section is about 4×4 m². The end of the fully-instrumented portion of the detector and beginning of the muon catcher is at 12.7 m.

Preselection cuts are analogous as those used at the far detector, though they are somewhat looser, as the preselection cuts must cover both the far detector core and peripheral samples.

- $20 < N_{\text{hit}} < 200$
- $0 < E_\nu < 4.5$ GeV
- $100 < \text{longest reconstructed prong} < 500$ cm

The containment cuts in the near detector are very similar to those at the far detector, largely based on the reconstructed distances between the detector faces and the reconstructed prong endpoints. However, though containment requirements at the far detector are principally designed to reject cosmic events, in the near detector, these cuts are imposed to minimize the escape of visible energy from a neutrino interaction from the detector. Since electromagnetic showers are much easier to contain than muons, the ν_e selection is restricted to the fully-instrumented portion of the detector to avoid the relatively poor energy reconstruction of events passing through the muon catcher.

- No activity in the first six planes of the detector
- All prongs must reconstruct > 100 cm from the upstream detector face
- All prongs must reconstruct > 30 cm from any other detector face
- All prongs must reconstruct > 50 cm from end of the fully-instrumented portion of the detector

Last, there is the PID cut on CVNe as in the far detector. This is the only cut that depends on whether the beam was configured to run neutrinos or anti-neutrinos. We require $\text{CVNe} > 0.84(0.89)$ for the $\nu_e(\bar{\nu}_e)$ selections, as in the far detector. Also, the sample is subdivided into two CVNe analysis bins with the same bin boundaries as in the far detector.

12.3 Near Detector ν_e Decomposition

The near detector ν_e sample is, by definition, a representation of the beam backgrounds for the appearance analysis, in both neutrino and anti-neutrino mode. As shown in Figure 12.8, the ν_e sample is somewhat equally composed of ν_μ CC, NC, and ν_e CC interactions. This is a major disparity with the $\bar{\nu}_e$ sample, which is primarily composed of $\bar{\nu}_e$ CC and ν_e CC events.

An immediate question relates to how the ν_e -selected neutrino data sub-divides into these three different background classes. Additionally, the question has a direct impact on the precision of the appearance analysis. The far detector prediction for each of these background components is the amount of a given process observed in the near detector convolved with the energy-dependent survival probability and acceptance differences. The baseline was optimized to make the ν_μ disappearance probability ≈ 1 , so that this component will be suppressed in the far detector relative to its near detector representation. For comparison, the NC component will be unaffected by oscillations.

We have three strategies for separating these three components in the ν_e -selected near detector sample. One relies on comparing the number of reconstructed Michel electrons in data and MC. This effectively separates the ν_μ CC component from the NC+ ν_e CC component, but offers very little information about splitting the NC and ν_e CC components. To lift this degeneracy, the ν_e CC component is estimated according to observed differences in flux ratios from ν_μ CC events. The Combo Decomposition using these two algorithms together is used for our neutrino sample.

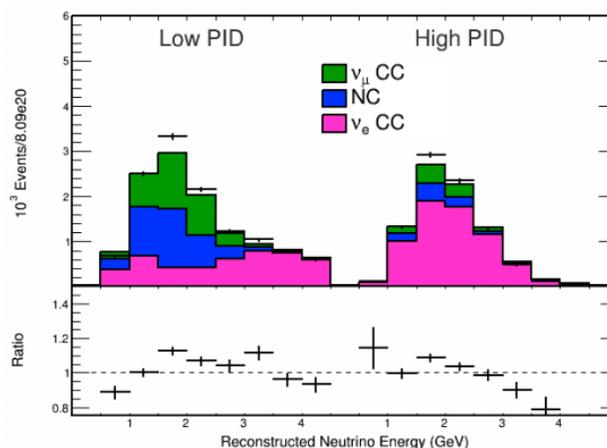


Figure 12.8: The ν_e CC simulated prediction at the near detector, decomposed by ν_μ CC, NC, and ν_e CC events.

The third is a trivial decomposition which does not constrain the relative proportions of each of these subcomponents but scales all MC subcomponents by the same scale so that total MC matches the observed data. The scale factor is applied separately in each analysis bin. This is called the Proportional Decomposition and is applied to the current anti-neutrino analysis. This dataset is far smaller than the neutrino dataset. Therefore, induced statistical errors of more sophisticated decomposition algorithms are larger than the reduction in systematic errors they offer.

12.4 BEam Nue Decomposition (BENDecomp)

The intrinsic ν_e flux in the beam is constrained by observed ν_μ events. The majority of ν_e flux comes from the $\pi^+ \rightarrow \mu^+ \nu_\mu \rightarrow e^+ \nu_e \bar{\nu}_\mu \nu_\mu$ decay chain. But, there is a significant contribution from K mesons. Though K mesons are much less commonly produced in the target than π 's, the fraction of ν_e flux that comes from K decays is relatively high as the $\pi^+ \rightarrow e^+ \nu_e$ branching ratio is so low, $\sim 10^{-4}$. Thus, the ν_e CC component is adjusted according to the observed ν_μ CC selected in the near detector. The ν_e CC prediction at the near detector is shown in Figure 12.9 differentiated by the ancestor particle type. Events with a K ancestor meson dominate at high energy while in the 1-4 GeV range, where oscillations occur, there is a significant contribution from both K and π ancestors.

In BENDecomp, the ν_μ -like events are separated into contained and uncontained subsamples[127]. A comparison of data and simulation is shown in Figure 12.10. As seen in the figure, the contained ν_μ CC subsample has negligible contamination from K ancestors. The uncontained subsample, however, is dominated by K -ancestor

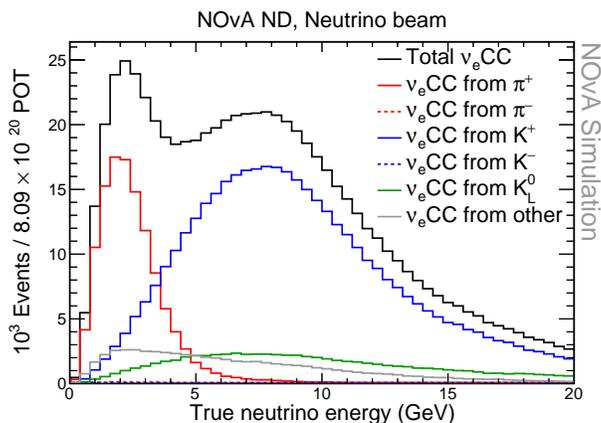


Figure 12.9: The ν_e CC simulated prediction at the near detector, decomposed by the meson ancestor produced in the target. K ancestors are shown in blue with π ancestors in red.

events over a reconstructed neutrino energy of 4.5 GeV. BENDecomps constrains the simulated ν_e CC population in the sample according to which meson parent produced it.

Events with a π ancestor are weighted as a function of reconstructed energy according to

$$w(E_\nu) = \frac{N_{\text{data}} - (N_{\text{MC}} - N_{\nu_\mu}^K \text{CC} - N_{\nu_\mu}^\pi \text{CC})}{N_{\nu_\mu}^K \text{CC} + N_{\nu_\mu}^\pi \text{CC}}. \quad (12.1)$$

That is, any simulated event that is not a ν_μ CC event with a π or K ancestor is first subtracted from the observed data, and the discrepancy is evenly attributed to those ν_μ CC events from π and K ancestors.

Rather than applying this weight directly to the simulated ν_e CC in reconstructed energy, these weights are applied to the ancestor π kinematic space. The ν_e CC will then be constrained according to the weights in their simulated distribution in ancestral kinematic space. The relevant kinematic variables are p_T and p_z of the ancestor π . This distribution is shown in Figure 12.11 for both ν_μ CC and ν_e CC simulated events.

The ν_μ CC's ancestor's p_T/p_z distribution is weighted as a function of the neutrino's reconstructed energy according to the weights in Eqn. 12.1 determined by the contained ν_μ CC data. The simulation is then weighted in the three-dimensional space, $p_T/p_z/E_\nu$. Weights in p_T/p_z space are determined by projecting this space down to p_T/p_z . This weight, $w(p_T, p_z)$, is then applied to ν_e CC events according to their ancestor's kinematics. This produces an energy dependent correction factor for ν_e CC events which originate from π ancestors.

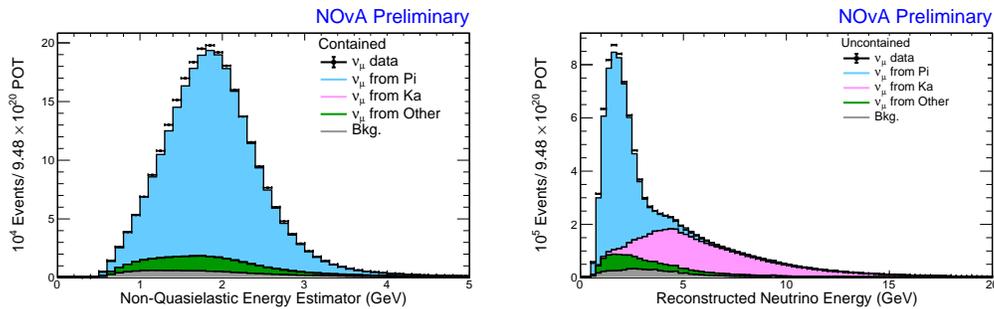


Figure 12.10: The contained, left, and uncontained, right, ν_μ CC selected distributions. For each, the sample has been decomposed into ν_μ ancestor and backgrounds. For contained ν_μ CC events, nearly all events come from an initial π in the target with a negligible K component. The high energy tail of the uncontained sample, however, is dominated by neutrinos from ancestral K .

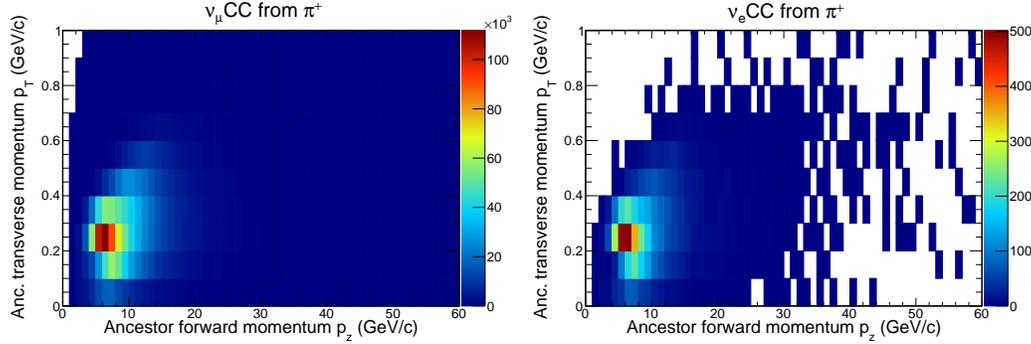


Figure 12.11: The p_T/p_z kinematic distribution for simulated ν_μ CC, left, and ν_e CC, right, events in the near detector.

Since the only notable population of ν_μ CC events with K ancestors is uncontained, the energy estimate of K events is tenuous. Thus, only an averaged scale is determined. First, the determined weights for simulated ν_e CC events with π ancestors are used to calculate an aggregate π weight, S_π

$$S_\pi = \frac{\sum N_{\text{data}} - (N_{\text{MC}} - N_{\nu_\mu \text{ CC}}^K - N_{\nu_\mu \text{ CC}}^\pi)}{\sum N_{\nu_\mu \text{ CC}}^K + N_{\nu_\mu \text{ CC}}^\pi}. \quad (12.2)$$

This average scale is then applied to the simulated events in the uncontained sample with K ancestors. An analogous average S_K can be computed for events with a K ancestor:

$$S_K = \frac{\sum N_{\text{data}} - (N_{\text{MC}} - N_{\nu_\mu \text{ CC}}^K)}{\sum N_{\nu_\mu \text{ CC}}^K}, \quad (12.3)$$

where N_{MC} has been adjusted with S_π described above and the sum only runs over energy bins > 4.5 GeV, where K daughters dominate.

This process is iterated until the derived S_π and S_K are stable under additional applications of the process. In practice, this can take a few iterations as there is quite a large π component in the uncontained sample.

Though the π scale is energy dependent, the derived scale is relatively flat with energy, increasing the component by $\approx 2.5\%$ in neutrino mode. The averaged K scale is 0.937.

12.5 The Michel Decomposition

MichelDecomp uses the distribution of the number of selected Michel electrons in an event to further pick apart the selected ν_e sample in the near detector into its subcomponents. This is particularly suited for separating the ν_μ CC component

from the rest of the sample. Those ν_μ CC events that are selected by the ν_e PID will typically have a muon decay. The muon capture rate is estimated to be 14% averaged over materials in the detector. Michel electrons are relatively infrequent in NC and ν_e CC events and come from $\pi^+ \rightarrow \mu^+ \nu_\mu \rightarrow e^+ \nu_e \bar{\nu}_\mu \nu_\mu$. Thus, the ν_μ CC component in the sample can be statistically separated from the remaining events in the selected sample by examining the distribution of the number of electrons. Michels are a robust handle for separating these components. Michel electrons are time-delayed from the parent cluster so that they are not reconstructed in the primary physics slice. Thus, the presence of a Michel electron does not implicitly correlate with PID's. And, as shown below, Michel electrons have a precisely predicted energy distribution and lifetime so that they are faithfully modeled in Geant4.

Michel Reconstruction

MichelDecomp uses the MEFinder reconstruction package and the MID identifier, described earlier, to reconstruct Michel electrons. This algorithm outputs both TrkME clusters, which reconstruct near the end of some Kalman track, and SlcME's which reconstruct elsewhere in the event. To include a Michel cluster in the decomposition, we require $MID > 0$. This cut was tuned to maximize the ν_μ separating power under systematic distortions in the simulation. Also, the $\Delta t_{\mu e}$ must be > 800 ns. A comparison of simulation and data for the MID inputs is shown in Figure 12.13 while the MID distribution is shown in Figure 12.14. The MID distribution shows some disagreement at low MID score. But, this is a region dominated by background clusters, mostly from neutron capture. And this discrepancy lies com-

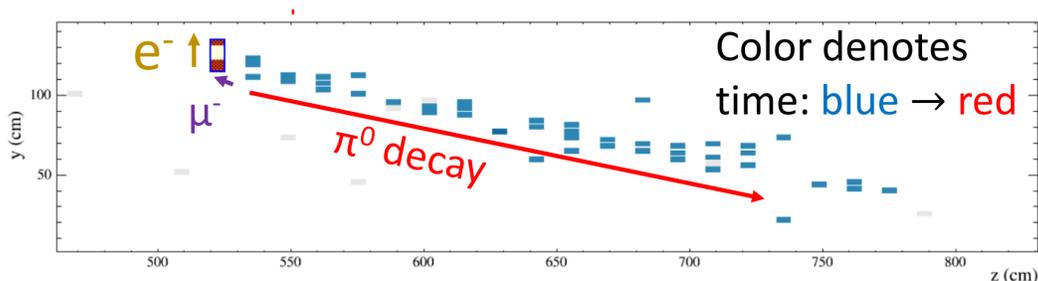


Figure 12.12: A simulated ν_μ CC event selected into the near detector ν_e sample. By virtue of being selected, ν_μ CC events in this sample tend to have short, hard-to-track muons with an energetic π^0 decaying to a large electromagnetic shower. This is the case here, where the muon had too little energy to leave a single hit and a high energy π^0 decayed asymmetrically to give a single electromagnetic shower. There is however, a Michel electron reconstructed in the red hits. The rate of such events can be estimated by the number of Michels in the selected sample.

fortably below the MID cut used in the analysis. In the region dominated by true Michel electrons, the data and MC agree well, as the input variables after a MID cut.

MichelDecomp Algorithm

MichelDecomp uses a template log-likelihood fit to data to separate the ν_μ CC component from the NC and ν_e CC components. Also, this fit is done independently in each bin of PID and reconstructed neutrino energy. This allows the decomposition to separately extrapolate energy-dependent discrepancies between data and MC.

As mentioned earlier, the distribution of number of Michels is nearly degenerate in NC and ν_e CC interactions. Thus, the MichelDecomp algorithm can not separate these two components and an additional constraint must be added to the problem to sufficiently reduce the number of degrees of freedom. We use the BENDecomp algorithm described earlier to break this degeneracy. The ν_e CC component is fixed

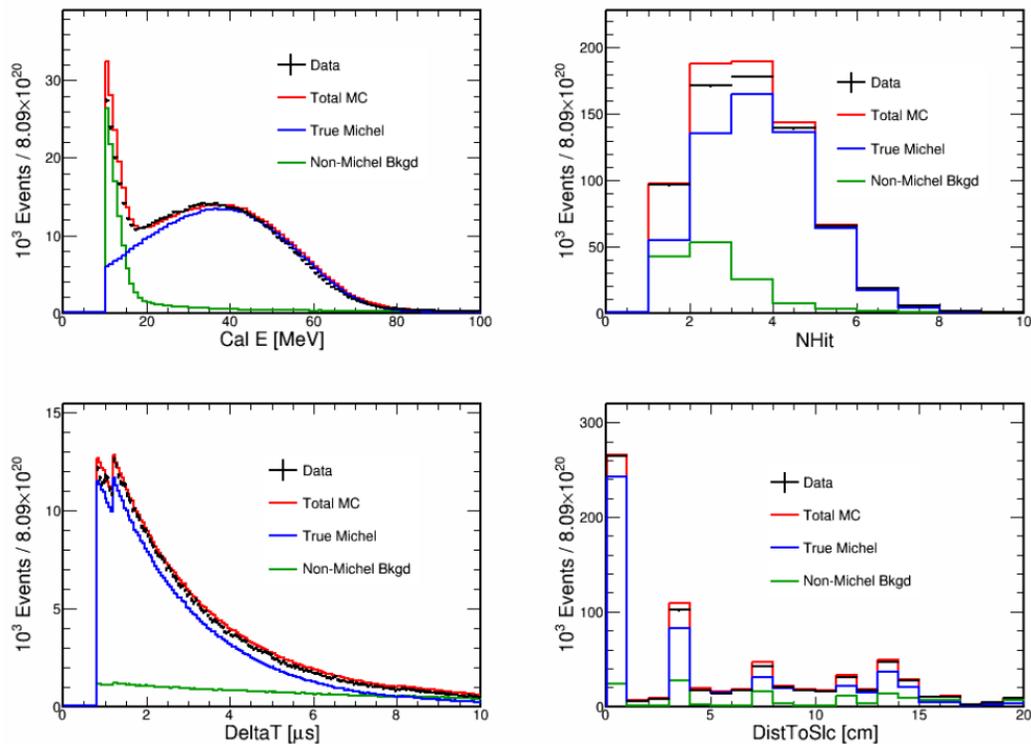


Figure 12.13: A comparison of data and MC for the four input variables used in MID for all Michel clusters reconstructed in the ν_e preselected sample. Most power comes from the calorimetric energy distribution, top left, and $\Delta t_{\mu e}$ distribution, bottom left, which are precisely known from theory. In each plot, the blue and green curves are the true Michel electron and non-Michel background, respectively.

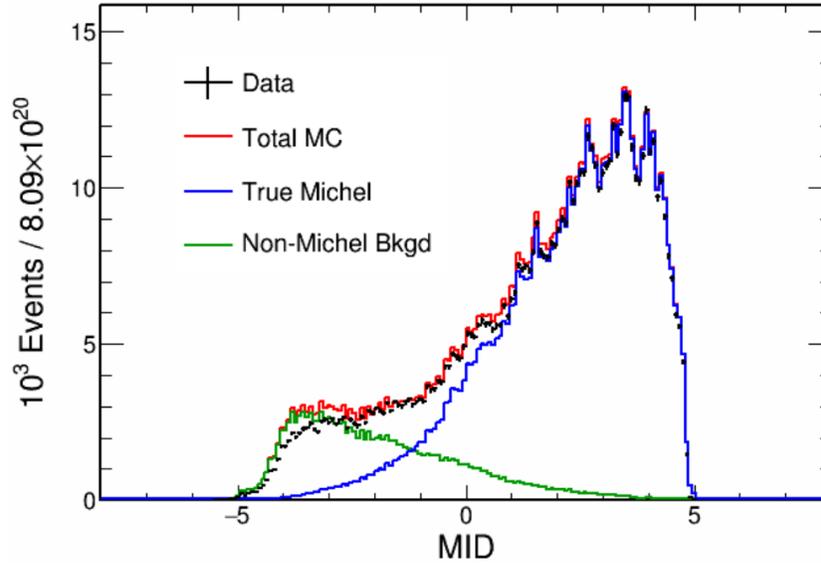


Figure 12.14: A comparison of data and MC for MID for all Michel clusters reconstructed in the ν_e preselected sample. The blue and green curves are the true Michel electron and non-Michel background, respectively.

to the output of BENDecomp in the Michel template fit. This Michel+BENDecomp combined decomposition is referred to as the “combo” decomposition.

Given an estimate for the ν_e CC component, this decomposition algorithm separates ν_μ CC and NC events using a Log-Likelihood fit of data to MC templates. Each analysis bin is decomposed independently of all the others. Let $(cc)_i$, $(nc)_i$, and $(\widehat{ne})_i$ represent the contents of the MC templates for ν_μ CC and NC interactions, and the ν_e CC estimate in an analysis bin, i . The following equation forces the sum of the predictions for the true components to exactly match the data in every analysis bin:

$$h \sum_i (cc)_i + g \sum_i (nc)_i + \sum_i (\widehat{ne})_i = \sum_i D_i, \quad (12.4)$$

where the sums are over bins of N_{ME} . This can be rearranged to reduce the number of degrees of freedom in the LL expression by writing the NC scale factor in terms of the ν_μ CC scale factor.

$$g = \frac{\sum_i (D_i - h(cc)_i - (\widehat{ne})_i)}{\sum_i (nc)_i}. \quad (12.5)$$

If we make this substitution, we’re left with a one parameter log-likelihood for a Poisson experiment,

$$LL = \sum_i \left(D_i - D_i \ln \frac{D_i}{\langle MC_i \rangle} - \langle MC_i \rangle \right). \quad (12.6)$$

Here, D_i is the data observed in each bin of N_{ME} and $\langle MC_i \rangle$ gives the expected number of events in the MC in each N_{ME} bin as a function of h , the ν_μ CC scale factor. The algorithm solves for h using a binary search.

Restricting Decomposition to Bins with Robust Fit

If the fraction of events that are ν_μ CC in a particular analysis bin is small, the MichelDecomp algorithm becomes ineffective. Even in the infinite-statistics cases, small systematic distortions in the large components in the bin can have a larger effect on the distribution of the observed number of Michels than a change in ν_μ CC rate. The analysis quantified in which bins it makes sense to use combo. In bins where combo is not used, the analysis falls back on using the BENDecomp estimate for the ν_e CC component while proportionally scaling the ν_μ CC and NC components.

The bin-by-bin background uncertainty in the far detector was studied using random shifts of systematic uncertainties. The uncertainties for both the combo and the proportional decompositions are shown in Figure 12.15. In simulation, the fraction of events in a given bin that are ν_μ CC is a strong indicator of which decomposition procedure leads to the more precise background prediction. In particular, in all but one case, in the bins in which combo performs better, the ν_μ CC fraction is greater than 20%. In this plot, a low-PID sideband outside the analysis sample is included with $CVNe > 0.75$. This sample was included in the study since the ν_e CC component was under-represented, so that almost all of the decomposition procedure relied on MichelDecomp.

With the 20% ν_μ CC fraction restriction, only five bins out of twelve are decomposed with combo. These are all in the low-PID bin with $0.5 < E_\nu < 3$ GeV. In other bins, the three-component paradigm breaks down, and the near detector sample is dominated by background ν_e CC events so that such a sophisticated decomposition procedure becomes less important.

Decomposition Result

The nominal near detector simulation, compared to the observed data, for the neutrino sample is shown in Figure 12.16. In both PID bins, there is an excess in data near the center of the beam peak along with a deficit in the high and low tails. The

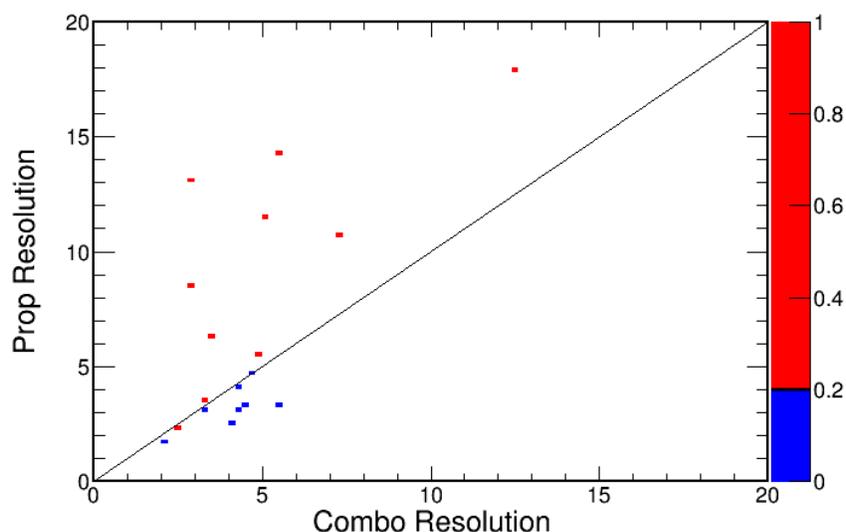


Figure 12.15: A comparison of the bin-by-bin performance of the proportional and combo decomp. The x -axis gives the RMS width of a gas of systematic shifts when using the combo decomposition. The y -axis shows the same for the proportional decomposition. The color indicates whether the fraction of near detector events that are ν_μ CC in each bin is greater or less than 20%. Only bins with $1 < E_\nu^{reco} < 3$ GeV with relatively high statistics are shown. From the plot, bins with a larger fraction of ν_μ CC clearly have more of a lever arm to accurately separate the ν_μ CC component from the others. In the only bin for which the combo decomp performs more accurately, the ν_μ CC fraction is at 21%, and the combo / proportional decompositions perform similarly.

decomposed result is shown in Figure 12.17. After the decomposition, the total constrained MC matches the observed data by construction. In the five energy bins of the low-PID region, there is an energy dependent trend for the ν_μ CC and NC scales. At lower energy, the sample looks more ν_μ CC like under the MichelDecomp templates while at higher energies, the NC component is enhanced while the ν_e CC component is reduced. Summed over the analysis bins, the combo decomp predicts the sample proportions to be very similar to the simulated prediction, as shown in Table 12.1, with an enhancement of ν_μ CC events and reduction of NC events. In each bin, there are two scales determined from three bins, giving a single degree of freedom per bin. Each of these five bins decomposed with combo decomp are fit completely independently, and so a total χ^2/N_{dof} of 5.5/5 is achieved. There is an outlier in χ^2 at 4.84. But, from sampling the χ^2 distribution with $\nu = 1$ five independent times, there is an 18% of measuring a χ^2 of at least 4.84 in one of the trials.

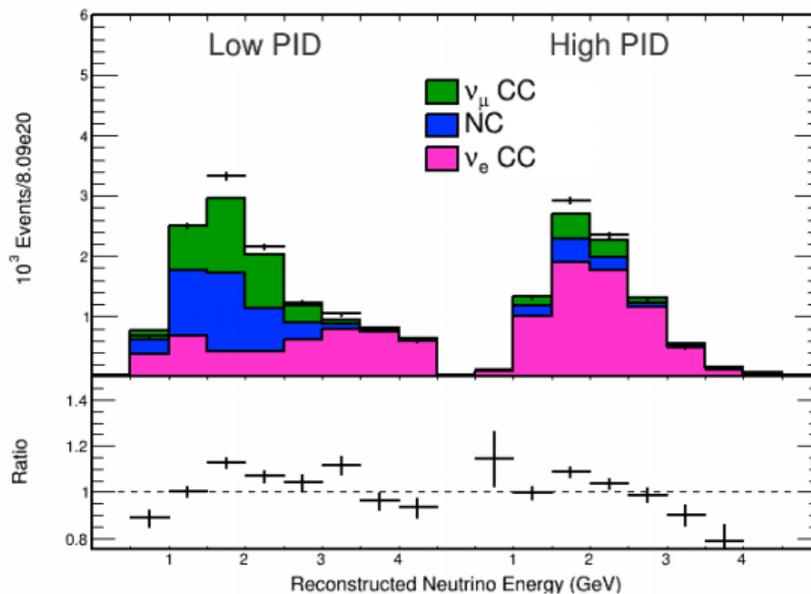


Figure 12.16: The raw, un-decomposed near detector sample in neutrino running with data overlaid. The bottom pane shows the total data to MC ratio.

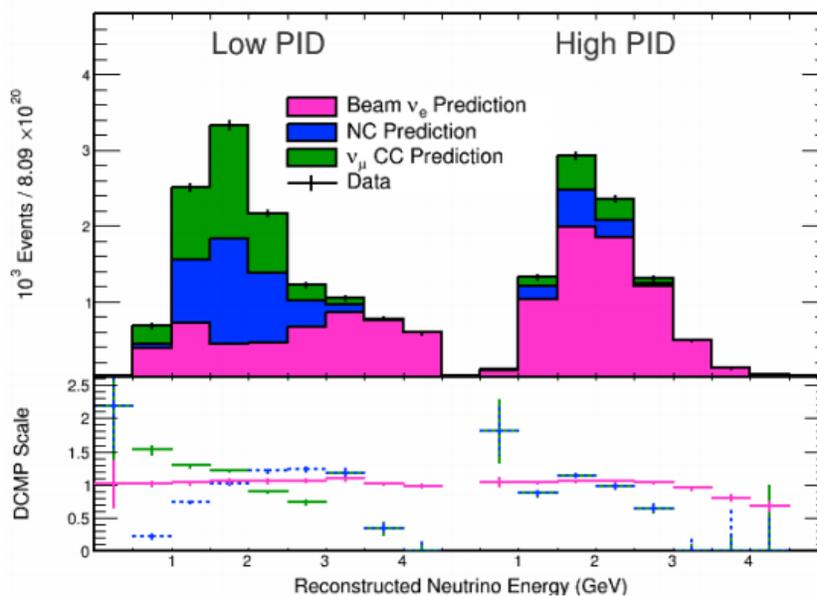


Figure 12.17: The decomposed near detector sample in neutrino running with data overlaid. The bottom pane shows the decomposition scales applied to each interaction component in each analysis bin.

| Sample | ν_μ CC | NC | ν_e CC | Total |
|----------------|--------------|------------|-------------|-------|
| Data | - | - | - | 21755 |
| Uncorrected MC | 4542 (21%) | 5032 (24%) | 11678 (55%) | 21252 |
| PropDecomp | 4732 (22%) | 5216 (24%) | 11792 (54%) | 21740 |
| ComboDecomp | 4871 (22%) | 4829 (22%) | 12055 (55%) | 21755 |

Table 12.1: The number of ν_μ CC, NC, and ν_e CC events predicted in the near detector neutrino dataset. The uncorrected simulation is 2% lower than the data when integrated over the entire sample. The decompositions do not significantly change the proportions of each component from the simulated predictions.

12.6 Data/MC Comparison for the Decomposed sample

The calorimetric energy on electromagnetic and non-electromagnetic reconstructed prongs is shown in Figures 12.18 and 12.19.

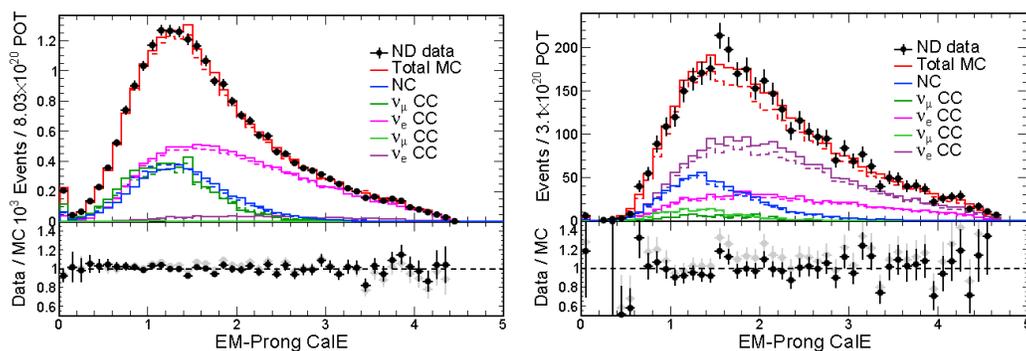


Figure 12.18: A comparison of the calorimetric energy on electromagnetic prongs in near detector data and MC. The left(right) plot shows the comparison in (anti-)neutrino data. The dashed curves give the simulated distributions prior to decomposition while the solid curves have had decomposition weights applied in analysis bins.

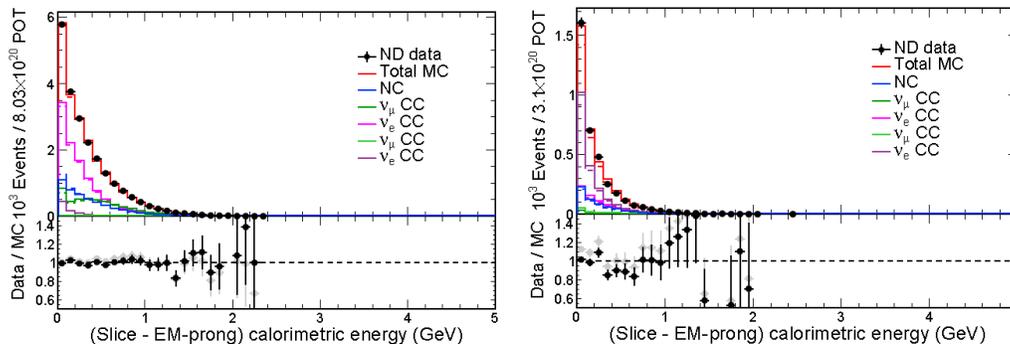


Figure 12.19: A comparison of the calorimetric energy on hadronic, non-electromagnetic prongs in near detector data and MC. The left(right) plot shows the comparison in (anti-)neutrino data. The dashed curves give the simulated distributions prior to decomposition while the solid curves have had decomposition weights applied in analysis bins.

12.7 Near-to-Far Extrapolation

As a two detector experiment, NOvA has an opportunity to use near detector data to correct the far detector simulation event rates. As already shown, this near detector dataset is quite large and gives a sample for detailed studies of systematics and modeling. The “extrapolation” is a loose term to describe incorporating the observed near detector data into the analysis to constrain the simulated predictions at the far detector. There are two closely related constructs here: a prediction and a flavor transition. A flavor transition corresponds to a physical oscillation channel: $\nu_\alpha \rightarrow \nu_\beta$. A prediction refers to the oscillated spectrum our simulation estimates at the far detector. Each oscillation channel we measure, $\nu_\mu \rightarrow \nu_\mu$, $\nu_\mu \rightarrow \nu_e$, $\bar{\nu}_\mu \rightarrow \bar{\nu}_\mu$, and $\bar{\nu}_\mu \rightarrow \bar{\nu}_e$, will be predicted by summing all allowed flavor transitions, while folding in the constraint from near detector data. In the $\nu_\mu \rightarrow \nu_e$ case, for example, the $\nu_\mu \rightarrow \nu_e$ flavor transition represents the analysis signal, while other transitions such as $\nu_\mu \rightarrow \nu_\mu$ and $\bar{\nu}_\mu \rightarrow \bar{\nu}_e$ constitute the background. The 14 possible transition channels are

- $\nu_e \rightarrow \nu_e$ CC
- $\nu_e \rightarrow \nu_\mu$ CC
- $\nu_e \rightarrow \nu_\tau$ CC
- $\nu_\mu \rightarrow \nu_e$ CC
- $\nu_\mu \rightarrow \nu_\mu$ CC
- $\nu_\mu \rightarrow \nu_\tau$ CC
- $\nu_\alpha \rightarrow \nu_\beta$ NC
- $\bar{\nu}_e \rightarrow \bar{\nu}_e$ CC
- $\bar{\nu}_e \rightarrow \bar{\nu}_\mu$ CC
- $\bar{\nu}_e \rightarrow \bar{\nu}_\tau$ CC

- $\bar{\nu}_\mu \rightarrow \bar{\nu}_e$ CC
- $\bar{\nu}_\mu \rightarrow \bar{\nu}_\tau$ CC
- $\bar{\nu}_\mu \rightarrow \bar{\nu}_\mu$ CC
- $\bar{\nu}_\alpha \rightarrow \bar{\nu}_\beta$ NC

The list is the same for both neutrino and anti-neutrino running, due to intrinsic wrong-sign backgrounds in the beam. The combinatorics are decreased since ν_x NC events are indistinguishable in our detector. With the four oscillation channels and the 14 flavor transitions allowed per prediction, we must account for 56 total flavor transitions over all oscillation channels. How this is accomplished for each of the oscillation channels measured is detailed in the following.

Strategies for Constraining Individual Flavor Transitions

In general, NOvA uses three methods for constraining various flavor transitions during building each of the oscillated predictions. These are referred to as “reco-true-reco” extrapolation, “reco-reco” extrapolation, and no extrapolation[128].

The “reco-true-reco” case is best explained schematically, Figure 12.20. The diagram shows the $\nu_\mu \rightarrow \nu_\mu$ flavor transition channel for the $\nu_\mu \rightarrow \nu_\mu$ disappearance measurement. It starts with plotting the near detector ν_μ -selected spectrum in reconstructed energy, on the top left. This is then translated to true energy via the simulated migration matrix, shown in the bottom left pane. During this process, the observed data to MC discrepancy in each energy bin is propagated to the histogram in true energy, \hat{n}_i^{true} :

$$\hat{n}_i^{\text{true}} = \sum_j \frac{n_{j,DCMP}^{\text{reco}}}{n_{j,MC}^{\text{reco}}} M_{ji}^{\text{ND}}, \quad (12.7)$$

where M_{ji}^{ND} gives the simulated migration matrix in the near detector. Here, $n_{j,DCMP}^{\text{reco}}$ gives the number of ν_μ CC events estimated in the near detector by the decomposition. As written, \hat{n}_i^{true} is a constrained prediction for what the true energy distribution in our near detector looks like, shown in the bottom, second pane from the left.

At this point, the far-to-near ratio is applied as predicted in the simulation. This encompasses efficiency and acceptance differences between the two detectors, differences in detector mass, and differences in the flux at the two detector locations. A back-of-the-envelope calculation gives this should on average be the differences in fiducial volumes \times the differences in fluxes:

$$\frac{M_{\text{FD}}}{M_{\text{ND}}} \times \frac{\Phi_{\text{FD}}}{\Phi_{\text{ND}}} = \frac{10^7 \text{kg}}{3 \times 10^4 \text{kg}} \times \frac{\frac{\Phi_0}{(810 \text{km})^2}}{\frac{\Phi_0}{(1 \text{km})^2}} = 0.0005. \quad (12.8)$$

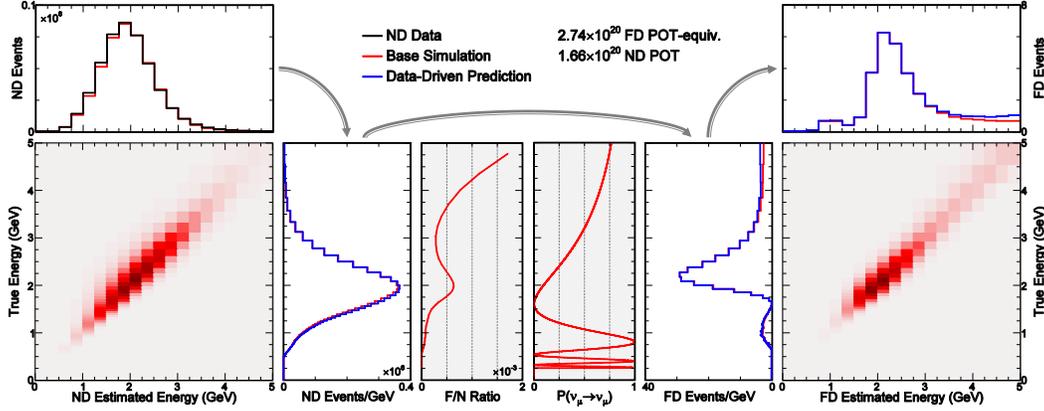


Figure 12.20: The NOvA “reco-true-reco” strategy for extrapolating discrepancies between near detector simulation and data to the far detector. This plot is explained in detail in the text.

For comparison, the far-to-near is about 0.0006 at 2 GeV as shown in the bottom, third-to-left panel. This results in an un-oscillated prediction of the far detector histogram, \hat{f}_i^{true} :

$$\hat{f}_i^{\text{true}} = \left[\sum_j \frac{n_{j,DCMP}^{\text{reco}}}{n_{j,MC}^{\text{reco}}} M_{ji}^{\text{ND}} \right] \frac{f_i^{\text{true}}}{n_i^{\text{true}}}. \quad (12.9)$$

Multiplying by an oscillation probability for the $\nu_\mu \rightarrow \nu_\mu$ flavor transition, shown in the bottom, fourth-from-the-left, the oscillated, constrained far detector event count in true energy, \hat{e}_i^{true} , is

$$\hat{e}_i^{\text{true}} = \left[\sum_j \frac{n_{j,DCMP}^{\text{reco}}}{n_{j,MC}^{\text{reco}}} M_{ji}^{\text{ND}} \right] \frac{f_i^{\text{true}}}{n_i^{\text{true}}} P_{\mu \rightarrow \mu}(E_\nu), \quad (12.10)$$

where E_ν is determined by the center of bin i . This is the distribution in the bottom, fifth-from-the-left.

After this, the last remaining step is to fold the true energy distribution back into reconstructed energy bins using the migration matrix M_{ik}^{FD} , shown in the bottom right. This gives us our expectation in each analysis bin, \hat{e}_k :

$$\hat{e}_k = \sum_i \left[\sum_j \frac{n_{j,DCMP}^{\text{reco}}}{n_{j,MC}^{\text{reco}}} M_{ji}^{\text{ND}} \right] \frac{f_i^{\text{true}}}{n_i^{\text{true}}} P_{\mu \rightarrow \mu}(E_\nu) M_{ik}^{\text{FD}}. \quad (12.11)$$

This final prediction is shown on the top right. When performing a fit, the predicted spectra at several test oscillation parameters must be generated. To be compatible with this, for each flavor transition constraint, it is possible to predict the component for an arbitrary choice of oscillation parameters.

The other two constraint schemes are a subset of the above description. In the “reco” case, the migration matrices are never used to pull the samples to an estimated true energy. Instead, the prediction formula for each component simplifies to:

$$\hat{e}_i = n_{i,DCMP}^{\text{reco}} \frac{f_i^{\text{reco}}}{n_i^{\text{reco}}} P_{\alpha \rightarrow \beta}(E_\nu). \quad (12.12)$$

The most notable way this differs from above is that the oscillation probability, which is of course defined for a specific true energy, is here averaged over the range of true energies that contributed to the particular bin of reconstructed energy.

The remaining scheme, “no extrapolation” is by far the simplest. It is estimated as

$$\hat{e}_i = f_i^{\text{reco}} P_{\alpha \rightarrow \beta}(E_\nu). \quad (12.13)$$

That is, the estimate in each bin is just based on the far detector simulation with the proper oscillation probability applied. This scheme is only applied to negligible flavor transitions, such as $\nu_e \rightarrow \nu_\mu$, which is always suppressed in flux, oscillation probability, and efficiency. Also, ν_τ appearance is always predicted with this scheme as $m_\tau > m_\pi, m_K$ so that there is no ν_τ contamination in the beam to use for constraining our reconstruction efficiency and the contribution to the selected far detector samples is negligible.

12.8 The Prediction for the Disappearance Spectrum

For the disappearance analysis, the signal $\nu_\mu \rightarrow \nu_\mu$ and $\bar{\nu}_\mu \rightarrow \bar{\nu}_\mu$ flavor transitions are the only significant contributions to the prediction. The simulated beam background is predicted to be only 1.3% of the selected far detector sample[129]. This beam background comes from NC events, which represent less than 0.34% of the simulated near detector sample, and ν_τ appearance events, which are absent from the near detector sample. Thus, there is very little hope for directly constraining backgrounds. However, their negligible representation in the near sample imply that even poor modeling of their normalization should have a small effect on the analysis. We do, however, have a data-driven NC constraint using a separately selected sample whose derived NC normalization error band is propagated through the analysis as a systematic[100]. On a similar note, as NOvA detectors are not magnetized, there is not a robust method to separate ν_μ and $\bar{\nu}_\mu$ events. So, we assume the simulated “wrong-sign” background ratio of $\bar{\nu}_\mu$ in the ν_μ disappearance spectrum is fixed to the simulated prediction, and similarly for the $\bar{\nu}_\mu$ disappearance. Thus, observed discrepancies with simulation in the near detector adjust both the $\nu_\mu \rightarrow \nu_\mu$ CC and $\bar{\nu}_\mu \rightarrow \bar{\nu}_\mu$ transitions in the same way.

Thus, the two relevant flavor transitions, $\nu_\mu \rightarrow \nu_\mu$ and $\bar{\nu}_\mu \rightarrow \bar{\nu}_\mu$, are predicted with the reco-true-reco extrapolation scheme while all other flavor transitions are predicted with no extrapolation, relying solely on the far detector simulation. This is true for both the neutrino and anti-neutrino predictions. As described earlier, the disappearance measurement bins the far detector spectrum in HadEFrac to effectively improve the energy resolution. Each HadEFrac bin in the far detector is predicted using the ν_μ -selected spectrum in near detector with the same HadEFrac range. Since the HadEFrac bins were determined to separate the far detector sample into quartiles, the total prediction in each bin is very nearly 25% of the total sample in each bin. The extrapolation procedure can change this number by $\pm 1\%$. Numbers for each flavor transition's predicted contribution to the sample are shown in Table 12.2 and 12.3 for the neutrino and anti-neutrino predictions separately. The neutrinos have been oscillated at the best fit from the 2017 analysis: $\sin^2 \theta_{23} = 0.56$, $\delta_{CP} = 1.21\pi$, and $\Delta m_{32}^2 = 0.0024 \text{ eV}^2$. There are 125(61) events predicted in the (anti-)neutrino dataset on a background of 3.2(1.3).

The extrapolation in HadEFrac bins is a helpful tool for reducing systematics. Systematic uncertainties will be described in more detail in the following chapter; the importance of the HadEFrac extrapolation is expressed here with a short study. The HadEFrac binning is a proxy for grouping neutrino interactions with similar production mechanisms. For example, the first HadEFrac quartile is dominated by QE events in both the near and far detector samples, while the highest HadEFrac is dominated by DIS events. Thus, if the simulation under-predicts the QE representation in the sample, the extrapolation procedure will see a large excess in the first quartile data which will in turn increase the far detector simulated prediction in the

| Flavor Transition | Predicted Events | Flavor Transition | Predicted Events |
|--------------------------------|------------------|--|------------------|
| $\nu_e \rightarrow \nu_e$ | 0.022 | $\bar{\nu}_e \rightarrow \bar{\nu}_e$ | 0.000 |
| $\nu_e \rightarrow \nu_\mu$ | 0.261 | $\bar{\nu}_e \rightarrow \bar{\nu}_\mu$ | 0.013 |
| $\nu_e \rightarrow \nu_\tau$ | 0.001 | $\bar{\nu}_e \rightarrow \bar{\nu}_\tau$ | 0.000 |
| $\nu_\mu \rightarrow \nu_e$ | 0.091 | $\bar{\nu}_\mu \rightarrow \bar{\nu}_e$ | 0.001 |
| $\nu_\mu \rightarrow \nu_\mu$ | 114.064 | $\bar{\nu}_\mu \rightarrow \bar{\nu}_\mu$ | 7.449 |
| $\nu_\mu \rightarrow \nu_\tau$ | 0.343 | $\bar{\nu}_\mu \rightarrow \bar{\nu}_\tau$ | 0.046 |
| $\nu + \bar{\nu}$ NC | 1.195 | Cosmic Evt | 1.42 |
| Total Events: 125.266 | | Total Background: 3.199 | |

Table 12.2: The prediction for each flavor transition's contribution to the far detector ν_μ sample. The prediction is made at the 2017 analysis best fit: $\sin^2 \theta_{23} = 0.56$, $\delta_{CP} = 1.21\pi$, and $\Delta m_{32}^2 = 0.0024 \text{ eV}^2$.

| Flavor Transition | Predicted Events | Flavor Transition | Predicted Events |
|--------------------------------|------------------|--|------------------|
| $\nu_e \rightarrow \nu_e$ | 0.005 | $\bar{\nu}_e \rightarrow \bar{\nu}_e$ | 0.003 |
| $\nu_e \rightarrow \nu_\mu$ | 0.036 | $\bar{\nu}_e \rightarrow \bar{\nu}_\mu$ | 0.070 |
| $\nu_e \rightarrow \nu_\tau$ | 0.000 | $\bar{\nu}_e \rightarrow \bar{\nu}_\tau$ | 0.000 |
| $\nu_\mu \rightarrow \nu_e$ | 0.004 | $\bar{\nu}_\mu \rightarrow \bar{\nu}_e$ | 0.001 |
| $\nu_\mu \rightarrow \nu_\mu$ | 15.002 | $\bar{\nu}_\mu \rightarrow \bar{\nu}_\mu$ | 45.025 |
| $\nu_\mu \rightarrow \nu_\tau$ | 0.125 | $\bar{\nu}_\mu \rightarrow \bar{\nu}_\tau$ | 0.120 |
| $\nu + \bar{\nu}$ NC | 0.459 | Cosmic Evt | 0.59 |
| Total Events: 61.447 | | Total Background: 1.312 | |

Table 12.3: The prediction for each flavor transition’s contribution to the far detector $\bar{\nu}_\mu$ sample. The prediction is made at the 2017 analysis best fit: $\sin^2 \theta_{23} = 0.56$, $\delta_{CP} = 1.21\pi$, and $\Delta m_{32}^2 = 0.0024 \text{ eV}^2$.

first quartile to correct for this under-prediction. The remaining quartiles will be less significantly impacted. However, if all HadEFrac bins were grouped together in the extrapolation, the QE deficit would be attributed to the full range of simulated events and a smaller correction factor would apply to all HadEFrac bins. In this case, the first quartile would be under-predicted while all others would be over-predicted. A concrete example of this phenomenon is shown in Figure 12.21 where a systematic, mostly affecting the first quartile, is properly extrapolated using the near detector first quartile spectrum. But, the extrapolation procedure under-predicts the signal if the entire near detector sample is used to predict the given quartile.

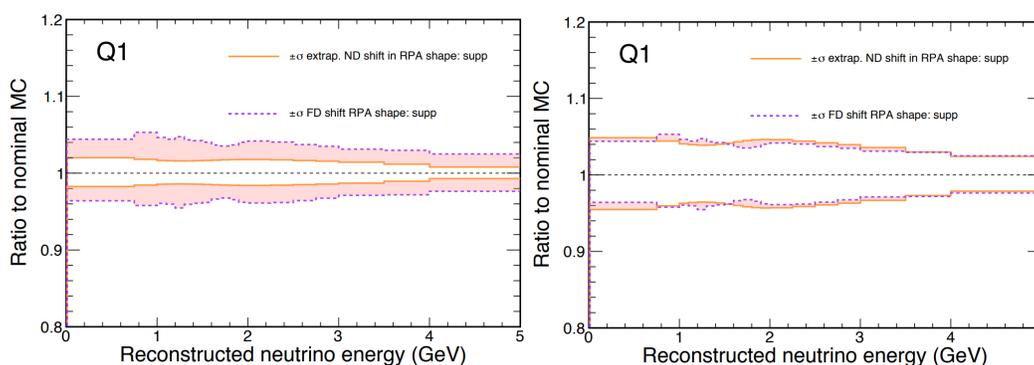


Figure 12.21: A neutrino interaction systematic knob is adjusted, yielding an excess in the prediction for quartile 1 in the far detector. This systematic is targeted to influence events in quartile 1. So, if the entire near detector sample is used to predict this quartile, the excess from this systematic gets spread over all quartiles leading to a characteristic under-prediction in quartile 1, as shown on the left. However, if events with the same HadEFrac range are used to extrapolate, the prediction almost perfectly accounts for the systematic variation, illustrated on the right.

12.9 The Prediction for the Appearance Spectrum

As in the ν_μ disappearance case, the flavor transitions that generate the analysis signal, $\nu_\mu \rightarrow \nu_e$ and $\bar{\nu}_\mu \rightarrow \bar{\nu}_e$, are estimated with a reco-true-reco extrapolation. Additionally, the same near detector ν_μ -selected spectrum is used to predict these transitions as we're measuring ν_e appearance in a ν_μ beam. Unlike in the disappearance prediction, the appearance prediction does not split the near detector into HadEFrac bins.

There are also very small components of the sample that are not extrapolated, falling back on the far detector simulation for each flavor transition. These are:

- $\nu_e \rightarrow \nu_\mu$ CC
- $\nu_e \rightarrow \nu_\tau$ CC
- $\nu_\mu \rightarrow \nu_\tau$ CC
- $\bar{\nu}_e \rightarrow \bar{\nu}_\mu$ CC
- $\bar{\nu}_e \rightarrow \bar{\nu}_\tau$ CC
- $\bar{\nu}_\mu \rightarrow \bar{\nu}_\tau$ CC

In contrast with the disappearance prediction, the list of non-extrapolated backgrounds does not exhaust all possible backgrounds. These background flavor transitions are constrained with a reco-reco extrapolation:

- $\nu_e \rightarrow \nu_e$ CC
- $\nu_\mu \rightarrow \nu_\mu$ CC
- $\nu_\alpha \rightarrow \nu_\beta$ NC
- $\bar{\nu}_e \rightarrow \bar{\nu}_e$ CC
- $\bar{\nu}_\mu \rightarrow \bar{\nu}_\mu$ CC
- $\bar{\nu}_\alpha \rightarrow \bar{\nu}_\beta$ NC

These are exactly the processes that the decomposition constrains in the near detector. The decomposition picks the near detector ν_e -selected sample into ν_μ CC, NC, and ν_e CC estimates. Since these backgrounds are estimated in the near detector, they

are extrapolated to constrain the far detector prediction. These events do not use a full reco-true-reco extrapolation primarily due to limitations in the number of MC events thrown in the near detector. There is insufficient statistics to populate the reco-true migration matrices since these events are relatively rare compared to ν_μ CC. Additionally, since those ν_μ CC and NC events that are selected by the ν_e PID are typically inelastic DIS events with significant hadronic energy. Therefore, an extrapolation in true energy strongly depends on the neutrino interaction model in the simulation.

As already explained, the near detector ν_e -selected sample is decomposed with the combo decomposition, including BENDecomp and MichelDecomp, while the $\bar{\nu}_e$ sample is decomposed proportionally. As in the signal prediction, the $\nu/\bar{\nu}$ split is always tied to the simulation prediction. This extrapolation procedure predicts 49.89 signal ν_e CC events and a background of 15.15 in neutrino mode and 7.04 $\bar{\nu}_e$ CC events on a background of 5.19 in anti-neutrino mode. The neutrinos have been oscillated at the best fit from the 2017 analysis: $\sin^2 \theta_{23} = 0.56$, $\delta_{CP} = 1.21\pi$, and $\Delta m_{32}^2 = 0.0024 \text{ eV}^2$. This is further broken down into the contribution from each flavor transition in Tables 12.4 and 12.5 for neutrino and anti-neutrino running. The predicted distribution is shown in the proper analysis bins in Figure 12.22. This plot shows the prediction that will be given to the fitter, for a certain choice of oscillation parameters, after making the far detector measurement.

The peripheral sample is not binned in energy for this analysis, though it is extrapolated in energy. The peripheral sample CVN cut is forced to be the CVN cut on the high-PID sample. This allows us to use the same high-PID sample in the near detector to predict both the high-CVN events in the core and the peripheral sample. The same decomposition ratio features in the peripheral extrapolation as for the core

| Flavor Transition | Predicted Events | Flavor Transition | Predicted Events |
|--------------------------------|------------------|--|------------------|
| $\nu_e \rightarrow \nu_e$ | 6.448 | $\bar{\nu}_e \rightarrow \bar{\nu}_e$ | 0.376 |
| $\nu_e \rightarrow \nu_\mu$ | 0.001 | $\bar{\nu}_e \rightarrow \bar{\nu}_\mu$ | 0.000 |
| $\nu_e \rightarrow \nu_\tau$ | 0.001 | $\bar{\nu}_e \rightarrow \bar{\nu}_\tau$ | 0.000 |
| $\nu_\mu \rightarrow \nu_e$ | 49.886 | $\bar{\nu}_\mu \rightarrow \bar{\nu}_e$ | 0.438 |
| $\nu_\mu \rightarrow \nu_\mu$ | 0.394 | $\bar{\nu}_\mu \rightarrow \bar{\nu}_\mu$ | 0.007 |
| $\nu_\mu \rightarrow \nu_\tau$ | 0.347 | $\bar{\nu}_\mu \rightarrow \bar{\nu}_\tau$ | 0.048 |
| $\nu + \bar{\nu}$ NC | 3.019 | Cosmic Evt | 4.07 |
| Total Events: 65.03 | | Total Background: 15.15 | |

Table 12.4: The prediction for each flavor transition's contribution to the far detector ν_e sample, at the 9.48×10^{20} POT analysis exposure.

| Flavor Transition | Predicted Events | Flavor Transition | Predicted Events |
|--------------------------------|------------------|--|------------------|
| $\nu_e \rightarrow \nu_e$ | 0.596 | $\bar{\nu}_e \rightarrow \bar{\nu}_e$ | 1.909 |
| $\nu_e \rightarrow \nu_\mu$ | 0.00 | $\bar{\nu}_e \rightarrow \bar{\nu}_\mu$ | 0.000 |
| $\nu_e \rightarrow \nu_\tau$ | 0.001 | $\bar{\nu}_e \rightarrow \bar{\nu}_\tau$ | 0.000 |
| $\nu_\mu \rightarrow \nu_e$ | 1.299 | $\bar{\nu}_\mu \rightarrow \bar{\nu}_e$ | 7.040 |
| $\nu_\mu \rightarrow \nu_\mu$ | 0.023 | $\bar{\nu}_\mu \rightarrow \bar{\nu}_\mu$ | 0.022 |
| $\nu_\mu \rightarrow \nu_\tau$ | 0.072 | $\bar{\nu}_\mu \rightarrow \bar{\nu}_\tau$ | 0.095 |
| $\nu + \bar{\nu}$ NC | 0.608 | Cosmic Evt | 0.57 |
| Total Events: 12.34 | | Total Background: 5.19 | |

Table 12.5: The prediction for each flavor transition’s contribution to the far detector $\bar{\nu}_e$ sample, at the 6.90×10^{20} POT analysis exposure.

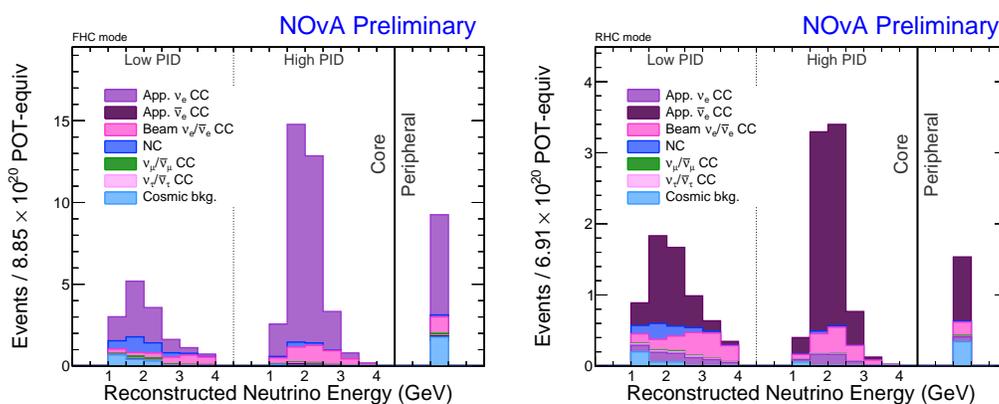


Figure 12.22: The extrapolated far detector predictions for the ν_e appearance measurement, left, and the $\bar{\nu}_e$ appearance measurement, right. In each plot, the left two peaks are the energy distribution in the low and high-PID bins. The right-most entry is the peripheral sample, whose energy distribution is integrated into a single event-count bin.

sample.

As a final note, the wrong-sign background flavor transitions, $\bar{\nu}_\mu \rightarrow \bar{\nu}_e$ in the ν_e appearance analysis and $\nu_\mu \rightarrow \nu_e$ in the $\bar{\nu}_e$ appearance analysis, are particularly harmful for this analysis. The appearance probability anti-correlates for $\nu/\bar{\nu}$ for several physical parameters. Notably, the CP-violating angle, δ_{CP} , will suppress the rate of $\bar{\nu}_e$ appearance if the ν_e appearance rate is enhanced. This makes oscillation parameters less clearly resolved as these two effects cancel to some degree. From Table 12.5, the wrong-sign background is 18% of the size of the signal $\bar{\nu}_e$, averaged over the sample. This gives a large effective reduction in the number of signal events seen in the analysis.

As such, knowing the wrong sign background in the $\bar{\nu}_e$ sample is important for a

precise prediction. We've developed three independent estimates for estimating the near detector wrong sign component in the ν_e - and ν_μ -selected spectra[130]. Each of these agrees, to within the simulated error band, for the wrong sign fraction in simulation. For this analysis, we do not amend the wrong sign prediction, but rather use it as a validation of our simulated prediction of our anti-neutrino flux.

A diagrammatic overview of the selected samples is shown in Figure 12.23. In these diagrams, the sample is segregated into selection bins in the core and peripheral samples and broken into signal, beam background, and cosmic background. In these plots, each unit of area corresponds to the same number of events, so the horizontal length of a box gives the fraction of selected events that reconstructs in a given selection bin. And, the vertical length of the red gives the signal purity. The predictions were made at the best fit from our 2017 dataset: $\sin^2 \theta_{23} = 0.556$, $\Delta m_{32}^2, \delta_{CP} = 1.21\pi$. In neutrino mode, over half of the selected events predicted are signal events in the high-PID bin of the core. The anti-neutrino plot is much more background-dominated as the 2017 fit suggests a preference for ν_e appearance over $\bar{\nu}_e$.

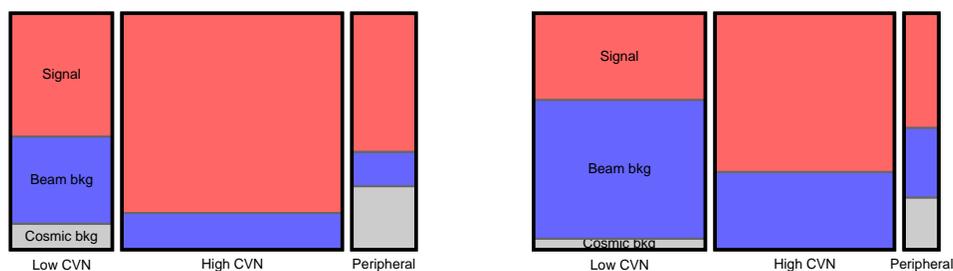


Figure 12.23: The far detector prediction, divided into CVN bins in the core and peripheral and separated into signal, beam background, and cosmic background. The neutrino prediction is on the left with the anti-neutrino on the right. Each unit of area corresponds to the same number of events. So, the vertical height of each bar gives the proportion of each sample corresponding to signal, etc., while the horizontal length corresponds to the fraction of the sample that is estimated in each subsample.

Chapter 13

SYSTEMATIC UNCERTAINTIES

An appropriate enumeration of systematic uncertainties is important for any analysis. Even though our sample is limited to a few dozen events over several bins for the appearance measurement, systematic uncertainties degrade our sensitivities to physics parameters notably. The treatment here has significant overlap with the $CC\pi^0$ analysis, described in Chapter 9. The sources of error are again divided into five different classes: flux errors, neutrino interaction errors, errors from near/far acceptance differences, detector response errors, and normalization errors. The sources of systematic error, along with their relative influence on our predictions, are detailed below. The prediction, as already mentioned, is capable of predicting oscillated far detector spectra at an arbitrary choice of oscillation parameters. Additionally, predictions know about the shape and normalization distortions brought on by each of the systematic effects detailed below. In the fit, each systematic is treated with a pull term, adding a single nuisance parameter. The influence these systematics have on our physics sensitivity are shown in Chapter 14.

We'll first detail each of the various sources of systematic error along with the methods used for calculating how it influences our analysis. At the end of the chapter, we will compare the total effect of all our systematics to show which has the biggest impact on our measurement.

13.1 Flux Systematic Uncertainty

The analysis relies heavily on MINERvA's PPFX[68] framework. The PPFX analysis featured an extensive survey of hadron production data applied to the NuMI target to significantly reduce the error stemming from hadron production. PPFX readjusts the NuMI central value for the flux prediction, along with a reduced error band estimate.

More than an error band, the PPFX analysis released a set of 100 "universes", in which each of MINERvA's systematics were thrown from the appropriate Gaussian distribution. This is clearly more helpful than a quoted error band, as it allows us to create a covariance matrix within our experimental setup for both detectors. In particular, we need the covariance matrix for the far-over-near, since we will always

use our near detector selected events to predict our far detector spectra. Figure 13.1 shows this covariance matrix between all four flux components, ν_μ , $\bar{\nu}_\mu$, ν_e , and $\bar{\nu}_e$, for both neutrino and anti-neutrino running. Apart from the PPFX universes which only cover systematic uncertainties associated with hadron production, NOvA's estimate for focus uncertainties in the beam line are also accounted for in this covariance matrix.

Apart from the hadron production systematics determined with the PPFX error band there are errors on the neutrino flux originating from insufficient modeling of the beam's optical model. By varying the beam's incident position on the target and horn currents, a covariance matrix similar to that shown in Figure 13.1 can be calculated. The optical model also accounts for bending of the secondary meson beam in the decay pipe by the Earth's magnetic field. Optical modeling systematic uncertainties are relatively minor compared to the flux uncertainties from hadron production uncertainties.

In the covariance matrix calculation, the PPFX universes have been smoothed to reduce sudden un-physical discontinuities. A comparison of several raw and

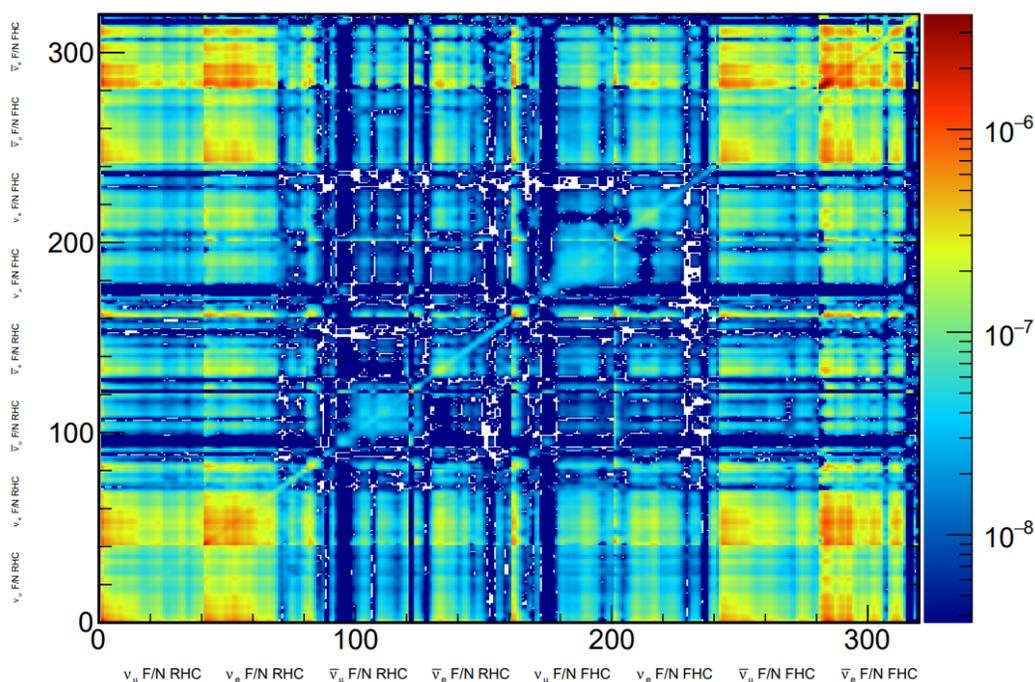


Figure 13.1: The covariance matrix for flux systematic errors relating to hadron production in the target. The first 160 bins represent the flux in anti-neutrino mode while the final 160 represent the neutrino mode flux. Within each polarity are the ν_μ , ν_e , and $\bar{\nu}_\mu$, $\bar{\nu}_e$ far-over-near ratios. Each has 40 bins of true neutrino energy.

smoothed universes is shown in Figure 13.2. Further, it was shown that the error band determined by the smoothed universes was identical to the nominal error band.

To take advantage of the covariance matrix we can construct from the PPFX universes without including all universes individually in our fit, we perform a Principal Component Analysis (PCA)[131] to determine the dominant systematic shapes in the covariance matrix. A PCA relies on finding the eigenvectors of a covariance matrix. These are analogous to principal axes in the n -dimensional error ellipse being studied. The eigenvectors are then sorted by eigenvalue, so that those eigenvectors that contribute most to the overall error are listed first in the list. These vectors with high eigenvalues usually represent low-frequency structure in the covariance matrix – either largely flat or with a low number of periodic oscillations over the total number of bins. Then, only the first m eigenvectors are accepted. m is chosen sufficiently large so that nearly all of the structure of the original covariance matrix is maintained after the truncation. The distribution of eigenvalues is shown in Figure 13.3. As shown, after the first few eigenvectors are included, the eigenvalues, a measure of the given vector’s contribution to the total error, fall very quickly.

We saw that 85% of bins are covered to 99% with keeping just five PCA eigenvectors, so this was chosen for inclusion in the analysis fit. The error band for each flux component is shown in Figure 13.4 using all PPFX universes and only five PCA vectors. As the PCA under-quotes the necessary error, the absolute error associated with each PCA is scaled by 25% to conservatively over-cover the systematic.

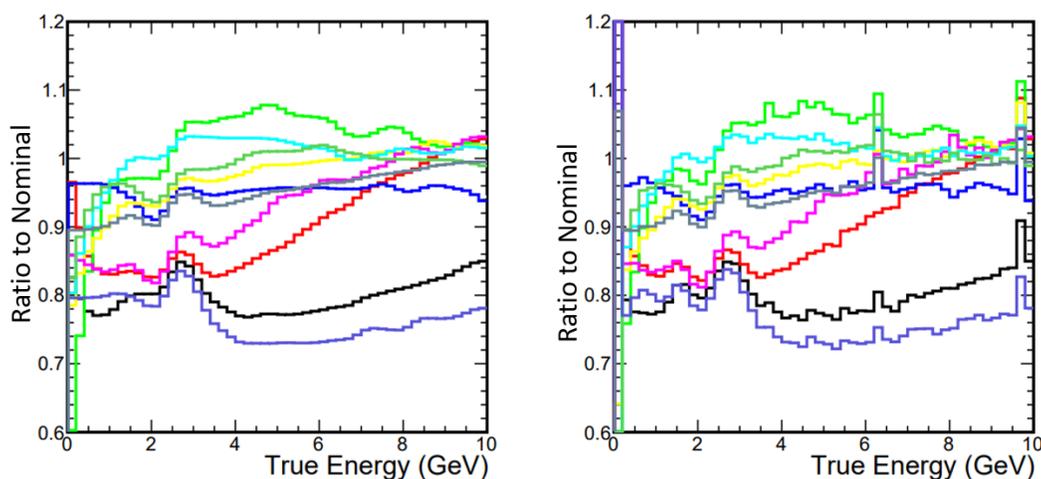


Figure 13.2: The ratios compared to the nominal flux for ten separate PPFX universes, each in a different color. The left shows the smoothed version, input into the above covariance matrix, and the right shows the raw distributions.

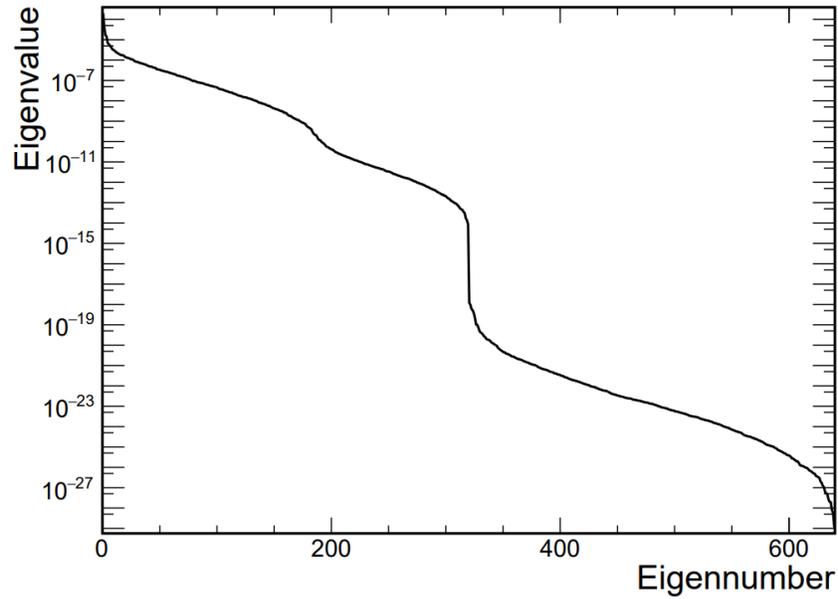


Figure 13.3: The distribution of ordered eigenvalue contribution to the flux error band. With just a few eigenvectors, the fraction of the error band not captured is already orders of magnitude down from the total.

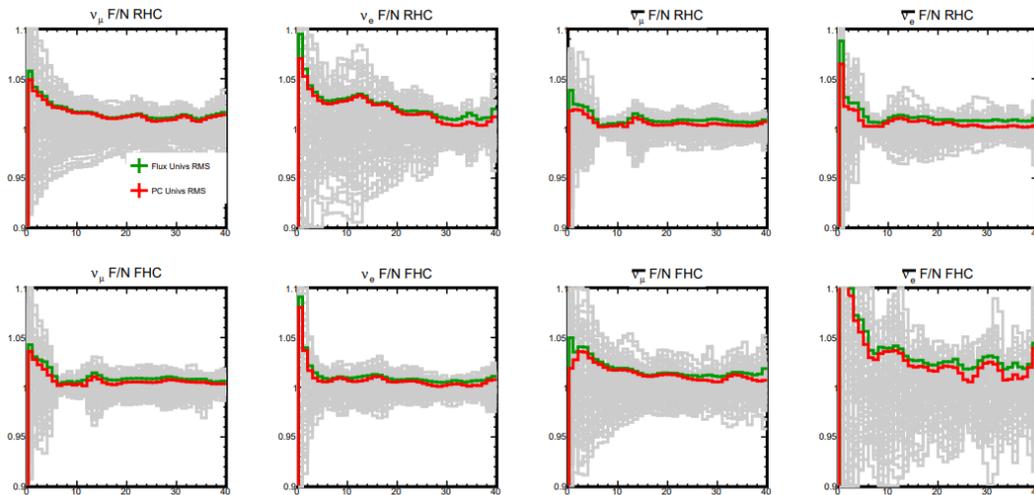


Figure 13.4: The error band from all PPFX universes, green, and the first five PCA eigenvectors, red, for anti-neutrino mode, top, and neutrino mode, bottom. The summed PCA eigenvectors fall short of the total curve for some energies. therefore, the eigenvectors are all scaled by 25% to compensate in a conservative manner.

13.2 MC Tuning for the Oscillation analysis

To accommodate recent advances in our understanding of neutrino scattering at relevant energies, we employ an in-situ tune of our near detector simulation to improve

agreement in several important kinematic variables. Incorporating these improved measurements in simulation allows us to reduce our systematic uncertainties relating to neutrino cross section uncertainties, which are among the largest sources of error in our oscillation analyses. This tune is primarily focused on adjusting the normalization and kinematic shape of GENIE predicted MEC events.

The axial mass for quasi-elastic events was slightly increased compared to the nominal simulation[132]. A recent re-analysis of old neutrino-Deuterium scattering data was performed[133]. With an updated understanding of analysis backgrounds, a more accurate error band was made of 1.04 ± 0.04 GeV by adding measurements from different bubble chamber experiments. We use a slightly more conservative uncertainty of 0.05 GeV in our analysis to account for potentially higher correlations between experiments than assumed in the calculation.

Additionally, the analysis incorporates a theoretical calculation of the Random Phase Approximation (RPA) influence on our QE and resonant simulated events. RPA is a charge-screening effect that reduces the cross section at very low- Q^2 [134]. Such a characteristic deficit in low- Q^2 has often been seen, even within NOvA's $CC\pi^0$ cross section analysis, whose relevant result is shown in Figure 9.40. Work is currently underway to improve calculations of this effect. We apply the RPA suppression as a function of Q^2 according to calculation. We've seen this significantly improve agreement between data and simulation of low Q^2 , forward events. For a conservative approach, we draw an error band on the underlying RPA model by taking the difference between our simulation with that model applied, and the simulation with no RPA correction made. Selected events in our near detector simulation with and without an RPA correction applied are shown in Figure 13.5 along with the comparison to observed data.

Additionally, we increase the normalization of GENIE's DIS prediction for events with $W > 1.7$ GeV[132]. At the $W = 1.7$ GeV boundary, GENIE switches the DIS model used. The underlying model GENIE uses in these events is designed to predict the cross section and very high neutrino energies. Thus, for events just over the 1.7 GeV cut, the model is unreliable. We increase the normalization of this component by 40% to match our observed data at very low track lengths. This change is accompanied with a 50% normalization uncertainty used in the simulation.

The final piece of the tune is treatment of MEC events, which also has the most dramatic effect on our simulated sample. We developed a fit in true $|\vec{q}|$, the three-vector component of the four-momentum transfer to the hadronic shower in the lab

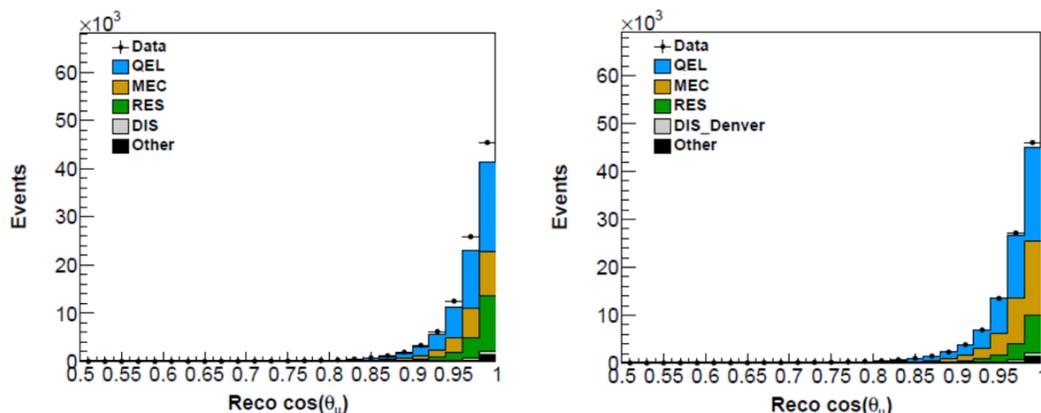


Figure 13.5: Candidate ν_μ CC events in the near detector in simulation compared to data. The left(right) plot shows a simulation without(with) an RPA correction applied with a MEC tune applied in both cases. This variable is not explicitly considered in the MEC tune. The disagreement for forward going muons without the RPA correction suggests our data prefers the inclusion of this low- Q^2 suppression.

frame, and true q_0 , the visible hadronic energy in the event. To optimize the fit, we compare the reconstructed data / MC ratio we observe in reconstructed $|\vec{q}|$ and q_0 with MEC events reweighted by the hypothesized weights.

The results of this fit are shown in Figures 13.6 and 13.7 for neutrino and anti-neutrino events, respectively. In general, the fit suggests that MEC events within our sample lie at a notably lower distribution of q_0 . This is true for both the neutrino and anti-neutrino tune.

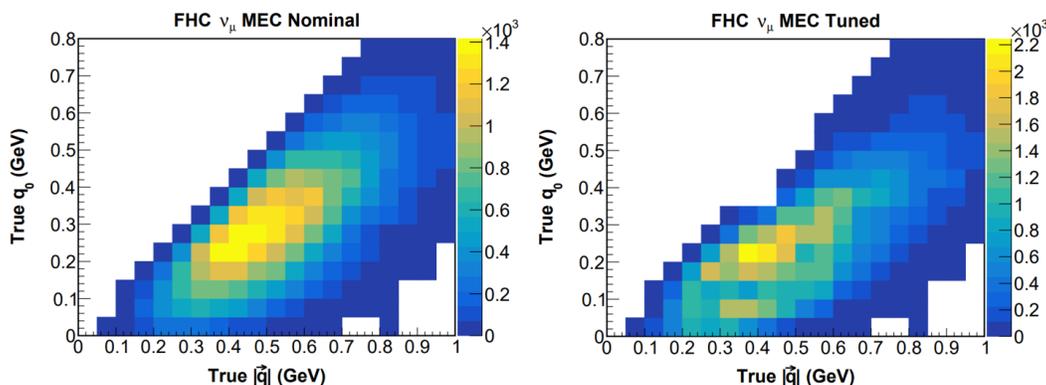


Figure 13.6: The true $|\vec{q}|$ vs q_0 distribution in the nominal MEC simulation from GENIE is shown on the left for neutrino mode. The right plot gives the distribution after it has been fitted in reconstructed space. The shape of the underlying MEC model has no theoretical or experimental justification, but by fitting in this space to our data, we can constrain the model to empirically agree with observed data.

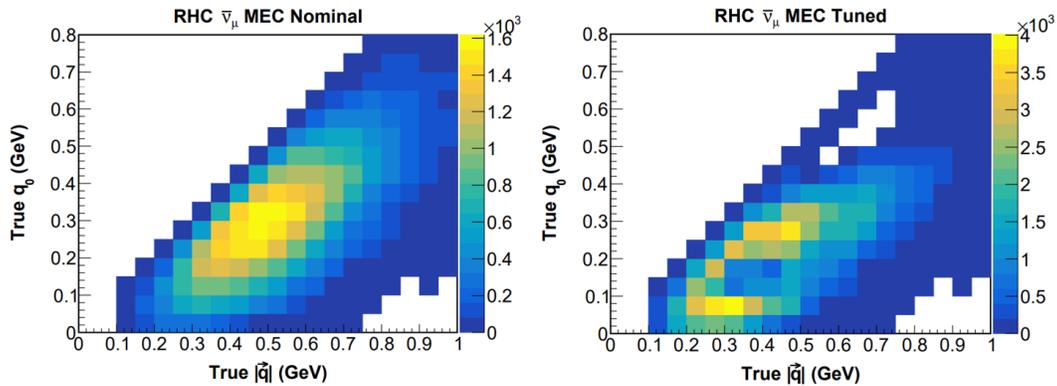


Figure 13.7: The true $|\vec{q}|$ vs q_0 distribution in the nominal MEC simulation from GENIE is shown on the left for anti-neutrino mode. The right plot gives the distribution after it has been fitted in reconstructed space. The shape of the underlying MEC model has no theoretical or experimental justification, but by fitting in this space to our data, we can constrain the model to empirically agree with observed data.

A systematic error band is drawn for our MEC events by systematically modifying the non-MEC components of our simulation and re-evaluating the fit. MEC events are known to lie between QE and resonant peaks in the hadronic system. Our systematic error band provides a procedure for coercing the MEC fit to span these two extremes. We jointly move eight systematic knobs in our simulation for this procedure:

- M_A for QE events: $+1\sigma$
- RPA suppression of QE events: $+1\sigma$
- RPA enhancement of QE events: $+1\sigma$
- Pauli-suppression of QE events: -1σ
- M_A for resonant events: -1σ
- M_V for resonant events: -1σ
- RPA suppression for resonant events: -1σ
- Calibration scale shifted: -1σ

Each of these knobs makes our simulation more “QE-like”, enhancing the QE component in the sample while suppressing. Thus, the MEC fit in the systematic

universe will be forced to mimic resonant events. Alternatively, we can make the MEC fit pull more “Res-like”, by applying each systematic knob in the opposite way. The spread in the hadronic energy distribution from the fitted MEC is shown in Figure 13.8. Indeed, the fits force the determined MEC distribution to be biased toward a Res-like or QE-like kinematic distribution. We thus evaluate a data-driven systematic error band from the shape of our MEC modeling by switching between weighting MEC events according to the fitted scale factors in true $|\vec{q}|$, q_0 space.

The MINERvA experiment developed a MEC tune that uses a philosophically very different approach on their own data[135]. However, the MINERvA tune, when applied to NOvA events, are in reasonable agreement. In fact, the difference between the NOvA and MINERvA tuned weights is just within the 1σ error band developed by our fitting procedure. This is general true in both neutrino and anti-neutrino data, shown in Figures 13.9 and 13.10, respectively. This is the first time that two independent experiments have been shown to approximately agree on a MEC tuning, suggesting that our understanding of MEC scattering is making notable progress.

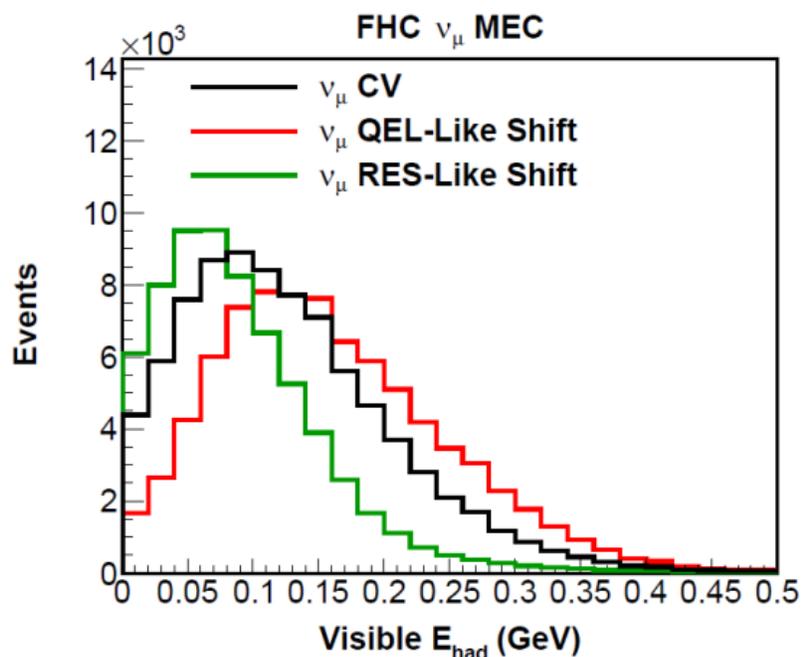


Figure 13.8: The hadronic energy distribution for fitted MEC events in the nominal procedure, black, with the sample shifted more QE-like, red, and more res-like, green.

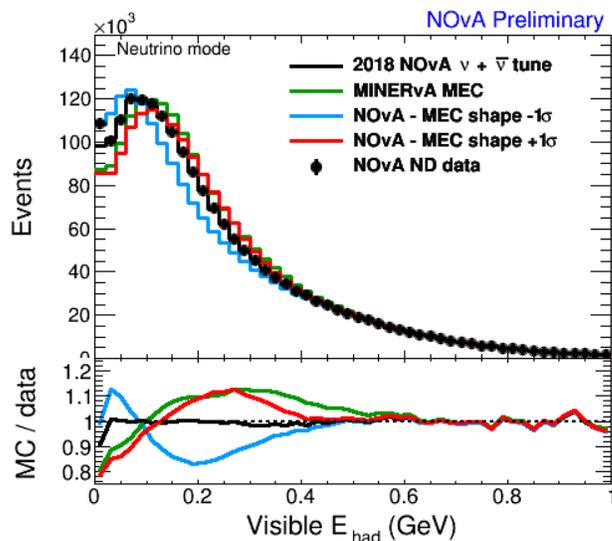


Figure 13.9: Our resulting hadronic energy distribution for neutrino running. Our nominal tune is in black, which agrees quite well with data. The red and blue curves span our error band due to the MEC-shape systematic described above. The green curve gives the prediction for NOvA simulation using the MINERvA MEC tune. The MINERvA tune lies quite close to the $+1\sigma$ curve.

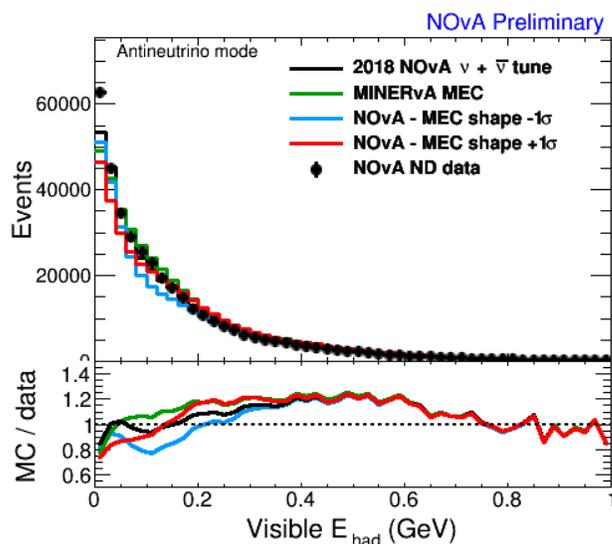


Figure 13.10: Our resulting hadronic energy distribution for anti-neutrino running. Our nominal tune is in black, which agrees quite well with data. The red and blue curves span our error band due to the MEC-shape systematic described above. The green curve gives the prediction for NOvA simulation using the MINERvA MEC tune. The consistency of the two weights is less clear than in the neutrino case, but the MINERvA tune never lies too far outside of the NOvA error band.

13.3 Neutrino Interaction Systematic Uncertainty

Neutrino interaction systematics for the oscillation analyses are calculated in an identical procedure to that of the $CC\pi^0$ cross section analysis. Using GENIE's reweighting functionality, for each systematic uncertainty, the simulation is adjusted to represent a distinct physical universe, in both detectors, and used to extrapolate the same near detector dataset. The spread in the predictions between the shifted universe's prediction and the nominal prediction is a systematic uncertainty. Unlike in the $CC\pi^0$ analysis, pull terms corresponding to each systematic error will be included in the likelihood fit used for constraining oscillation parameters. To encapsulate systematic information in the predictions, the bin-by-bin prediction is calculated for the nominal, ± 1 , 2 , and 3σ systematic shifts. A polynomial is then fit through these seven points for each bin, and the resulting polynomial coefficients are stored so that the predicted content in each bin at an arbitrary value of the systematic pull can be computed.

All systematic uncertainties provided by GENIE are included in the fit. Also included is the MEC shape systematic detailed in Section 13.2. For effects that influence the predicted spectra by more than the 1% level, the systematic uncertainty is calculated directly. For all other barely influential systematic sources, a PCA treatment, as in the flux uncertainty case, is used[136]. Thus, their influence is captured in the fit without adding many more fit parameters.

As an example of the process, Figure 13.11 shows the effect of adjusting the axial mass for CC resonant events for the ν_e appearance measurements. The effect on each HadEFrac quartile on the disappearance analysis shows similar cancellation through the extrapolation procedure.

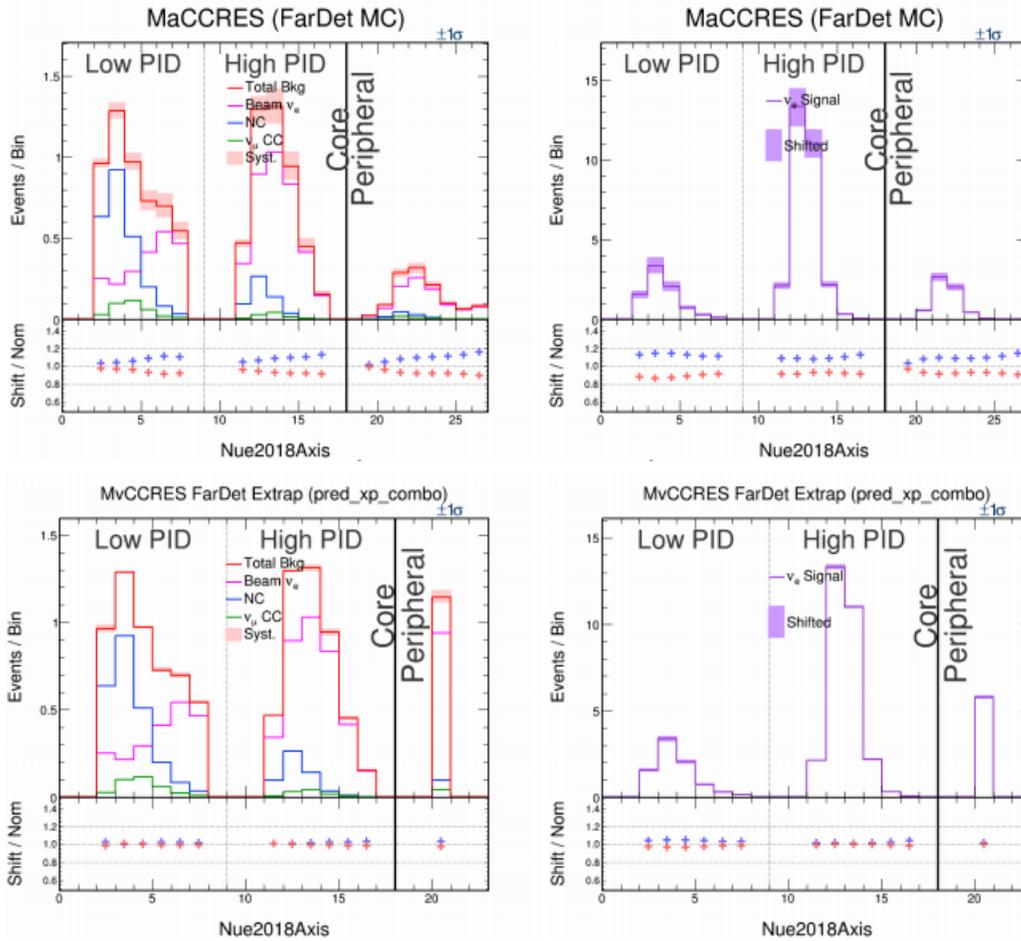


Figure 13.11: The effect of modifying the axial mass for resonant CC events by $\pm 1\sigma$ for ν_e appearance backgrounds, left, and signal, right. The top plots show the effect determined from the far detector simulation without extrapolation. The bottom shows the effect after our extrapolation procedure.

13.4 Acceptance Systematic

The near and far detectors are built of the same material and based on the same technology so that many systematic uncertainties are reduced through the extrapolation procedure. But they are very different sizes. The analysis acceptance is thus significantly different between the two detectors, implying an unfortunate model dependence in our prediction scheme. The most significant concern is in extrapolating the $\nu_\mu \rightarrow \nu_e$ signal, since the ν_e and ν_μ PID's sculpt the kinematics of selected events in different ways. Thus, a notable, data-driven systematic uncertainty is evaluated to cover our reliance on the neutrino generator model causing these acceptance artifacts for the ν_e appearance analysis.

To test geometric acceptance effects, the far detector predictions are constrained

with a near detector subsample, limiting the vertex range of selected distribution to the front/back, east/west, top/bottom, inner/outer halves of the detector. This gives eight different sub-regions used in the extrapolation. For neutrino running, the greatest difference between the prediction in any two sub-samples is 1.5(1.2)% for the core(peripheral) samples and 4.0(2.3)% for anti-neutrino running.

The kinematic differences between the near detector selected ν_μ and far detector selected ν_e events are studied with a reweighting procedure. The reweighting is done in three variables, Q^2 , $\cos \theta_{lep}$, and p_T/p , for neutrino and anti-neutrino running. For Q^2 , the migration matrices for ν_e 's in the far detector and ν_μ 's in the near detector are quite different due to differing reconstruction resolutions of the two samples. Thus, the weights are calculated by the ratio of true Q^2 in both populations, and applied to reconstructed Q^2 data and MC in the near detector when evaluating the systematic. First, near detector data and MC are weighted in each variable to match the far detector kinematic distribution for selected ν_e events. The extrapolation procedure is then rerun with these modified near detector data and MC events and a far detector prediction is made. The difference between the nominal and adjusted predictions is then taken as a systematic uncertainty. Only the reweighting variable that caused the largest bias relative to the nominal is quoted in the systematic error budget. This is $\cos \theta_{lep}$ for neutrino running, which leads to a 1.5% uncertainty, and Q^2 for anti-neutrino running, which gives a 2.8% uncertainty.

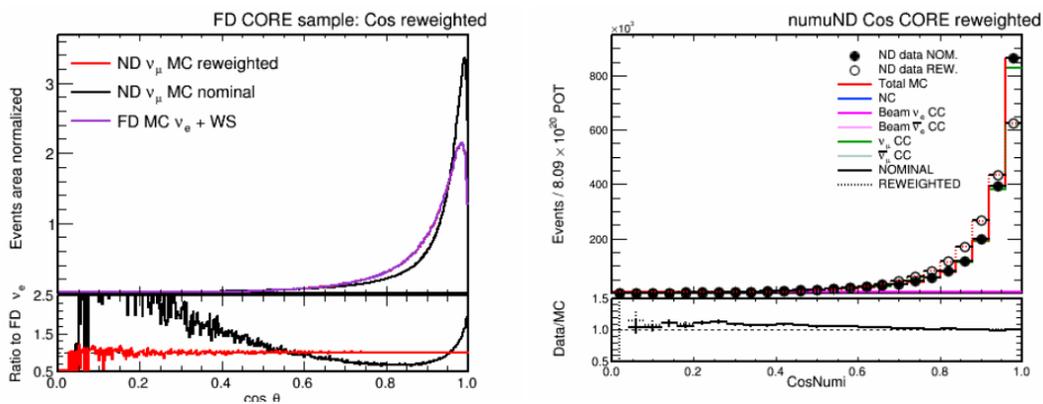


Figure 13.12: The kinematic reweighting procedure that gives the largest bias for anti-neutrino data, in weighting the event $\cos \theta_{lep}$. The left plot shows the $\cos \theta_{lep}$ distribution for near detector $\bar{\nu}_\mu$ selected events, black, the far detector $\bar{\nu}_e$ selected events, purple, and near detector reweighted events, red. The right shows the near detector ν_μ data / MC comparison before and after applying the reweights. The near detector is then extrapolated with these weights applied, and a systematic is taken as the difference between this prediction and the nominal

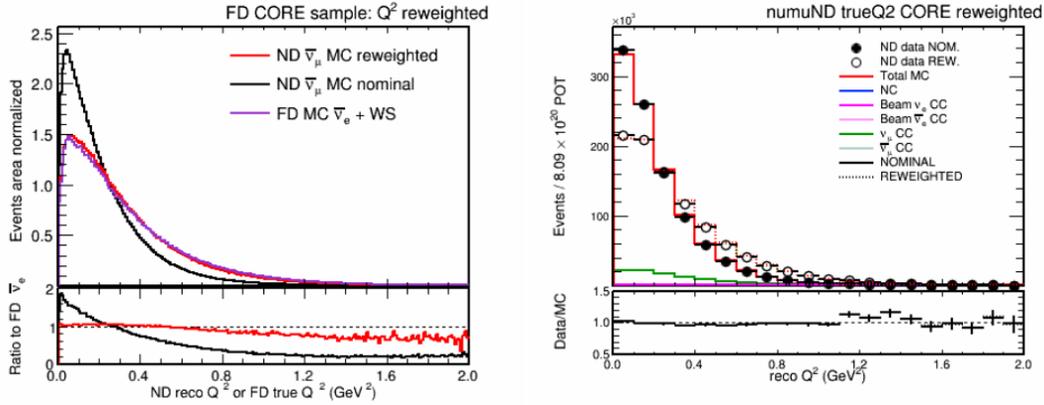


Figure 13.13: The kinematic reweighting procedure that gives the largest bias for anti-neutrino data, in weighting the event Q^2 . The left plot shows the Q^2 distribution for near detector $\bar{\nu}_\mu$ selected events, black, the far detector $\bar{\nu}_e$ selected events, purple, and near detector reweighted events, red. The right shows the near detector ν_μ data / MC comparison before and after applying the reweights. The near detector is then extrapolated with these weights applied, and a systematic is taken as the difference between this prediction and the nominal

13.5 Detector Response Systematics

A set of detector response and light level systematics are evaluated by adjusting our detector and brightness conditions and re-simulating events. In a similar strategy to other systematics, the prediction is re-evaluated in this modified simulation.

The absolute calibration is adjusted by $\pm 5\%$. This number is set by discrepancies in dE/dx distributions for proton tracks in the near detector. Additionally, our calibration procedure involves explicit functional forms, which do not perfectly reflect reality. Therefore, we evaluate systematic samples where we distort this functional form, in both x and y views separately, to cover all attenuation calibration discrepancies we see in the far detector.

There is also a systematic uncertainty stemming from uncertainty in the muon energy estimate. Since our muon energy is estimated from reconstructed tracklength, this uncertainty is much more related to our knowledge of our detector geometry and density than light yield. A survey of the simulation of our detector geometries gives a 0.94% uncertainty on the muon tracklength, and thus energy. Additionally, a 0.27% relative uncertainty is applied to account for differences in the near and far detectors.

Additionally, we apply a light level systematic where we shift the true light yield by $\pm 10\%$, which is then re-calibrated. The size of this shift is taken from the

measured spread in cell-by-cell brightness relative to the simulated value. This gives us an idea of how mismodeling true energy deposition of different charged particles can bias our reconstruction and PID algorithms. Importantly, this is also a proxy for estimating the effect of adjusting our thresholds. To first order, the light level reduction is compensated by increasing the calibration scale factor. But, if the light yield reduction is enough to push the hit below the detection threshold, it will be lost from the simulation entirely.

13.6 Normalization Systematics

A number of additional systematic uncertainties that do not influence the shape of our predictions are included as normalization systematics. There is a 0.5% systematic on the total accumulated POT and 0.27% error on the detector mass that influence our quoted exposure. There is also a 100% uncertainty on our estimate of non-fiducial events generated in the rock around our detector that is selected. This is a negligible part of the ν_μ spectrum and $\sim 0.3\%$ of the ν_e appearance spectrum. Additionally, there is a 60% normalization taken on the ν_τ cross section. This systematic is determined from the OPERA cross section measurement of $\nu_\mu \rightarrow \nu_\tau$ oscillations[29].

There is a normalization systematic attributed to our mis-modeling of our reconstruction efficiency due to event pileup. To calculate this, single simulated neutrinos are overlaid onto spills in data and simulation. The reconstruction efficiency of these overlaid simulated events is then compared in the two scenarios. This gives a 1.3(0.3)% difference in neutrino and anti-neutrino mode[137], which is included in quadrature with the above normalization systematics.

13.7 Overview of Systematic Effects

An overview of the effect of the error sources that influence our analysis is shown in Figure 13.14. Compared to the statistical error, our systematic errors are still small, but in neutrino mode, they are large enough to slightly expand the total error estimate for the analysis. The largest systematic errors come from neutrino interaction uncertainties. The acceptance systematic uncertainties are also quite large for predicted signal rates. The influence these systematics have on our physics sensitivity is shown in Chapter 14.

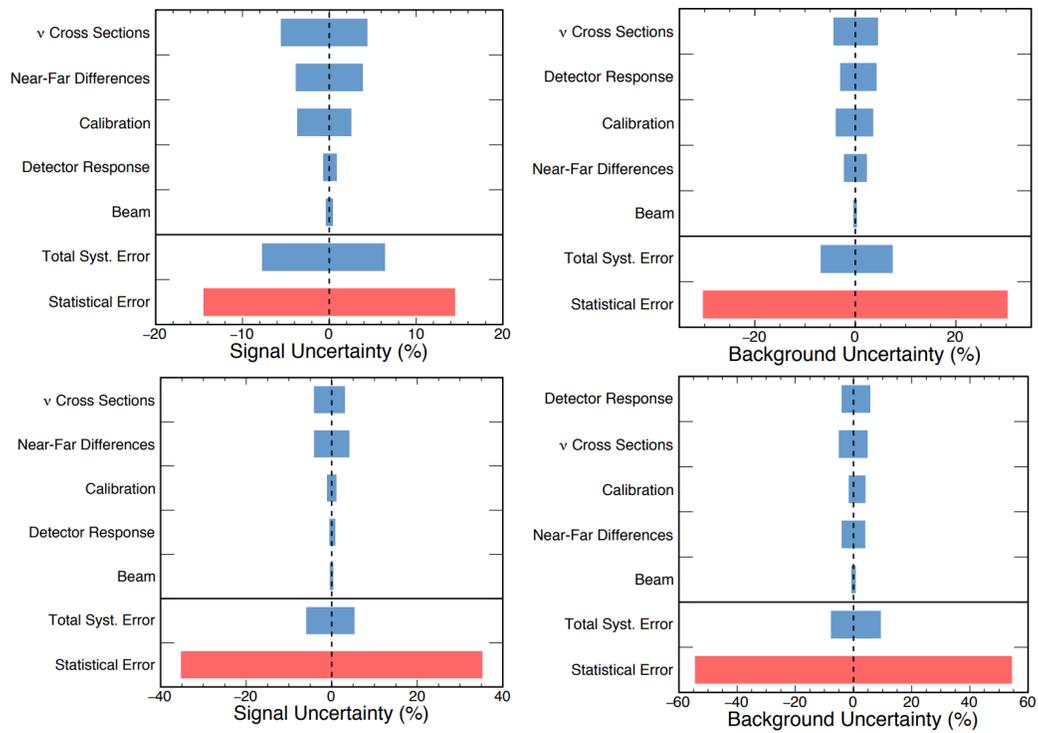


Figure 13.14: Systematic, blue, and statistical, red, errors on the total signal and background rates for the ν_e , top, and $\bar{\nu}_e$, bottom appearance spectra. The left(right) plots compare the errors on the signal(background) event rate. Particularly for $\bar{\nu}_e$ appearance, statistical errors dominate.

Chapter 14

FITTING, SENSITIVITY, AND CONFIDENCE INTERVAL CONSTRUCTION

Now, after detailing the far detector selections and how we use near detector data to constrain our simulated far detector predictions, we can describe the fitting and confidence region determination for the joint oscillation analysis. The ν_μ disappearance channel is sensitive to $\sin^2 2\theta_{23}$ and $|\Delta m_{32}^2|$. The ν_e appearance channel is also sensitive to these parameters, but to a lower precision. However, the appearance channel is uniquely sensitive to the neutrino mass hierarchy, the CP violating phase, and the θ_{23} octant. Thus, the simultaneous analysis of both of these channels gives the best handle to resolve the degeneracies in the oscillation probabilities and thus measure the mixing properties of interest.

We first perform a log-likelihood fit between oscillated predictions and the observed data. The analysis profiles over all nuisance parameters in the system. A map of $\Delta\chi^2$ values over possible oscillation parameter space is constructed. Then, a pseudo-experiment method for constructing confidence intervals is employed. This is necessary due to low event counts and the presence of non-Gaussian fit parameters: a binary parameter, the neutrino mass hierarchy; a periodic parameter, δ_{CP} ; and a parameter that lies near a physical boundary, $\sin^2 2\theta_{23} < 1$. We follow the Feldman-Cousins algorithm[138] for determining suitable confidence intervals, as outlined below.

14.1 Log-Likelihood Fit

The physical parameters of interest measured by the analysis are $\sin^2 \theta_{23}$, Δm_{32}^2 , and δ_{CP} . Additionally, the oscillation parameter, $\sin^2 2\theta_{13}$, affects the oscillation probabilities, but is constrained according to the latest world-fit on the value, 0.082 ± 0.004 [82], and included in the fit as a nuisance parameter with a pull term. The same is true for the solar oscillation parameters: $\Delta m_{12}^2 = 7.59 \pm 0.21 \times 10^{-5}$ eV² and $\sin^2 \theta_{12} = 0.297 \pm 0.017$. But, as these oscillations occur on much longer baselines than relevant for our measurements, the solar oscillation parameters have a minimal impact on our measurement. Systematic uncertainties are included in the fit as nuisance parameters with a pull term parameterized by their prior error.

The fit maximizes the Poisson log-likelihood function as defined in Eqn. 14.1[82]. Here, the sum runs over all bins in each of the four oscillation channels, $\boldsymbol{\theta}$ is a vector of fit parameters including systematic nuisance parameters, ν_i gives the predicted number of events in each bin for a choice of $\boldsymbol{\theta}$, and n_i gives the observed number of events in each bin. The second sum gives the contribution to $\log \mathcal{L}$ from systematic pull terms with δ_i the fitted value of the systematic and σ_i the prior error on that systematic.

$$2 \log \mathcal{L}(\boldsymbol{\theta}) = -2 \sum_{i=1}^N \left[\nu_i(\boldsymbol{\theta}) - n_i + n_i \ln \frac{n_i}{\nu_i(\boldsymbol{\theta})} \right] - \sum_{i=1}^M \frac{\delta_i^2}{\sigma_i^2} \quad (14.1)$$

The log-likelihood is re-cast as a χ^2 via $\chi^2(\boldsymbol{\theta}) = -2 \log \mathcal{L}(\boldsymbol{\theta})$. The vector of parameters that minimize $\chi^2(\boldsymbol{\theta})$, and maximize $\log \mathcal{L}(\boldsymbol{\theta})$, are written as $\hat{\boldsymbol{\theta}}$. Using this notation, a $\Delta\chi^2(\boldsymbol{\theta})$ relative to the minimum χ^2 can be defined as

$$\Delta\chi^2(\boldsymbol{\theta}) = \chi^2(\boldsymbol{\theta}) - \chi^2(\hat{\boldsymbol{\theta}}) \quad (14.2)$$

for an arbitrary choice of the parameters, $\boldsymbol{\theta}$. As defined, $\Delta\chi^2 \geq 0$, and is a well-ordered test statistic in that if $\Delta\chi^2(\boldsymbol{\theta}_1) > \Delta\chi^2(\boldsymbol{\theta}_2)$, $\boldsymbol{\theta}_2$ is more likely than $\boldsymbol{\theta}_1$ given the measurement. With this $\Delta\chi^2$ function, our observed oscillated spectra, and their predictions, a $\Delta\chi^2$ map relative to the best fit is calculated for each possible vector of fit parameters.

14.2 Sensitivities

Before making a far detector measurement and drawing corresponding contours, it is useful to make sensitivity plots to understand how well our experiment can limit certain sets of oscillation parameters. As the ν_e appearance probability depends on so many parameters, we can not make a human-readable contour plot that encompasses all oscillation parameter space. We thus draw 2D contours and 1D “slice” plots that treat certain relevant oscillation parameters as nuisance parameters. For example, we will draw the contour in $\sin^2 \theta_{23}$ and Δm_{32}^2 , often called the “atmospheric” contour for historical reasons. Though δ_{CP} is physically interesting parameter measured by the analysis, it is treated as a nuisance parameter for the atmospheric contour.

Figures 14.1 - 14.4 show our projected sensitivity to oscillation parameters for our accumulated exposure in neutrino and anti-neutrino running. Fake-data at the Asimov prediction for the NOvA 2017 best fit has been used in place of data for the fits. The Asimov prediction gives predicted spectra without any statistical

fluctuations, a measure of the analysis’s median sensitivity. It is named for Isaac Asimov, whose short story, “Franchise”, describes an election procedure where a nation’s most typical person singularly decides the outcome. The plots show both the statistics-only fit and the fit including systematic effects. None of these plots are drawn with significances determined by Feldman-Cousins, due to the significant computing resources involved. Thus, the contours here are drawn at a constant $\Delta\chi^2$ up-value.

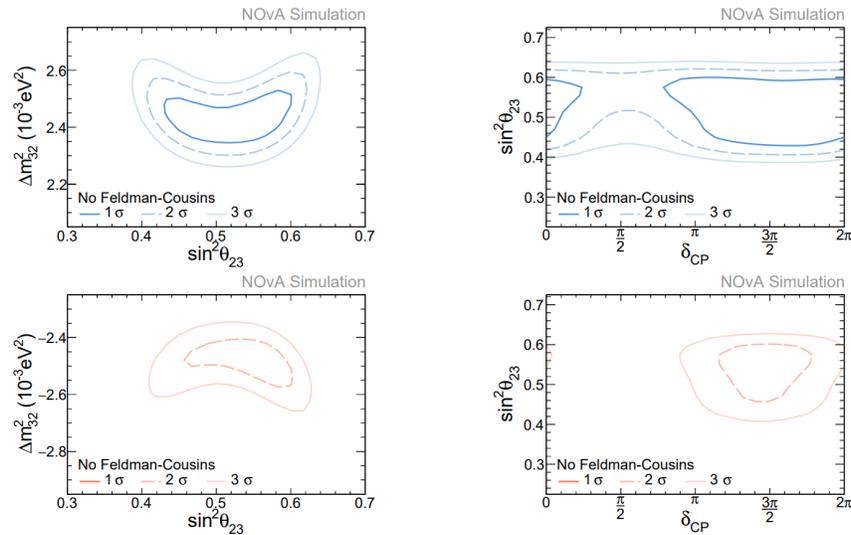


Figure 14.1: Our statistics-only sensitivity to our atmospheric contour, left, and δ_{CP} vs $\sin^2 \theta_{23}$, right. The top plot shows the NH allowed parameter space with the bottom the IH region. Contours are drawn at the corresponding two-dimensional Gaussian $\Delta\chi^2$ up-values.

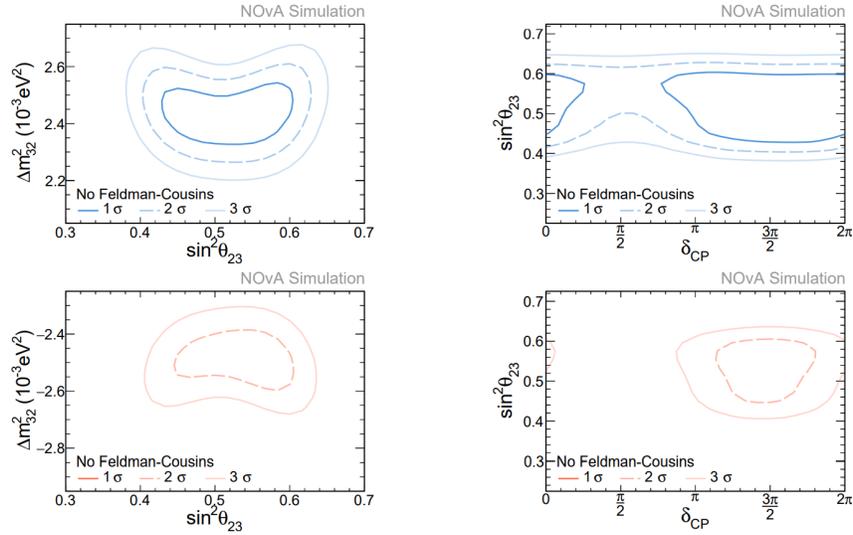


Figure 14.2: Our sensitivity to our atmospheric contour, left, and δ_{CP} vs $\sin^2 \theta_{23}$, right, including systematic effects. The top plot shows the NH allowed parameter space with the bottom the IH region. Contours are drawn at the corresponding two-dimensional Gaussian $\Delta\chi^2$ up-values.

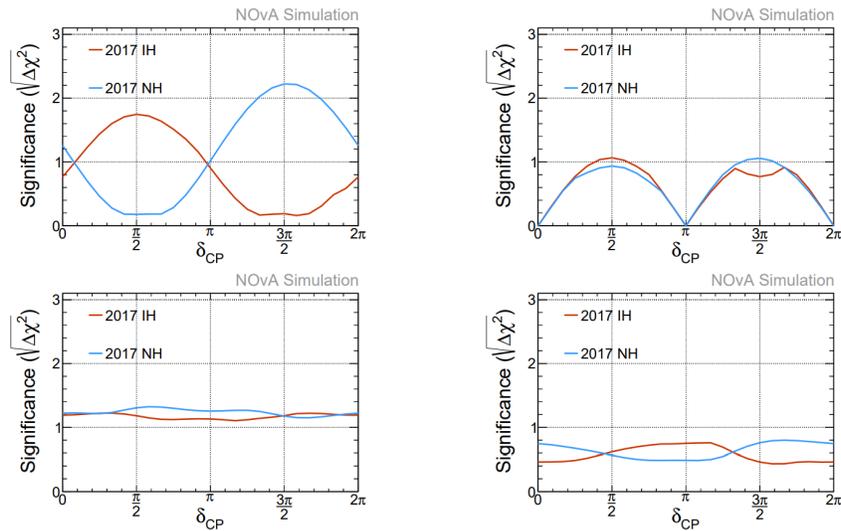


Figure 14.3: Our statistics-only sensitivity to wrong hierarchy rejection, top left, CP violation, top right, the $\theta_{23} = \pi/4$ hypothesis, bottom left, and the θ_{23} octant, bottom right as a function of δ_{CP} . These slices are drawn at the corresponding one-dimensional Gaussian $\Delta\chi^2$ up-values.

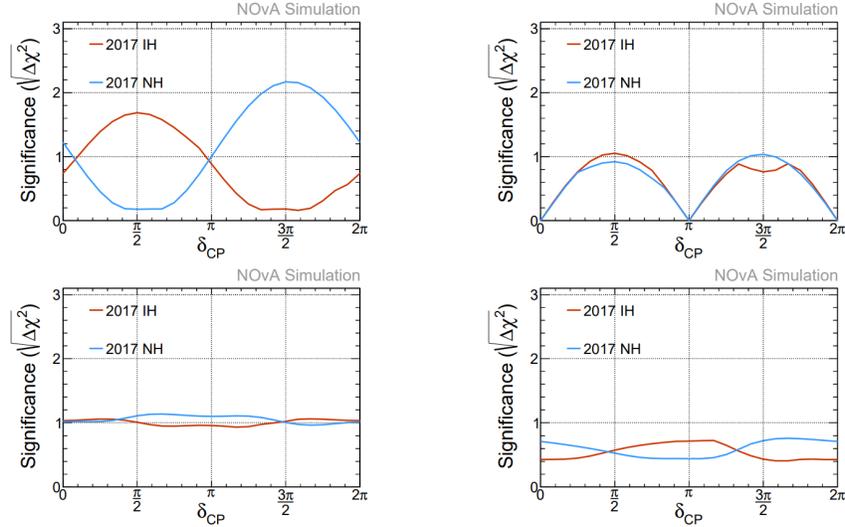


Figure 14.4: Our sensitivity to wrong hierarchy rejection, top left, CP violation, top right, the $\theta_{23} = \pi/4$ hypothesis, bottom left, and the θ_{23} octant, bottom right as a function of δ_{CP} . The fits include systematic uncertainties. These slices are drawn at the corresponding one-dimensional Gaussian $\Delta\chi^2$ up-values.

14.3 Feldman-Cousins Pseudo-experiments

The unified approach from Feldman-Cousins[138] gives a way of estimating a confidence of inclusion at each point in the measured parameter space. The method requires an input parameter, c , which gives the desired level of confidence. It then gives a way of constructing a confidence interval with this coverage for some true parameter, α , with a list of nuisance parameters, η .

The procedure prescribes that each point ζ that α can take be directly tested for inclusion in the confidence interval. A series of pseudo-experiments are performed to derive the mapping between a measured $\Delta\chi^2_\zeta$ and a p -value for any choice of ζ . In each pseudo-experiment, a Poisson fluctuated “mock-data” histogram is thrown from the Asimov prediction with α set to ζ , and two fits are performed. The first is minimized over all parameters, (α, η) , and gives a minimum χ^2 labeled as χ^2_{best} . The second fit holds α fixed to ζ , but all parameters in the vector, η , are allowed to float. This fit gives a minimum χ^2 labeled as χ^2_{true} . The distribution of $\Delta\chi^2_\zeta = \chi^2_{\text{true}} - \chi^2_{\text{best}}$ is recorded for many pseudo-experiments. A critical $\Delta\chi^2_{\zeta, \text{crit}}$ is determined by finding the smallest value of $\Delta\chi^2_\zeta$ that is higher than the fraction, c , of thrown pseudo-experiments. After calculating this critical value, the minimum $\Delta\chi^2(\zeta, \hat{b}\hat{m}\hat{\eta})$ is found from the $\Delta\chi^2$ surface determined from the fit to the observed experimental data with α fixed to ζ but with $\hat{\eta}$ profiled. Finally, if this observed $\Delta\chi^2(\zeta, \eta)$ is

less than $\Delta\chi_{\zeta,\text{crit}}^2$, it is included in the confidence interval at the confidence level, c ; otherwise it is excluded.

Unfortunately, the $\Delta\chi_{\zeta,\text{crit}}^2$ determined depends slightly on the values of the nuisance parameters used to seed the pseudo-experiments. Thus, if thrown from values that do not represent the physical universe, the procedure no longer produces confidence intervals with proper coverage determined by construction. Generally, the most conservative nuisance parameters to use in seeding the pseudo-experiments are those at the best fit to the experimental data, $\hat{\eta}$, for every tested point, ζ [139].

It is hypothesized that this seed will give the most conservative confidence interval. However, this need not be true, and an experiment ought to explicitly show their contours drawn by the Feldman-Cousins procedure does not under-cover. We follow the Berger and Boos prescription for testing coverage[140]. In this method, the coverage is tested for a wide range of values away from the best fit. For each test, nuisance parameters are no longer thrown at the best fit to data, $\hat{\eta}$, but at some arbitrary tested vector, $\hat{\eta} + \hat{\delta}$. But, for most variables, this involves checking parameters over an infinite range of hypotheses for the nuisance parameters. Further, these hypotheses become arbitrarily incompatible with the measurement when testing values many σ away from their best fit value. So, Berger and Boos argue that the nuisance parameter must only be tested for possible values more consistent with the measurement than some arbitrary precision, p , which is much less than the significance an experiment quotes at a particular point. The statistical rejection for the given hypothesis is then degraded by p , the fraction of the PDF for the nuisance parameter that is not tested.

Balancing the value of the Berger and Boos statistical cross check and the computational expense of running Feldman Cousins pseudo-experiments, we run this check on our most important single number, the rejection of the hierarchy we do not prefer, for the nuisance oscillation parameters: $\sin^2 \theta_{23}$, δ_{CP} , and Δm_{32}^2 . Results from this cross check confirm our intuition that throwing pseudo-experiments from our best fit to our data gives the most conservative rejection value we could quote for all possible values the oscillation values could have taken[141]. Parameters more than 3σ away from the best fit were not tested, and thus the confidence we quote in rejecting the disfavored hierarchy is reduced by 0.25%.

OSCILLATION RESULTS

Having defined our procedure for selecting events and predicting the far detector spectra with the near detector constraint, we are now ready to make a measurement of the far detector oscillated spectra. The observed spectra and fits to simulated predictions are shown below. In summary, the number of selected events in each oscillation channel is shown in Figure 15.1.

| Osc. Channel | Observed | Expected Bkg |
|---|----------|--------------|
| $\nu_\mu \rightarrow \nu_\mu$ | 113 | 3.76 |
| $\nu_\mu \rightarrow \nu_e$ | 58 | 15.06 |
| $\bar{\nu}_\mu \rightarrow \bar{\nu}_\mu$ | 65 | 1.31 |
| $\bar{\nu}_\mu \rightarrow \bar{\nu}_e$ | 18 | 5.32 |

Table 15.1: A list of the observed event counts in each analysis channel studied. Also, the expected background contribution to the sample is shown.

This is the first observation of $\bar{\nu}_e$ appearance at $> 4\sigma$ confidence. Using a pseudo-experiment method accounting for systematic distortions and statistical fluctuations on the background, including wrong-sign $\nu_\mu \rightarrow \nu_e$ appearance, a p -value of 3.3×10^{-5} was determined for the background-only hypothesis fluctuating to 18 or more events. This corresponds to 4.1σ evidence for $\bar{\nu}_e$ appearance. We disfavor the maximal mixing hypothesis, $\theta_{23} = \pi/4$ to 1.8σ . Additionally, we favor the upper octant for θ_{23} at 1.7σ . Our best fit of δ_{CP} is at 0.166π , consistent with CP conservation, and we prefer the normal mass ordering at 1.77σ .

15.1 Observed Far Detector Spectra

The ν_e and $\bar{\nu}_e$ distributions are shown in Figures 15.1 and 15.2, respectively. The predictions for signal, background, and total event counts are shown in Table 15.2 for various choices of oscillation parameters. The total number of events expected for ν_e and $\bar{\nu}_e$ appearance at different combinations of oscillation parameters are compared to the observed data and shown in Figure 15.3. Many of these combinations of oscillation parameters are disfavored to a high degree given the event distributions observed. The observed event count in both samples is higher than the median of those predictions listed in the table.

Comparisons of our selected ν_e and $\bar{\nu}_e$ appearance events, compared to constrained predictions for relevant kinematic and selection variables, are shown in Figures 15.4-15.7. The observed data is consistent with the predictions.

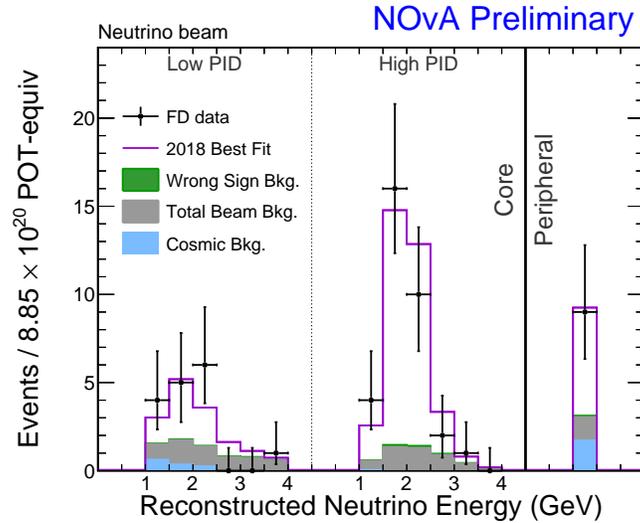


Figure 15.1: The observed ν_e CC appearance spectrum in each selection bin. The blue and gray stacked histograms give the cosmic and total background, respectively. The excess gives the total prediction at the best fit to all data.

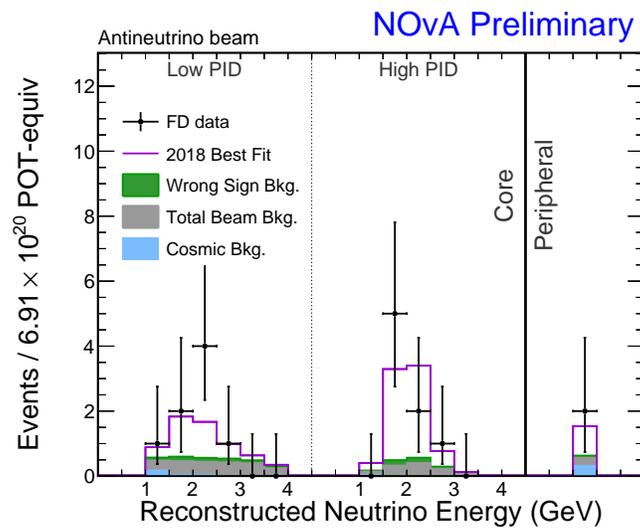


Figure 15.2: The observed $\bar{\nu}_e$ CC appearance spectrum in each selection bin. The blue and gray stacked histograms give the cosmic and total background, respectively. The purple gives the total prediction at the best fit to all data.

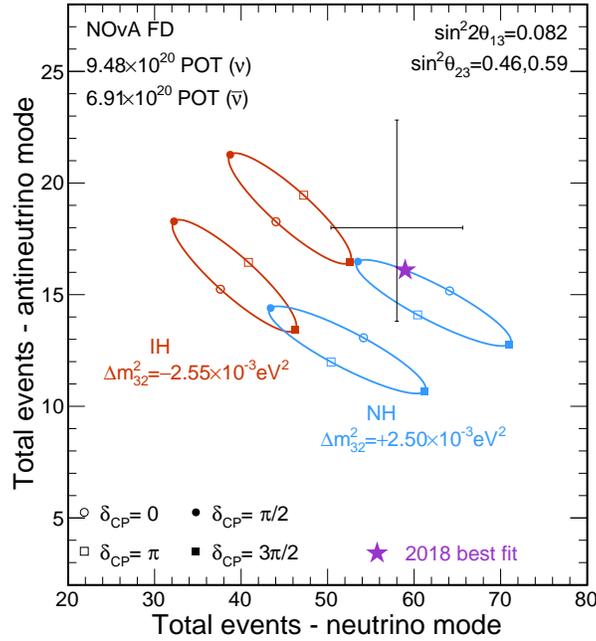


Figure 15.3: The number of ν_e , x -axis, and $\bar{\nu}_e$, y -axis, selected events are shown in black. The colored ellipses show predictions for various values of δ_{CP} and the θ_{23} octant. The purple star gives the number of events in both modes predicted at the 2018 best fit.

| Osc. Params | neutrino mode | | | anti-neutrino mode | | |
|---|---------------|-------|-------|--------------------|-------|------|
| | Total | Sig | Bkg | Total | Sig | Bkg |
| NH, $\delta = \pi/2$, $\sin^2\theta_{23} = 0.4$ | 38.35 | 23.52 | 14.82 | 13.08 | 8.36 | 4.72 |
| NH, $\delta = \pi/2$, $\sin^2\theta_{23} = 0.5$ | 45.91 | 31.12 | 14.79 | 14.81 | 9.89 | 4.91 |
| NH, $\delta = \pi/2$, $\sin^2\theta_{23} = 0.6$ | 54.03 | 39.08 | 14.95 | 16.46 | 11.34 | 5.12 |
| NH, $\delta = 3\pi/2$, $\sin^2\theta_{23} = 0.4$ | 55.58 | 41.00 | 14.58 | 9.38 | 4.22 | 5.17 |
| NH, $\delta = 3\pi/2$, $\sin^2\theta_{23} = 0.5$ | 63.50 | 48.96 | 14.54 | 11.03 | 5.66 | 5.37 |
| NH, $\delta = 3\pi/2$, $\sin^2\theta_{23} = 0.6$ | 71.26 | 56.56 | 14.71 | 12.76 | 7.19 | 5.57 |
| IH, $\delta = \pi/2$, $\sin^2\theta_{23} = 0.4$ | 28.94 | 13.72 | 15.22 | 16.32 | 11.88 | 4.44 |
| IH, $\delta = \pi/2$, $\sin^2\theta_{23} = 0.5$ | 33.68 | 18.44 | 15.24 | 18.75 | 14.18 | 4.56 |
| IH, $\delta = \pi/2$, $\sin^2\theta_{23} = 0.6$ | 38.90 | 23.44 | 15.45 | 21.08 | 16.38 | 4.70 |
| IH, $\delta = 3\pi/2$, $\sin^2\theta_{23} = 0.4$ | 42.42 | 27.51 | 14.92 | 11.59 | 6.79 | 4.79 |
| IH, $\delta = 3\pi/2$, $\sin^2\theta_{23} = 0.5$ | 47.44 | 32.51 | 14.93 | 13.91 | 8.99 | 4.92 |
| IH, $\delta = 3\pi/2$, $\sin^2\theta_{23} = 0.6$ | 52.39 | 37.23 | 15.15 | 16.34 | 11.29 | 5.05 |
| Observed Data | 58 | - | - | 18 | - | - |

Table 15.2: Expected event counts for a variety of values of δ_{CP} , $\sin^2\theta_{23}$, and both hierarchies. Here, Δm_{32}^2 is set to ± 0.0025 eV^2 , depending on the hierarchy, and $\sin^2\theta_{23}$ is set to the reactor constraint central value at 0.0082.

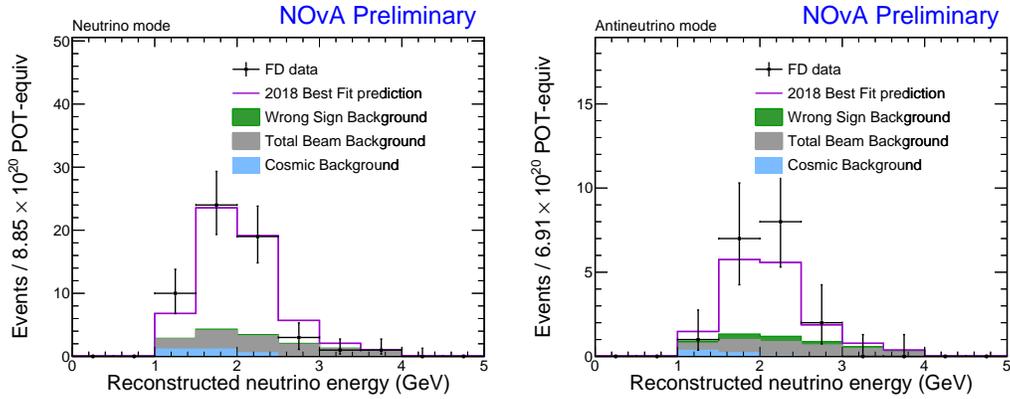


Figure 15.4: The reconstructed energy distribution observed in the neutrino, left, and anti-neutrino, right, samples. The spectra are compared to the best fit prediction.

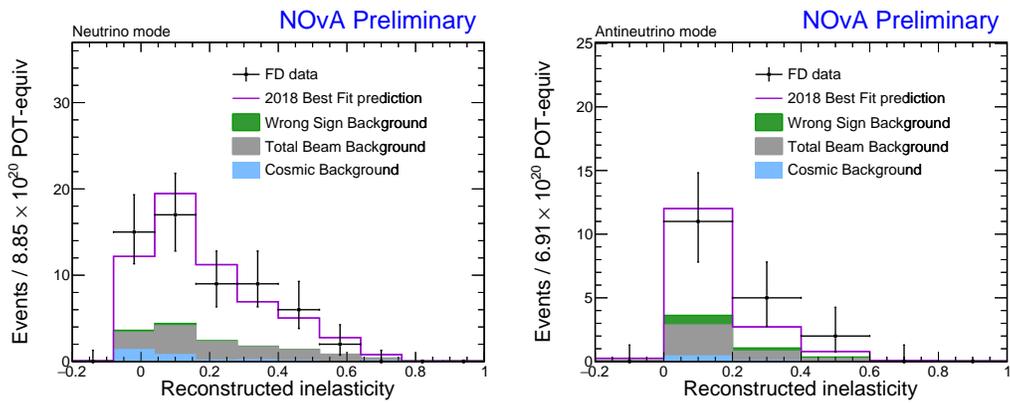


Figure 15.5: The reconstructed inelasticity, $1 - E_{\text{lep}}/E_{\nu}$, distribution observed in the neutrino, left, and anti-neutrino, right, samples. The spectra are compared to the best fit prediction.

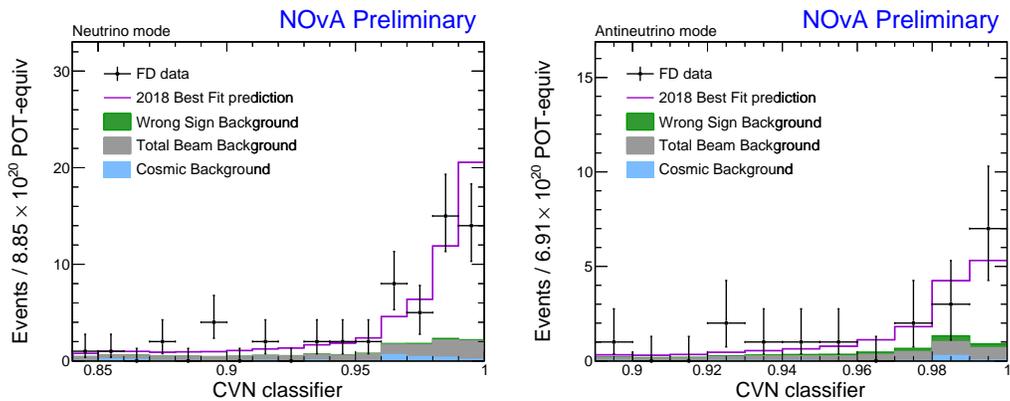


Figure 15.6: The CVN distribution of observed in the neutrino, left, and anti-neutrino, right, samples. Bins correspond to the PID bins used in the analysis. The spectra are compared to the best fit prediction.

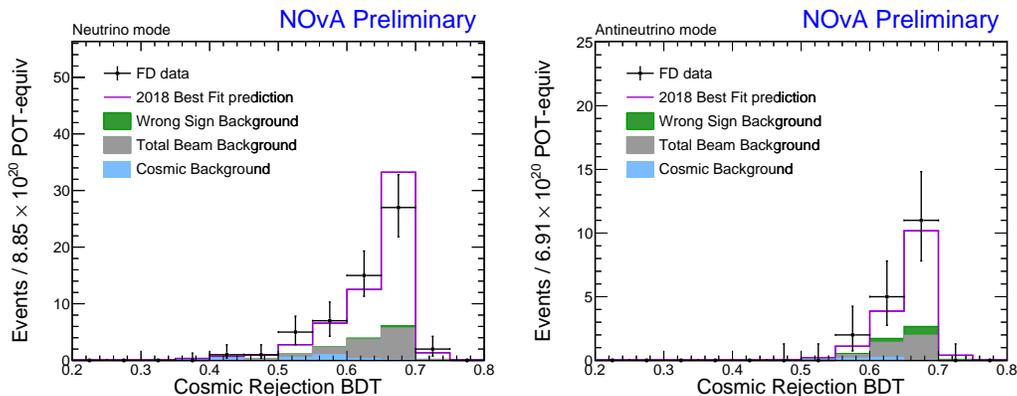


Figure 15.7: The cosmic rejection BDT distribution, used in the peripheral sample definition, observed in the neutrino, left, and anti-neutrino, right, samples. The spectra are compared to the best fit prediction.

Meanwhile, in the disappearance channels, the 113 and 65 events observed in the ν_μ and $\bar{\nu}_\mu$ spectra are shown along with their appropriate best fit predictions in Figures 15.8 and 15.9. The same distributions, summed over all four HadEFrac quartiles, are shown in Figures 15.10 and 15.11. Table 15.3 shows the number of ν_μ and $\bar{\nu}_\mu$ disappearance events observed in each quartile of HadEFrac in comparison to the best fit. Without oscillations, we would have expected a total of 730(266) events in (anti-)neutrino mode.

The data / MC distributions at the far detector for two interesting kinematic variables, the reconstructed muon energy and HadEFrac, is shown in Figures 15.12 and 15.13. The observed data is consistent with the best fit predictions.

A Komlogorov-Smirnov test[142] between the observed events in each sample against the accumulated exposure was performed, shown in Figure 15.14. This results in consistency between the two distributions to 20.5(75.8)% for the ν_μ ($\bar{\nu}_\mu$) distributions and 22.6(43.1)% for ν_e ($\bar{\nu}_e$) distributions.

| | ν_μ Sample | $\bar{\nu}_\mu$ Sample |
|---------------|------------------|------------------------|
| All Quartiles | 113 (130.90) | 65 (54.14) |
| Quartile 1 | 32 (33.58) | 17 (14.55) |
| Quartile 2 | 25 (30.67) | 14 (13.52) |
| Quartile 3 | 26 (31.14) | 19 (12.62) |
| Quartile 4 | 30 (35.51) | 15 (13.45) |

Table 15.3: The observed(predicted) number of events at the best fit in each quartile for the ν_μ and $\bar{\nu}_\mu$ disappearance spectra.

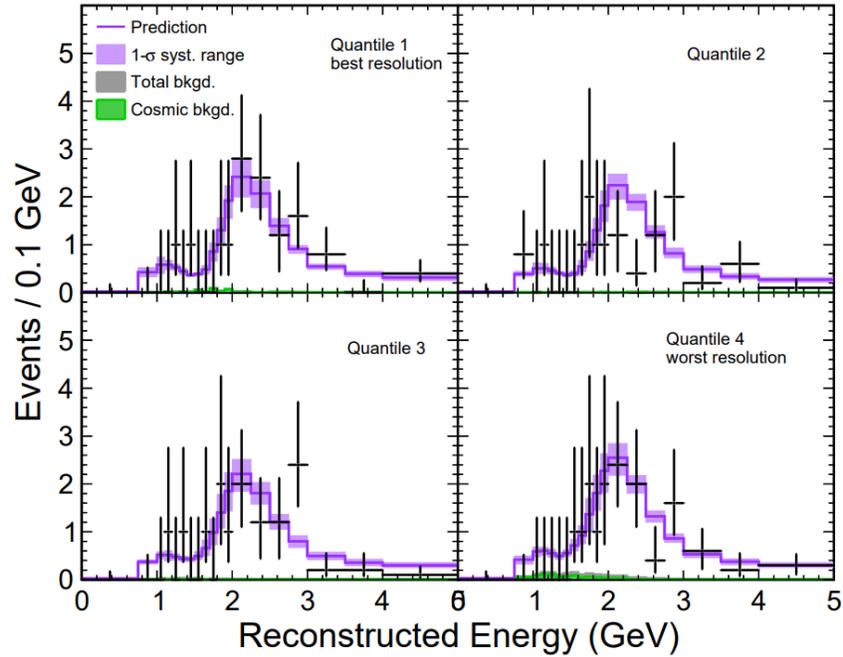


Figure 15.8: The observed ν_μ CC disappearance spectrum in each of the HadEFrac quartiles. The green and gray stacked histograms give the cosmic and total background, respectively. The purple band on the simulated prediction gives the total systematic error.

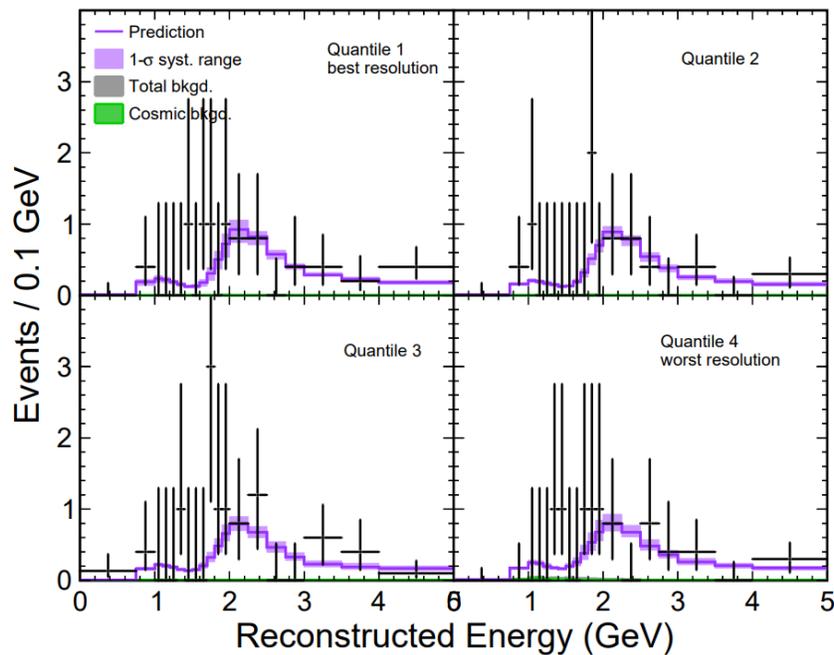


Figure 15.9: The observed $\bar{\nu}_\mu$ CC disappearance spectrum in each of the HadEFrac quartiles. The green and gray stacked histograms give the cosmic and total background, respectively. The purple band on the simulated prediction gives the total systematic error.

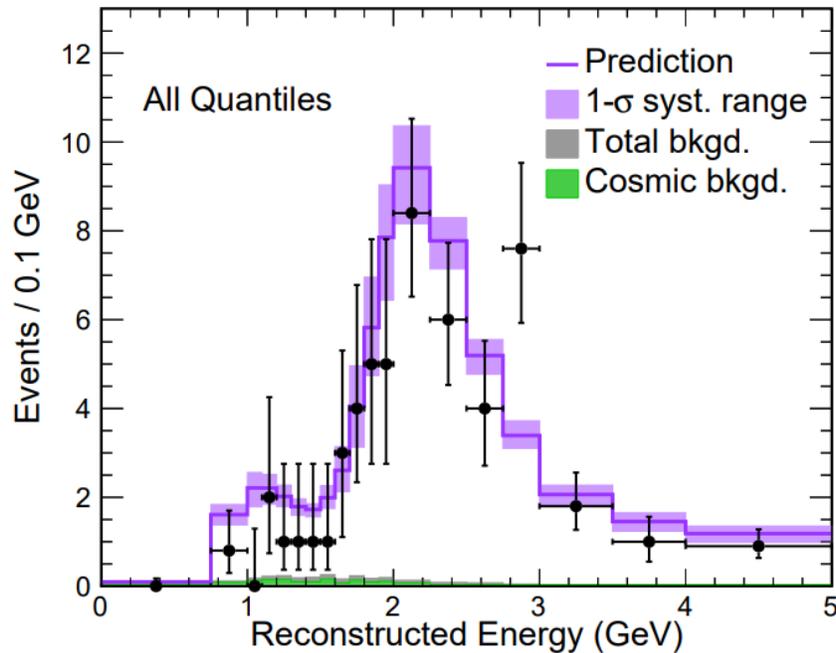


Figure 15.10: The observed ν_μ CC disappearance spectrum summed over all HadE-Frac quartiles. The green and gray stacked histograms give the cosmic and total background, respectively. The purple band on the simulated prediction gives the systematic error band.

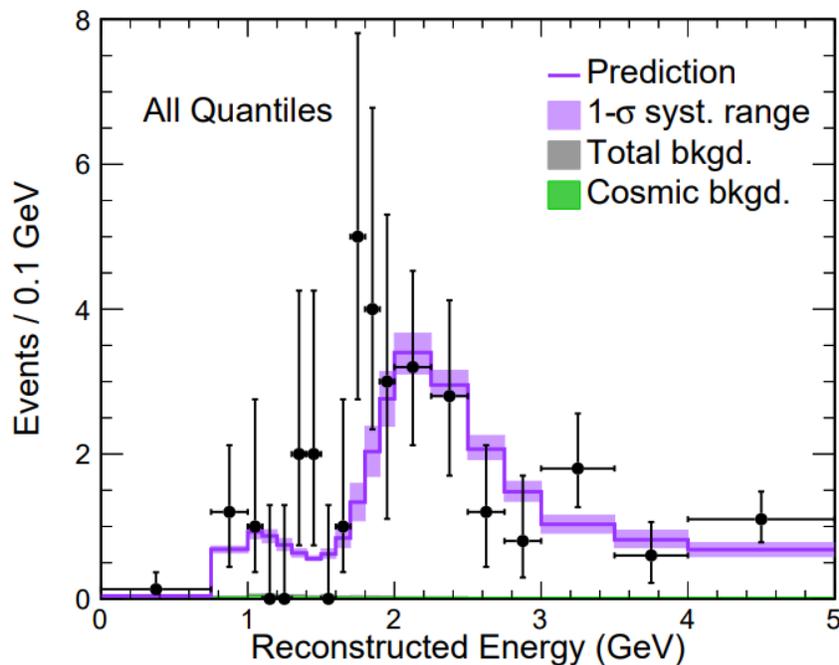


Figure 15.11: The observed $\bar{\nu}_\mu$ CC disappearance spectrum summed over all HadE-Frac quartiles. The green and gray stacked histograms give the cosmic and total background, respectively. The purple band on the simulated prediction gives the systematic error band.

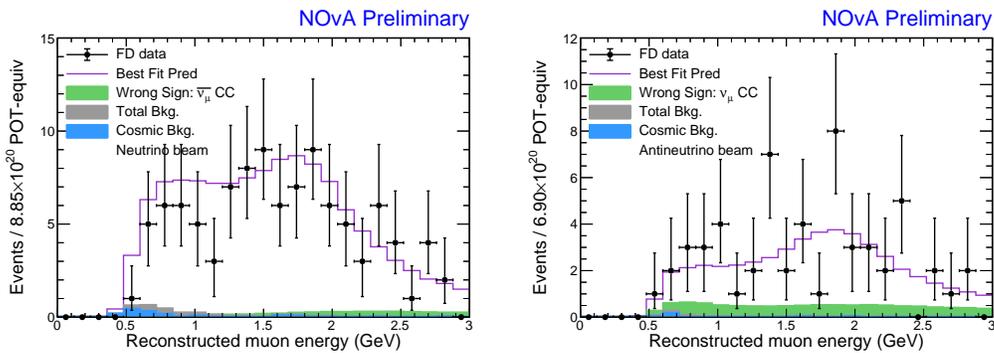


Figure 15.12: The observed data / MC distribution in reconstructed muon energy for events selected in the ν_μ , left, and $\bar{\nu}_\mu$, right, disappearance samples.

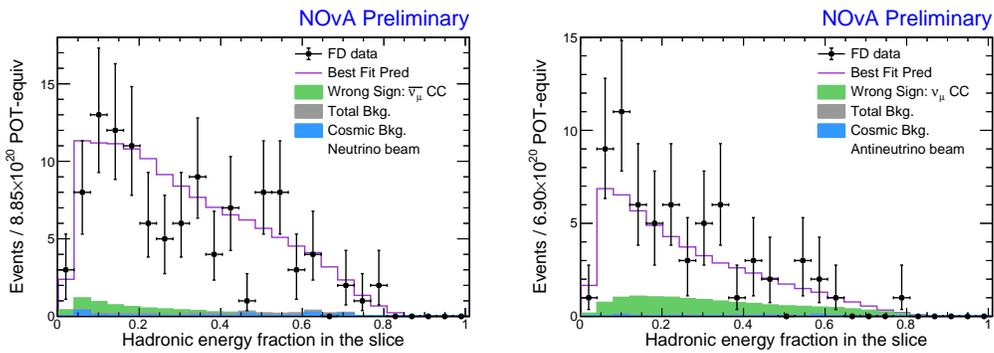


Figure 15.13: The observed data / MC distribution in reconstructed HadEfrac for events selected in the ν_μ , left, and $\bar{\nu}_\mu$, right, disappearance samples. This is the kinematic variable used to separate the analysis sample into quartile bins.

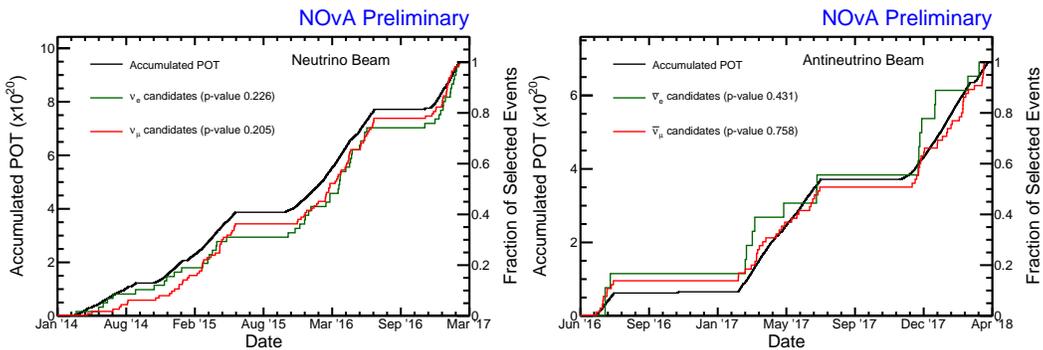


Figure 15.14: The observed ν_μ , red, and ν_e , green, events at the far detector compared to the total accumulated exposure in neutrino data, left, and anti-neutrino data, right. A KS consistency p -value was determined for each case. Out of the four curves, the minimum p -value was 0.205.

15.2 Oscillation Fits and Constraints on Physics Parameters

With observed far detector spectra, we can now simultaneously fit these to their respective predictions. As already mentioned in Chapter 14, a $\Delta\chi^2$ surface is drawn using a log-likelihood fit for both 2D contours of pairs of oscillation parameters and 1D slices that have had all but one oscillation parameter profiled away. In each plot, oscillation parameters that are not explicitly shown are profiled in the fit, along with all other nuisance parameters.

The allowed values of $\sin^2 \theta_{23}$ and δ_{CP} for each hierarchy hypothesis are shown in Figure 15.15. In this plot, the influence of Δm_{32}^2 has been profiled away. The 1σ contour lies entirely within the upper octant. All but a very small region of the 1σ contour is in the NH, centering around $\delta_{CP} = \pi/2$. The 2σ allowed region accepts the majority of the NH parameter space allowed by the disappearance channel, but disfavors most of the IH region in the lower octant and with δ_{CP} between 0 and π . The best fit is in the NH with $\sin^2 \theta_{23} = 0.584$, $\delta_{CP} = 0.166\pi$, and $\Delta m_{32}^2 = 0.00251 \text{ eV}^2$.

Figure 15.16 shows how consistent our data is with a given δ_{CP} and hierarchy hypothesis. Here, the effect of $\sin^2 \theta_{23}$ has been profiled away. As suggested by the 2D contours, the data is most consistent with the NH with δ_{CP} between 0 and π . We disfavor the δ_{CP} range outside of 1 and 2π in the IH to more than 2σ . The overall, profiled $\Delta\chi^2$ between the two hierarchy hypothesis is 2.38. With a Feldman-Cousins procedure, this corresponds to a 1.77σ preference for the normal mass ordering.

The contour in Δm_{32}^2 and $\sin^2 \theta_{23}$ is shown in Figure 15.17. Here, δ_{CP} has been profiled. The variables $\sin^2 \theta_{23}$ and Δm_{32}^2 are drawn individually in slice plots in Figures 15.18 and 15.19, respectively. As shown in the plot, both the maximal mixing hypothesis and entire lower octant are inconsistent with the observed data at nearly the 2σ level.

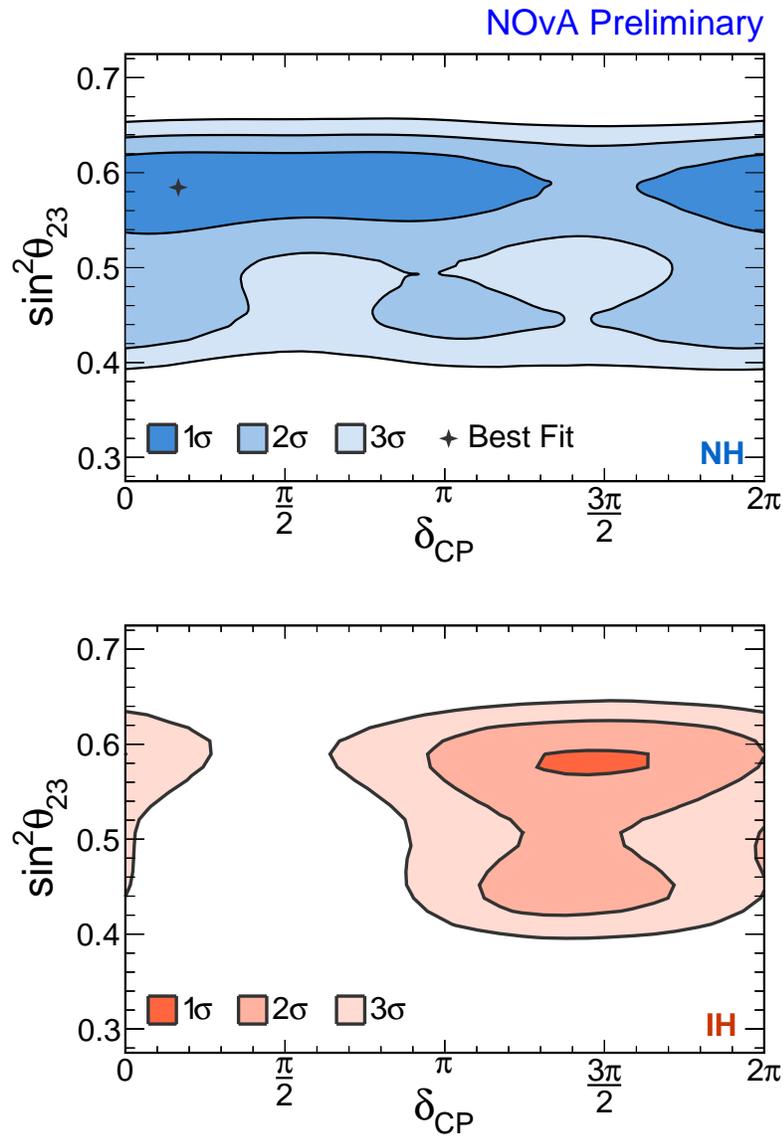


Figure 15.15: The allowed parameter space in $\sin^2 \theta_{23}$ and δ_{CP} according to our joint fit of disappearance and appearance in neutrino and anti-neutrino data. The 2σ allowed region is confined to about π to 2π in the IH, while just limiting the parameter space in the lower octant near $\pi/2$ and $3\pi/2$ for the NH.

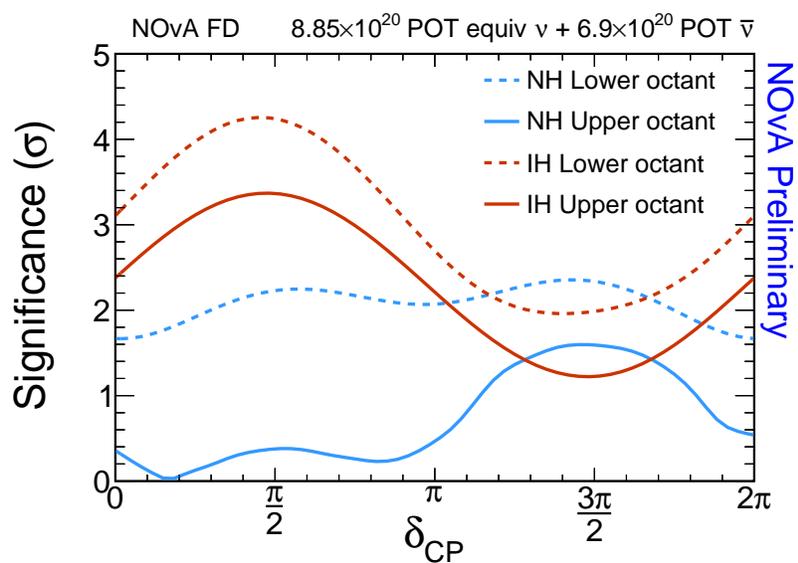


Figure 15.16: The significance at which we are consistent with values of δ_{CP} in the NH, blue, and IH, red. We disfavor $\delta_{CP} = 3\pi/2$ in both hierarchies to a modest degree. An overall $\Delta\chi^2 = 2.38$ is measured between the two hierarchy hypotheses with a preference for the normal mass hierarchy. The data gives a Feldman-Cousins preference of 1.77σ for the normal mass hierarchy.

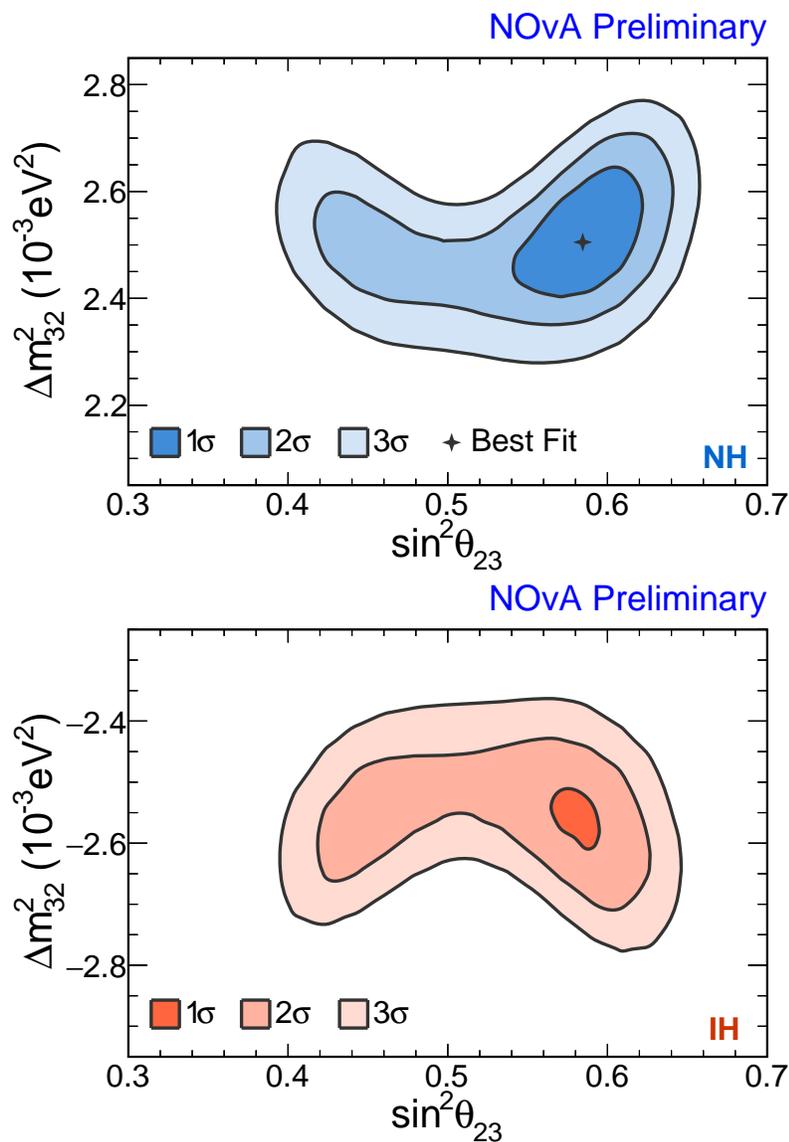


Figure 15.17: The contours in Δm_{32}^2 and $\sin^2 \theta_{23}$ according to our joint fit of disappearance and appearance in neutrino and anti-neutrino data. The observed data is clearly more consistent with oscillation parameters in the upper octant. Similarly, there is a preference for the NH, as already seen.

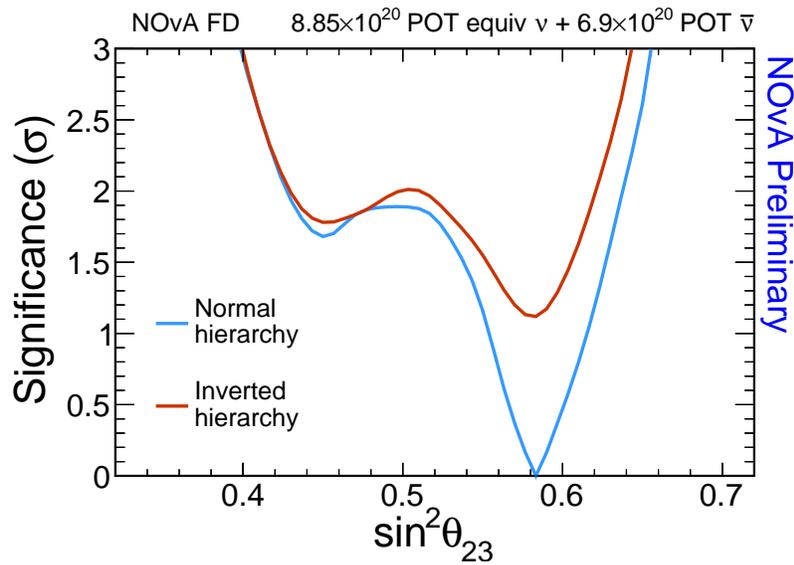


Figure 15.18: The significance at which we are consistent with values of $\sin^2 \theta_{23}$ in the NH, red, and IH, blue. Due to relatively large rates of ν_e and $\bar{\nu}_e$ appearance relative to the median expectations and the ν_μ and $\bar{\nu}_\mu$ disappearance differences discussed in the text, there is a relatively strong preference for the upper octant. The lower octant is preferred at 1.7σ and the maximal mixing hypotheses is ruled out to 1.8σ .

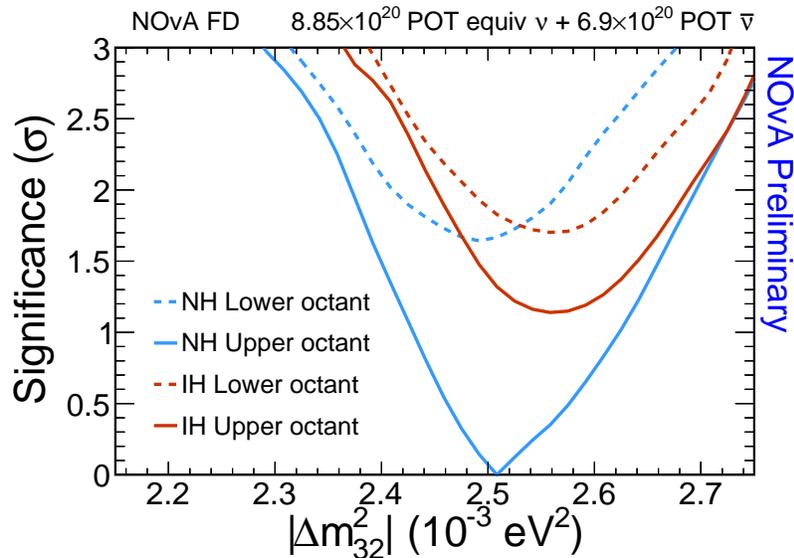


Figure 15.19: The significance at which we are consistent with values of Δm_{32}^2 in the NH, red, and IH, blue, showing a preference for the NH.

Chapter 16

SUMMARY

NOvA has placed leading measurements on neutrino oscillation parameters that influence the $\nu_\mu \rightarrow \nu_\mu$, $\bar{\nu}_\mu \rightarrow \bar{\nu}_\mu$, $\nu_\mu \rightarrow \nu_e$, and $\bar{\nu}_\mu \rightarrow \bar{\nu}_e$ oscillation channels. In the appearance channels, 58 ν_e and 18 $\bar{\nu}_e$ candidate events are seen in the sample. This corresponds to 4.1σ evidence for $\bar{\nu}_e$ appearance. For disappearance, 113 ν_μ and 65 $\bar{\nu}_\mu$ candidates were identified. The combined fit of these four oscillation channels is sensitive to many open questions in the neutrino sector: Is there CP violation in leptons? Is ν_3 more massive than $\nu_{1,2}$? Is θ_{23} exactly $\pi/4$? If not, is it greater than or less than $\pi/4$? These questions are intimately connected with phenomenological models that use leptogenesis to produce a preference of matter in the universe[39] and predict perturbations to the tri-bi-maximal neutrino mixing scheme[143][144].

The inverted mass hierarchy is disfavored at 1.77σ . The data suggests $\theta_{23} > \pi/4$ at 1.7σ , a suggestive departure from maximal mixing. At the accumulated exposure representing the data analyzed herein, the statistical significance for answering these questions remains tenuous. As statistical errors in the far detector dominate the analysis uncertainties, continued running will increase the experiment's sensitivity to these questions. The parameter space currently allowed puts NOvA in position to resolve some of these questions within its scheduled running period over the next six years. In this time, NOvA is expected to triple and quadruple its total exposure in neutrino and anti-neutrino running, respectively. By 2024, NOvA is projected to measure CP-violation and the mass hierarchy for about 40% and 70% of possible δ_{CP} values, respectively, to 95%. This is shown in Figure 16.1.

At the moment, there is tension with the oscillation results from NOvA and T2K[40], setting the stage for important results in the future. The T2K best fit lies at $\delta_{CP} = 3\pi/2$ in the normal mass hierarchy which is disfavored by NOvA data to $\Delta\chi^2 = 2.2$. Increasing exposure collected by both experiments and a continued understanding of systematic uncertainties, particularly those that affect these experiments very differently, will be vital to understanding this apparent discrepancy. A detailed study of neutrino cross sections in the few-GeV region, especially involving exclusive final state topologies, is becoming more important for proper measurement of neutrino oscillations. Results from NOvA's $CC\pi^0$ analysis will improve neutrino generators'

prediction of the neutrino samples in this energy regime. This channel is especially important for understanding π^0 -producing background neutrino interactions in both ν_μ CC and NC events. This and future cross section measurements will be important for mitigating systematic uncertainties as T2K, NOvA, and eventually DUNE and Hyper-K accumulate larger far detector neutrino samples.

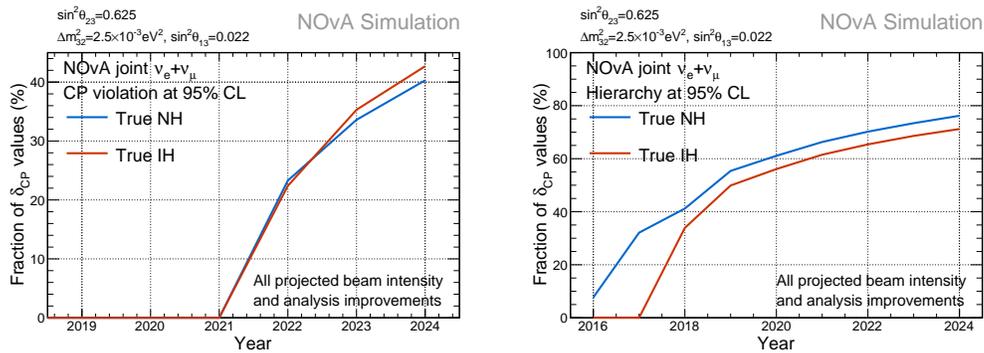


Figure 16.1: The fraction of possible δ_{CP} values for which NOvA can reject leptonic CP conservation, left, and determine the mass hierarchy, right, projected as a function of time with assumed exposure. These sensitivities are made assuming the true parameters lie in the normal mass hierarchy and upper θ_{23} octant.

BIBLIOGRAPHY

- [1] W. Pauli. “A letter to a meeting on Radioactivity”. 1930.
- [2] J. Chadwick. “Intensitätsverteilung im magnetischen Spektren der β -Strahlen von Radium B + C”. In: *Verhandlungen der Deutschen Physikalischen Gesellschaft* 16 (1914), pp. 383–391.
- [3] E. Fermi. “Versuch einer Theorie der beta-Strahlen”. In: *Zeitschrift für Physik* 88.161 (1934).
- [4] C. Cowan et al. “Detection of the Free Neutrino: A Confirmation”. In: *Science* 124.3212 (1956).
- [5] M. Goldhaber, L. Grodzins, and A. W. Sunyar. “Helicity of Neutrinos”. In: *Phys. Rev.* 109 (3 Feb. 1958), pp. 1015–1017. doi: 10.1103/PhysRev.109.1015. URL: <https://link.aps.org/doi/10.1103/PhysRev.109.1015>.
- [6] G. Danby et al. “Observation of High-Energy Neutrino Reactions and the Existence of Two Kinds of Neutrinos”. In: *Phys. Rev. Lett.* 9 (1 July 1962), pp. 36–44. doi: 10.1103/PhysRevLett.9.36. URL: <https://link.aps.org/doi/10.1103/PhysRevLett.9.36>.
- [7] K. Kodama et al. “Observation of tau neutrino interactions”. In: *Physics Letters B* 504.3 (2001), pp. 218–224. ISSN: 0370-2693. doi: [https://doi.org/10.1016/S0370-2693\(01\)00307-0](https://doi.org/10.1016/S0370-2693(01)00307-0). URL: <http://www.sciencedirect.com/science/article/pii/S0370269301003070>.
- [8] V.M. Hannen. “The KATRIN neutrino mass experiment”. In: *PoSICHEP2012* (2013), p. 374.
- [9] Benjamin Monreal and Joseph A. Formaggio. “Relativistic cyclotron radiation detection of tritium decay electrons as a new technique for measuring the neutrino mass”. In: *Phys. Rev. D* 80 (5 Sept. 2009), p. 051301. doi: 10.1103/PhysRevD.80.051301. URL: <https://link.aps.org/doi/10.1103/PhysRevD.80.051301>.
- [10] D.M. Asner et al. “Single electron detection and spectroscopy via relativistic cyclotron radiation”. In: *Phys. Rev. Lett.* 114.16 (2015), p. 162501. doi: 10.1103/PhysRevLett.114.162501. arXiv: 1408.5362 [physics.ins-det].
- [11] A. Gando et al. “Search for Majorana Neutrinos Near the Inverted Mass Hierarchy Region with KamLAND-Zen”. In: *Phys. Rev. Lett.* 117 (8 Aug. 2016), p. 082503. doi: 10.1103/PhysRevLett.117.082503. URL: <https://link.aps.org/doi/10.1103/PhysRevLett.117.082503>.

- [12] Ziro Maki, Masami Nakagawa, and Shoichi Sakata. “Remarks on the Unified Model of Elementary Particles”. In: *Progress of Theoretical Physics* 28.5 (1962), pp. 870–880. DOI: 10.1143/PTP.28.870. eprint: /oup/backfile/content_public/journal/ptp/28/5/10.1143/ptp.28.870/2/28-5-870.pdf. URL: +%20http://dx.doi.org/10.1143/PTP.28.870.
- [13] S.P. Mikheyev and A. Smirnov. “Resonance enhancement of oscillations in matter and solar neutrino spectroscopy”. In: *Soviet Journal of Nuclear Physics* 32.6 (1985).
- [14] L. Wolfenstein. “Neutrino oscillations in Matter”. In: *Phys Rev D* 17.9 (1978).
- [15] K. Zuber. *Neutrino Physics*. CRC Press, 2011.
- [16] J. A. Formaggio and G. P. Zeller. “From eV to EeV: Neutrino cross sections across energy scales”. In: *Rev. Mod. Phys.* 84 (3 Sept. 2012), pp. 1307–1341. DOI: 10.1103/RevModPhys.84.1307. URL: https://link.aps.org/doi/10.1103/RevModPhys.84.1307.
- [17] D. Perkins. *Introduction to High Energy Physics*. Cambridge University Press, 2000.
- [18] Nicola Cabibbo. “Unitary Symmetry and Leptonic Decays”. In: *Phys. Rev. Lett.* 10 (12 June 1963), pp. 531–533. DOI: 10.1103/PhysRevLett.10.531. URL: https://link.aps.org/doi/10.1103/PhysRevLett.10.531.
- [19] Makoto Kobayashi and Toshihide Maskawa. “CP-Violation in the Renormalizable Theory of Weak Interaction”. In: *Progress of Theoretical Physics* 49.2 (1973), pp. 652–657. DOI: 10.1143/PTP.49.652. eprint: /oup/backfile/content_public/journal/ptp/49/2/10.1143/ptp.49.652/2/49-2-652.pdf. URL: +%20http://dx.doi.org/10.1143/PTP.49.652.
- [20] C. S. Wu et al. “Experimental Test of Parity Conservation in Beta Decay”. In: *Phys. Rev.* 105 (4 Feb. 1957), pp. 1413–1415. DOI: 10.1103/PhysRev.105.1413. URL: https://link.aps.org/doi/10.1103/PhysRev.105.1413.
- [21] R. P. Feynman and M. Gell-Mann. “Theory of the Fermi Interaction”. In: *Phys. Rev.* 109 (1 Jan. 1958), pp. 193–198. DOI: 10.1103/PhysRev.109.193. URL: https://link.aps.org/doi/10.1103/PhysRev.109.193.
- [22] Raymond Davis, Don S. Harmer, and Kenneth C. Hoffman. “Search for Neutrinos from the Sun”. In: *Phys. Rev. Lett.* 20 (21 May 1968), pp. 1205–1209. DOI: 10.1103/PhysRevLett.20.1205. URL: https://link.aps.org/doi/10.1103/PhysRevLett.20.1205.

- [23] John N. Bahcall. “Solar neutrinos: Where we are, where we are going”. In: *Nuclear Physics B - Proceedings Supplements* 48.1 (1996). Proceedings of the Fourth International Workshop on Theoretical and Phenomenological Aspects of Underground Physics, pp. 309–316. ISSN: 0920-5632. DOI: [https://doi.org/10.1016/0920-5632\(96\)00266-6](https://doi.org/10.1016/0920-5632(96)00266-6). URL: <http://www.sciencedirect.com/science/article/pii/0920563296002666>.
- [24] Florian Kaether et al. “Reanalysis of the GALLEX solar neutrino flux and source experiments”. In: 685 (Feb. 2010), pp. 47–54.
- [25] Q. R. Ahmad et al. “Measurement of the Rate of $\nu_e + d \rightarrow p + p + e^-$ Interactions Produced by 8B Solar Neutrinos at the Sudbury Neutrino Observatory”. In: *Phys. Rev. Lett.* 87 (7 July 2001), p. 071301. DOI: [10.1103/PhysRevLett.87.071301](https://doi.org/10.1103/PhysRevLett.87.071301). URL: <https://link.aps.org/doi/10.1103/PhysRevLett.87.071301>.
- [26] S. Abe et al. “Precision Measurement of Neutrino Oscillation Parameters with KamLAND”. In: *Phys. Rev. Lett.* 100 (22 June 2008), p. 221803. DOI: [10.1103/PhysRevLett.100.221803](https://doi.org/10.1103/PhysRevLett.100.221803). URL: <https://link.aps.org/doi/10.1103/PhysRevLett.100.221803>.
- [27] A. Yu. Smirnov. “The MSW effect and solar neutrinos”. In: *Neutrino telescopes. Proceedings, 10th International Workshop, Venice, Italy, March 11-14, 2003. Vol. 1+2*. 2003, pp. 23–43. arXiv: [hep-ph/0305106](https://arxiv.org/abs/hep-ph/0305106) [hep-ph].
- [28] Marco G. Giammarchi. “Solar and geoneutrino physics with Borexino”. In: *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment* 742 (2014). 4th Roma International Conference on Astroparticle Physics, pp. 250–253. ISSN: 0168-9002. DOI: <https://doi.org/10.1016/j.nima.2013.11.088>. URL: <http://www.sciencedirect.com/science/article/pii/S0168900213016537>.
- [29] N. Agafonova et al. “Discovery of τ Neutrino Appearance in the CNGS Neutrino Beam with the OPERA Experiment”. In: *Phys. Rev. Lett.* 115 (12 Sept. 2015), p. 121802. DOI: [10.1103/PhysRevLett.115.121802](https://doi.org/10.1103/PhysRevLett.115.121802). URL: <https://link.aps.org/doi/10.1103/PhysRevLett.115.121802>.
- [30] Andrew John Lowe. “Neutrino Physics and The Solar Neutrino Problem”. In: (2009). arXiv: [0907.3658](https://arxiv.org/abs/0907.3658) [hep-ex].
- [31] M. G. Aartsen et al. “Measurement of Atmospheric Neutrino Oscillations at 6–56 GeV with IceCube DeepCore”. In: *Phys. Rev. Lett.* 120.7 (2018), p. 071801. DOI: [10.1103/PhysRevLett.120.071801](https://doi.org/10.1103/PhysRevLett.120.071801). arXiv: [1707.07081](https://arxiv.org/abs/1707.07081) [hep-ex].
- [32] F. P. An et al. “New Measurement of Antineutrino Oscillation with the Full Detector Configuration at Daya Bay”. In: *Phys. Rev. Lett.* 115 (11 Sept. 2015), p. 111802. DOI: [10.1103/PhysRevLett.115.111802](https://doi.org/10.1103/PhysRevLett.115.111802). URL: <https://link.aps.org/doi/10.1103/PhysRevLett.115.111802>.

- [33] J. H. Choi et al. “Observation of Energy and Baseline Dependent Reactor Antineutrino Disappearance in the RENO Experiment”. In: *Phys. Rev. Lett.* 116 (21 May 2016), p. 211801. DOI: 10.1103/PhysRevLett.116.211801. URL: <https://link.aps.org/doi/10.1103/PhysRevLett.116.211801>.
- [34] Y. Abe et al. “Indication of Reactor $\bar{\nu}_e$ Disappearance in the Double Chooz Experiment”. In: *Phys. Rev. Lett.* 108 (13 Mar. 2012), p. 131801. DOI: 10.1103/PhysRevLett.108.131801. URL: <https://link.aps.org/doi/10.1103/PhysRevLett.108.131801>.
- [35] M. H. Ahn et al. “Measurement of neutrino oscillation by the K2K experiment”. In: *Phys. Rev. D* 74 (7 Oct. 2006), p. 072003. DOI: 10.1103/PhysRevD.74.072003. URL: <https://link.aps.org/doi/10.1103/PhysRevD.74.072003>.
- [36] P. Adamson et al. “Measurement of Neutrino Oscillations with the MINOS Detectors in the NuMI Beam”. In: *Phys. Rev. Lett.* 101 (13 Sept. 2008), p. 131802. DOI: 10.1103/PhysRevLett.101.131802. URL: <https://link.aps.org/doi/10.1103/PhysRevLett.101.131802>.
- [37] P. Adamson et al. “Electron neutrino and antineutrino appearance in the full MINOS data sample”. In: *Phys. Rev. Lett.* 110.17 (2013), p. 171801. DOI: 10.1103/PhysRevLett.110.171801. arXiv: 1301.4581 [hep-ex].
- [38] A. Cervera et al. “Golden measurements at a neutrino factory”. In: *Nucl. Phys.* B579 (2000). [Erratum: *Nucl. Phys.*B593,731(2001)], pp. 17–55. DOI: 10.1016/S0550-3213(00)00606-4, 10.1016/S0550-3213(00)00221-2. arXiv: hep-ph/0002108 [hep-ph].
- [39] S. T. Petcov. “Leptonic CP violation and leptogenesis”. In: *Int. J. Mod. Phys.* A29 (2014). [179(2014)], p. 1430028. DOI: 10.1142/S0217751X14300282, 10.1142/9789814590112_0009. arXiv: 1405.2263 [hep-ph].
- [40] K. Abe et al. “Combined Analysis of Neutrino and Antineutrino Oscillations at T2K”. In: *Phys. Rev. Lett.* 118.15 (2017), p. 151801. DOI: 10.1103/PhysRevLett.118.151801. arXiv: 1701.00432 [hep-ex].
- [41] J. Huang. “NOvA FD Cosmic MC Validation”. NOvA Internal Note 11331. 2014.
- [42] S. Mufson et al. “Liquid scintillator production for the NOvA experiment”. In: *Nucl. Instrum. Meth.* A799 (2015), pp. 1–9. DOI: 10.1016/j.nima.2015.07.026. arXiv: 1504.04035 [physics.ins-det].
- [43] A. Aurisano. “Tech Note: 2017 Light Model”. NOvA Internal Note 23228. 2017.
- [44] NOvA Collaboration. “FERMILAB-DESIGN-2007-01”. 2007.

- [45] R. L. Talaga et al. “PVC Extrusion Development and Production for the NOvA Neutrino Experiment”. In: *Nucl. Instrum. Meth.* A861 (2017), pp. 77–89. DOI: 10.1016/j.nima.2017.03.004. arXiv: 1601.00908 [physics.ins-det].
- [46] O. Oshinowo et. al. “Near Detector Muon Catcher Steel Measurements”. NOvA Internal Note 16937. 2017.
- [47] E. Song. “Technical note for Fast Monopole Trigger with a Neural Network”. NOvA Internal Note 19323. 2017.
- [48] A. Antoshkin et. al. “Slow Monopole Technical Note”. NOvA Internal Note 22610. 2017.
- [49] B. Wang. “Muon-neutrino electron elastic scattering and a search for the muon-neutrino magnetic moment in the NOvA near detector”. PhD thesis. Southern Methodist University, 2017.
- [50] J. Huang. “Neutron-Antineutron Oscillation: An Exploratory DDT Trigger”. NOvA Internal Note 26343. 2018.
- [51] C. Backhouse. “Light Dark Matter in the Near Detector”. NOvA Internal Note 15118. 2016.
- [52] P. Adamson et al. “The NuMI Neutrino Beam”. In: *Nucl. Instrum. Meth.* A806 (2016), pp. 279–306. DOI: 10.1016/j.nima.2015.08.063. arXiv: 1507.06690 [physics.acc-ph].
- [53] Bruce C. Brown et al. “The Fermilab Main Injector: high intensity operation and beam loss control”. In: *Phys. Rev. ST Accel. Beams* 16.7 (2013), p. 071001. DOI: 10.1103/PhysRevSTAB.16.071001. arXiv: 1307.2934 [physics.acc-ph].
- [54] P. Adamson. “Reuse Recycler: High Intensity Proton Stacking at Fermilab”. In: *Proceedings, 57th ICFA Advanced Beam Dynamics Workshop on High-Intensity and High-Brightness Hadron Beams (HB2016): Malmö, Sweden, July 3-8, 2016*. 2016, THAM1X01. DOI: 10.18429/JACoW-HB2016-THAM1X01. URL: <http://lss.fnal.gov/archive/2016/conf/fermilab-conf-16-277-ad.pdf>.
- [55] D. Kalra. “Optimization of neutrino flux using NOvA target design”. NOvA Internal Note 22328. 2017.
- [56] R. Schroeter. “Wrong sign neutrino contamination and beam nue content for FHC/RHC beam configurations - Fluka 2011”. NOvA Internal Note 13097. 2015.
- [57] C. Backhouse. “Multipoint readout and fine timing”. NOvA Internal Note 10821. 2014.
- [58] A. Norman. “NOvA DAQ Data Formats - Definitions”. NOvA Internal Note 4390. 2009.

- [59] E. Niner. “TDU/DCM delay calibration tech note”. NOvA Internal Note 11568. 2014.
- [60] M. Judah. “DDT Operations Update and Michel Trigger Status”. NOvA Internal Note 17003. 2017.
- [61] L. Mualum. “Per-pixel Threshold Setting”. NOvA Internal Note 5606. 2011.
- [62] J. Liu. “Quantifying Crosstalk in Avalanche Photodiodes Used in the NOvA Detector”. NOvA Internal Note 7992. 2012.
- [63] L. Mualem. “Caltech APD Sag Measurements”. NOvA Internal Note 13008. 2015.
- [64] A. Nikolay. “Studies of the NOvA APD Sag vs different Bypass Capacitance”. NOvA Internal Note 13192. 2015.
- [65] GDML Engineers. “GDML Manual”. 2016. URL: <http://gdml.web.cern.ch/GDML/doc/GDMLmanual.pdf>.
- [66] M. Strait. “Tech Note: 2016 geometry and composition changes”. NOvA Internal Note 23132. 2017.
- [67] 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.
- [68] L. Aliaga et al. “Neutrino Flux Predictions for the NuMI Beam”. In: *Phys. Rev. D* 94.9 (2016). [Addendum: *Phys. Rev. D* 95, no. 3, 039903 (2017)], p. 092005. DOI: 10.1103/PhysRevD.94.092005, 10.1103/PhysRevD.95.039903. arXiv: 1607.00704 [hep-ex].
- [69] J. M. et. al. Paley. “Measurement of charged pion production yields off the NuMI target”. In: *Phys. Rev. D* 90 (3 Aug. 2014), p. 032001. DOI: 10.1103/PhysRevD.90.032001. URL: <https://link.aps.org/doi/10.1103/PhysRevD.90.032001>.
- [70] C. Alt et al. “Inclusive production of charged pions in p+C collisions at 158-GeV/c beam momentum”. In: *Eur. Phys. J. C* 49 (2007), pp. 897–917. DOI: 10.1140/epjc/s10052-006-0165-7. arXiv: hep-ex/0606028 [hep-ex].
- [71] N Abgrall et al. “Measurements of Cross Sections and Charged Pion Spectra in Proton-Carbon Interactions at 31 GeV/c”. In: *Phys. Rev. C* 84 (2011), p. 034604. DOI: 10.1103/PhysRevC.84.034604. arXiv: 1102.0983 [hep-ex].
- [72] C. et. al. Andreopoulos. “The GENIE Neutrino Monte Carlo Generator: Physics and User Manual”. In: (2015). arXiv: 1510.05494 [hep-ph].

- [73] C.H. Llewellyn Smith. “Neutrino reactions at accelerator energies”. In: *Physics Reports* 3.5 (1972), pp. 261–379. ISSN: 0370-1573. DOI: [https://doi.org/10.1016/0370-1573\(72\)90010-5](https://doi.org/10.1016/0370-1573(72)90010-5). URL: <http://www.sciencedirect.com/science/article/pii/0370157372900105>.
- [74] GEANT4 Collaboration. “GEANT4 Physics Reference Manual”. 2017. URL: http://geant4-userdoc.web.cern.ch/geant4-userdoc/UsersGuides/PhysicsListGuide/html/reference_PL/index.html.
- [75] Hugo W. Bertini. “Low-Energy Intranuclear Cascade Calculation”. In: *Phys. Rev.* 131 (4 Aug. 1963), pp. 1801–1821. DOI: [10.1103/PhysRev.131.1801](https://doi.org/10.1103/PhysRev.131.1801). URL: <https://link.aps.org/doi/10.1103/PhysRev.131.1801>.
- [76] P. Chernekov. “Visible Luminescence of the Pure Fluids Induced by Gamma Rays”. In: *Dokl. Akad. Nauk SSSR* 2 (1934), pp. 451–454.
- [77] I. M. Frank and I E Tamm. “Coherent Visible Radiation of Fast Electrons Passing Through Matter”. In: *Dokl. Akad. Nauk SSSR* 14 (1937), pp. 10–113.
- [78] J B Birks. “Scintillations from Organic Crystals: Specific Fluorescence and Relative Response to Different Radiations”. In: *Proceedings of the Physical Society. Section A* 64.10 (1951), p. 874. URL: <http://stacks.iop.org/0370-1298/64/i=10/a=303>.
- [79] C N Chou. “The Nature of the Saturation Effect of Fluorescent Scintillators”. In: *Physical Review* 87 (1952), p. 904.
- [80] William G. Jacoby. “LOESS: a nonparametric, graphical tool for depicting relationships between variables”. In: 19 (Dec. 2000), pp. 577–613.
- [81] T. Alion. “Third Analysis Calorimetric Energy Scale in the NOvA Detectors”. NOvA Internal Note 23372. 2017.
- [82] C. Patrignani et al. “Review of Particle Physics”. In: *Chin. Phys.* C40.10 (2016), p. 100001. DOI: [10.1088/1674-1137/40/10/100001](https://doi.org/10.1088/1674-1137/40/10/100001).
- [83] D. Pershey. “Technical Note for the CCPi02017 Analysis”. NOvA Internal Note 18768. 2017.
- [84] P.V.C. Hough. *Method and means for recognizing complex patterns*. US Patent 3069654. 1962.
- [85] J.C. Dunn. “A Fuzzy Relative of the ISODATA Process and Its Use in Detecting Compact Well-Separated Clusters”. In: *Journal of Cybernetics* 3257 (1973).

- [86] M. Ester et al. “A density-based algorithm for discovering clusters in large spatial databases with noise”. In: *KDD'96 Proceedings of the Second International Conference on Knowledge Discovery and Data Mining*. 1996, pp. 226–231.
- [87] M. Baird. “Technote for Slicer4D”. NOvA Internal Note 9195. 2013.
- [88] Andraž Mehle, Boštjan Likar, and Dejan Tomaževič. “In-line recognition of agglomerated pharmaceutical pellets with density-based clustering and convolutional neural network”. In: *IPSJ Transactions on Computer Vision and Applications* 9.1 (Mar. 2017), p. 7. DOI: 10.1186/s41074-017-0019-2. URL: <https://doi.org/10.1186/s41074-017-0019-2>.
- [89] M. Baird. “Technote for the Multi-Hough Transform”. NOvA Internal Note 8241. 2012.
- [90] M. Baird. “Vertex Reconstruction Based on Elastic Arms”. NOvA Internal Note 7530. 2012.
- [91] R. Krishnapuram and J.M. Keller. “A Probabilistic Approach to Clustering”. In: *IEEE* 98110 (1993).
- [92] R.E. Kalman. “A New Approach to Linear Filtering and Prediction Problems”. In: *J. Basic Engineering* 82.1 (1960), pp. 35–45.
- [93] N. Raddatz. “KalmanTrack Technical Note”. NOvA Internal Note 13545. 2015.
- [94] N. Raddatz. “ReMId Technical Note”. NOvA Internal Note 11206. 2014.
- [95] E. Niner J. Bian and K. Sachdev. “Likelihood Based Nue Identifier (LID)”. NOvA Internal Note 13590. 2015.
- [96] A. Aurisano et al. “A convolutional neural network neutrino event classifier”. In: *Journal of Instrumentation* 11.09 (2016), P09001. URL: <http://stacks.iop.org/1748-0221/11/i=09/a=P09001>.
- [97] Christian Szegedy et al. “Going Deeper with Convolutions”. In: *CoRR* abs/1409.4842 (2014). arXiv: 1409.4842. URL: <http://arxiv.org/abs/1409.4842>.
- [98] D. Pershey. “MEFinder Tech Note”. NOvA Internal Note 14789. 2015.
- [99] D. Pershey. “APD Dead Time”. NOvA Internal Note 12295. 2014.
- [100] M. Strait. “Tech Note on Using Neutrons to Find NC and RHC Numu Contamination”. NOvA Internal Note 22955. 2017.
- [101] A. Rodriguez and A. Laio. “Clustering by fast search and find of density peaks”. In: *Science* 344.6191 (2014), pp. 1492–1496. ISSN: 0036-8075. DOI: 10.1126/science.1242072. eprint: <http://science.sciencemag.org/content/344/6191/1492.full.pdf>. URL: <http://science.sciencemag.org/content/344/6191/1492>.

- [102] R. C. Prim. “Shortest Connection Networks and Some Generalizations”. In: *Bell System Technical Journal* 36.6 (Nov. 1957), pp. 1389–1401.
- [103] ANL Collaboration. “Study of single-pion production by weak charged currents in low-energy νd interactions”. In: *Phys. Rev. D* 25 (5 Mar. 1982), pp. 1161–1173. DOI: 10.1103/PhysRevD.25.1161. URL: <https://link.aps.org/doi/10.1103/PhysRevD.25.1161>.
- [104] BNL Collaboration. “Charged-current exclusive pion production in neutrino-deuterium interactions”. In: *Phys. Rev. D* 34 (9 Nov. 1986), pp. 2554–2565. DOI: 10.1103/PhysRevD.34.2554. URL: <https://link.aps.org/doi/10.1103/PhysRevD.34.2554>.
- [105] A. A. Aguilar-Arevalo et al. “Measurement of ν_μ -induced charged-current neutral pion production cross sections on mineral oil at $E_\nu \in 0.5 - 2.0$ GeV”. In: *Phys. Rev. D* 83 (2011), p. 052009. DOI: 10.1103/PhysRevD.83.052009. arXiv: 1010.3264 [hep-ex].
- [106] C. Mariani. “Neutral pion cross section measurement at K2K”. In: *AIP Conf. Proc.* 967 (2007), pp. 174–178. DOI: 10.1063/1.2834471.
- [107] O. Altinok et al. “Measurement of ν_μ charged-current single π^0 production on hydrocarbon in the few-GeV region using MINERvA”. In: *Phys. Rev. D* 96.7 (2017), p. 072003. DOI: 10.1103/PhysRevD.96.072003. arXiv: 1708.03723 [hep-ex].
- [108] A.N. Tikhonov and V.Y. Arsenin. “Solution of Ill-Posed Problems”. In: *Winston and Sons* (1977).
- [109] G. D’Agostini. “A Multidimensional unfolding method based on Bayes’ theorem”. In: *Nucl. Instrum. Meth.* A362 (1995), pp. 487–498. DOI: 10.1016/0168-9002(95)00274-X.
- [110] Andreas Hocker and Vakhtang Kartvelishvili. “SVD approach to data unfolding”. In: *Nucl. Instrum. Meth.* A372 (1996), pp. 469–481. DOI: 10.1016/0168-9002(95)01478-0. arXiv: hep-ph/9509307 [hep-ph].
- [111] K. Sachdev. “Target count for cross-sections”. NOvA Internal Note 15787. 2016.
- [112] Costas Andreopoulos et al. “The GENIE Neutrino Monte Carlo Generator: Physics and User Manual”. In: (2015). arXiv: 1510.05494 [hep-ph].
- [113] D. Ashery et al. “True Absorption and Scattering of Pions on Nuclei”. In: *Phys. Rev. C* 23 (1981), pp. 2173–2185. DOI: 10.1103/PhysRevC.23.2173.
- [114] E. S. Pinzon Guerra et al. “Measurement of σ_{ABS} and σ_{CX} of π^+ on carbon by the Dual Use Experiment at TRIUMF (DUET)”. In: *Phys. Rev. C* 95 (4 Apr. 2017), p. 045203. DOI: 10.1103/PhysRevC.95.045203. URL: <https://link.aps.org/doi/10.1103/PhysRevC.95.045203>.

- [115] D. Pershey. “Pion Scattering Cross Sections”. NOvA Internal Note 19281. 2017.
- [116] D. Pershey. “Charged Pion CEx Systematic for CCPi0”. NOvA Internal Note 19470. 2017.
- [117] D. Pershey. “Birks-Chou Parameter Selection”. NOvA Internal Note 13158. 2015.
- [118] D. Pershey. “SA ND ME Calib Checks”. NOvA Internal Note 14959. 2016.
- [119] G. Davies. “ π^0 Invariant Mass in SA”. NOvA Internal Note 15005. 2016.
- [120] S. Lein. “DE/DX Along Tracks”. NOvA Internal Note 15028. 2016.
- [121] N. Raddatz. “Fiducial Mass Systematic”. NOvA Internal Note 13237. 2015.
- [122] D. Pershey. “Data Stability for the Second ν_e Appearance Analysis”. NOvA Internal Note 15388. 2016.
- [123] G. Davies. “Pi0 Invariant Mass in SA Files”. NOvA Internal Note 15005. 2016.
- [124] B. Howard. “Exposure Accounting for the 2018 RHC Electron Antineutrino Appearance Analysis”. NOvA Internal Note 28346. 2018.
- [125] T.K. Warburton. “Near Detector Data/MC Comparisons ν_μ 2018 Analysis”. NOvA Internal Note 27824. 2018.
- [126] D. Pershey. “SA ND Cuts”. NOvA Internal Note 15059. 2016.
- [127] E Cataño-Mur. “BEN Decomposition Technote 2017”. NOvA Internal Note 23174. 2017.
- [128] J. Lozier. “Modular Extrap Technical Note”. NOvA Internal Note 12563. 2014.
- [129] J. Hartnell and A. Radovic. “Summary of the 2018 FHC+RHC ν_μ Disappearance Analysis”. NOvA Internal Note 26702. 2018.
- [130] N. Nayak B. Howard and E. Niner. “Wrong Sign Identification Techniques for the nue Appearance Analysis”. NOvA Internal Note 27248. 2018.
- [131] Karl Pearson F.R.S. “LIII. On lines and planes of closest fit to systems of points in space”. In: *The London, Edinburgh, and Dublin Philosophical Magazine and Journal of Science* 2.11 (1901), pp. 559–572. doi: 10.1080/14786440109462720. eprint: <https://doi.org/10.1080/14786440109462720>. URL: <https://doi.org/10.1080/14786440109462720>.
- [132] K Bays et al. “NOvA 2018 cross-section tune technote”. NOvA Internal Note 27755. 2017.

- [133] Aaron S. Meyer et al. “Deuterium target data for precision neutrino-nucleus cross sections”. In: *Phys. Rev. D* 93 (11 June 2016), p. 113015. DOI: 10.1103/PhysRevD.93.113015. URL: <https://link.aps.org/doi/10.1103/PhysRevD.93.113015>.
- [134] J Nieves et al. “Quasi-elastic scattering, RPA, 2p2h and neutrino–energy reconstruction”. In: *15th International Workshop on Neutrino Factories, Super Beams and Beta Beams (NuFact2013) Beijing, China, August 19-24, 2013*. 2013. arXiv: 1310.7091 [hep-ph]. URL: <https://inspirehep.net/record/1262343/files/arXiv:1310.7091.pdf>.
- [135] P. A. Rodrigues et al. “Identification of Nuclear Effects in Neutrino-Carbon Interactions at Low Three-Momentum Transfer”. In: *Phys. Rev. Lett.* 116 (7 Feb. 2016), p. 071802. DOI: 10.1103/PhysRevLett.116.071802. URL: <https://link.aps.org/doi/10.1103/PhysRevLett.116.071802>.
- [136] M. Groh. “GENIE Systematics for the 2018 Oscillation Analysis”. NOvA Internal Note 27914. 2018.
- [137] T. Alion. “FHC and RHC ND Overlay Normalization Systematic”. NOvA Internal Note 27230. 2018.
- [138] Gary J. Feldman and Robert D. Cousins. “A 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. arXiv: physics/9711021 [physics.data-an].
- [139] G. Feldman. “Notes on the Inclusion of Nuisance Parameters in the Unified Approach”. NOvA Internal Note 15884. 2016.
- [140] “P values maximized over a confidence set for the nuisance parameter”. English (US). In: *Journal of the American Statistical Association* 89.427 (), pp. 1012–1016. ISSN: 0162-1459. DOI: 10.1080/01621459.1994.10476836.
- [141] D. Pershey. “Osc Parameter Seeding Check for IH Rejection - 9e20 FHC Result”. NOvA Internal Note 28865. 2018.
- [142] A. Kolmogorov. “Sulla determinazione empirica di una legge di distribuzione”. In: *G. Ist. Ital. Attuari* 4 (1933), pp. 83–91.
- [143] Paul D. Carr and Paul H. Frampton. “Group Theoretic Bases for Tribimaximal Mixing”. In: (2007). arXiv: hep-ph/0701034 [hep-ph].
- [144] Biswajoy Brahmachari and Amitava Raychaudhuri. “Perturbative generation of θ_{13} from tribimaximal neutrino mixing”. In: *Phys. Rev. D* 86 (2012), p. 051302. DOI: 10.1103/PhysRevD.86.051302. arXiv: 1204.5619 [hep-ph].
- [145] et al. H. L. Anderson. “Search for the Electronic Decay of the Positive Pion”. In: *Nuovo Cim* 6.6 (1957).

- [146] T. Fazzini et al. “Electron Decay of the Pion”. In: *Phys. Rev. Lett.* 1 (7 Oct. 1958), pp. 247–249. DOI: 10.1103/PhysRevLett.1.247. URL: <https://link.aps.org/doi/10.1103/PhysRevLett.1.247>.
- [147] J. C. Hardy and I. S. Towner. “Standard-Model Tests with Superaligned β -Decay: An Important Application of Very Precise Mass Measurements”. In: *Hyperfine Interactions* 132.1 (Jan. 2001), pp. 115–126. DOI: 10.1023/A:1011999930767. URL: <https://doi.org/10.1023/A:1011999930767>.
- [148] Stephen L. Adler. “Tests of the Conserved Vector Current and Partially Conserved Axial-Vector Current Hypotheses in High-Energy Neutrino Reactions”. In: *Phys. Rev.* 135 (4B Aug. 1964), B963–B966. DOI: 10.1103/PhysRev.135.B963. URL: <https://link.aps.org/doi/10.1103/PhysRev.135.B963>.
- [149] M. L. Goldberger and S. B. Treiman. “Decay of the Pi Meson”. In: *Phys. Rev.* 110 (5 June 1958), pp. 1178–1184. DOI: 10.1103/PhysRev.110.1178. URL: <https://link.aps.org/doi/10.1103/PhysRev.110.1178>.
- [150] Kevin Scott McFarland. “Neutrino Interactions”. In: *Neutrinos in particle physics, astrophysics and cosmology. Proceedings, 61st Scottish Universities Summer School in Physics, SUSSP61, St. Andrews, UK, August 8-23, 2006*. 2008, pp. 65–90. arXiv: 0804.3899 [hep-ex]. URL: <https://inspirehep.net/record/784224/files/arXiv:0804.3899.pdf>.
- [151] J. Park et al. “Measurement of Neutrino Flux from Neutrino-Electron Elastic Scattering”. In: *Phys. Rev. D* 93.11 (2016), p. 112007. DOI: 10.1103/PhysRevD.93.112007. arXiv: 1512.07699 [physics.ins-det].
- [152] Melanie Day and Kevin S. McFarland. “Differences in Quasi-Elastic Cross-Sections of Muon and Electron Neutrinos”. In: *Phys. Rev. D* 86 (2012), p. 053003. DOI: 10.1103/PhysRevD.86.053003. arXiv: 1206.6745 [hep-ph].
- [153] Veronique Bernard, Latifa Elouadrhiri, and Ulf-G. Meissner. “Axial structure of the nucleon: Topical Review”. In: *J. Phys.* G28 (2002), R1–R35. DOI: 10.1088/0954-3899/28/1/201. arXiv: hep-ph/0107088 [hep-ph].
- [154] M. Breidenbach et al. “Observed Behavior of Highly Inelastic Electron-Proton Scattering”. In: *Phys. Rev. Lett.* 23 (16 Oct. 1969), pp. 935–939. DOI: 10.1103/PhysRevLett.23.935. URL: <https://link.aps.org/doi/10.1103/PhysRevLett.23.935>.
- [155] C. G. Callan and David J. Gross. “High-Energy Electroproduction and the Constitution of the Electric Current”. In: *Phys. Rev. Lett.* 22 (4 Jan. 1969), pp. 156–159. DOI: 10.1103/PhysRevLett.22.156. URL: <https://link.aps.org/doi/10.1103/PhysRevLett.22.156>.

- [156] Dieter Rein and Lalit M Sehgal. “Neutrino-excitation of baryon resonances and single pion production”. In: *Annals of Physics* 133.1 (1981), pp. 79–153. ISSN: 0003-4916. DOI: [https://doi.org/10.1016/0003-4916\(81\)90242-6](https://doi.org/10.1016/0003-4916(81)90242-6). URL: <http://www.sciencedirect.com/science/article/pii/0003491681902426>.
- [157] M. Martini et al. “Unified approach for nucleon knock-out and coherent and incoherent pion production in neutrino interactions with nuclei”. In: *Phys. Rev. C* 80 (6 Dec. 2009), p. 065501. DOI: [10.1103/PhysRevC.80.065501](https://doi.org/10.1103/PhysRevC.80.065501). URL: <https://link.aps.org/doi/10.1103/PhysRevC.80.065501>.
- [158] A. A. Aguilar-Arevalo et al. “First measurement of the muon neutrino charged current quasielastic double differential cross section”. In: *Phys. Rev. D* 81 (9 May 2010), p. 092005. DOI: [10.1103/PhysRevD.81.092005](https://doi.org/10.1103/PhysRevD.81.092005). URL: <https://link.aps.org/doi/10.1103/PhysRevD.81.092005>.
- [159] J. Nieves, I. Ruiz Simo, and M. J. Vicente Vacas. “Inclusive charged-current neutrino-nucleus reactions”. In: *Phys. Rev. C* 83 (4 Apr. 2011), p. 045501. DOI: [10.1103/PhysRevC.83.045501](https://doi.org/10.1103/PhysRevC.83.045501). URL: <https://link.aps.org/doi/10.1103/PhysRevC.83.045501>.
- [160] Dieter Rein and Lalit M. Sehgal. “Coherent π^0 production in neutrino reactions”. In: *Nuclear Physics B* 223.1 (1983), pp. 29–44. ISSN: 0550-3213. DOI: [https://doi.org/10.1016/0550-3213\(83\)90090-1](https://doi.org/10.1016/0550-3213(83)90090-1). URL: <http://www.sciencedirect.com/science/article/pii/0550321383900901>.

Appendix A

NEUTRINO-NUCLEUS SCATTERING PHENOMENOLOGY

At energies relevant for NOvA, $E_\nu \approx 2 \text{ GeV} \ll M_W^2$, the W propagator can be simplified so that the weak interaction Lagrangian can be written as

$$L = i \frac{G_F^2}{2} J_\alpha^\mu (J_\beta)_\mu. \quad (\text{A.1})$$

Neutrino-nucleus scattering is the dominant scattering process for NOvA. This Lagrangian is only suitable for scattering of bare fermions. As a nucleus is made of constituent bound hadronic states, the full story is more complicated than this Lagrangian suggests. Theoretical and experimental descriptions of processes involving correlations between nucleons are just beginning.

From here, we will only consider interactions with $E_\nu \sim 1 \text{ GeV} \ll M_W^2$. In this approximation, the W propagator is simplified so that the interaction Lagrangian from Eqn. 1.34 can be written as

$$L = i \frac{G_F^2}{2} J_\alpha^\mu (J_\beta)_\mu, \quad (\text{A.2})$$

where G_F^2 is the familiar Fermi constant from the four-body theory of weak interactions. The universality of G_F is hypothesized with the Conservation of the Vector Current (CVC) theorem[21]. This was developed as a consequence of the correspondence between the governing coupling factors in both β and μ decay. CVC was initially on weak footing, as the decay $\pi^+ \rightarrow e^+ \nu_e$ had not been experimentally observed with branching ratios limited below 1×10^{-5} [145] while the theory predicted a value of 1.3×10^{-4} . The experimental detection of this decay channel was thus prioritized to test CVC and was observed[146] within a year of the hypothesis. Ever since, the observed branching ratio has agreed with the predicted value.

This description refers specifically to the action on bare fermions. Quarks are only observed in bound hadronic states; at distances greater than $\approx 1 \text{ fm}$, the confinement term in the QCD potential makes free quarks energetically forbidden. In a bound hadronic state, renormalization can amend the $V-A$ structure of the weak interaction. To be explicit, an effective interaction current connecting a proton and neutron takes the form

$$J_{\text{eff}}^\mu = \bar{p} \gamma^\mu (c_V - c_A \gamma^5) n, \quad (\text{A.3})$$

where c_V and c_A can depend on the Lorentz-invariant four-momentum transfer. Distortions to these coefficients with Q^2 represent a weak charge-screening in the scattering process. Experiment shows $c_V = 1$ with $c_A = 1.24$ at $Q^2 \approx 0$ from β decay[147]. The observation of $c_V = 1$ at $Q^2 = 0$ is a direct test of the CVC hypothesis. Additionally, it's observed that $|c_A - 1| \ll 1$, related to Adler's Partial Conservation of the Axial Current (PCAC) theorem. This describes the axial component as only depending on the divergence of the pion field as described below.

Adler's theorem gives a relationship between the hadronic current of the neutrino scattering process, $\nu + \alpha \rightarrow l^- + \beta$, and the purely hadronic process, $\pi^+ + \alpha \rightarrow \beta$ [148]. This theorem is valid in cases where the lepton and neutrino momenta are parallel. This ensures that $Q^2 = 0$ and CVC applies. A proof is given by Adler, which is only outlined below, highlighting the physical reasoning. Starting with the general matrix transition element that shows the leptonic and hadronic currents factored,

$$\mathcal{M} = \frac{G_F}{\sqrt{2}} \bar{l} \gamma^\mu (1 - \gamma^5) \nu \langle \beta | V^\mu + A^\mu | \alpha \rangle. \quad (\text{A.4})$$

Adler showed that in the parallel configuration, the spin-averaged squared matrix element is

$$\langle |\mathcal{M}^2| \rangle = 2G_F^2 \frac{E_\nu E_l}{q_0^2} |\langle \beta | \partial_\mu A^\mu | \alpha \rangle|^2, \quad (\text{A.5})$$

showing the cross section only depends on the divergence of the axial vector term. With the Goldberger-Treiman relation[149], the last factor can be related to the pion scattering process and the form factor, f_π , determined from π^\pm decay:

$$|\langle \beta | \partial_\mu A^\mu | \alpha \rangle|^2 = f_\pi^2 |\mathcal{M}(\pi + \alpha \rightarrow \beta)|^2, \quad (\text{A.6})$$

which, together with Eqn. A.5, gives

$$\langle |\mathcal{M}^2| \rangle = 2G_F^2 \frac{E_\nu E_l}{q_0^2} f_\pi^2 |\mathcal{M}(\pi + \alpha \rightarrow \beta)|^2. \quad (\text{A.7})$$

In other words, the neutrino interaction can be factored so that the only weak process involved is a π^\pm decay. This introduces the $\pi + \alpha \rightarrow \beta$ process, but such processes are more precisely known than the neutrino interaction. This equivalence is illustrated in Figure A.1. This result is the basis for neutrino interaction models which approximate the running of the π form factor as a dipole expansion around the $Q^2 = 0$ prediction from PCAC.

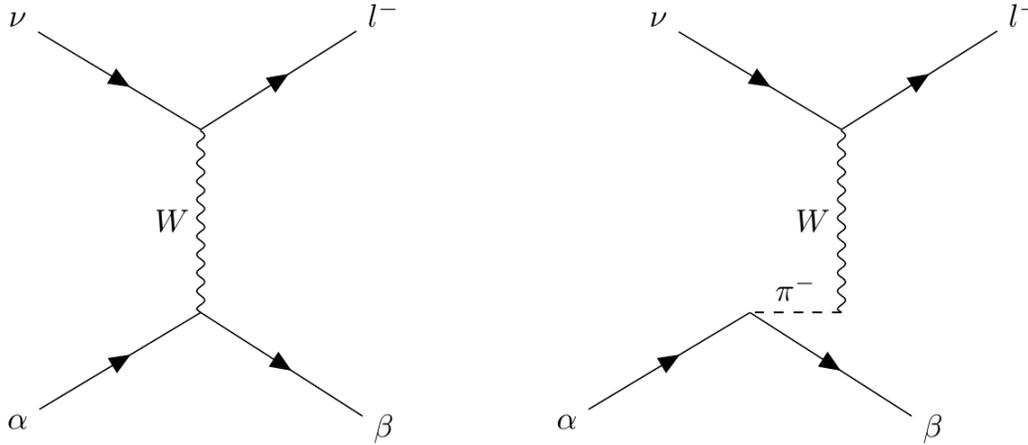


Figure A.1: The equivalency predicted from Adler's theorem, at $Q^2 = 0$. The neutrino interaction on the left can be thought of as the product of the strong interaction $\alpha \rightarrow \pi\beta$ and the divergence of the π shown on the right.

Leptonic Cross Sections

Leptonic cross sections involving neutrinos are readily calculated. The matrix element for such a process is given by

$$L = i \frac{G_F^2}{2} g_{\mu\nu} J_1^\mu J_2^\nu. \quad (\text{A.8})$$

For the process $\sigma(\nu_\mu e^- \rightarrow \mu^- \nu_e)$, only possible through the CC s -channel, the total cross section for a ν_μ incident on a stationary electron is

$$\sigma = \frac{2G_F^2}{\pi} m_e E_\nu \quad (\text{A.9})$$

at tree level[150]. Numerically, this is $1.7 \times 10^{-41} (E_\nu/\text{GeV}) \text{ cm}^2$. Neutrino cross sections on nuclei are much less understood. Experimentally, the total charged current cross section for ν_μ on nucleons is about $1 \times 10^{-38} \text{ cm}^2$ at $E_\nu = 1 \text{ GeV}$ in the lab. Clearly, the cross section on nuclei is the dominant contributor to the total neutrino cross section. There are measurements of the rare $\nu - e$ elastic scattering process [151]. These are principally used as an in-situ constraint for a neutrino flux within an experiment. This is useful constraining the flux-induced systematic uncertainties on other cross section measurements within the detector.

Quasi-Elastic Scattering

Below 1 GeV, the ν_μ and ν_e charged current neutrino cross sections are dominated by the quasi-elastic (QE) scattering, the processes:

$$\nu + n \rightarrow l^- + p, \quad (\text{A.10})$$

$$\bar{\nu} + p \rightarrow l^+ + n. \quad (\text{A.11})$$

At these low neutrino energies, the four-momentum transfer to the hadronic system, given by

$$Q^2 = -q^2 = -(p_l - p_\nu)^2 = 2E_\nu(E_l - p_l \cos \theta_{\nu l}) + m_l^2, \quad (\text{A.12})$$

is too low to resolve any of the nucleon's substructure, and the nucleon can be reasonably approximated as a fundamental particle as in Fermi's four-body theory.

Theoretically, interactions with nuclei are unmanageably complicated to model from a microscopic approach. Thus, the QE process is parameterized according to the most general possible Lorentz-invariant description in terms of six Lorentz-invariant hadronic form factors[73]:

$$\frac{d\sigma}{dQ^2} = \frac{G_F^2 m_N^2 |V_{ud}^{\text{CKM}}|^2}{8\pi E_\nu^2} \left(A \pm \frac{s-u}{m_N^2} B + \frac{(s-u)^2}{m_N^4} C \right), \quad (\text{A.13})$$

where m_N gives the nucleon mass, and s and u are the familiar Mandelstam variables and the $-(+)$ applies to $\nu(\bar{\nu})$ scattering. A , B , and C are functions of the six possible form factors, themselves a function of Q^2 . One of these form factors violates CVC, and thus is assumed to be 0. Another violates T symmetry and is assumed to be sub-dominant. This leaves an axial-vector form factor, a pseudoscalar form factor and two vector form factors. The vector form factors are estimated from electron scattering data via the CVC hypothesis. The axial and pseudoscalar form factors, by the PCAC hypothesis and the Goldberger-Treiman relation, are proportional[152]. This leaves the QE cross section parametrized by only a single form factor which must be measured by neutrino scattering experiments[152]. Data shows that for $Q^2 < 1 \text{ GeV}^2$, this can be suitably described as a dipole term:

$$F_A(Q^2) = \frac{F_A(0)}{m_A^2 + Q^2}. \quad (\text{A.14})$$

Fitting bubble chamber data from multiple experiments to a dipole term gives an average $m_A = 1.026 \pm 0.021 \text{ GeV}$ [153].

Deep-Inelastic Scattering

For energies above 5 GeV, neutrino cross sections are dominated by Deep-inelastic scattering (DIS) off a quark within a nucleon according to $\nu + q \rightarrow l^- + X$. In this process, a quark is freed from the nucleus, forming a hadronization jet along with inelastic fragmentation of the nucleon remnant. This is deeply connected to analogous electron scattering processes. The constituent quarks are described using Feynman's parton model to explain measurements taken at SLAC by Bjorken.

Q^2 is the most influential kinematic variable in DIS scattering. As Q^2 increases, shorter lived sea quarks become available to contribute to the scattering. The cross section is best described in the Bjorken x and y scaling variables:

$$y = 1 - \frac{E_l}{E_\nu}, \quad (\text{A.15})$$

$$x = \frac{Q^2}{2m_n E_\nu y}. \quad (\text{A.16})$$

At SLAC, form factors were found to depend only on the Bjorken x variable, rather than both Q^2 and y [154], which Feynman identified with the fraction of the nucleon momentum carried by the struck parton. The cross section, differential in x and y , is shown in Eqn. A.17.

$$\frac{d^2\sigma}{dx dy} = \frac{G_F^2 m_N E_\nu}{\pi(1 + Q^2/M_W^2)^2} \left[y^2 x F_1(x, Q^2) + (1 - y - \frac{m_N}{2E} xy) F_2(x, Q^2) \pm y(1 - \frac{y}{2}) x F_3(x, Q^2) \right]. \quad (\text{A.17})$$

The cross section depends on three form factors. The Callan-Gross relation[155] gives $2xF_1 = F_2$. Both of these form factors are measurable in Q^2 from electron-proton scattering data, giving an experimental constraint on this rule. Parity forces $F_3 = 0$ for electron scattering, and so this is entirely measured with neutrino scattering experiments. For high energies, $E_\nu > 10$ GeV, neutrino experiments give high precision measurements of these form factors. At lower energies in the transitional region, there are fewer measurements in no small part due to ambiguity with other processes.

Resonant Scattering

The high and low energy CC cross sections are dominated by QE and DIS, respectively. Now, in the transitional region illustrated in Figure 1.2, there are additional qualitative scattering processes that appear in tandem with QE and DIS. Among these, resonant scattering describes processes such as

$$\nu + N \rightarrow N^* l \rightarrow N' \pi l, \quad (\text{A.18})$$

where N and N' are both nucleons and N^* is a higher-mass Baryon resonance. The lowest-mass N^* is Δ_{1232} , but there are a total of 18 resonances that contribute to the cross section with mass < 2 GeV. The Rein-Seghal model[156] incorporates each of these resonances. This model includes correlations among the possible N^* . For instance, the $(\Delta^{++}, \Delta^+, \Delta^0, \Delta^-)$ span an isospin 3/2 triplet and 1/2 singlet state, which lead to relationships among scattering rates based on the triplet amplitude, A_3 , and the singlet amplitude, A_1 :

$$\sigma(\nu_{\mu}p \rightarrow \mu^{-}p\pi^{+}) = \sqrt{2}A_3, \quad (\text{A.19})$$

$$\sigma(\nu_{\mu}n \rightarrow \mu^{-}p\pi^{0}) = \frac{2}{3}(A_3 - A_1), \quad (\text{A.20})$$

$$\sigma(\nu_{\mu}n \rightarrow \mu^{-}n\pi^{+}) = \frac{\sqrt{2}}{3}(A_3 + 2A_1). \quad (\text{A.21})$$

Scattering models include such interference between resonances attaching to left and right polarizations. Doing this, the cross section for a given resonance mass is written as

$$\frac{d^2\sigma}{dQ^2 d\nu} = \frac{G_F^2 Q^2}{4\pi|\vec{q}|^2} \left(\frac{M^2 - m_n^2}{2m_n} \right) \left(u^2\sigma_L + v^2\sigma_R + 2uv\sigma_S \right), \quad (\text{A.22})$$

where $\sigma_{L,R}$ are the cross sections derived from left and right-polarized bosons and σ_S gives the scalar polarization contribution[156]. Here, \vec{q} gives the three-momentum transfer by the weak boson, M the resonance mass, and

$$\nu = 1 - E_l/E_\nu, \quad (\text{A.23})$$

$$u = \frac{E_\nu + E_l - |\vec{q}|}{2E_n u}, \quad (\text{A.24})$$

$$v = \frac{E_\nu + E_l + |\vec{q}|}{2E_\nu}. \quad (\text{A.25})$$

Meson Exchange Current (MEC) Scattering

In recent years, it's become evident that the QE picture used to describe scatters off hydrogen in bubble chamber data is insufficient in a nuclear medium. The MEC picture accounts for correlations between nucleons within a target nucleus[157], often mediated by a π field. Necessity of MEC's inclusion in the nuclear cross section was demonstrated by MiniBooNE[158] who measured an QE $m_A = 1.35$ GeV in a detector insensitive to nucleons. Later, the community identified this discrepancy as MEC events contaminating MiniBooNE's QE analysis as the experiment could not discriminate between events with single and multiple knockout nucleons. A calculation of the MEC cross section was shown to resolve the tension between the

MiniBooNE and past experiments[159] to within the experimental and theoretical error bands.

$\nu - A$ Coherent π Production

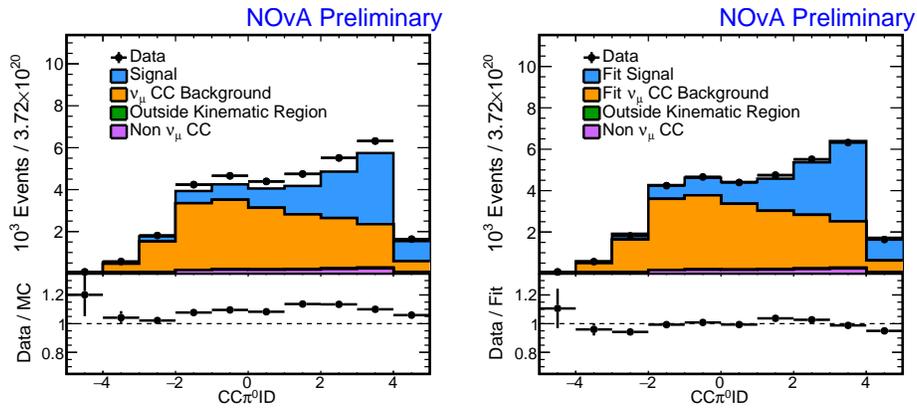
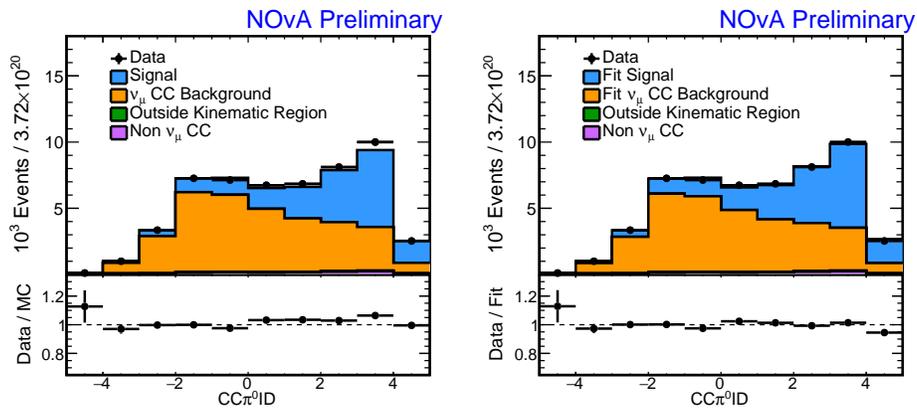
Additionally, a neutrino can scatter off an entire nucleus in such a way that the nucleus does not fragment. As a nucleus is very heavy compared to E_ν , the nuclear recoil is very slight, leaving the final state products constrained to low Q^2 . Both a CC and NC process, governed by $\nu + A \rightarrow l^- + \pi^+ + A$ and $\nu + A \rightarrow \nu + \pi^0 + A$, respectively, along with analogous anti-neutrino processes, are possible.

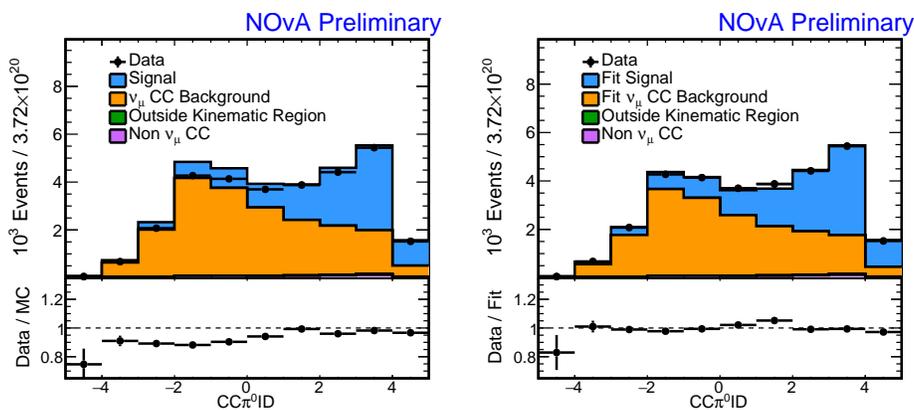
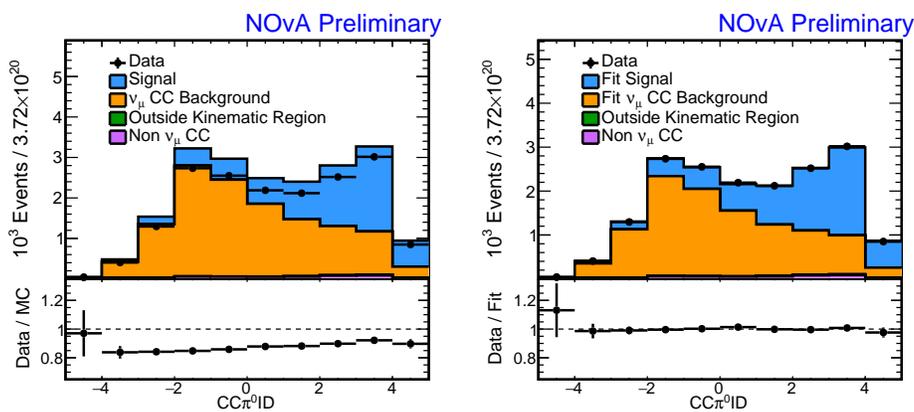
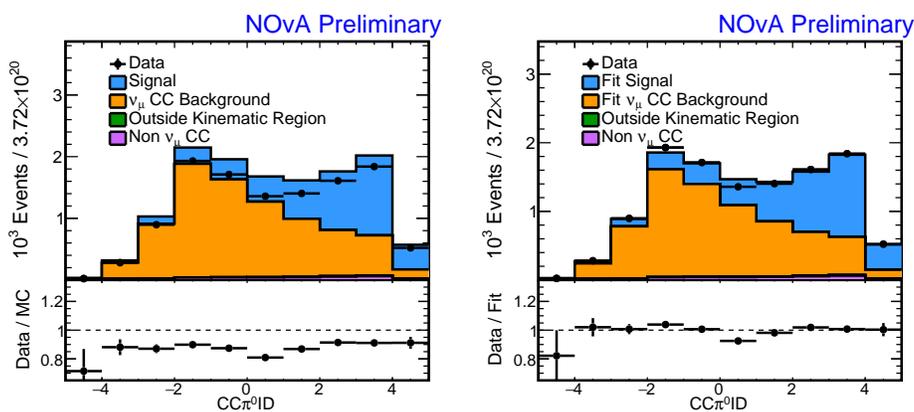
The cross section is described in terms of a kinematic variable unique to coherent pion production, $|t| = (q - p_\pi)^2$. As the process is limited to $Q^2 \approx 0$, the cross section can be calculated based on the PCAC prediction[160]

$$\frac{d^3\sigma}{dx dy d|t|} = \frac{G_F^2 m_n f_\pi^2 E_\nu (1-y)}{2\pi^2} \left(\frac{m_A^2}{Q^2 + m_A^2} \right)^2 \frac{d\sigma(\pi^0 A \rightarrow \pi^0 A)}{d|t|}. \quad (\text{A.26})$$

The coherent π^\pm cross section is about 1% of the total CC cross section at a few GeV. Due to its low cross section, measurements of this channel are predominantly motivated theoretically, as a test of PCAC in such a scenario.

Appendix B

CC π^0 ID FITS FOR CROSS SECTION ANALYSISB.1 Measurement in p_μ Figure B.1: CC π^0 ID fit to data for events with $0 < p_\mu < 0.5$ GeV/cFigure B.2: CC π^0 ID fit to data for events with $0.5 < p_\mu < 0.75$ GeV/c

Figure B.3: $CC\pi^0ID$ fit to data for events with $0.75 < p_\mu < 1 \text{ GeV}/c$ Figure B.4: $CC\pi^0ID$ fit to data for events with $1 < p_\mu < 1.25 \text{ GeV}/c$ Figure B.5: $CC\pi^0ID$ fit to data for events with $1.25 < p_\mu < 1.5 \text{ GeV}/c$

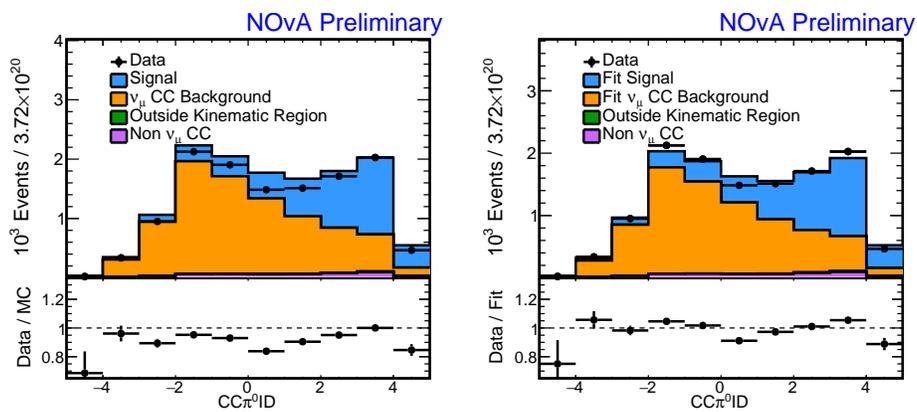


Figure B.6: $CC\pi^0$ ID fit to data for events with $1.5 < p_\mu < 2$ GeV/c

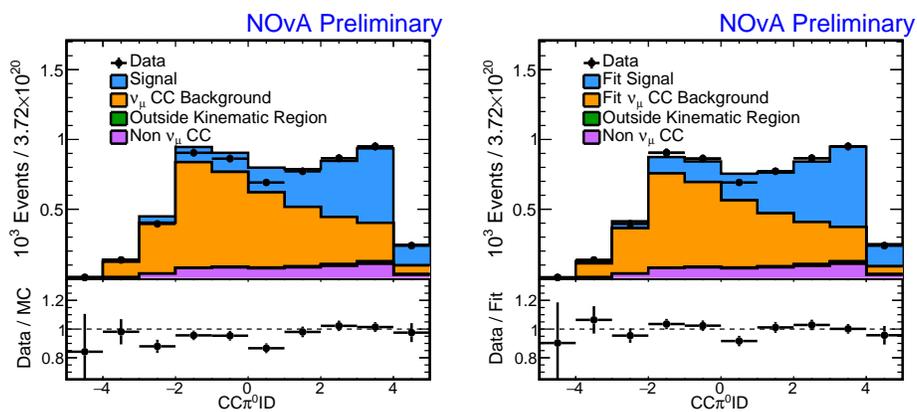


Figure B.7: $CC\pi^0$ ID fit to data for events with $2 < p_\mu < 4$ GeV/c

B.2 Measurement in $\cos \theta_\mu$

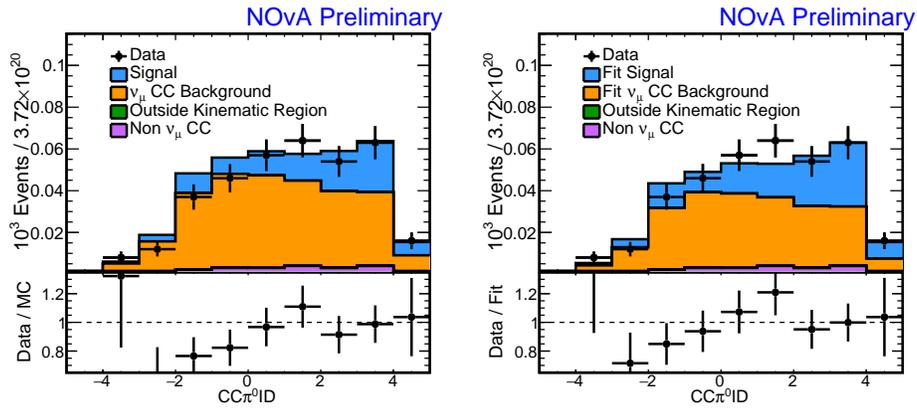


Figure B.8: $CC\pi^0ID$ fit to data for events with $-1 < \cos \theta_\mu < 0$

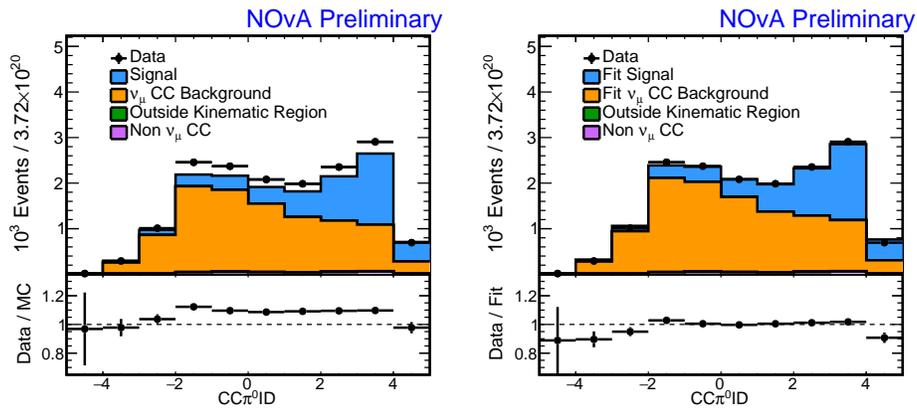


Figure B.9: $CC\pi^0ID$ fit to data for events with $0 < \cos \theta_\mu < 0.5$

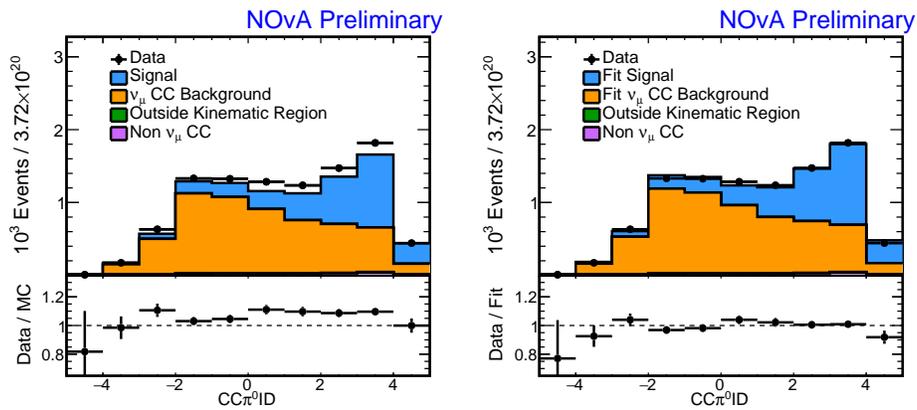
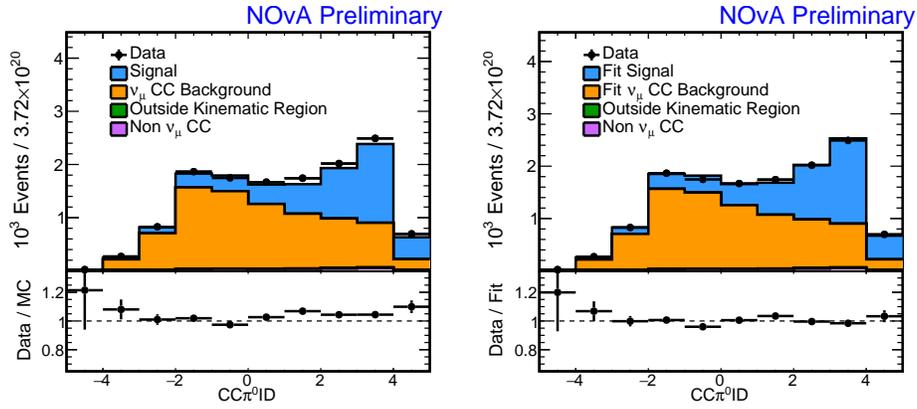
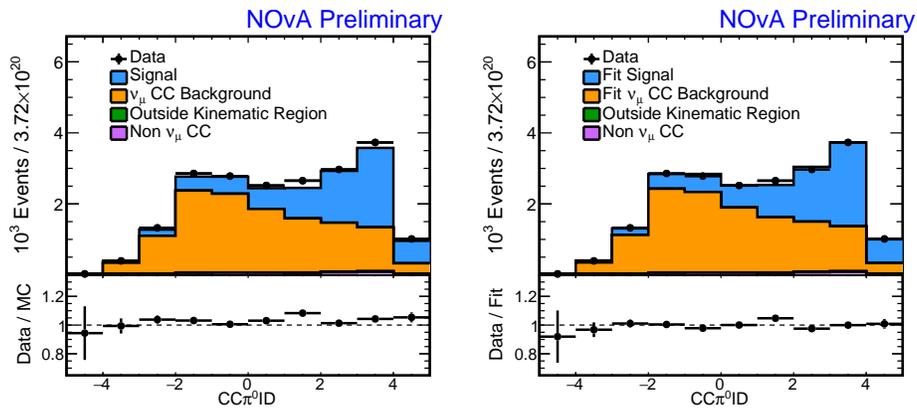
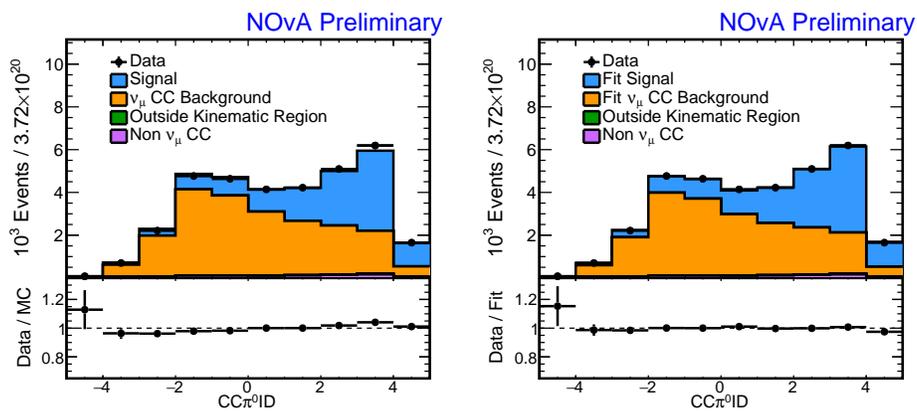


Figure B.10: $CC\pi^0ID$ fit to data for events with $0.5 < \cos \theta_\mu < 0.6$

Figure B.11: $CC\pi^0ID$ fit to data for events with $0.6 < \cos \theta_\mu < 0.7$ Figure B.12: $CC\pi^0ID$ fit to data for events with $0.7 < \cos \theta_\mu < 0.8$ Figure B.13: $CC\pi^0ID$ fit to data for events with $0.8 < \cos \theta_\mu < 0.9$

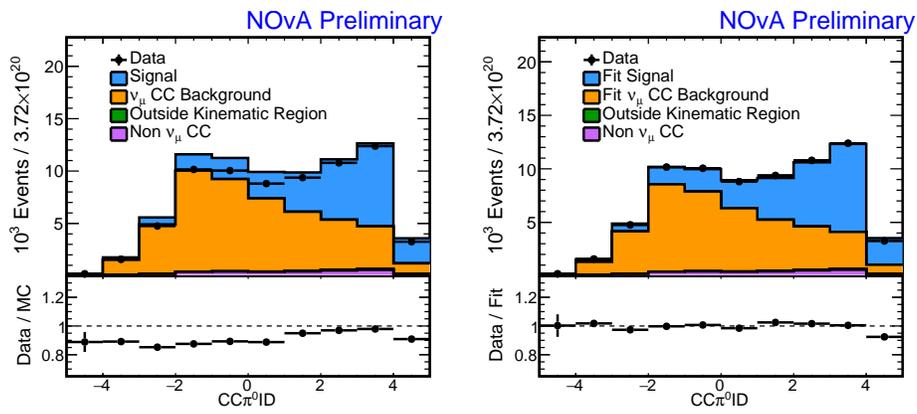


Figure B.14: $CC\pi^0ID$ fit to data for events with $0.9 < \cos \theta_\mu < 1$

B.3 Measurement in p_π

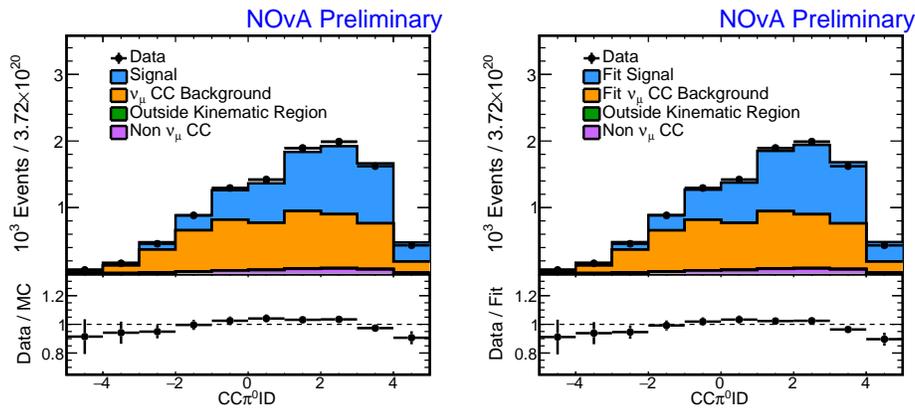


Figure B.15: $CC\pi^0ID$ fit to data for events with $0 < p_\pi < 0.2$ GeV/c

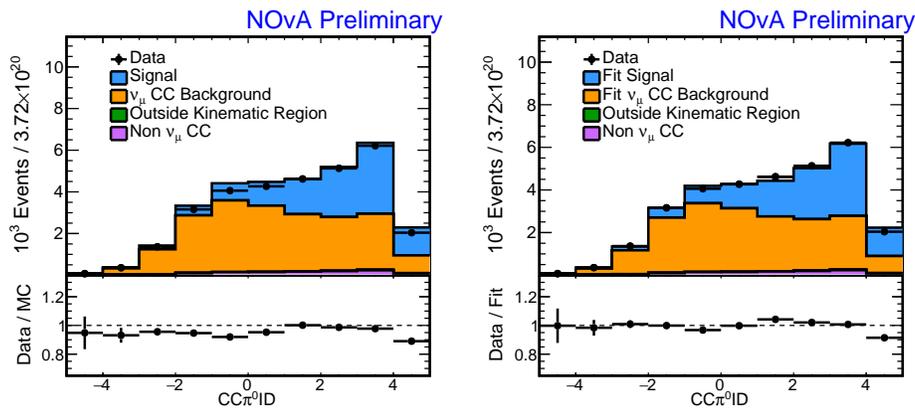


Figure B.16: $CC\pi^0ID$ fit to data for events with $0.2 < p_\pi < 0.3$ GeV/c

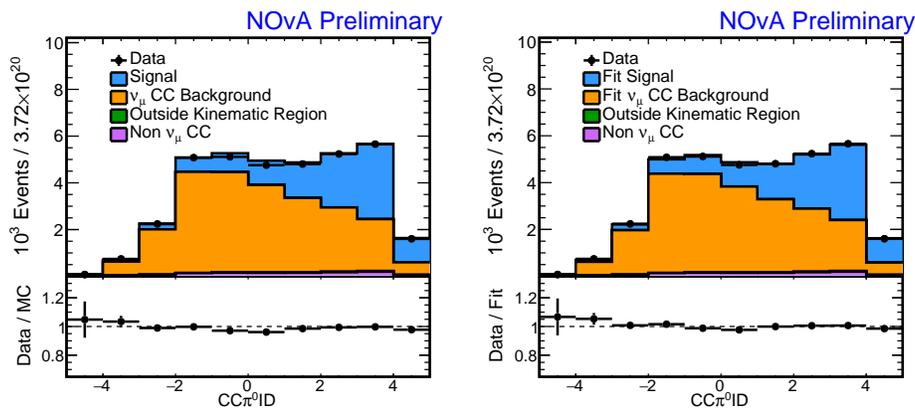
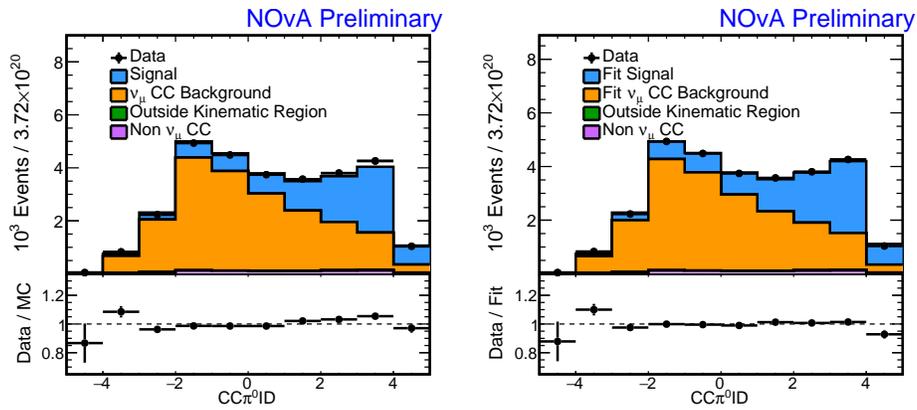
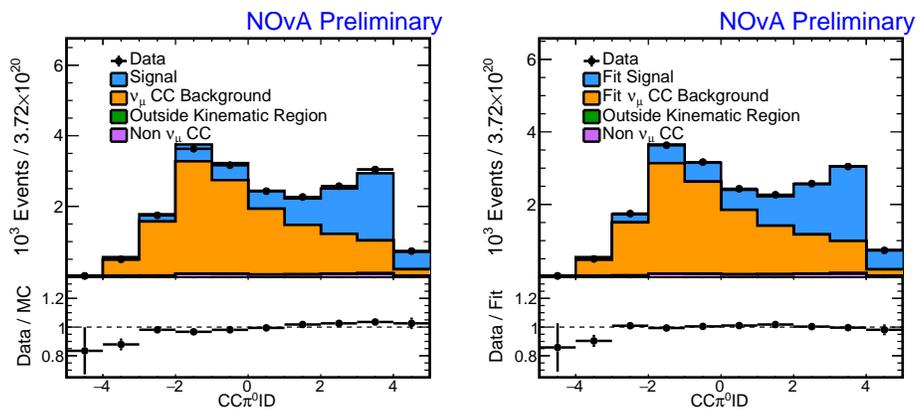
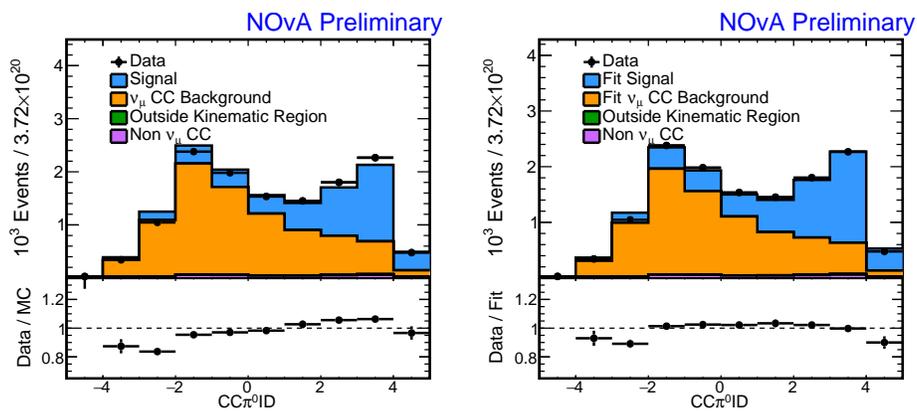
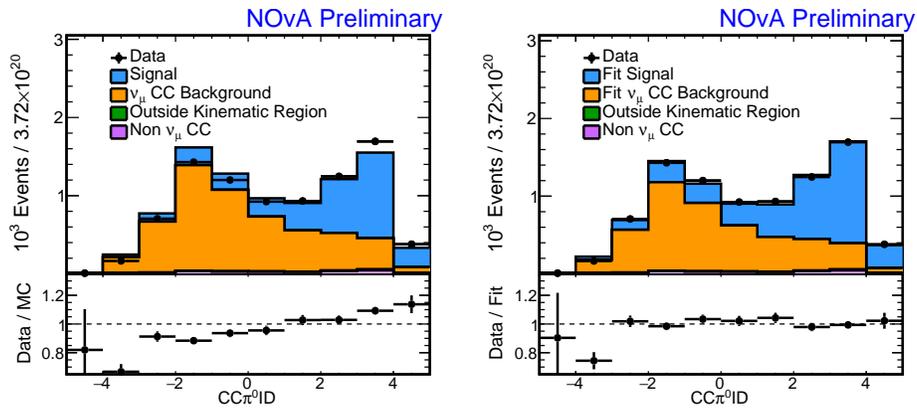
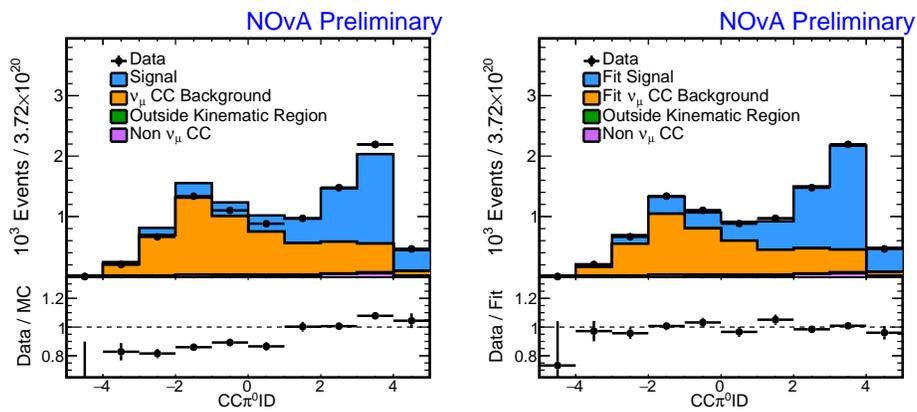
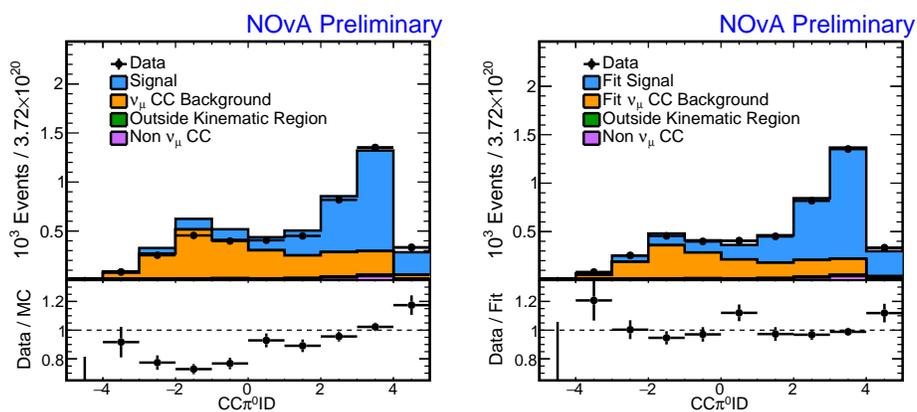


Figure B.17: $CC\pi^0ID$ fit to data for events with $0.3 < p_\pi < 0.4$ GeV/c

Figure B.18: $CC\pi^0ID$ fit to data for events with $0.4 < p_\pi < 0.5$ GeV/cFigure B.19: $CC\pi^0ID$ fit to data for events with $0.5 < p_\pi < 0.6$ GeV/cFigure B.20: $CC\pi^0ID$ fit to data for events with $0.6 < p_\pi < 0.7$ GeV/c

Figure B.21: $CC\pi^0ID$ fit to data for events with $0.7 < p_\pi < 0.8$ GeV/cFigure B.22: $CC\pi^0ID$ fit to data for events with $0.8 < p_\pi < 1$ GeV/cFigure B.23: $CC\pi^0ID$ fit to data for events with $1 < p_\pi < 1.25$ GeV/c

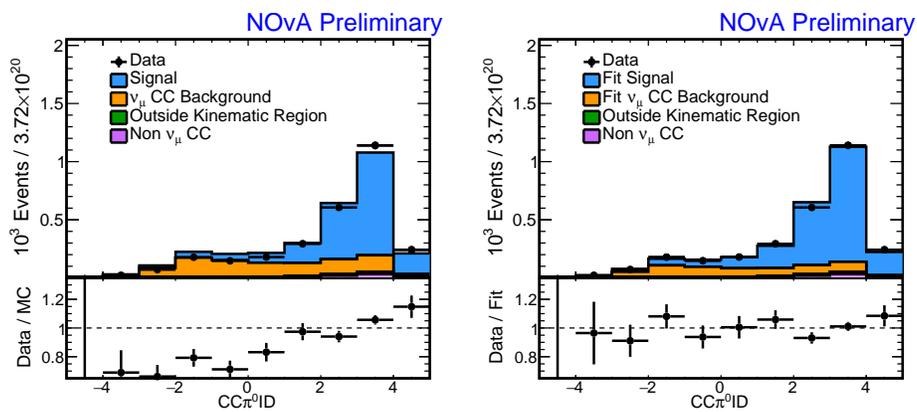


Figure B.24: $CC\pi^0$ ID fit to data for events with $1.25 < p_\pi < 2$ GeV/c

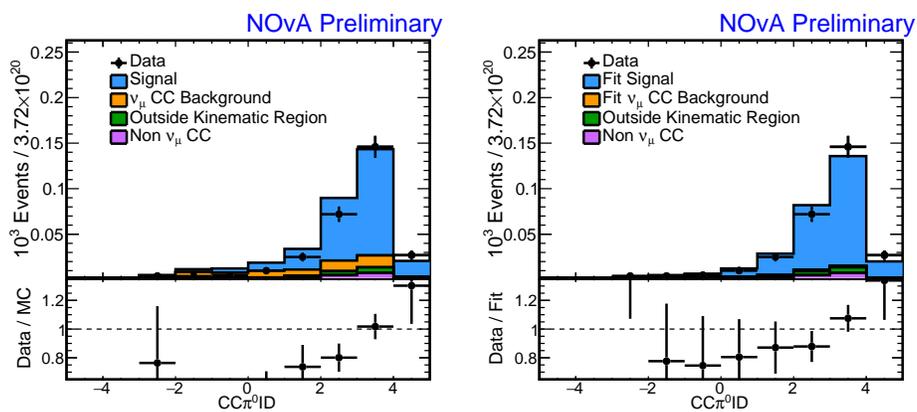


Figure B.25: $CC\pi^0$ ID fit to data for events with $2 < p_\pi < 3$ GeV/c

B.4 Measurement in $\cos \theta_\pi$

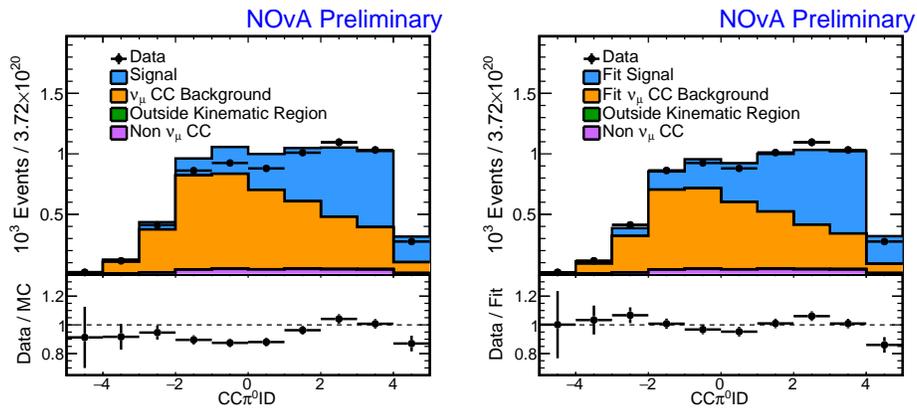


Figure B.26: $CC\pi^0ID$ fit to data for events with $-1 < \cos \theta_\pi < -0.5$

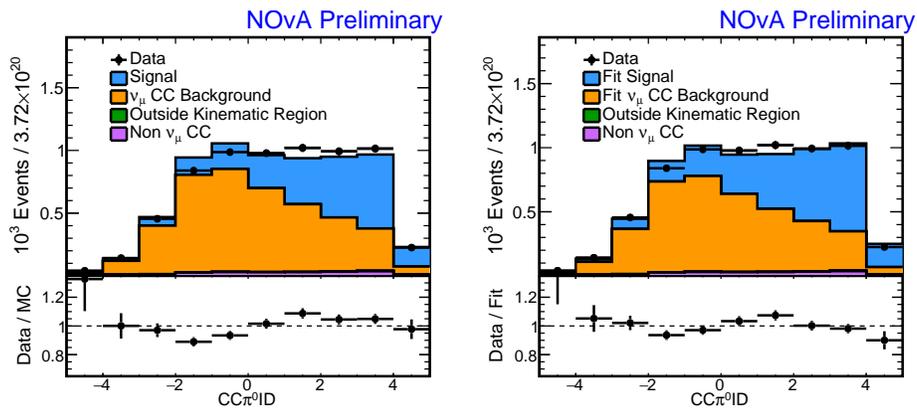


Figure B.27: $CC\pi^0ID$ fit to data for events with $-0.5 < \cos \theta_\pi < 0$

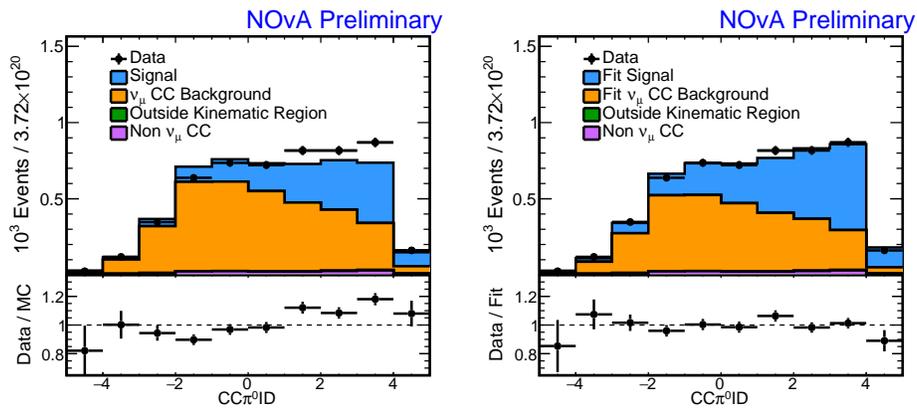
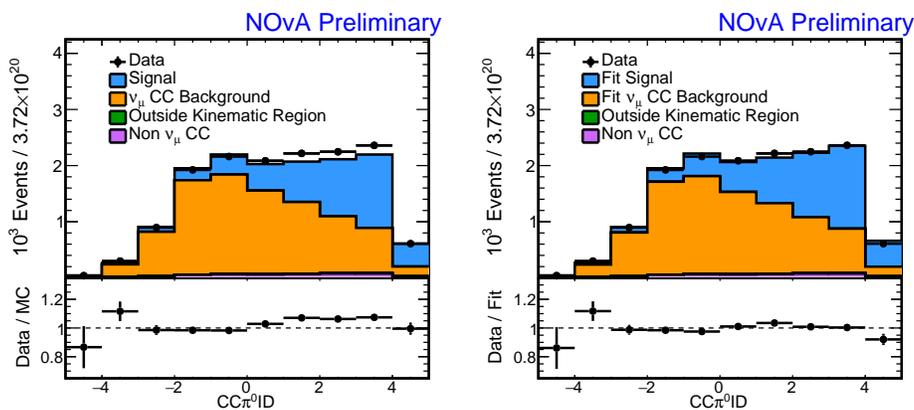
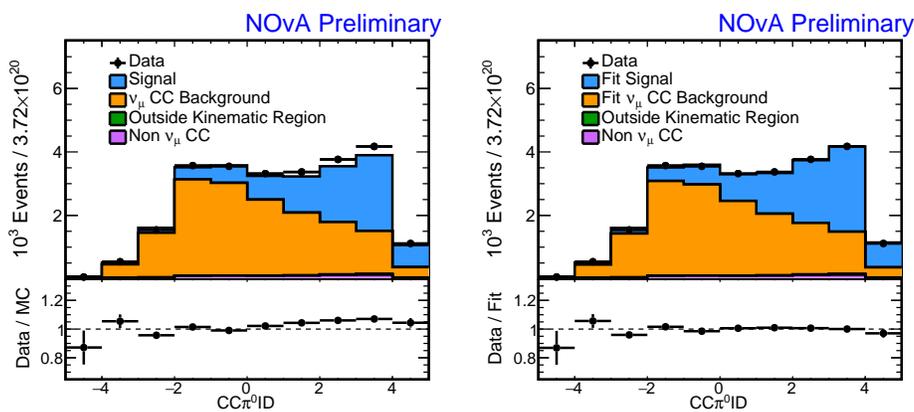
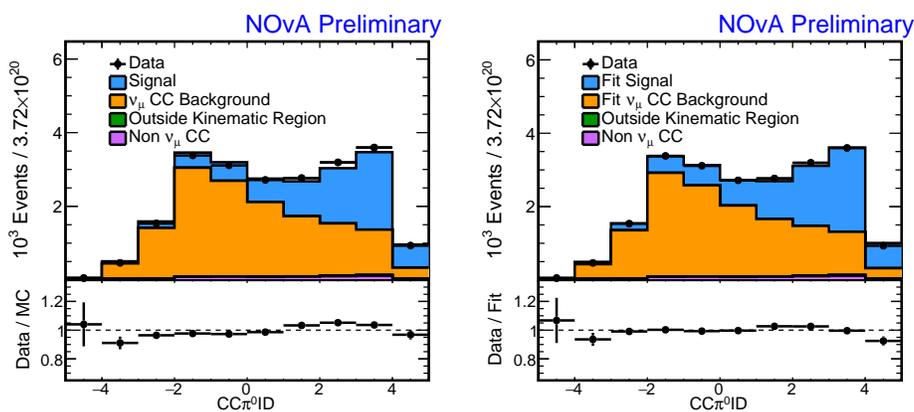


Figure B.28: $CC\pi^0ID$ fit to data for events with $0 < \cos \theta_\pi < 0.25$

Figure B.29: $CC\pi^0$ ID fit to data for events with $0.25 < \cos \theta_\pi < 0.5$ Figure B.30: $CC\pi^0$ ID fit to data for events with $0.5 < \cos \theta_\pi < 0.7$ Figure B.31: $CC\pi^0$ ID fit to data for events with $0.7 < \cos \theta_\pi < 0.8$

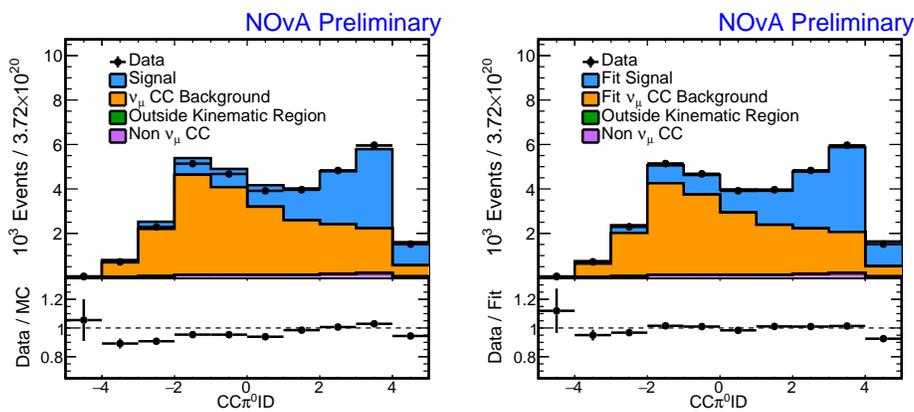


Figure B.32: $CC\pi^0ID$ fit to data for events with $0.8 < \cos \theta_\pi < 0.9$

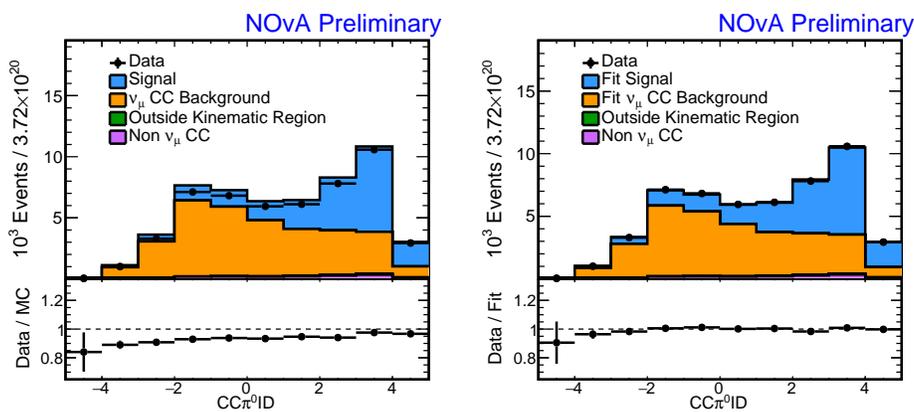


Figure B.33: $CC\pi^0ID$ fit to data for events with $0.9 < \cos \theta_\pi < 1$

B.5 Measurement in Q^2

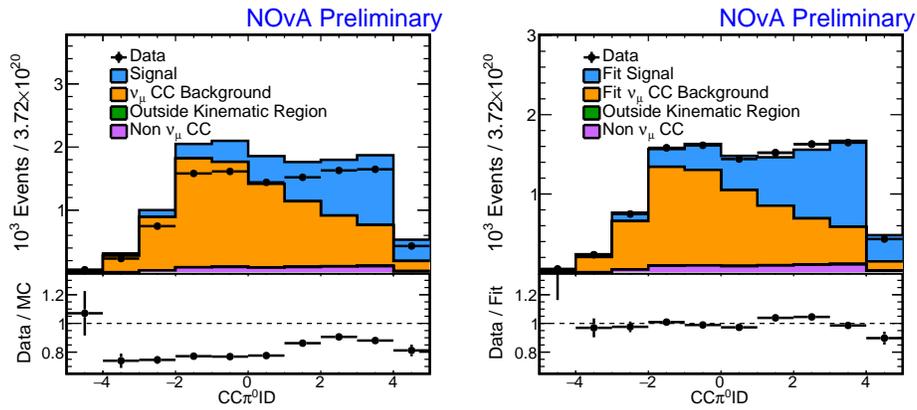


Figure B.34: $CC\pi^0ID$ fit to data for events with $0 < Q^2 < 0.1 \text{ GeV}^2/c^2$

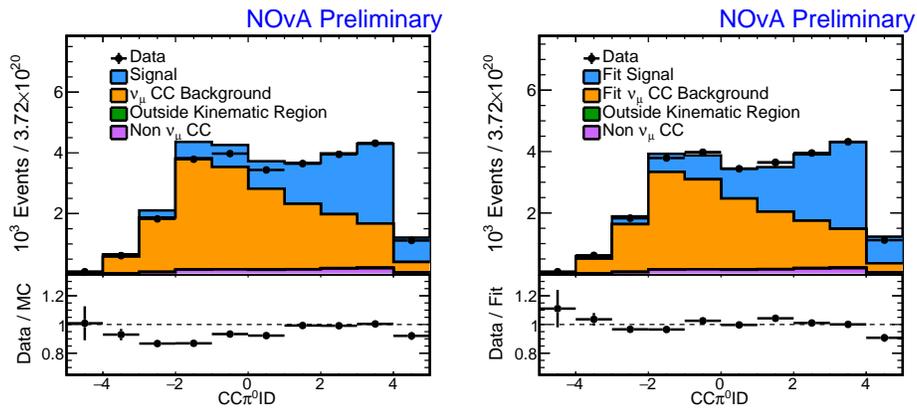


Figure B.35: $CC\pi^0ID$ fit to data for events with $0.1 < Q^2 < 0.25 \text{ GeV}^2/c^2$

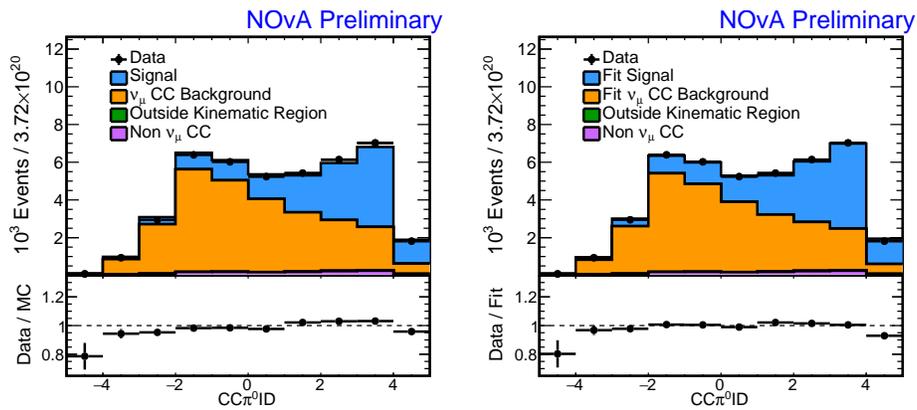


Figure B.36: $CC\pi^0ID$ fit to data for events with $0.25 < Q^2 < 0.5 \text{ GeV}^2/c^2$

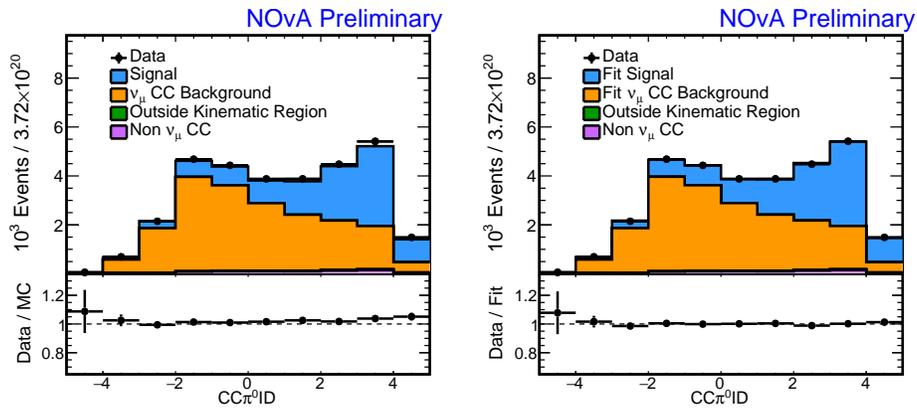


Figure B.37: $CC\pi^0ID$ fit to data for events with $0.5 < Q^2 < 0.75 \text{ GeV}^2/c^2$

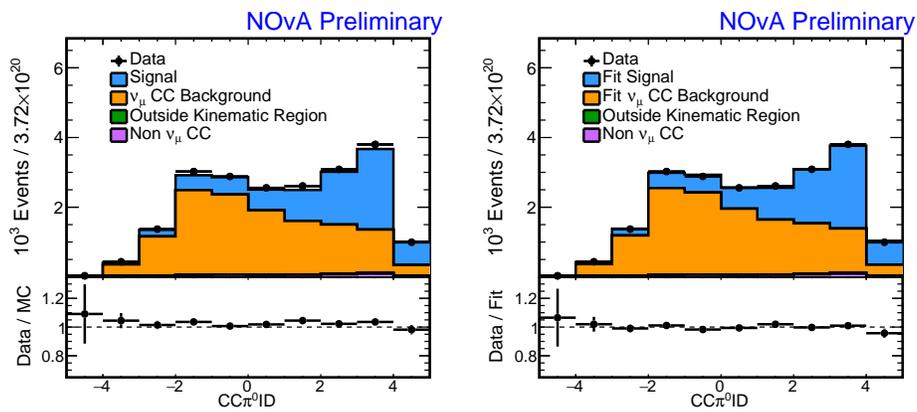


Figure B.38: $CC\pi^0ID$ fit to data for events with $0.75 < Q^2 < 1 \text{ GeV}^2/c^2$

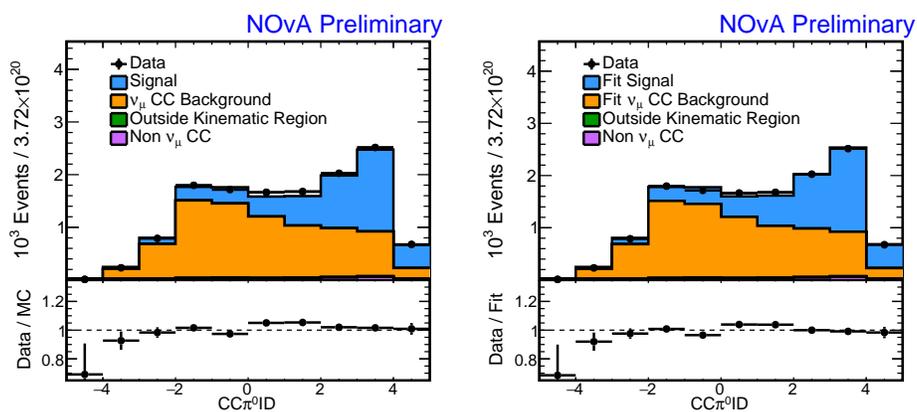
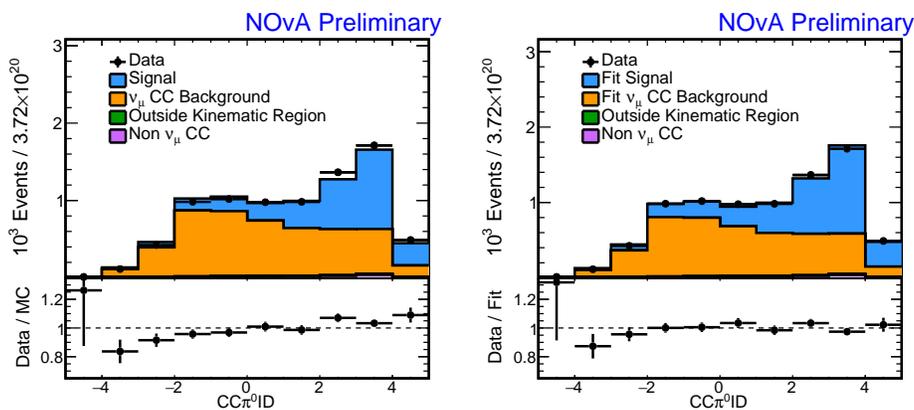
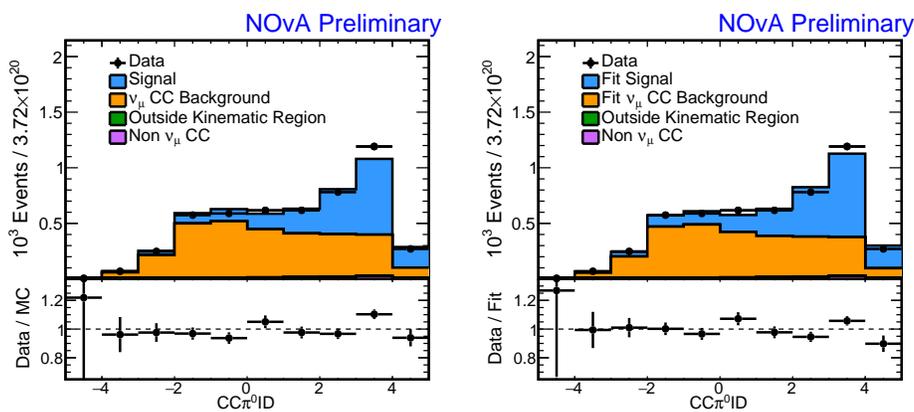
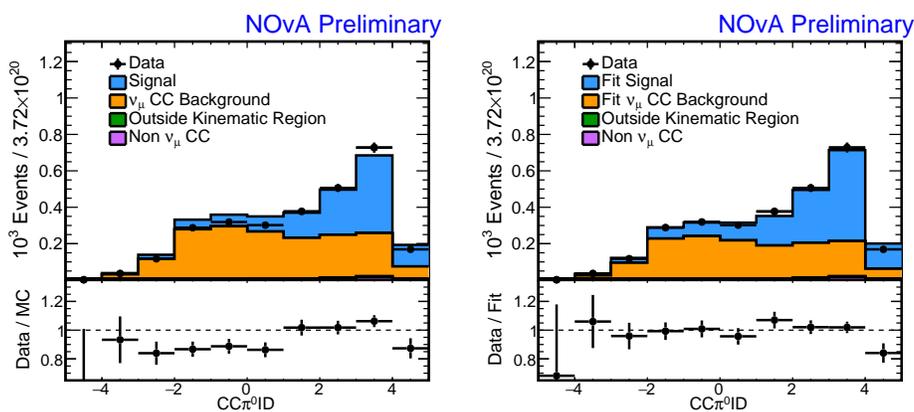


Figure B.39: $CC\pi^0ID$ fit to data for events with $1 < Q^2 < 1.25 \text{ GeV}^2/c^2$

Figure B.40: CCπ⁰ID fit to data for events with $1.25 < Q^2 < 1.5 \text{ GeV}^2/c^2$ Figure B.41: CCπ⁰ID fit to data for events with $1.5 < Q^2 < 1.75 \text{ GeV}^2/c^2$ Figure B.42: CCπ⁰ID fit to data for events with $1.75 < Q^2 < 2 \text{ GeV}^2/c^2$

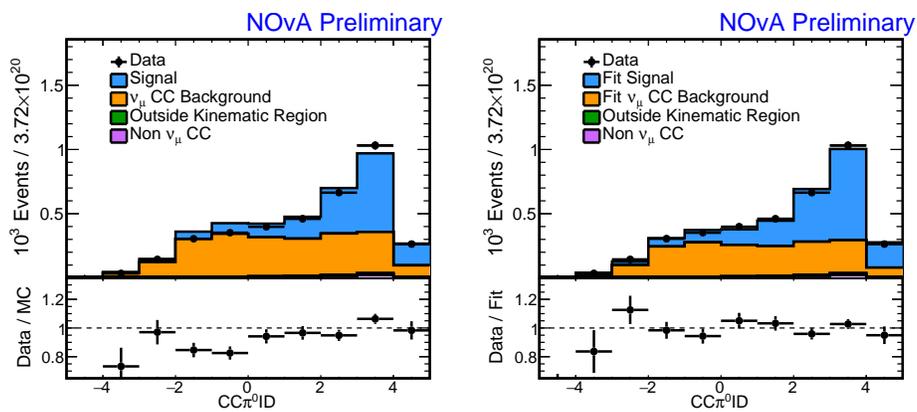


Figure B.43: $CC\pi^0ID$ fit to data for events with $2 < Q^2 < 3 \text{ GeV}^2/c^2$

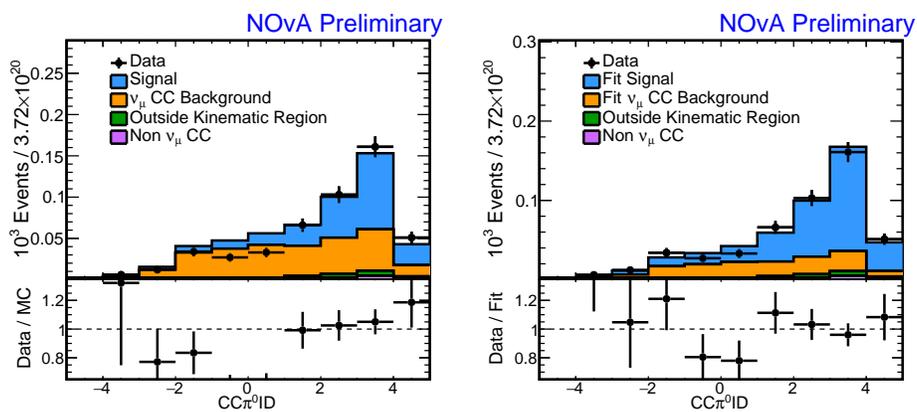


Figure B.44: $CC\pi^0ID$ fit to data for events with $3 < Q^2 < 4 \text{ GeV}^2/c^2$

B.6 Measurement in W

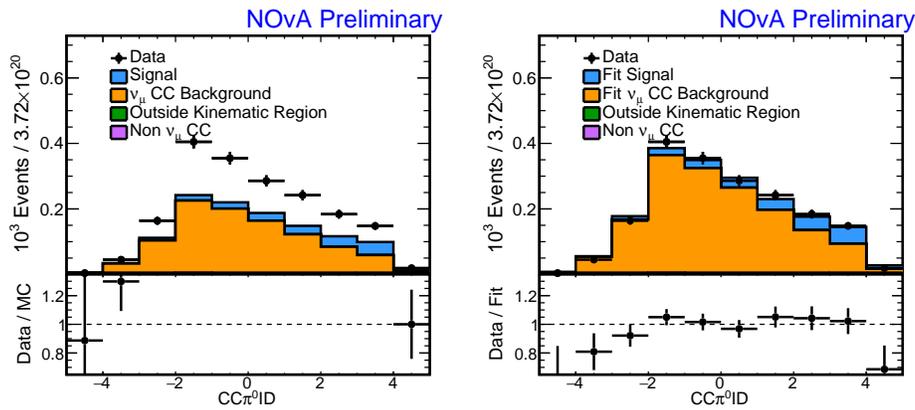


Figure B.45: $CC\pi^0ID$ fit to data for events with $1 < W < 1.1 \text{ GeV}/c^2$

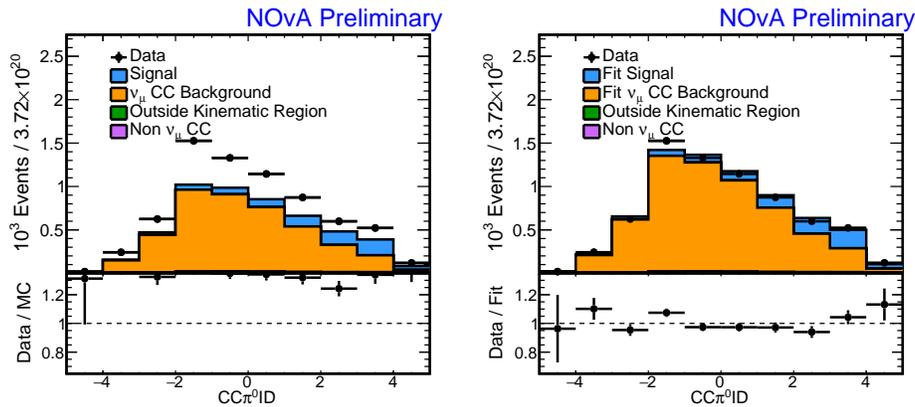


Figure B.46: $CC\pi^0ID$ fit to data for events with $1.1 < W < 1.2 \text{ GeV}/c^2$

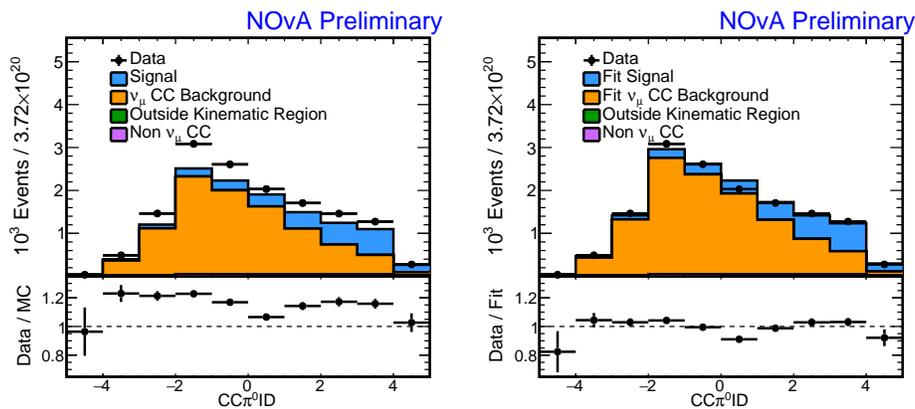
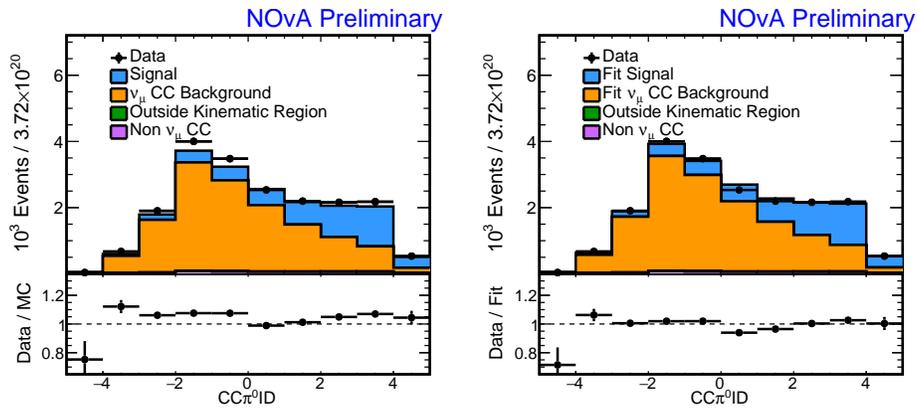
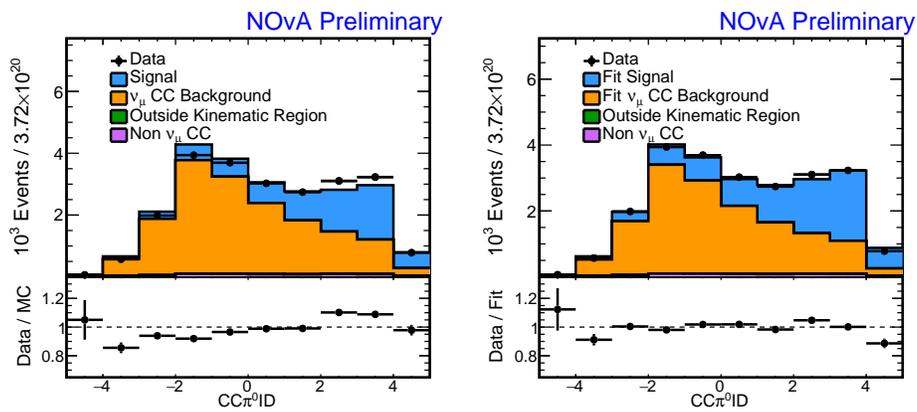
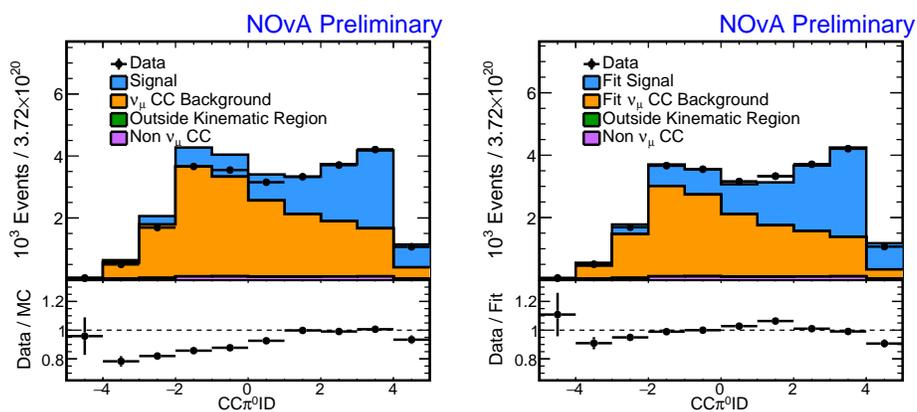
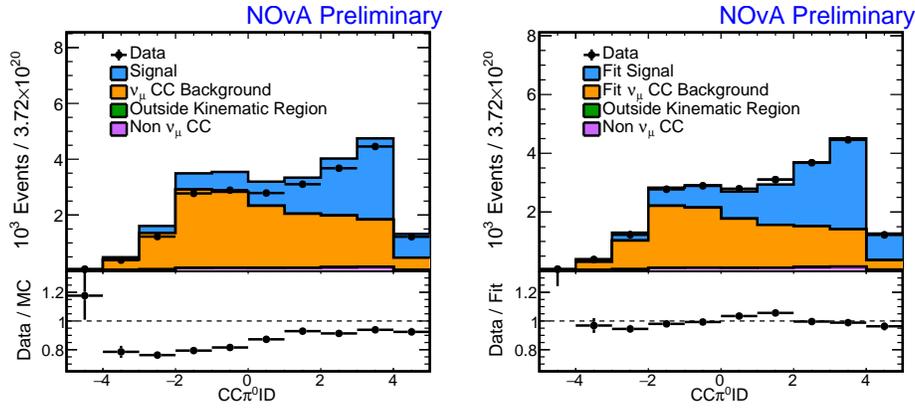
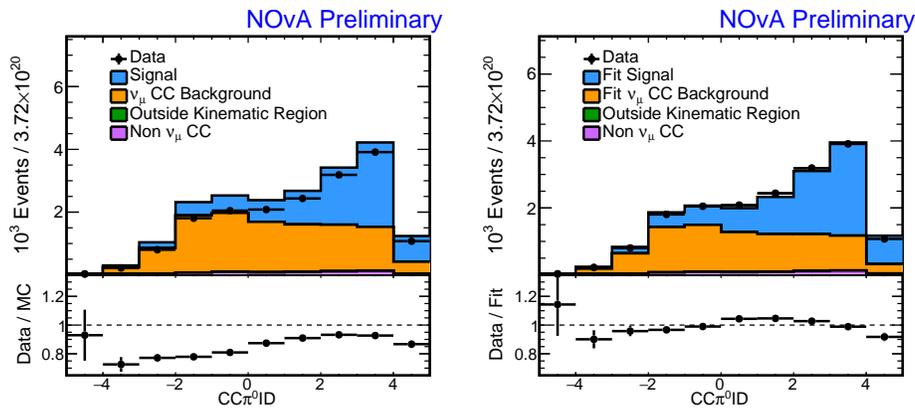
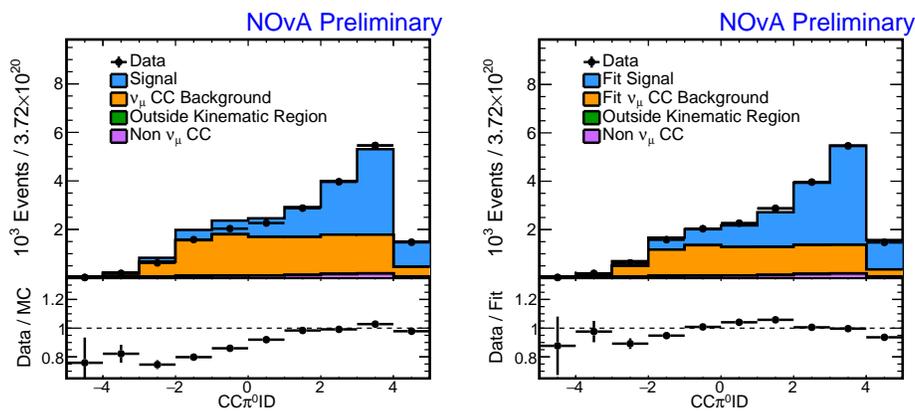
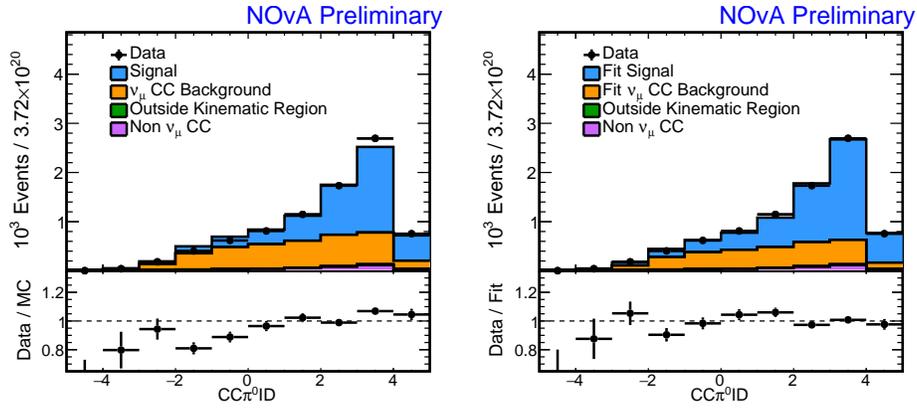
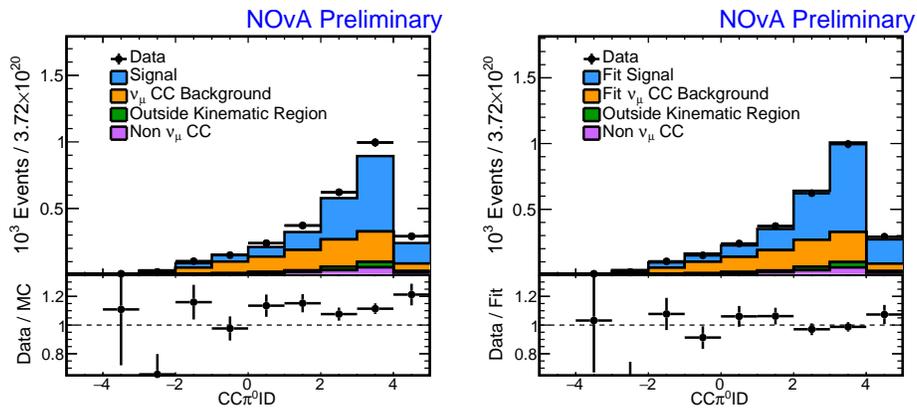
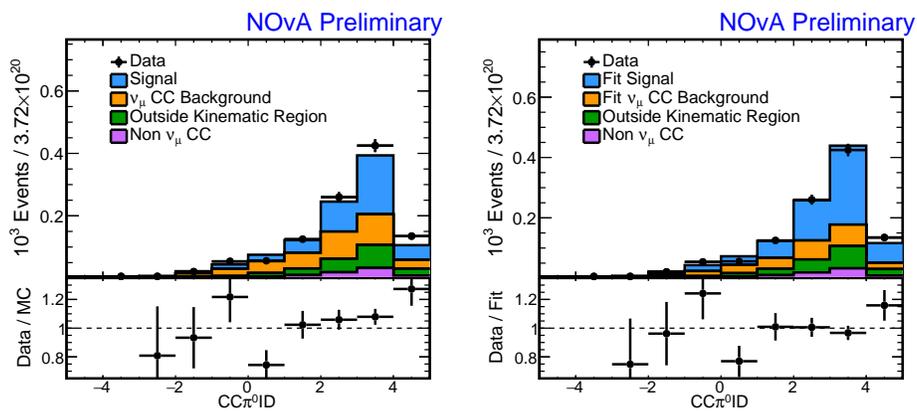


Figure B.47: $CC\pi^0ID$ fit to data for events with $1.2 < W < 1.3 \text{ GeV}/c^2$

Figure B.48: $CC\pi^0ID$ fit to data for events with $1.3 < W < 1.4 \text{ GeV}/c^2$ Figure B.49: $CC\pi^0ID$ fit to data for events with $1.4 < W < 1.5 \text{ GeV}/c^2$ Figure B.50: $CC\pi^0ID$ fit to data for events with $1.5 < W < 1.6 \text{ GeV}/c^2$

Figure B.51: $CC\pi^0ID$ fit to data for events with $1.6 < W < 1.7 \text{ GeV}/c^2$ Figure B.52: $CC\pi^0ID$ fit to data for events with $1.7 < W < 1.8 \text{ GeV}/c^2$ Figure B.53: $CC\pi^0ID$ fit to data for events with $1.8 < W < 2 \text{ GeV}/c^2$

Figure B.54: $CC\pi^0$ ID fit to data for events with $2 < W < 2.25 \text{ GeV}/c^2$ Figure B.55: $CC\pi^0$ ID fit to data for events with $2.25 < W < 2.5 \text{ GeV}/c^2$ Figure B.56: $CC\pi^0$ ID fit to data for events with $2.5 < W < 3 \text{ GeV}/c^2$

*Appendix C***CC π^0 CROSS SECTION AND COVARIANCE TABLES**

In the following pages, there is a summary table for each of the differential measurements for CC π^0 cross section measurement. Within these tables, we compare the central value of the extracted cross section with NOvA data along with the GENIE prediction averaged over the NuMI flux. Along with this, there is a covariance and correlation matrix provided for convenience when adding this result with other measurements.

| | | | | | | | | | | | | | | | | |
|--------------------------|-------|-------|-------|-------|-------|-------|-------|-------|------|-------|------|-------|-------|--------|--------|--------|
| | 0.00 | 0.10 | 0.20 | 0.30 | 0.40 | 0.50 | 0.60 | 0.70 | 0.80 | 0.90 | 1.00 | 1.25 | 1.50 | 1.75 | 2.00 | 2.50 |
| $\frac{d\sigma}{dp_\pi}$ | 2.03 | 5.51 | 4.86 | 4.62 | 4.06 | 3.20 | 2.51 | 1.93 | 1.53 | 1.12 | 0.74 | 0.39 | 0.21 | 0.11 | 0.041 | 0.012 |
| GENIE pred | 1.92 | 5.33 | 4.81 | 4.55 | 3.89 | 2.98 | 2.26 | 1.70 | 1.33 | 0.98 | 0.65 | 0.35 | 0.18 | 0.10 | 0.040 | 0.012 |
| 0.00 | 1.13 | 2.20 | 1.70 | 1.42 | 1.06 | 0.76 | 0.58 | 0.40 | 0.27 | 0.17 | 0.23 | 0.10 | 0.048 | 0.017 | -0.006 | -0.014 |
| 0.10 | 0.80 | 6.76 | 5.92 | 4.95 | 3.50 | 2.18 | 1.69 | 1.07 | 0.68 | 0.40 | 0.44 | 0.21 | 0.11 | -0.002 | -0.073 | -0.077 |
| 0.20 | 0.68 | 0.96 | 5.58 | 4.68 | 3.25 | 1.94 | 1.53 | 0.94 | 0.58 | 0.36 | 0.39 | 0.21 | 0.14 | 0.020 | -0.063 | -0.069 |
| 0.30 | 0.65 | 0.93 | 0.96 | 4.23 | 3.01 | 1.84 | 1.45 | 0.91 | 0.58 | 0.38 | 0.46 | 0.28 | 0.17 | 0.057 | -0.023 | -0.045 |
| 0.40 | 0.65 | 0.88 | 0.90 | 0.95 | 2.36 | 1.52 | 1.19 | 0.76 | 0.52 | 0.32 | 0.45 | 0.27 | 0.16 | 0.058 | 0.008 | -0.021 |
| 0.50 | 0.68 | 0.79 | 0.78 | 0.85 | 0.94 | 1.11 | 0.82 | 0.56 | 0.41 | 0.24 | 0.38 | 0.21 | 0.12 | 0.048 | 0.024 | -0.008 |
| 0.60 | 0.67 | 0.80 | 0.80 | 0.87 | 0.95 | 0.96 | 0.66 | 0.43 | 0.31 | 0.19 | 0.31 | 0.18 | 0.093 | 0.042 | 0.023 | -0.002 |
| 0.70 | 0.68 | 0.74 | 0.72 | 0.80 | 0.90 | 0.96 | 0.96 | 0.31 | 0.22 | 0.14 | 0.23 | 0.13 | 0.065 | 0.028 | 0.018 | -0.003 |
| 0.80 | 0.59 | 0.61 | 0.58 | 0.66 | 0.80 | 0.91 | 0.90 | 0.95 | 0.18 | 0.11 | 0.19 | 0.11 | 0.058 | 0.026 | 0.021 | 0.002 |
| 0.90 | 0.56 | 0.54 | 0.54 | 0.64 | 0.72 | 0.80 | 0.84 | 0.90 | 0.89 | 0.082 | 0.14 | 0.076 | 0.041 | 0.023 | 0.015 | 0.002 |
| 1.00 | 0.42 | 0.33 | 0.33 | 0.44 | 0.58 | 0.71 | 0.74 | 0.82 | 0.89 | 0.93 | 0.26 | 0.15 | 0.082 | 0.047 | 0.038 | 0.012 |
| 1.25 | 0.30 | 0.24 | 0.27 | 0.41 | 0.53 | 0.61 | 0.66 | 0.70 | 0.76 | 0.81 | 0.91 | 0.11 | 0.060 | 0.036 | 0.027 | 0.011 |
| 1.50 | 0.21 | 0.20 | 0.27 | 0.39 | 0.47 | 0.53 | 0.53 | 0.54 | 0.62 | 0.65 | 0.74 | 0.85 | 0.047 | 0.025 | 0.015 | 0.006 |
| 1.75 | 0.12 | -0.01 | 0.06 | 0.19 | 0.27 | 0.32 | 0.37 | 0.36 | 0.44 | 0.56 | 0.65 | 0.78 | 0.82 | 0.020 | 0.013 | 0.007 |
| 2.00 | -0.05 | -0.24 | -0.23 | -0.10 | 0.04 | 0.20 | 0.24 | 0.28 | 0.42 | 0.44 | 0.65 | 0.72 | 0.61 | 0.77 | 0.013 | 0.007 |
| 2.50 | -0.16 | -0.36 | -0.35 | -0.27 | -0.16 | -0.09 | -0.04 | -0.06 | 0.06 | 0.11 | 0.28 | 0.42 | 0.35 | 0.57 | 0.72 | 0.007 |

Table C.1: A summary of the extracted $CC\pi^0$ cross section, differential in p_π . The top row, and the left column give the lower-edges of each analysis bin. The second row gives the central value of the cross section measurement in $10^{-39}\text{cm}^2/(\text{GeV}/c)$ for each kinematic bin while the third row gives the GENIE prediction. The remaining matrix gives the covariance matrix, the numbers above and including the diagonal entries of the matrix in $10^{-81}\text{cm}^4/(\text{GeV}/c)^2$, while entries below the diagonal give the correlation matrix.

| | -1.00 | -0.75 | -0.50 | -0.25 | 0.00 | 0.10 | 0.20 | 0.30 | 0.40 | 0.50 | 0.60 | 0.70 | 0.80 | 0.90 |
|-----------------------------------|-------|-------|-------|-------|------|------|------|------|------|------|------|------|------|-------|
| $\frac{d\sigma}{d\cos\theta_\pi}$ | 0.47 | 0.53 | 0.64 | 0.81 | 1.00 | 1.15 | 1.32 | 1.52 | 1.78 | 2.15 | 2.71 | 3.56 | 5.13 | 9.25 |
| GENIE pred | 0.45 | 0.49 | 0.58 | 0.72 | 0.87 | 0.98 | 1.12 | 1.32 | 1.57 | 1.92 | 2.44 | 3.27 | 4.86 | 9.31 |
| -1.00 | 0.30 | 0.36 | 0.42 | 0.46 | 0.26 | 0.23 | 0.28 | 0.31 | 0.26 | 0.38 | 0.52 | 0.62 | 1.02 | 1.89 |
| -0.75 | 0.97 | 0.46 | 0.51 | 0.55 | 0.33 | 0.29 | 0.35 | 0.39 | 0.32 | 0.48 | 0.66 | 0.76 | 1.23 | 2.32 |
| -0.50 | 0.96 | 0.95 | 0.63 | 0.64 | 0.39 | 0.34 | 0.42 | 0.47 | 0.37 | 0.57 | 0.79 | 0.90 | 1.46 | 2.81 |
| -0.25 | 0.98 | 0.94 | 0.94 | 0.75 | 0.42 | 0.38 | 0.44 | 0.48 | 0.42 | 0.61 | 0.81 | 0.98 | 1.57 | 2.88 |
| 0.00 | 0.92 | 0.93 | 0.95 | 0.93 | 0.27 | 0.22 | 0.26 | 0.31 | 0.23 | 0.37 | 0.50 | 0.55 | 0.88 | 1.70 |
| 0.10 | 0.95 | 0.95 | 0.95 | 0.97 | 0.95 | 0.20 | 0.24 | 0.27 | 0.22 | 0.34 | 0.45 | 0.52 | 0.81 | 1.53 |
| 0.20 | 0.95 | 0.96 | 0.96 | 0.93 | 0.93 | 0.97 | 0.30 | 0.33 | 0.28 | 0.41 | 0.56 | 0.64 | 1.02 | 1.94 |
| 0.30 | 0.87 | 0.91 | 0.92 | 0.87 | 0.92 | 0.94 | 0.94 | 0.41 | 0.31 | 0.49 | 0.66 | 0.71 | 1.09 | 2.17 |
| 0.40 | 0.87 | 0.87 | 0.85 | 0.87 | 0.80 | 0.91 | 0.93 | 0.88 | 0.30 | 0.40 | 0.54 | 0.64 | 0.98 | 1.84 |
| 0.50 | 0.88 | 0.90 | 0.90 | 0.89 | 0.90 | 0.95 | 0.95 | 0.96 | 0.92 | 0.63 | 0.82 | 0.92 | 1.40 | 2.73 |
| 0.60 | 0.90 | 0.92 | 0.94 | 0.88 | 0.91 | 0.94 | 0.97 | 0.97 | 0.92 | 0.97 | 1.13 | 1.24 | 1.93 | 3.80 |
| 0.70 | 0.92 | 0.91 | 0.92 | 0.92 | 0.86 | 0.94 | 0.96 | 0.90 | 0.95 | 0.95 | 0.95 | 1.51 | 2.34 | 4.42 |
| 0.80 | 0.95 | 0.93 | 0.94 | 0.93 | 0.87 | 0.93 | 0.96 | 0.87 | 0.91 | 0.90 | 0.93 | 0.98 | 3.81 | 7.10 |
| 0.90 | 0.93 | 0.92 | 0.95 | 0.90 | 0.88 | 0.92 | 0.96 | 0.91 | 0.90 | 0.93 | 0.96 | 0.97 | 0.98 | 13.77 |

Table C.2: A summary of the extracted $CC\pi^0$ cross section, differential in $\cos\theta_\pi$. The top row, and the left column give the lower-edges of each analysis bin. The second row gives the central value of the cross section measurement in 10^{-39}cm^2 for each kinematic bin while the third row gives the GENIE prediction. The remaining matrix gives the covariance matrix, the numbers above and including the diagonal entries of the matrix in 10^{-81}cm^4 , while entries below the diagonal give the correlation matrix.

| | | | | | | | | | | |
|--------------------------|------|-------|------|------|-------|------|------|------|------|--------|
| | 0.00 | 0.25 | 0.50 | 0.75 | 1.00 | 1.25 | 1.50 | 1.75 | 2.00 | 3.00 |
| $\frac{d\sigma}{dp_\mu}$ | 1.21 | 2.71 | 2.52 | 2.16 | 1.62 | 1.20 | 0.87 | 0.56 | 0.21 | 0.047 |
| GENIE pred | 1.08 | 2.37 | 2.31 | 2.07 | 1.72 | 1.30 | 0.90 | 0.57 | 0.20 | 0.043 |
| 0.00 | 2.37 | 4.46 | 3.12 | 2.02 | 0.86 | 1.28 | 0.95 | 0.51 | 1.16 | 0.29 |
| 0.25 | 0.88 | 10.73 | 7.76 | 5.11 | 2.14 | 3.11 | 2.27 | 1.15 | 2.62 | 0.79 |
| 0.50 | 0.81 | 0.95 | 6.22 | 4.11 | 1.77 | 2.55 | 1.96 | 0.98 | 2.37 | 0.70 |
| 0.75 | 0.75 | 0.89 | 0.94 | 3.08 | 1.53 | 1.89 | 1.35 | 0.71 | 1.51 | 0.43 |
| 1.00 | 0.42 | 0.49 | 0.53 | 0.65 | 1.79 | 0.53 | 0.39 | 0.27 | 0.24 | -0.082 |
| 1.25 | 0.63 | 0.72 | 0.78 | 0.82 | 0.30 | 1.73 | 1.10 | 0.57 | 1.27 | 0.41 |
| 1.50 | 0.64 | 0.72 | 0.82 | 0.80 | 0.30 | 0.87 | 0.93 | 0.42 | 1.06 | 0.32 |
| 1.75 | 0.63 | 0.67 | 0.75 | 0.78 | 0.39 | 0.82 | 0.82 | 0.27 | 0.52 | 0.15 |
| 2.00 | 0.61 | 0.64 | 0.76 | 0.69 | 0.15 | 0.78 | 0.88 | 0.80 | 1.55 | 0.46 |
| 3.00 | 0.34 | 0.45 | 0.52 | 0.46 | -0.11 | 0.58 | 0.61 | 0.53 | 0.68 | 0.29 |

Table C.3: A summary of the extracted $CC\pi^0$ cross section, differential in p_μ . The top row, and the left column give the lower-edges of each analysis bin. The second row gives the central value of the cross section measurement in $10^{-39}\text{cm}^2/(\text{GeV}/c)$ for each kinematic bin while the third row gives the GENIE prediction. The remaining matrix gives the covariance matrix, the numbers above and including the diagonal entries of the matrix in $10^{-81}\text{cm}^4/(\text{GeV}/c)^2$, while entries below the diagonal give the correlation matrix.

| | -1.00 | -0.50 | 0.00 | 0.25 | 0.50 | 0.60 | 0.70 | 0.80 | 0.85 | 0.90 | 0.95 |
|-----------------------------------|-------|-------|------|------|------|------|------|------|------|-------|-------|
| $\frac{d\sigma}{d\cos\theta_\mu}$ | 0.17 | 0.35 | 0.66 | 1.19 | 1.98 | 2.79 | 4.08 | 5.96 | 7.95 | 10.71 | 13.62 |
| GENIE pred | 0.16 | 0.33 | 0.63 | 1.11 | 1.79 | 2.54 | 3.85 | 5.55 | 7.42 | 10.33 | 13.19 |
| -1.00 | 0.40 | 0.63 | 0.57 | 0.96 | 0.47 | 0.80 | 0.89 | 0.53 | 0.74 | 0.98 | 0.88 |
| -0.50 | 0.97 | 1.05 | 0.92 | 1.49 | 0.72 | 1.23 | 1.47 | 0.88 | 1.21 | 1.59 | 1.43 |
| 0.00 | 0.93 | 0.93 | 0.94 | 1.83 | 0.83 | 1.44 | 1.53 | 0.88 | 1.24 | 1.64 | 1.55 |
| 0.25 | 0.73 | 0.70 | 0.90 | 4.39 | 1.80 | 3.27 | 2.82 | 1.55 | 2.36 | 3.16 | 3.18 |
| 0.50 | 0.76 | 0.72 | 0.87 | 0.87 | 0.97 | 1.45 | 1.52 | 0.88 | 1.29 | 1.74 | 1.84 |
| 0.60 | 0.79 | 0.75 | 0.93 | 0.97 | 0.92 | 2.58 | 2.40 | 1.36 | 2.03 | 2.70 | 2.79 |
| 0.70 | 0.77 | 0.78 | 0.87 | 0.74 | 0.85 | 0.82 | 3.34 | 1.85 | 2.51 | 3.09 | 3.35 |
| 0.80 | 0.78 | 0.80 | 0.84 | 0.69 | 0.83 | 0.79 | 0.95 | 1.15 | 1.50 | 1.89 | 2.06 |
| 0.85 | 0.79 | 0.80 | 0.87 | 0.77 | 0.89 | 0.86 | 0.93 | 0.95 | 2.16 | 2.77 | 3.07 |
| 0.90 | 0.77 | 0.77 | 0.84 | 0.75 | 0.88 | 0.84 | 0.84 | 0.88 | 0.94 | 4.04 | 4.00 |
| 0.95 | 0.61 | 0.62 | 0.71 | 0.67 | 0.83 | 0.77 | 0.81 | 0.85 | 0.93 | 0.88 | 5.09 |

Table C.4: A summary of the extracted $CC\pi^0$ cross section, differential in $\cos\theta_\mu$. The top row, and the left column give the lower-edges of each analysis bin. The second row gives the central value of the cross section measurement in 10^{-39}cm^2 for each kinematic bin while the third row gives the GENIE prediction. The remaining matrix gives the covariance matrix, the numbers above and including the diagonal entries of the matrix in 10^{-81}cm^4 , while entries below the diagonal give the correlation matrix.

| | 0.00 | 0.10 | 0.25 | 0.50 | 0.75 | 1.00 | 1.25 | 1.50 | 1.75 | 2.00 | 3.00 |
|------------------------|-------|------|------|------|------|------|------|------|------|------|--------|
| $\frac{d\sigma}{dQ^2}$ | 1.69 | 2.87 | 3.09 | 2.53 | 1.92 | 1.40 | 0.98 | 0.65 | 0.43 | 0.17 | 0.035 |
| GENIE pred | 1.71 | 2.68 | 2.91 | 2.44 | 1.84 | 1.27 | 0.87 | 0.58 | 0.38 | 0.15 | 0.031 |
| 0.00 | 0.50 | 0.99 | 1.69 | 1.28 | 1.08 | 0.82 | 0.52 | 0.33 | 0.17 | 0.29 | -0.008 |
| 0.10 | 0.82 | 2.89 | 4.23 | 3.21 | 3.06 | 2.07 | 1.14 | 0.71 | 0.38 | 0.69 | 0.004 |
| 0.25 | 0.85 | 0.88 | 7.96 | 7.14 | 6.05 | 4.49 | 3.00 | 2.21 | 1.42 | 2.28 | 0.20 |
| 0.50 | 0.66 | 0.69 | 0.92 | 7.50 | 5.97 | 4.60 | 3.32 | 2.65 | 1.82 | 2.82 | 0.34 |
| 0.75 | 0.66 | 0.78 | 0.92 | 0.94 | 5.38 | 3.89 | 2.67 | 2.10 | 1.36 | 2.20 | 0.22 |
| 1.00 | 0.66 | 0.69 | 0.91 | 0.96 | 0.96 | 3.08 | 2.14 | 1.73 | 1.12 | 1.80 | 0.19 |
| 1.25 | 0.56 | 0.52 | 0.82 | 0.94 | 0.89 | 0.94 | 1.67 | 1.34 | 0.92 | 1.44 | 0.17 |
| 1.50 | 0.43 | 0.39 | 0.72 | 0.89 | 0.83 | 0.91 | 0.96 | 1.17 | 0.79 | 1.23 | 0.15 |
| 1.75 | 0.31 | 0.29 | 0.65 | 0.86 | 0.76 | 0.83 | 0.92 | 0.94 | 0.59 | 0.88 | 0.13 |
| 2.00 | 0.35 | 0.34 | 0.69 | 0.87 | 0.81 | 0.87 | 0.95 | 0.96 | 0.97 | 1.38 | 0.19 |
| 3.00 | -0.05 | 0.01 | 0.32 | 0.56 | 0.43 | 0.47 | 0.59 | 0.63 | 0.74 | 0.72 | 0.050 |

Table C.5: A summary of the extracted $CC\pi^0$ cross section, differential in Q^2 . The top row, and the left column give the lower-edges of each analysis bin. The second row gives the central value of the cross section measurement in $10^{-39}\text{cm}^2/(\text{GeV}/c)^2$ for each kinematic bin while the third row gives the GENIE prediction. The remaining matrix gives the covariance matrix, the numbers above and including the diagonal entries of the matrix in $10^{-81}\text{cm}^4/(\text{GeV}/c)^4$, while entries below the diagonal give the correlation matrix.

| | 1.00 | 1.10 | 1.20 | 1.30 | 1.40 | 1.50 | 1.60 | 1.70 | 1.80 | 2.00 | 2.25 | 2.50 |
|----------------------|-------|-------|-------|-------|-------|-------|-------|-------|------|------|--------|--------|
| $\frac{d\sigma}{dW}$ | 0.61 | 2.04 | 3.52 | 3.92 | 4.35 | 4.82 | 4.68 | 3.78 | 2.51 | 1.07 | 0.38 | 0.095 |
| GENIE pred | 0.53 | 1.83 | 3.30 | 3.44 | 3.74 | 4.31 | 4.38 | 3.50 | 2.19 | 0.89 | 0.30 | 0.068 |
| 1.00 | 0.39 | 0.94 | 0.67 | 0.55 | 0.61 | 0.71 | 0.99 | 0.76 | 0.60 | 0.35 | -0.088 | -0.246 |
| 1.10 | 0.93 | 2.63 | 1.95 | 1.33 | 1.45 | 1.77 | 2.71 | 2.19 | 1.98 | 1.26 | -0.016 | -0.361 |
| 1.20 | 0.73 | 0.82 | 2.16 | 1.68 | 1.51 | 1.66 | 2.15 | 1.58 | 1.65 | 1.37 | 0.27 | -0.005 |
| 1.30 | 0.63 | 0.58 | 0.81 | 2.00 | 2.00 | 2.12 | 1.95 | 1.37 | 1.24 | 0.62 | -0.108 | -0.443 |
| 1.40 | 0.60 | 0.55 | 0.63 | 0.87 | 2.64 | 2.93 | 2.58 | 1.95 | 1.61 | 0.23 | -0.491 | -0.927 |
| 1.50 | 0.60 | 0.58 | 0.59 | 0.79 | 0.95 | 3.62 | 3.23 | 2.44 | 1.98 | 0.35 | -0.591 | -1.170 |
| 1.60 | 0.81 | 0.86 | 0.75 | 0.71 | 0.81 | 0.87 | 3.79 | 3.03 | 2.75 | 1.04 | -0.285 | -0.809 |
| 1.70 | 0.75 | 0.83 | 0.66 | 0.59 | 0.74 | 0.79 | 0.95 | 2.66 | 2.51 | 0.83 | -0.179 | -0.538 |
| 1.80 | 0.55 | 0.71 | 0.65 | 0.51 | 0.57 | 0.60 | 0.81 | 0.89 | 3.01 | 1.14 | 0.31 | 0.11 |
| 2.00 | 0.41 | 0.56 | 0.68 | 0.32 | 0.10 | 0.13 | 0.39 | 0.37 | 0.48 | 1.90 | 0.79 | 0.68 |
| 2.25 | -0.17 | -0.01 | 0.21 | -0.09 | -0.35 | -0.36 | -0.17 | -0.13 | 0.21 | 0.67 | 0.73 | 0.84 |
| 2.50 | -0.36 | -0.20 | -0.00 | -0.29 | -0.52 | -0.57 | -0.38 | -0.30 | 0.06 | 0.45 | 0.91 | 1.18 |

Table C.6: A summary of the extracted $CC\pi^0$ cross section, differential in W . The top row, and the left column give the lower-edges of each analysis bin. The second row gives the central value of the cross section measurement in $10^{-39}\text{cm}^2/(\text{GeV}/c^2)$ for each kinematic bin while the third row gives the GENIE prediction. The remaining matrix gives the covariance matrix, the numbers above and including the diagonal entries of the matrix in $10^{-81}\text{cm}^4/(\text{GeV}/c^2)^2$, while entries below the diagonal give the correlation matrix.